(cover sheet)

Parallel Matched-Field Tracking (MFT) for Distributed Deployable Systems

J. Han, B. Koh, A. George, and K. Kim

High-performance Computing and Simulation (HCS) Research Laboratory
Department of Electrical and Computer Engineering, University of Florida
P.O. Box 116200, Gainesville, FL 32611-6200

Corresponding Author

Dr. Alan D. George
High-performance Computing and Simulation (HCS) Research Laboratory
Department of Electrical and Computer Engineering
University of Florida
216 Larsen Hall, P.O. Box 116200
Gainesville, FL 32611-6200

Phone: (352)392-5225
Fax: (352)392-8671
e-mail: george@hcs.ufl.edu

Parallel Matched-Field Tracking (MFT) for Distributed Deployable Systems

J. Han, B. Koh, A. George, and K. Kim

High-performance Computing and Simulation (HCS) Research Laboratory
Department of Electrical and Computer Engineering
University of Florida, Gainesville FL

Quiet submarine threats and high clutter in the littoral undersea environment demand the development and use of enhanced and new acoustic processing algorithms with increased sophistication. These algorithms exhibit high levels of computational complexity and memory utilization, making implementation in real-time sonar array systems a significant challenge. Concomitant with the increase in demand for computing resources implied by new acoustic processing algorithms, mission requirements continue to transition toward the goal of autonomous, in-situ processing with minimal off-array communication and battery power consumption. Taken together, these trends make imperative the development and use of advanced distributed and parallel processing techniques in terms of algorithm, architecture, network, and system design. In that regard, this presentation focuses on the design and analysis of several novel parallel algorithms for a prominent algorithm in sonar array processing, Matched-Field Tracking (MFT), and includes promising experimental results from a distributed array testbed comprised of a network of SHARC processors.

In a shallow-water acoustic environment, sonar signals propagate as a waveguide and the sounds at the boundaries are measured with hydrophones. Matched-Field Processing (MFP) is a method to exploit this dispersive part of the wave in order to estimate the source position. The general approach involves correlating pressure fields at the receivers and matching them with calculated fields based on an appropriate mathematical model of environment. However, since MFP algorithms search all possible locations for an unknown acoustic source within a surveillance region, implementation for real-time applications can be extremely challenging because of their high computational complexity and memory requirements.

The Matched Field Tracking (MFT) algorithm was devised by Bucker et al. [1-2] to reduce the computation and memory requirements of MFP in real-time applications. MFT correlates the values of possible grid points and computes the location of the target track based on information obtained by processing the data on a wide time window. One of the more recent variants of the MFT algorithm is the Hydra algorithm, which is devised for a sonar processing system consisting of a horizontal line array of hydrophones. This algorithm serves as the basis for the parallel algorithms developed for this research. Processing in the Hydra algorithm takes place in four stages, those being the frequency selection stage, the replica vector generation stage, the initial tracking stage, and the tracking adjustment stage. First, the averaging and selecting of the strongest frequencies are performed in the frequency selection stage through each track period. Next, to estimate the sound source location, the expected field data from the model and the measured field data from the sensors are exploited. The replica vectors, which represent the modeled acoustic pressure field, are generated from a normal-mode underwater acoustic propagation model. After the replica vector table has been computed, the initial tracking stage is performed in order to estimate multiple track locations using a coarse grid of data points. Finally, in the tracking adjustment stage, the tracks obtained are corrected with the purpose of optimizing the accuracy on a fine grid, and the result is a fixed set of best tracks for the movement of the source.

Of course, as with any effective parallel algorithm designed for high-performance embedded computing (HPEC), the target architecture and the mapping of the algorithm(s) to the target are of key importance. For sensor arrays and other systems where it is desirable to disperse the processing and memory demands of the application across multiple nodes, a distributed architecture can be constructed by networking together multiple digital signal processor (DSP) nodes. The distributed architecture developed and employed in this research as the HPEC testbed consists of multiple floating-point DSP development boards connected to one another in a ring topology. Each board includes a single ADSP-21062 Super Harvard ARChitecture (SHARC) processor from Analog Devices as well as additional hardware for links to other nodes, off-chip memory, etc. These links are used to build a ring network of SHARC nodes, and a lightweight network transport and parallel coordination service known as MPI-SHARC was designed, implemented and optimized to support this distributed architecture.

Since Hydra uses an array of sensors to extract track information, by coupling each transducer node with one or several DSPs and networking them together the computational burden can be distributed among the computing nodes. Hence, parallel algorithms that effectively exploit the maximum capacity of all the processors by distributing fragments of the computation on different processors can be developed to diminish execution times. Conversely, by achieving significant parallel speedup, the parallel algorithms can make it possible for the Hydra and other MFT algorithms to operate with an enhanced mathematical model, larger problem size, and higher precision while maintaining a fixed overall execution time required for matching the real-time constraints of the application. Thus, the tradeoff exists with parallel MFT algorithms for distributed, deployable, and autonomous sonar-array systems to compute results faster and/or compute better results.

Four parallel algorithms for Hydra MFT are developed and presented, two based on coarse-grained decompositions and two based on medium-grained decompositions. The coarse-grained parallel algorithms (XY-GPD/TD and Z-GPD/TD) decompose the grid points and selected tracks at the two most dominant of the stages in the Hydra algorithm, those being the initial tracking stage to compute the estimated tracks and the track adjustment stage to correct the computed tracks. By contrast, in both of the medium-grained parallel algorithms (DPD and FD), the decompositions are focused not on stages but instead on the correlation function, a focal point of Hydra computation that is repeatedly invoked in terms of track data points and strongest frequency bins for DPD and FD, respectively.

These four parallel algorithms were implemented in MPI-C code and executed on both the HPEC testbed of networked SHARC processors (using MPI-SHARC) as well as on a general-purpose cluster of networked PCs. A series of experiments was undertaken on both platforms to determine average execution time, computation time, communication time, and memory utilization. Furthermore, speedup and parallel efficiency were also determined using the sequential Hydra algorithm implemented in C code as a baseline. The results of these experiments and an analysis of the results will be featured in the presentation.

In general, the coarse-grained parallel algorithms are observed to perform better than the medium-grained methods. A significant advantage of the coarse-grained algorithms is their relative independence from the network performance, making them suitable for networks with only modest data rates and average latencies. However, in the case of XY-GPD/TD, workload distribution and thus overall efficiency are heavily dependent upon the data provided by the transducers, and thus the performance variance can be large for different input datasets. Moreover, in the case of Z-GPD/TD, constraints must be enforced to achieve a reasonable amount of load balancing, such as a requirement that the number of best tracks and depth grid points must be a multiple of the number of processors.

By contrast, with an adequate problem size, the medium-grained algorithms are observed to achieve a higher inherent degree of load balancing with more flexibility for variations in the sizes of the domains of the problem size. However, by their very nature, they require a faster communication network where network latency is low to achieve reasonable performance. Since the DSP array with the MPI-SHARC transport provides this capability, these algorithms perform well in an HPEC environment but poorly on a traditional PC cluster.

Acknowledgements

The support provided by D. Davison of the Office of Naval Research on grant N00014-99-1-0278 is acknowledged and appreciated. We also acknowledge and appreciate the support provided by H. Bucker and J.M. Stevenson of the Space and Naval Warfare Systems Command (SPAWAR) in terms of FORTRAN code and data for the baseline Hydra MFT algorithm, and by S. Neshvad, J. Kohout, and K. Cho at the University of Florida for their preliminary work on our parallel MFT algorithms and the MPI-SHARC service.

References

1. H. Bucker, "Matched-Field Tracking in Shallow Water," *Journal of the Acoustical Society of America*, Vol. 96, No. 6, Dec. 1994, pp. 3809-3811.
2. H. Bucker and P.A. Baxley, "Automatic Matched-field Tracking with Table Lookup," *Journal of the Acoustical Society of America*, Vol. 106, No. 6, Dec. 1999, pp. 3226-3230.