Phase I summary for NASA NNX11CH10P: A Rapid Model Fitting Tool Suite

Michael Galloy, Tech-X Corporation

Objectives

We have achieved our objectives in the Phase I project:

1. *Demonstrate the feasibility of a parameter fitting framework on GPUs.*
2. *Demonstrate performance benefits of using GPUs for parameter fitting problems.*

The Phase I results warrant Phase II continuation of this project.

Tasks

1. *Prototype forward models on GPUs.*

   From conversations with SAGE III team, we focused our efforts on the line-by-line code. CPU optimizations improved performance of the line-by-line code by a factor of 6.4x. Further investigations with using a GPU to accelerate routines offered an additional factor of 10.7 on specific sections of code.

2. *Prototype a GPU finite difference gradient approximation.*

   While some CPU optimizations were done for this tasks, analysis of the profile timings of the routine indicated that the function evaluations were more critical bottlenecks, hence we concentrated on optimizations that effected the entire line-by-line code.

3. *Prototype Levenberg-Marquardt based minimization on a GPU.*

   We used the results from tasks 1 and 2 to prototype the Levenberg-Marquardt algorithm on a GPU. We used routines from the MAGMA library for the dense linear solve.

4. *Analyze performance of GPU accelerated parameter fitting.*

   CPU optimizations took the runtime of the line-by-line code from 72 sec to 11.3 sec, about a 6.4x speed up. GPU calculation of dense linear solve improves performance by a factor of 28, moving the overall speed up to about an 8x speed up. Further investigations with using a GPU to accelerate routines offered an additional factor of 10.7 on specific sections of code.

5. *Write final report and establish NASA contacts.*

   We have submitted a final report as a separate document. We have collaborated with the SAGE III team for work on the line-by-line code.