



An Introduction to Digital Filters

Authors: Dr. David B. Chester, Geoff Phillips, Stan Zepp

Introduction

Digital Signal Processing (DSP) affords greater flexibility, higher performance (in terms of attenuation and selectivity), better time and environment stability and lower equipment production costs than traditional analog techniques. Additionally, more and more microprocessor circuitry is being displaced with cost-effective DSP techniques and products: an example of this is the emergence of DSP in cellular base stations. Components available today let DSP extend from baseband to intermediate frequencies (IFs). This makes DSP useful for tuning and signal selectivity, and frequency up and down conversion.

These new DSP applications result from advances in digital filtering. This Application Note will overview digital filtering by addressing concepts which can be extended to baseband processing on programmable digital signal processors.

Essentials of Digital Filtering

A digital filter is simply a discrete-time, discrete-amplitude convolver. An example is shown in Figure 1 for three-bit amplitude quantization. Basic Fourier transform theory states that the linear convolution of two sequences in the time domain is the same as multiplication of two corresponding spectral sequences in the frequency domain. Filtering is in essence the multiplication of the signal spectrum by the frequency domain impulse response of the filter. For an ideal lowpass filter the pass band part of the signal spectrum is multiplied by one and the stopband part of the signal by zero.

Analog Filters, Software-Based and Hard-wired Digital Filters

Owing to the way that analog and digital filters are physically implemented, an analog filter is inherently more size- and power-efficient, although more component-sensitive, than its digital counterpart - if it can be implemented in a straightforward manner. In general, as signal frequency increases, the disparity in efficiency increases.

Characteristics of applications where digital filters are more size and power efficient than analog filters are: linear phase, very high stop band attenuation, very low pass band ripple; the filter's response must be programmable or adaptive; the filter must manipulate phase and; very low shape factors (a digital filter's shape factor is the ratio of the filter's pass band width plus the filter's transition band width to the filter's pass band width).

General-purpose digital signal microprocessors, now commodity devices, are used in a broad range of applications and can implement moderately complex digital filters in the audio frequency range. Many standard signal processing algorithms, including digital filters, are available in software packages from digital signal processor and third party vendors. As a result, software development costs are trivial when amortized over production quantities.

The architectures of digital signal microprocessors are usually optimized to perform a sum-of-products calculation with data from RAM or ROM. They are not optimized for any specific DSP function. However, to get extended sampling rate performance from a digital filter requires hardware designed to perform the intended filter function at the desired sampling frequencies.

For example, Harris Semiconductor offers a family of standard digital filter products with several others in development. Some hardware-specific digital filters can now sample at rates approaching 75 Megasamples per second (MSPS). Higher performance is possible for high volume applications by limiting the range of parameters. Standard filter products strike a balance between optimized filter architectures and programmability by offering a line of configurable filters. That is, these products are function-specific, with optimized architectures and programmable parameters.

Conceptual Differences

Frequency-Domain Versus Time Domain Thinking

Thinking about analog filters, most engineers are comfortable in the time domain. For example, the operation of an RC lowpass filter can easily be envisioned as a capacitor charging and discharging through a resistor. Likewise, it is easy to envision how a negative-feedback active filter uses phase shift as a function of frequency, which is a time domain operation.

A digital filter is better conceptualized in the frequency domain. The filter implementation simply performs a convolution of the time domain impulse response and the sampled signal. A filter is designed with a frequency domain impulse response which is as close to the desired ideal response as can be generated given the constraints of the implementation. The frequency domain impulse response is then transformed into a time domain impulse response which is converted to the coefficients of the filter.




FIGURE 1A. $x(n)$ (SIGNAL SAMPLES)




FIGURE 1B. $h(n)$ (FILTER IMPULSE RESPONSE SAMPLES)




FIGURE 1C. $y(n) = h(n) \oplus x(n)$ (OUTPUT SAMPLES) = $\sum_{m=0}^2 h(m) \times (n-m)$

FIGURE 1. THE DISCRETE TIME AND DISCRETE AMPLITUDE CONVOLUTION OF $h(n)$ WITH $x(n)$

Mathematics Versus Physics

The characteristics of an analog filter are directly attributable to the physics of the device that implements it. In contrast, the characteristics of a digital filter are only indirectly attributable to the physics of the device that implements it.

A digital filter's passband ripple, shape factor, stopband attenuation, and phase characteristics are all functions of the order and type of polynomial used to approximate the ideal impulse response, the number of bits used in performing the arithmetic, and the type of architecture used to implement the arithmetic. Actual frequencies have no meaning in a digital filter except in their relation to the sampling frequency. This is because the impulse response is generated as a function of z^{-1} , the sample interval (the time between samples). For a smaller shape factor, the order of the filter and the number of bits in the arithmetic can be increased. The only physical limitation is the amount of arithmetic processing that can be integrated on a device or devices given the filter order and the input sampling rate.

Digital Filter Types

There are two basic types of digital filters, finite impulse response (FIR) and infinite impulse response (IIR) filters. The general form of the digital filter difference equation is

$$y(n) = \sum_{i=0}^N a_i x(n-i) - \sum_{i=1}^N b_i y(n-i)$$

where $y(n)$ is the current filter output, the $y(n-i)$'s are previous filter outputs, the $x(n-i)$'s are current or previous filter inputs, the a_i 's are the filter's feed forward coefficients corresponding to the zeros of the filter, the b_i 's are the filter's feedback coefficients corresponding to the poles of the filter, and N is the filter's order. IIR filters have one or more non-zero feedback coefficients. That is, as a result of the feedback term, if the filter has one or more poles, once the filter has been excited with an impulse there is always an output. FIR filters have no non-zero feedback coefficient. That is, the filter has only zeros, and once it has been excited with an impulse the output is present for only a finite (N) number of computational cycles.

Figures 2 and 3 show common IIR architectures. Figures 4 through 6 show common FIR architectures.

Strengths and Weaknesses

Because an IIR filter uses both a feed-forward polynomial (zeros as the roots) and a feedback polynomial (poles as the roots), it has a much sharper transition characteristic for a given filter order. Like analog filters with poles, an IIR filter usually has nonlinear phase characteristics. Also, the feedback loop makes IIR filters difficult to use in adaptive filter applications.

Due to its all zero structure, the FIR filter has a linear phase response when the filter's coefficients are symmetric, as is the case in most standard filtering applications. An FIR's implementation noise characteristics are easy to model,

especially if no intermediate truncation is used. In this common implementation, the noise floor is at $-6.02 B + 6.02 \log_2 N$ dB where B is the number of actual bits used in the filter's coefficient quantization and N is again the filter order. This is why most Harris filter ICs have more coefficient bits than data bits.

An IIR filter's poles may be close to or outside the unit circle in the Z plane. This means an IIR filter may have stability problems, especially after quantization is applied. An FIR filter is always stable. FIR filters also allow development of computationally efficient architectures in decimating or interpolating applications, which will be described in more detail later.

Filter Response Design Methods

The total specification of the ideal filter includes the location of passbands and stopbands, the minimum stopband attenuation, the maximum passband ripple, the filter order, and perhaps the shape of the response in some of the specified bands.

Typically, there are three stages to the design of digital filter responses for passband filters. First, the ideal filter response is specified. Next, a floating point response is designed. Finally, the floating point coefficients are quantized to yield a fixed point response.

Creating a floating-point IIR filter response starts with a prototype analog filter. Then an s-domain to z-domain transformation is used to generate a set of digital filter coefficients. Common methods of designing floating-point FIR filter responses are windowing, frequency sampling, and optimal. All are described in general digital signal processing texts and are standard in most commercially available digital filter design software packages.

Converting floating-point coefficients into fixed-point coefficients requires quantizing the coefficients and calculating the frequency domain impulse response of the filter or filter model to verify that the filter meets the required specification. If it does not, either the number of coefficient bits can be increased, the filter response can be redefined and step two repeated, or the filter arithmetic can be redesigned, or a combination of these procedures can be performed. When filter hardware is at a premium, sophisticated simulated annealing techniques can be used for both fixed-point FIR and IIR filters to produce the best set of filter coefficients, given a fixed filter order and coefficient width.

Decimation and Interpolation

Decimation (Figure 7) is reducing the output sampling rate by ignoring all but every M th sample. When a digital filter reduces the bandwidth of a signal of interest so the filter output is over-sampled if the input sample rate is preserved, it is inefficient to compute outputs that will be ignored in the decimation process. Thus, there is a one-to-one correspondence between decimation rate and gain in computational efficiency. However, this computational efficiency can not be fully realized in an IIR filter because the feedback path must be computed for every input cycle.

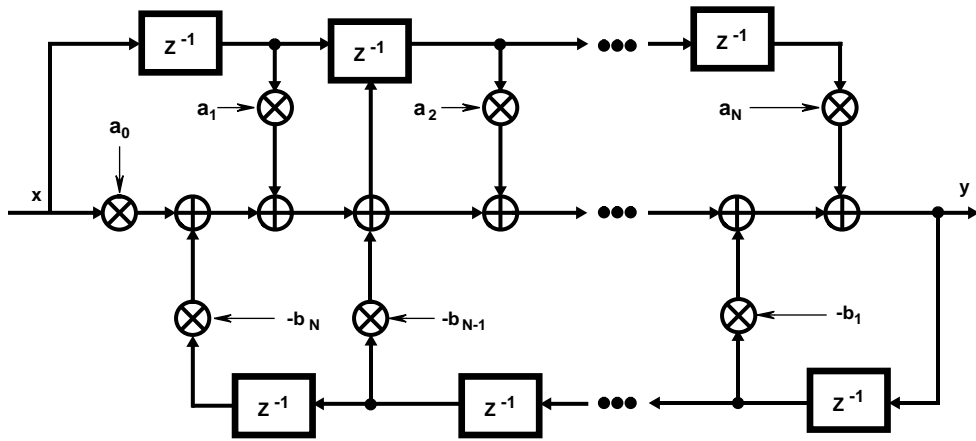


FIGURE 2. IIR FILTER DIRECT FORM 1




FIGURE 3. IIR FILTER DIRECT FORM 2




FIGURE 4. TRANSVERSAL IMPLEMENTATION OF AN FIR FILTER

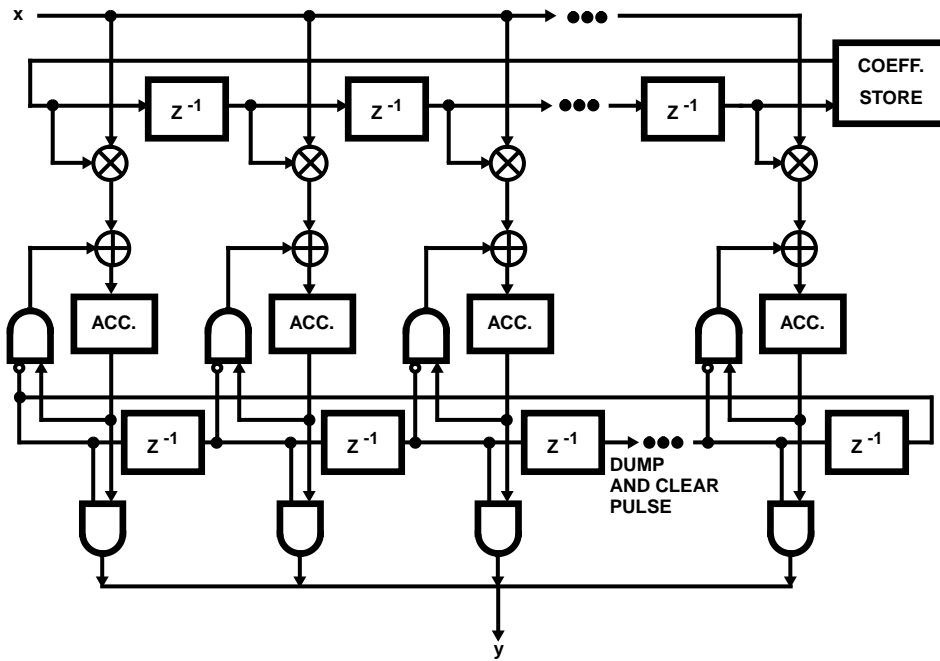


FIGURE 5. PARALLEL MULTIPLIER/ACCUMULATOR CELL FIR FILTER IMPLEMENTATION




FIGURE 6. TRANSPOSED FIR FILTER IMPLEMENTATION




FIGURE 7. BLOCK DIAGRAM AND SPECTRAL REPRESENTATION OF THE DECIMATION PROCESS

Interpolation (Figure 8) means increasing the sampling rate. It is most often used when a narrowband signal will be combined with a signal that requires a higher sampling rate. Conceptually, the first step in interpolation is to stuff $L-1$ zero-valued samples between each valid input sample, expanding the sampling rate by L , and causing the original signal spectrum to be repeated $L-1$ times. To perform the actual interpolation, the zero-valued input samples must be converted to approximations of signal samples. This is equivalent to preserving the original signal spectrum. Effectively, the zero-stuffed input stream is filtered by a lowpass filter with a pass band at the original spectrum location. This filters out all of the repeated spectra.

In FIR filters, interpolation computational efficiency is possible because only every L th data value is non zero and thus requires an actual multiply-add operation. Consequently, there is a one-to-one correspondence between interpolation rate and gain in computational efficiency. The efficient interpolation structure is called polyphase. Again, because of the feedback loop in an IIR filter, this computational efficiency can not be realized.

Hardware Versus Software

There is no difference between the mathematics and filter response design of a hardware versus a software implementation of a digital filter. The only differences lie in the implementation architectures.

Optimal Architectures

A digital filter implemented in software is typically run on a processor with a single computational element. This forces the filter to be implemented in a serial sum-of-products. An FIR is implemented as a single sum-of-products and an IIR is typically a sum of products for the feed-forward section and another sum-of-products for the feedback section.

Hardware can be optimized for each application. Low speed and low power can be achieved using a bit-serial implementation. Moderate speed and power may call for a single parallel multiplier-accumulator. High speed applications can utilize multiple multiplier-accumulator structures with specialized memory address schemes. Decimating or interpolating filters can use optimized polyphase structures, multiple stages and specialized memory addressing schemes. Coefficient memory can be designed to take advantage of coefficient characteristics and accumulators can be custom-designed to take advantage of unity gain properties.

Speed/Flexibility Tradeoffs

When a digital filter is hard-wired for one specific task, its performance can be optimized because there is no control overhead associated with reconfiguring the IC.

Even with a hard-wired filter, however, some degree of flexibility can be achieved without sacrificing much computational efficiency. For example, an FIR filter with programmable coefficients and filter order and a limited selectable decimation rate requires only a modified memory and memory addresser added to a fixed FIR filter.

Harris Semiconductor's line of standard digital filtering devices focuses on reconfigurable function-specific devices. The standard products are specialized FIR filters; IIR filter implementations are limited to application-specific designs. Most of the Harris standard filters have programmable decimation rates. Some can interpolate. Most have programmable coefficients and filter orders and some provide up/or down-conversion.

General Applications

Traditionally, most digital filter applications have been limited to audio and high-end image processing. With advances in process technologies and digital signal processing methodologies, digital filters are now cost-effective in the IF range and in almost all video markets.

Digital filters are commonly used for audio frequencies for two reasons. First, digital filters for audio are superior in price and performance to the analog alternative. Second, audio analog-to-digital converters (A/Ds) and digital-to-analog converters (DACs) can be manufactured with high accuracy and are available at low cost. Thus, the combined cost of filtering and conversion (if necessary) is low. The cost trades are much more difficult in the 1MHz to 100MHz signal range, such as the IF ranges of many radio receivers.

While digital signal processing technology can now produce cost effective digital filters for IF, the cost or even the availability of data conversion products are limiting factors. Many IF digital filtering applications are band-limiting and decimating. In these cases the design engineer must not only know digital filters, but also understand the effects of narrow-band-filtering processing-gain on A/D requirements. Additionally, power dissipation must be considered. Currently, digital IF filter solutions are excluded from low power applications such as personal communication devices. In contrast, audio frequency digital filters are essential.




FIGURE 8. BLOCK DIAGRAM AND SPECTRAL REPRESENTATION OF THE INTERPOLATION PROCESS

Application Note 9603

IF Processing

Harris' HSP43216 Halfband Filter IC (Figure 9) can perform a quadrature split on a real signal. In this example, the input signal undergoes anti-alias filtering and is digitized and passed to the Halfband Filter IC. The quadrature $f_s/4$ local oscillator (LO) and mixer circuits on the IC center the upper sideband of the real signal spectrum at DC. The Halfband Filter itself operates on the resultant complex signal to filter out the lower sideband, forming a quadrature signal (a characteristic of which is a single-sided spectrum). Other circuits in the IC then decimate the real and imaginary output by two to eliminate the unused spectral region.

To cite a more specific application, IF processing can be implemented in a cellular base station using the HSP50016 Digital Down Converter (DDC) and two HSP43124 Serial I/O Filters (Figure 10). In this example a 4MHz band of the GSM spectrum has already been mixed to a near baseband IF, has been appropriately anti-alias filtered, and digitized by an A/D. The spectral plot makes the point that aliasing may take place as long as it does not encroach on the band of interest.

Because this is a channelizing application, there is a signal-to-noise ratio (SNR) processing gain of 3dB for every factor of 2 that the noise bandwidth is reduced. In this example, the processing gain is approximately 17dB. As a result, the SNR

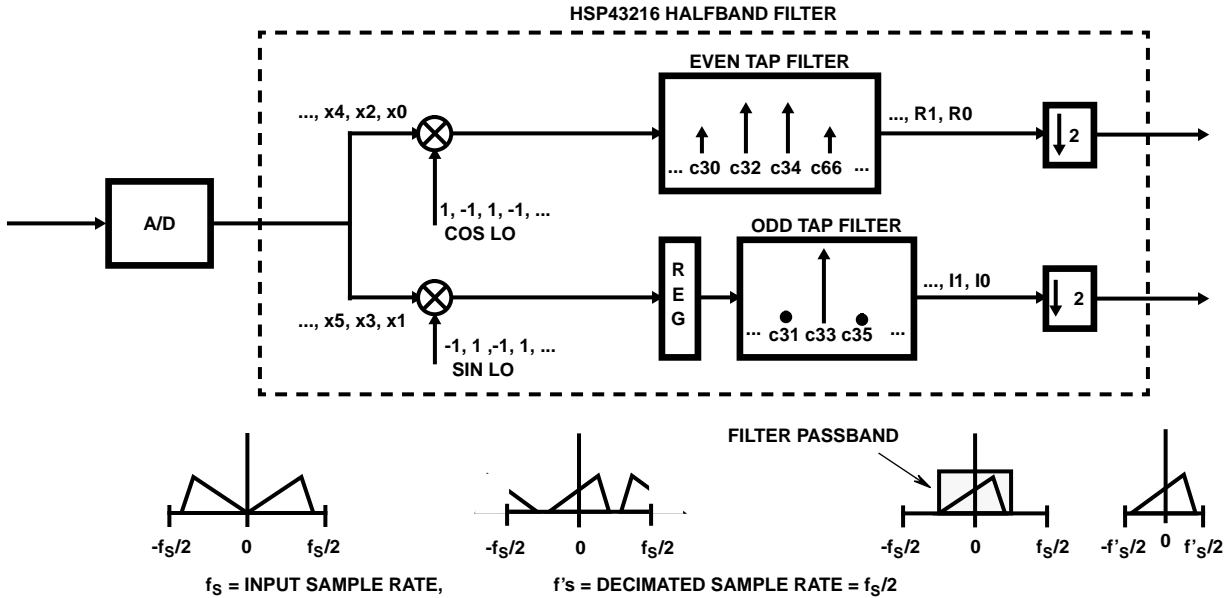


FIGURE 9. REAL TO QUADRATURE CONVERSION USING THE HSP43216 HALFBAND FILTER

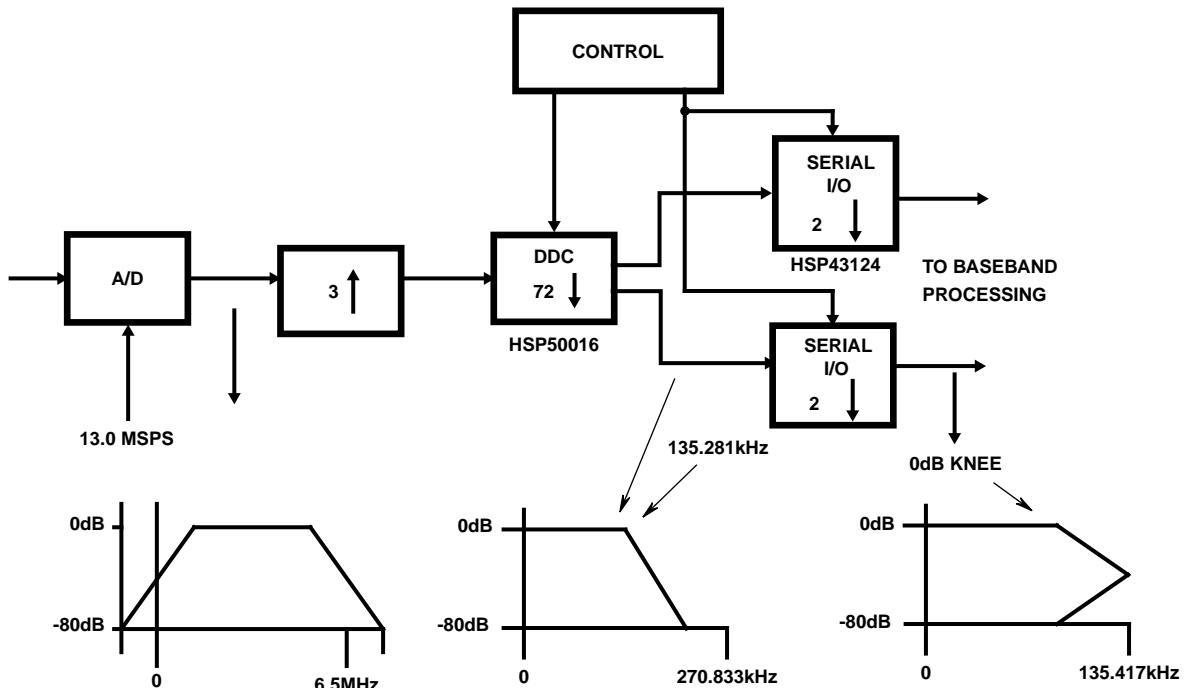


FIGURE 10. GSM BASE STATION EXAMPLE

of the A/D can be 17dB lower than the desired output SNR. The A/D's spurious free dynamic range (SFDR), however, must be equal to the desired output SFDR since there is no gain effect on in-band frequency spurs.

The DDC is a single-chip quadrature down-converter and two stage filter. It is used to tune the channel of interest to base-band, and to perform narrow-band filtering and decimation. The Serial I/O Filters are used to apply the GSM filter shape to the output quadrature data stream. This process includes decimation by two to modify the output sampling rate to that required by the GSM specification. This requirement forces the DDC to decimate at a rate below its minimum of 64 in quadrature-output mode. This is dealt with by placing a three-times sample rate expander in front of the DDC. The HSP50016 and HSP43216 data sheets provide greater detail.

Application-Specific Architectures

As can be seen in the example above, while reconfigurable function-specific digital filters provide viable solutions to specific applications, they may not be cost-effective in high volumes (Annual quantities in the 10 to 100 thousand range). This is where application-specific architectures play a role. In high volume applications, reconfigurable function-specific digital filters can be used initially for prototyping and refining an efficient architecture. Then, that architecture can be implemented in an application-specific device. In most cases, because of the sophisticated nature of digital filters, efficient architectures are likely to require semicustom or custom designs.

Imaging/Video Systems With Separable Dimensions

Most digital signal processing applications involving filter operations on two-dimensional data assume that the horizontal and vertical filtering operations can be separated. That is, (Figure 11) the image data can be filtered in one dimension (for example, the horizontal or row dimension) followed by filtering in the other dimension (the vertical or column dimension). This separability greatly simplifies the filter design compared to a two-dimensional filter.

One- and Two-Dimensional Filters

One-dimensional DSP operations are widely used in imaging or video systems to perform down-conversion or filtering. In video, purely one-dimensional operations are almost always in the horizontal dimension because this is the dimension in which video is sampled in real time.

Two-dimensional operations are used primarily to alter the size and shape of the image or to filter in two dimensions. The latter operations include highpass filters, to sharpen edges in all directions, or lowpass filters, to limit high-frequency noise or to deliberately soften edges. An important case is image resizing, where the input image is resampled to a different sized output image. In reducing the image size, filtering is needed because simply down-sampling (throwing away pixels) vertically and horizontally produces unacceptable aliasing. A two-dimensional filter can be made from one-dimensional filters (Figure 12). Here, an HSP43168 Dual FIR Filter provides horizontal band-limiting prior to horizontal down-sampling. Its multi-rate capabilities allow it to perform the entire decimation operation. Then a HSP48908 Two-Dimensional Convolver is used as a three-coefficient vertical filter to reduce vertical bandwidth prior to vertical down-sampling.

Row and Column Filters

In terms of algorithms, there is no basic difference between row and column digital filters. The implementation difference is that horizontal filters have delay elements of z^{-1} and vertical filters, while they have the same structure, are implemented with delay elements of z^{-L} , where L is the line length in pixels. The importance of this is that vertical line delays (z^{-L}) are more costly than horizontal sample delays, because lines are hundreds of pixels long. The result is that in low cost implementations, vertical filter functions tend to be very rudimentary (three coefficients or less) or nonexistent. Note that although vertical filters operate on pixels widely separated in time, they are still required to produce a new output at the horizontal pixel rate because new sets of inputs are presented at the horizontal rate.




FIGURE 11. GENERAL SEPARABLE FILTER BLOCK DIAGRAM




FIGURE 12. TOP LEVEL BLOCK DIAGRAM OF A RESIZER

Application Note 9603

Video Applications

A video A/D, for example the Harris HI5702, may sample the signal at twice the required sampling rate to ease analog anti-aliasing filter requirements. This would be followed by a halfband filter similar to the HSP43216 for decimation by two (Figure 13). To understand the usefulness of this operation, assume the required output rate is 13.5 MSPS (a CCIR 601 standard rate). With 4.5MHz input video bandwidth, the required anti-aliasing filter cutoff with a 13.5 MSPS sampling rate is 6.75MHz. This requires a filter shape factor of $6.75/4.5 = 1.5$. If 2:1 oversampling is followed by a digital halfband filter with a decimation rate of 2, the required shape factor of the anti-aliasing filter is $13.5/4.5 = 3.0$. Given the generally stringent passband requirements on video filters, this is a much more cost-effective solution than the more severe analog anti-aliasing filter. It allows the analog filter design to concentrate on passband performance rather than a sharp transition.

Another area that relies extensively on one-dimensional digital filtering techniques is NTSC or PAL decoding, which require operations like chroma/luma separation filters, quadrature down-conversion of the chroma information and chroma decimation by two filtering.

Conclusion

This application note has presented a brief overview of digital filtering and has highlighted the similarities and differences of filters implemented in software on general purpose digital microprocessors, in function-specific hardware, and in application specific hardware. Digital filter characteristics and performance have been compared to analog filters. Applications of high performance Harris function-specific digital filters in IF and Video signal processing were illustrated as examples of cost effective uses of digital filters with high throughput rates.




FIGURE 13. TOP LEVEL BLOCK DIAGRAM OF AN OVERSAMPLER TO REDUCE ANTI-ALIASING FILTER REQUIREMENTS

All Harris Semiconductor products are manufactured, assembled and tested under **ISO9000** quality systems certification.

Harris Semiconductor products are sold by description only. Harris Semiconductor reserves the right to make changes in circuit design and/or specifications at any time without notice. Accordingly, the reader is cautioned to verify that data sheets are current before placing orders. Information furnished by Harris is believed to be accurate and reliable. However, no responsibility is assumed by Harris or its subsidiaries for its use; nor for any infringements of patents or other rights of third parties which may result from its use. No license is granted by implication or otherwise under any patent or patent rights of Harris or its subsidiaries.

Sales Office Headquarters

For general information regarding Harris Semiconductor and its products, call **1-800-4-HARRIS**

NORTH AMERICA
Harris Semiconductor
P. O. Box 883, Mail Stop 53-210
Melbourne, FL 32902
TEL: 1-800-442-7747
(407) 729-4984
FAX: (407) 729-5321

EUROPE
Harris Semiconductor
Mercure Center
100, Rue de la Fusee
1130 Brussels, Belgium
TEL: (32) 2.724.2111
FAX: (32) 2.724.22.05

ASIA
Harris Semiconductor PTE Ltd.
No. 1 Tannery Road
Cencon 1, #09-01
Singapore 1334
TEL: (65) 748-2000
FAX: (65) 748-0400

