A Low-Power Acoustic Periodicity Detector Chip for Voice and Engine Detection

Hisham Abdalla, Student Member, IEEE, Timothy Horiuchi, Member, IEEE
Department of Electrical and Computer Engineering, Institute for Systems Research
University of Maryland, College Park
College Park, MD, 20472 USA

***Abstract*—The detection of voices or the rumble of engines is a desirable function in many different devices from toys to smart homes to military applications. Typical approaches involving frequency-domain computation are quite computationally intensive and require a significant power and computational budget. In an effort to construct a very low-power detector capable of acting as a wake-up signal for other systems, we have designed a simple, low-power (1.5 μW) analog VLSI circuit that detects periodicity in the time-domain envelope of the acoustic signal. The circuit was fabricated in a commercially 1.5 μm available CMOS process.**

***Index Terms*— inter-spike intervals, periodicity detection, subthreshold**

I. Introduction

Voices and engines typically have a strong harmonic acoustic signature that drives most spectrally-based voice and engine detection algorithms, however, in our system we utilize the observation that the periodicity of the time-domain envelope occurs at the fundamental frequency. By carefully applying a peak-detector and measuring the period of the envelope peaks, a sufficient number of consecutively *matching* periods are interpreted as the presence of a periodic signal of interest. By limiting the range of frequencies that are acceptable for the periodicity, detection specific to human voices or other periodic signals is possible. This *algorithm* rejects impulse events, wind noise, and is intensity-independent. Clearly, as this type of algorithm will trigger on pure tones, impulses, or non-voice or non-vehicle harmonic sounds in the right range, this is not a foolproof system and is better suited for use as a wake-up signal for more sophisticated algorithms or when false positives are acceptable.

II. System Block Diagram

Fig. 1 shows the block diagram of the circuit. A microphone signal, externally-amplified and lowpass filtered, is sent into the chip to a peak-detector/spike-generator circuit. The 'spike generator' generates a voltage spike at the onset of a peak in the envelope. The 'Interval to Voltage' stage linearly converts the inter-spike intervals into a voltage. This interval voltage is sampled and compared ('Comp' block) with the previous interval that is stored in 'Memory'.

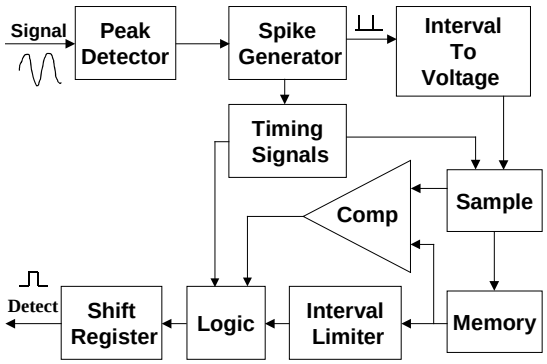

Fig. 1. Block diagram of the peak-periodicity detection chip.

The sampled interval is subsequently transferred to memory in the period between spikes. An 'Interval Limiter' circuit checks whether the interval being processed falls in its acceptance range and therefore constitutes a "valid interval". Consecutive valid intervals that match generate a "hit" pulse that is used to advance a shift-register-based counter. When non-matching intervals are encountered, the counter is reset. In the experimental data we show, three or five consecutive hits are interpreted to be a detection event. The 'Timing Signals' circuit is responsible for generating the signals necessary to sequence the sampling, voltage conversion and decision making processes following each peak in the signal.

III. Circuits

In an effort to reduce power and to match the longer time intervals of interest in the audio frequency range using smaller capacitors, we bias most of the circuits in the subthreshold region [1] of operation where the saturation drain current equation is exponential and the drain-source saturation voltage is about 100mV.

A. Peak Detector

The first stage of the circuit is the peak detector shown in Fig. 2. As the capacitor's voltage drops below the input, the transconductance amplifier (transamp) [1] turns M1 on, charging the capacitor. When the input drops below the

capacitor's voltage, M1 is turned off and the capacitor is linearly discharged by the current source.

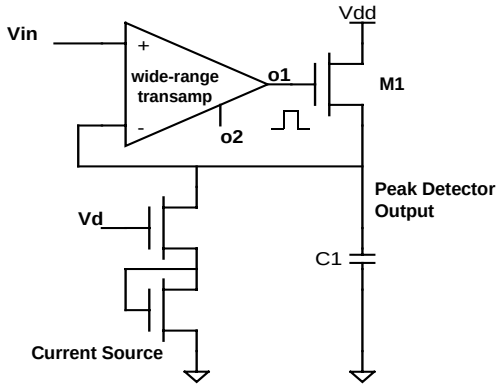

Fig. 2. Peak Detector

The output pulse o2 is high as long as the peak detector's output tracks the input. The rising edge of o2 marks the onset of the peak of the envelope. The wide-range transconductance amplifier used in the circuit is shown in Fig. 3.

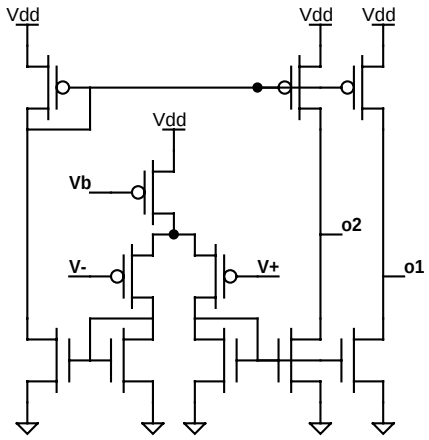

Fig. 3. Dual-output, wide output-range transconductance amplifier

B. Spike Generator

The edge detector circuit shown in Fig. 4 uses the rising edge of the peak detector's output o2 to generate a spike.

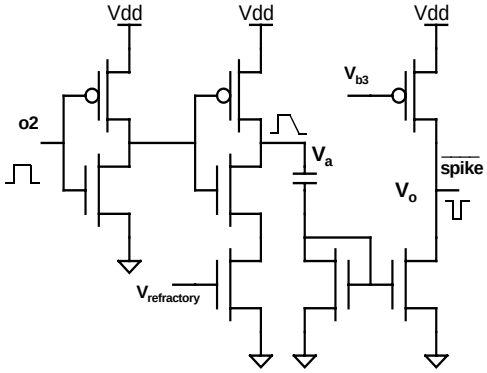

Fig. 4. Edge Detector/Spike Generator

With the rising edge of o2, V_a swings from ground to V_{dd}, the capacitive current sends a current down the current mirror pulling V_o low thus generating our voltage spike. Similar to the "refractory period" of a biological neuron, this circuit has the property that after a spike has been generated, no other spike can be generated for a certain period of time. This refractory period can be controlled by the bias voltage $V_{refractory}$. Fig. 5 is an oscilloscope trace that shows the response of the peak detector and the spike generator to an input signal.

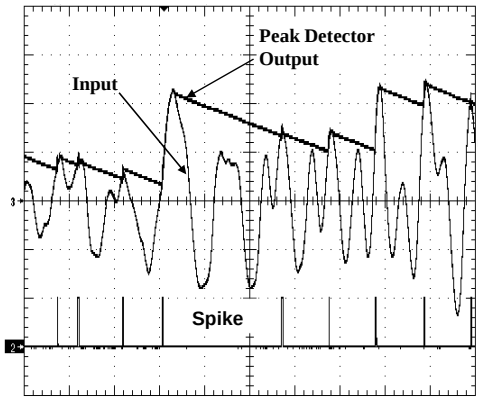

Fig. 5. Oscilloscope trace showing the input to the chip, the corresponding peak detector output and the generated spike train

C. Voltage Conversion, Sampling and Storage

An inter-spike interval is linearly transformed into a voltage, sampled, and stored for comparison with the subsequent interval. All three operations are shown in Fig. 6. The current source M1 linearly charges C3. After a spike is triggered, the sampling pulse turns on the follower tr_amp1, thus sampling the voltage ramp Vc3 and the sampled "interval" $\mathbf{V_n}$ is temporarily stored in C4. A reset pulse follows, discharging C3. The current source continuously charges C3 and so the voltage Vc3 linearly increases with time until a new spike is generated. Tr_amp2 is a follower that is weakly biased so that

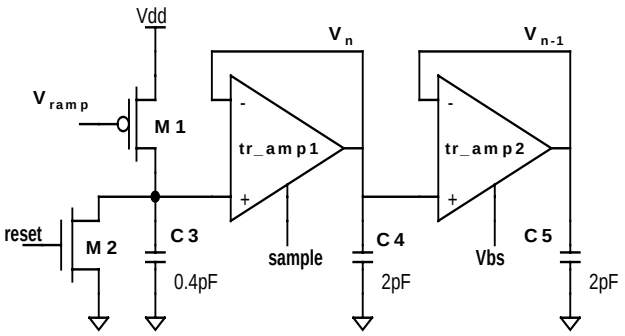

Fig. 6. Interval to Voltage Converter, Sampling and Storing of intervals

its output slowly follows its input. This allows the sampled "interval" to be slowly stored as $\mathbf{V_{n-1}}$ in C5. Because of the finite dynamic range as determined by the power rails, the interval-to-voltage transformation saturates for long intervals. Tr_amp1 and tr_amp2 are wide-range transconductance amplifiers. Fig. 7 is an oscilloscope trace which shows the

interval-to-voltage conversion and the sampling processes. The ramp is sampled with every new spike and the voltage V_n is held constant until the next sampling.

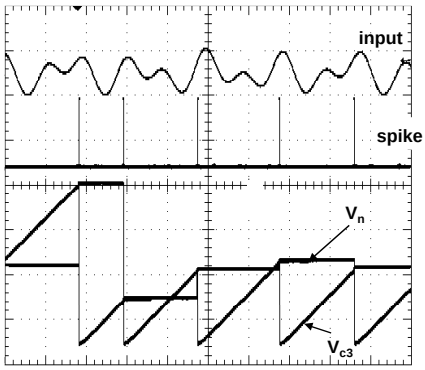

Fig. 7. Oscilloscope trace shows the interval-to-voltage conversion process as well as the sampled waveform V_n.

D. Timing Signals

To ensure the proper sequencing of the sampling and reset events, two timing waveforms, reset and sample, are generated with a guard interval. The guard interval is necessary to ensure that sampling has been completed before the reset of C3 takes place. It also serves to provide the comparator stage with enough time to "reach its decision". Fig. 8 shows the two timing signals as captured from an oscilloscope trace.

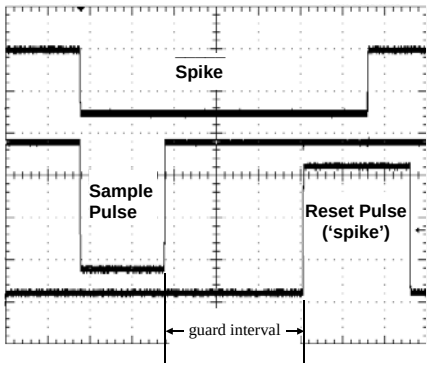

Fig. 8. Oscilloscope trace of Timing Signals.

The approach to generating the timing signals is to trigger three independent ramps on the negative edge of $\overline{spike}$. As each ramp crosses the threshold of an inverter, a pulse is generated and is used in generating the timing pulses with the required guard interval. Fig. 9 shows the circuit for generating the sample pulse. Because of the large charging current, C6 charges almost instantaneously with the negative edge of $\overline{spike}$. V_{to} controls the slope of the ramp voltage V_{c7} thereby controlling the onset of the "Turn off" pulse which terminates the sampling pulse. Fig. 10 shows the circuit for generating the reset pulse. V_r controls the slope of the ramp voltage V_{c8} thereby controlling the onset of the reset pulse. By controlling the voltage rbias, the discharge of C9 can be slowed down extending the width of the reset pulse as shown in Fig. 8.

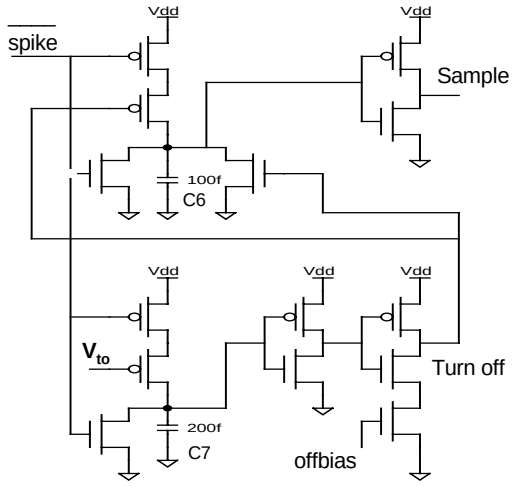

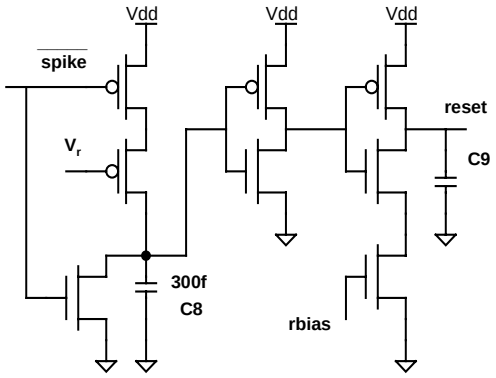

Fig. 10. Generating the reset pulse

E. Interval Limiter

It is desirable to have a circuit that defines a range of "valid intervals"; in the case of a noisy input, it is possible to get very small inter-spike intervals that match producing a false alarms. A second concern it that long intervals may match because of the saturation of the interval-to-voltage transformation. This validation circuit is shown in Fig. 11. As long as the stored

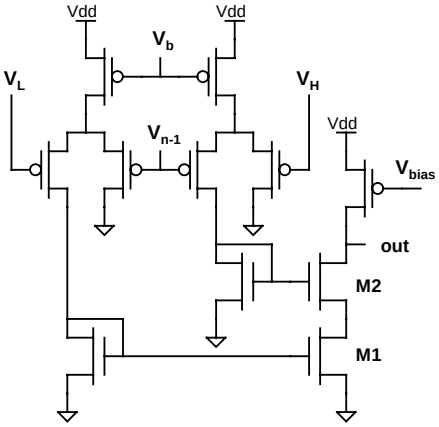

Fig. 11. Interval Limiter

voltage interval V_{n-1} satisfies the inequality $V_L < V_{n-1} < V_H$, the current mirrors will ensure that the two output transistors, M1 and M2, are on pulling the output of the circuit low indicating a "valid interval".

F. Interval Comparator

The interval comparator, shown in Fig. 12, has to decide whether the two consecutive inter-spike "intervals", V_n and V_{n-1} satisfy the inequality $|V_n - V_{n-1}| < \varepsilon$, that is, if they *match*. V_b sets a current I_b in the differential pair which is split into I_1 and I_2 which are mirrored to the two outer branches. V_{th} sets a current $I_{th} < I_b/2$. When the input voltages match, both I_1 and I_2 are strong enough to hold V_x and V_y low, thereby turning off M1 and M2 and the output is high, indicating matched intervals. When the inputs do not match, most of I_b is steered in one of the two branches, I_{th}, in turn, pulls either V_x or V_y high and the output voltage is pulled down, indicating non-matching intervals. The decision must be made before the onset of the reset pulse, i.e. in less than 200 μs. This causes ε to be a function of V_{th2}. Table 1 shows the (measured) threshold ε as a function of $(V_{th}-V_b)$.

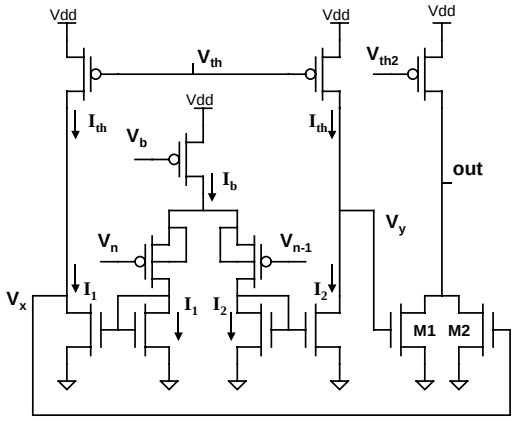

Fig. 12. Interval Comparator

ΔV (mV)	49	61	72	83	96	105	114
ε (mV)	30	50	73	86	106	125	144

Table 1. Shows the accuracy (ε) of the comparator (for a fixed V_b and V_{th2}) as a function of $\Delta V = V_{th} - V_b$.

IV. Results

Consecutive *valid intervals* that *match* generate a "hit" pulse that is used to advance a shift-register-based counter. The chip can count up to five consecutive hits. Table 2 shows the *maximum* number of hits (false alarms) in response to a white Gaussian noise signal as a function of noise power (P). Table 3 shows the *minimum* number of hits (detects) as a function of signal-to-noise ratio (SNR) for a 20 Hz sinusoidal input with a background noise level of –28 dB. In both cases, the signal duration was four minutes, lowpass filtered at 1000 Hz and repeated five times.

P (dBW)	-32	-28	-24	-20	-16
3 hits	29	80	45	18	9
5 hits	6	9	5	4	0

Table 2. Number of hits (false alarms) as a function of noise power

SNR	-10	-5	0	5	10
3 hits	76	203	449	589	638
5 hits	9	51	216	384	498

Table 3. Number of hits (detects) as a function of SNR

Fig. 13 shows the response of the chip to the speech utterance 'zero' by a female speaker. The pitch in the voiced region can be directly estimated from the sampled interval voltage.

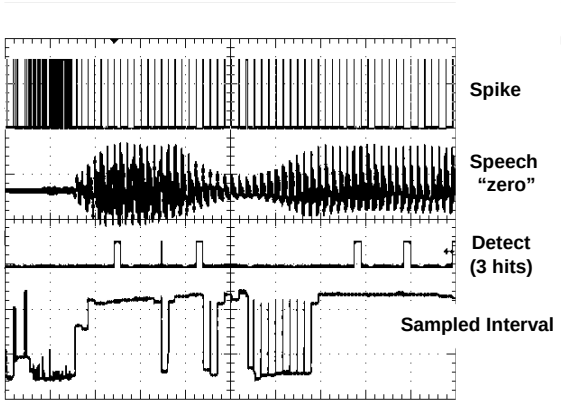

Fig. 13. Response of the chip to a speech utterance ('zero')

V. Conclusion

We have designed and fabricated a time-domain based, low-power, analog VLSI chip that is capable of detecting periodicities in the envelope of acoustic signals. Our experiments show detections for signal-to-noise ratios as low as –5 dB.

Acknowledgment

We thank Shihab Shamma for early discussions of this work and the Signal Systems Corporation for providing various acoustic recordings during the design period of this project. We also thank the MOSIS fabrication service for its continuing support in providing fabrication facilities. This work was supported by a contract with DARPA (Air Coupled Microsensors -0001400C0315).

References

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