

VLSI Implementation of an Energy-Aware Wake-Up Detector for an Acoustic Surveillance Sensor Network

DAVID H. GOLDBERG and ANDREAS G. ANDREOU

Johns Hopkins University

PEDRO JULIÁN

Universidad Nacional del Sur

PHILIPPE O. POULIQUEN

Johns Hopkins University

and

LAURENCE RIDDLE and RICH ROSASCO

Signal Systems Corporation

We present a low-power VLSI wake-up detector for a sensor network that uses acoustic signals to localize ground-based vehicles. The detection criterion is the degree of low-frequency periodicity in the acoustic signal, and the periodicity is computed from the “bumpiness” of the autocorrelation of a one-bit version of the signal. We then describe a CMOS ASIC that implements the periodicity estimation algorithm. The ASIC is fully functional and its core consumes 835 nanowatts. It was integrated into an acoustic enclosure and deployed in field tests with synthesized sounds and ground-based vehicles.

Categories and Subject Descriptors: B.7.1 [**Integrated Circuits**]: Types and Design Styles—*Algorithms implemented in hardware*; C.3 [**Special-Purpose and Application-Based Systems**]—*Signal processing systems*

General Terms: Algorithms, Performance, Design

Additional Key Words and Phrases: Wake-up, vehicle detection, power management, acoustic surveillance, sensor networks, periodicity, VLSI implementation

Portions of this article appeared in the proceedings of IPSN 2004. This work was supported by DARPA/ONR contract N00014-00-C-0315. Chip fabrication was provided by MOSIS. The authors acknowledge the team of engineers at Signal Systems Corporation who assisted in data collection in the field tests.

Authors' addresses: D. H. Goldberg, Department of Physiology and Biophysics, Weill Medical College of Cornell University, New York, NY 10021; email: dhgolberg@ieee.org; A. G. Andreou, and P.O. Pouliquen, Dept. of Electrical and Computer Engineering, Johns Hopkins University, Baltimore, MD 21218; P. Julián, Departamento de Ingeniería Eléctrica y Computadoras, Universidad Nacional del Sur, Bahía Blanca, CP 8000, Argentina; L. Riddle and R. Rosasco, Signal Systems Corporation, Severna Park, MD 21146.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701, USA, fax +1 (212) 869-0481, or permissions@acm.org.
© 2006 ACM 1550-4859/06/1100-0594 \$5.00

1. INTRODUCTION

Sensor network nodes are subject to strict power budgets, as dictated by the need to prolong battery life. This requirement necessitates power-conscious design from high-level algorithms down to circuit implementation. One way to reduce power consumption is to employ a power management strategy. For a surveillance application, one effective power management strategy requires the nodes to operate at full functionality only when a novel object is present, and to persist in a “sleep” state for the remainder of the time. In the sleep state, the node must act as a “wake-up” detector—detecting novel objects, and then arousing the node to full functionality. It is crucial that the wake-up subsystem consume very little power relative to the system as a whole if the power management strategy is to be effective.

1.1 Overview

In this article, we describe a wake-up detector for an acoustic surveillance sensor network. This sensor network detects and localizes ground-based vehicles such as jeeps and tanks. The wake-up criterion is based on the presence of low-frequency periodicity in the acoustic signal, a feature that is characteristic of sounds generated by vehicle engines.

In Section 2, we summarize a maximum-likelihood approach to periodicity estimation and detection. In Section 3, we describe an algorithm for periodicity estimation and detection that maps to a low-power VLSI implementation. The periodicity is determined from the “bumpiness” of the autocorrelation over a specified range of lags. We relate this algorithm to the maximum likelihood estimator for periodicity. In Section 4, we describe the application specific integrated circuit (ASIC) that implements the wake-up detection algorithm. In Section 5, we show experimental results from field tests with synthesized sounds and ground-based vehicles. In Section 6, we summarize the article and discuss the factors that limit the performance of the system.

1.2 Previous Work

The application of sensor networks to vehicle detection and tracking is an area of active research [Li et al. 2002; Ding et al. 2004; Arora et al. 2004; He et al. 2002]. Detection has been done on the basis of radar, magnetic, video, seismic, and acoustic signals. Some advantages of acoustic detection include robust sensors, low cost, low power, and freedom from line-of-sight restrictions.

The inclusion of a wake-up detector on a sensor network node makes many power management schemes possible. In the simplest scheme, each node is responsible for waking itself up. In a more complex scheme, nodes are divided into two classes—sentries and nonsentries. Sentry nodes continuously monitor the environment for novel stimuli, while non-sentry nodes periodically listen for a wake-up signal from the sentry nodes [Hui et al. 2003]. In yet another approach, the nodes sleep on a rotating schedule [Cao et al. 2005]. Recently, Gu and Stankovic [2005] have proposed powering RF wake-up detectors with energy from the stimulus itself. This scheme cannot be used for an acoustic

wake-up function, but could be useful for nodes that receive wake-up signals from sentry nodes.

Periodicity estimation and detection have been actively researched in the speech processing community for decades [Rabiner et al. 1976]. The advent of sensor networks technology, particularly networks that analyze acoustic and seismic signals, has widened interest in this topic. There are several ways to assess the periodicity of a signal; for a recent review of periodicity estimation techniques, see Klapuri [2000]. Perhaps the most straightforward technique is to compute the Fourier transform and search for a peak. Because transforming to the frequency domain can be computationally expensive, time domain methods have been developed and adopted. One such method is envelope periodicity detection, where the fundamental frequency is found from the period of the signal envelope. An analog VLSI implementation of this method has recently been reported [Abdalla and Horiuchi 2005]. In this article, we employ a technique based on autocorrelation [Rabiner 1977].

2. MAXIMUM LIKELIHOOD ESTIMATOR FOR PERIODICITY

The performance of an estimation algorithm can be gauged by comparison with the maximum likelihood estimator (MLE), which can be thought of as a “gold standard.” In this section, we describe the MLE for periodicity, which can be combined with a threshold to make a periodicity detector. In the next section, we will present an alternative periodicity measure that maps to a low-power VLSI implementation and compare it to the MLE for periodicity.

The problem of periodicity estimation is closely related to that of pitch estimation. Pitch estimation asks the question, “What is the period of the signal,” whereas periodicity estimation asks the question, “What is the degree of periodicity in the signal?” Much work has been done on pitch estimation, particularly in the context of speech recognition applications [Rabiner et al. 1976]. A maximum likelihood approach has been developed by several authors [Noll 1969; Wise et al. 1976; Friedman 1977]. We will briefly summarize that approach, and then we will relate it to the MLE for periodicity.

In a pitch estimation task, the sound is assumed to consist of a distinct repeating pattern; the goal is to extract the length of this pattern. We start with a received signal $x[k]$, and we analyze the signal in blocks of K samples. The block extends from $x[0]$ to $x[K-1]$, and it is assumed that samples for $k > K-1$ are available from the next block. The received signal can be decomposed into a periodic signal $s[k]$ plus noise $n[k]$:

$$x[k] = s[k] + n[k]. \quad (1)$$

The signal $s[k]$ can be considered a repeating version of a subsegment $q[k]$, which is P samples long and repeated $L = K/P$ times¹ such that

$$s[k] = q[k \bmod P], \quad (2)$$

¹For simplicity, we will only describe the case where the length of $x[k]$ is an integral number of periods ($K = LP$).

where $q[k]$ and P are unknown and must be estimated by maximizing the likelihood of the received signal. For a given value of P , $q_P[k]$ is given by

$$q_P[k] = \frac{1}{L} \sum_{l=0}^{L-1} x[k + lP]. \quad (3)$$

The energy of $q_P[k]$ is given by²

$$\sigma_Q^2(P) = \frac{1}{P} \sum_{k=0}^{P-1} (q_P[k])^2. \quad (4)$$

The MLE for the period \hat{P} is given by the value of P that maximizes $\sigma_Q^2(P)$.

Once we obtain the MLE for the period, we only need go one step further to find the MLE for periodicity. Our estimate of the underlying repeating sequence is given by

$$\hat{q}[k] = q_{\hat{P}}[k]. \quad (5)$$

We can use Equation 2 to obtain

$$\hat{s}[k] = \hat{q}[k \bmod \hat{P}]. \quad (6)$$

The degree of periodicity, in a maximum likelihood sense, is given by the estimated signal-to-noise ratio:

$$\mathcal{P}_{\text{MLE}} = \frac{\sigma_{\hat{S}}^2}{\sigma_N^2} = \frac{\sum_{k=0}^{K-1} (\hat{s}[k])^2}{\sum_{k=0}^{K-1} (x[k] - \hat{s}[k])^2}. \quad (7)$$

In order to demonstrate the algorithm, we apply it to recorded data. The data consists of an ambient recording provided by Signal Systems Corporation (Severna Park, Md.) and recordings drawn from the Army Research Laboratory's Acoustic-Seismic Classification Identification Data Set (ACIDS). A detailed description of this data set is given in Munich [2004]. The data consists of recordings from light wheeled, heavy wheeled, light tracked, and heavy tracked vehicles in normal, desert, and arctic conditions. The vehicles were traveling at constant speeds ranging from 5 km/h to 40 km/h, with the closest point of approach ranging from 25 m to 100 m. Microphone recordings were high-pass filtered (cut-off frequency of 25 Hz) to reduce the effects of wind noise and low-pass filtered (cut-off frequency of 400 Hz) to prevent aliasing, and then digitized by sampling ~ 1024 Hz and quantizing with 16-bit precision. Each recording was approximately 200 seconds long. A sample rate of 1024 Hz is adequate because the audible harmonics of the engine noise are on the order of 100 Hz. The analysis block size must be short enough to limit the response latency, but long enough to provide a robust estimate of the periodicity. We found that a block size of 1 second ($K = 1024$ samples) satisfies this trade-off.

Figure 1 demonstrates analysis with the MLE method. Figures 1(a) and (b) show the periodicity SNR (in dB) as a function of time (x axis; 1 second-blocks) and candidate periods (y axis) for the ambient and vehicle recording,

²This notation assumes that $s[k]$ is zero mean.

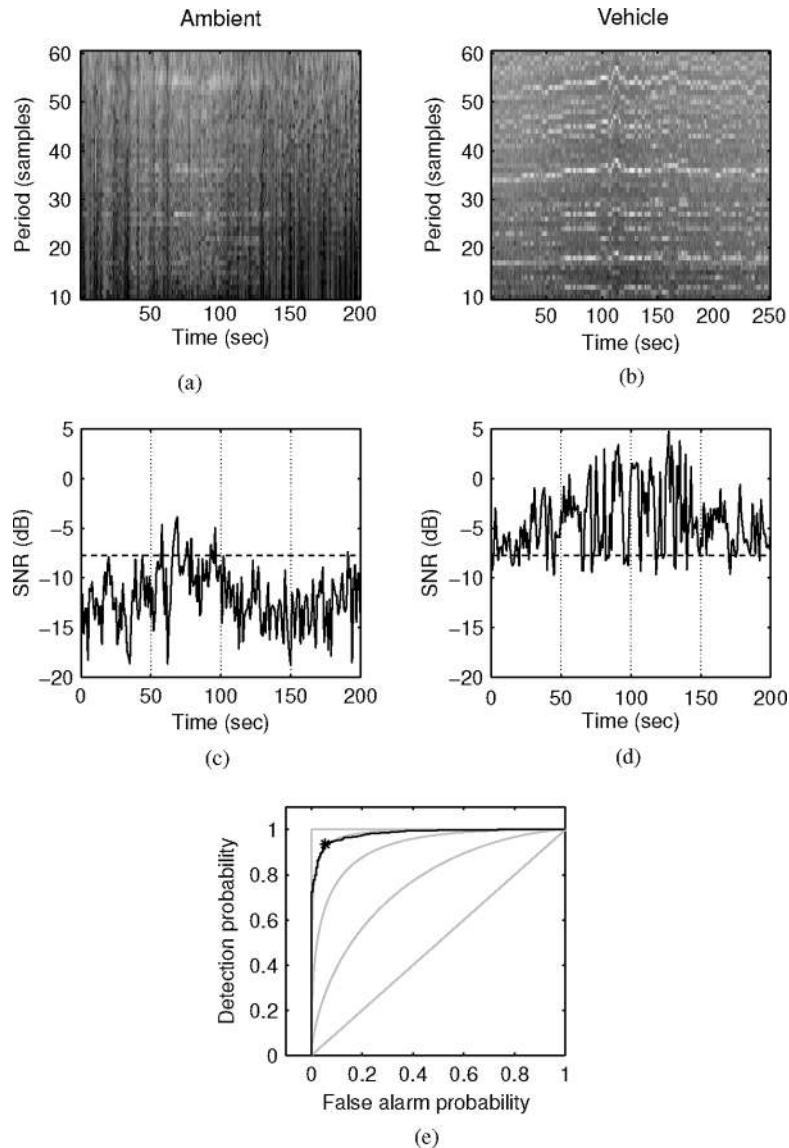


Fig. 1. Detection results for the measure based on the maximum likelihood estimator for periodicity. The periodicity measure was computed for 1 second blocks ($K = 1024$ samples). Periods covering the range of [10, 60] samples were examined. (a) Periodicity estimates over all periods for an ambient recording. White corresponds to 5 dB and black corresponds to -40 dB. (b) Periodicity estimates over all periods for a vehicle recording. (c) MLE for periodicity (the periodicity estimate at \hat{P}) for an ambient recording. The dotted line corresponds to the threshold that gives the minimum probability of error on the entire data set. (d) MLE for periodicity for a vehicle recording. (e) ROC curve for MLE periodicity measure generated from 16 vehicle recordings and one ambient recording. Lines of constant d' , corresponding to $d' = [0 \ 1 \ 2 \ 3 \ \infty]$, shown in gray. The point on the ROC curve that minimizes the probability of error is given by the *.

respectively. During each block, the period that maximizes the SNR is the MLE for the pitch period. The maximum SNR is the MLE for periodicity, which is plotted in Figures 1(c) and (d).

To rigorously quantify the discrimination ability of the detector, we constructed a receiver operating characteristic (ROC) curve [Van Trees 2001]. The ROC curve is generated by systematically varying the detection threshold and plotting the detection probability versus false alarm probability. The curve, shown in Figure 1(e), was compiled from the ambient recording, and 16 vehicle recordings, from ACIDS. The area under the ROC curve is 0.977 (where 0.5 is chance detection and 1 is perfect detection), and the minimum probability of error $P[\text{miss}] + P[\text{false alarm}]$ is 11.9%, which corresponds to a threshold of -7.71 dB.³ A detector's discriminability can be characterized independently of the threshold by the discriminability index d' , which quantifies the distance between the probability distribution of the periodicity conditioned on a vehicle's presence and the probability distribution of the periodicity conditioned on a vehicle's absence. Chance discrimination corresponds to $d' = 0$ and perfect discrimination corresponds to $d' = \infty$. The MLE for periodicity gives a nearly constant value of $d' = 3$.

3. ALTERNATIVE PERIODICITY MEASURE

The MLE for periodicity is computationally intensive, as it requires the estimation of the period of the underlying signal P , which in turn requires the computation of $q_P[k]$ for each value of P . In this section, we develop an alternative periodicity measure (PM), which does not require the estimation of P . Because we have specific prior knowledge of the problem—the periodic signal generated by the vehicle engine is embedded in outdoor ambient noise—we can make assumptions that dramatically reduce the computational requirements without sacrificing detection performance.

A straightforward way to assess the periodicity of a signal $x[k]$ is to examine the signal's autocorrelation function (ACF), given by

$$R_{XX}[n] = \frac{1}{K} \sum_{k=0}^{K-1} x[k]x[k+n]. \quad (8)$$

It is useful to normalize the ACF by the signal power:

$$\bar{R}_{XX}[n] = R_{XX}[n]/R_{XX}[0]. \quad (9)$$

Figures 2(a) and (b) depict the normalized ACF for 1 second segments of ambient and vehicle recordings, respectively. The ACF of the noise-like ambient recording has peaks that are far apart, and the peaks decay as the lag increases. The ACF of the periodic vehicle recording has peaks that are close together and do not decay. We formulate an alternative PM that quantifies this distinction: First we compute the discrete derivative of the normalized ACF (Figures 2(c)

³It should be kept in mind that minimizing the probability of error will not necessarily correspond to optimal performance. The true optimal operating point will be determined by the relative costs of misses and false alarms.

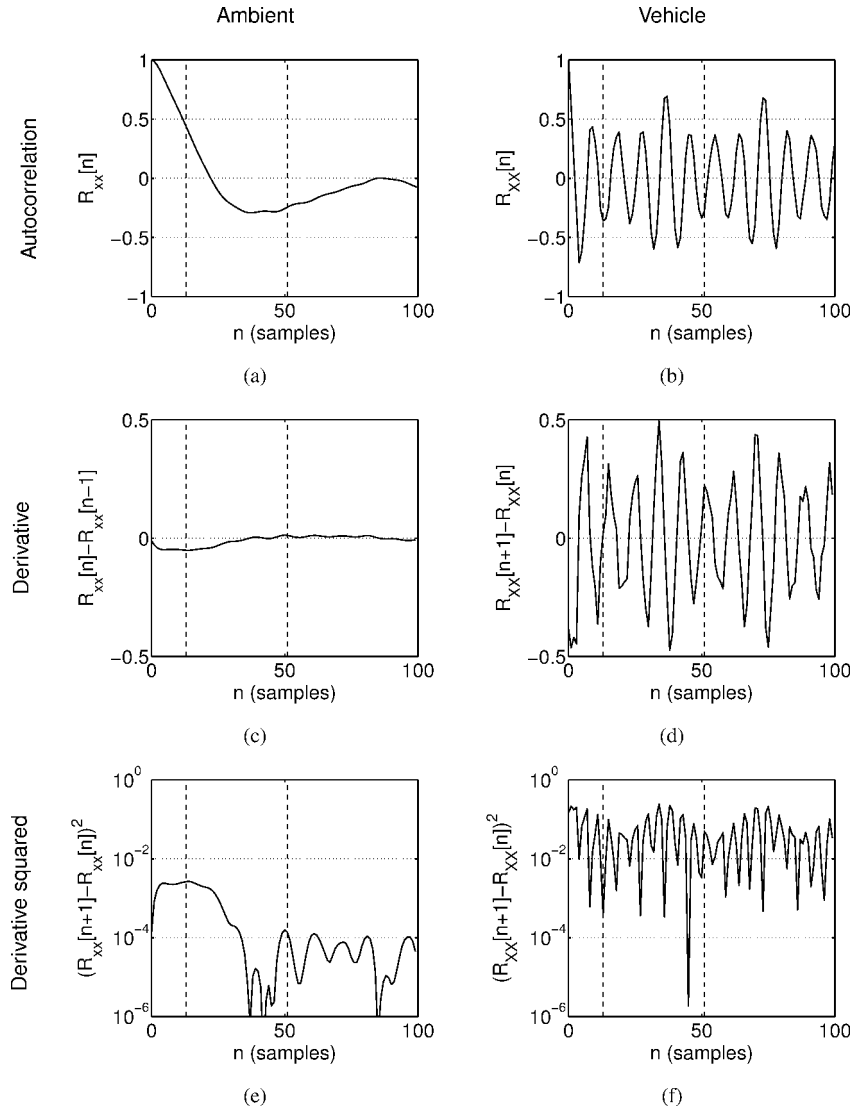


Fig. 2. Development of an alternative periodicity measure. Left column: ambient recording. Right column: Vehicle recording. (a) and (b) Normalized autocorrelation functions. The dashed vertical lines correspond to $N_{\min} = 13$ samples and $N_{\max} = 52$ samples. (c) and (d) Derivative of normalized autocorrelation function. (e) and (f) Square of derivative of normalized autocorrelation function.

and (d)), we square them (Figures 2(e) and (f)), and then we sum over a range of lags of interest, $[N_{\min}, N_{\max}]$. This process is summarized by the equation

$$\mathcal{P}_{\text{alt}} = \sum_{n=N_{\min}}^{N_{\max}-1} (\bar{R}_{XX}[n+1] - \bar{R}_{XX}[n])^2. \quad (10)$$

$N_{\max} = 52$ samples was chosen because it corresponds to $f_{\min} = f_{\text{samp}}/N_{\max} \approx$

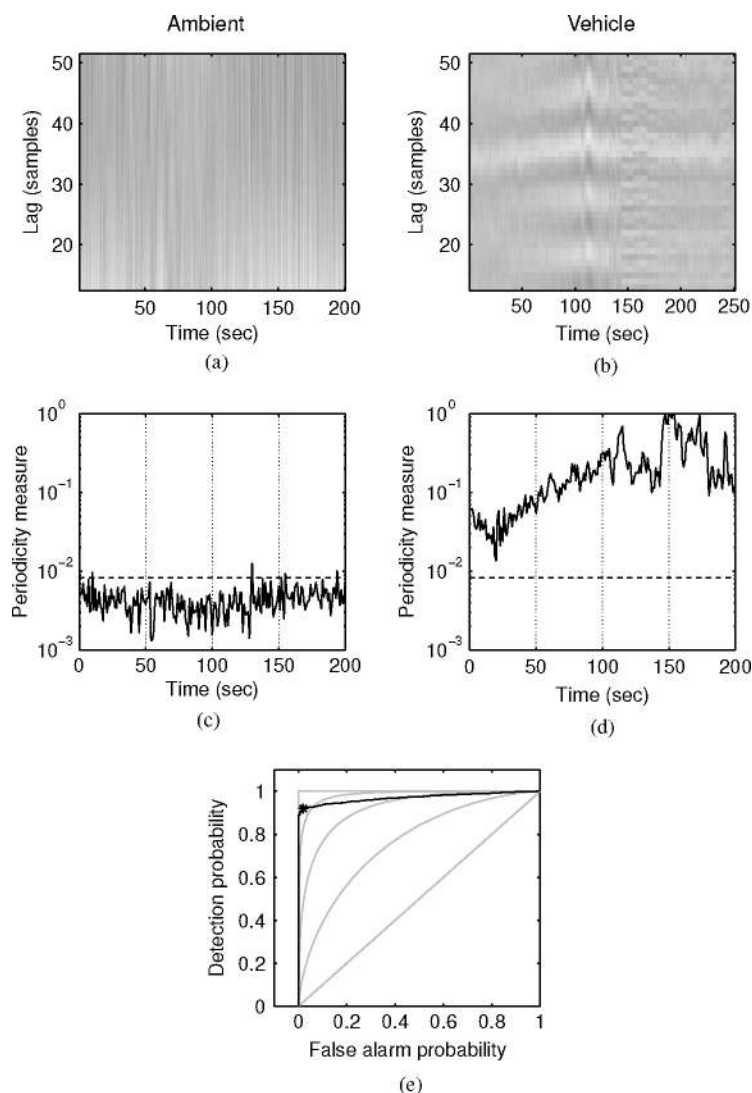


Fig. 3. Detection results for the alternative periodicity measure. $N_{\min} = 13$ samples and $N_{\max} = 52$ samples. (a) Autocorrelation function for an ambient recording. White corresponds to +1 and black corresponds to -1. (b) Autocorrelation function for vehicle recording. (c) Periodicity measure for an ambient recording. The detection threshold value was set to the value that minimizes the probability of error. (d) Periodicity measure for vehicle recording. (e) ROC curve for simplified periodicity measure generated from 16 vehicle recordings and one ambient recording. Same format as Figure 1(e).

20 Hz, which is the expected low-frequency cut-off of the microphone. $N_{\min} = 13$ samples corresponds to $f_{\max} = f_{\text{samp}}/N_{\min} = 75$ Hz, which captures the first two harmonics of a sound that has a fundamental frequency of 25 Hz.

Figure 3 shows the analysis of the two typical signals with the alternative PM. Figures 3(a) and (b) show the ACF as a function of time (x axis; 1 second

blocks) and lags (y axis). Figures 3(c) and (d) show the PM as a function of time. The ROC curve for this measure is shown in Figure 3(e)). The area under the ROC curve is 0.970 and the minimum probability of error is 10.1%, which corresponds to a threshold of 8.34×10^{-3} . This performance is comparable to the MLE detector. It should be noted that the ROC curve in this case is somewhat skewed, enabling very high detection probabilities with essentially no false alarms.

3.1 Relationship Between Alternative Periodicity Measure and Maximum Likelihood Estimator for Periodicity

To demonstrate the relationship between the PM and the MLE for periodicity, we employ a simple example: a pure sine wave embedded in noise with a bandpass power spectrum. The sine wave is given by

$$s[k] = \sqrt{2}\sigma_S \sin[2\pi(f_S/f_{\text{samp}})k]. \quad (11)$$

The frequency f_S could correspond to that of a prominent harmonic in an engine's acoustic signature. This signal has the autocorrelation

$$R_{SS}[k] = \frac{1}{T_S f_{\text{samp}}} \sum_{k=0}^{T_S f_{\text{samp}}} \sqrt{2}\sigma_S \sin[2\pi(f_S/f_{\text{samp}})k] \\ \times \sqrt{2}\sigma_S \sin[2\pi(f_S/f_{\text{samp}})(k+n)] \quad (12)$$

$$= \sigma_S^2 \cos[2\pi(f_S/f_{\text{samp}})n], \quad (13)$$

where $T_S = 1/f_S$. The bandpass characteristic of the ambient noise is due to the combination of the low frequency cut-off of the microphone and the $1/f^n$ characteristic inherent in wind noise [Morgan and Raspet 1992]. The autocorrelation of the ambient noise is

$$R_{NN}[k] = \sigma_N^2 e^{-(\alpha/f_{\text{samp}})|n|} \cos[2\pi(f_N/f_{\text{samp}})n], \quad (14)$$

where f_N is the center frequency and α is a parameter that determines the bandwidth.

Recall that the received signal is given by $x[k] = s[k] + n[k]$. Assuming that the vehicle signal and ambient noise are independent and zero-mean, the autocorrelation of X is

$$R_{XX}[n] = R_{SS}[n] + R_{NN}[n] \quad (15)$$

$$= \sigma_S^2 \cos[2\pi(f_S/f_{\text{samp}})n] + \sigma_N^2 e^{-(\alpha/f_{\text{samp}})|n|} \cos[2\pi(f_N/f_{\text{samp}})n] \quad (16)$$

$$\bar{R}_{XX}[n] = \frac{1}{\sigma_S^2 + \sigma_N^2} \{ \sigma_S^2 \cos[2\pi(f_S/f_{\text{samp}})n] + \sigma_N^2 e^{-(\alpha/f_{\text{samp}})|n|} \cos[2\pi(f_N/f_{\text{samp}})n] \}. \quad (17)$$

The MLE for periodicity in this case is $\mathcal{P}_{\text{MLE}} = \sigma_S^2/\sigma_N^2$ and the proposed PM is given by plugging Equation 17 into Equation 10. While analytical expressions for the PM are not particularly informative, this simple model enables us to vary the parameters and explore how the PM compares to the MLE, as shown in Figure 4(a). We see that for the simple model and the parameters we

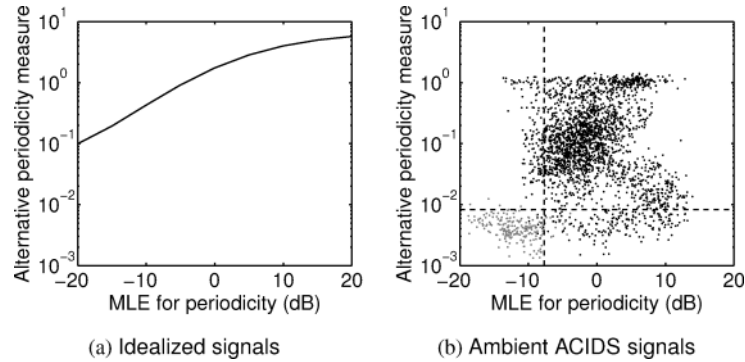


Fig. 4. Comparison of alternative periodicity measure (PM) with MLE for periodicity. (a) Plot of alternative PM versus MLE for a pure sine wave (σ_S^2 variable, $f_S = 100$ Hz) embedded in bandpass noise ($\sigma_N^2 = 1$, $\alpha = 10$ rad/sec and $f_N = 10$ Hz). (b) Plot of PM versus MLE for the ambient signal (gray dots) and the ACIDS signals (black dots). The dashed lines correspond to the thresholds that minimize the probabilities of error.

have chosen, the relationship between the MLE and the PM is monotonic. As the signal-to-noise ratio decreases, the PM will fall below the threshold and detection fails. A comparison for real signals and noise is shown in Figure 4(b). The deviation from the ideal monotonic relationship can be attributed to the complexity of the recorded signals as compared to the idealized model signals.

3.2 Algorithm Simplifications for Implementation

The alternative PM can be simplified a great deal and still perform satisfactorily in the detection task. These simplifications have been made with a hardware implementation in mind, whether it be in an embedded processor or in a full-custom ASIC.

Up until this point, the acoustic signals have been quantized with 16-bit precision. It has been shown, however, that the ACF of an infinitely clipped signal $\tilde{x}[n]$ is related to the ACF of the original signal $x[n]$ by

$$R_{\tilde{X}\tilde{X}}[n] = \frac{2}{\pi} \sin^{-1}(R_{XX}[n]) \quad (18)$$

as the infinitely clipped signal retains the original signal's periodic structure [Weinreb 1963; Van Vleck and Middleton 1966]. This implies that one-bit precision is sufficient to give acceptable results for the detection task.

The use of one-bit input signals greatly simplifies the hardware implementation. For example, provided that we encode the signal with zeros and ones, the multiply operation in the correlation computation is reduced to an XNOR. The one-bit signal has the additional advantage that the normalization step in the ACF computation (the division by $R_{XX}[0]$) is eliminated, as the amplitude information in the signal has already been discarded. One caveat is that we must ensure that the signal is centered around the quantizer threshold point. The PM computation can be further simplified by computing the sum of the absolute value of the discrete differences rather than the sum of the squares.

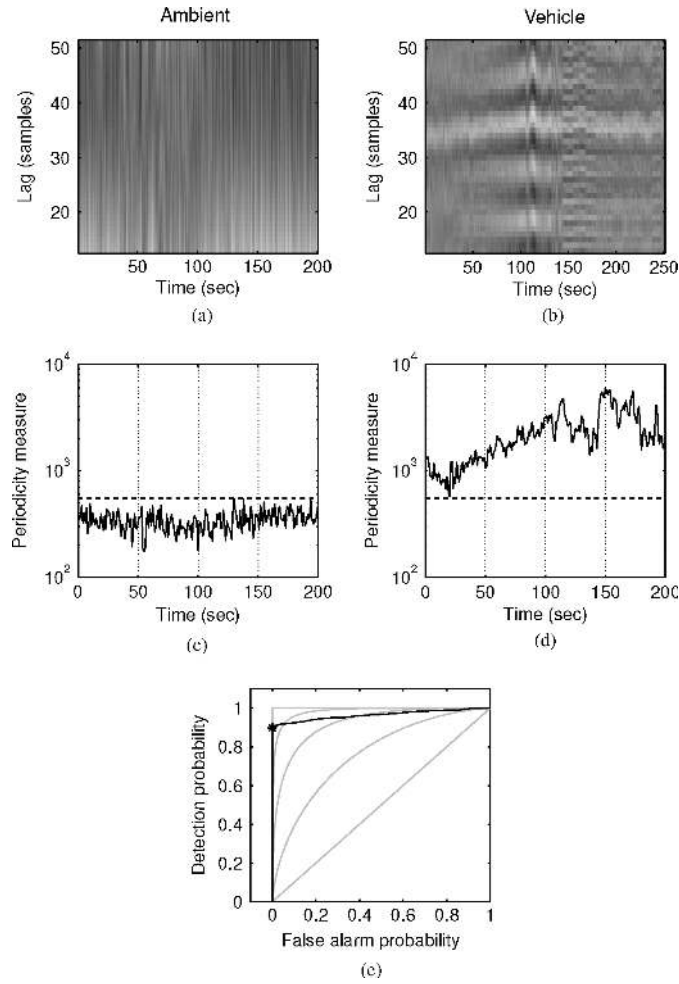


Fig. 5. Detection results for the simplified periodicity measure. Same parameters as Figure 3. (a) Autocorrelation function for an ambient recording. White corresponds to 1024 and black corresponds to 0. (b) Autocorrelation for vehicle recording. The detection threshold value was set to 551. (c) Periodicity measure for an ambient recording. (d) Periodicity measure for vehicle recording. (e) ROC curve for simplified periodicity measure generated from 16 vehicle recordings and one ambient recording. Same format as Figure 1(e).

The PM equations become

$$\tilde{R}_{\bar{x}\bar{x}}[n] = \sum_{k=0}^{K-1} \overline{x[k] \oplus x[k+n]} \quad (19)$$

$$\mathcal{P}_{\text{simp}} = \sum_{n=N_{\min}}^{N_{\max}} |\tilde{R}_{\bar{x}\bar{x}}[n+1] - \tilde{R}_{\bar{x}\bar{x}}[n]| \quad (20)$$

where \oplus represents the XOR operation and the bar represents negation.

Figure 5 shows the results of the simplified PM. The format of the figure is the same as that of Figure 3. The area under the ROC curve is 0.964 and




Fig. 6. Micrograph of the wake-up detector CMOS ASIC.

the minimum probability of error is 10.3%, which corresponds to threshold of 551. We see that the use of the one-bit signal has no significant effect on the discrimination ability of the detector.

4. VLSI IMPLEMENTATION

We designed a full-custom CMOS ASIC to implement the simplified wake-up detection algorithm described in the previous section. The chip was fabricated on a $3 \text{ mm} \times 1.5 \text{ mm}$ die in a $0.5 \mu\text{m}$ -process available from the MOSIS service. It operates from a power supply voltage of 3.3 V. A micrograph of the chip is shown in Figure 6.

Figure 7(a) shows a block diagram of the ACF accumulation circuitry and Figure 7(b) gives a timing diagram for a single detection operation. The 1024 clock cycle detection operation consists of two main phases: a 983 clock cycle ACF phase and a 41 clock cycle PM phase. The chip constantly receives a one-bit input stream, regardless of its operation phase. A 52-sample history of the input is stored in an input register INP. The 40 most recent samples (INP[0 : 39]) in the register are correlated with the oldest sample (INP[51]); this realizes a lag range of [13, 52]. As discussed in the previous section, the correlation consists of an XNOR operation. During the ACF phase, the output of each correlation operation goes to a 10-bit accumulator, ACC. After the ACF has been computed for 983 clock cycles, the PM phase begins. A state machine generates select signals such that one even ACF lag and one odd ACF lag are driven onto their respective buses in the proper sequence, enabling the computation of the discrete derivative of the ACF. The results are accumulated in the PMACC register. At the end of the PM phase, the PM is compared to a user-settable detection threshold. If the PM exceeds the threshold, a wake-up signal is generated. At this point, the ACF registers are reset, and the detection operation begins anew.

The wake-up ASIC was connected to a microphone and signal conditioning circuitry (including a comparator for one-bit A/D conversion) and tested in a laboratory setting. The chip consumes $6.3 \mu\text{W}$ during operation. Because

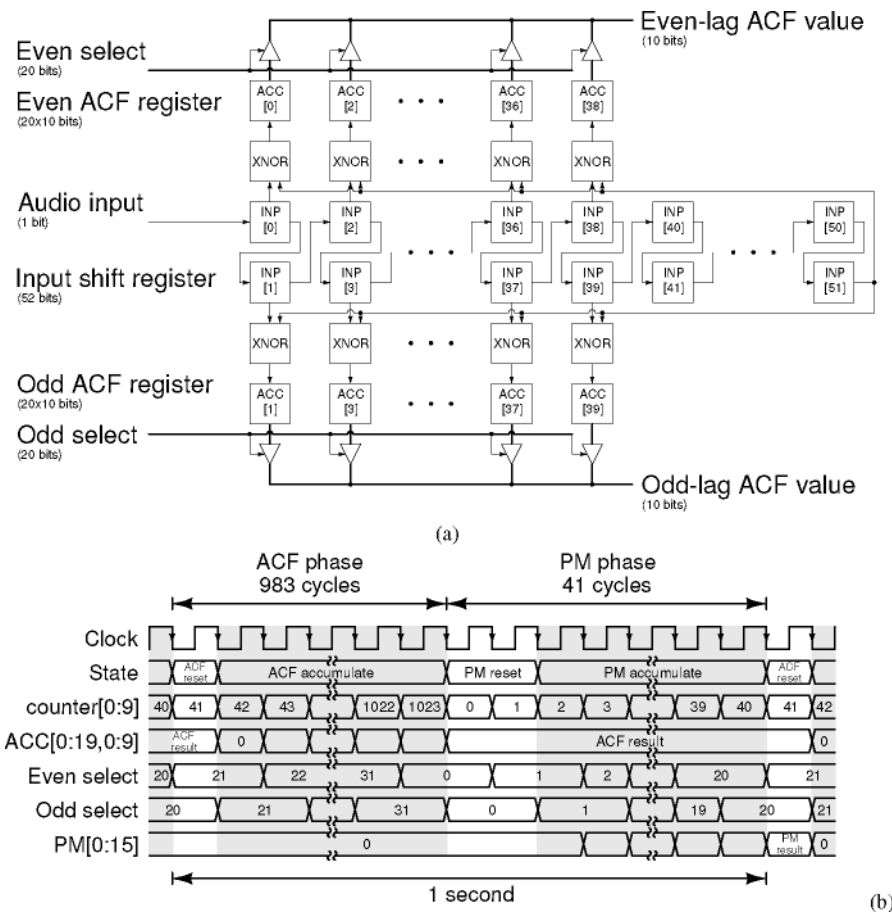


Fig. 7. VLSI implementation of the wake-up detection algorithm. (a) Block diagram of the circuitry that computes the autocorrelation. INP represents the input register and ACC represents a 10-bit accumulator and corresponding register. (b) Timing diagram of the detection operation. The two main phases of the operation are the autocorrelation function (ACF) accumulation phase and the periodicity measure (PM) accumulation phase.

we use separate power supply pins for the I/O pads and the core, we can divide the total power consumption into its constituent parts. The I/O pads consume $5.5 \mu\text{W}$ and the core consumes 835 nW . The large power consumption of the pads is attributable to two factors: First, the clock input switches at 32 kHz , and is subsequently internally divided down to 1 kHz . This design decision was made in order to facilitate integration with a COTS oscillator. Second, approximately 40 nonessential pads were included for debugging purposes. Therefore, it is reasonable to assume that a next generation chip in the same process would consume $1 \mu\text{W}$. This power consumption level is far smaller than that of the microphones and signal conditioning circuitry, which draw $300 \mu\text{W}$. At these levels, the system will run for well over a year on 3 AA batteries.




Fig. 8. Acoustic surveillance unit node. (a) Photograph of acoustic surveillance unit enclosure. (b) Simplified block diagram of key components.

5. FIELD TEST RESULTS

The wake-up ASIC was integrated into an acoustic surveillance unit (ASU) enclosure along with an array of four Knowles SiSonic MEMS microphones, signal conditioning circuitry, two bearing estimation ASICs [Stanačević and Cauwenberghs 2005; Julián et al. 2006], and microprocessor and communication modules. A photograph of the enclosure is depicted in Figure 8(a) and a simplified block diagram of the ASU node is depicted in Figure 8(b). One of the four microphones is connected to the wake-up ASIC. When the wake-up signal

is asserted, the bearing estimation ASICs are activated. The bearing estimates are processed by the microprocessor, and ultimately transmitted to a central processor where they are combined with other bearing estimates to localize the source. Several wake-up configurations are possible depending on the relative cost of misses and false alarms—a wake-up signal from a signal node could wake up the network, or wake-up signals from all nodes could be required for wake-up. Further details about the ASU and its behavior in a network setting are available in Cauwenberghs et al. [2005].

We conducted two field tests, one in a public park in Severna Park, Maryland with synthesized sounds (field test 1), and another at Aberdeen Proving Ground, Aberdeen, Maryland with a selection of ground-based military vehicles (field test 2). In both tests, the detection threshold was set to 1024. This setting corresponds to the point at which the ROC curve in Figure 5(e) departs from the y axis and attains the highest value of $P[\text{detection}]$ for which $P[\text{false alarm}] = 0$.

5.1 Field Test 1: Synthesized Sounds

In the first field test, synthesized sounds were played from a subwoofer placed in an open field. In this setting, we had strict control of the frequency content and magnitude of the sounds. In all trials, the ASU was placed 30 feet from the subwoofer at an angle of 90° . We first performed a series of trials with a signal consisting of three simultaneous time-varying, harmonically related tones (125 Hz, 150 Hz, 175 Hz) to model engine noise (Figures 9(a) and (d)). The wake-up detector was reliably triggered down to a narrowband SNR of 13 dB (Figures 9(b) and (e)). In trials with broadband white noise, the loudest possible volume (50 dB SPL) did not elicit a trigger (Figures 9(c) and (f)).

5.2 Field Test 2: Ground-Based Vehicles

In the second field test, vehicles were driven around a $662 \text{ m} \times 108 \text{ m}$ oval-shaped track, and three ASUs were placed at various points around the oval. One ASU, at one end of the oval, contained the wake-up detector. When the detector was triggered, the bearing estimation circuit on all of the ASUs localized and tracked the sound source. Table I summarizes the wake-up results on an assortment of vehicles in terms of the maximum distance that elicited a sustained detection. Also, numerous unscripted targets were detected, such as helicopters, powerboats, and trucks. At one instance, an F/A-18 fighter jet flew overhead at approximately 10,000 ft and a detection was not elicited. Unfortunately, the logistics of the field test did not permit a more detailed evaluation of the detector's performance.

6. CONCLUSION

We have presented a wake-up detector algorithm based on the degree of periodicity in acoustic signals generated by ground-based vehicles. This algorithm was developed with a low-power VLSI implementation in mind. We designed and tested a CMOS ASIC that implements this algorithm. The core of the ASIC consumes 835 nW, and we expect the next-generation entire chip to consume

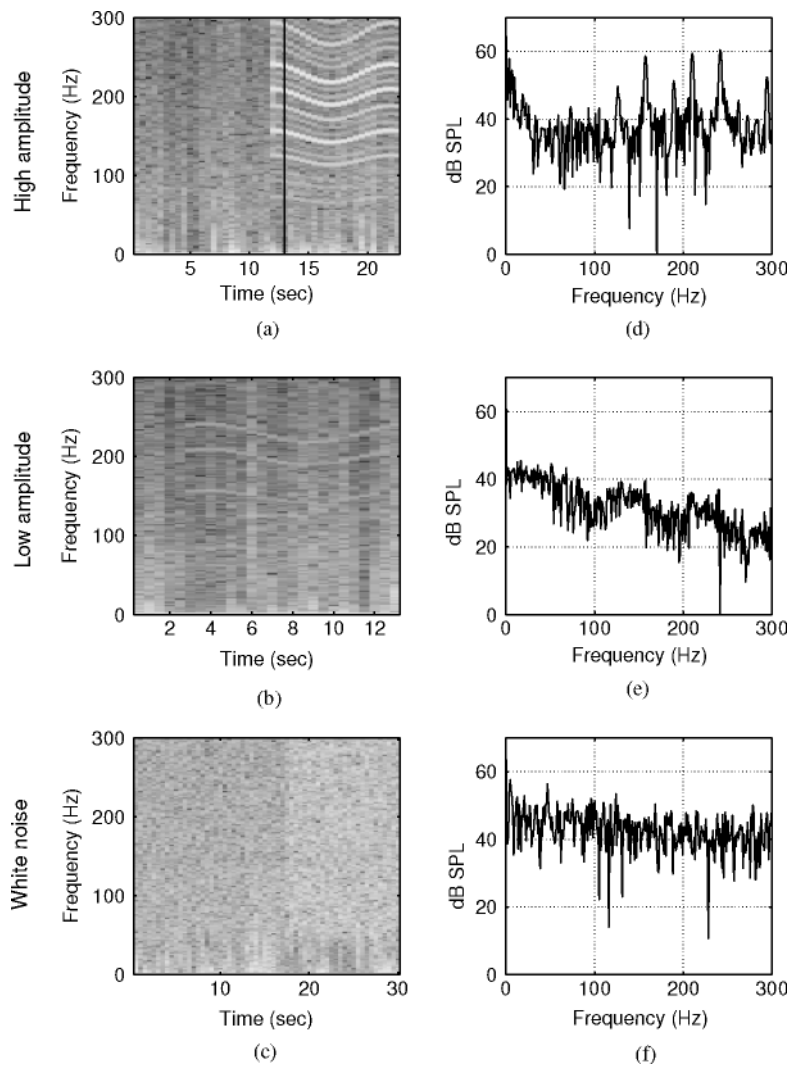


Fig. 9. Field test 1 results for a signal with three simultaneous time-varying, harmonically related tones (125 Hz, 150 Hz, 175 Hz). (a) Spectrogram representation of a loud sound. The vertical line indicates the moment that the detection signal is triggered. (b) Spectrogram representation of a lower amplitude sound, corresponding to the limits of the system's operation. (c) Spectrogram representation of ambient noise. (d) Power spectral density for loud sound over $t = [11, 12]$ seconds. (e) Power spectral density for lower amplitude sound over $t = [2, 3]$ seconds. (f) Power spectral density for lower amplitude sound over $t = [2, 3]$ seconds.

on the order of $1 \mu\text{W}$. The wake-up detector has been field tested and its performance is robust.

The performance of the wake-up system is primarily limited by the signal acquisition and conditioning subsystem rather than the algorithm itself. Because the algorithm requires a one-bit input, it relies on a high-quality comparator at the output of the microphone. In order to provide a useful signal, the comparator

Table I. Performance of the Wake-Up System on Ground-Base Vehicles

Vehicle	Description	Distance
M60	Heavy Tracked	>500 m
HEMET	Heavy Wheeled	~250 m
M548	Light Tracked	~400 m
HMMWV	Light Wheeled	~55 m

threshold must be able to track the DC level of the acoustic signal. In the case of low-amplitude sounds, the microphone output is on the order of the comparator offset, and the comparator fails to trigger. The design of a low-power, low-offset comparator for integration with a MEMS microphone is a focus of our current research in this area.

As discussed in Section 3, the algorithm's performance is limited under conditions where the SNR is low. This observation draws attention to the algorithm's weakness—detection can be avoided if the engine noise is masked with a loud broadband signal. Our modeling indicates that such a masking sound would have to be several times louder than the engine itself, making the vehicle easily detectable by other means.

REFERENCES

- ABDALLA, H. AND HORIUCHI, T. K. 2005. An analog VLSI low-power envelope periodicity detector. *IEEE Trans. Circuits Syst.* 52, 9 (Sept.), 1709–1720.
- ARORA, A., DUTTA, P., BAPAT, S., KULATHUMANI, V., ZHANG, H., NAIK, V., MITTAL, V., CAO, H., DEMIRBAS, M., GOUDA, M., CHOI, Y., HERMAN, T., KULKARNI, S., ARUMUGAM, U., NESTERENKO, M., VORA, A., AND MIYASHITA, M. 2004. A line in the sand: A wireless sensor network for target detection, classification, and tracking. *Comput. Netw.* 46, 5, 605–634.
- CAO, Q., ABDELZAHER, T., HE, T., AND STANKOVIC, J. 2005. Towards optimal sleep scheduling in sensor networks for rare-event detection. In *4th International Symposium on Information Processing in Sensor Networks (IPSN'05)*. IEEE Press, 20–27.
- CAUWENBERGHS, G., ANDREOU, A., WEST, J., STANACEVIC, M., CELIK, A., JULIÁN, P., TEIXEIRA, T., DEIEHL, C., AND RIDDLE, L. 2005. A miniature low-power intelligent sensor node for persistent acoustic surveillance. In *Unattended Ground Sensor Technologies and Applications VII*, E. M. Carapezza, Ed. *Proceedings of SPIE*, vol. 5796. SPIE, Bellingham, WA, 294–305.
- DING, J., CHEUNG, S.-Y., TAN, C.-W., AND VARAIYA, P. 2004. Signal processing of sensor node data for vehicle detection. In *Proceedings of the 7th International IEEE Conference on Intelligent Transportation Systems*. IEEE, 70–75.
- FRIEDMAN, D. H. 1977. Pseudo-maximum-likelihood speech pitch extraction. *IEEE Trans. Acoustics, Speech, Sig. Proc.* ASSP-25, 3 (June), 213–221.
- GU, L. AND STANKOVIC, J. A. 2005. Radio-triggered wake-up for wireless sensor networks. *Real-Time Syst.* 29, 2-3, 157–182.
- HE, T., KRISHNAMURTHY, S., LUO, L., YAN, T., GU, L., STOLERU, R., ZHOU, G., CAO, Q., VICAIRE, P., STANKOVIC, J. A., ABDELZAHER, T. F., HUI, J., AND KROGH, B. 2002. Vigilnet: An integrated sensor network system for energy-efficient surveillance. *ACM Trans. Sensor Netw.* 2, 1, 1–38.
- HUI, J., REN, Z., AND KROGH, B. H. 2003. Sentry-based power management in wireless sensor networks. In *Proceedings of Information Processing in Sensor Networks*. Lecture Notes in Computer Science, vol. 2634. Springer-Verlag, Berlin, Germany 458–472.
- JULIÁN, P., ANDREOU, A. G., AND GOLDBERG, D. 2006. A low-power correlation-derivative CMOS VLSI circuit for bearing estimation. *IEEE Trans. VLSI Syst.* 14, 2, 207–212.
- KLAPURI, A. 2000. Qualitative and quantitative aspects in the design of periodicity estimation algorithms. In *Proceedings of the European Signal Processing Conference (EURASIP)*, Tampere, Finland.

- LI, D., WONG, K. D., HU, Y. H., AND SAYEED, A. M. 2002. Detection, classification, and tracking of targets. *IEEE Signal Processing Magazine* 19, 2 (Mar.), 17–29.
- MORGAN, S. AND RASPET, R. 1992. Investigation of the mechanisms of low-frequency wind noise generation outdoors. *J. Acous. Society Amer.* 92, 2, 1180–1183.
- MUNICH, M. E. 2004. Bayesian subspace methods for acoustic signature recognition of vehicles. In *Proceedings of the European Signal Processing Conference (EURASIP)* Vienna, Austria.
- NOLL, A. M. 1969. Pitch determination of human speech by the harmonic product spectrum, the harmonic sum spectrum, and a maximum likelihood estimate. In *Proceedings of the Symposium on Computer Processing in Communications*, Vol. 19. J. Fox, Ed. Wiley-Interscience, Chichester, Sussex, UK, 779–797.
- RABINER, L. 1977. On the use of autocorrelation analysis for pitch detection. *IEEE Trans. Acoustics Speech, Sig. Proc.* 25, 1 (Jan.), 24–33.
- RABINER, L. R., CHENG, M. J., ROSENBERG, A. E., AND MCGONEGAL, C. A. 1976. A comparative performance study of several pitch detection algorithms. *IEEE Trans. Acoustics Speech, Sig. Proc.* 24, 5 (Oct.), 399–418.
- STANAČEVIĆ, M. AND CAUWENBERGHS, G. 2005. Micropower gradient flow acoustic localizer. *IEEE Trans. Circuits Syst.* 52, 10 (Oct.), 2148–2157.
- VAN TREES, H. L. 2001. *Detection, Estimation, and Modulation Theory: Part I. Detection, Estimation, and Linear Modulation Theory*. Wiley-Interscience, New York, NY.
- VAN VLECK, J. H. AND MIDDLETON, D. 1966. The spectrum of clipped noise. *Proceedings of the IEEE* 54, 1 (Jan.), 2–19.
- WEINREB, S. 1963. A digital spectral analysis technique and its application to radio astronomy. Ph.D. thesis, MIT, Cambridge, MA.
- WISE, J. D., CAPRIO, J. R., AND PARKS, T. W. 1976. Maximum likelihood pitch estimation. *IEEE Trans. Acoustics Speech, Sig. Proc.* 24, 5 (Oct.), 418–423.

Received February 2005; revised February 2006; accepted August 2006