Precursors in Signal Processing

Jelena Kovačević
Carnegie Mellon University

I. Introduction

What a treat to look back and observe what happened within the field of wavelets over the past 20 years. We witnessed tremendous advances, wavelets became a commonplace technique, several standards are wavelet-based, wavelet software packages such as Matlab are used regularly. As icing on the cake, wavelets brought people from many disparate areas together. It is commonplace today to attend a wavelet meeting and sit in a room with mathematicians, engineers, statisticians and physicists.

Where did it all start? This volume gives the answer: Many places. One of those many is signal processing. If you are not familiar with signal processing, do not immediately start thinking of DSP chips. Signal processing is the mathematical framework for acquisition, representation, analysis of signals (and many other tasks). It is heavily motivated by applications and is (mostly) discrete-time oriented; your signal is a sequence of numbers representing a speech or audio signal, an EEG signal, a heart rate signal. It can also be a matrix of numbers representing illumination of individual picture elements (pixels), or even a three-dimensional volume representing moving images.

The beautiful collection of papers that follows is a glimpse into the history of the birth of wavelets. The beauty of these solutions comes from trying to solve real problems and ingenious ways of doing it. Looking back, you will see that the multiresolution ideas kept on coming up, showing that multiresolution is inherent within a certain class of problems.

So sit down, relax, and enjoy the trip to yesteryear...

The selection of papers in this chapter could be split into two uneven parts. The first is the work of Burt and Adelson, while the other contains the rest. The reason for this is that historically, the second set developed more or less in a sequence, and the papers were influenced by the previous ones. The work of Burt and Adelson was picked up later by the signal processing people but developed initially on its own and did not influence the early signal processing papers.

A small aside about notation: Discrete-time sequences are denoted by h_n, with $n \in \mathcal{Z}$. Their discrete-time Fourier transform is denoted by $H(e^{j\omega})$, a 2π-periodic function. Finally, for the nonDSP people: the z-transform is defined as $H(z) = \sum_{n\in\mathcal{Z}} h_n z^{-n}$ and reduces to the discrete-time Fourier transform on the unit circle, that is, for $z = e^{j\omega}$. If you are wondering where the z-transform came from, you can think of it as a counterpart of the Laplace transform in continuous time.

A. The Laplacian Pyramid as a Compact Image Code

The paper by Burt and Adelson designs an efficient image coding system based on the novel idea of pyramid coding. The key idea is to produce a prediction of the image by lowpass filtering and downsampling the original image and then calculate the difference between the original and the prediction based on that lowpass image. This produces an instant compression system since the difference image is of low energy which can be appropriately quantized. Further compression can be achieved by iterating the process on the prediction resulting in a "pyramid" of difference images and a final lowpass signal.

Contact author: Jelena Kovačević, `jelenak@cmu.edu`

If you are not completely wavelet-immune, you will, in the preceding paragraph, immediately recognize elements of wavelet theory. What is more amazing, though, is that Burt and Adelson, without the benefit of our 20 years of wavelet knowledge base, did the same – albeit not calling it wavelets. To illustrate this point, consider the following few quotations from their paper:

"It (pyramid) represents an image as a series of quasi-bandpassed images, each sampled at successively sparser densities. The resulting code elements, which form a self-similar structure, are localized in both space and spatial frequency." Moreover, "The scale of the Laplacian operator doubles from level to level of the pyramid, while the center frequency of the passband is reduced by an octave."

It is worth noting that, since the difference signals are not sampled, the scheme is oversampled (overcomplete) and thus corresponds to a wavelet frame rather than a wavelet basis.

B. Digital Coding of Speech in Sub-bands

This work by Crochiere, Webber and Flanagan truly starts the area of subband coding which later lead to connections with wavelets. What is also worth noting, is that the reason subbands are used is to allow perception to play a part in coding – a concept much used in compression starting with late the 1980s.

This is a precursor of a precursor. It does divide a speech signal in subbands, by bandpass filtering the original speech signal. These bandpass filters are narrow to allow for efficient quantization and to eliminate the effect that noise in one subband might have on another. So sampled bands are then quantized according to perceptual criteria.

In wavelet terms, this work would amount to doing a discrete wavelet transform with one level but not quite using the correct mother wavelet to allow the scheme to be inverted. Moreover, one might consider it undersampled since guard bands are allowed, amounting to parts of the spectrum not being reconstructed. In signal processing terms, aliasing – overlapping of the spectrum due to downsampling – is still present.

The quest for a perfectly reconstructed signal starts.

C. Application of Quadrature Mirror Filters to Split Band Voice Coding Schemes

This work by Esteban and Galand continues seamlessly along the lines of the previous article. It, quite naturally, addresses the problem of how to eliminate aliasing. They consider the two-channel filter bank (see Figure I-C), where the channels (subbands) are obtained by lowpass (H_0) and highpass (H_1) filtering the original signal followed by downsampling by two. The authors come up with an ingenious solution to remove aliasing, which became famous in the community as QMF (quadrature mirror filters). Many later solutions which would eventually construct bases (perfect reconstruction in signal processing terms), would still be called QMF. The deceptively simple solution is to form the highpass filter $H_1(z)$ as

$$H_1(z) \;=\; -H_0(-z).$$

The quest for a perfectly reconstructed signal continues...

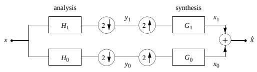

Fig. 1. Two-channel filter bank.

D. A Procedure for Designing Exact Reconstruction Filter Banks for Tree-Structured Subband Coders AND Filters for Distortion-Free Two-Band Multirate Filter Banks

Why the titles of two papers in this subsection? The first is by Smith and Barnwell, and the second by Mintzer. The reason is that these two papers solved the same problem independently. Namely, they both found the solution to having perfect reconstruction, two-channel filter banks. Moreover, their solution is orthogonal (corresponds to a unitary transform) and the filter design method is by spectral factorization. For you waveleters, spectral factorization is exactly what Ingrid Daubechies used in her famous paper to design "Orhthonormal bases of compactly supported wavelets" (reprinted later in this volume), except that the factors she chose were slightly different.

In short, the solution is to form the filter bank with the following filters: $H_0(z), H_1(z) = H_0(-z^{-1}), G_0(z) = H_0(z^{-1}), G_1(z) = H_0(-z)$. The condition for perfect reconstruction then boils down to

$$H_0(z)H_0(z^{-1}) + H_0(-z^{-1})H_0(-z) = P(z) + P(-z) = 1.$$

Thus, to solve for the filter $H(z)$, one has to find an appropriate $P(z)$ under very mild constraints, and then factor it into its spectral factors. Given the above, $P(z)$ can be factored into a product of reciprocally paired zeros. The distribution of zeros between $H(z)$ and $H(z^{-1})$ can be performed many ways. Smith and Barnwell choose one solution, while Daubechies is governed by the smoothness of the corresponding wavelets and chooses to put the maximum number of zeros at $z = -1$ into $H(z)$ and the rest inside the unit circle.

The quest for a perfectly reconstructed signal is over. Or, is it?

E. Filter Banks Allowing Perfect Reconstruction

As you might have guessed, humans have a unsatiable intellect, and as soon one solved how to obtain perfect reconstruction filter banks for two channels (dyadic case), the quest for doing it with the least harsh constraints on the filters (biorthogonal solution) and for the same solutions when there are more than two channels, started. In his work, Vetterli solves both problems. Moreover, he establishes for many years to come, the polyphase analysis as a standard for analyzing filter banks, thus transforming a linear, shift-variant system (because of downsampling) into a multi-input, multi-output linear shift-invariant system. The polyphase domain – the domain of the cosets of the regular lattice – becomes the preferred tool. It brings about the convenience and elegance of matrix notation and analysis into play. Finally, Vetterli also solved the dual problem – transmultiplexers, a communications application.

F. Theory And Design Of M-Channel Maximally Decimated Quadrature Mirror Filters With Arbitrary M, Having The Perfect Reconstruction Property

The final paper in this section introduces an important concept into the filter bank community – that of losslessness. The idea coming from circuit theory, deals with systems in which there is no loss of energy. In z-domain, these systems become paraunitary on the unit circle and thus correspond to orthonormal wavelet bases.

Moreover, Vaidyanathan proposes efficient schemes for factoring and implementing perfect-reconstruction filter banks. These are based on lossless building blocks – lattice structures.

G. In Conclusion

So how does this all tie in? Why do we say these papers are precursors to wavelets in signal processing?

Suppose we take the two-channel filter bank from Figure I-C, and iterate on the lowpass filter. That is, at the output of the lowpass branch we add another two-channel filter bank. We continue the process J times. What we get is the *discrete-time wavelet transform* with J levels. Unless you are involved in wavelet research, this is the only wavelet transform you will actually use. This is what people use in Matlab to implement the wavelet transform.

The discrete-time wavelet transform illustrates many of the concepts we mentioned until now. It is critically sampled, that is, the number of samples at any point in the system is the same. If the polyphase matrix of the initial analysis filter bank is invertible, we do have a transform – and we call it a biorthogonal transform (it is the most general solution to the problem). The filters used can have compact support (FIR filters) or infinite support (IIR) filters. If the polyphase matrix is paraunitary, our transform is orthonormal. In practical terms, this means that the norm (energy) is preserved throughout the system, and the synthesis filters are the same as the analysis ones (within reversal). The outputs of the last lowpass branch are scaling (coarse) coefficients, while the outputs of all the bandpass/highpass branches are wavelet (detail) coefficients. If we drew the equivalent filters through each branch of the system, we get wavelet filters, which are (discretely) streched versions of the same basic mother wavelet (highpass filter). If we increase the number of channels to N and iterate on the lowpass branch as before, we will obtain a wavelet transform with a stretching (dilation) factor of N. The wavelet filters are now N-times streched versions of the same mother wavelet. If we relax the constraint of critical sampling and sample with the number smaller then N, we get an overcomplete system (frame).

This was all discrete. What about the real – continuous-time wavelet bases? Well, if the lowpass filter is smooth enough and we iterate to infinity, we will obtain the continuous-time wavelet bases. This construction is the one Daubechies has in her paper. The other way around, getting the discrete-time version from the continuous-time one is always possible: just assign the coefficients in the two-scale equations for the scaling function and the wavelet to the lowpass and highpass filters, respectively.

I hope you will enjoy reading the historic papers in this section and the rest of the volume.