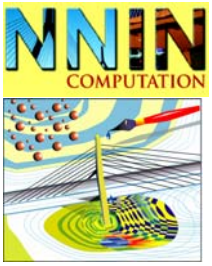

Nano by Numbers, M. Stopa

National Nanotechnology Infrastructure Network Computation (NNIN/C) Project Hands-on Workshop

# Programming the GPU: Introduction to CUDA

Harvard University Center for Nanoscale Systems

**Dates: August 12-14, 2009**

**Location: Harvard University**

Graphical Processing Units (or GPUs) are dedicated graphics devices whose ancestors date to the 1980s and the graphical chips at the heart of Atari Game consoles and others. GPUs are optimized to perform linear algebra functions, such as matrix multiplication, with a high degree of parallelism. Recently it has been discovered that the highly parallel functionality of GPUs makes them potentially useful for accelerating scientific calculation and simulation. NVIDIA Corporation, a maker of GPUs and the developer of the CUDA language for GPU programming, has recently named Harvard University a CUDA Center of Excellence. The principal research computer for the Center is the "Orgoglio" cluster[†]; a joint resource of NNIN/C and the NSF-funded Cyber Discovery Initiative at Harvard.


**Workshop**

NNIN/C will hold a hands-on tutorial workshop on CUDA programming for the Graphical Processing Unit at Harvard University from August 12 to 14, 2009. This workshop will be limited to approximately 20-30 researchers/students and will focus on the basic concepts and examples from the physical sciences of parallel programming for the GPU. The workshop will include lectures each morning from accomplished computational scientists who have developed algorithms for code on GPUs, followed by afternoon, hands-on sessions in the computational lab at the Harvard Science Center.

**Lecturers**

Confirmed lecturers include: Richard Edgar, IIC, Harvard; Alan Aspuru-Guzik, Chemistry, Harvard; Mike Clark, Physics, BU; Hanspeter Pfister, SEAS, Harvard; Mark Watson, Chemistry, Harvard; Michael Stopa, NNIN/C.

**Participation**

For registration, see: http://cns.fas.harvard.edu/nnin/cuda.php. The registration fee is $250 and includes a one year account on the Orgoglio cluster (NNIN nodes). Deadline for applications is August 7, 2009. Additional questions: stopa@cns.fas.harvard.edu. A block of dorm rooms at Suffolk University have been reserved for thrifty out of town participants (see registration page).

[†] The Orgoglio cluster consists of 24 Nodes, each with:

- Single quad-core Xeon 'Harpertown' processors at 3 GHz
- 16 GB of EEC DDR2 800 RAM
- Two Tesla C1060 GPUs (each with 4GB of RAM)

(total of 24 nodes/motherboards, 96 cores, 192 GB RAM, 48 S1070 cards).
QLogic 24-Port 9024 DDR InfiniBand networking between the nodes.