Accelerating an Imaging Spectroscopy Algorithm Using GPUs

Matthew Sellitto and David Kaeli
Department of Electrical and Computer Engineering
Northeastern University, Boston, MA USA
matthew.sellitto@gmail.com; kaeli@ece.neu.edu

James Goodman
HySpeed Computing
PO Box 431824, Miami, FL USA
jgoodman@hyspeedcomputing.com

Abstract

Graphics Processing Units (GPUs) have proven to be highly effective at accelerating processing speed for a large range of scientific and general purpose applications. As data needs increase, and more complex data analysis methods are used, the processing requirements for solving scientific problems also correspondingly increase. The massive parallel processing power of GPUs can be harnessed and used alongside multi-core CPUs to address these increased needs and allow acceleration of scientific algorithms to open up new realms of possibilities. As an example, there are many scientific problems that require solving non-linear optimization problems of multiple variables across large arrays of data. These types of problems are classified as highly difficult and require a great deal of computational time to solve using traditional techniques. By utilizing modern local optimization techniques, such as the iterative quasi-Newton algorithms, and combining them with the computational throughput of a CPU-GPU heterogeneous computing platform, we can greatly decrease the processing time required to solve scientific problems of this form.

Introduction and Overview

Remote Sensing

- Remote sensing is utilized across a wide array of disciplines, including resource management, disaster relief planning, environmental assessment, civil infrastructure management, homeland security, and climate change impact analysis.
- Remote sensing and image processing represent ideal problems to be addressed using GPU acceleration techniques, since many algorithms are easily parallelized across the image space.
- The data volume and processing requirements associated with remote sensing are rapidly expanding as a result of the increasing number of satellite and airborne sensors, greater data accessibility, and expanded utilization of data intensive technologies, such as imaging spectroscopy.
- As an example, imaging spectroscopy, i.e. hyperspectral remote sensing, measures numerous, narrow, contiguous portions of the spectrum, which generates significantly larger data volumes than traditional multispectral sensor systems.

Numerical Optimization

- Many scientific problems require solving non-linear optimization problems of multiple variables across large arrays of data.
- Optimization problems are often classified as highly difficult and require a great deal of computational time to solve.
- The massive parallel processing power of GPUs and multiple CPU cores can be combined with modern programming techniques to accelerate complex optimization problems and allow application across large arrays of data.
- We demonstrate these performance gains by accelerating an imaging spectroscopy algorithm for shallow coastal environments.
- Results show that this technology has enormous potential for continued growth in exploiting high performance computing, and provides the foundation for significantly enhanced remote sensing capabilities.

GPU Computation Methods

- We employ a heterogeneous CPU-GPU approach to accelerate a numerical optimization algorithm.
- Work focused on using a quasi-newton projected-BFGS method L-BFGS-B to find the search direction and perform the line search to determine the optimal step size in the solution space of the problem.
- The objective is to balance tradeoffs between accuracy and speed, and ideally achieve the goal of real-time processing.
- We developed a parallel optimization solver, BFGS-B CL, that uses the OpenCL framework to execute the projected-BFGS algorithm over all image elements in parallel by utilizing GPUs and multiple CPU threads.
- The function and gradient evaluations of the function are performed by the GPU while the CPU threads perform the line search calculations.

Projected-BFGS Method

$$s_k = x_{k+1} - x_k \qquad y_k = \nabla f_{k+1} - \nabla f_k;$$

$$B_{k+1} = B_k - \frac{B_k s_k s_k^T B_k}{s_k^T B_k s_k} + \frac{y_k y_k^T}{y_k^T s_k}.$$

$$\phi(\alpha) := f(\mathbf{x}^0 + \alpha \mathbf{d}^0)$$

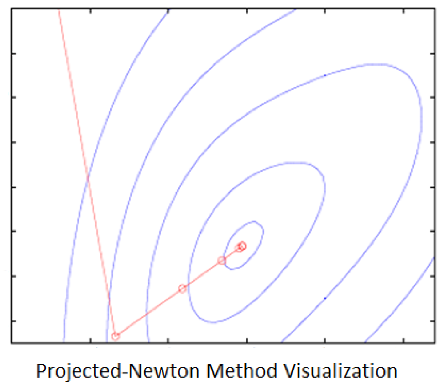
Projected-Newton Method Visualization

Imaging Spectroscopy Algorithm

Visual Overview

- The algorithm derives estimates of water depth, water properties, and bottom albedo from hyperspectral imagery.

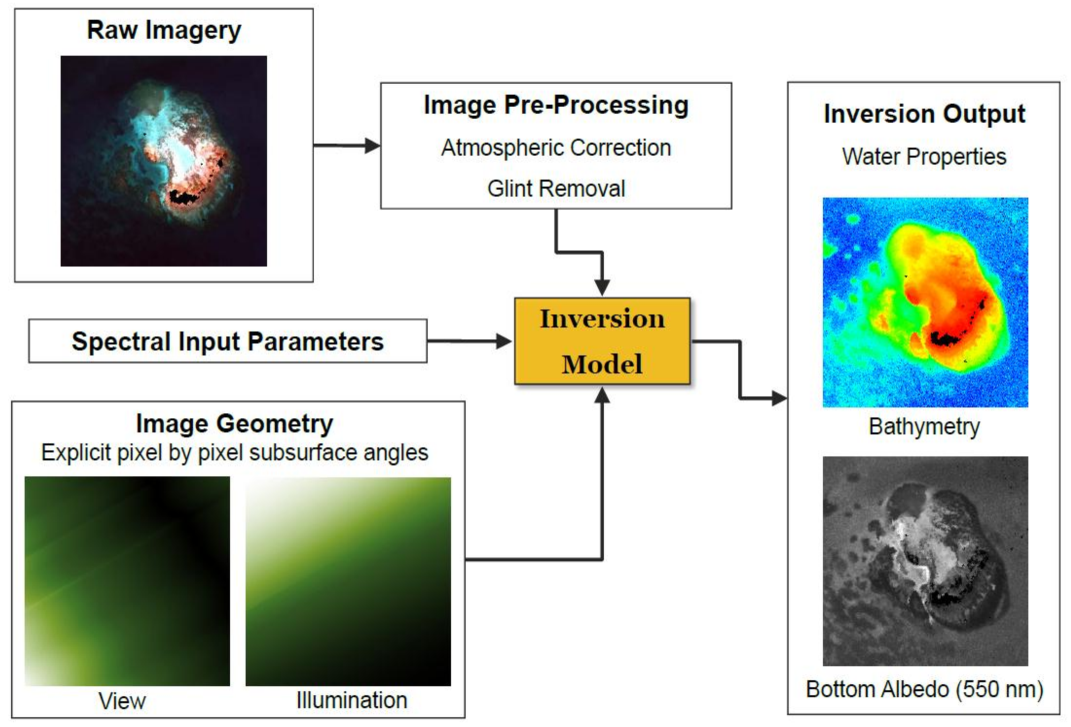

Optimization

- The foundation of the spectroscopy algorithm is an inverse semi-analytical model that relies on a constrained non-linear optimization method.

$R_{rs} = f\,(P, G, BP, B, H)$

$R_{rs} = 0.5\, r_{rs} / (1 - 1.5\, r_{rs})$

$r_{rs} \approx r_{rs}^{dp}\, (1 - \exp[-\kappa H\, (1/\cos(\theta_w) + D_u^C/\cos(\theta)]$ *water column contribution*

$+ (\rho_b B/\pi) \exp[-\kappa H\, (1/\cos(\theta_w) + D_u^B/\cos(\theta)]$ *bottom contribution*

Compute-Time Performance

- The following graphs show the runtime performance of the BFGS-B CL Optimization Solver relative to a serial CPU implementation for three different hyperspectral images. All tests were run on an Nvidia Tesla M2070 GPU with a 6-core Intel Xeon 2.9 GHz CPU. Performance results show as much as a ***>25x speedup*** compared to a serial version.

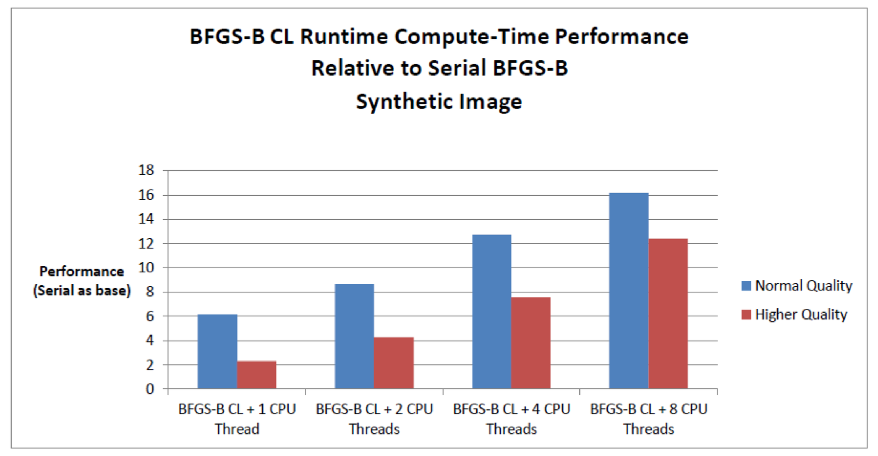

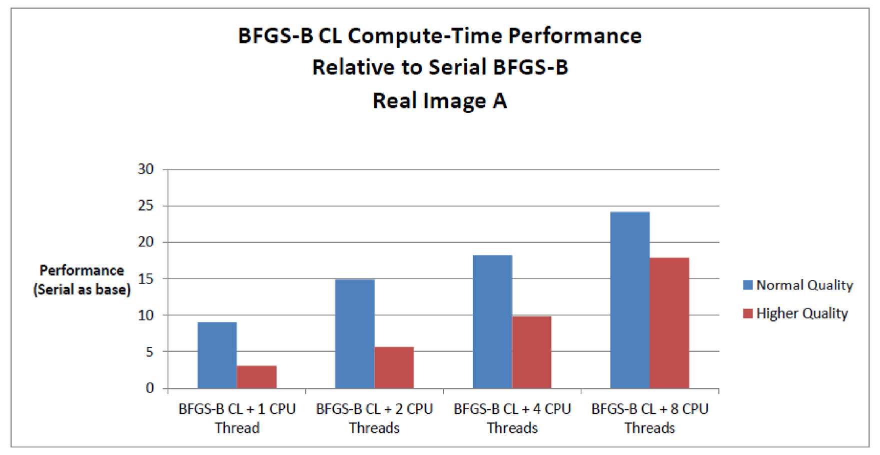

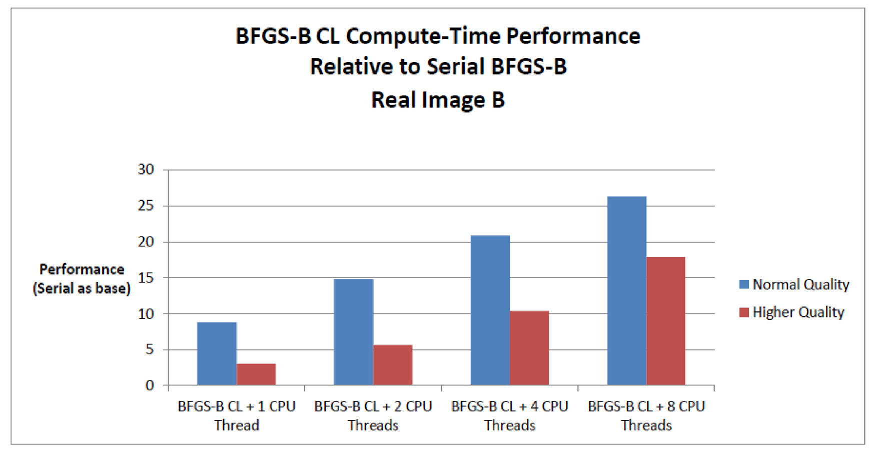

Research to Reality

- GPU technology provides enormous potential for the acceleration of data-intensive applications.
- The BFGS-B CL Solver uses modern optimization techniques to be able to solve problems quickly from a wide variety of different scientific areas.
- GPU technology and the techniques used in this research have potential for application to a vast array of image processing and numerical optimization problems.
- The hyperspectral algorithm for coastal environments demonstrates the advantages and capabilities of a GPU optimization framework.

References

James A. Goodman, Dave Kaeli, and Dana Schaa, 2011, "Accelerating an Imaging Spectroscopy Algorithm for Submerged Marine Environments Using Graphics Processing Units", Journal of Selected Topics in Earth Observation and Remote Sensing.

James A. Goodman and Susan L. Ustin, 2007, "Classification of Benthic Composition in a Coral Reef Environment Using Spectral Unmixing", Journal of Applied Remote Sensing, Vol. 1, 011501.

Accuracy Performance

- The following graphs show the accuracy results (R2 and Standard Error) of the water properties obtained from processing the three hyperspectral images compared to the output of a commercial optimization solver.

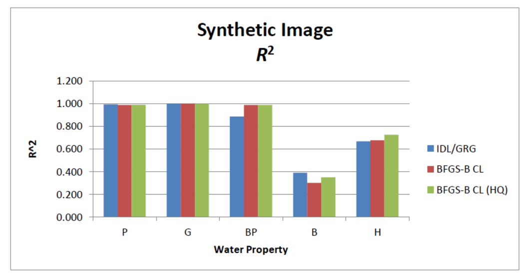

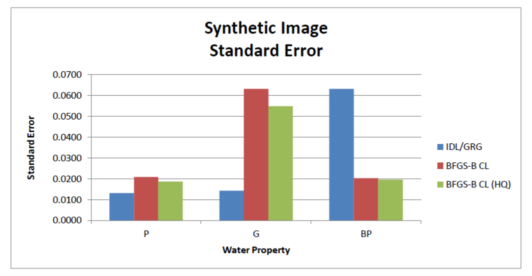

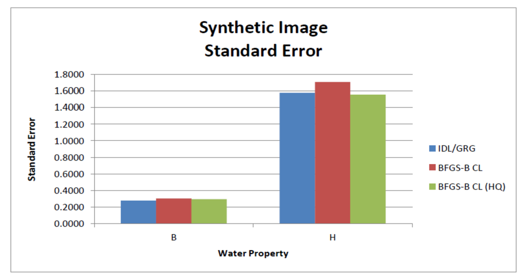

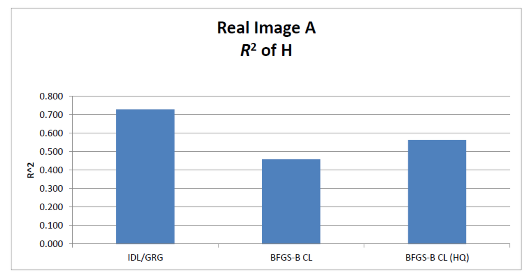

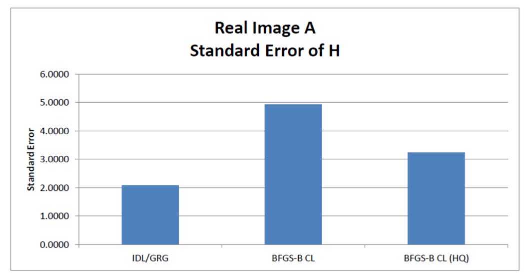

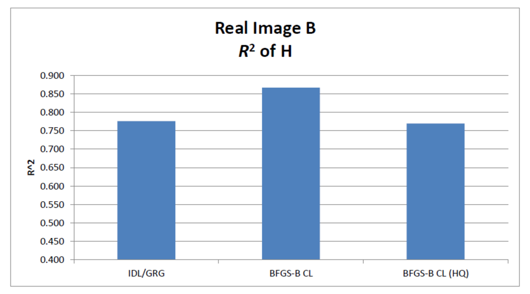

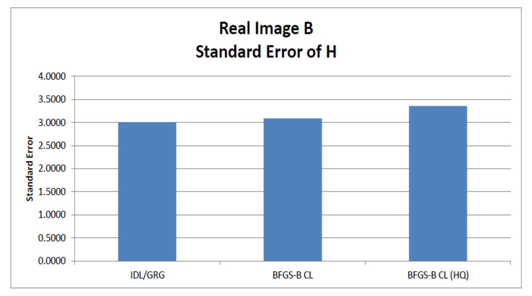

Discussion

- Results from the parallel projected-BFGS solver demonstrate that GPU version of the algorithm implemented on a single GPU system obtains up to ***>25x speedup*** over a comparable serial version on the CPU.
- The current implementation of the projected-BFGS method produces reasonable accuracy for estimation of water properties (P, B, BP, H), but slightly reduced accuracy for absorption coefficient (G) in the synthetic image compared to a commercial solver.
- The BFGS-B CL Solver is significantly more efficient in executing the hyperspectral algorithm. Running on a single GPU and multi-core CPU, the parallel solver outperforms a serial version of the solver by an order of magnitude and does so with comparable accuracy to commercial optimization solvers.
- Future work will focus on pipelining the GPU and CPU code to better leverage computing resources, incorporating an initial global search to identify better initial values and/or limit the possible solution space, and implementing the optimization algorithm to other application areas.

Acknowledgement

This work was supported in part by the Bernard M. Gordon Center for Subsurface Sensing and Imaging Systems, under the Engineering Research Centers Program of the National Science Foundation (Award Number EEC-9986821).