A Second-Generation Analog Wavelet Chip for Speech Compression

Chip Specifications

R. Timothy Edwards and Michael D. Godfrey

December 15, 1992

Abstract

This paper describes the theory and implementation in CMOS technology of a circuit which performs analog wavelet decompositions of a one-dimensional (e.g., sound) input. The circuit is based on the Analog VLSI techniques described by Carver Mead in *Analog VLSI and Neural Systems* [1]. The analog wavelet outputs are the output of a logarithmically scaled set 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. This chip has been designed by Tim Edwards under the direction of Michael Godfrey, and fabricated by MOSIS. Zilog Corporation of Campbell, California, provided the funding for this project.

1 Overview

This paper describes the Analog Wavelet Transform chip, a CMOS subthreshold circuit design intended for use in an experimental speech compression system. This chip was built as an alternative to using a computer or DSP system to perform the same calculations; 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 problem 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 Wavelet transform performed by the chip is an analog filtering function and is similar to what is known as the Gabor spectrogram (and is not to be confused with the Discrete Wavelet Transform)[2], which divides a signal into a set of frequency bands (in this case, six), often on a logarithmic scale, using a gaussian-shaped filter as the bandpass function. We have used complex demodulation so that the center frequency of each filter band is zero [3]. Two outputs for each filter band preserve phase information about the signal in order to provide the ability to reconstruct the signal from its decomposed outputs. The total combined bandwith of the outputs is approximately equal to the bandwidth of the original signal.

The original version of this paper, for the first-generation chip, accurately describes the theory behind the continous wavelet transform and its implementation on the chip. The chip specifically addressed by that paper has been designated **WaveChip2**, and the current second-generation chip has been designated **WaveChip4**. WaveChip1 was the same design as WaveChip2 but with a different aspect ratio for the layout. WaveChip3 is essentially the same as WaveChip4 but was fabricated by Zilog without a poly1/poly2 process. Also, WaveChip3 did not contain the latest version of the oscillator circuit. All chips were fabricated at a 2.0 micron technology. The expected completion date for the fabrication of WaveChip4 at MOSIS is January, 1993.

2 New Features

The new features of WaveChip4 which are improvements over the original are as follows:

- Sampling and inverse sampling of the signal is performed on-chip, with a time-division multiplexing scheme to fit all of the input and output data onto two lines each, representing the real and imaginary parts of the wavelet description.
- Where possible, all wide-range transconductance amplifiers have been replaced with simple transconductance amplifiers. The simple amplifiers inject less noise into the signal, and since the signal along the main data path has a narrow range about the power supply midpoint, there is no need to use the wide-range amplifier.
- The Gaussian filter has been increased to 8 stages from 5 on the basis of the statistical analysis of the electronic cochlea design, which suggests that there will be less noise in the signal for greater numbers of stages, a result which is due to the Central Limit Theorem.
- A new design is used for the voltage-controlled oscillator. The double oscillator circuit has been abandoned in favor of a single oscillator, and the hysteretic inverter has been redesigned. The $90°$ phase-shifted oscillating signal is automatically produced by the oscillator using a feed-forward adaptive technique.
- The flip-flops have been replaced with a more robust design. The phase-shifted oscillator signal is used with the normal signal in order to produce a two-phase clock.
- Digital circuitry has been placed on its own power supply which is disconnected from Analog Vdd and GND. In addition, most of the digital circuitry has been isolated from the analog circuitry by surrounding certain sections with wells and guard rings.
- Diffusion capacitors have been replaced with poly1/poly2 capacitors in the analog section.
- Custom pads and a custom padframe were used, so there are no poly resistors connected between the routing wiring and the pads.
- All wires from the core to the padframe were hand-routed to prevent unnecessary coupling of wires.

3 Frequency-Time Description

The continuous-time outputs of the gaussian filter represent a redundant signal description. The most efficient (non-redundant) description of the signal is achieved by sampling the wavelet outputs.

Figure 1 shows how proper sampling of the outputs of the chip creates a set of points which best describe the input signal as a function of both time and frequency. If the filter

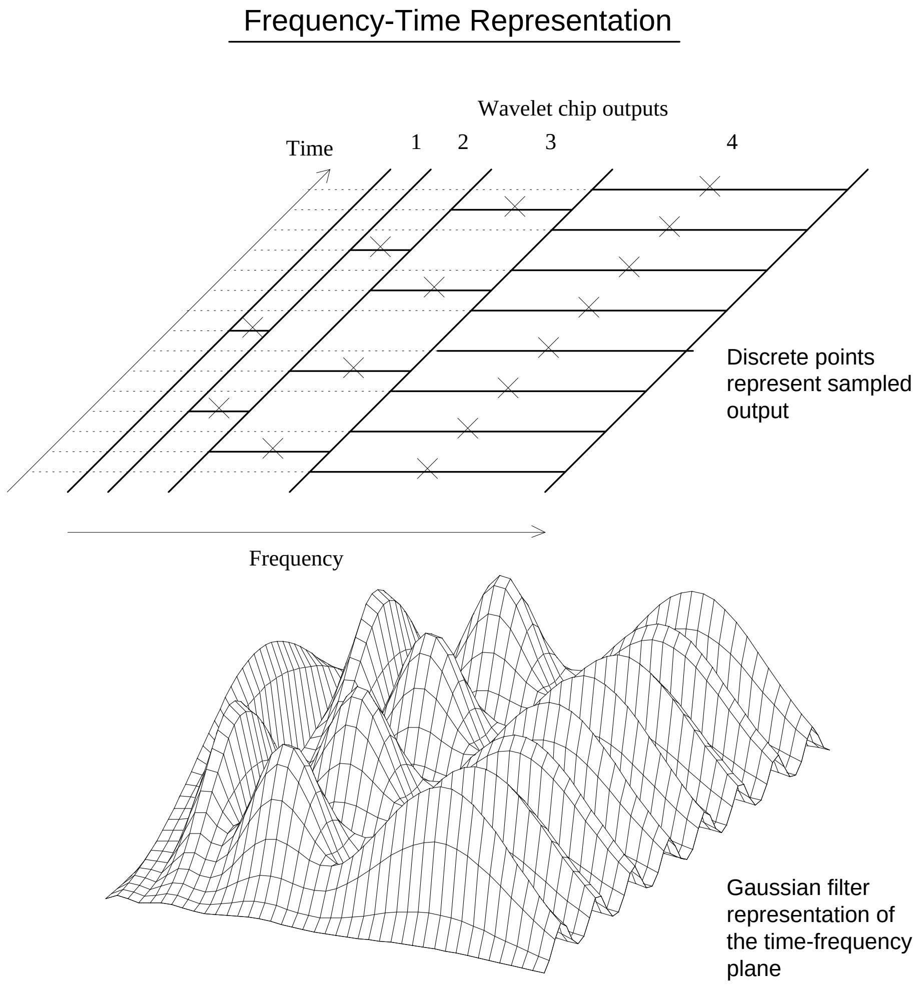

Figure 1: Output sampling and frequency-time representation of the input.

bands are sampled in the binary-tree fashion shown, then each band can be sampled at a rate proportional to its bandwidth, and all of the samples can be time-division multiplexed into an output stream.

Each rectangle in the time-frequency plane has an equal area; this is the area over which the filter calculation is made at the time and in the filter band that the sample was taken. The bottom part of the figure shows another representation of the gaussian filters which produce each sampled value.

4 Wavelet Transformer Chip I/O

The current circuit design requires a total of 54 pins, as documented below. The package type for this chip will be a 84-pin PGA package.

The pinouts and their names are desribed in the Table 1. Pins are for the 84-pin PGA package, and entries are listed in clockwise order of chip pads, beginning in the upper left-hand corner on the top. The remaining pins in the package are devoted to another test circuit by Boyd Fowler which is isolated from WaveChip4 by a separate pad ring (with diffusion guard rings) and a separate power supply.

WaveChip4 Pin List		
Pin No.	*Name*	*Description*
1	A_GND	Analog and padring Ground, 0V
2	A_GND	Analog and padring Ground, 0V
3	Sync	Output synchronization signal (phi1 of first flipflop output)
4	TESTbar	Signal to create loopback of outputs to inputs, disabling the pins
5	Iout	Multiplexed decomposition output (imaginary)
6	Rout	Multiplexed decomposition output (real)
7	Vdriver	Bias to output follower-amplifiers
8	VgaussH	Gaussian filter width control, high end
9	Vout	Reconstructed signal output
10	Vin	Signal input
11	MultTest	Output from the first signal multiplier
12	VtauH	Modulator LPF width control, high
13	CosWaveTest	Output from the smoothed (sinusoid) CosOscTest signal
14	VQH	Modulator LPF resonance control, high
15	SineOscTest	Oscillator phase-shifted output (modulator section)
16	CosOscTest	Oscillator normal output (modulator section)
17	HystBias2	Bias voltage to oscillator hysteretic inverter (sampler)
18	OscLowBias	Bias voltage to oscillator time-average circuit (both)
19	Vref	Input/Ouput voltage reference (typ. about 2.5V)
20	A_Vdd	Analog and padring Power, +5V
21	A_Vdd	Analog and padring Power, +5V
22	HystBias1	Bias voltage to oscillator hysteretic inverter (modulator)
23	Resetbar	System reset signal
24	ExtCLK	External clock override (Input to chip is driven by CLK input) (modulator)
25	Vf	Frequency-control voltage input
26	CLK	External clock input (also part of reset mechanism) (modulator)
27	SinOut1Test	Output from the last divide-down flipflop, modulator section
28	VtauL	Modulator LPF width control, low
29	VQL	Modulator LPF resonance control, low
30	Vb1	Transamp bias for attenuator
31	Vb2	Transamp bias for wide-range multiplier
32	Vb3D	Transamp bias for resistor (decomposition)
33	Vb3R	Transamp bias for resistor (reconstruction)
34	D_GND	Digital Ground (isolated), 0V
72	D_Vdd	Digital Power (isolated), +5V
73	VgaussL	Gaussian filter width control, low end
74	Vtau2L	Sample input LPF width control, low
75	VQ2L	Sample input LPF resonance control, low
76	Vbias	Bias voltage for non-critical transamps (aggregators, etc.)

Table 1: Pin list for the WaveChip 84-pin PGA package.

WaveChip4 Pin List (continued)		
Pin No.	*Name*	*Description*
77	CosOut2Test	Output from the last divide-down flipflop (sampler)
78	CLK2	External clock input (also part of reset mechanism) (sampler)
79	ExtCLK2	External clock override (Input to chip is driven by CLK input) (sampler)
80	Vf2	Frequency-control voltage input (for the sampler)
81	VQ2H	Sample input LPF resonance control, high
82	VQ2L	Sample input LPF resonance control, low
83	Rin	Multiplexed reconstruction input (real)
84	Iin	Multiplexed reconstruction input (imaginary)

Table 2: Pin list for the WaveChip 84-pin PGA package.

5 Chip Specifications

- Power Supply: +5V DC ±5%
- Input mean value: 2.5V ±1.0V
- Input p-p amplitude: 0.4V
- Input frequency range: 180Hz to 12kHz
- Number of output channels: 2 (1 pair)
- Center frequencies of outputs: Adjustible (see above)
- Wavelet function: Gaussian bandpass filter
- Input and Output SNR: To be determined
- Silicon area: (insert here)
- Fabrication: Orbit 2.0μm SCPE process
- Chip package: 84-pin PGA

6 Wavelet Chip block diagram

A block diagram of the Wavelet Transform chip is shown in Figure 2. The oscillator and test circuitry are not shown.

7 Analysis of Analog VLSI Circuits

Figures 3 through 8 show the building-block circuits of the wavelet transform. A detailed analysis of some of the basic circuits can be found in Mead [1] and a partail description can be found in the paper for the first-generation WaveChip. These include the wide-range transconductance amplifier, the wide-range Gilbert multiplier, the second-order section, and the follower-aggregator. These can be found on pages 67–80, 90–96, 179–192, and 105–107,

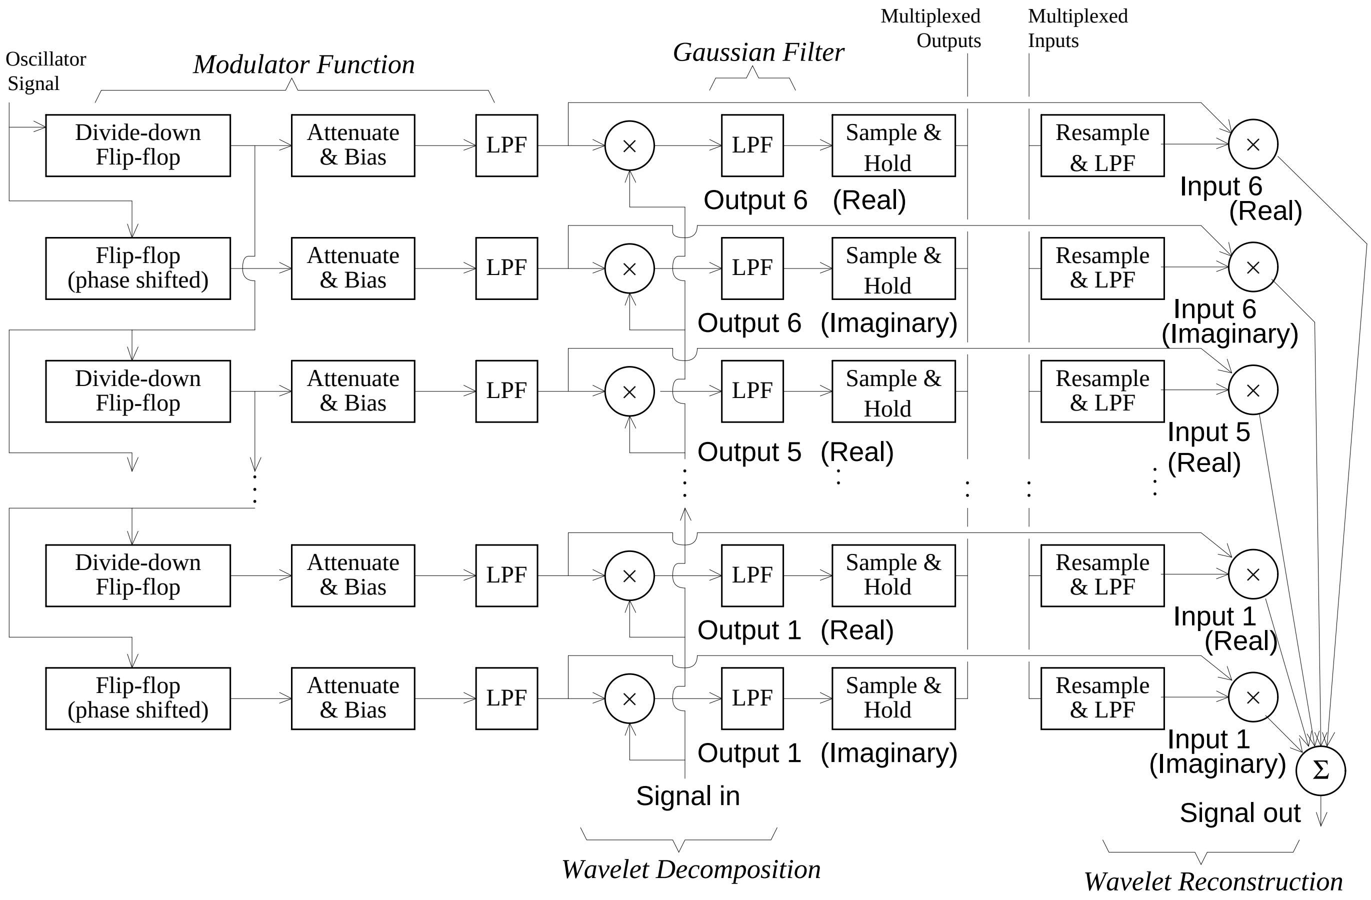

Figure 2: Wavelet Transform Chip block diagram.

respectively, of *Analog VLSI and Neural Systems*. Novel circuits designed for this chip are presented below.

7.1 The Hysteretic Inverter

This circuit, shown in Figure 3, is the low-power analog VLSI equivalent of the latch. In essence, it consists of a transconductance amplifier; its input is the amplifier's negative differential voltage input, and its output is the amplifier output, tied to the positive differential voltage input for a positive feedback configuration. In order to prevent the latch circuit from locking onto one of the voltage rails, the differential output opposite the usual single-ended output is tied to a capacitor (400fF), which prevents large voltage swings on that node. Consequently, the latch circuit is able to be flipped into the opposite state when the input becomes high or low enough. The bias voltage **Vbias** is set to place the transistor in the subthreshold region. The limited amount of current prevents the circuit from switching very fast, so it should 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 is on the order of 20 kHz. Where necessary, raising the bias voltage will decrease the switching time at the expense of power consumption. The switching points of the hysteretic inverter are close to power and ground, and are a weak function of **Vbias**. This circuit is an improvement over the hysteretic inverter used in WaveChip1 through WaveChip3, and is

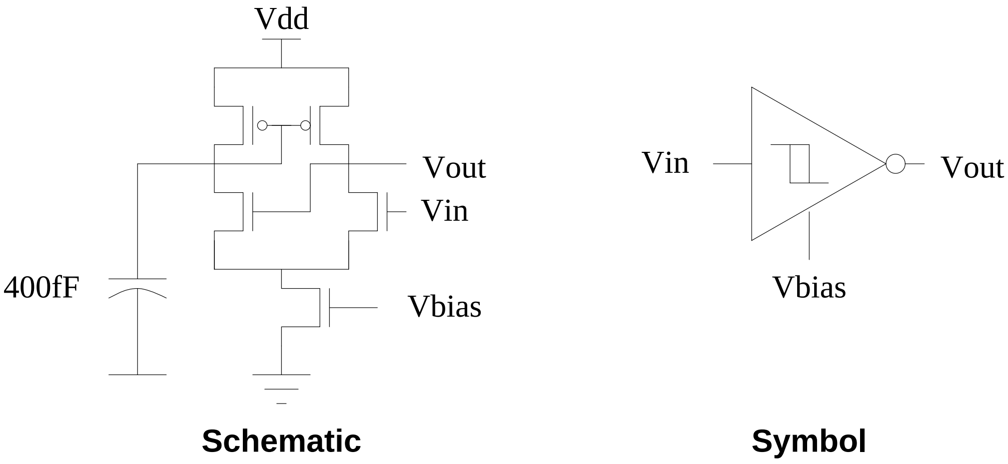

Figure 3: The hysteretic inverter/latch circuit.

robust under a normal range of bias voltages. The hysteretic inverter is used in the Voltage Controlled Oscillator, described below.

7.2 Voltage Controlled Oscillator (VCO)

Figure 4 shows the voltage-controlled oscillator used in this chip. The VCO operates using the principle that the saturated current output of wide-range amplifier (**1**) 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 (**1**). 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, which is controlled by the bias voltage **Vf**. The two inverters in the main part of the circuit (**a**) ensure a clean square wave at the output.

The $+90°$ phase-shifted square wave output is produced by a simple feed-forward adaptive circuit (**b**). A follower-integrator amplifier (**2**) 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 (**3**) 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 about 0.3 to 0.4 volts, which ensures a slow adaptation. The noise produced by setting this amplifier bias to such a low value will not significantly affect the square wave output. 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 **Vf**. Empirical results from simulation give an equation for frequency vs. bias voltage as:

$$\mathbf{Vf} \approx (0.039)\ln(f) + (0.394),$$

which holds for the normal operating range of **Vf**: $0.4 < \mathbf{Vf} < 0.9$.

The desired voltage **Vf** is that which produces a frequency f at the desired highest center frequency (the frequency of the first demodulating cosine/sine waves).

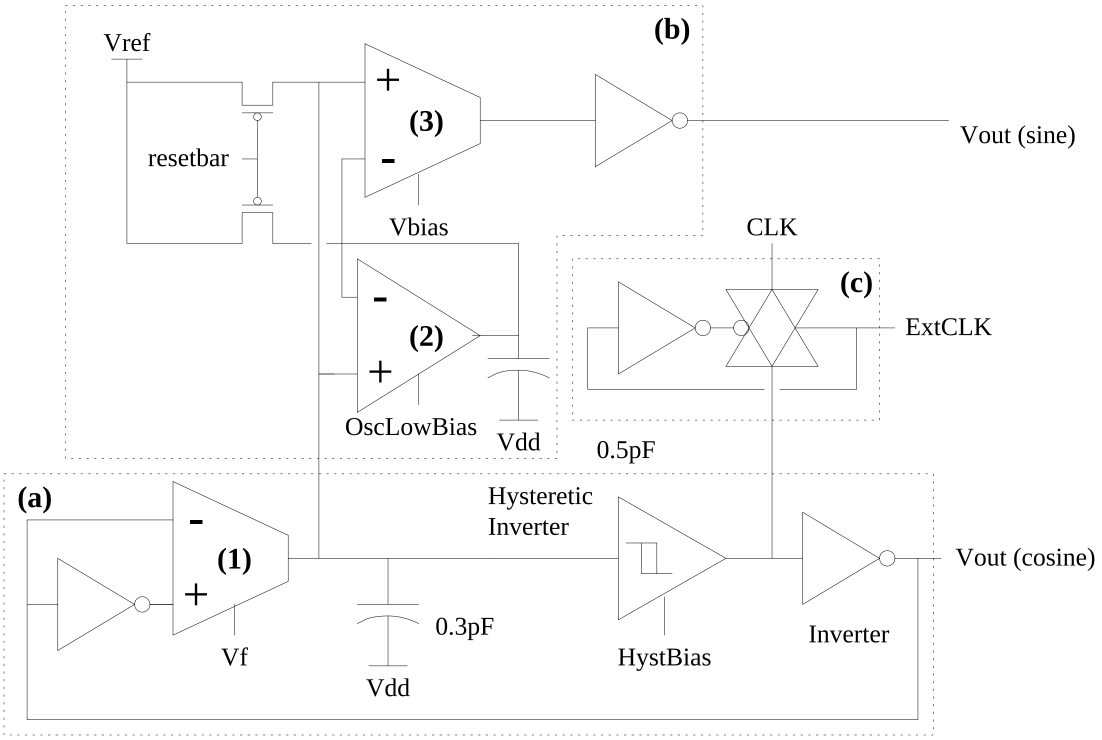

Figure 4: Low-power voltage-controlled oscillator.

A second VCO is used to control the timing of the input and output sampling. This VCO should be run at the desired highest sampling frequency; typically, this will be 4/3 times faster than the other oscillator.

The driver circuit (**c**) allows the VCO to operate in slave mode wherein the output of the main oscillator (**a**) is forced to follow an external clock signal **CLK**. Provided that the input **Vf** is set to a value approximately consistent with the external clock frequency, the circuit will continue to perform the function of generating the phase-shifted signal. This is a good way to ensure temperature independence of the oscillator at the expense of inability to control the frequency value.

7.3 Low-Gain Resistor-Amplifier

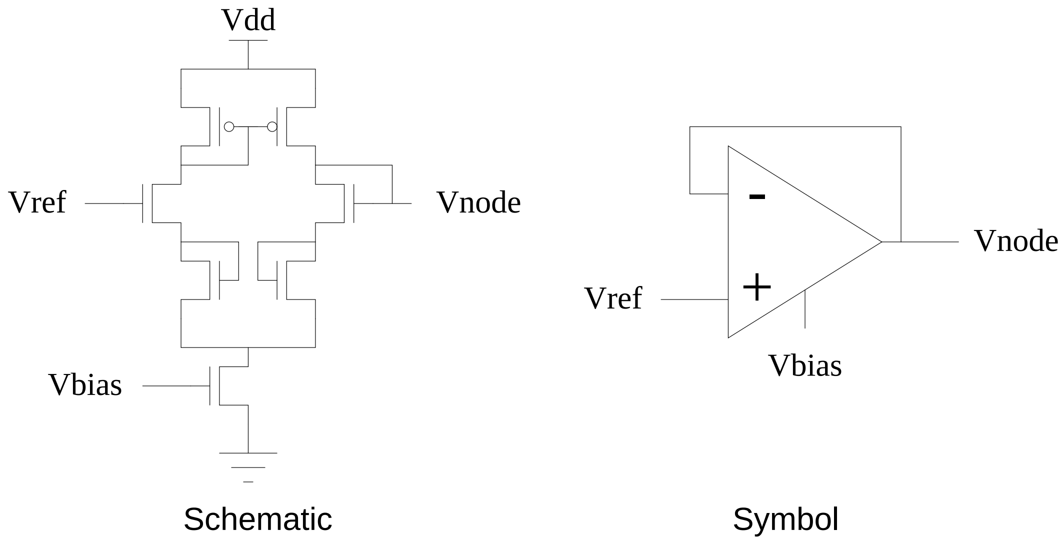

Figure 5: The resistor amplifier-follower circuit.

One of the difficulties of Analog VLSI is translating Voltages to Currents and vice versa in a linear fashion. The main problem is that the value of resistors needed is too large for a conveniently short poly line drawn on the chip, and too small for the natural resistance of a transistor fixed in subthreshold. However, in cases for which a resistance is needed from any point to a fixed reference voltage, in order to transform a current output from an amplifier or similar device into a voltage, a transconductance amplifier with fairly low gain can be acceptable. Such a circuit is shown in Figure 5. This amplifier is connected as a follower and thus acts like a follower-aggregator. The purpose of forcing the amplifier to have low gain increases the region in which the effective resistance is linear. The two diode-connected transistors in the circuit add to the g_0 conductance of the input transistors which reduces the small signal voltage gain.

7.4 Divide-Down Flip-Flops

The function of this circuit (Figure 6), which is digital in nature, is easily seen from the circuit diagram, which shows a standard two-phase clock controlled trigger flip-flop. The original divide-down flip-flop derives its two-phase clock signal from the VCO outputs, and every other flip-flop derives its two-phase clock signal from the flip-flop above it. The circuit

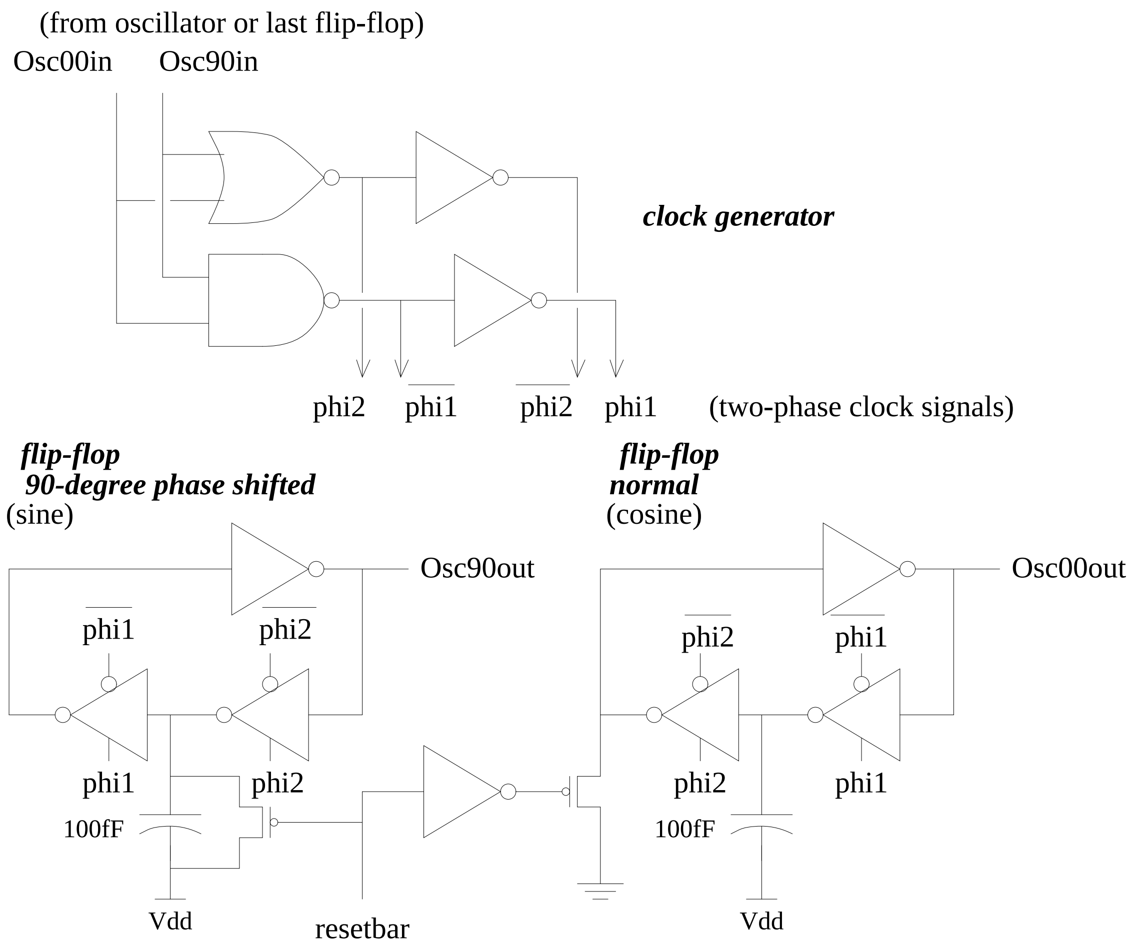

Figure 6: Divide-down flip-flop circuit.

on the bottom right differs from that on the bottom left in that the output of one is 90° out of phase with the other, which is necessary to produce the sine and cosine signals for the demodulation process, and to generate the proper clock signals for the next flip-flop circuit.

7.5 Gaussian filter section

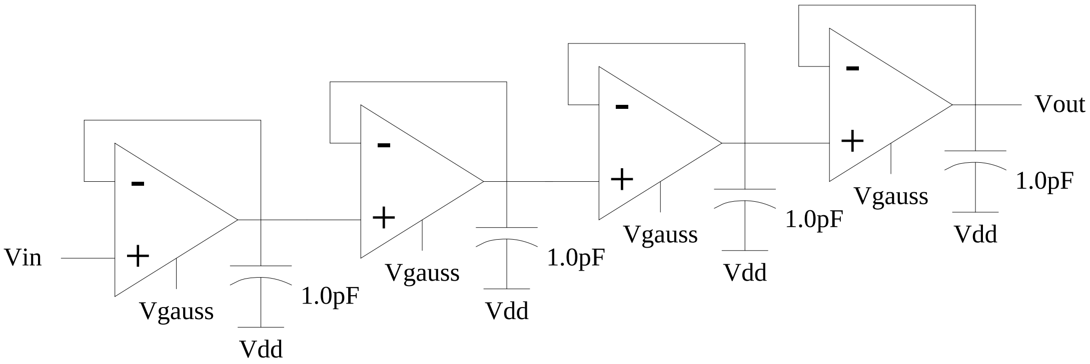

Figure 7: A four-stage filter approximating a half-gaussian function.

The parameter τ relates to the width (in frequency space) of the gaussian filter. The approximate gaussian function is the usual normal distribution

$$H(s) = e^{-\frac{\omega^2}{2\sigma^2}}$$

for which we can define a bandwidth

$$\mathrm{BW} = \sqrt{-2\sigma^2 \ln\left(\frac{1}{2}\right)}.$$

This parameter is the same τ as is computed for any follower-integrator, $\tau = C/g_m$. The approximate bandwidth calculation is found in [1] and is

$$\mathrm{BW} = \frac{1}{\tau}\sqrt{\frac{2}{n}}$$

for n cascaded stages. Thus we can compute σ in terms of τ,

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

For the eight stages of this filter, we assume that the amplifiers and capacitors are matched, such that the τ for the whole filter is the same as the τ for any one amplifier. Since we want to achieve a desired bandwidth BW, and we can compute τ for any given BW, then we need to know **Vgauss**, which is

$$\mathbf{Vgauss} = \frac{V_T}{\kappa}\ln\left(\frac{2CV_T}{I_0\kappa\tau}\right),$$

given the process-dependent parameters κ and I_0, the thermal voltage V_T, and the capacitance value C. The Central Limit Theorem assures that more stages produce a closer

Gaussian approximation. The limiting factors are **Vgauss**, which should stay above about 0.6V to keep the amplifiers out of the noisy region, and the size of the filter section, which is not an insignificant portion of the chip area. Figure 7 shows four cascaded stages of the filter.

7.6 Sample and Hold circuits

(To be completed)

7.7 Wavelet chip slice

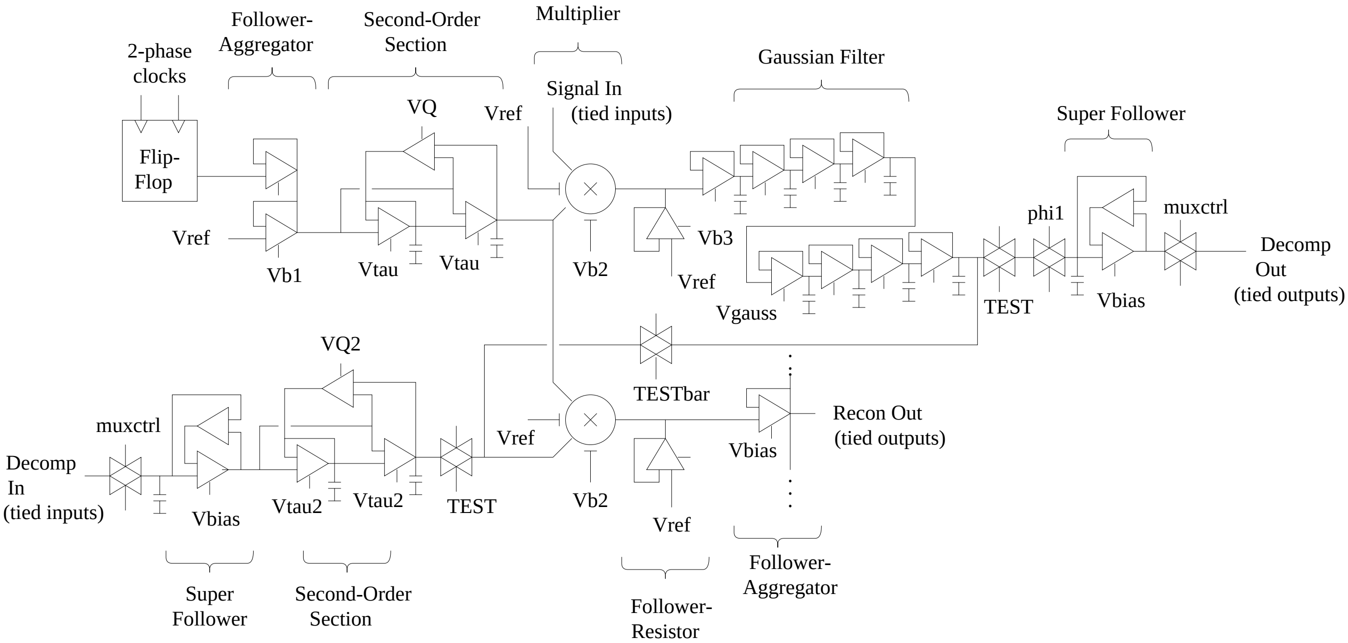

Figure 8: One of twelve slices of the wavelet transformer.

Figure 8 shows a single "slice" of the wavelet chip. 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 decomposition part of the transform, the other for the reconstruction half), low-pass 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 reconstruction input is retrieved from a multiplexed input stream by a sampling and lowpass filter circuit. The decomposition output is multiplexed into a stream using a simple sampling circuit. The slices of the transformer alternate between those producing sine wave modulating signals and cosine wave modulating signals (only defined by their phase relation to one another). Each sine-cosine pair of slices receives the same value of **VQ** and **Vtau** for the modulator low-pass filter, the same value of **VQ2** and **Vtau2** for the sampling LPF, and the same value of **Vgauss** for the gaussian filter. These three values are connected along the length of the wavelet transform chip by long polysilicon resistors; when the value of **VQ**, **Vtau**, and **Vgauss** are set properly for the first and last slices, the values for the intermediate slices are automatically linearly spaced between the endpoint values. In the chip pin description, the addition of **H** or **L** subcript on the above signal names means the high or low endpoint value (for the high-frequency modulator output or low-frequency modulator output, respectively).

8 Using the Wavelet Chip

Since the theoretical results are only approximations, we expect to set the operating voltages on the chip piece by piece, according to given test procedures.

1. **Setup procedure:** Until otherwise notified, **TESTbar** should be set high (5.0V). This signal works as described above. The other lines which must be set before beginning the test procedures are **A_Vdd** and **D_Vdd** (5.0V), **A_GND** and **D_GND** (0.0V), **Vdrive** (2.6V), **resetbar** (5.0V), and **Vref** (2.5V).

2. **Testing the Flipflops:** This test will confirm the clock range over which the flipflops operate correctly. Set **ExtCLK1** high, and attach input **CLK1** to a function generator generating a variable-frequency square wave between 0 and 5 volts. Display the signals **SineOscTest** and **SinOut1Test** on an oscilloscope in addition to the input function generator signal. **SineOscTest** should be at the same frequency as the input but lagging one-quarter of a period, and **SinOut1Test** should be oscillating at 1/32 the rate of the **SineOscTest**. Determine the frequency range of correct operation. Repeat for the other oscillator, using signals **CLK2**, **ExtCLK2**, and **CosOut2Test**. **CosOut2Test** should be oscillating at 1/32 the rate of the input, and there is no equivalent signal to **SineOscTest**.

3. **Testing the Oscillator:** Set **OscLowBias** to 0.35V, **HystBias1** to 0.8V, and **Vf1** to a voltage generator. Attach **CosOscTest** to an oscilloscope. Set **ExtCLK1** low to disable the external clock. Set **Vf1** to a value of 0.7V. This should produce a square wave on the oscilloscope. Vary **Vf1** and measure frequency at a few different points. Find the maximum and minimum frequencies at which the circuit will behave properly. Make sure that **SinOut1Test** is still producing a square wave at 1/32 the rate of **SineOscTest**. Repeat for the other oscillator using signals **Vf2**, **ExtCLK2**, and **HystBias2**. **OscLowBias** operates on both oscillators simultaneously. Test the reset mechanism by setting **resetbar** low, and then high again. If the internal clock is being used, then **CLK1** must be set to 0.0V, and **ExtCLK1** must be the logical inverse of **resetbar**. If the external clock is beign used, then reset must transition high while the input clock is in the middle of its low (0.0V) state. Upon transitioning high, both oscillator outputs should become high, and **Sync** should become high. The oscillator cosine output should fall low after one-quarter period.

4. **Testing the Modulator LPF:** Continue to generate a square wave input to the circuit using a signal on the external clock as outlined in test number 2 above. Set **Vb1** to 0.8V. Set the external clock rate to the rate at which the wavelet chip will be expected to operate. Set **VtauH** to about 0.61V and **VQH** to about 0.60V. Leave **VtauL** and **VQL** disconnected. Attach the signal **CosWaveTest** to an oscilloscope. Adjust the values of **VtauH** and **VQH** until the output best matches a sinusoid. Generally, **VtauH** should remain about 0.01V higher than **VQH**. Modify the gain using **Vb1** (and also **VQH**

and **VtauH** if necessary) to set the peak-to-peak amplitude of **CosWaveTest** to about 0.4V. Repeat the test for an external clock rate of 1/32 of the first test. When done, set **VtauH** and **VQH** to the values determined from the first test, and **VtauL** and **VQL** to the values determined from the second test.

5. **Testing the Modulator Multiplier:** Initially set pins **Vb2** and **Vb3D** to 0.8V each. All other pins should remain as they have been set from the previous tests. This test is important to keep the output of the multiplier from overflowing. Attach pin **Vin** to **VrefD**, and the pin **MultTest** to an oscilloscope. The output **MultTest** should be nearly constant and equal to 2.5V. If there is a significant signal through the multiplier due to mismatched transistors in the current mirrors, adjust **Vref** until it disappears. Record the new **Vref** value. Now, generate a sine wave centered at 2.5V (**Vref**), having a peak-to-peak amplitude of 0.4V. The output **MultTest** should be a signal containing the sum and difference of the function generator sinusoid and the **CLK** input (or the output from the voltage-controlled oscillator, if that is being used instead). Adjust **Vb3D** until the multiplier output is well-behaved, and if the output is greater than 0.4V peak-to-peak, adjust **Vb3D** until it is limited to that amount.

6. **Testing the Gaussian Filter:** Make sure **TESTbar** is high. Set **VgaussL** to 0.4V, and do not connect **VgaussH**. Set **ExtCLK2** to 5.0V, and connect **CLK2** to ground. Toggle **resetbar** to reset the second oscillator. **SYNC** should be high. The output of **Rout** should now reflect the value at the output of the first gaussian filter. Set **ExtCLK1** to 5.0V, and **CLK1** to a function generator with square wave output, and **Vin** to a constant, about 2.7V or so. For several values of **VgaussH** between about 0.4V and 0.7V, measure the amplitude of **Rout** and find the -3dB cutoff frequency (of **CLK1**) for each. When done, set **VgaussH** to a value appropriate for the normal operating frequency of **CLK1**, and **VgaussL** to a value appropreate for a frequency 1/32 times **CLK1**.

7. **Testing the Recomposition Output:** Set the **TESTbar** line low (0.0V) in order to attach the decomposition outputs directly to the reconstruction inputs. Connect **Vb3R** to **Vb3D**. Set **Vbias** to about 0.8V. Attach **Vin** to a function generator and **Vout** to an oscilloscope. Generate a sine wave of 0.4V p-p amplitude at the input, centered around **Vref**. Confirm that for all inputs in the range of the transform bandwidth, the reconstructed output is approximately equal to the input. If the output is highly nonlinear, separate **Vb3R** from **Vb3D** and decrease its value until the output becomes linear. Note that there may be a constant gain difference between the original and reconstructed signals.

8. **Testing the Sampling Circuits:** (To be completed)

References

[1] Mead, C. *Analog VLSI and Neural Systems*, Reading, MA, Addison-Wesley (1989).

[2] 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.

[3] C. Bingham, M. D. Godfrey, and J. W. Tukey, "Modern Techniques of Power Spectrum Estimation," *IEEE Transactions on Audio and Electroacoustics*, Vol. AU-15, No. 2, June 1967, pp. 56–66.