High-Performance Social Dynamics Simulations on Graphics Processors

Jason Rupert Roop
North Carolina Central University, Durham, North Carolina
Research Alliance in Math and Science (RAMS)

Mentor: Dr. Kalyan S. Perumalla
Computational Sciences and Engineering Division
Oak Ridge National Laboratory

http://www.csm.ornl.gov/Internships/rams_06/abstracts/j_roop.pdf

Abstract

Graphics processors, although intended for accelerating visual graphics, have become more general-purpose in nature, and are now termed General Purpose Graphics Processing Units (GPGPUs). Several applications have been able to perform better on GPGPUs than on conventional computers. The high-performance potential of the GPGPUs is exploited to simulate social models of large populations at high-fidelity, down to the individual level. Intuitively, if each individual is mapped to a pixel, an image (called "texture" in graphics) becomes a collection of people. Manipulating a texture corresponds to simulating the evolution of individuals' behaviors. Since GPGPUs are highly optimized for such image operations, simulations can be performed very rapidly. This approach is used on a social model, by Schryver and Perumalla, which tracks the spread of mood of a large population and the influence of external local/global media news on propagation of the mood. Results are presented on populations with up to 1 million individuals, using NVIDIA graphics processors and OpenGL/Cg software environments.

Problem

- Population dynamics simulations find many applications
  - Defense, Home Land Security, Business Economics
- High resolution of models needed for accuracy
  - At the level of each individual's behavior
- Computational needs are very high for high resolution
  - Conventional CPUs are ill-suited

Mood Diffusion Model

- Individual moods influenced by neighbours and events/media Homeostatic effects try to pull mood back to "normal"

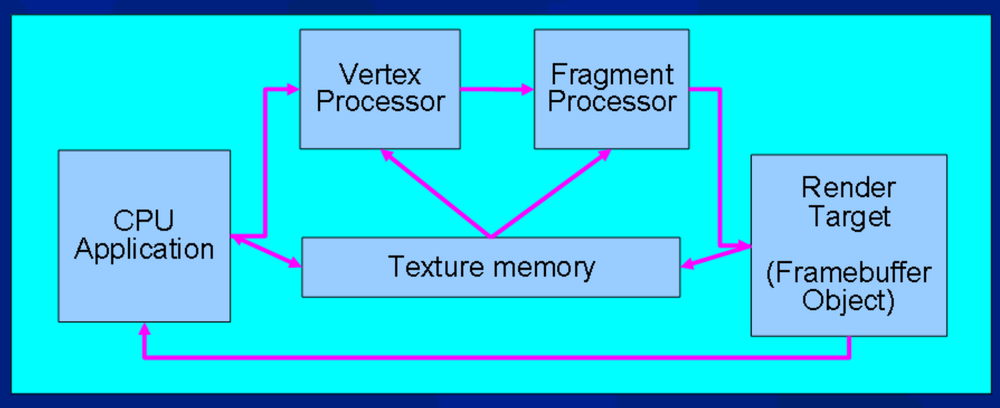

Figure 1. GPGPU pipeline model – Texture data streamed through internal processors to target textures.

GPGPU Computation Paradigm

Array => Texture
Kernel => Fragment Program
Feedback => Copy To Texture
Data Stream => Draw Graphic

Implementation Methods

Implemented application runtime in C
Used OpenGL for graphics operations
Wrote fragment program in NVIDIA Cg
Made visualization with retrieved data

Approach

- Utilize emerging, novel computation platforms
  - Graphical processors for non-graphics computation
- Realize capability for high-resolution modelling
  - For nation-scale populations (>10 million individuals)
- Demonstrate with mood tracking application
  - Used to model effects of global and local events, news, etc. on people in nations around the world

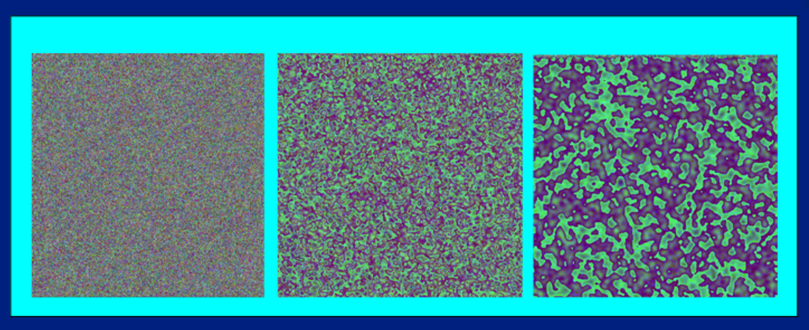

Figure 2. Screen shots of 1024 × 1024 texture evolution over three 20-step time intervals.

Achieved 1500-fold improvement in simulation speed! Intel® Pentium® D Processor 930 compared to NVIDIA 7950 GX2

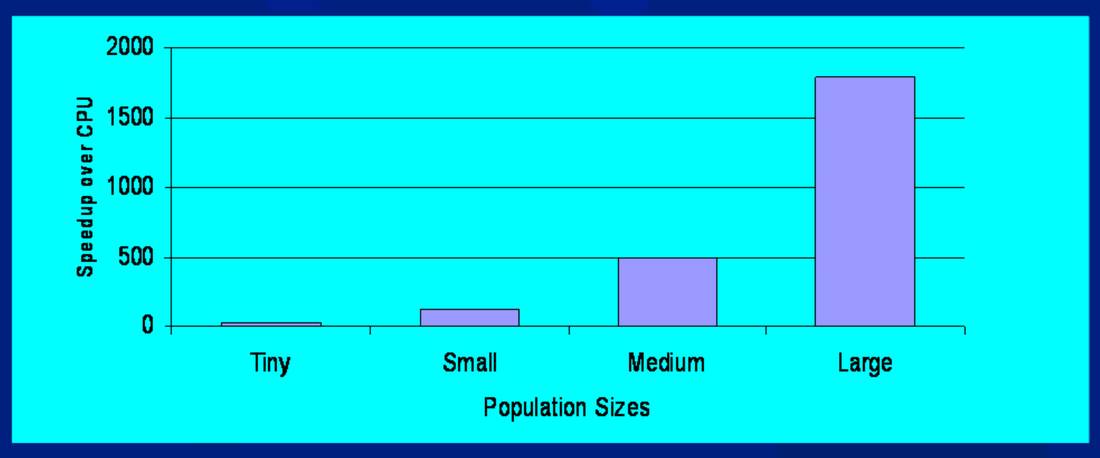

Figure 3. Speedup of GPGPU implementation over identical CPU calculation for 1000 iterations.

Conclusion

- GPGPU significantly faster for large population sizes
- Fast computation for less cost than comparable processors
- Enables real-time simulations for analysts requiring immediate results on desktop or laptop computers

The Research Alliance in Math and Science program is sponsored by the Mathematical, Information, and Computational Sciences Division, Office of Advanced Scientific Computing Research, U.S. Department of Energy. The work was performed at the Oak Ridge National Laboratory, which is managed by UT-Battelle, LLC under Contract No. De-AC05-00OR22725. This work has been authored by a contractor of the U.S. Government, accordingly, the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes.

Special thanks to my mentor, Dr. Kalyan Perumalla, for helping me with this research. An extreme thank you to Debbie McCoy and Judy Burns for their invaluable contributions the RAMS program this summer. Your time and effort were very much appreciated.

Office of Science U.S. DEPARTMENT OF ENERGY

UT-BATTELLE

OAK RIDGE NATIONAL LABORATORY
U.S. DEPARTMENT OF ENERGY