Statistical Computing on Graphics Cards: Testing a New R Library

Yawa Adonsou, Jennifer Griffiths,
Darius McDaniel, Florencia Tolentino
Mentor: Mary Kathryn Cowles, Ph.D.

The University of Iowa
Iowa Summer Institute in Biostatistics (ISIB)

Summer 2010

Challenge

- In order to analyze large data sets like those in spatial statistics, linear algebra needs to be performed on large matrices.
- Current statistical software, more specifically R, is prohibitively slow in these computations on large matrices.
  - Matrix multiplication on a matrix of size 5000 by 5000 takes approximately 4.5 minutes.
  - 4.5 minutes $\times 1000 = 4500$ minutes $= 75$ hours

R

- A programming language
- An environment for statistical computing and graphics
- Created by Ross Ihaka and Robert Gentleman at the University of Auckland, New Zealand
- Today, developed and maintained by an international team

Why R is Popular

- R is available as free software
- R is open source
- Effective data handling and storage facility
- Has a suite of functions with fast and easy accessibility
- Graphical facilities
- It compiles and runs on a wide variety of UNIX platforms, Windows and Mac Os
- Object-oriented programming
- Packages that make it more specialized

Packages in R

- Collection of R functions, data and compiled code in a well-defined format
- Capabilities that allows specialized statistical techniques
- R users can write and contribute packages which makes R flexible

CPU and GPU

- Central Processing Unit (CPU)
  - Brain of the computer
  - CPUs process information serially
  - Most of calculations take place here
  - Smaller number of cores (one to four)
- Graphical Processing Unit (GPU)
  - Electronic processor dedicated to provide a high-performance, visually rich, interactive 3D experience
  - GPUs process information in parallel
  - Larger number of cores (from 100)
  - Able to process billions of calculations per second

GPU Computing

- The use of a GPU to do general purpose scientific and engineering computing
- The CPU and the GPU are used together in heterogeneous computing model
- Sequential part of the application runs on the CPU and the computationally-intensive part runs on the GPU

Advantages of GPU Computing

- Can be used to accelerate a wide range of science and engineering applications, in many cases offering dramatically increased performance compared to CPUs.
- GPUs are very efficient when it comes to matrix arithmetic and other parallel data operations.
- Analysis results can be delivered much faster.
- A cheap GPU upgrade is sufficient. The need for expensive replacement of current computing hardware is eliminated.

GPU Computing

- Not any type of GPU can allow GPU computing
- GPUs are specialized and cannot replace replace the function of a CPUs, they are only meant to complement CPUs
- GPUs are especially good at parallel computations
- Require skills programmers to map data on GPUs because it is difficult to do so
- The purpose of this present section is to show the role played by the GPU computing in the new package. The speed of the GPU will be the key to improve R's performance on linear algebra computations.

GPU vs. CPU Computing

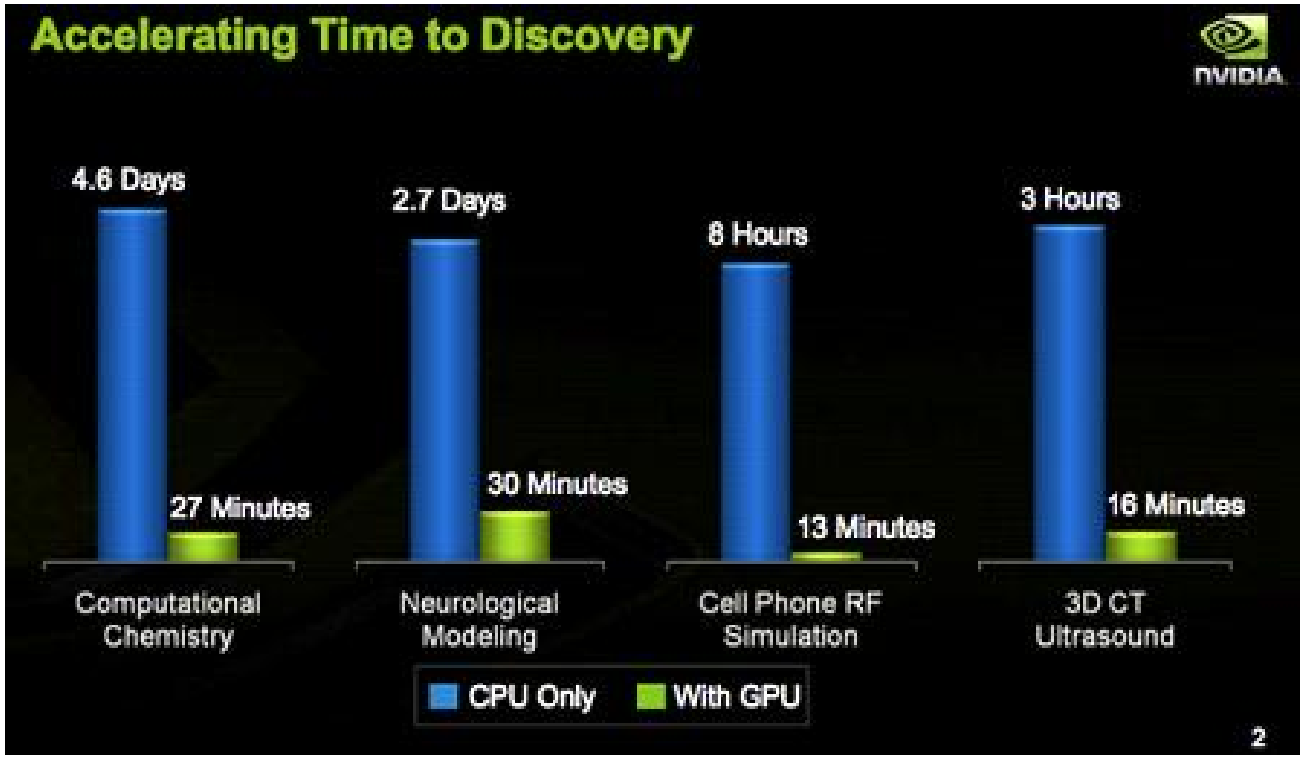

↪ Graphic from http://gizmodo.com/5252545/giz-explains-gpgpu-computing-and-why-itll-melt-your-face-off

MAGMA

- Matrix Algebra on GPU and Multicore Architectures
- Software developed at the University of Tennessee
- Designed for linear algebra algorithms using multicore (CPU) and GPU systems
- Enables linear algebra computations on large data sets to be computed quickly
- However, MAGMA is right now only available in the C and Fortran language

MAGMA in R

- Dr. Brian Smith of the University of Iowa's Biostatistics department has created a new package in R that calls the MAGMA library.
- Dr. Smith's new package is expected to perform linear algebra operations on large data sets, such as spatial statistical data, more quickly than the current linear algebra functions in R.
- Dr. Smith's new package is currently in the beta testing stage.

Beta Testing

- Alpha testing is the first testing done by the programming team themselves.
- Beta versions of software are released to a limited audience outside of the programming team.
- Allows for detection of software failures so that defects may be uncovered and corrected.

Beta Testing a New R Package

- 3 matrix operations:
  - Matrix multiplication
  - Cholesky decomposition
  - Matrix inversion
- Looked at the computing time of the matrix operations on matrix dimensions ranging from 100 by 100 to 5000 by 5000 comparing the CPU time and GPU time.
- Made sure that Dr. Smith's package that utilizes the MAGMA library gives the same results as the existing linear algebra functions in R.
- Performed testing on two separate computers.

Beta Testing Definitions

- The existing R functions use CPU
- Dr. Smith's new R package uses GPU
- Speedup = CPU time $\div$ GPU time

Beta Testing Results

		Matrix Multiplication			
Dim	Seed	CPU	GPU	Speedup	Match
100	15	.002	.002	.800	1.000
200	15	.009	.005	1.785	1.000
300	15	.029	.011	2.676	1.000
400	15	.104	.022	4.698	1.000
500	15	.208	.030	6.911	1.000
600	15	.386	.037	10.424	1.000
700	15	.624	.052	12.046	1.000
800	15	.931	.067	13.837	1.000
900	15	2.583	.165	16.390	1.000
1000	15	5.268	.285	18.394	1.000
1500	15	6.394	.293	22.222	1.000
2000	15	14.079	.559	25.381	1.000
2500	15	31.141	1.092	28.214	1.000
3000	15	47.043	1.583	29.727	1.000
3500	15	87.263	2.294	37.955	1.000
4000	15	113.081	3.217	35.192	1.000
4500	15	183.608	4.397	41.864	1.000
5000	15	254.308	5.693	44.665	1.000

Beta Testing Results

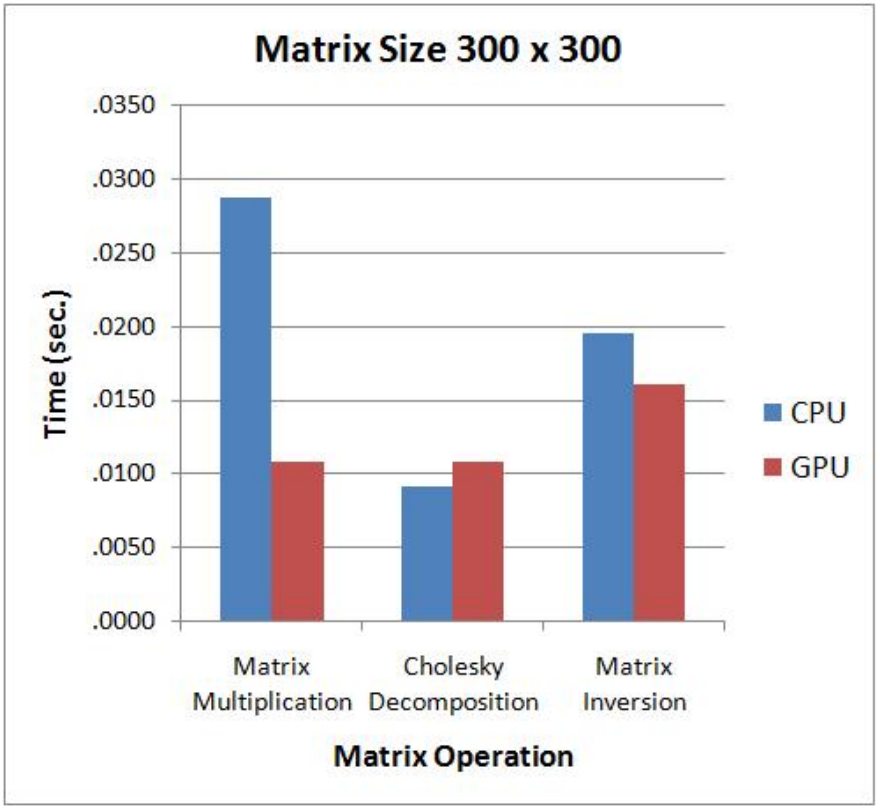

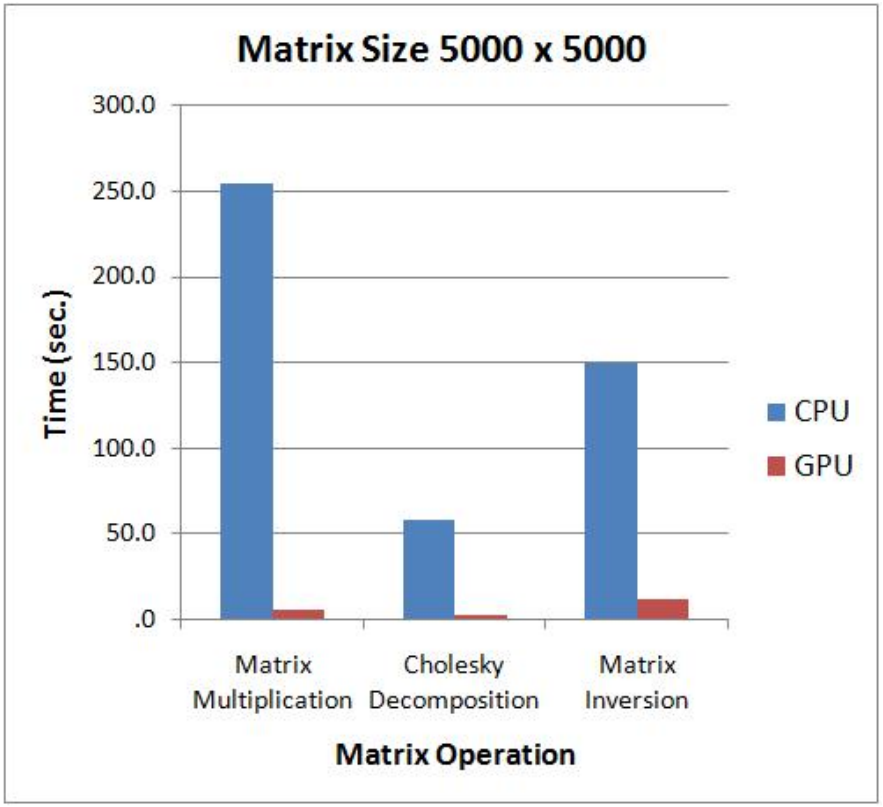

Beta Testing Results

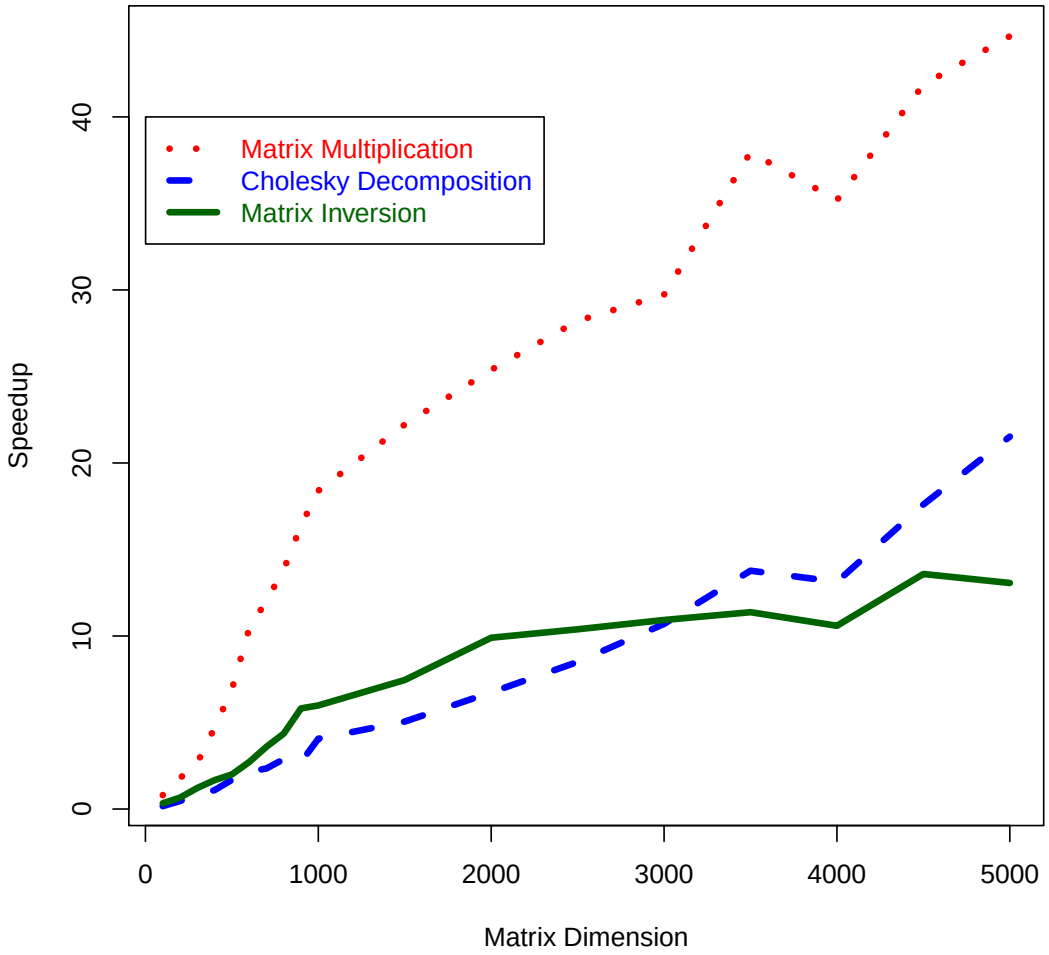

Beta Testing Results Interpretation

- Our beta testing shows that Dr. Smith's new R package is effectively utilizing the MAGMA library.
- This causes the time it takes to compute linear algebra operations on large matrices to decrease when compared to the existing R functions.
- Further testing needs to be done, especially on the matrix inversion, to be certain that the new package has no errors.

Conclusion

- Compared to existing linear algebra functions in R that use CPUs, Dr. Smith's package speeds-up computations on large matrices because it utilized GPUs.
- The new package is much more accessible because users don't need to know how to program in C.
- After further testing, Dr. Smith's package will effectively and efficiently perform linear algebra computations on large matrices in R.

References

- GPGPU <http://gpgpu.org/>.
- GPU Computing <http://www.uweb.ucsb.edu/yichuwang/ecv/paper/gpu_computing.pdf>.
- Magma <http://icl.cs.utk.edu/magma/>.
- NCSA Answers Questions about GPU Computing <http://www.ncsa.illinois.edu/News/Stories/GPUcomputing/>.
- Quick R <http://www.statmethods. net/interface/packages.html>.
- R <http://www.r-project.org/>.
- Software testing <http://en.wikipedia.org/wiki/Software_testing>.