

Embedded Ultra Low-Power Digital Signal Processing

1.0 Introduction

The wide-spread and growing use of portable, battery-powered devices like cellular telephones, audio-capable personal digital assistants (PDAs), MP3 players and similar applications has resulted in an increasing demand for miniature, ultra low-power digital signal processing (DSP) technology. Many of these devices make heavy use of digital signal processing techniques like modulation, demodulation, filtering, automatic gain control, equalization and subband coding and decoding. In these devices, users expect a range of DSP-based features to be delivered with little impact on battery life and in miniature, portable packages.

The conflicting demands of ultra-low power consumption and increasing DSP functionality have led to a number of advances in algorithms, semiconductor technologies and system architectures. Based on research for digital hearing aids that started in the early 1990's, we have developed a new DSP system that has benefited from advances in all of these areas. It offers miniature size, ultra low-power consumption and is sufficiently flexible to support a wide range of applications.

This technology will result in a new range of devices where ultra low-power, miniature DSP technology is embedded into a system or sub-system and invisibly performs a useful task. By embedding ultra low-power, miniature signal processing capabilities, we expect improved performance in everything from embedded sensors to digital hearing aids, especially in adverse signal conditions.

This paper presents an overview of the requirements for ultra low-power embedded DSP systems, the technology that was developed for our signal processing system, and a detailed look at a demanding application: a digital, frequency-domain, beamforming hearing aid.

2.0 System Overview

2.1 Requirements

The requirements for embedding DSP systems into miniature, ultra low-power applications are challenging (Table 1). These requirements were driven by our initial application, digital hearing aids. In this application, size and power consumption are particularly restrictive.

Table 1: Requirements for a miniature, ultra low-power DSP system

Size	<ul style="list-style-type: none"> Miniature size (hearing aids require a complete DSP system in less than 3 x 5 x 3 mm)
Power	<ul style="list-style-type: none"> Single-battery operation; operates to 0.9 volts Less than 1 mA system current consumption (< 0.1 mW/MIPS for DSP platform)
Performance	<ul style="list-style-type: none"> At least 5 MIPS of signal processing capability Flexibility to support a wide range of applications Broadcast quality fidelity (minimum 8 kHz bandwidth) Less than 10 ms group delay More than 50 dB of gain adjustment

3.0 System Design

Figure 1 shows a block diagram of the system. It consists of three major components:

- Weighted overlap-add (WOLA) filterbank coprocessor,
- RCORE DSP core, and
- Input-output processor (IOP).

by *Todd Schneider and Robert Brennan,*
Dspfactory Ltd., Waterloo, ON
Edward Chau, University of Guelph, Guelph, ON

Abstract

This paper presents an overview of the requirements for ultra low-power embedded DSP systems, the technology that was developed for our signal processing system, and a detailed look at a demanding application: a digital, frequency domain, beamforming hearing aid.

Sommaire

Cet article présente un sommaire des exigences des systèmes embarqués de très faible puissance pour le traitement numérique du signal. Cette technologie a été développée pour notre système de traitement du signal. L'article présente une analyse détaillée d'une application particulièrement exigeante, soit un appareil acoustique numérique pour malentendants atténuant les bruits de fond, amplifiant les conversations selon leur direction et oeuvrant dans le domaine fréquence.

A mixed-signal sub-system contains the analog-to-digital converters (A/D), a digital-to-analog converter (D/A) and other interface circuitry. Both the RCORE and the WOLA coprocessor can run concurrently providing approximately 5 MIPS performance on a 1 MHz system clock.

Figure 2 shows the processing model for the system. A time domain input signal, $x(n)$, is transformed into the frequency domain by the analysis filterbank, the RCORE can then manipulate the gains applied to the complex output from the filterbank. The synthesis filterbank transforms data back to a time-domain signal, $y(n)$. In essence, the design is an over-sampled, subband CODEC. The output from the WOLA is complex and contains both magnitude and phase information.

3.1 WOLA Filterbank

The vast majority of DSP algorithms, everything from subband CODECs to directional processing, can be cast into a filtering paradigm. Thus, our design incorporates an efficient, hardware-based filtering coprocessor, the weighted overlap-add (WOLA) filterbank [1,3,8]. The WOLA is implemented in hardware and this results in:

- Greatly reduced power consumption because a signal processing architecture optimized for filtering is more power efficient than a general purpose architecture doing the same processing, and
- Reduced chip size because less memory is required.

To provide the flexibility required for a range of applications, the WOLA filterbank has a number of adjustable parameters. The fast Fourier transform (FFT) Size (N), window length (L) and input block step size (R) are all adjustable. Two key innovations in the WOLA filterbank design are the incorporation of adjustable oversampling and the provision for two filterbank stackings, even and odd. Adjustable oversampling allows a user selectable trade-off to be made between fidelity, group-delay and power consumption [1]. Results for some configurations are shown in Table 2. Note how reduced group delay (greater oversampling and/or a smaller window length) can be "traded" for increased power consumption, reduced fidelity (a lower spurious-free dynamic range, SFDR) or both. The WOLA filterbank can be configured from 4 to 128 bands.




Figure 1: Block diagram of DSP system




Figure 2: System processing model

3.2 RCORE DSP Core

The RCORE DSP core provides the flexibility needed to implement a wide range of signal processing algorithms. It has access to the frequency domain data (output from the analysis section of the WOLA filterbank) and the time-domain data (in the WOLA filterbank input and output FIFO buffers).

The RCORE is a fully software programmable, 16-bit, dual-Harvard DSP core. It performs a single-cycle multiply accumulate with simultaneous update of two address pointers. It has instructions that are specialized for audio processing (e.g., single-cycle normalization and denormalization) and a 40-bit accumulator. It interfaces with the WOLA filterbank

and the IOP through shared memory.

Table 2: Sample filterbank configurations (SFDR: spurious-free dynamic range; relative power for filterbank only)

Bands (=N/2)	OS (=N/R)	Delay (ms)	Rel. Power	SFDR (dB)
16	2	14	1	65
16	4	6	1.5	50
32	4	12	1.6	45
128	1	27	2	40

3.3 Input-Output Processor (IOP)

The IOP is a block-based direct-memory access controller that is tightly coupled to the WOLA filterbank. It operates on blocks of data and only interrupts the DSP core when necessary. This reduces power consumption because the DSP core can switch to a low-power sleep mode when it is not needed for calculations.

The IOP incorporates decimation and interpolation filters that work in conjunction with oversampling A/D and D/A converters. The decimation filter has an integral DC removal filter.

3.4 System Implementation

Further reductions in power consumption are provided by (1) operating directly at single battery voltage (the system will operate down to 0.9 volts) and (2) using low-power, deep submicron semiconductor technology [7].

The entire system (Figure 1) is implemented on three integrated circuits. The WOLA filterbank, RCORE, IOP and associated peripherals are fabricated using 0.18 μ technology on a die that is less than 10 mm².

The design also incorporates an ultra low-power integrated circuit that has two 14-bit A/Ds and a 14-bit D/A converter. This subsystem also has programmable input and output gain blocks as well as an on-chip oscillator and charge-pump. The entire mixed-signal subsystem is under software control via a low-speed, single-wire synchronous serial interface. This circuit is fabricated using 1.0 μ semiconductor technology on a die that is less than 8 mm². A third, off-the-shelf EEPROM die provides non-volatile memory for the system.

Figure 4 shows packaged versions of the system that incorporate the digital die, the mixed-signal die and the EEPROM die (128 kbits).

Figure 3 shows frequency response plots for even and odd stackings. For the configurations shown, 16 bands of frequency equalization are available, each with over 40 dB of gain adjustment. Both stackings (even and odd) have a group delay (τ) of only 6 milliseconds, including the blocking delay introduced by the IOP (which simultaneously inputs and outputs blocks of data while the WOLA filterbank is running).

Even stacking uses a traditional FFT and provides N/2-1 (where N is the FFT size) full bands and two half bands (at DC and the Nyquist frequency). Odd stacking provides N/2 equal width bands. Having two stackings provides for more precise equalization because there are twice the number of band edges.

Finally, the WOLA filterbank can operate in stereo mode and simultaneously convert two time-domain signals into the frequency domain. This feature, along with the complex output signal from the filterbank, makes the WOLA filterbank ideal for the implementation of frequency-domain directional processing algorithms and demodulators.

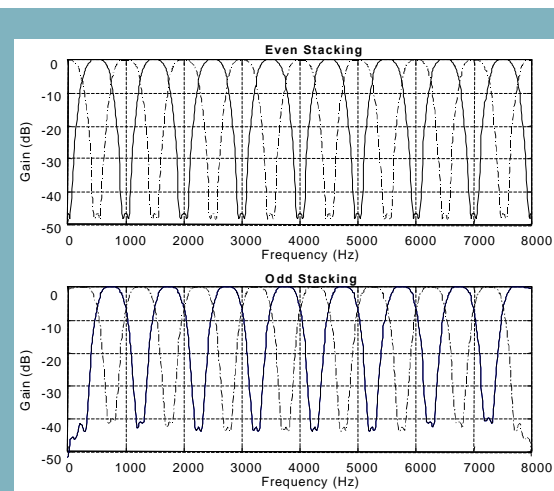


Figure 3: Frequency responses for even and odd stacking (16-channels, $\tau = 6$ ms, $f_s = 16$ kHz)




Figure 4: Packaged DSP system: (a) hybrids for hearing aid applications, and (b) a multi-chip module for PDA and other portable applications

4.0 Applications

Our DSP system has a wide range of applications. It is already implemented in digital hearing aids [6], speech recorders (as a subband CODEC) and PDA applications.

We are actively working on several directional processing algorithms, everything from simple two-microphone delay-and-sum systems to advanced frequency domain beamforming. The stereo processing mode of the WOLA filterbank greatly simplifies the implementation of these algorithms. The remainder of this paper discusses these interesting applications in more detail.

4.1 Beamforming Hearing Aid

Background noise amplified by a hearing aid makes it very difficult for many hearing aid users to under-

stand speech. A proven approach to improve speech intelligibility for these users in background noise is to employ a beamformer [5]. A beamformer is a spatial filter that allows filtering of signals depending on the direction-of-arrival (DOA) of the signals. Assuming that the user tends to face with the desired signal source, a beamformer can be used to suppress sounds that are not originating from this look-forward direction, thereby improving speech intelligibility.

In order to resolve the signal DOA, a beamformer needs to employ an array of two or more sensors (microphones). Generally, the more sensors that are available in the array, the better the beamformer performance. Some beamformers developed for speech intelligibility enhancement have employed arrays of five or more microphones [5]. However, with the small size of typical hearing aids, it is often impractical to implement an array of more than two microphones.

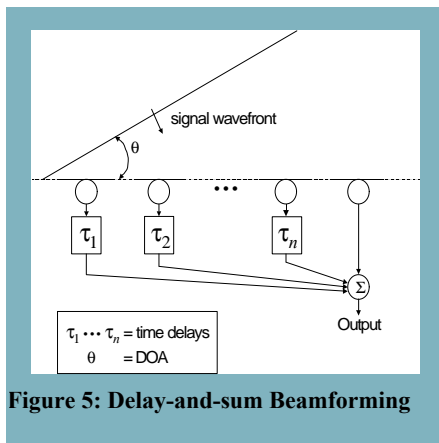


Figure 5: Delay-and-sum Beamforming

from a specific DOA and frequency [4]. Essentially, the time delay is applied such that the signals from each microphone will be time-aligned. The time-aligned signals are then summed together so that the power of the signal components originating from a particular DOA is enhanced relative to the power of those from other directions (see Figure 5).

The gain response of the classical delay-and-sum beamformer is both frequency- and DOA-dependent. Consider an array of two microphones separated by a distance d . Let $\omega_m = \pi c/d$, where c is the speed of sound. Figure 6 shows the beam patterns (polar plot of the beamformer gain response) of a beamformer aimed at 0 degree DOA for signals at various frequencies. As can be seen in the figure, at frequencies lower than ω_m , the nulls are degraded and, at higher frequencies, spatial aliasing causes additional main lobes to appear. This occurs because, while the propagation delay of the signal wavefront remains the same for all frequencies, the corresponding phase delays are different at different frequencies.

Frequency-dependent gain response is clearly undesirable in hearing aid applications, where the gain response of the beamformer should be consistent over all frequencies of interest. Fortunately, with the use of a powerful DSP platform and a stereo filterbank, the problem of a frequency-dependent beam pattern can be easily alleviated by applying a frequency-domain extension to the classical delay-and-sum algorithm.




Figure 6: Beam patterns at different frequencies for 0-degree steering

Assuming again an array of two microphones, the new algorithm introduces two additional frequency-dependent delays so that, in effect, besides applying the constant beamforming delay, it also applies a variable delay (as a function of frequency) to both of the received signals at the microphones. The variable delays compensate for the different phase delays at each frequency component, so that the resulting phase delay over all frequencies is the same as that at ω_m . This provides the same beam patterns over all frequencies. However, to avoid spatial aliasing, ω_m must be set at the highest frequency of interest. Figure 7




Figure 7: Extension to Classical Delay-and-Sum Beamformer

shows the new beamformer for the case of a two-microphone array. In the figure, τ_1 is the constant beamforming delay (for aiming towards a particular DOA), and $\tau_1^*(\omega)$ and $\tau_2^*(\omega)$ are the two frequency-dependent delays for compensation. The summation sign in the figure actually denotes the “butterfly” operation instead of the simple arithmetic summation. Note that this beamformer can be implemented only in the frequency domain, because the actual phase delay between the two received signals at each frequency must be known for all times.

In theory, this beamformer will produce exactly the same beam pattern for any frequency component $\omega \leq \omega_m$. In practice, however, the beam pattern is subject to “maladjustment” because of the finite bandwidth of the filterbank subbands. Clearly, the effect of this maladjustment is more apparent with wider subbands. We have found that with a 64-band WOLA filterbank, the effect of this maladjustment is negligible.

Another potential cause for maladjustment in this beamformer is that the determination of the phase delay at each subband assumes that the dominant energy in the subband comes from a single signal source only. The reason for this is that for signal sources with different DOA, different compensation is needed to produce the consistent beam patterns. Hence, as long as the dominant energy in each subband is contributed by one signal source only, the compensations will be accurate.

For simulation, this beamformer has been implemented in C, using the WOLA filterbank structure [1] with 16-band and 64-band implementations. A 10-second male speech utterance (target) is mixed in white noise at various SNR for use as the test signal. In general, the simulation has shown that an average of 10 dB improvement can be obtained using the frequency-domain beamformer. While it has been found that the performance of this beamformer tends to degrade quickly when more than one noise source is present, overall, with an efficient filterbank and DSP platform, this beamformer is a simple yet effective way of providing background noise reduction in digital hearing aids.

Finally, since the beamformer described here is a relatively simple algorithm that performs well under favorable conditions, the development of our ultra low-power DSP platform offers the computing resources for more complex algorithms. One novel approach we are investigating involves a neural network based system that supplements the frequency-domain beamformer to provide better background noise suppression. Figure 8 shows the block diagram of the overall system.

Assuming that the neural network module will operate without on-line adaptation, a static neural network is simply a sequence of multiply-sum operations, with the activation function easily approximated by a look-up table. We expect that this system can be implemented easily on our DSP platform, provided a satisfactory neural network solution can be found.




Figure 8: Block Diagram of Novel System

5.0 Conclusions

Software programmable, ultra low-power, miniature DSP systems will result in whole new range of DSP applications such as digital hearing aids, audio enabled personal digital assistants and portable audio playback devices. Our design demonstrates that ultra low-power DSP systems can offer sufficient computational capability and flexibility to be used in a range of applications.

We believe our experience in this area can be generalized to other ultra low power, miniature applications: the greatest savings in power and size comes from having an efficient algorithm that is targeted at a specific algorithm or a class of algorithms. Our design is an application specific signal processor (ASSP) that incorporates very efficient, yet flexible filtering

The specific example of a beamforming hearing aid illustrates that even complex, two-input frequency domain algorithms can be supported on such miniature, ultra low-power platforms. In the near future, such algorithms will bring much needed benefit to hearing aid users and possibly find application in other systems (e.g., speech recognition front-end processing) where an improvement in SNR will result in more robust system operation.

6.0 References

- [1]. Brennan, R.L., and Schneider, T., "Filterbank Structure and Method for Filtering and Separating an Information Signal into Different Bands, Particularly for Audio Signals in Hearing Aids", PCT Patent Application PCT/CA98/00329.
- [2]. Brennan, R.L., Schneider, T., "A Flexible Filterbank Structure for Extensive Signal Manipulations in Digital Hearing Aids," Proc. ISCAS-98, Monterey, CA.

- [3]. Crochiere, R.E. and Rabiner, L.R., Multirate Digital Signal Processing. Prentice-Hall Inc., 1983.
- [4]. Orfanidis, S., "Optimum Signal Processing", Macmillan, 1988, pp. 341-343.
- [5]. Saunders, G., and Kates, M., "Speech Intelligibility enhancement using hearing-aid array processing", Journal of Acoustical Society of America, 1997, pp. 1827-1837.
- [6]. Schneider, T., Brennan R.L., "A Multichannel Compression Strategy for a Digital Hearing Aid," Proc. ICASSP-97, Munich, Germany, pp. 411-415
- [7]. Schneider, Todd et al., "An Ultra Low-Power Programmable DSP System for Hearing Aids and Other Applications," Proc. ICSPAT-99, Orlando, FL.
- [8]. Vaidyanathan, P.P., Multirate Systems and Filter Banks. Prentice-Hall Inc., 1993.

7.0 Glossary

DOA	- Direction of Arrival
IOP	- Input-Output Processor
CODEC	- Coder/Decoder
RCORE	- DSP Core
WOLA	- Weighted Overlap-Add
DSP	- Digital Signal Processor
PDA	- Personal Digital Assistant
SNR	- Signal to Noise Ratio
ASSP	- Application Specific Signal Processor
FFT	- Fast Fourier Transform
SFDR	- Spurious-Free Dynamic Range
FIFO	- First In - First Out
EEPROM	- Erasable Programmable Read Only Memory

About Dspfactory Ltd. : Dspfactory is a rapidly growing, dynamic company with expertise in digital signal processing (DSP) architectures and algorithms for miniature, ultra low-power audio and baseband applications. Its mission is to embed ultra low-power, miniature DSPs invisibly into a wide range of products. Target products for its technology include hearing aids, baseband wireless, personal digital assistants, personal digital audio players, cellular telephones and embedded sensors -- in short, any DSP-based products that are portable and battery-powered. More information is available at www.dspfactory.com

About the Authors

Todd Schneider graduated from the University of Waterloo with a B.A.Sc. (1989) and a M.A.Sc. (1991), both in Electrical Engineering.



He is now the VP Technology at Dspfactory. His technical interests include DSP algorithms, system architectures for efficient DSP systems, DSP tools and Linux.

He is a member of the IEEE and the Audio Engineering Society.

Robert Brennan graduated with a doctoral degree in electrical engineering from the University of Waterloo in 1991 investigating low bit-rate speech coders.



As VP Research at Dspfactory, he continues work on filterbank speech decomposition methods and speech enhancement/processing strategies.

He is a member of the IEEE and the Acoustical Society of America.

Edward Chau is currently pursuing his M.Sc. Degree in Engineering Systems and Computing at the University of Guelph. He obtained his B.A.Sc. in Electrical Engineering at University of Waterloo in 1999.



His primary research interests include Neural Network & Evolution Computing approaches to Digital Signal Processing, particularly in audio signal processing. He is currently developing a neural network based noise reduction module for digital hearing aids. He would like to thank Dspfactory and the Natural Science and Engineering Research Council for their support in his graduate research.