



INSTITUTE FOR DEFENSE ANALYSES

HHT Sifting and Adaptive Filtering

Reginald N. Meeson

August 2003

Approved for public release;
distribution unlimited.

IDA Paper P-3766

Log: H 03-000428

This work was conducted under IDA's independent research program. The publication of this IDA document does not indicate endorsement by the Department of Defense, nor should the contents be construed as reflecting the official position of that Agency.

© 2003, 2004 Institute for Defense Analyses, 4850 Mark Center Drive, Alexandria, Virginia 22311-1882 • (703) 845-2000.

This material may be reproduced by or for the U.S. Government.

INSTITUTE FOR DEFENSE ANALYSES

HHT Sifting and Adaptive Filtering

Reginald N. Meeson

Executive Summary

Time-frequency analysis is the process of determining what frequencies are present in a signal, how strong they are, and how they change over time. Most of the information carried by analog signals is contained in their time-varying, dynamic, and transient frequency spectra. Understanding how the frequencies in a signal change with time can also explain much about the physical processes that generate or influence the signal. Better resolution of details of frequency changes provides better visibility into these underlying physical processes.

The Hilbert/Huang Transform (HHT) is a time-frequency analysis technique that offers higher frequency resolution and more accurate timing of transient and non-stationary signal events than conventional integral transform techniques. The HHT separates complex signals into simpler component signals, each of which has a single, well-defined, time-varying frequency. Real-time HHT algorithms enable this enhanced signal analysis capability to be used in process monitoring and control applications.

“Sifting” is the central signal separation process of the HHT algorithm. This paper compares the component signal separations of Huang’s sifting process with those produced by adaptive filtering techniques. Initially, we conjectured that adaptive filtering, with appropriate real-time adjustments to parameters, could substitute for Huang’s sifting process, but this was found not to be the case. Five case studies present HHT and adaptive filtering results for stationary amplitude- and frequency-modulated signals, as well as signals with more dynamic transient behavior. These examples show that, in general, HHT sifting and adaptive filtering separate signal components quite differently.

Our experiments with example signals led to the discovery of aliasing in the HHT sifting algorithm. Aliasing is a condition in sampled-data signals where high-frequency content is misinterpreted as lower-frequency content. In movies, for example, the illusion of spoked wheels that appear to spin backwards is caused by aliasing. Aliasing is usually considered undesirable and a form of signal corruption. We are continuing to investigate how adaptive filtering might be combined with the HHT sifting process to avoid aliasing and improve the signal separations that result.

Preface

This paper was prepared by the Institute for Defense Analyses (IDA) under an internal Central Research Project. This work was aimed at exploring evolving signal analysis technologies with potential applications in Command, Control, Communications, Computers, Intelligence, Surveillance, and Reconnaissance (C4ISR) and information systems.

IDA research staff members Dr. Alfred E. Brenner, Dr. Kevin E. Foltz, Dr. J. Michael Hanratty, and Dr. L. Roger Mason, and summer intern Mr. Homer E. Ong reviewed this paper.

Contents

Executive Summary	ES-1
1. Introduction	1
2. Objectives of HHT Sifting	3
2.1 Restrictions on Amplitude and Phase Functions	4
3. Huang's Algorithm	9
4. Incremental, Real-Time HHT Sifting	11
4.1 Testing for Iteration Convergence	12
4.2 Time-Warp Analysis	13
4.3 Calculating Warp Filter Characteristics	14
4.4 Separating Amplitude and Phase	16
5. Filtering in Standard Time	17
6. Case Studies	19
6.1 Simple Reference Example	19
6.2 Amplitude Modulated Example	20
6.3 Frequency Modulated Example	24
6.4 Amplitude Step Example	27
6.5 Frequency Shift Example	32
7. Summary and Conclusions	37
7.1 Summary of Case Study Findings	37
7.2 Research Directions	38
References	Refs-1

Figures

Figure 1. Example Multicomponent Signal	6
Figure 2. Block Diagram of the HHT Signal Separation Process	9
Figure 3. Block Diagram of One Iteration of $\square$	12
Figure 4. Example Signal Containing Hidden Peaks	12
Figure 5. Frequency Response of HHT Trend-Estimating Process.....	14
Figure 6. Warped Filter Transfer Functions	15
Figure 7. Adaptive Filter Transfer Functions	17
Figure 8. Simple Two-Component Example Signal	20
Figure 9. Example Amplitude-Modulated Signal.....	20
Figure 10. High-Frequency Component Separated from the AM Signal by Adaptive Filtering	21
Figure 11. Instantaneous Frequency of the AM Signal Component Separated by Adaptive Filtering	22
Figure 12. High-Frequency Component Separated from the AM Signal by the HHT Sifting Process	23
Figure 13. Instantaneous Frequency of the AM Signal Component Separated by HHT Sifting.....	23
Figure 14. Example Frequency-Modulated Signal.....	24
Figure 15. Instantaneous Frequency of the FM Signal Component Separated by HHT Sifting	25
Figure 16. High-Frequency Component Separated from the FM Signal by Adaptive Filtering	26

Figure 17. Instantaneous Frequency of the FM Signal Component Separated by Adaptive Filtering	27
Figure 18. Amplitude Step Example Signal	27
Figure 19. HHT Component and Trend Results for the Amplitude Step Signal	28
Figure 20. Adaptive Filter High- and Low-Pass Results for the Amplitude Step Signal	28
Figure 21. Instantaneous Frequency of the Amplitude Step Component Separated by HHT Sifting	29
Figure 22. Instantaneous Frequency of the Amplitude Step Component Separated by Adaptive Filtering	29
Figure 23. Fourier Transform (Magnitude) of the Amplitude Step Signal	30
Figure 24. Adaptive Filter High- and Low-Pass Spectra for the Amplitude Step Signal	31
Figure 25. Spectra of the HHT Trend and Separated Component for the Amplitude Step Signal	32
Figure 26. Frequency Shift Example Signal	32
Figure 27. Instantaneous Frequency of the Frequency Shift Component Separated by HHT Sifting	33
Figure 28. Adaptive Filter High- and Low-Pass Results for the Frequency Shift Signal	34
Figure 29. Instantaneous Frequency of the Frequency Shift Component Separated by Adaptive Filtering	34
Figure 30. Fourier Transform (Magnitude) of the Frequency Shift Signal	35
Figure 31. Adaptive Filter High- and Low-Pass Spectra for the Amplitude Shift Signal	36

1. Introduction

One way to describe a timed series of measurements, referred to as a signal, is in terms of the frequencies in its variations. The process of determining what frequencies are present, how strong they are, and how they change over time is called time-frequency analysis. Conventional time-frequency analysis techniques use integral calculus transforms to map time-based signals into frequency-based or joint time- and frequency-based representations [1]. Examples of these techniques include Fourier transforms, windowed Fourier or Gabor transforms, wavelet transforms, and joint time-frequency distributions.

The Hilbert/Huang Transform (HHT) is a new time-frequency analysis technique that offers higher frequency resolution and more accurate timing of transient and non-stationary signal events than conventional Fourier and wavelet transform techniques [2]. Conventional techniques assume signals are stationary, at least within the time window of observation. Fourier analysis assumes further that the signal is harmonic and repeats itself with a period that exactly matches the width of the sampling window. These analysis techniques are employed widely even though their (theoretically necessary) enabling conditions rarely hold for signals of interest.

In addition, integral transform techniques suffer from an uncertainty problem similar, mathematically, to Heisenberg's uncertainty principle in physics. This uncertainty limits their ability to accurately measure timing and frequency at the same time. That is, after a point, higher-resolution frequency measurements cannot be achieved without sacrificing timing accuracy, and vice versa. The HHT is able to resolve frequencies accurately and time them precisely without this limiting uncertainty.

The original HHT algorithm was formulated as a "batch" computation where a complete data set is collected and then processed as a whole. An incremental algorithm that transforms evolving input data streams into streams of HHT results has also been developed [3]. Modern microprocessors and signal processing chips offer sufficient performance for this incremental algorithm to be used in many real-time applications. The terms "incremental" and "real-time" are, therefore, used interchangeably to describe this algorithm.

"Sifting" is the central signal separation process of the HHT algorithm. In the seminal work on the HHT [2], Huang described sifting informally as analogous to an adaptive

filtering process, but then developed a different algorithmic procedure to separate signal components. This led us to conjecture that adaptive filters, with parameters appropriately adjusted in real time, could mimic the HHT sifting process. Results from adaptive filtering seemed natural to analyze and compare along side the HHT results. Huang's original HHT sifting algorithm was the starting point for this comparison. Results from the incremental HHT algorithm and adaptive filtering were used in this analysis.

Adaptive filtering, for this discussion, means conventional finite impulse response (FIR) digital filtering where filter coefficients can be changed on a sample-by-sample basis. Our experiments with these signal analysis techniques revealed new insights into the mathematical properties of the HHT signal separation process that may help refine HHT processing techniques.

In Section 2 we describe the objectives of the HHT signal separation process and the desired attributes of separated components. Huang's original empirical mode decomposition algorithm, which later became known as the HHT, is described in Section 3. Section 4 describes the incremental HHT algorithm and analyzes a special case where an analogy to conventional digital filtering techniques can be used. In Section 5 we describe the shift from special-case static filtering to a general method using adaptive filtering. HHT and adaptive filtering results for five example signals are compared in Section 6. Section 7 concludes with a summary and some directions for future research.

2. Objectives of HHT Sifting

The HHT sifting process separates a signal into a series of amplitude- and frequency-modulated component signals in the form:

- $s(t) = \sum a_i(t) \cos(\phi_i(t))$

where the $a_i(t)$ terms represent the amplitude modulation characteristics and the $\phi_i(t)$ terms are the phase functions that represent the frequency modulation characteristics of each component.

There are numerous possible solutions to this separation scheme. One familiar solution is the Fourier series [4], which is made up of constant amplitude and constant frequency (linear phase) functions. The solution the HHT seeks is quite different. Rather than trying to represent a signal by predetermined basis functions, the HHT tracks and adapts dynamically to transient, non-stationary, and nonlinear changes in component frequencies and amplitudes as the signal evolves over time.

Windowed Fourier and wavelet signal analysis techniques are also able to track slowly changing signal behavior but, as described above, they suffer from an uncertainty problem that can limit the accuracy of the frequency (scale for wavelets) and timing information they yield [1]. The product of the frequency (scale) variance and the timing variance for results from these techniques has a positive lower bound. This means that once this limit is reached, increasing the accuracy of frequency measurements can only be achieved by sacrificing timing accuracy, and vice versa.

Earthquake data, for example, contain short-duration transients that are difficult to analyze because of this uncertainty limitation. Using conventional analysis techniques, it is not possible to accurately time when specific frequencies were present. Transient events can be timed accurately but accurate frequency information cannot be resolved within that narrow time window.

HHT signal separations are not subject to this limitation and provide both accurate frequency and accurate timing simultaneously. This is a unique advantage of the HHT over conventional time-frequency analysis techniques. HHT analysis of earthquake data [5], for example, shows a very different distribution of frequencies over time than

conventional Fourier analysis, which may prove tremendously important in analyzing the strength of buildings, bridges, and other structures.

2.1 Restrictions on Amplitude and Phase Functions

In order to extract the desired amplitude and frequency information, without conflicting interpretations or paradoxical results, restrictions must be imposed on the amplitude and phase functions, $a_i(t)$ and $\phi_i(t)$. The primary requirement for HHT components is that they be sufficiently well behaved to allow extraction of well-defined amplitude and phase functions. Such functions are called “monocomponent” functions and we distinguish them from “multicomponent” functions, from which amplitude and phase cannot be cleanly extracted. Although there seems to be no generally accepted mathematical definition of “monocomponentness,” there is little debate over one primary criteria, which is that at any time a monocomponent signal must have a single, well-defined, positive instantaneous frequency represented by the derivative of its phase function.

The first approach suggested for finding necessary conditions for a separated component’s “monocomponentness” was to look at the component’s analytic signal, which is given by:

- $\mathcal{A}[c(t)] = c(t) + j\mathcal{H}[c(t)]$

where $c(t) = a(t) \cos(\phi(t))$ and $\mathcal{H}$ is the Hilbert transform. (See [1] for a thorough discussion of analytic signals and the Hilbert transform.) The analytic signal is a complex function whose Fourier transform is twice that of $c(t)$ over the positive frequencies and zero over negative frequencies. The spectrum of this signal, therefore, contains only positive frequencies. This does not guarantee, however, that the signal’s instantaneous frequency (the derivative of its phase) will always be positive. Cohen [1] shows examples of analytic signals that have paradoxical instantaneous frequency characteristics, including some with negative instantaneous frequencies. The analytic signal, therefore, by itself, does not appear to provide sufficient criteria for separating monocomponent signals.

A second approach suggested for finding monocomponent conditions was to consider the function’s quadrature model, which is:

- $\mathcal{Q}[c(t)] = a(t) e^{j\phi(t)}$

Another formulation of the quadrature signal is:

- $\mathcal{Q}[c(t)] = a(t) [\cos(\phi(t)) + j \sin(\phi(t))]$

Using the additional knowledge about the Hilbert transform that

- $\mathcal{H}[\cos(\varphi(t))] = \sin(\varphi(t))$

the quadrature model can be compared with the analytic signal. The two are the same when the amplitude function can be factored out of the signal's Hilbert transform; that is, when

- $\mathcal{H}[a(t) \cos(\varphi(t))] = a(t) \mathcal{H}[\cos(\varphi(t))] = a(t) \sin(\varphi(t))$

The conditions under which this relationship holds were established by Bedrosian [6] and elaborated by Nuttall [7]. The conditions are, for some positive frequency φ_0 :

- a. The spectrum of the amplitude function is restricted to frequencies below φ_0 , and
- b. The spectrum of the cosine term is restricted to frequencies above φ_0 .

An example function that does not satisfy these conditions is:

- $s(t) = 1.25 \cos(t) - \cos(2t)$

The analytic signal of this function is similar to one of Cohen's problematic signals,

- $\mathcal{A}[s(t)] = 1.25 e^{jt} - e^{j2t}$

which cannot be expressed in the form $a(t) e^{j\varphi(t)}$ without either $a(t)$ oscillating rapidly or $\varphi(t)$ turning negative periodically. As can be seen in the graph shown in Figure 1, the real signal $s(t)$ has local minima with positive values. Such a signal cannot be expressed in the form $a(t) \cos(\varphi(t))$ with a slowly varying amplitude and an increasing phase function. If we assume a slowly varying amplitude, to satisfy Bedrosian's first condition, then the $\cos(\varphi(t))$ term would have to turn and go back up again without going negative. The phase function, therefore, would have to decrease for a time, resulting in a negative instantaneous frequency. This violates Bedrosian's spectral separation conditions, since the amplitude function would have to have a negative upper frequency bound. If we stipulate an increasing phase function, then the amplitude must peak near $t=(2n+1)\varphi$ and dip to a minimum near $t=2n\varphi$, giving it an average frequency of $\varphi=L$, the same as the average change in phase. Either way, Bedrosian is not satisfied.




Figure 1. Example Multicomponent Signal

Bedrosian's conditions are a bit too restrictive for our needs, however. Purely frequency-modulated signals with constant amplitude can have spectra that extend down to zero frequency. Any amplitude modulation imposed on such a “carrier” signal would violate Bedrosian's conditions – even though the signal would make a perfectly good HHT component. The case studies below show that it is important for solutions to allow phase functions that exhibit this sort of frequency-modulated behavior.

Teager's energy operator, $\square$, was suggested as a possible non-linear approach for restricting amplitude and phase functions for combined amplitude-modulated (AM) and frequency-modulated (FM) signals [8].

- $\square(s(t), t) = [s'(t)]^2 - s(t) s''(t)$

For component signals of the form $a(t) \cos(\phi(t))$, $\square$ can be expanded as:

- $\square[a(t) \cos(\phi(t)), t] = [a(t) \phi'(t)]^2 + 0.5 a^2(t) \sin(2\phi(t)) \phi''(t) + \cos^2(\phi(t)) \square[a(t), t]$

If a signal has a dominant high-frequency component, the first term in this formula will dominate the others. Maragos [8] describes the secondary terms as “error” terms and shows how they can be minimized by constraining the AM and FM indexes of modulation, and the modulating signal bandwidth.

The integrals of the two terms in Teager's $\square$ operator are both equal to the signal's total energy times its average square frequency. That is,

- $\int [s'(t)]^2 dt = - \int s(t) s''(t) dt = \int \omega^2 |S(\omega)|^2 d\omega = E \int \omega^2 d\omega$

where $S(\omega)$ is the signal's Fourier transform, E is its total energy,

- $E = \int [s(t)]^2 dt = \int |S(\omega)|^2 d\omega$

and $\int \omega^2 d\omega$ is the average square frequency.

Instantaneously, though, Teager's two terms are quite different. ω may not even yield positive results. For the signal in Figure 1, for example, values of ω are negative in the vicinity of $t=2n\pi$ (where $s'(t) \approx 0$, $s(t) > 0$, and $s''(t) > 0$).

For lightly modulated signals, ω produces a stable output dominated by $[a(t)\omega(t)]^2$. As long as the "error" terms are sufficiently small, Maragos [9] showed that ω can be used to demodulate the signal and extract approximate values for $a(t)$ and $\omega(t)$ by applying ω to the signal and its derivative:

- $\omega[s(t), t] = \omega[a(t) \cos(\omega(t)), t] \approx [a(t) \omega(t)]^2$
- $\omega[s'(t), t] \approx a^2(t) [\omega(t)]^4$

Teager's formula appears to offer possibilities for identifying signals that would satisfy our general notion of monocomponentness. Turning these results into algorithms for separating monocomponent signals from more complex ones, however, is still an open problem.

We proceed from this point without a concrete definition of monocomponentness, but recognizing that it implies constraints on phase monotonicity ($\omega'(t) > 0$), amplitude and "carrier" signal bandwidth, and degrees of amplitude and frequency modulation.

3. Huang's Algorithm

Huang's sifting process separates the highest-frequency component embedded in a multi-component signal from all the lower-frequency components. This separated component is well behaved, although the mathematical monocomponentness criteria it satisfies is not easily determined. The remaining lower-frequency components together make up the signal trend. A signal, described in terms of its first component and residual trend functions, is:

- $s(t) = a_1(t) \cos(\varphi_1(t)) + r_1(t)$

The sifting process for a single component is repeated using the trend output from one stage as the input to the next, producing the series of $a_i(t) \cos(\varphi_i(t))$ terms that sum to reconstruct the original signal, $s(t)$. A block diagram of this process is shown in Figure 2.




Figure 2. Block Diagram of the HHT Signal Separation Process

To determine $r(t)$, Huang fit smooth envelope curves (using cubic splines) to the local maxima of the signal and to the local minima. The average of these two envelopes provides a rough estimate of $r(t)$. (Local maxima are referred to as positive peaks even though the signal values at those points may be positive or negative. Local minima are similarly referred to as negative peaks.) Huang then applied an iteration scheme to refine the estimated trend. The iteration scheme can be formulated as:

- $r_{(n+1)}(t) = r_{(n)}(t) + \square(c_{(n)}, t)$

where $c_{(n)}(t) = s(t) - r_{(n)}(t)$

The function $\square$ represents the spline curve fitting and averaging process applied to the peaks of function $c_{(n)}(t)$. (Subscripts in parentheses indicate the iteration count.) This calculation is repeated (starting with $r_{(0)}(t)=0$) until a fixed point is reached and $\square(c_{(n)}, t)$ converges to zero (within some small $\square$). Once the residual or trend function is determined, the difference between it and the input signal is the highest-frequency separated component, $c_i(t) = a_i(t) \cos(\square(t))$.

Huang called this separation technique “empirical mode decomposition,” and the individual component signals “intrinsic mode functions” [2]. His colleagues later named the method the Hilbert/Huang Transform.

To separate the $a_i(t)$ and $\square_i(t)$ functions, Huang computed the component’s analytic signal using Fourier transforms. The Fourier transform of a function’s Hilbert transform satisfies the relation:

- $\mathcal{F}[\mathcal{H}[s(t)]] = -j \text{sign}(\square) \mathcal{F}[s(t)]$

where $\mathcal{F}$ is the Fourier transform and $\mathcal{H}$ is the Hilbert transform. The Fourier transform of a function’s analytic signal can then be formulated as:

- $\mathcal{F}[\mathcal{A}[s(t)]] = \mathcal{F}[s(t)] + \text{sign}(\square) \mathcal{F}[s(t)]$

which is zero for all negative frequencies and double the input signal’s values for all positive frequencies.

Taking a separated component’s Fourier transform, zeroing its negative-frequency terms and doubling its positive-frequency terms, and then applying the inverse Fourier transform, produces the component’s complex analytic signal. The magnitude of the analytic signal (theoretically) is $a(t)$ and the argument is $\square(t)$.

4. Incremental, Real-Time HHT Sifting

In Huang’s original HHT algorithm, the data passed between processing blocks in Figure 2 are arrays containing entire time series. The incremental algorithm [3] turns these batch-processing blocks into pipeline processes that operate incrementally on streams of data, passing one data sample at a time.

The first step in sifting is to identify signal peaks. Calculating peak values and times in the incremental HHT algorithm is the same as in Huang’s original algorithm, except that peak value and time pairs, $[y_p, t_p]$ are produced incrementally as the input stream evolves. The resulting stream of peak values corresponds to sampling the input signal at its peak times rather than at uniform intervals.

Spline interpolation uses global information to calculate the derivative of the positive envelope at each positive peak, and similarly for the negative envelope at each negative peak. For incremental processing only local information is available, so we must rely on Hermite interpolation [10], which is very similar to spline interpolation but uses derivative values estimated from local signal behavior.

Using the spline parameters derived for each segment of the positive peak envelope, values are calculated at points corresponding to the signal’s original sample times. This resampling process produces a stream of uniformly sampled envelope values, although with some latency from the peak detection and spline interpolation process. The same process is applied to the negative-peak data. The two resampled envelope streams are then averaged to produce a stream of trend values. This process, diagrammed in Figure 3, represents one application of Huang’s $\square$ function. Each stage of this process is performed incrementally, so the calculation of $\square$ is achieved incrementally.




Figure 3. Block Diagram of One Iteration of $\hat{c}$

4.1 Testing for Iteration Convergence

Huang's test for iteration convergence is a global test that spans the entire signal duration, which is not consistent with our incremental processing objectives. One reason for the global test is that removing the residual or trend component occasionally exposes new peaks that appeared only as inflections in the original signal. An example where this occurs is illustrated by the signal:

- $s(t) = \cos(t) - 0.167 \cos(5t)$

The graph of this function is shown in Figure 4. The trend function produced by $\hat{c}$, also shown in the figure (dashed line), cuts through the inflection points in the signal as it crosses the axis, which produces new peaks in the next iteration $c(t)$ that were not present in the input for this iteration. These new peaks are included in all further iterations.




Figure 4. Example Signal Containing Hidden Peaks

Discovery of new peaks introduces highly nonlinear disturbances in the trend that may require several additional iterations to smooth out. This can occur even when the trend has nearly converged to its fixed point. Any strictly local test for convergence of the iteration process, therefore, is likely to give occasional false indications. We have not yet found a satisfactory incremental test for convergence. We therefore use fixed-length chains of $\square$ operations, and make them long enough so that errors from terminating the iteration too early are rare. Unnecessary iterations could be short-circuited if we could devise a reliable incremental test for convergence.

4.2 Time-Warp Analysis

If the peaks of an input signal are uniformly spaced, a number of simplifying assumptions can be made in the sifting process. These assumptions do not apply in general, so this approach cannot be used to process arbitrary signals, but the analysis provides insights that can be generalized.

Disregard, for the moment, the timing information that accompanies the incremental stream of peak values described above, and assume these peak values had been sampled at some uniform rate. The distortion this introduces is referred to as a “time warp,” since the actual peak times in general are not uniformly spaced. Although all of the nonlinear phase information between peaks in the original signal is lost (for the moment), the trend of the warped signal can be easily calculated using standard low-pass digital filtering techniques.

In the warped world, one iteration of Huang’s fixed-point function, $\square$, for a series of warped peak values v at time t_p , corresponds to the following expression:

$$\bullet \quad \square(v, t_p) = 1/2 v_p - 1/32 v_{p-3} + 9/32 v_{p-1} + 9/32 v_{p+1} - 1/32 v_{p+3}$$

This is the average of the two envelopes, one of which is represented by v_p and the other is interpolated from the spline curve derived from the neighboring opposite-sign peaks (v_{p-3} , v_{p-1} , v_{p+1} , and v_{p+3}) at time t_p . This expression corresponds to a simple low-pass digital filter, which has the frequency response shown in Figure 5. As can be seen in this graph, the transition band for one pass through this filter crosses at approximately one-half of the warped signal’s Nyquist frequency.

If the timing of peaks does not change from iteration to iteration, multiple iterations correspond to passing the signal through this filter multiple times. (The timing of peaks may change slightly, usually in the initial iterations.) Multiple passes through a simple filter are equivalent to a single pass through a larger filter [11]. Huang’s iteration scheme is formulated so that it is the high-pass filter that is iterated, which successively reduces

the resulting pass band. The corresponding longer low-pass filters have wider pass-band regions and sharper transitions to the stop band. Examples of transfer functions for filters representing different iterations of $\square$ are shown in Figure 5.




Figure 5. Frequency Response of HHT Trend-Estimating Process

The original HHT algorithm uses the shrinking corrections of the iteration process to judge when it has converged. This corresponds to choosing the characteristics of filters dynamically, based on the signal's behavior. As can be seen in Figure 5, each iteration of $\square$ shifts the filter transfer function to a higher-frequency cutoff point. Note also that successive iterations have less and less effect on the size of the frequency shift. Rather than iterate the simple filter corresponding to $\square$, we wish to determine the filter characteristics necessary to directly satisfy the monocomponent criteria and separate the component from the trend in a single pass.

4.3 Calculating Warp Filter Characteristics

Consider that the separated warped signal can be described by:

- $s_p = a_p + r_p$ for all positive peaks, and
- $s_p = -a_p + r_p$ for all negative peaks

where a_p is the absolute value of the high-pass filter output and r_p is the low-pass filter output for each peak. The a_p values are interpreted as approximating a warped sampling of the amplitude function, $a(t)$. The r_p values are similarly interpreted as a warped sampling of the residual function, $r(t)$.

The spectrum of the warped residual function is controlled by the low-pass filtering effects of multiple iterations of $\square$. This same filtering process also controls the spectrum of the warped amplitude function. The spectrum of the series of a_p values is shifted upward by the modulating effects of the warped “carrier” signal, $\cos(\square p)$. The spectrum captured by the high-pass filter, therefore, is that of the amplitude function shifted upward by $\square$. If $R(\square)$ is the low-pass filter transfer function for the r_p values, then the corresponding transfer function for the a_p values is:

- $A(\square) = 1 - R(\square - \square)$

This relationship, for an idealized separation filter, is shown in Figure 6. (The transfer function for the high-pass filter is shown as $C(\square)$.) From these graphs we can see that, to satisfy Bedrosian and keep the $\cos(\square p)$ and a_p spectra from overlapping, the stop band breakpoint for the high-pass filter must be no lower than half the warped Nyquist frequency.




Figure 6. Warped Filter Transfer Functions

Ten iterations of this filter would reduce the effective filter throughput at one-half the warped Nyquist frequency to approximately 2^{-10} , which should satisfy Bedrosian’s separation criteria for many practical purposes. Iterating the simple warped filter or substituting a more efficient filter, however, will not discover any new peaks. In practice, we have often encountered signals that require 25 to 30 iterations of Huang’s $\square$ operator to converge. Much of this disparity in iteration counts is attributable to the nonlinear disturbances caused by the discovery of new peaks.

4.4 Separating Amplitude and Phase

To separate amplitude and phase functions incrementally we substituted a Hilbert transform filter for the batch Fourier transform process described earlier for calculating analytic signals. A Hilbert transform filter has a transfer function that approximates the Fourier transform of a signal's Hilbert transform: $H(\omega) = -j \text{sign}(\omega)$ (see, for example, [11]). For a monocomponent signal, $a(t) \cos(\phi(t))$, this filter approximates:

- $$h(t) * [a(t) \cos(\phi(t))] = a(t) \sin(\phi(t))$$

where $h(t)$ represents the Hilbert transform filter coefficients and “*” represents convolution. The amplitude and phase functions are easily separated using this result:

- $$a(t) = \sqrt{[a(t) \sin(\phi(t))]^2 + [a(t) \cos(\phi(t))]^2}$$
- $$\phi(t) = \text{atan2}(a(t) \sin(\phi(t)), a(t) \cos(\phi(t)))$$

Once the phase function is extracted, the signal's instantaneous frequency is calculated by passing $\phi(t)$ through a differentiating filter (after compensating for the discontinuities in the atan2 results). All of these calculations are done incrementally.

The band-limiting effects of warp filtering on the amplitude envelope indicate that $a(t)$ should be relatively smooth. That is, we expected $a(t)$ to look like the smooth spline-connected envelopes calculated in the final iteration of ϕ in the sifting process, with all of the high-frequency content captured by the phase function, $\phi(t)$. Both the Hilbert transform filter and the Fourier batch technique, however, were found to introduce a high-frequency “ripple” in the amplitude results for some signals.

The explanation for this seeming anomaly is that, within certain limits, the spectral energy of a combined amplitude- and frequency-modulated signal can be freely exchanged between the amplitude and phase functions. While we expected a band-limited amplitude, the Hilbert transform appears to split the difference, sharing the high-frequency content between the amplitude and phase functions. The result, therefore, is sometimes a bit different from what we expected, but is an equivalent representation of the signal.

We experimented with a number of different possible techniques for separating amplitude and phase, including Teager's energy operator. None of these other techniques were as successful as the Hilbert transform filter. Teager's operator worked fine for the signal itself, but occasionally produced negative results for the derivative of the signal, spoiling Maragos's demodulation approach [9]. Boashash [12] provides an extensive discussion of additional techniques for extracting a signal's instantaneous frequency.

5. Filtering in Standard Time

The next objective, to test our conjecture about adaptive filtering substituting for HHT sifting, was to reproduce the effects of Huang's $\square$ operation in standard time, without resampling the original input signal. In the process, we wanted to avoid the unreasonable time-warp analysis assumptions about uniformly spaced peaks. The question posed was: Is there a corresponding standard-time filter that will isolate a comparable (unwarped) trend function and, if so, what are its characteristics? Any filter that approximates this response will have to change its attributes over time (possibly every few samples) to track transient and non-stationary changes in the signal.

The transfer function for this low-pass filter is shown schematically in Figure 7 as $R(\square)$. The transfer function for the complementary high-pass filter for the $a(t) \cos(\square(t))$ term is shown as $C(\square)$. This filtering should also leave the spectrum of the amplitude function as shown by $A(\square)$, maintaining Bedrosian's separation from the minimum frequency of the $\cos(\square(t))$ term. All we have to do is determine the breakpoint frequencies, $\square$ and $\square/2$, for these filters and calibrate the horizontal scale.




Figure 7. Adaptive Filter Transfer Functions

The spectrum of the $a(t) \cos(\omega_c t)$ term will, in general, contain both AM and FM components. Amplitude modulation of a constant-frequency “carrier” signal shifts the spectrum of the amplitude signal from the origin to the carrier frequency. If $A(\omega)$ is the spectrum of $a(t)$, then the spectrum of $a(t) \cos(\omega_c t)$ will be $A(\omega \pm \omega_c)$, where ω_c is the carrier frequency. Frequency modulation redistributes the spectrum of its modulating signal in much more complex ways.

In a combined AM and FM signal, the FM spectrum overlaps and mixes with the AM spectrum so that separating the two components using a simple linear process (like conventional filtering) does not appear promising. The HHT process, however, is able to make a separation, although not always in exactly the same form as used to formulate sample inputs. (Remember, solutions satisfying the HHT monocomponent separation criteria are not unique.)

As a first approximation for the breakpoint for the high-pass filter pass band, the minimum peak-to-peak frequency of the signal over the time span covered by the filter¹ was used. This frequency is marked as ω_{min} on the axis in Figure 7. The pass band breakpoint for the high-pass filter was set to this frequency. The stop band breakpoint, based on our experience with warped filtering, was set to one-half this frequency. As signals pass through the filter their peak-to-peak frequencies are monitored and the filter coefficients are adjusted to track any changes.

¹ We note that Bedrosian’s spectral separation criteria, being based on integral transform analysis, must hold (theoretically) for all time, not merely for the time span covered by the filter. We conjecture that this rather severe constraint can be relaxed using more modern tight-frame analysis. We have not completed the analysis to formally confirm this, however, and proceed, taking it as an assumption.

6. Case Studies

In this section we present five case studies that illustrate and compare the results produced by the HHT and adaptive filtering approaches. The first example is a simple composite signal that serves as a reference for comparison with the second example. The second example is a steady-state AM signal. The third example is a steady-state FM signal. The fourth and fifth examples contain unit step changes in amplitude and frequency, respectively, and begin to explore the dynamic capabilities of the HHT and adaptive filtering mechanisms.

6.1 Simple Reference Example

The first example is a simple combination of constant amplitude sinusoids defined by:

- $s(t) = \cos(t) + 0.5 \cos(t/2)$

The graph of this function is shown in Figure 8, along with the signal trend (dotted line). The maximum timing between peaks is slightly greater than π , indicating the need for high- and low-pass filters with upper breakpoint frequencies at $\omega=0.97$. The result of filtering this signal, because of our selection of filter breakpoints, produces a nearly perfect separation of the two components, namely:

- $c_I(t) = \cos(t)$
- $r_I(t) = 0.5 \cos(t/2)$

The HHT sifting process produces nearly identical results. One difference is that HHT sifting approximates the trend using splines, so its trend is represented by a series of cubic polynomials pieced together at the peaks. These small differences are of little concern here. Our primary interest in this simple signal is its similarity to the next example.




Figure 8. Simple Two-Component Example Signal

6.2 Amplitude Modulated Example

The second example is a stationary amplitude-modulated signal defined by:

- $s(t) = (1 + 0.5 \cos(t/2)) \cos(t)$

The graph of this function is shown in Figure 9 along with its positive and negative envelope functions (dotted lines). Note that a very similar envelope could also be constructed for the previous example.




Figure 9. Example Amplitude-Modulated Signal

The differences between this and the first example are that the tall positive peaks are a little narrower and the shorter positive peaks are a little broader. The positive peaks have exactly the same values and timing. The negative peaks extend slightly lower (to -1.03) and their timing is shifted slightly toward the tall positive peaks. Another way to

examine these signals is to expand this example's definition and apply a trigonometric identity for the product of two cosines:

- $$s(t) = \cos(t) + 0.5 \cos(t/2) \cos(t)$$

$$= \cos(t) + 0.25 \cos(t/2) + 0.25 \cos(3t/2)$$

This shows that the difference between this and the previous example is a smaller coefficient for the $\cos(t/2)$ term and an additional higher-frequency term, $0.25 \cos(3t/2)$.

The maximum timing between peaks is again slightly greater than $\bar{\Delta}$, indicating the need for filters with upper breakpoint frequencies at $\bar{\Delta}=0.93$. The result of filtering this signal separates the lower-frequency ($\cos(t/2)$) term from the two higher frequency components; that is:

- $$c_{adapt}(t) = \cos(t) + 0.25 \cos(3t/2)$$
- $$r_{adapt}(t) = 0.25 \cos(t/2)$$

The high-frequency component produced by adaptive filtering, $c_{adapt}(t)$, is shown in Figure 10, along with its amplitude envelope. The instantaneous frequency of the adaptive filtering solution ranges from approximately 0.83 to 1.10, as shown in Figure 11.




Figure 10. High-Frequency Component Separated from the AM Signal by Adaptive Filtering




Figure 11. Instantaneous Frequency of the AM Signal Component Separated by Adaptive Filtering

The result produced by HHT sifting is quite different. The HHT sifting algorithm produces nearly the same low-pass trend result as in the first example, which has the same frequency but twice the expected amplitude:

- $c_{HHT}(t) = \cos(t) + 0.25 \cos(3t/2) - 0.25 \cos(t/2) + 0.0563$
- $r_{HHT}(t) = 0.5 \cos(t/2) - 0.0563$

The small constant terms in the HHT formulas offset the frequency modulation effects that result when the three cosine terms in $c_{HHT}(t)$ are combined. These effects are discussed in the next example.

The high-frequency component produced by the HHT sifting process, $c_{HHT}(t)$, is shown in Figure 12, along with the trend (dotted line). The amplitude envelope for this signal is constant, which makes the frequency modulation effects in the signal more prominent. The instantaneous frequency of this signal, shown in Figure 13, has a larger range than that for the adaptive filter solution. The instantaneous frequency of the HHT sifting solution ranges from approximately 0.69 to 1.19.




Figure 12. High-Frequency Component Separated from the AM Signal by the HHT Sifting Process




Figure 13. Instantaneous Frequency of the AM Signal Component Separated by HHT Sifting

Both solutions produce monocomponent high-pass components and band-limited trend signals, which is how the HHT objectives were characterized earlier. The adaptive filter produces a mixed AM and FM component with a smaller-amplitude trend signal. HHT sifting produces a purely FM component with larger frequency variations, and a larger-amplitude trend signal.

In this example, the HHT result also illustrates a classic example of signal aliasing. The HHT and warped filtering processes, being based on peak values, under-sample the input signal and misinterpret the energy from the higher-frequency ($\cos(3t/2)$) component, attributing it to the lower-frequency ($\cos(t/2)$) term. The extra energy in the HHT trend for this signal does not accurately reflect the energy contained in the input signal. Aliasing too often has unintended consequences, though, and is generally considered best avoided, if possible.

6.3 Frequency Modulated Example

The third example is a stationary frequency-modulated signal defined by:

- $s(t) = \cos(t + 0.5 \sin(t))$

The amplitude of this signal is constant but its phase increases nonlinearly. The graph of this function, shown in Figure 14, shows sharpened positive peaks and rounded negative peaks, much like solutions to Stokes's equation [2] (although this is not a solution to Stokes's equation).




Figure 14. Example Frequency-Modulated Signal

HHT analysis of this signal finds evenly spaced constant-valued positive and negative peaks. The trend function is a constant zero, and the separated component captures the entire signal. The instantaneous frequency derived from the HHT results, as shown in Figure 15, matches our expectations.

- $\varpi(t) = 1 + 0.5 \cos(t)$




Figure 15. Instantaneous Frequency of the FM Signal Component Separated by HHT Sifting

The adaptive filtering results are a bit more complicated to explain. The coefficients of the Fourier series for a frequency-modulated signal are defined in terms of Bessel functions. (See for example [13] or [14].) If the form of the signal is generalized to:

- $s(t) = A \cos(\omega_c t + \beta \sin(\omega_m t))$

where A represents the signal's constant amplitude, ω_c is its “carrier” frequency, β is the index of modulation, and ω_m is the modulating frequency, then the equivalent Fourier series is:

- $s(t) = A \sum J_n(\beta) \cos((\omega_c + n \omega_m) t)$

where J_n is the Bessel function (first kind) of order n . The summation, theoretically, ranges over integral values of n from $-\infty$ to ∞ . Bessel function values for small values of β , however, are essentially zero for all but a few terms. An approximate Fourier series for this signal is:

- $s(t) \approx -0.242 + 0.969 \cos(t) + 0.242 \cos(2t) + 0.031 \cos(3t)$

This representation of the signal shows that its nonlinear phase gives it a constant “DC” term as well as higher-frequency harmonic components. The filter breakpoint frequencies for this signal, determined from the signal's peak-to-peak timing, were $\omega=1$ and $\omega=1/2$. This produced the separation:

- $c_{adapt}(t) = 0.969 \cos(t) + 0.242 \cos(2t) + 0.031 \cos(3t)$

- $r_{adapt}(t) = -0.242$

The results for the high-pass component are shown in Figure 16, along with a smooth amplitude envelope connecting the absolute values of the peaks (dashed lines).




Figure 16. High-Frequency Component Separated from the FM Signal by Adaptive Filtering

These results differ from the monocomponent signal we started out with, although the basic shape of the input signal is preserved. The oscillating amplitude appears problematic, since the input signal contained no amplitude modulation. Furthermore, the amplitude oscillations have the same average frequency as the signal, which violates Bedrosian’s spectral separation conditions. These amplitude oscillations appeared because the adaptive filtering process removes the constant term in the signal’s Fourier series. Our earlier time-warp analysis showed that the amplitude envelope should be band limited to below one-half of the signal’s “carrier” frequency. The observed higher-frequency content, therefore, is an unexpected artifact that must be attributed to the adaptive filtering process.

Similar nonlinear signal behavior was encountered in the previous (AM) example. The high-frequency component separated by HHT sifting (shown in Figure 12) contains alternating narrow and wide positive peaks. This nonlinear phase behavior gives this signal a constant term similar to that described here. As these examples show, any signals with nonlinear phase behavior can potentially introduce similar artifacts in adaptive filter results.

The instantaneous frequency derived from the high-pass adaptive filter output is shown in Figure 17. This signal has a smaller frequency range than the HHT component ($\omega=0.79$ to 1.28) and the variations are not purely sinusoidal.




Figure 17. Instantaneous Frequency of the FM Signal Component Separated by Adaptive Filtering

6.4 Amplitude Step Example

The preceding examples are all stationary signals that could be handled by static filtering techniques (if the frequencies are known in advance). The signal shown in Figure 18 begins to exercise the dynamic capabilities of the HHT and adaptive filtering processes. This signal contains a step discontinuity in its amplitude at time $t=0$. That is,

- $s(t) = \sin(t)$ for $t \leq 0$
- $s(t) = 2 \sin(t)$ for $t \geq 0$




Figure 18. Amplitude Step Example Signal

Both the HHT and adaptive filtering processes are expected to smooth out this amplitude transition because of the bandwidth limitations on component amplitude envelopes suggested by the monocomponentness considerations. The results plotted in Figures 19

and 20 show that this is indeed the case. The differences in smoothing are a result of the differing filter transfer functions and, in the case of the HHT, its signal aliasing behavior. The trend signals in both cases are shaped somewhat like sampling functions. The HHT trend has considerably higher amplitude than the adaptive filter low-pass signal.

There is also a time delay of approximately 24 time units for the incremental HHT result and 25 time units for the adaptive filtering results. These delays are necessary to collect data on the signal's future behavior, which both processes need before they can produce their results.




Figure 19. HHT Component and Trend Results for the Amplitude Step Signal




Figure 20. Adaptive Filter High- and Low-Pass Results for the Amplitude Step Signal

The instantaneous frequencies, derived numerically, for the two separated high-pass components are shown in Figures 21 and 22. In both cases, the effect of smoothing out the amplitude step transient has created transient frequency modulations. This suggests the presence of a “conservation of transient energy” law that allows amplitude transients to be transformed into frequency transients.




Figure 21. Instantaneous Frequency of the Amplitude Step Component Separated by HHT Sifting




Figure 22. Instantaneous Frequency of the Amplitude Step Component Separated by Adaptive Filtering

Although our understanding of this frequency behavior is incomplete, we can explain the behavior of the two signal separation processes using their representation in the frequency domain. The Fourier transform of the amplitude step signal is:

$$S(\omega) = j3\omega[\omega\omega+1) - \omega\omega-1)/2 + 1/(1-\omega^2)$$

Figure 23 shows the magnitude of this transform. It has complex poles at $\omega=\pm 1$, which reflects the $\sin(t)$ term in the signal. The bandwidth contributed by the amplitude step is distributed smoothly over the entire frequency spectrum.




Figure 23. Fourier Transform (Magnitude) of the Amplitude Step Signal

Figure 24 shows how adaptive filtering separates the amplitude step signal in the frequency domain. The low-pass (solid) curve shows the spectrum of the signal trend and the high-pass results (dashed) curve shows the spectrum of the separated component. Inverting these transforms back into the time domain reproduces the trend and component signals shown in Figure 20.²

² Care must be taken with numerical Fast Fourier Transform (FFT) tools in analyzing these signals and spectra. The results presented here are for continuous infinite-integral transforms of one-time transient events. Numerical techniques that operate on finite-duration numerical representations of signals and their spectra can easily generate different results. For example, a finite representation of the signal shown in Figure 18 will be presumed to repeat periodically. While the graph still looks like a one-time unit step amplitude change, the transform produced will be for a repeating square-wave modulated signal.




Figure 24. Adaptive Filter High- and Low-Pass Spectra for the Amplitude Step Signal

In practice, the results shown back in Figure 20 are produced by direct convolution of the signal with the digital filter coefficients, not by applying transforms. The results, however, are the same by either process.

Explanation of the HHT results requires introducing the effects of the warped filter resampling process and the filter's transfer function. Figure 25 shows the spectra of the signals separated by the warped filter. The low-frequency “hump” (solid line) is the trend's spectrum. The second curve (dashed line) is the spectrum of the separated high-frequency component. Transforming these spectra back into the time domain reproduces the signal trend and separated component shown in Figure 19.

The third curve in Figure 25 (dotted line) shows the spectrum of the warped signal that was derived by resampling the input signal at its peaks. This is a direct effect of aliasing. Because the peak sampling rate is below the signal's original sampling rate, aliasing creates overlapping replicas of the spectrum shown in Figure 23. As can be seen by comparing the results shown in Figure 24, aliasing has a significant effect on the apparent spectrum processed by the warped filter and it imparts considerable energy to the trend that is not part of the input signal. The separated high-frequency component is calculated by resampling the trend at its original sample times and subtracting that result from the original input signal. This component, therefore, is only affected by the trend signal, not by the aliased spectrum.




Figure 25. Spectra of the HHT Trend and Separated Component for the Amplitude Step Signal

6.5 Frequency Shift Example

The final example signal to be explored contains a step discontinuity in frequency at time $t=0$. A graph of this signal is shown in Figure 26.

- $s(t) = \sin(t)$ for $t \leq 0$
- $s(t) = \sin(2t)$ for $t \geq 0$




Figure 26. Frequency Shift Example Signal

Because the signal amplitude is constant, the HHT trend remains constant (zero) through this frequency shift. There is no effect from aliasing because the trend is zero. The first separated HHT component captures the entire input signal. It seems clear from this and the earlier frequency-modulated example that the HHT will separate any constant-amplitude, monotonically increasing phase signal as a single component.

The instantaneous frequency extracted from the signal, which is the HHT-separated component, is shown in Figure 27. It tracks the signal nearly perfectly through the transition. While the HHT produced considerable smoothing of the amplitude step in the previous example, it makes no attempt to smooth out the frequency shift here.




Figure 27. Instantaneous Frequency of the Frequency Shift Component Separated by HHT Sifting

Adaptive filtering produces quite different results, as shown in Figure 28. The high-pass signal (solid line) shows a clear disturbance, although it is difficult to characterize. The low-pass signal (central dotted line) looks something like an inverted sampling function, centered at the point where the frequency shift takes place. The amplitude envelope around the high-frequency signal (upper and lower dotted lines) also reflects the disturbance.




Figure 28. Adaptive Filter High- and Low-Pass Results for the Frequency Shift Signal

The instantaneous frequency, derived numerically, for the high-pass component signal is shown in Figure 29. As with the previous example, we do not fully understand why the frequency behavior should take this shape.




Figure 29. Instantaneous Frequency of the Frequency Shift Component Separated by Adaptive Filtering

As with the previous example, we turn to the frequency domain to explain the behavior of the adaptive filter. The Fourier transform of the frequency step signal is:

$$\bullet \quad S(\omega) = j\omega[\omega(\omega+1) - \omega(\omega-1)]/2 + j\omega[\omega(\omega+2) - \omega(\omega-2)]/2 - 1/(1-\omega^2) + 2/(4-\omega^2)$$

The magnitude of this transform is shown in Figure 30. The complex poles at $\omega = \pm 1$ and $\omega = \pm 2$ reflect the signal's two sinusoidal frequencies. The bandwidth contributed by the frequency transition is distributed smoothly over the entire spectrum.




Figure 30. Fourier Transform (Magnitude) of the Frequency Shift Signal

Figure 31 shows how adaptive filtering separates the frequency shift signal in the frequency domain. The low-pass curve (solid line) shows the spectrum of the signal trend. The high-pass curve (dashed line) shows the spectrum of the separated component. Inverting these transforms reconstructs the signals shown in Figure 28. The results shown in Figure 28, though, were calculated by direct convolution of the signal with the digital filter coefficients, not by applying transforms.




Figure 31. Adaptive Filter High- and Low-Pass Spectra for the Amplitude Shift Signal

Because the filter breakpoint frequencies are determined by the lowest peak-to-peak frequency within the span of the filter, these results are effectively the same as for static filters with breakpoints at $\Omega=1/2$ and $\Omega=1$. Once the last low-frequency peak passes through the filter, its coefficients are adjusted to move the breakpoint frequencies to $\Omega=1$ and $\Omega=2$. This has no effect on the filter outputs because in both cases the signal resides completely within the high-pass pass band.

7. Summary and Conclusions

The HHT component separation, or “sifting,” process has been compared with an adaptive filtering process that was intended to mimic HHT behavior. The conjecture that conventional digital filters, by adapting dynamically to signal frequency content, could substitute for the HHT process was found to be incorrect. Results from several example signals showed that under most conditions the two techniques produce distinct results.

The experiments we conducted to compare the HHT and adaptive filtering processes led to the discovery of aliasing in the HHT sifting algorithm. The process of sampling a signal at its peak times results in a classic example of under-sampling that leads to misinterpretation of signal frequency content. Specifically, signal content at frequencies above the peak-to-peak sampling rate is misinterpreted as lower-frequency content.

The question of whether aliasing is a problem or a “feature” in terms of HHT signal separations has not yet been completely resolved. Results from both aliased (HHT) and non-aliased (adaptive filter) processes appear to satisfy the requirements for “monocomponentness,” so separated components are expected to have well-defined instantaneous frequencies. Ordinarily, though, aliasing is considered a form of signal corruption that is best avoided whenever possible. Further investigation is needed to determine if unaliased filtering results are indeed “better,” or if the aliasing is in some unusual way a necessary aspect of the HHT sifting process.

7.1 Summary of Case Study Findings

For signals with a dominant highest frequency (case study #1), the HHT and adaptive filtering were found to produce equivalent separations.

For stationary amplitude-modulated signals with a dominant central “carrier” frequency (case study #2), adaptive filtering separates the lower sidebands as the trend, and the carrier and upper sidebands as the high-frequency component. The HHT, because of aliasing, misinterprets the upper sideband energy as lower-frequency energy, effectively doubling the lower sideband amplitude. This gives the high-frequency component a nearly constant amplitude and larger variations in instantaneous frequency.

For signals with transient amplitude changes (case study #4), HHT sifting produced a broad smoothing of the amplitude transition and, because of aliasing, a large trend amplitude. Adaptive filtering also smoothed out the amplitude transition, but not as broadly as the HHT. Its trend amplitude was small compared to the HHT trend.

For frequency-modulated signals with monotonically increasing phase (case studies #3 and #5), the HHT high-frequency component captures the entire signal, leaving a zero-valued residual trend. The extracted phase function, $\varphi(t)$, and instantaneous frequency, $\varphi'(t)$, for these signals tracked the signal behavior very closely, even with significant transients in frequency (case study #5). Adaptive filtering had considerably more difficulty with FM signals. Signals with nonlinear phase functions often have significant low-frequency content. Conventional filtering separates the high- and low-frequency energy, disrupting the input signal's monocomponent characteristics.

7.2 Research Directions

Although this paper investigated a key step in separating signal components, there are additional aspects of the overall problem that need attention. The following research areas have been identified as areas still to be explored.

Resolving the question about aliasing is of high priority. Our preference for a solution would be an algorithm that separates complex signals into components without aliasing, and without the amplitude disturbances adaptive filtering causes with FM signals.

Second on our list is finding a better way to separate amplitude and phase information from monocomponent signals. Although the Hilbert transform is the obvious theoretical solution, current finite numerical approximations produce anomalous results.

Episodes of signals with only very low-frequency content compared to their sampling rate (that is, with many samples between peaks) would require excessively long filters to achieve the separations we propose. This corresponds to shifting φ way to the left in Figure 7. To process such signals a method is needed for adaptively down-sampling or decimating the signal, and automatically restoring higher sampling rates when higher-frequency content returns. Static down-sampling is used extensively in wavelet transform processing [15]. To our knowledge, the idea of a dynamic down-sampling mechanism is yet to be explored.

The residual trend signals that are passed to successive stages of sifting have their high-frequency content removed, resulting in signals with lower and lower frequency content. This is a prime example of where signal down sampling is needed. Non-uniform sampling techniques [16] may be useful here, although they appear to require more complex up-

sampling procedures to restore their original sampling rates than do uniformly sampled signals.

Real-world signals often contain components that turn on and off intermittently, like the telephone that rings while you are listening to your favorite music or eating dinner. Huang developed a technique for dealing with such intermittent components that attempts to minimize the disturbance in analysis of more continuous “background” components [17]. Although there is a clear need for this capability, it has not yet been addressed in our real-time algorithms.

References

- [1] Cohen, L., *Time-Frequency Analysis*, Prentice Hall, 1995.
- [2] Huang, N. E., Z. Shen, S. Long, M. Wu, H. Shih, Q. Zheng, N. C. Yen, C. C. Tung, and H. Liu, "Empirical mode decomposition and the Hilbert spectrum for nonlinear and non-stationary time series analysis," *Proc. Royal Soc. London, A*, 454, 1998, pp. 903-995.
- [3] Meeson, R., *An Incremental, Real-Time Algorithm for the Hilbert/Huang Transform*, IDA Paper P-3656, Institute for Defense Analyses, 2002.
- [4] Oppenheim, A. and R. Schafer, *Discrete-Time Signal Processing*, Prentice Hall, 1989.
- [5] Huang, N. E., C. C. Chern, K. Huang, L. W. Salvino, S. R. Long, and K. K. Fan, "A new spectral representation for earthquake data: Hilbert spectral analysis of station TCU129, Chi-Chi, Taiwan, 21 September 1999," *Bulletin of the Seismological Society of America*, vol. 91, no. 5, October, 2001, pp. 1310-1338.
- [6] Bedrosian, E., "A product theorem for Hilbert transforms," *Proc. of the IEEE*, vol. 51, no. 5, May 1963, pp. 868-869.
- [7] Nuttall, A. H., "On the quadrature approximation to the Hilbert transform of modulated signals," *Proc. of the IEEE*, vol. 54, no. 10, October 1966, pp. 1458-1459. (with reply by Bedrosian)
- [8] Maragos, P., J. F. Kaiser, and T. F. Quatieri, "On amplitude and frequency demodulation using energy operators," *IEEE Trans. On Signal Processing*, vol. 41, no. 4, April 1993, pp. 1532-1550.
- [9] Maragos, P., J. F. Kaiser, and T. F. Quatieri, "On separating amplitude from frequency modulations using energy operators," in *Proc. IEEE Intl. Conf. on Acoustics, Speech, and Signal Processing*, vol. 2, March 1992, pp. 1-4.
- [10] Kincaid, D. and W. Cheney, *Numerical Analysis*, Brooks/Cole Publ., 1991.
- [11] Parks, T. W. and C. S. Burrus, *Digital Filter Design*, John Wiley, 1987.

- [12] Boashash, B., “Estimating and interpreting the instantaneous frequency of a signal—Part 1: Fundamentals,” and “—Part 2: Algorithms and applications,” *Proc. of the IEEE*, vol. 80, no. 4, April 1992, pp. 520-568.
- [13] Lathi, B. P., *Signals, Systems and Communication*, John Wiley & Sons, 1965.
- [14] Schwartz, M., *Information Transmission, Modulation, and Noise*, 4th Ed., McGraw-Hill, 1990.
- [15] Daubechies, I., *Ten Lectures on Wavelets*, SIAM, Philadelphia, PA, 1992.
- [16] Aldroubi, A. and K. Gröchenig, “Nonuniform sampling and reconstruction in shift-invariant spaces,” *SIAM Review*, vol. 43, no. 4, 2001, pp. 585-620.
- [17] Huang, N. E., Z. Shen, and S. R. Long, “A new view of nonlinear water waves: the Hilbert spectrum,” *Annual Review of Fluid Mechanics*, vol. 31, 1999, pp. 417-457.

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YY) August 2003		2. REPORT TYPE Study (Final)		3. DATES COVERED (From – To)	
4. TITLE AND SUBTITLE HHT Sifting and Adaptive Filtering				5a. CONTRACT NUMBER DASW01-98-C-0067/ DASW01-02-C-0012	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBERS	
6. AUTHOR(S) Meeson, Reginald, N.				5d. PROJECT NUMBER	
				5e. TASK NUMBER IDA Central Research Program C5052	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESSES Institute for Defense Analyses 4850 Mark Center Drive Alexandria, VA 22311-1882				8. PERFORMING ORGANIZATION REPORT NUMBER IDA Paper P-3766	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Institute for Defense Analyses 4850 Mark Center Drive Alexandria, VA 22311-1882				10. SPONSOR'S / MONITOR'S ACRONYM IDA	
				11. SPONSOR'S / MONITOR'S REPORT NUMBER(S) IDA Paper P-3766	
12. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; unlimited distribution: 17 December 2003.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT The Hilbert/Huang Transform (HHT) is a time-frequency analysis technique that offers higher frequency resolution and more accurate timing of transient and non-stationary signal events than conventional integral transform techniques. Real-time HHT algorithms enable this enhanced signal analysis capability to be used in process monitoring and control applications. This paper compares real-time sifting with adaptive filtering, which we conjectured might be an efficient substitute. While HHT sifting is analogous to filtering, the two techniques produce distinct results. Sifting, for example, will pass purely frequency-modulated signals with low-frequency content, where adaptive filtering will separate the high- and low-frequency content. Sifting was found to introduce aliasing, which is often considered a form of signal corruption.					
15. SUBJECT TERMS Time-Frequency Analysis, Hilbert/Huang Transform, HHT Sifting Algorithm, Empirical Mode Decomposition, Real-Time HHT Sifting, Adaptive Filtering, Transient Signals, Signal Aliasing.					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Unlimited	18. NUMBER OF PAGES 58	19a. NAME OF RESPONSIBLE PERSON Dr. Reginald N. Meeson
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified			19b. TELEPHONE NUMBER (Include Area Code) (703) 845-6619

