Instrumentation and Controls Division

A HARDWARE IMPLEMENTATION OF MULTIREOLUTION FILTERING FOR
BROADBAND INSTRUMENTATION*

Stephen W. Kercel and William B. Dress

RECEIVED
DEC 08 1995
OSTI

Paper to be presented at the
SPIE 1995 International Symposium
on Aerospace/Defense Sensing and Dual Use Photonics
Orlando, Florida
April 17-21, 1995

"The submitted manuscript has been authored by a contractor of the U.S. Government under contract No. DE-AC05-84OR21400. Accordingly, the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes."

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

*Research sponsored by the U.S. Nuclear Regulatory Commission under Interagency Agreement No. DOE 1866-8195-1L and performed at Oak Ridge National Laboratory, managed by Martin Marietta Energy Systems, Inc., for the U.S. Department of Energy under contract DE-AC05-84OR21400.

MASTER

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

et

A hardware implementation of multiresolution filtering for broadband instrumentation

Stephen W. Kercel and William B. Dress

Oak Ridge National Laboratory
Instrumentation and Controls Division
P.O. Box 2008, Oak Ridge, Tennessee 37831-6318

ABSTRACT

The authors have constructed a wavelet processing board that implements a 14-level wavelet transform. The board uses a high-speed analog-to-digital (A/D) converter, a hardware queue, and five fixed-point digital signal processing (DSP) chips in a parallel pipeline architecture. All five processors are independently programmable.

The board is designed as a general purpose engine for instrumentation applications requiring near real-time wavelet processing or multiscale filtering. The present application is the processing engine of a magnetic field monitor that covers 305 Hz through 5 MHz. The monitor is used for the detection of peak values of magnetic fields in nuclear power plants.

This paper describes the design, development, simulation, and testing of the system. Specific issues include the conditioning of real-world signals for wavelet processing, practical trade-offs between queue length and filter length, selection of filter coefficients, simulation of a 14-octave filter bank, and limitations imposed by a fixed-point processor. Test results from the completed wavelet board are included.

Keywords: multirate, multiresolution, wavelet, filter bank, magnetic field measurement, real time, environmental sensing

1. BACKGROUND

The problem being addressed in this work is the measurement of broadband electromagnetic and radio-frequency interference (EMI/RFI). These are known to cause upsets and malfunctions in safety related instrumentation and control (I&C) systems.^{1,2} EMI/RFI becomes a particular concern as industrial I&C systems designers begin to use high-performance digital circuitry.³

The nature of progress in digital circuitry is illustrated in Fig. 1. As the digital state of the art progresses, clock rates become faster and logic levels become lower. In practical terms, the more modern the digital system, the smaller is the logic pulse, in both the time and voltage dimensions, and, consequently, the greater the likelihood that a pulse can be corrupted by EMI/RFI. To minimize the effects of EMI/RFI, good electromagnetic compatibility (EMC) design and installation procedures must be used.

However, what is good EMC design? A good design is one that is sufficient to withstand typical ambient levels. Unfortunately, typical EMI ambient levels in industrial settings are unknown. They have never been measured. Suitable equipment to make this type of measurement is not even commercially available.⁴

This problem is of particular concern in the nuclear power industry. To assist the process of establishing good EMC engineering practices in nuclear power plants, the U.S. Nuclear Regulatory Commission (NRC) has concluded that acceptance criteria are needed.⁵ The Oak Ridge National Laboratory (ORNL) is assisting NRC to establish a technical basis for acceptance criteria. Since long-term observations of EMI ambient effects have never been conducted in nuclear power plants, a major element of this assistance consists of extended on-site surveys to determine ambient EMI/RFI levels.

Since the equipment to make this type of measurement is not available commercially, ORNL had to develop it before the ambient surveys could be conducted. Two instruments were developed.⁶ One is for the measurement of electric fields from 5 MHz to 8 GHz and does not use advanced signal processing techniques. The other is for the measurement of magnetic fields from 305 Hz to 5 MHz and includes a near real-time implementation of the wavelet transform.

Fig. 1. EMI disruption of digital signals.

The magnetic spectral receiver is intended for automatic, unattended, long-term (several weeks or months) recording of EMI/RFI levels at industrial plant sites. It is unattended in the sense that no human operator is required while the receiver is making its observations. The receiver is housed in an unadorned gray box with no external displays and no external controls except a key lock "off/on" switch whose key can be removed in either setting. The receiver is shown in Fig. 2. The video monitor and keyboard are used only when the receiver is being set up or shut down.

The output of the receiver is a two-dimensional matrixed histogram. One dimension consists of octave-wide bins of frequency. The other dimension consists of octave-wide bins of peak field strength. The entries in the histogram matrix are the running totals of the number of instances for which the ambient EMI/RFI level falls within the intersection of a strength bin and a frequency bin. The receiver does a wavelet analysis of the entire 14-octave spectrum about every half-second. The histograms and their time tags are stored to disk every 6 hours. Thus the counts for any observation period can be determined by subtracting the running total counts through the previous period from the running total counts through the period of interest.

2. THE DYADIC TREE STRUCTURE

The heart of the receiver is a system of 26 finite impulse response (FIR) filters in a dyadic tree, as shown in Fig. 3. The tree has a single input, a band-limited broadband signal spanning the range of 305 Hz to 5 MHz. The tree has a set of 14 outputs. Assuming that each filter is a half-band filter, each output can be interpreted at an octave-wide band extracted from the original input.⁷

It is reasonable to ask, "Why bother with a dyadic tree?" If the idea is to split the 14-octave-wide signal into its 14 constituent octaves, why not simply feed the signal into a parallel bank of 14 different bandpass filters? There are two answers. First, the dyadic

Fig. 2. Photo of magnetic spectral receiver.

tree provides the same level of performance as a parallel filter bank, but with significantly fewer multiplications. Second, the filters in the dyadic tree are much easier to implement than a conventional bandpass filter.

Fig. 3. Dyadic tree filter structure.

A FIR filter, $H_0(z)$, followed by a decimate by n operation, followed by another FIR filter, $H_1(z)$, can be replaced by an "equivalent circuit" consisting of a FIR filter, $H_0(z)H_1(z^n)$, followed by a decimate by n operation.⁸ If the high-pass and low-pass filters are each L coefficients long and if n is 2, the operation through both requires $2L$ multiplications. However, the "equivalent circuit" requires $3L - 2$ multiplications.

Thus, if a low pass followed by a decimate followed by a high pass is replaced by a bandpass filter followed by a decimate and if the two circuits have exactly the same transfer function, the bandpass filter is more costly to implement. For example, assume that the low-pass and high-pass filters each have 16 coefficients. The "equivalent circuit" bandpass filter, to generate the second output from the bottom, has 46 coefficients.

The same principle can be applied twice at the next higher level, and the "equivalent circuit" would require 106 coefficients. The "equivalent circuit" at the lowest frequencies would require tens of thousands of coefficients. Using 16 coefficients for every filter, the dyadic tree in Fig. 3 can be implemented with 416 multiplications. An equivalent parallel filter bank would require over a quarter of a million multiplications.

The second reason that the dyadic tree is superior to a parallel filter bank is that it can be implemented on hardware that operates over a lower dynamic range. In this context the dynamic range is defined by the ratio of the largest coefficient to the smallest coefficient for a given filter. In the filters of the dyadic tree, all the transitions occur near half the Nyquist frequency, and the filters require a relatively low dynamic range. The practical consequence is that the filters can be implemented on 16-bit fixed-point hardware, and the resulting quantizing error is tolerable.

On the other hand, a filter whose transitions occur many octaves from the Nyquist frequency requires a large dynamic range. The filter coefficients must be expressed with extreme precision. Consequently, the implementing hardware must be able to represent this extreme precision. It needs many more than 16 bits to operate successfully.

Thus, from an engineering standpoint, for multioctave analysis, the dyadic tree is clearly superior to a parallel filter bank. It requires far fewer multiplications, and each of those multiplications can be performed at relatively low precision. The other overpowering advantage of the dyadic tree is that it can be easily adapted to other operations.

Consider the subband coder scheme shown in Fig. 4.⁹ H_0 is a half-band low-pass filter and is followed by a decimate by 2 operation. H_1 is a half-band high-pass filter and is followed by a decimate by 2 operation. Together, these four operations constitute the analysis stage. The output consists of the upper octave in the high-pass channel and everything else in the other channel. Although there are two channels, neither carries redundant information.

Fig. 4. Perfect reconstruction subband coder.

The other four operations constitute the synthesis stage. The interpolators insert a zero between each of the data points, and the interpolating filters G_0 and G_1 produce the missing values. By summing the two outputs, an estimate of the input signal is produced. If the filters meet the paraunitary constraint, then the subband coder is said to have the attribute of perfect reconstruction (PR). In practical terms, a PR filter bank produces an output that is a scaled and delayed version of the input, with the wave shape perfectly preserved.

A consequence of the paraunitary constraint is to impose the following relationship on the impulse responses of the four filters.

$$\begin{aligned} h_1(n) &= (-1)^{(n+1)} h_0(N-1-n) \\ g_0(n) &= h_0(N-1-n) \\ g_1(n) &= (-1)^n h_0(n) \end{aligned} \quad (1)$$

In these formulas, N is the length of the filter and n is the position from 0 to $N-1$. These are Eq. (3.141) from Akansu and Haddad.¹⁰

If a filter bank is paraunitary PR and if $H_0(-1) = 0$, where H_0 is the Z-transform of h_0 , then the analysis bank constitutes a two-level wavelet transformer.¹¹ That is, the list of numbers emerging from the high-pass output of the analysis bank is the list of details or a list of the weights which, when multiplied by the translates of the basis wavelet at its finest scale, gives the projection of the original signal onto that wavelet scale.

The list of numbers emerging from the low-pass output is the list of approximations. The approximations can be processed by an identical analysis bank whose high-pass output becomes the next finest set of details and whose low-pass output becomes the next coarsest approximation. The process can be repeated as many times as needed. Thus the dyadic tree is a generalization of the analysis bank, producing wavelet coefficients at as many scales as necessary.

The synthesis bank of the two-band PR subband coder acts as an inverse wavelet transformer. That is, assuming $H_0(-1) = 0$, it takes the two lists of wavelet coefficients generated by the analysis bank and reconstructs the original signal. Again, the synthesis bank can be generalized to as many levels as needed.

Thus, assuming a wavelet filter design, such as the Daubechies extremal phase wavelet,¹² the 14 outputs of the filter bank can be interpreted as the lists of wavelet coefficients. Each level represents a different scale of the basis wavelet. Within each scale, each coefficient represents the contribution to the overall signal of the basis wavelet at that particular offset in time and that particular scale.

The dyadic tree as implemented in the magnetic spectral receiver is a near real-time wavelet transformer. Depending on how the wavelet coefficients are processed, the device can be used as the heart of any kind of system that requires wavelet analysis. For example, wavelet spaces are often very good pattern recognition feature spaces,¹³ and this kind of system can be adapted as a pattern recognition engine.

3. QUEUE AND FILTER LENGTH

Having settled on a receiver operating paradigm in which an A/D converter quickly fills a hardware queue and a dyadic tree of 26 filters empties the queue a little less quickly, two more questions arise. How long should the queue be? How long should the filters be? It turns out that the lengths depend on each other.

The dominant consideration is the problem of end effects. In practice, a decimating filter of length N operates as follows. The filtering process holds a queue vector of N input data and another vector with the N filter coefficients. The inner product of the two vectors is computed and becomes the next output datum. The two data at the tail of the queue are discarded, and the next two input data are pushed to the head of the queue. Then the process repeats.

There is a problem about how the process starts and ends. The crudest and riskiest method is to use zero padding. Assume that when the first two data are pushed to the input data queue vector, all other elements are zero. Likewise, at the end of the process, when there are no more input data, push zeros to the head of the queue to carry on the next step. The problem is that the resulting wavelet transform is the transform of the input data and its zero pads. Assuming that the data at either end of the input list are far from zero, this can lead to extremely large values being reported in the details when they are not really part of the details. Other end padding tricks include extending the last datum at each end of the list or assuming that the end pads are mirror images of the N data at each end of the list. All these methods make an assumption about the data that may or may not be true and can produce to misleading results.

We adopted another strategy in designing the magnetic spectral receiver. Instead of assuming what data the end pads might contain, simply let the input data queue vector fill up before taking the first inner product. When the last two input data are pushed to the head of the queue, compute the inner product with the full queue and then stop. The cost of this strategy is that some of the wavelet information implicit in the input data is discarded. In this application, based on our previous experience, we considered the strategy to be less risky than making assumptions about the end pads.

Given this assumption, the relationship between the filter length and the minimum allowable queue length is as follows. The very last filters in the dyadic tree must output two data points each. The last point requires L input data, where L is the length of the filter. N additional output data points require $2N$ additional input points. Any given filter must supply the number of points needed by its successor. For a dyadic tree of any length, the minimum number of data points that must be fed into the input is given by the following recursive relationship:

$$N_{in}(L, \text{pos}, N_{out}) = L + 2N_{out}, \quad \text{pos} = 1, \quad (2)$$

$$N_{in}(L, \text{pos}, N_{out}) = L + 2N_{in}(L, \text{pos} - 1, N_{out}), \quad \text{pos} > 1.$$

In this relationship, N_{in} is the number of data points required at the input of a dyadic tree consisting of pos filter pairs. L is the filter length. N_{out} is the number of data points required out of the last filter in the tree.

In this design, pos is 13 (see Fig. 3). N_{out} is 2. Assuming that L is 16, the resulting queue length is 131,058 data points. A queue made up of four 8-bit 32-K memory chips has a capacity of 131,072. Because of addressing problems, it was difficult to implement a longer queue with the hardware that was available to us. Thus, we settled on a 128-K hardware queue. In general,

the more coefficients, the better the filter; but this hardware constraint imposes a strict upper limit of 16 for the filter length and uses virtually all the physical capacity of the queue.

For prototyping purposes, we did not want to use a design that required virtually all the physical capacity of the queue. To give some margin, we decided to use 8 coefficient filters for the last two filter pairs in the tree. With $\text{pos} = 2$, $L = 8$, and $N_{\text{out}} = 2$ and using Eq. (2), N_{in} , the number of data points that must be fed into the next-to-last filter pair is 26. The other 11 filter pairs use 16 coefficient filters and must produce 26 output data points. With $\text{pos} = 11$, $L = 16$, and $N_{\text{out}} = 26$, N_{in} , the minimum number of input data points, is 81,906.

For the prototype, the design was slightly conservative. We used 11 pairs of 16 coefficient filters, 2 pairs of 8 coefficient filters, and a hardware queue with a capacity of 131,072. In a future modification, we intend to improve the performance at the lowest frequency octaves by using 16 coefficient filters for all pairs and using the full capacity of the queue.

4. PERFECT RECONSTRUCTION FILTER DESIGN

Having settled on filter lengths, the next question is what filter design to use. That is, what are the filter coefficients? Because of the paraunitary constraint, we need find only the coefficients of the low-pass filter of the analysis stage. Any of the other three filters of a PR filter bank can be obtained by Eq. (1).

There are many possible FIR filter designs. However, two of particular interest arise from PR filter bank theory. One is the Daubechies filter, which both meets the PR constraint and constitutes a wavelet.¹⁴ The other is the cascaded lattice design,¹⁵ which meets the PR constraint but does not constitute a wavelet.

Even if the FIR filter designer were not interested in the special attributes of wavelets (simultaneous time-frequency localization, etc.), the Daubechies filter would still be of interest. It is the theoretical maximally flat FIR digital filter,¹⁶ just as the Butterworth design is the theoretical maximally flat analog filter.¹⁷ Where L is filter length, the Daubechies filter can be implemented with $L/2$ multiplications,¹⁸ whereas other FIR filters, including most wavelet types, require L multiplications.

In addition, the Daubechies filter is not especially difficult to design. Akansu and Haddad note that the Daubechies filter is a weighted sum of binomial filters and provide straightforward equations for computing filter coefficients. They offer a recursive method for obtaining the binomial filter coefficients in Eqs. (2.132) and (2.138).^{18,19} Equations (4.3) and (4.8) provide a system of $L/2$ quadratic equations in terms of the binomial coefficients and the $L/2$ unknown weights. These can be solved for the weights, which can be fed into Eq. (4.1) to obtain the filter coefficients.

The whole process can be implemented in about 12 lines of Mathematica code.²⁰ The computed results agree exactly with Table 6.1 in Daubechies.¹² It is fair to note that for more than eight weights, the system of quadratic equations has hundreds or thousands of possible solutions, and the process of finding the minimal phase solution becomes unwieldy. The relevant Daubechies filters are listed in Table 1 as ep4 (4 weights, 8 coefficients) and ep8 (8 weights, 16 coefficients). For each entry, $h_0(0)$ is at the top of the list. The entries in Table 1 are normalized such that their sum is 1 instead of $2^{1/2}$, as in Daubechies.

The cascaded lattice is another design of particular interest to the receiver designer. In this design, the analysis stage of the PR filter bank, with filter length L , is represented as a cascade of $L/2$ lattices. One numerical weight completely characterizes each lattice. By specifying a desired cutoff frequency and constraining the system to minimize the energy at frequencies beyond cutoff, Vaidyanathan and Hoang have listed weights for various filter specifications in Table II and III of their paper.¹⁵ Using Eq. (3.191) in Akansu and Haddad, the polyphase matrix of the analysis stage can be easily computed from the published weights. Equation (3.79) in Akansu and Haddad recovers the filter transfer functions from the polyphase matrix.

This process is easily implemented in Mathematica. Coefficients for the cascaded lattice filters of interest in the design of the magnetic spectral receiver were computed from the published weights by using a simple Mathematica application. They are listed in Table 1. The designations 8a, 16b, 16c, and 16f correspond to table identifiers from Vaidyanathan and Hoang.¹⁵

The cascaded lattice filter has several interesting attributes. It remains PR under coefficient quantization. This is very important when trying to implement a filter on fixed-point hardware. Compared to other filter designs with the same cutoff frequency, it has superior out-of-band rejection. This is very desirable in a receiver design if strong out-of-band signals are present.

Table 1. PR Filter Coefficients

Daubechies wavelet		Vaidyanathan and Hoang			
ep4	ep8	8a	16b	16c	16f
<i>8 Coefficient filters</i>					
0.1629017140256491		-0.01257799389099265			
0.5054728575459143		0.03318107491227977			
0.4461000691304508		0.02508363218573078			
-0.01978751311782235		-0.1377949837464604			
-0.1322535836845199		-0.03568162902405254			
0.02180815023708858		0.4114734923427839			
0.02325180053549088		0.519398871045011			
-0.007493494665180714		0.196889216044501			
<i>16 Coefficient filters</i>					
	0.03847781105407849	-0.000853605457568049	-0.003353511229926781	-0.004837871318288043	
	0.2212336235761366	0.003317412984441417	0.00994820447762271	0.01325438378091221	
	0.4777430752138947	-0.00155598953123399	-0.003664862130719162	-0.004412277826112802	
	0.4139082662112024	-0.01070619447921235	-0.01980737814115445	-0.02334901944929056	
	-0.01119286766690355	0.01262557445102279	0.01660654731998661	0.01748015717086685	
	-0.2008293163905141	0.02171171160310362	0.03309154090202154	0.03679509552296786	
	0.000334097046219494	-0.03902440517691352	-0.03844828094070893	-0.03724779112556963	
	0.0910381784236654	-0.03490052963330355	-0.05195070917668688	-0.05647450704162396	
	-0.01228195052284944	0.0865871163499388	0.0713990975159501	0.06537312344433804	
	-0.03117510332514108	0.05537125487093934	0.0849925033819879	0.0912730518255076	
	0.00988607964835204	-0.1623349179383235	-0.1172517460760082	-0.102371201070543	
	0.006184422409816434	-0.1227156094055434	-0.1664542183056222	-0.1733527334064618	
	-0.00344385962844211	0.2703386639051493	0.1608182451332496	0.1293032034419242	
	-0.000277002274479426	0.5028753106205964	0.4784068284636719	0.4645868171175648	
	0.0004776148556496479	0.3334602191819811	0.4080482657535146	0.4277066012802498	
	-0.000083068630686634	0.085802842248025	0.1375519014144307	0.1561135948059724	

Remarkably, the cascaded lattice filter does not constitute a wavelet. Orthonormal wavelets with compact support must meet two conditions. The generating filter bank must be PR, and the low-pass filter transfer function must have a zero at $z = -1$ ($\omega = \pi$).¹¹ The magnitude of the stop-band ripple of the cascaded lattice filter has a nonzero local maximum at $\omega = \pi$. While this is an excellent filter for receiver applications, it may not be a good choice for pattern recognizers or other devices whose operating principle depends on the properties of the wavelet transform.

Comparative characteristics of the Daubechies and cascaded lattice filters are shown in Figs. 5 and 6. The Daubechies filter does not cut off as sharply as any of the cascaded lattice designs of the same length, but it does not have ripple in either the stopband or the passband. All the cascaded lattice designs have passband ripple, but it is so small as to be unresolvable on a full-scale plot. All the cascaded lattice designs have stopband ripple whose severity increases as the cutoff becomes sharper. The peaks of the stopband ripples of the 16c and 16f designs are on the order of 30 dB down and are very objectionable for receiver design. The 16b design appears to offer the best of both worlds. It has a significantly sharper cutoff than the Daubechies filter, but the stopband ripple is below 50 dB down, quite acceptable for a receiver with 36-dB dynamic range.

It is important to realize that Figs. 5 and 6 do not tell the whole story. The graphs in Figs. 5 and 6 were computed with extremely high floating point precision and do not reflect quantizing effects. A more practical comparison of the different filter designs in their intended application is shown in Fig. 7. In Fig. 7, the filter consists of the first low-pass decimated filter followed by the second high-pass decimated filter in the dyadic tree of Fig. 3. The horizontal axis is frequency in hertz. The vertical axis

Fig. 5. Eight-coefficient PR low-pass filters.

Fig. 6. Sixteen-coefficient PR low-pass filters.

Fig. 7. Floating point versus quantized performance.

is insertion loss in decibels. The floating point calculations are done at the same level of precision as those for Figs. 5 and 6. For the quantized calculations, filter coefficients are quantized to 2^{16} discrete levels, and the input signal is quantized at 2^8 discrete levels.

The quantizing effects are significant. The quantized Daubechies filters retain their smooth performance until they reach about 65 dB of stopband rejection. In the regions where the floating point Daubechies filters produce extremely small signals, the quantized version produces noise. The noise may or may not be acceptable, depending on the intended application. The quantized 16c and 16f filters have such big ripples that the quantizing noise is not especially significant. The quantized 16b design produces about the same amount of noise in the 65-dB rejection region as the Daubechies filter, but it has a sharper cutoff and its stopband rejection is generally below 50 dB.

The magnetic spectral receiver was originally developed using the Daubechies filters. The out-of-band signal rejection was acceptable, and the receiver was deployed to meet a tight schedule. Following deployment, we conducted the investigation that led to the data plotted in Fig. 7. These data suggest that by using the 16b filter design, we should see a small overall improvement in the receiver's performance. At some future date, when we have the receiver out of service for maintenance, we will modify the DSP filter programs to use the 16b design.

5. IMPLEMENTATION

The multiresolution filter is configured as a receiver to observe low-frequency magnetic fields. Frequency coverage is from 305 Hz to 5 MHz. The conceptual diagram is shown in Fig. 8, and a photograph of the processing electronics is shown in Fig. 9. The receiver captures the peak magnetic field strength in each of the 14 octaves during each sampling cycle.

The magnetic spectral receiver repeatedly cycles through the 14 octaves about every half-second. The queue captures 128-K samples at a sampling rate of 10 million samples per second, thus filling up in 13,107 μ s. At a processing duty cycle of about 3%, it takes the five DSP chips a little more than 30 times as long to extract the information from the data in the queue as it does to collect it. Thus, the whole process will take about 30 times 13,107 μ s, or 393,000 μ s. At the conclusion of the processing, the histogram is updated, and the cycle repeats. Any signal that occurs while the queue is filling will be captured. Since the queue fills for only 13,107 μ s out of each half-second cycle, only signals lasting more than a half-second are certain to be captured.

The magnetic spectral receiver uses a minimal hardware configuration, consisting of an antenna, a preprocessing filter, broadband amplifier, A/D converter, hardware queue, multirate filter bank, and controller. Because of the extreme bandwidth (14 octaves), the desire to minimize hardware complexity, and the desire to maximize the amount of time each octave is observed, signal analysis is done digitally.

The analog part of the system operates as follows. The receiver uses a passive loop antenna as a broadband magnetic field probe. The antenna produces a current proportional to the incident magnetic field in the band of 20 Hz to 5 MHz. The current is forced through a 50- Ω load, producing a voltage proportional to the current. The voltage across the 50- Ω load is typically in the range of 20 mV peak to peak.

Between the antenna and the 50- Ω load is a set of passive filters. A 7-element analog (inductive-capacitive) Butterworth high-pass filter is used to suppress the power frequency, its third harmonic, and any other undesired signals below 305 Hz. An analog low-pass filter is used to prevent aliasing of signals above 5 MHz.

A system of video operational amplifiers provides signal conditioning. The antenna delivers a signal to the load that swings between -10 mV and +10 mV. A two-stage amplifier with a gain of 20 dB per stage increases voltage swing to the range of ± 1 V. A unity gain stable differential amplifier changes the voltage swing to the range of 0 to -2 V. This is the level needed at the input of the A/D converter.

Eight-bit A/D conversion is done with a self-contained (A/D) chip. While this imposes an upper limit of about 36 dB on the useful dynamic range, it is sufficient for this application. A sampling rate of 10 million samples per second provides the Nyquist limit at the highest input frequency of 5 MHz.

Fig. 8. Block diagram of magnetic spectral receiver.

Fig. 9. Photo of magnetic receiver electronics.

The processing is done on a sample-and-hold basis. At the start of the operational cycle, the queue controller enables the A/D converter, selects the queue addresses, and allows the queue to fill. When the queue is full, the queue controller disables the A/D converter and sends a "queue is full" signal to DSP0.

The five DSP chips are used in a parallel pipeline structure to implement the dyadic tree filter bank. Figure 3 indicates which filter elements are implemented on each chip. DSP0 first obtains 16 data points from the queue. From then forward, it does the high- and low-pass filter multiplications, sends the resulting high-pass value to the peak detector, sends the low-pass value to

the next filter pair, discards the oldest two signal data points, and pops the next two values from the queue. It repeats the process until the queue is empty. Finally, it sends the upper octave peak value down the pipeline.

Since the dyadic tree both filters and decimates, each succeeding high-pass, low-pass pair gets half as many data as the one before. Nevertheless, all DSP chips operate at the same clock rate of 40 MHz. Thus, each chip down the pipeline is able to implement twice as many filters as the one before. In the original design, we intended to implement a single low-pass filter on DSP0 and a single high-pass filter on DSP1 and to have DSP0 pass on the first-level approximation data to DSP2 via the parallel data bus. DSP2 would implement both filters for the next stage of the tree, and the succeeding chips would each implement twice as many filters as the one before.

Unfortunately, we ran into an unexpected contention problem when the parallel bus control lines for DSP0, DSP1, and DSP2 were tied together. Although this can be overcome in a future modification of the receiver, we were facing a tight deadline for deployment. Thus, we decided to take a 50% reduction in overall throughput and settle for a 3% duty cycle rather than the hoped-for 6% duty cycle, remove DSP1 from the circuit, and implement both a high-pass and a low-pass filter on DSP0.

The other DSP chips implement the rest of the dyadic tree. DSP2 obtains the peak value for the next highest octave, strips out the bottom 12 octaves (the next approximation), and feeds the resulting signal to DSP3 via a serial bus. DSP3 obtains the peak value for the two next highest octaves, strips out the bottom 10 octaves, and feeds the filtered signal to DSP4 via a parallel bus. DSP4 obtains the peak value for the four next highest octaves, strips out the bottom six octaves, and feeds the filtered signal to DSP5 via a serial bus. DSP5 obtains the peak value for the last six octaves and feeds all the peak readings to the controller via a parallel bus.

Several practical issues are involved in the design of the digital processing circuit. For DSP2 through DSP5, alternating serial and parallel communication links preclude the problem of contention for communications resources. The board was designed to operate at a clock rate of 80 MHz. In the prototype, the 80-MHz clock chip would not produce enough power to feed five DSP chips and the frequency divider for the A/D converter. The system worked without difficulty at 40 MHz. In a future modification, an 80-MHz clock and amplifier will be used. Doubling the clock rate would double the throughput of the system. There are no read-only memory chips fast enough to use with the DSP chips. Consequently, the programs for each DSP chip are stored on battery-backed static random access memory chips.

6. THEORY VERSUS EXPERIMENT

To report actual-versus-predicted performance for all 14 octaves would become rather tedious. Since all the octaves perform fairly much the same, it is sufficient to look at a typical octave. In the following data, the fourth octave counting from the highest frequency is considered. What is actually being observed is the output from the fourth high-pass filter in the dyadic tree. The passband is 312,500–625,000 Hz.

The experimental data were obtained by feeding a signal of known voltage and frequency into the receiver input and observing the output voltage for the fourth octave. The input signal was set such that a voltage of 0.5 V peak to peak appears at the A/D converter input. Actually, this is translated such that -1 ± 0.25 V appears on the A/D converter input pin. The A/D converter encodes the peak value as 32. The input frequency is swept from 156,250 Hz to 850,494 Hz in increments of 1/9 octave. The (dB) listing in Table 2 is based on the ratio of the observed output value and the known input value of 32.

The Daubechies prediction is based on a quantized version of the Daubechies coefficients. The input sinusoid is given a peak value of 32 and quantized at 2^8 discrete levels. The Daubechies coefficients for 16 coefficient low-pass and high-pass filters are quantized at 2^{16} discrete levels. The sinusoid is run through the low-pass filter, followed by the high-pass filter. The data list emerging from the high-pass filter is fitted to a function of the form $A\cos(2\pi fx) + B\sin(2\pi fx)$, and the peak value is obtained. Table 2 shows these data as decibels of magnitude compared to 32.

We note that there is relatively close agreement between observed performance and the Daubechies prediction. There are some small variations in the neighborhood of 200 kHz which can be attributed to quantizing uncertainty and slight nonlinearity in the A/D converter.

The cascade prediction is based on a quantized version of the cascade coefficients for the 16b design. The input sinusoid is given a peak value of 32 and quantized at 2^8 discrete levels. The cascade coefficients for 16b low-pass and high-pass filters are quantized at 2^{16} discrete levels. The sinusoid is run through the low-pass filter, followed by the high-pass filter. The data list emerging from the high-pass filter is fitted to a function of the form $ACos(2\pi fx) + B\sin(2\pi fx)$, and the peak value is obtained. Table 2 shows these data as decibels of magnitude compared to 32.

The cascade prediction is a prediction of how the receiver should operate using the cascade filter coefficients instead of the Daubechies coefficients. Inside the passband, the two filters are nearly identical. Outside the passband, the cascade filter shows superior rejection.

Table 2. Filter Performance at 1/9-Octave Frequency Intervals

Frequency (Hz)	Daubechies Prediction (dB)	Observed (dB)	Cascade Prediction (dB)
156,250	-32.3	-30.1	-48.1
168,759	-28.6	-24.1	-54
182,270	-24.3	-24.1	-51
196,863	-20.4	-20.6	-34.6
212,623	-16.6	-14.5	-25.5
229,646	-13.2	-11	-18.7
248,031	-10	-9.3	-13.3
267,889	-7.3	-7.2	-8.9
289,336	-4.9	-5	-5.5
312,500	-3	-3.3	-3
337,519	-1.6	-1.8	-1.3
364,540	-0.7	-0.9	-0.4
393,725	-0.3	-0.6	0
425,247	-0.1	-0.3	0
459,292	-0.2	-0.6	0
496,063	-0.5	-0.6	-0.2
535,777	-0.9	-1.2	-0.6
578,672	-1.7	-1.8	-1.4
625,000	-3	-3.3	-3
675,037	-5.1	-5	-5.8
729,081	-8.2	-8.5	-10.4
787,451	-12.9	-13.2	-18.2
850,494	-19.8	-20.6	-32.5

7. CONCLUSIONS

The major conclusion is that the receiver works and is performing its intended function. It has produced 4 months of data at two nuclear plant sites. It will be deployed at three other sites over the next 6 months. These data, along with the data from the electric spectral receiver, will make up the bulk of the data for our survey of nuclear power plant EMI/RFI ambient levels.

I&C engineers in other industries are also concerned with EMI/RFI effects. After the nuclear site surveys are finished, we expect to do similar surveys for other industries. Both semiconductor manufacturers and textile loom operators have expressed an interest.

The receiver as currently implemented is a first prototype. There are several obvious modifications that would improve its performance. An 80-MHz clock with sufficient output power would double the throughput of the dyadic tree. Overcoming the

bus contention problem, enabling the implementation of one filter each on DSP0 and DSP1, would also double the throughput. In addition, changing from Daubechies coefficients to cascade coefficients would slightly improve the out-of-band signal rejection of each filter.

The receiver design exploits wavelet theory in two respects. The dyadic tree structure (or fast wavelet transformer) is a highly efficient structure for implementing a multiband filter bank. It requires significantly fewer multiplications than other structures. In addition, the PR filter bank (necessary, but not sufficient for implementing a fast wavelet) leads to the cascaded lattice and Daubechies filter designs. For a given filter length, the cascaded lattice provides optimal rejection of undesired signals and is highly desirable for some receiver applications, including this one. The Daubechies filter is the maximally flat FIR filter and implements an orthonormal compactly supported wavelet.

In its present configuration, the receiver implements a slightly noisy 14-level Daubechies wavelet transform. With a suitable change of transducer and modifications to the DSP programs, it could be used for many applications. In particular, it could be used as a real-time pattern recognition device.

8. ACKNOWLEDGMENTS

Research sponsored by the U.S. Nuclear Regulatory Commission under Interagency Agreement No. DOE 1866-8195-1L and performed at Oak Ridge National Laboratory, managed by Martin Marietta Energy Systems, Inc., for the U.S. Department of Energy under contract DE-AC05-84OR21400.

9. REFERENCES

1. J. Cirillo and M. Prussel, "Electromagnetic Compatibility in Nuclear Power Plants," presented at the WATtec Conference, Knoxville, Tennessee, February 11-14, 1986.
2. P. D. Ewing and K. Korsah, *Technical Basis for Evaluating Electromagnetic and Radio Frequency Interference in Safety Related I&C Systems*, NUREG/CR-5941, April 1994, available from National Technical Information Service, Springfield, Va.
3. P. D. Ewing, K. Korsah, and C. Antonescu, "Technical Basis for Acceptance Criteria on the Susceptibility of Digital Systems to Electromagnetic Interference," presented at the 20th Water Reactor Safety Information Meeting, Rockville Md., October 21-23, 1992.
4. S. W. Kercel, *Comparison of the Magnetic Spectral Receiver with Conventional Technologies*, ORNL/TM-12916, Oak Ridge Natl. Lab., March 1995, available from National Technical Information Service, Springfield, Va.
5. "Review of Vendors' Test Programs to Support the Design Certification of Passive Light Water Reactors," NRC Policy Issue SECY-91-273, Nuclear Regulatory Commission, Rockville Md., 1991.
6. S. W. Kercel, "A Confirmatory Research Approach to the Measurement of EMI/RFI in Commercial Nuclear Power Plants," presented at the 22nd Water Reactor Safety Information Meeting, Bethesda, Md., October 24-26, 1994.
7. A. N. Akansu and R. A. Haddad, *Multiresolution Signal Decomposition*, Academic Press, San Diego, 1992, p. 133.
8. Ref. 7, pp. 103-110.
9. Ref. 7, pp. 123-126.
10. Ref. 7, p. 159.
11. Ref. 7, pp. 321-326.
12. I. Daubechies, *Ten Lectures on Wavelets*, Society for Industrial and Applied Mathematics, Philadelphia, 1992, pp. 194-195.
13. W. B. Dress and S. W. Kercel, "Wavelet-Based Acoustic Recognition of Aircraft," in *Wavelet Applications*, Harold H. Szu, Ed., Proc. SPIE 2242, pp. 778-791, 1994.
14. Ref. 7, pp. 338-339.
15. P. P. Vaidyanathan and P. Q. Hoang, "Lattice Structures for Optimal Design and Robust Implementation of Two-Band Perfect Reconstruction QMF Banks," *IEEE Trans. Acoustics, Speech, Signal Process.* 36(1), 81-94 (January 1988).
16. Ref. 7, pp. 248-250.
17. L. C. Ludemann, *Fundamentals of Digital Signal Processing*, John Wiley and Sons, New York, 1986, p. 122.
18. Ref. 7, pp. 241-243.
19. Ref. 7, pp. 47-53.
20. Wolfram Research, Inc., *Mathematica*, Version 2.2, Wolfram Research, Inc., Champaign, Ill., 1992.