Implementation/Acceleration of Vectorized Acoustic Source Localization Algorithms on the CELL Multi-core Processor

Neena Imam, Jacob Barhen, and Travis Humble

Center for Engineering Science Advanced Research, Oak Ridge National Laboratory, Oak Ridge, Tennessee 37830, USA

imamn@ornl.gov, barhenj@ornl.gov, humblets@ornl.gov

***Abstract*— This paper describes the optimized design and implementation of sonar signal processing algorithms (matched filter) on the CELL multi-core processor. The algorithm is modified to achieve maximum parallelism to enhance time performance on the multi-core platform. In this proof-of-concept effort, we show a factor of 10 speedup by parallelizing the matched filter code on 16 separate computing nodes on the CELL processor for broadband matched filter implementation. This demonstration of considerably faster signal processing capability should be of substantial interest to the signal processing community in general, and the oceanic engineering community in particular.**

***Index Terms* — Doppler Sensitive Waveform (DSW), Matched Filter (MF), Multi-core, Fast Fourier Transform (FFT), Ambiguity Function (AF).**

I. Introduction

For real-time acoustic source localization applications, one of the primary challenges is the considerable growth in computational complexity associated with the emergence of ever larger, active or passive, distributed sensor networks. The complexity of the calculations needed to achieve accurate source localization increases dramatically with the size of sensor arrays, resulting in substantial growth of computational requirements that cannot be met with standard computational resources. One option to meet this challenge builds upon the emergence of multi-core computing platforms. The objective of this work is to explore the implementation of key building block algorithms used in underwater source localization on the CELL multi-core processor. We investigate key concepts of threat-detection algorithms such as Matched Filter (MF) techniques with the purpose of implementation on the CELL multi-core processor. The major accomplishments of this research, in terms of computational speedup achieved via the utilization of multi-core computing technology, should be of substantial interest to the signal processing community.

II. Technical Background and Related Work

The Center for Engineering Science Advanced Research (CESAR) at the Oak Ridge National Laboratory is involved in the development and demonstration of exciting unconventional technologies for computing and signal processing. The CESAR efforts in the area of advanced signal processing are driven by the emergence of powerful new processors such as the IBM CELL multi-core platform, Intel Polaris, Coherent Logix the HyperX, etc. The algorithm under consideration in the present study is MF correlation processing via Fast Fourier Transform (FFT) of broadband Doppler-Sensitive Waveforms (DSW). The specific goals of this proof-of-concept effort is to demonstrate the ability to achieve required accuracy in the computations and to quantify the speed-up achieved via parallelization on the CELL multi-core platform as compared to a leading-edge conventional processor (Intel-Xeon or DSP).

It is possible to achieve superior performance of active sonar systems via proper waveform selection, accurate signal/system modeling, and efficient real-time signal processing via MF bank implementation. The common waveforms used in active sonars have diverse and complimentary characteristics. For example, Constant Frequency (CF) pulses provide superior range-rate estimation but poor range resolution capabilities [1]. The reverberation clutter power vs. Doppler shift of a CF pulse is also more concentrated than that of Linear Frequency Modulated (LFM) signals, another common candidate waveform in active sonar systems. To resolve the inherent conflict between reliable detection and good range resolution, signals other than the simple CF pulse have to be used [2]. The next subsection discusses the major characteristics of broadband sonar signals that offer desirable features such as better range accuracy, improved minimum range of detection, greater gains in the signal to noise ratio, and better Doppler tolerance. However, all these features come at the costs of increased signal processing challenges and broadband signal modeling complexity. The following subsections present the major considerations of broadband sonar signal processing such as Doppler stretch factor for broadband signals, signal processing via MF implementation, and acceleration of the MF processing via multi-core processor utilization.

A. Broadband Doppler-Sensitive Waveform

The prime motivation for using broadband signals in active sonar platforms is to improve the range resolution capabilities while preserving reliable detection performance. This is determined by the total amount of energy in the echo signal reflected from the target. This can be achieved by employing signals with increased bandwidth and significant energy distribution over a wide range of frequencies. Signals with large time bandwidth products are classified as broadband signals. In most long-range active sonar receivers, a MF is used to reject ambient noise. The MF is a correletor whose output is characterized by the AF. The two dimensional AF, $|\chi_s(\tau,\varphi)|^2$, maps the correlator output power against

978-1-4244-5222-4/10/$26.00 ©2010 IEEE

range (determined by time delay τ) and Doppler frequency shift φ. The expression $\chi_s(\tau,\varphi)$ for an arbitrary waveform *s(t)* is

$$\chi_s(\tau,\varphi) = \int_{-\infty}^{\infty} s(t)s^*(t+\tau)e^{j2\pi\varphi t}dt. \qquad (1)$$

The AF specifies the output of the matched filter in the absence of noise. Thus the two parameter correlation function of the signal waveform, with respect to both time displacement τ and frequency displacement φ, is a measure of the relative response of the correlator to the echo and reverberation signals. The region of the τ and ϕ domain, within which the AF function still has a significant peak relative to the central peak, is called the region of signal ambiguity. The AF is a useful tool for sonar system analysis and provides a starting point for waveform synthesis. The waveform synthesis problem, which consists of finding the appropriate signal given a specified ambiguity function, is a challenging problem, particularly for broadband signals. For narrowband signals, all signals are modified replicas of the original signal with a frequency displacement. For broadband signals, the effect of target velocity can no longer be approximated by a simple frequency shift since the bandwidth of the transmitted signal is a considerable fraction of the mean signal frequency. In addition to the frequency shift, the Doppler effect also induces a time compression/stretching of the transmitted broadband signal. Kelly and Wishner defined the Wideband Ambiguity Function (WAF) as a function of time delay τ and the Doppler scale factor η. Thus, in terms of η, one can write

$$\chi_s(\tau,\varphi) = \sqrt{\eta}\int_{-\infty}^{\infty} s(t)s^*[\eta(t-\tau)]dt \qquad (2),$$

where $\eta = (1+\upsilon/c)/(1-\upsilon/c)$, c is the speed of signal propagation, and v is the target velocity [3]. The sonar operating environment is extremely hostile due to ambient noise and reverberation interference. Reverberation is the sum of signals returned from unwanted paths, producing large amplitudes in the correlator output at the τ and φ values where the target echo signal is to be manifested. The sonar system response to reverberation is very much dependent on the waveform design. Range and Doppler resolutions, as well as the distribution sidelobes of the AF, affect the reverberation level, limiting the performance of active sonars in shallow water/littoral environment. The relative interference power levels of the different waveforms can be approximated from the proportions of the reverberation distribution that is enclosed by the AF of the target echo. The problem of echo to reverberation gain maximization is to choose a waveform such that a minimum of the reverberation signal distribution lies within the ambiguity region, while the central peak of the correlation function intercepts the parameter space of the expected echo signals. High reverberation processing gain may be achieved by minimizing the area under the square of the modulus of the Doppler cross-power spectrum along a line of constant Doppler scaling [4]. A common scheme for doing this is to spread the energy of the transmitted pulse over a broad bandwidth in the case of Frequency Modulated (FM) signals. If the FM function of a waveform is periodic, the spectrum of the waveform will consist of many individual spectral lobes spaced at multiples of the repetition frequency. The energy of the pulse will therefore be spread out in the frequency domain. The periodic sidelobes of the AF will enclose a smaller portion of the reverberation distribution than a CF pulse, thus leading to a corresponding decrease in the reverberation level. One of the most common types of comb-spectrum is a periodic Sinusoidal Frequency Modulated (SFM) waveform. The AF function for a SFM signal is shown in Figure 1. The simplest case of a signal with comb-like spectrum is the Sinusoidal Frequency Modulated (SFM) signal where the modulating waveform itself is also a sinusoid $s(t) = w(t)\exp^{j2\pi f_o t + j\beta\sin(2\pi f_m t)}$

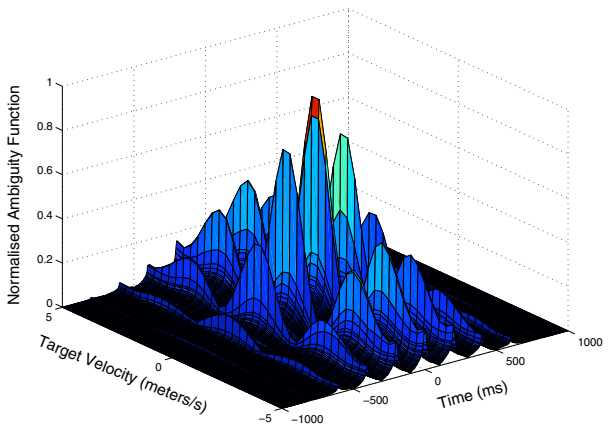

Figure 1. Wideband ambiguity function of an SFM signal.

B. Matched-Filter for Broadband Doppler-Sensitive Waveform

The MF is central to sonar signal processing [5]. As discussed earlier, the MF compares the received signal with a hypothesized signal. The output of the MF gives a measure of how well the hypothesized signal matches the received signal as functions of a set of parameters, usually the range and velocity of targets. A matched filter isolates a signal in signal space as a space array isolates a direction in real space. Modern day sonar platforms employ a bank of MFs, the ith of which is tuned to a particular Doppler frequency. The target range and radial velocity can be obtained by passing the received signal through a bank of MFs, where each filter in the array is matched to a different target velocity, as shown in Figure 2. The output of each filter is fed to a threshold detector. When the threshold is exceeded, the time delay provides target range estimates. The estimated velocity is that for which the correlation peak magnitude of the filter output has the maximum value. The output of the correlator is calculated via Fast Fourier Transform (FFT), followed by

inverse FFT. The vectors representing the discrete replica and echo signals can have considerable size in the case of broadband signals and the size of the FFT can easily exceed *500K* samples. Hence, it is anticipated that the very computationally expensive implementation of broadband matched filtering will be expedited on multi-core processors.

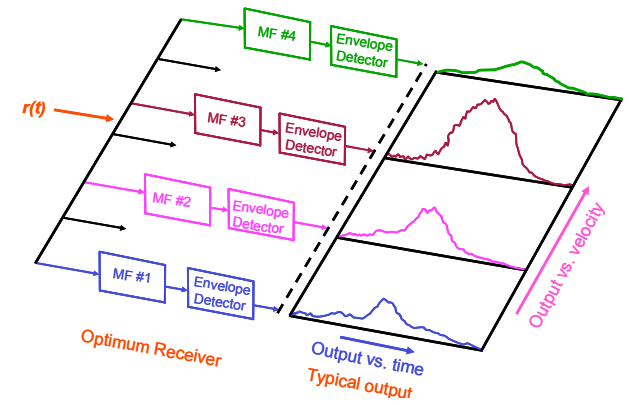

Figure 2. Schematic of a matched filter

For this benchmark study, the source was assumed to transmit an SFM waveform at a center frequency f_0=1200 Hz, with a 400 Hz bandwidth B. The pulse duration used was T = 1s, the frequency was modulated at 5 Hz, and the Doppler scale ranged from -5 to +5 m/s. The sampling frequency f_s assumed was 5000 Hz, resulting in a sampling interval length of T_s = 0.0002 s. These pulse parameters are chosen to fall within the range of operational parameters used in active sonar systems employing DSW signals. For this benchmark study, the assumed target range and velocity were 3 Km and –

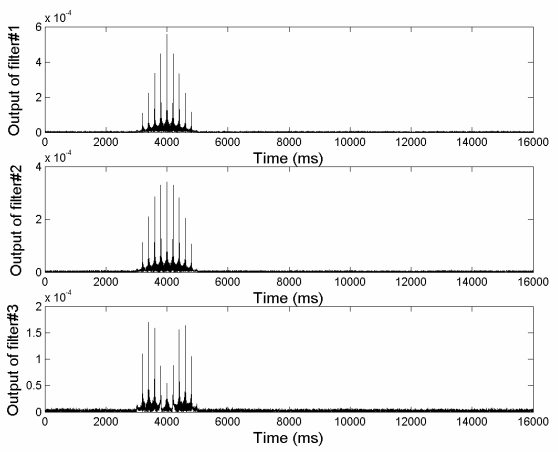

Figure 3. Matched filter output of the 32 bin MF

5 m/s respectively. A MF bank with 16 Doppler bins was implemented with each filter performing an 8K-sample complex FFT to calculate the cross-power spectrum and an 8K-sample inverse FFT to obtain the cross-correlation output of the filter. Figure 3 shows the outputs of three filters of the MF bank obtained via MATLAB simulation (all 16 filters are not shown due to space consideration). The output of filter #1 has the closest velocity match to the received signal. The estimated target delay is 4 s and hence the estimated target range is 3 Km, with incoming direction. To continue on with our research in signal processing via leading-edge unconventional processors, we will present the implementation of the MF algorithm on the CELL multi-core processor and compare the results with those obtained via the sequential code

III. Implementation on the CELL Multi-core Platform

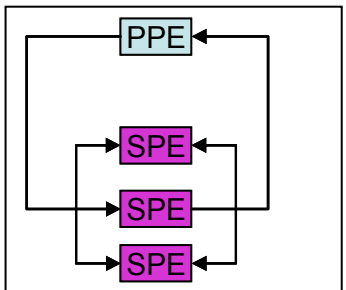

Figure 4. The parallel pipeline model

One of the multi-core processors commanding significant attention recently is the CELL Broadband Engine (CBE). The CBE is a joint product of Sony, Toshiba, and IBM, also known as STI. Although originally intended for use in game consoles, the CBE has enabled fundamental advances in processor performance. The CBE is a single-chip multi-processor consisting of nine processors sharing memory. The processors fall into two distinct groups: the power PC processor element (PPE), and the synergistic processor element (SPE). The PPE is intended to run the operating systems and the top level control thread of the applications. The SPE is the workhorse of the platform and is performs compute-intensive tasks. The SPEs are endowed with a rich instruction set that includes single-instruction, multiple data functionality (SIMD). SIMD processing exploits data-level parallelism meaning that the operations required to transform a set of vector elements can be performed on all elements of the vector at the same time. One level of parallelism of the CBE architecture is due to the fact that is a single instruction can be applied to multiple data elements in parallel. Another level of parallelism arises from the fact that programs running on the CELL processor elements can be partitioned among the available processor elements. In determining the optimized distribution of the workload to achieve maximum level of parallelism one needs to take into account algorithm structure, program data flow, data access pattern, and cost of code/data movement among processors, as presented in our previous work [6-8]. For the present research, we have employed the PPE centric application model to run the main application on the PPE and to off-load the compute-intensive tasks to the SPEs. For the MF algorithm, there is a large amount of data that can be partitioned and acted on at the same time. Therefore it makes sense to use SPEs to process different parts of the data in parallel. This parallel stage pipeline model is shown in Figure 4. In order to execute the acoustic sensor signal processing code, we needed to modify the algorithm

and parallelized the code for execution across multiple SPEs. It is also possible to vectorize/SIMDize the code to enhance performance by employing SIMD strategy appropriate for our algorithm. Another strategy is to employ communication-computation concurrency. For example, we can start communication for the next piece of data while processing the current piece as shown in Figure 5. After a parallel program executes without error on the PPE and the SPEs, further optimization through parameter-tuning is possible. Algorithms can be analyzed to eliminate data stalls and instruction dependency stalls. SPE timing analyzer performs a static timing analysis to determine what program sections may be susceptible to stalls. For our MF example, we partitioned the FFT and inverse FFT calculations for the 16 Doppler banks into 16 parallel threads. These threads were then executed on the 16 available SPEs. We found a factor of 10 speedup for the parallel code in comparison with the sequential version (all the calculations were done on the PPE).

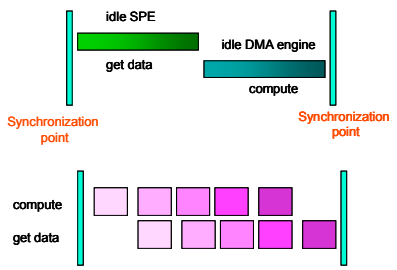

Figure 5. Compute-communicate concurrency

IV. DISCSSION

In our manuscript we aim to motivate a paradigm shift in sonar signal processing via utilization of computing platforms with multiple cores. We described the efficient and optimized design of acoustic sensor signal processing algorithms that are specifically designed for the CELL multi-core compute engine. These algorithms can be both vectorized and parallelized to be able to fully utilize the tremendous computing power of the CELL multi-core architecture. Scalability analysis and various optimization techniques for MF algorithms implemented on multi-core processors will be presented in our future works.

REFERENCES

[1] W. Burdick, *Underwater Acoustic System Analysis*, New York, United States: Prentice Hall, 1984.

[2] J. L. Stewart and E. C. Westerfield, “A theory of active sonar detection," Proc. IRE, vol. 47, pp. 872-881, May 1959.

[3] E. J. Kelly and R. P. Wishner, “Matched filter theory for high velocity accelerating targets," IEE Trans. Mil. Electron., vol. 9, pp. 56-69, 1965.

[4] T. Collins and P. Atkins, “Doppler-sensitive active sonar pulse designs for reverberation processing," IEE Proceedings - Radar, Sonar, Navigation, vol. 145, pp. 347-353, December 1998.

[5] G. L.Turin, “An introduction to digital matched filter”, Proc. IEEE, vol. 64, pp. 1092-1112, July 1976.

[6] J. Barhen and N. Imam, “Sensor data processing for tracking underwater threats using terascale optical core devices, *NATO ARW Book Series,* IOS Press (in press).

[7] N. Imam and J. Barhen, “Acoustic Source Localization via Distributed Sensor Networks using Tera-scale Optical-Core Devices”, Acoustics 2008 Conf, June, 2008, Paris, France.

[8] N. Imam and J. Barhen, “Efficient Replica Generation for Space Time Adaptive Processing of Acoustic Signals”, Acoustics 2008 Conf, June, 2008, Paris, France.