OPTIMAL AND ROBUST SHOCKWAVE DETECTION AND ESTIMATION

Brian M. Sadler *Laurel C. Sadler* *Tien Pham*
Army Research Laboratory, 2800 Powder Mill Road, Adelphi, MD 20783-1197

ABSTRACT

We consider detection and estimation of aeroacoustic shockwaves generated by supersonic projectiles. The shockwave is an N-shaped acoustic wave. The optimal detection/estimation scheme is considered based on an additive white Gaussian noise model. The introduction of an invertible linear transformation, such as the Fourier transform or the wavelet transform, does not improve detection performance under this model. However, if unknown interference and/or model mismatch is present, linear transforms may be of use. In addition, they may significantly reduce complexity at the cost of sub-optimality. We consider the use of the wavelet transform as a means of detecting the very fast rise and fall times of the shockwave, resulting in a 1-D edge detection problem. This method is effective at moderate to high SNR and is robust with respect to unknown environmental interference that will generally not exhibit singularities as sharp as the N-wave edges.

1. INTRODUCTION

We consider optimal and robust detection and estimation of aeroacoustic shockwaves generated by supersonic projectiles. This problem arises in military, law enforcement, and other cases. It is desired to detect the presence of a bullet or other projectile, and to estimate the parameters of the shockwave. Detection will be useful in a variety of scenarios with application in sniper location as well as on vehicles and aircraft. Of particular interest are robust methods that will work at moderate SNR in the presence of platform noise.

The shockwave is an "N-shaped" wave emanating in the form of an acoustic cone trailing the projectile [1]. Letting A denote the peak amplitude of the shockwave then [2]

$$A \propto d \left(\frac{c}{v}\right)^{1/4} (l)^{-1/4}(x)^{-3/4}, \tag{1}$$

where d, v, and l are the projectile diameter, velocity, and length, respectively, c is the velocity of sound in air, and x is the perpendicular distance from the projectile trajectory to the sensor. Denoting the length of the N-wave as L, then

$$L \propto d \left(\frac{c}{v}\right) (l)^{-1/4}(x)^{1/4}, \tag{2}$$

with A decreasing and L increasing as the shockwave propagates (and dissipates) in the atmosphere. From (1) and (2) we see that the amplitude and length of the N-wave depend linearly on the diameter of the projectile, but are otherwise weakly dependent on its overall shape and velocity. Although somewhat complex in nature when first formed, the shockwave assumes the N-shape after propagating ≈ 50 projectile diameters [2]. The leading and trailing edges have extremely fast rise and fall times (≈ 100 nsec), leading to the observed N-wave characteristic. Thus, the observed shockwave shape is largely independent of the projectile shape and velocity after a short propagation distance. This in turn implies that a general purpose detector can be developed that is applicable to a wide variety of projectiles.

The N-wave can be parameterized in terms of time of arrival τ, amplitude A, and length L. An idealized constant slope N-wave is shown in Figure 1 and described by

$$f(t;\theta) = Af\left(\frac{t-\tau}{L}\right), \quad \tau \le t \le \tau + L, \tag{3}$$

where

$$f(t) = 1 - 2t, \quad 0 \le t \le 1, \tag{4}$$

is the amplitude and length-normalized signal, and $\theta = [\tau, A, L]'$ denotes the parameter vector. Acceptable ranges for θ are assumed to be known from context, based on (1) and (2). Next we consider optimal schemes for detecting $f(t)$ and estimating θ in Gaussian noise.

2. OPTIMAL DETECTION AND ESTIMATION

Consider the binary hypothesis test

$$\begin{aligned} H_1 : r(t) &= f(t;\theta) + n(t), \quad 0 \le t \le T \gg L \\ H_0 : r(t) &= n(t), \end{aligned} \tag{5}$$

where $n(t)$ is white Gaussian noise with variance N_0. We assume that $f(t;\theta)$ is completely contained in the interval T. The Bayes optimal decision rule is based on the likelihood ratio

$$\frac{\int_\theta p_1(\mathbf{r}|\theta) w_\theta(\theta) d\theta}{p_0(\mathbf{r})} \underset{H_0}{\overset{H_1}{\gtrless}} \lambda_0, \tag{6}$$

where $w_\theta(\theta)$ is the a priori joint probability density of θ, and $p_i(\cdot)$ is the likelihood function under the ith hypothesis.

We further assume that the unknown parameters in θ are independent. This last assumption is not strictly true: A and L both depend on the parameters in (1). However, we are assuming the quantities d, v, l, and x are unknown.

Next we consider the form of the optimal detection receiver. Suppose that L is random with τ and A known, and assume a uniform prior on L, so that $L \sim \mathcal{U}[L_0, L_1]$, with $0 < L_0 < L_1$. Now,

$$p_1(\mathbf{r}) = c_0 \int_{L_0}^{L_1} \exp\left\{\frac{-1}{N_0}\int_0^T [r(t) - f(t;L)]^2 dt\right\} \frac{dL}{L_1 - L_0}, \tag{7}$$

while under H_0

$$p_0(\mathbf{r}) = c_0 \exp\left\{\frac{-1}{N_0}\int_0^T [n(t)]^2 dt\right\}, \tag{8}$$

with c_0 a constant. Defining the signal energy

$$E_f = \int_0^T [f(t)]^2 dt = \frac{LA^2}{2}, \quad (9)$$

and also defining

$$q(L) = \int_0^T r(t)f(t)dt, \quad (10)$$

then we can write

$$\lambda(\mathbf{r}) = \int_L \lambda(\mathbf{r}|L)w(L)dL = \int_{L_0}^{L_1} \exp\left\{\frac{-E_f}{N_0} + \frac{2A}{N_0}q(L)\right\} \frac{dL}{L_1 - L_0}. \quad (11)$$

For the purposes of implementation we partition the uniform density for L into a discrete set of equally likely lengths L_i, $i = 1, \ldots, M$, so that we may replace the integration of (11) by the summation

$$\lambda(\mathbf{r}) \approx \frac{1}{M}\sum_{i=1}^{M} \lambda(\mathbf{r}|L_i). \quad (12)$$

A similar argument for the time of arrival τ may be applied, where we take $\tau \sim \mathcal{U}[0, \tau_1]$. For L and τ random and assuming A known, then

$$\lambda(\mathbf{r}) = \int_0^{\tau_1} \lambda(\mathbf{r}|\tau)\frac{d\tau}{\tau_1} = \int_0^{\tau_1}\int_{L_0}^{L_1} \exp\left\{\frac{-E_f}{N_0} + \frac{2A}{N_0}q(L)\right\} \frac{dL}{L_1 - L_0}\frac{d\tau}{\tau_1}, \quad (13)$$

where $\lambda(\mathbf{r}|\tau)$ is given by (11). Partitioning the delays τ into an equally likely set τ_j, $j = 1, \ldots, N$, then

$$\lambda(\mathbf{r}) \approx \frac{1}{MN}\sum_{i=1}^{M}\sum_{j=1}^{N} \lambda(\mathbf{r}|L_i, \tau_j), \quad (14)$$

with

$$\lambda(\mathbf{r}|L_i, \tau_j) = \exp\left\{\frac{-E_f}{N_0} + \frac{2A}{N_0}q(L)\right\}. \quad (15)$$

Note that E_f depends on L and A and is therefore not constant from realization to realization under our assumptions. Thus, in the implementation based on (14) the correction term $-E_f/N_0$ is applied in each branch for normalization.

Finally, consider the effects of A, L, and τ random. Now $\lambda(\mathbf{r}|A)$ is given by (13), and we note that $\lambda(\mathbf{r}|A)$ is maximized for any fixed $A > 0$ if q is maximized. Thus, a decision may be made by comparing the correlation q to a threshold, and q provides a uniformly most powerful test with respect to amplitude A.

Without knowledge of the prior probabilities of H_1 versus H_0 it is prudent to select the decision threshold λ_0 via the Neyman-Pearson criterion so as to maximize the probability of detection for a fixed probability of false alarm.

An alternative to (14) is the "maximum-likelihood" detector, which is an approximation to (14). This detector proceeds by taking the maximum of the M paths, as shown in Figure 2, and corresponds to a bank of matched filters. The maximum-likelihood detector, or generalized likelihood-ratio test (GLRT), corresponds to forming the maximum-likelihood estimates of the parameters and then using these in the likelihood-ratio as if they were the true θ. Thus the detector of Figure 2 is appealing for our problem because it simultaneously yields estimates $\hat{\tau}$ and $\hat{L}$. Given $\hat{\tau}$ and $\hat{L}$ an optimal estimate of A is easily obtained via linear regression over $\hat{\tau} \le t \le \hat{\tau} + \hat{L}$. Because E_f can change, a normalization is required before applying the threshold λ_0. Alternatively, the statistic q can be employed requiring a separate threshold for each channel, and estimation of A can be avoided altogether.

3. WAVELET-BASED DETECTION AND ESTIMATION

The use of wavelets for solving an hypothesis test must be carefully motivated. Under the assumptions on (5), application of an invertible linear transformation such as a wavelet transform (WT) or Fourier transform does not provide any benefit. This is because the optimal theory will lead us to invert the transformation, thus recovering the original hypothesis test in white Gaussian noise. Two possible motivations for using wavelets are (i) lower complexity implementation, and (ii) robust performance when assumptions are violated. In the case of (i) we are trading off performance versus complexity because the use of the WT will lead to a sub-optimal detector. In the case of (ii) our assumptions may be violated in such a way that an optimal solution is unknown, including such possibilities as an unknown additive interference or signal model mismatch.

The continuous wavelet transform (CWT) of $f(t)$ is given by

$$\tilde{f}(s,t) = \int_{-\infty}^{\infty} s^{-1/2}\psi\left(\frac{t-\tau}{s}\right) f(t)dt, \quad (16)$$

where s is the scale and $\psi(t)$ is the basic wavelet. It is interesting to note that the CWT can be reinterpreted as a filter that is matched to the various scalings s and time of arrivals τ of the signal $f(t)$. In radar the scaling is due to Doppler shifts in the return signal, and the CWT is equivalent to the wideband ambiguity function. In the present problem the wavelet $\psi = \psi(t; \tau, s)$ is exactly of the form of $f(t; \tau, L)$ with L playing the role of scale (see (3)). Thus, under our assumptions, finding the maximum of $\tilde{f}(s,t)$ leads to estimates of L and τ corresponding to those obtained in the scheme of Figure 2 (ignoring quantization error). In practice we compute samples of $\tilde{f}(s,t)$, leading to the same quantization issues. It may be possible to reduce the search space over scale, depending on the form of $f(t;\theta)$. If dyadic scales are sufficient then fast discrete wavelet transforms (DWT) may be used, and the issue becomes one of choosing the appropriate (dyadic) scale sampling rate [3]. This last approach is for complexity reduction and is strictly sub-optimal under the assumptions on (5). Computation of samples of $\tilde{f}(s,t)$ for arbitrary scales can also be accomplished efficiently by employing the chirp z-transform [4].

For the specific problem of shockwave detection we consider a scheme based on the ability of the DWT to characterize the local regularity of signals. It is well known that wavelets may be used for detecting and characterizing singularities [5]. This has been applied to edge detection in images by analysis across scale-space [6]. Here it is desired to detect a pair of 1-D singularities, the rising and falling edges of the N-wave. The extremely fast rise and fall times produce singularities even in the presence of other deterministic interferers, so that singularity-based detection will be reasonably robust to the presence of such unwanted contributions to the observed signal. Such interference is very likely in this application, due both to platform noise (e.g., vehicle engine noise) and all manner of other sounds produced in the environment. In addition, for miss distances x

on the order of meters, the relative SNR with respect to the random additive noise will be relatively high such that this sub-optimal approach may be nearly optimal in many scenarios. The use of the DWT in this manner reduces the 2-D matched filter approach into two 1-D estimation problems. First, τ is estimated based on the pair of observed singularities: estimates of L (and A) follow easily. This decoupling of estimation of τ and L results in a significant complexity reduction that depends on the range of L under test.

The particular wavelet of interest here consists of the first derivative of a smoothing function $u(t)$, given by $\psi(t) = du(t)/dt$. It is straightforward to show (e.g., [5])

$$\tilde{f}(s,t) = f(t) * \left(s\frac{du_s}{dt}\right)(t) = s\frac{d}{dt}(f * u_s)(t), \qquad (17)$$

where $u_s(t) = (1/s)u(t/s)$ and $*$ denotes convolution. Thus, for appropriate choice of $u(t)$, $\tilde{f}(s,t)$ can be interpreted as a derivative of a local average of $f(t)$ where the degree of smoothing depends on s. The result is estimation of the derivative of $f(t)$ at various levels of smoothing (scales). In [6] Mallat and Zhong developed a DWT based on $u(t)$ being a cubic spline, and show that finding the local maxima of the DWT modulus is equivalent to the Canny edge detector. In using the DWT here the discretization is dyadic in scale but not in time (shift), which simply corresponds to a filter bank with no downsampling.

Let s_n, $n = 1, 2, \ldots$, denote the dyadic scales of the DWT. Figure 3 (top) shows a measured shockwave time series $f(t)$, sampled at 250 KHz. Also shown in Figure 3 are $\tilde{f}(s_n,t)$ for the first three scales ($n = 1,2,3$). Note that this DWT is dyadic in scale but not in shift. The time series $f(t)$ exhibits multipath but is generally at high SNR. $\tilde{f}(s_1,t)$ yields accurate estimates of the time of arrival and time of departure of the shockwave, but also has multiple peaks that disguise the presence of the true ones. In [6] an approach for edge detection based on estimation of the Lipschitz exponent is developed, but this method is highly susceptible to additive noise. A robust approach is to analyze across scales, beginning with the highest scale (in practice seven or eight scales appear sufficient, depending on the original sampling rate). At the higher scales the smoothing reduces noise and produces fewer false peaks. With scale s_n fixed the presence of $f(t)$ will produce two positive peaks in the time series $\tilde{f}(s_n,t)$, whose spacing depends on L and $u_s(t)$.

Let $\tilde{f}(s_n,t_i)$ and $\tilde{f}(s_n,t_j)$, $t_j > t_i$, denote two values in the DWT at scale s_n. Acceptable peak pairs are selected based on a peak threshold p_0 and the spacing of the peaks. The most basic test conditions are (i) $\tilde{f}(s_n,t_i) > p_0$, $\tilde{f}(s_n,t_j) > p_0$, and (ii) $t_j - t_i \in [L_0, L_1]$. When a candidate peak pair at times (t_i,t_j) is detected in the higher scales, corresponding peaks are found in the lower scales. Finally, estimates of τ and $\tau + L$ are chosen in $\tilde{f}(s_1,t)$, by selecting maxima in a neighborhood of $\tilde{f}(s_1,t_i)$ and $\tilde{f}(s_1,t_j)$, where the neighborhood size depends on $\hat{L}$.

Numerous variations of this approach are possible via selection of the thresholds and choice of scales. For example, a detection can be declared if the majority of scales exhibit the appropriate peak pair, avoiding scales which may be contaminated by interference. A simple approach is to use two scales only, s_1 and some higher scale s_n. We illustrate this below.

4. SIMULATION AND EXPERIMENT

We compare maximum-likelihood (ML) estimation of τ and L in additive white Gaussian noise with the DWT approach. τ and L were varied randomly over 100 realizations for each SNR value with $\tau \in [100, 300]$ and $L \in [50, 200]$. Each realization was of length $T = 500$. ML estimates were obtained via the scheme of Figure 2.

The DWT approach was based on $\tilde{f}(s_1,t)$ and $\tilde{f}(s_5,t)$, as outlined above. For details of the DWT implementation see the appendices of [6]. The derivative threshold p_0 was selected data-adaptively as 60% of the maximum of $\tilde{f}(s_5,t)$, $0 \leq t \leq T$. The first acceptable peak pair (in time) were used. Estimates of τ and L were taken as the maxima in $\tilde{f}(s_1,t)$ closest to the peak locations in $\tilde{f}(s_5,t)$.

Results are shown in Figure 4, depicting experimental standard deviation of the estimation error versus SNR, with SNR defined as E_f/N_0 in dB. Note that perfect estimation (zero error) is possible due to the quantization of the true τ and L. Also, while the ML approach always yields an estimate, the DWT method may fail to yield a solution if the threshold tests fail.

The top of Figure 4 depicts results in white Gaussian noise. The bottom shows results with an additional interference made up of four odd harmonics with the fundamental at $f = 0.01$ (normalized sampling). The interference power was 20 dB with respect to the Gaussian noise, and roughly simulates engine noise. In both cases the simple DWT detector approaches optimal performance at about 32 dB SNR (E_f/N_0). Experimentally observed SNR is often significantly beyond this level (see Figure 3), and tests of the DWT algorithm using this and other experimental data has shown good results.

Of further interest are DWT algorithms using more than two scales, comparison with optimal Cramer-Rao estimation bounds, and further experimental results. Also of interest: study of multi-sensor scenarios for determining projectile direction of arrival, as well as conditions under which projectile classification is feasible. Results on these issues will presented elsewhere.

REFERENCES

[1] G. B. Whitham, "Flow patterns of a supersonic projectile," *Comm. on Pure and Applied Math*, Vol. 5, No. 4, 1952.

[2] G. B. Whitham, "The behaviour of supersonic flow past a body of revolution, far from the axis," *Proc. Royal Soc. of London, Series A*, Vol. 201, pp. 89–109, 1950.

[3] M. Frisch and H. Messer, "Detection of a known transient signal of unknown scaling and arrival time," *IEEE Trans. Sig. Proc.*, Vol. 42, No. 7, pp. 1859–1863, July 1994.

[4] D. Jones and R. Baraniuk, "Efficient approximation of continuous wavelet transforms," *Electronics Letters*, Vol. 27, No. 9, pp. 748–750, April 1991.

[5] S. Mallat and W. L. Hwang, "Singularity detection and processing with wavelets," *IEEE Trans. Info. Theory*, Vol. 38, pp. 617–643, 1992.

[6] S. Mallat and S. Zhong, "Characterization of signals from multiscale edges," *IEEE Trans. Pattern Anal. and Mach. Intell.*, Vol. 14, No. 7, pp. 710–732, July 1992.

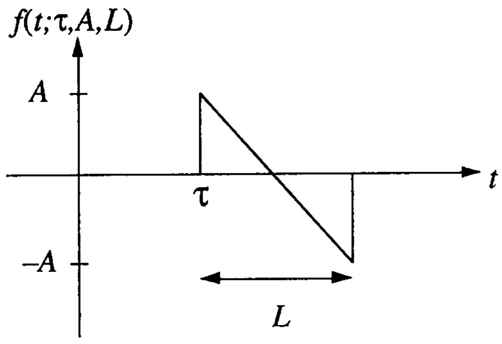

Figure 1. Ideal shockwave (N-wave) $f(t)$.

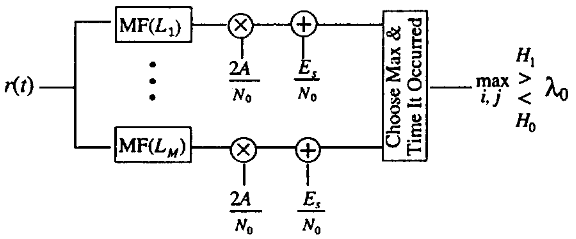

Figure 2. Maximum-likelihood detector in white Gaussian noise.

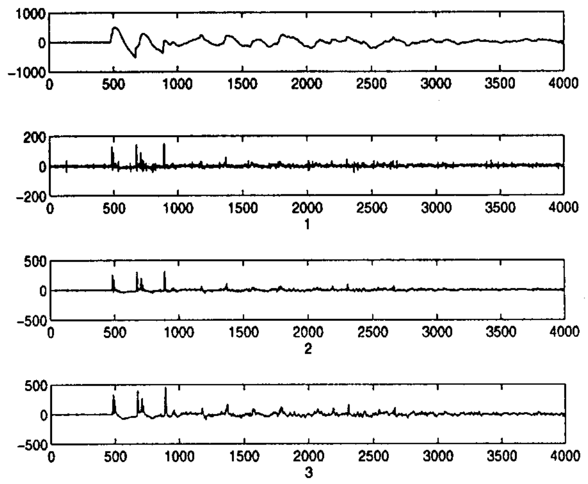

Figure 3. Top: measured shockwave time series $f(t)$, 1–3: DWT $\tilde{f}(s_n, t)$, $n = 1, 2, 3$.

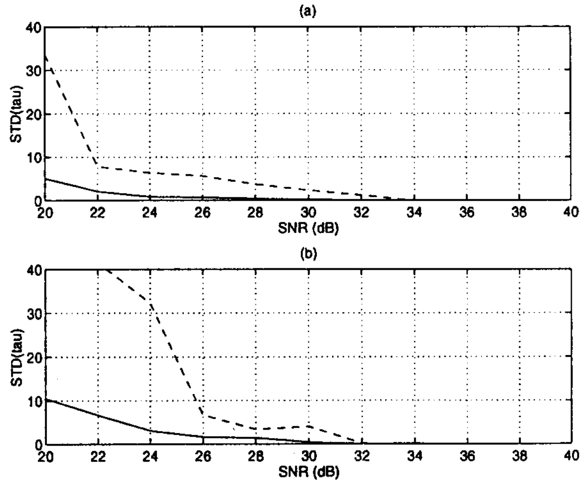

Figure 4. ML (solid) and DWT-based (dashed) estimation of τ. (a) white Gaussian noise, (b) plus simulated engine interference.