

This is a 3-day in-house training course for Quants, Risk Professionals and Programmers in Quantitative Finance. It is a hands-on class in which students will write GPU enhanced programs throughout.

Course at-a-glance:

### **Day 1**

9:00am	A brief history of GPGPU programming / Introduction to GPGPU platforms: CUDA, OpenCL, DirectX Compute and commercial offerings / Why choose CUDA?
9:45am	Overview of the CUDA architecture / The CUDA tool chain / Building and debugging CUDA programs / Profiling CUDA programs
10:30am	Break
11:00am	First steps in CUDA programming / Kernels and threads
11:45am	Control flow in CUDA / Numerics and floating point
12:30pm	Lunch
1:30pm	GPU hardware architecture in-depth / Stream processors / Memory from a hardware perspective
2:15pm	Thread / Memory hierarchies in CUDA
3:00pm	Break
3:30pm	Structuring CUDA programs to make best use of available hardware
4:15pm	Review and hands-on programming
5:30pm	Class ends

### **Day 2**

9:00am	Parallel programming patterns
9:45am	CUDA libraries (using and building)
10:30am	Break
11:00am	Advanced memory management / Large data sets
11:45am	Profiling & optimizing CUDA programs
12:30pm	Lunch
1:30pm	More advanced patterns for CUDA programs
2:15pm	More advanced host/device interaction
3:00pm	Break
3:30pm	Multi-GPU card and cluster programming with CUDA
4:15pm	Review and hands-on programming
5:30pm	Class ends

### **Day 3**

9:00am	Problems in finance and risk that are a good fit for the GPU
9:45am	Interfacing GPU programs with Excel
10:30am	Break
11:00am	Case study – Derivatives pricing
11:45am	Case study – Monte-Carlo simulation
12:30pm	Lunch
1:30pm	Case study - Risk
2:15pm	Case study – Tick data
3:00pm	Break
3:30pm	Case study – Real-time event processing/trading
4:15pm	Review and hands-on programming
5:30pm	Class ends