

## PRESENTATIONS & TUTORIALS

### P1: Which GPU Programming Model is Right for Me?

In this dynamic talk we compare and contrast CUDA™, OpenCL™, OpenACC and C++ AMP, examining hardware support, development tools and performance results. Drawing from our experience developing commercial applications, the presentation will outline key considerations to help you determine which model best meets your requirements.

### P2: Optimize your CUDA Code - 3 Essential Performance Tips

New to CUDA programming? Accelerate your applications with 3 must-know tips for optimizing your CUDA code. We will discuss optimizing data transfers between the host and the device, strategies for reducing global memory bandwidth and understanding the performance implications of branches in your code.

### P3: Evaluating Hardware Acceleration for your Computational Workload

Science can always make use of faster computers, but what hardware technology is appropriate? Will hardware accelerators speed up your applications? How will this affect software design and maintenance? What are the factors to consider when calculating the price/performance ratio? This talk provides an overview of the latest in hardware advancements and a discussion on the suitability of hardware acceleration for your workflow.

### P4: Programmability & Performance of the Intel® Xeon Phi™ Coprocessor

Reverse Time Migration (RTM) is the state-of-the-art in full wave seismic depth imaging and it requires significant computational power. The Intel Xeon Phi Coprocessor introduces an opportunity to reduce overall hardware and power requirements. We will present Acceleware's experience implementing the core RTM algorithm on the Intel Xeon Phi Coprocessor, discussing programmability and key performance results.

### P5: An Introduction to OpenCL for Altera FPGAs

A 7 slide introduction to using OpenCL to target Altera FPGAs! This high-level overview will cover the system configuration, compiler tool and it's throughput analysis. We will use sample kernel code to illustrate how the compiler converts OpenCL C code into a pipeline parallel FPGA circuit. Key optimization techniques will be highlighted.

### T1: C++ AMP: An introduction to heterogeneous programming with C++

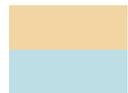
This half day tutorial is designed for programmers who are looking to develop skills in writing and optimizing applications using C++ AMP. Participants will be provided with an introduction to the programming model, the tools, and the knowledge needed to accelerate data-parallel algorithms by taking advantage of hardware such as GPUs.

Full schedule on reverse side

## PRESENTATION SCHEDULE

	Sun, Nov 11	Mon, Nov 12	Tue, Nov 13	Wed, Nov 14	Thu, Nov 15
10:30am			P1	P4	P4
11:00am			P5		
11:30am				P1	P1
12:30pm			P2	P3	P2
1:00pm	T1 1:30 - 5:30pm				
2:00pm			P5		
2:30pm				P2	
3:30pm			P3	P4	
4:30pm			P5	P1	
7:30pm		P1			
8:00pm		P2			

### Locations



Acceleware Booth #2233

Altera Booth #430



Intel Booth #2601

Conference Center, Room 251-E

Visit the Acceleware booth for your chance to win a new **iPad** mini!

