

# An Analog Wavelet Transform Chip

R. Timothy Edwards

Michael D. Godfrey\*

ISL, Dept. of Electrical Engineering  
Stanford University  
Stanford, California 94305

**Abstract**— This paper describes the theory and implementation in subthreshold analog CMOS technology of a circuit which performs continuous wavelet decompositions of a one-dimensional (e.g., audio) input. The analog wavelet outputs are the output of a logarithmically scaled bank of bandpass filters; each band is sampled at a rate proportional to the Nyquist rate of the highest frequency content of that band. The result is a matrix of discrete points describing the input signal as a function of both frequency and time. The filter function of each band is gaussian shaped in order to best resolve the uncertainty relation between time and frequency at each sampled point.

## I. INTRODUCTION

The authors have designed and built a chip which performs continuous wavelet transforms on its input using analog subthreshold CMOS circuits fabricated in a standard process.

The Continuous Wavelet Transform (CWT) performed by the chip is an analog filtering function similar to what is known as the Gabor spectrogram [1]. It is not to be confused with the Discrete Wavelet Transform (DWT), although the result of both transforms is a similar time-frequency description of a signal. That description is shown graphically in Fig. 1. The DWT produces this result by starting with a block of discrete data and performing successive high- and low-pass digital filtering. The filtering is repeated on the low-pass output in order to band-pass the signal in a series of stages called “dilations.” Each dilation divides the frequency space of the current interval in half while doubling the time span, thus keeping the time-frequency product constant. The CWT divides a signal into a set of logarithmically-scaled frequency bands by passing it through a bank of bandpass filters. The simplest and most natural form of the bandpass filter is a

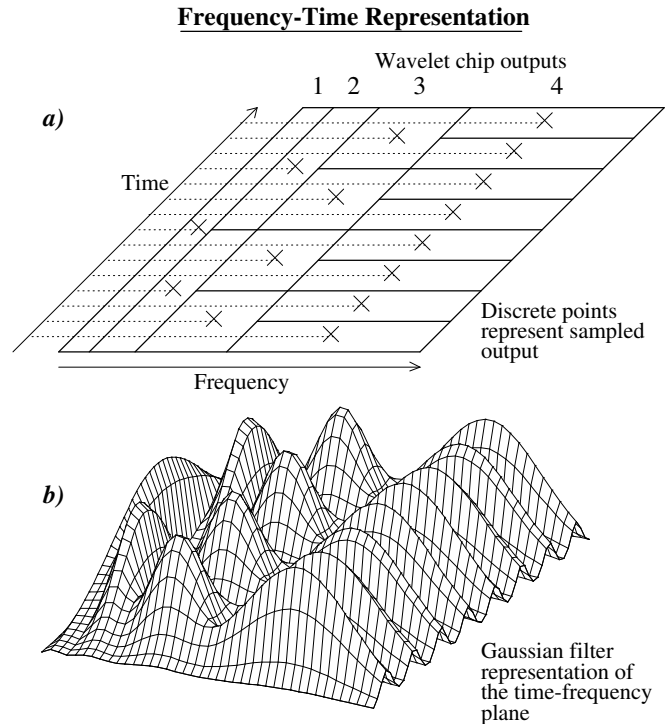


Fig. 1. Output sampling and frequency-time representation of the input.

gaussian-shaped function,

$$f(\omega) = e^{-\omega^2/2\sigma^2} \quad (1)$$

which is also gaussian-shaped (and therefore effectively bandlimited) in the time domain. A few other useful filter functions have been investigated in [1, 2]. The CWT as described is merely a band equalizer; in order to produce the most efficient description in terms of the time-frequency uncertainty relation, each frequency band is sampled at a rate proportional to its bandwidth. The result is that each rectangle in the time-frequency plane has an equal

\*This chip was designed by T. Edwards under the direction of M. D. Godfrey, and fabricated by Zilog Corporation of Campbell, California. Zilog provided the funding for this project.

area representing the effective time and frequency bandwidths of the gaussian filter (with some overlap).

Fig. 1a shows how proper sampling of the transform outputs creates a set of points which best describes the input signal as a function of both time and frequency. If the filter bands are sampled in the binary-tree fashion shown and the channels are centered on a  $\log_2$  scale, the samples can be easily time-division multiplexed into a single output stream. Fig. 1b is a representation of the gaussian filter functions which produce each sampled value.

This chip was built as an alternative to using a computer or DSP system to perform the same transform; the advantages of this approach are reduced size and power consumption. The resulting analog implementation is inflexible in terms of ability to reprogram the function performed, and requires dealing with the problems of temperature sensitivity, variable process parameters, and noise injection throughout the circuit. However, we believe that the circuit will be effective for speech signals and will be smaller and cheaper than the digital alternatives. The circuit is based on the Analog VLSI techniques described by Carver Mead in *Analog VLSI and Neural Systems* [3].

## II. COMPLEX DEMODULATION

In order to simplify the process of designing bandpass filters that maintain a gaussian shape while allowing variable center frequency and width, we use a demodulation process to convert the center frequency of each filter band to zero before filtering; the result is that each required bandpass filter is reduced to two lowpass filters which are relatively simple in design. This process is known as *complex demodulation* [4] and is described below.

The input function is designated by  $f_{\text{in}}(t)$ . In order to demodulate it with respect to some given frequency  $\omega_c$  (one of the center frequencies of the wavelet decomposition), we use the following multiplication:

$$f_{\text{out}}(t) = 2LPF(f_{\text{in}}(t)(\cos \omega_c t + i \sin \omega_c t)). \quad (2)$$

Since the low-pass filter  $LPF(\cdot)$  is assumed to be real, and  $f_{\text{in}}(t)$  is real, then the output  $f_{\text{out}}(t)$  must necessarily be complex-valued, and can be represented by

$$f_{\text{out}}(t) \equiv f_{\text{real}}(t) + i f_{\text{imag}}(t). \quad (3)$$

Therefore,

$$f_{\text{real}}(t) = 2LPF(f_{\text{in}}(t) \cos \omega_c t) \quad (4)$$

and

$$f_{\text{imag}}(t) = 2LPF(f_{\text{in}}(t) \sin \omega_c t). \quad (5)$$

In other words, the “real” and “imaginary” parts are both results which are easily obtained by multiplying the input

by two sinusoids which are  $90^\circ$  out of phase with each other.

Each part (sine and cosine) of the demodulation process produces a new signal which contains the sum and the difference of the original and “carrier” frequencies. Since we are demodulating the carrier frequency down to zero, we are interested only in the difference, so we low-pass filter the function to get rid of the part containing the sum of the two signal frequencies. The remaining difference does not differentiate between signals on one side or the other of  $\omega_c$  since  $\omega_c$  is now at zero and negative frequencies have no physical meaning. However, this information is preserved in the relationship between  $f_{\text{real}}(t)$  and  $f_{\text{imag}}(t)$ , as shown by the exact reconstruction below.

In order to remodulate the signal back to its original frequency, we perform the following multiplication:

$$\hat{f}_{\text{in}}(t) = f_{\text{out}}(t)(\cos \omega_c t - i \sin \omega_c t). \quad (6)$$

This is the same function as the demodulation (2) except for the change in sign. We multiply out the real and imaginary parts of this equation to get a purely real result, which is the reconstruction of the original input:

$$\hat{f}_{\text{in}}(t) = f_{\text{real}}(t) \cos \omega_c t + f_{\text{imag}}(t) \sin \omega_c t. \quad (7)$$

Here the signs have worked out such that the remodulating sinusoids have exactly the same phase relation as the demodulating sinusoids. Note that no low-pass filter is needed for reconstruction. In the instance of the continuous wavelet transform, the low-pass filter for the demodulation can be combined with the gaussian filter of the transformation.

## III. WAVELET GAUSSIAN FUNCTION

The design of the circuit which approximates a gaussian filter (as described by Grossman [1]) is based on a probability argument. Because the signals are first demodulated, it is only necessary to create one half of a gaussian centered around zero. When the signals are remodulated to their respective center frequencies, this low-pass filter will behave as if it were reflected symmetrically across the modulation frequency, becoming a band-pass filter with gaussian characteristics.

Equation (8) describes the filter we designed:

$$H(s) = \frac{V_{\text{out}}}{V_{\text{in}}} = \left( \frac{1}{\tau s + 1} \right)^n. \quad (8)$$

This filter consists of a cascade of  $n$  follower-integrator sections in series. Although (8) converges to a delta function in the limit as  $n \rightarrow \infty$  for constant  $\tau$ , it can be shown that when  $\tau$  is replaced by an expression which maintains constant bandwidth, the transfer function approaches the

gaussian function (1) as  $n \rightarrow \infty$ . A proof of this equation can be found in [5] which shows that the gaussian shape is an example of the central limit theorem; in other words, it is a result of the use of cascaded stages and is relatively independent of the shape of the filter.

For  $n$  sections in cascade, the relationship between  $\tau$  and  $\sigma$  from (1) is

$$\sigma = \frac{1}{\tau \sqrt{n \ln(2)}}. \quad (9)$$

Due to considerations of signal-to-noise ratio, and the ability to generate a given constant  $\tau$  from a voltage applied to the follower-integrators, we chose a cascade order of five sections. The circuit is shown in part in Fig. 2.




Fig. 2. Three cascaded stages of a filter approximating a half-gaussian function.

The center frequencies of the filters in the filter bank are spaced on a  $\log_2$  scale. Assuming speech-quality bandwidth for the input signal, we decided that six outputs would be sufficient, giving typical center frequencies of 9kHz, 4.5kHz, 2.25kHz, 1.125kHz, 562.5Hz, and 281.25Hz. The center frequency of the highest-frequency filter is determined by an oscillator which can be generated either on-chip for a voltage-controlled frequency, or off-chip for a stable frequency. The center frequencies of the rest of the filters are determined by dividing down the oscillator appropriately.

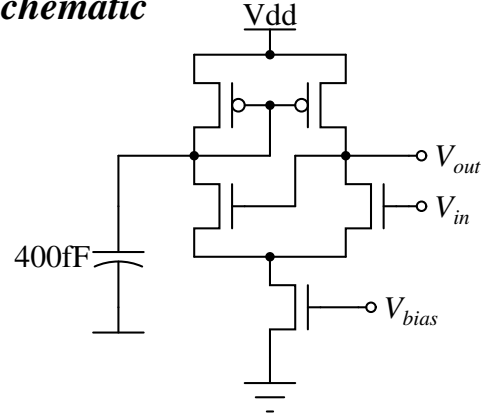
The bandwidth of each gaussian filter is set automatically with respect to the others with the exception of the first and last filters, which have widths adjustable using two control voltages **VgaussH** and **VgaussL**. In terms of the transfer function (8) for the filter, the parameters  $\tau$  of the highest- and lowest-frequency filters are fixed by these control inputs.  $\tau$  should be calculated to assure that the width of each gaussian is proportional to the value of the center frequency.

#### IV. THE VOLTAGE-CONTROLLED OSCILLATOR

##### A. The Hysteretic Inverter

This circuit, shown in Fig. 3, is the low-power analog VLSI equivalent of the latch. The circuit is a transconductance

##### Schematic



##### Symbol




Fig. 3. The hysteretic inverter/latch circuit.

amplifier with a positive feedback connection. The capacitor prevents large voltage swings on the side opposite the output, and so prevents the circuit from sticking in a latched state. The bias voltage **Vbias** is set appropriately to place the transistor in the subthreshold region. The limited amount of current prevents the circuit from switching very fast, so it can only be used in instances where slow switching time is acceptable. In the wavelet transform, it is used to produce the oscillatory output for the modulators, so its highest rate for intended applications is on the order of 10 kHz. The switching points of the hysteretic inverter are close to power and ground, and are a weak function of **Vbias**.

##### B. Voltage-Controlled Oscillator

Fig. 4 shows the voltage-controlled oscillator (VCO) used in this circuit. The VCO operates using the principle that the saturated current output of wide-range amplifier (**A1**) produces an almost perfectly linear rise in the voltage on the capacitor at its output. When that voltage reaches a critical value, the hysteretic inverter flips states, reversing the values on the inputs to amplifier (**A1**). Consequently, charge is pulled off of the capacitor and the voltage drops linearly until reaching the critical low value which causes the hysteretic inverter to change state again. The frequency of the output **Vout** (cosine) is dependent on the time required for the capacitor to charge and discharge, which is dependent on the current drive of the amplifier,

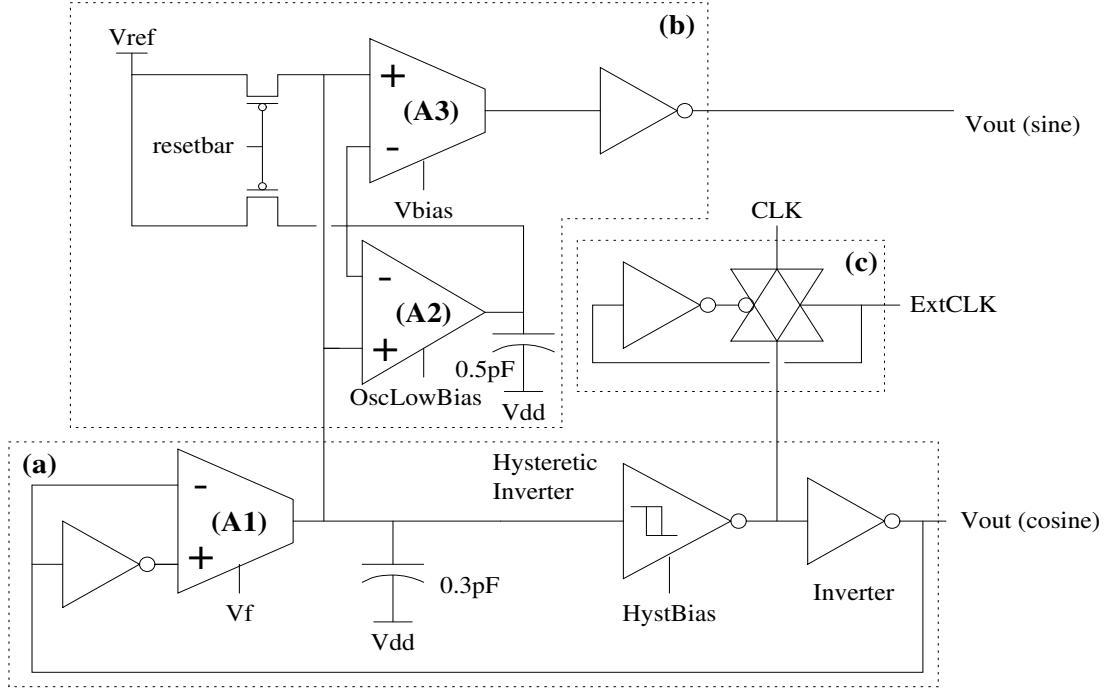


Fig. 4. Low-power voltage-controlled oscillator.

which is controlled by the bias voltage  $V_f$ . The two inverters in the main part of the circuit (a) ensure a clean square wave at the output.

The  $+90^\circ$  phase-shifted square wave output is produced by a simple feed-forward adaptive circuit (b). A follower-integrator amplifier (A2) tracks the average value of the input, which in this case is the triangle wave produced by the charging/discharging capacitor of the main oscillator (a). This midpoint value is compared to the actual value of the triangle wave using a high-gain amplifier (A3) which then acts as a signum function. The result is a square wave output which is quarter cycle behind the main oscillator output. The value of **OscLowBias** should be set to ensure a slow adaptation. The digital inverter at the output of the circuit ensures a clean square wave signal.

The frequency  $f$  of the oscillator is proportional to the exponential bias voltage  $V_f$ .

#### V. WAVELET CHIP SLICE

Due to the parallel nature of the analog wavelet computation, this chip is easily created using abutting slices of circuitry. There are twelve slices in all, each containing the logic to divide down an incoming clock (oscillator) signal, low-pass filter the oscillator to produce a smooth sine (or cosine) wave, multiply this modulating signal with two different inputs (one of them for the signal decompo-

sition part of the transform, the other for the reconstruction half), filter the decomposition result with a gaussian-shaped function, and aggregate (average) the reconstruction result with the result from all other slices to produce a final result. The slices alternate between those producing sine wave modulating signals and cosine wave modulating signals (defined by their phase relation to one another). Each sine-cosine pair of slices receives the same value of  $V_{\text{gauss}}$  for the gaussian filter. This value is tapped off of a long polysilicon resistor which reaches across all the filter banks and has fixed endpoint values.

A block diagram of the entire Wavelet Transform chip is shown in Fig. 5.

#### VI. CHIP SPECIFICATIONS

- Power Supply: +5V DC  $\pm 5\%$
- Input mean value: 2.5V  $\pm 0.5V$
- Input p-p amplitude: 0.4V
- Input frequency range: 80Hz to 10kHz
- Number of output channels: 12 (6 pairs)
- Silicon area:  $1.96 \times 10^6 \mu\text{m}^2$  (in 2.0 micron CMOS process)
- Chip package: 68-pin PLCC

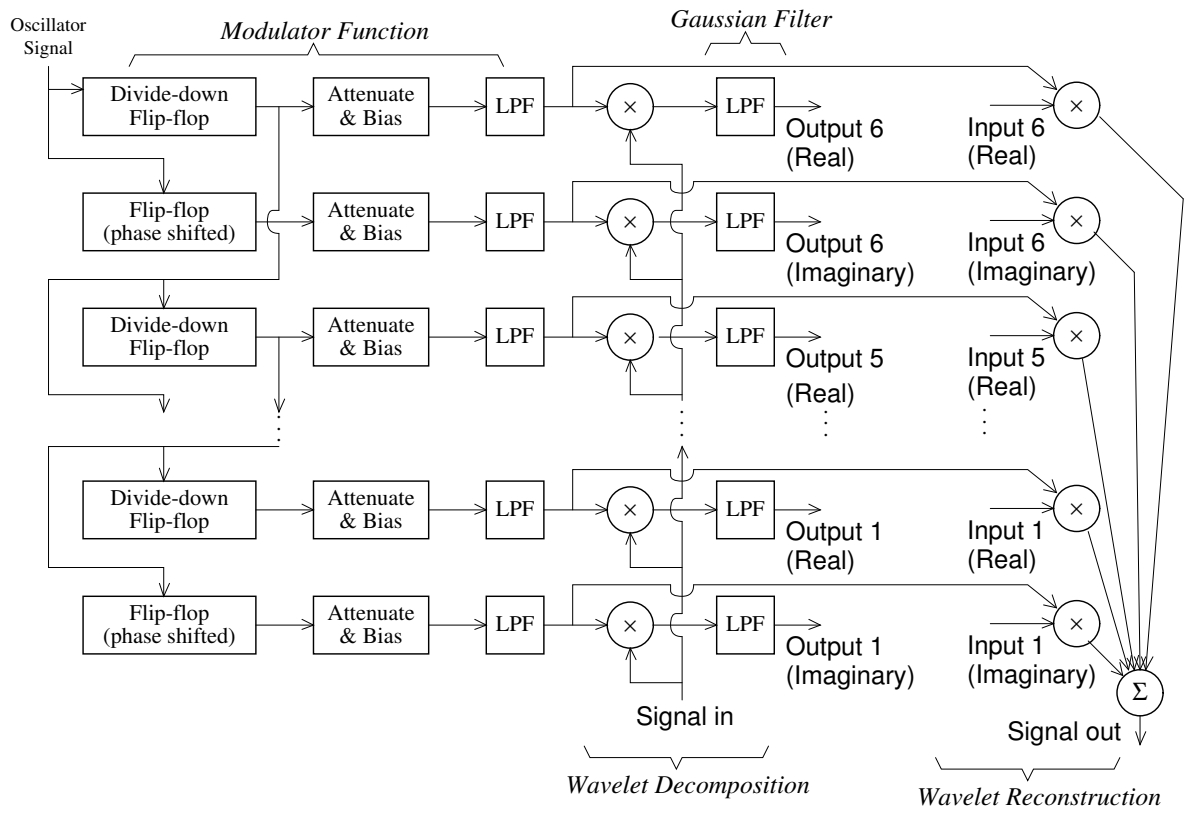


Fig. 5. Wavelet Transform Chip block diagram.

#### REFERENCES

- [1] A. Grossmann, R. Kronland-Martinet, and J. Morlet, "Reading and understanding continuous wavelet transforms," *Wavelets: Time-Frequency Methods and Phase Space*. Springer-Verlag, 1989, pp. 2–20.
- [2] R. Bracewell, "Adaptive chirplet representation of signals on time-frequency plane," *Electronics Letters*, vol. 27, no. 13, pp. 1159–1161, June 1991.
- [3] C. Mead, *Analog VLSI and Neural Systems*. Reading, MA: Addison-Wesley, 1989.
- [4] C. Bingham, M. D. Godfrey, and J. W. Tukey, "Modern techniques of power spectrum estimation," *IEEE Trans. Audio and Electroacoustics*, vol. 15, no. 2, pp. 56–66, June 1967.
- [5] W. McC. Siebert, *Circuits, Signals, and Systems*. Cambridge, MA: The MIT Press, 1986, pp. 496–497.