High Performance Computer Acoustic Data Accelerator (HPC-ADA): A New System for Exploring Marine Mammal Acoustics for Big Data Applications

Peter Dugan[1], John Zollweg[1], Herve Glotin[4], Marian Popescu[1], Denise Risch[3], Yann LeCun[2] and Christopher Clark[1]

(1) Bioacoustics Research Program (BRP), Cornell University, Ithaca, NY, USA
(2) The Courant Institute of Mathematical Sciences, New York University, NYC, USA
(3) Northeast Fisheries Science Center, NOAA, Woods Hole Oceanographic Institute, MA, USA
(4) Institut Universitaire de France, CNRS LSIS and USTV, FR

Abstract— **With continuing growth of the world's population and rapid economic development, our need to preserve the natural environment, especially our oceans, is becoming an increasing concern. For the past several decades scientists have been monitoring the oceans using a variety of sensors and tools. Passive acoustic monitoring is one of the primary methods used to investigate the behavior patterns of soniferous marine animals. Analyzing the vast amount of collected data poses an enormous challenge. This paper presents a new system designed for high speed acoustic processing called the High Performance Computer Acoustic Data Accelerator (HPC-ADA). Together with an appropriate software suite, the HPC-ADA is a powerful tool currently being used by the Bioacoustics Research Program (BRP) at the Cornell Lab of Ornithology, Cornell University. This paper provides a high level technical overview of the HPC-ADA system's architecture, software suite, and operation of the HPC-ADA. We also summarize the projects that have successfully used the HPC-ADA system; totaling over one million hours of processed sound to date.**

Keywords - Ocean acoustics, high performance computing, passive acoustic monitoring, big data, data science, biodiversity.

I. Introduction

Nearly every branch of science is experiencing an explosion in the amount of data collected and available for analysis. From *in situ* sensor networks to remote sensing satellites, enormous stores of ocean data are being amassed from a plurality of sources (e.g. see marinexplore.com). This includes acoustic sensors that are the mechanism by which passive acoustic data are acquired.

Passive acoustic monitoring has been used in terrestrial as well as in marine contexts to document changes in animal occurrence and acoustic environments. This resultant information can be used for improved management of endangered species and natural resources. The acoustic modality is particularly important for marine mammals, which live in an environment that is particularly well-suited for sound propagation. They produce sounds for communicating, foraging and navigating. Furthermore, acoustic monitoring methods are robust. They are not subject to visual sighting limitations imposed by weather, daylight or ocean environmental conditions. The main challenge for acoustic monitoring, as with many fields, is in processing and analyzing the vast amounts of collected data. In the past two decades, the bioacoustic sciences have made significant advances in software for collecting and analyzing both archived and real-time data [1-3]. Despite these advances, large amounts of acoustic data remain unprocessed. There are many challenges. First, sound collections are relatively large (i.e. TBs to PBs). The community lacks standards and methods for making the data available for easy access. There is a shortage of funding and human resources for analysis; and no capable systems for automatically processing data are available to the scientific community.

In an attempt to address Big Data concerns in Bioacoustics, we have designed and constructed a high performance computing (HPC) system for processing large stores of acoustic data. This paper describes that system. HPC can be used for a variety of tasks associated with time series data, such as compression and acoustic modeling. Our system has applications that use image processing techniques on spectrogram images to detect and classify vocal patterns.

The background section focuses on system architecture and provides an overview of several key aspects for high performance computing as it pertains to processing large sound archives. We provide a high level understanding for HPC computing as it relates specifically to a system developed at the Bioacoustics Research Program at the Cornell Lab of Ornithology, Cornell University (BRP).

The design section describes the hardware and software system developed at BRP. The hardware, developed for HPC processing, is referred to as the High Performance Computer, Acoustic Data Analyzer (*HPC-ADA*). The software, designed for sonic signal Detection and Localization using Machine learning Algorithms (called *DeLMA)* is capable of running on a variety of hardware, including the *HPC-ADA*, desktop servers and laptops.

The results section summarizes in tabular form projects and algorithms used in terrestrial and marine applications. A simple study for runtime performance is also presented, comparing throughput performance between a desktop computer and the HPC-ADA system.

II. APPROACH

In the past ten years software publishers have developed stable platforms for distributed and parallel processing [4, 5]; specifically those that provide high performance computing (HPC). Advances have allowed programmers to integrate commercial off-the-shelf technologies for speech recognition [6], seismic modeling [7], face recognition [8] and general signal and image processing [5, 6, 9]. HPC applications can be used for a variety of tasks associated with time series data, such as compression and acoustic modeling. In this paper we discuss applications that use image processing techniques on shapes within spectrogram images to detect and classify vocal patterns. We briefly discuss key programmatic requirements for introducing HPC

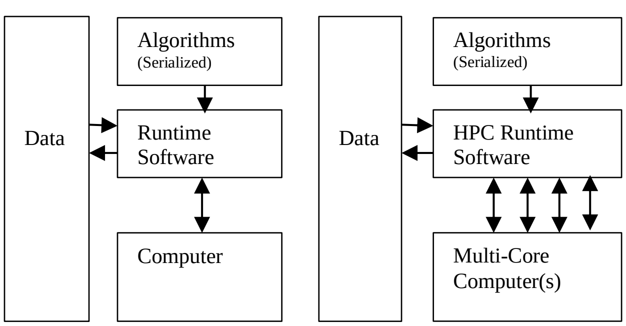

(a) Serial (b) HPC

Figure 1. System for data mining sound archives, components consisting of *Data, Algorithms*. (a) Serial system uses *Runtime Software* and standard *Computer*. (b) HPC System uses single or distributed *Multi-computer(s)* and *HPC Runtime Software*.

technology to bioacoustics; namely *affordability, usability, scalability* and *interoperability* as they pertain to designing HPC Software.

A. Serial versus HPC

The basic components for the HPC system are shown in Figure 1. Referring to Figure 1, there are four main components that make up the data mining system for processing sound archives. *Data* represents both input and output formats. Input data is stored in the way of sound archives, or time series, typically audio formats are used to support single or multiple channels. *Algorithms* are routines designed specifically to meet the processing requirements for the specific data mining task. A wide variety of research has been done, outlining the various recipes for building algorithms to extract various forms of animal sound. Typically Algorithms are designed using a serial processing model Figure 1(a), using a simple *computer* along with serial *runtime software*. Figure 1(b) shows the HPC data mining configuration, whereby specialized *HPC Runtime Software* has the capability to distribute workload across multiple processing cores and /or computers. Standard data and algorithms can be used to execute across all the cores of multi-core computer(s).

B. Challenges

Some of the major challenges in HPC reside in the ability to host applications on a flexible architecture, thereby providing stable software and hardware environments and allowing experienced users relative ease for debugging and modifying code. Furthermore, it should be easy to debug code and distribute workload according to the complexity of the job. One of the simplest approaches to accommodate software debugging and design on a distributed system is to use the spmd (single program, multiple data) approach where serial algorithms run as independent processes working on different portions of a single data stream. Thus, the analysis algorithm can be designed and tested under serial conditions, then executed on a distributed platform with less chance of runtime errors. Serial algorithms are also more widely available in the open source community.

Data management is a formidable problem when processing large sound archives. Improvements in collection technology have resulted in ever larger sound datasets, resulting in significantly higher storage costs. At the same time, the ability to accurately and efficiently

Resource	Requirement
Computer Nodes	4 - 64 core
Memory	4 - 192 GB
Disk Space	500GB - 18 TB (NAS, 12 TB usable)
Network	Gigabit
COTS Software	Mathworks Distributed-Parallel (MDCS) Custom HPC-ADA Control Node for sound

Table 1. Computing specifications for initial development of HPC-ADA system.

process these large datasets has lagged behind this growth.

Sound collection in the field is typically done in deployments, or batches. Sound archives are not pristine and have artifacts that result from device malfunction or environmental conditions. Channel synchronization errors, channel dropout, and recording duty cycling are intermittent conditions found in most acoustic datasets.

Furthermore, because the software algorithms are not yet fully reliable, accurate analysis relies on human inspection of sound files and spectrograms. This process is slow and labor-intensive, so only a small percentage of the data can be inspected. The full exploration of existing acoustic datasets, either by human inspection or more accurate algorithms, would allow better integration with metadata sources, such as from visual surveys, tags, biopsy samples or satellites. This would lead to richer understandings of animal ecology, population biology, biodiversity and anthropogenic impacts.

III. HPC DESIGN, HARDWARE AND SOFTWARE

A. *Hardware Design: HPC-ADA*

Initial phases of this work focused on developing cost effective hardware and software which is capable of processing large datasets, and/or can be scaled to larger systems. The initial design was developed using hardware summarized in Table 1. The specifications in Table 1 indicate that a range of computing specifications could be used for processing. Software for the HPC-ADA was developed to accommodate a range of systems, from a single standalone computer, such as a laptop, to a large cloud-type server hosted in a raised floor environment. The HPC hardware system was constructed utilizing a DELL Cloud Server C6220, with remote access to other platforms.

B. *Software Design: DeLMA*

Since bioacoustics has a relatively low funding ceiling, computer systems designed to handle acoustic datasets should be: (1) *affordable*, cost effective; (2) *usable*, accessible for use by skilled operators; (3) *scalable*, efficiently running on small, medium or large computer platforms; and (4) *interoperable*, interfaced to common software tools and data management systems.

Affordability is addressed by allowing HPC software to run on various platforms. *DeLMA* is designed to run on laptops and desktop servers, as well as distributed servers, such as the HPC-ADA machine.

Detection, Classification Algorithm	**Algorithm Type**	**Algorithm Source**
Right whale	1	Custom MATLAB
Right whale	2	Kaggle-2013 Whale Challenge
Right whale	3	
Seismic	4	Custom MATLAB
Brydes whale	5	xBat Data Template
Fin whale	5	xBat Data Template
Sperm whale	4	Custom MATLAB
Sei whale	5	xBat Data Template
Minke, Humpback, Haddock	4	Custom MATLAB
Elephant Rumble Detection	2	Custom MATLAB
Fin, Blue whale detection	4	Current
Right whale detection	6	

Table 2. Serialized data-mining algorithms used on selective BRP projects; Interfaces were built for the HPC *DeLMA* system, allowing algorithms to sun parallel or distributed without changing the algorithm software.

Usability is addressed by accommodating serial algorithms, HPC Software, Figure 1(b), uses lower level routines to distribute processing to parallel or distributed nodes. A host of serial algorithms have been incorporated through various BRP projects, as listed in Table 2. Referring to table 2, algorithms are described by two factors, one the target species and two the algorithm name, which describes type of technology used. Type 1 *(isRAT)* algorithm is associated with well published technique [10-13], type 2 and type 3 algorithm is based on shallow systems approach using histogram of oriented gradients (HOG) and connected region analysis (CRA). Work summarized in [14-16]. Type 4 (*PT*) is based on repeating signal structure, work shown in [17-20]. Type 5 algorithm is based on data-template and matched filtering concepts, (*DTD)* [21, 22] and type 6 *(CNN)* algorithms are deep learning based solutions [12, 23, 24].

Scalability happens in two ways. First, the HPC software is allowed to load balance, or use less than 100% of the entire processing capability for the given computer. Multiple jobs on one computer are considered parallel, and the user may select the maximum number of processors. This parameter is configurable using the UI menu, Figure 2. The software scales to allow multiple computers to be connected in a distributed configuration, these settings require an additional toolbox offered from Mathworks called the Matlab Distributed Computing

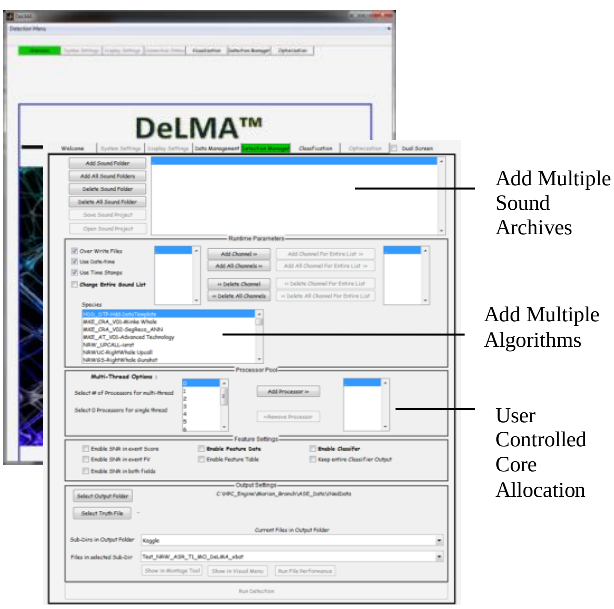

Figure 2. *DeLMA* software application, user interface with calls to the *DeLMA-HPC* routines written in MATLAB.

Server (MDCS), however interfaces are constructed within the software to utilize this capability without maintaining separate builds of *DeLMA*. Therefore *DeLMA* is smart enough to understand the constraints of the hardware system, automatically adjusting to accept parallel or distributed configurations based on processing limitations.

Deployment	**Est. Channel Hours**	**Est. Job Runs**	**Algorithms**
Excellerate	832k	1	Fin (DTD) NRW (isRAT)
GoMex	350k	3	Sperm (PT) Brydes(DTD)
Greenland	7.2k	5	Seismic (PT)
Mass CEC	25k	3	Minke (PT), NARW (isRAT, CRA, HOG)
Gulf of Maine	26.3k	2	Minke (PT) NARW (isRAT, CRA, HOG)
CCB 2008 Spring	21.6k	6	NARW (isRAT, CRA, HOG) Minke (PT) Fin (DTD)
CCB 2009 Spring	21.6k	6	
CCB 2010 Spring	21.6k	6	
CCB 2011 Spring	21.6k	6	
CCB 2012 Spring	21.6k	6	
SBNMS 08 Fall	18.9k	10	NARW (isRAT, CRA, HOG), Minke (PT), Humback (PT), Haddock (PT) Fin (DTD),
SBNMS 08 Winter	18.7k	10	
SBNMS 09 Spring	15.3k	10	
SBNMS 08 Summer	7.5k	10	
VA 2012 Spring	3.1k	2	NARW (isRAT, CRA, HOG)), Minke (PT), Fin (DTD)
VA 2012 Summer	13.2k	2	
VA 2012 Fall	7.2k	2	
VA 2012 Winter	NA	NA	
NAVFAC-01 (32 kHz)	4k	2	Sperm (PT) NARW (isRAT, CRA, HOG) Minke (PT) Fin (DTD)
NAVFAC-01 (32 kHz)	6k	2	

Table 3. Select projects that used HPC-ADA system and DeLMA software. Channel Hours, estimated times the jobs were run and algorithms used is also shown.

For HPC technology, performance and scalability are directly related[1]. Systems need to be fast and flexible, able to re-process large amounts of information, while

[1] MDCS limit is 512 distributed cores.

running custom algorithms for data mining and exploration. Input and output data need to be organized and relatively easy to work with. Therefore faster processing becomes a function of the cost for the hardware, with the maximum number of nodes limited by the parallel or MDCS license.

Interoperability is the fourth requirement consider for designing *DeLMA*. HPC Software can accommodate a variety of output formats, including standard database tables along with selection tables for publicly available tools like *Raven or XBAT*. The key for this process is to wrap the algorithms to deal with generic input and output data formats. The capability to plugin standard (and custom) designed data mining algorithms has not yet been addressed.

The user interface for processing sounds covers three basic configurations: stand-alone mode, console-mode, and client-server mode. Only the console mode has been used for the processing reported here.

C. HPC DeLMA Projects

The *HPC-ADA* hardware and *DeLMA* software have been used to analyze many datasets. Table 3 shows a selection[2] of the projects. Referring to Table 3, estimated channels hours, estimated job runs and algorithms used are summarized for each deployment. There are 19 deployments used in the example. In many situations, deployments were run several times. Not counting redundant runs, the total duration of the acoustic data processed totals over 1 million continuous hours of acoustic recordings [20, 25].

D. Performance

In order to make sufficient comparisons of the runtime performance, three methods need to be compared. Since the development of the HPC hardware and software technology; BRP realizes the fastest data processing with *DeLMA* running on the *HPC-ADA* computer. This is considered *Method One.* Since *DeLMA* is capable of efficiently running serial algorithms on multi-core computers and laptops. BRP often utilizes available multi-core desk top servers for processing smaller jobs. This is referred to as *Method Two*. Before the advent of the HPC, BRP used a series of workstations to process acoustic data. This is referred to as *Method Three*. For Method Three, human operators were required to manually distribute sounds across various desktop computers. The software for Method Three was serial MATLAB routines. This process was very time consuming and inefficient.

[2] Table does not include Artic, Terrestrial, European or Navy Projects.

IV. Results

A. Project Hours

Processing hours between 2011 and 2013 for the BRP are summarized in Table 4. Results represent channel hours, average number of jobs run per deployment and total hours run through the *HPC-ADA*, *DeLMA* system and are computed using Table 3 from the methods section.

B. HPC Performance

For this work, we present two performance scenarios. For both scenarios, two types of computers are used; runtime performance of the HPC-ADA, Intel Xeon E-2670 @ 2.6 GHz versus a Dell desktop workstation, Intel Xeon E5-2620 @ 2.0 GHz.

Scenario one, Table 5, is for current HPC processing, comparing the distributed ADA system to a parallel desktop server, *Method One* versus *Method Two*. This scenario, represents a runtime for the current "as is" status for executing jobs using the *DeLMA* software and the HPC-ADA computer for BRP. Referring to Table 5, runtime efficiency between the HPC-ADA and the desktop server ranges based on the archive size and sample rate by as much as 9:1 and 13:1. For this example, HPC-ADA is only using 48 cores, with a maximum of 128 virtual cores possible. The desktop servers are limited to running 12 cores in parallel; only 4 cores were used in this example.

East Coast Jobs Processed Between 2011-2013	Channel Hours (million hours)	Average Number (Jobs Run / Project)	Total Channel Hours Run (million hours)
19	1.44	~5.0	3.36

Table 4. Summary of East Coast Data, for example in Table 3, processed through the HPC system between 2011 and 2013.

Scenario two, Table 6, compares the fastest HPC configuration, HPC-ADA to a serialized desktop server,

Dataset		Method One (HPC-ADA)		Method Two (Desktop Server)		Run Time Efficiency.
Sample Rate	Total Hours (Size bytes)	Number of Cores	Runtime (HH:MM:SS)	Number of Cores	Runtime (HH:MM:SS)	
16 kHz	5,520 (592 GB)	48	12:46:40	4	162:00:00	**x13**
2 kHz	168 (11 GB)	48	00:29:10	4	04:53:00	**x10**
2 kHz	29,808 (380 GB)	48	03:57:08	4	36:00:00	**x9**

Table 5. Results showing the an example tradeoff between running *DeLMA* on a larger cloud server (*Method One*) or desktop workstation (*Method Two*). An increase efficiency factor ranging from 13:1 to 9:1 for various sample rate and archive sizes.

which is intended to model "pre-HPC" conditions for BRP, *Method One* versus *Method Three*. This scenario represents a runtime comparison of the *old way* of running data to the *new HPC* method. The *old way* consists of using several desktop servers, manually transferring data and the new method compares the HPC-ADA multi core server, whereby all data is served from a centralized disk system, requiring no manual movement; being accessed by 128 cores [3] running distributed processing.

Dataset		Method One (HPC-ADA)		Method Three (pre-HPC)		Run Time Efficiency.
Sample Rate	Total Hours (Size bytes)	Number of Cores	Runtime (HH:MM:SS)	Number of Cores	Runtime (HH:MM:SS)	
2-20 kHz	172,896 (4.38 TB)	128	06:00:30	1	528:00:00	**x88**

Table 6. Results showing an example tradeoff between running *DeLMA* on a larger cloud server (*Method One*) or desktop workstation (*Method Three*). An increase efficiency factor of 88:1 can been seen when the HPC-ADA machine analyzes data in parallel with 128 nodes versus a single node being used on the workstation.

Baseline performance is summarized in Table 6 which compares a workstation to the current HPC-ADA system. The basic measure of performance is an efficiency factor. In a practical sense, this factor would account for the time required to process the data and perform the map-reduce operation, which includes I/O from the storage device. This factor is expected to improve using faster, more distributed network storage. The workstation uses local storage, whereas the HPC-ADA utilized the NAS device outlined in Table 1. Referring to Table 6, runtime efficiency is improved by as much as a factor of 88:1 when comparing the HPC-ADA to the workstation for archives of 4.38 TB.

[3] Currently 64 real, 64 virtual, load balanced for 128 workers.

V. CONCLUSION

BRP has developed a software and hardware computing system for processing archival data based on HPC technology. The heart of this system is a specialized software application written in MATLAB called *DeLMA*. *DeLMA*, uses a series of MATLAB libraries to efficiently distribute serial algorithms across parallel or distributed computer architectures. Designed around *affordability*, *usability*, *scalability* and *interoperability; DeLMA* provides processing capability on laptops, desktop servers and cloud level systems. *DeLMA* software uses standard algorithms and specialized runtime to execute sound processing across parallel or distributed computer architectures; meaning sound archives can execute on laptops, desktop servers. Software allows users to balance jobs based on the number of processor cores, whereby multiple algorithms and sound archives can be staged, addressing Big Data application. *DeLMA* uses HPC libraries to distribute standard, serial algorithms; available through the open source community, making an attractive system for use in Bioacoustics or Machine learning communities. Algorithms are distributed across multiple processing nodes, without requiring parallelizing the routines. Robust operation is also designed within the software, whereby *DeLMA* can handle various problems in the datasets; such as breaks in the sound stream, channel drop out and duty cycle conditions.

The Hardware, *HPC-ADA*, uses distributed cloud computer architecture along with onboard NAS storage and centralized network access. The *HPC-ADA* hardware, which is designed for an enterprise network and Big Data environments, is capable of workload queuing, load balancing and distributed computing, accessing 128 worker threads consisting of real and virtual nodes.

Feasibility for utilizing the system for Big Data applications was tested using 19 deployments which contained a mixture of data formats. Example sets spanned 1.44 million channel hours of acoustic recordings taken from the BRP archive collection, focusing on U.S. east coast region. Fast processing provided the capability to interactively develop and re-run jobs; whereby the example datasets were processed, on average, roughly 5 times for each deployment; resulting in a total of 3.36 million channel hours of data for the 19 project deployments.

Performance tests were also conducted, comparing desktop servers to the *HPC-ADA* hardware. Various configurations were tested for full processor loading as an estimated throughput rate of 88 times faster using a 64 core cloud computer *HPC-ADA* over a standard office grade server, running a single core. Running multiple cores on the desktop server ranges from 13:1 to 9:1 times in comparing a 48 Node *HPC-ADA* machine. However desktop configurations have limitations in accessing Big Data archives, whereby the ADA machine uses a scalable NAS and distributed processing architecture. Despite performance differences in computer architectures, same *DeLMA* software can be used in either configuration, allowing the research program to adjust processing to meet the needs of the cost requirement available for the data science project.

Combining the HPC-ADA, with the *DeLMA* software; creates a formidable research tool capable of processing large sound archives ideal for big data environments. Future papers will provide examples where the HPC tools are ideal for running other high speed audio processing, such as acoustic modeling and noise analysis. The BRP HPC system is a flexible architecture designed for distributing resources based on several factors, including audio channel properties, such as sample rate and number of sensor channels; making the HPC system practical for general purpose as well as specialized acoustic processing.

Acknowledgment

Work made possible by National Oceanic Partnership Program (NOPP), Office of Naval Research (ONR) N000141210585, National Fish and Wildlife Foundation (NFWF) 0309.07.28515. Thanks to BRP research analyst team for data products, and to the rest of the Cornell BRP staff. Special thanks for those individuals who show continued collaboration and support for DCL Big Data Applications for Bioscoutsics, namely Dr. Sofie Van Parijs Northeast Fisheries Science Center, NOAA. New York University, Dr. Joan Bruna, Ross Goroshin.

References

[1] BRP and Cornell University, "Raven Pro", Interactive Sound Analysis Software (Version 1.4) Computer software, Ithaca, NY, Cornell Lab of Ornithology, Available from http://www.birds.cornell.edu/raven.

[2] D.K. Mellinger, '"Ishmael", Oregon State University", http://ishmael.software.informer.com/1.0/, 2010.

[3] D. Gillespie, "PAMGUARD", semiautomated, open source software for real-time acoustic detection and localisation of cetaceans, http://www.pamguard.org/, 2011- 2014..

[4] A. Behboodian, S. Grad-Freilich and G. Martin, "The Mathworks Distributed and Parallel Computing Tools for Signal Processing Applications," *Acoustics, Speech and Signal Processing, 2007. ICASSP 2007. IEEE International Conference on*, vol. 4, pp. IV-1185-IV-1188.

[5] E.S. Manolakos, D.G. Galatopoullos and A.P. Funk, '"Distributed Matlab based signal and image processing using JavaPorts," *Acoustics, Speech, and Signal Processing, 2004. Proceedings. (ICASSP '04). IEEE International Conference on*, vol. 5, pp. V-217-20 vol.5.

[6] Chung-Ching Shen, W. Plishker and S.S. Bhattacharyya, "Design and optimization of a distributed, embedded speech recognition system", *Parallel and Distributed Processing, 2008. IPDPS 2008. IEEE International Symposium on*, pp. 1-8.

[7] T. Danek, "Parallel and distributed seismic wave field modeling with combined Linux clusters and graphics processing units", *Geoscience and Remote Sensing Symposium, 2009 IEEE International, IGARSS 2009*, vol. 4, pp. IV-208-IV-211.

[8] ZheHuang Huang and Xiaodong Shi, "A distributed parallel AdaBoost algorithm for face detection", *Intelligent Computing and Intelligent Systems (ICIS), 2010 IEEE International Conference on*, vol. 1, pp. 147-150.

[9] C.W. Clark and P.J. Dugan, "Cornell Bioacoustics Scientists Develop a High-Performance Computing Platform for Analyzing Big Data", http://www.mathworks.com/company/user_stories/, *Mathworks Central User Stories*, October 1, 2013.

[10] I.R. Urazghildiieva and C.W. Clark, "Acoustic detection of North Atlantic right whale contact calls using spectrogram-based statistics", vol. 122, no. 2, pp. 769-776.

[11] I.R. Urazghildiiev and C.W. Clark, "Detection performances of experienced human operators compared to a likelihood ratio based detector", vol. 122, no. 1, pp. 200-204.

[12] P. Dugan, M. Pourhomayoun, Y. Shiu, M. Popescu, I. Urazghildiiev, X. Halkias, Y. LeCun and C. Clark, "Survey of methods: Comparison of Automated Recognition for Right Whale Contact Calls", *International Workshop on Detection Classification, Localization and Density Estimation (DCLDE)*, St. Andrews, Scotland, 2013.

[13] S.E. Parks, I. Urazghildiiev and C.W. Clark, "Variability in ambient noise levels and call parameters of North Atlantic right whales in three habitat areas," vol. 125, no. 2, pp. 1230-1239.

[14] P. Dugan, Y. LeCun, S. Van Parijs, D. Ponirakis, M. Popescu, M. Pourhomayoun, Y. Shiu, A. Rice and C. Clark, '"HPC and Bioacoustics, Practical Considerations for Detection Classification for Big Data", Key Note Talk, *ICML 2013, Workshop on Machine Learning for Bioacoustics*, Atlanta, Georgia; USA.

[15] P.J. Dugan, A.N. Rice, I.R. Urazghildiiev and C.W. Clark, "North Atlantic Right Whale acoustic signal processing: Part I. comparison of machine learning recognition algorithms," *Applications and Technology Conference (LISAT), 2010 Long Island Systems*, pp. 1-6.

[16] P. Dugan, M. Popescu, D. Risch, J. Zollweg, A. Mikolajczyk and C. Clark, " Exploring Passive Acoustic Data Using High Performance Computing. Case Study for Pulse Train Exploration: Stellwagen Bank National Marine Sanctuary", *International Workshop on Detection Classification, Localization and Density Estimation (DCLDE)*, St. Andrews, Scotland, 2013.

[17] M. Popescu, P. Dugan, M. Pourhomayoun, D. Risch, H. Lewis and C. Clark, '"Bioacoustical Periodic Pulse Train Signal Detection and Classification using Spectrogram Intensity Binarization and Energy Projection," *ICML 2013 Workshop on Machine Learning for Bioacoustics*, vol. 1, 2013, Atlanta, Georgia; USA.

[18] M. Pourhomayoun, P. Dugan, M. Popescu, D. Risch, H. Lewis and C. Clark, '"Classification for Big Dataset of Bioacoustic Signals Based on Human Scoring System and Artificial Neural Network", vol. 1, 2013, Atlanta, Georgia; USA.

[19] M. Popescu, P. Dugan, J. Zollweg, A. Mikolajczyk and C. Clark, "Large-scale Detection and Classification (DC): Four Case Studies
Using an Applied Distributed High Performance Computing (HPC) Platform", *International Workshop on Detection Classification, Localization and Density Estimation (DCLDE)*, poster, St. Andrews, Scotland, 2013

[20] D. Riesch, U. Seibert, P. Dugan, M. Popescu and S. Van Parijs, "Acoustic ecology of minke whales in the Stellwagen Bank National Marine Sanctuary", *Marine Ecology Progress Series*, 2013.

[21] D.K. Mellinger and C.W. Clark, "A method for filtering bioacoustic transients by spectrogram image convolution," *OCEANS '93. Engineering in Harmony with Ocean. Proceedings*, pp. 122-127, vol.3.

[22] H. Figueroa, '"XBAT", (version 6), Bioacoustics Research Program, Cornell University, October 3, 2012.

[23] Yann LeCun, K. Kavukcuoglu and C. Farabet, "Convolutional networks and applications in vision", *Circuits and Systems (ISCAS), Proceedings of 2010 IEEE International Symposium on*, pp. 253-256.

[24] M. Ranzato, Y.-.L. Boureau and Y. LeCun, "Sparse Feature Learning for Deep Belief Networks", Advances in Neural Information Processing (NIPS 2007), Cambridge, MA: MIT Press, vol. 20, pp. 1185- 1192, 2008.

[25] P. Dugan, M. Pourhomayoun, Y. Shiu, R. Paradis, A. Rice and C. Clark, "Using High Performance Computing to Explore Large Complex Bioacoustic Soundscapes: Case Study for Right Whale Acoustics", Procedia Computer Science **20**, 156–162.