



The Nineteenth International Conference on

# Parallel Architectures and Compilation Techniques (PACT 2010)

Vienna, Austria, September 11-15, 2010

## WORKSHOP AND TUTORIAL PROGRAM



### Workshops & Tutorials at PACT 2010

Workshops and tutorials take place on Saturday, September 11 and Sunday, September 12.

#### Program of Saturday, September 11

On Saturday we have a tutorial and two workshops in parallel and (optionally) a dinner in the evening:

##### Tutorial 1:

### Programming Ct – Porting Applications to Future Multicore

*Michael Klemm and Peter Hinsbeeck (Intel)*

**Saturday, September 11, Sitzungssaal, 8:30 am – 5:00 pm**

Intel® Ct Technology supports a high-level, generalized and portable programming model for data-parallel programming. It simplifies the efficient parallelization of computations over large data sets. Programmers do not need to focus on the implementation details of their data-parallel program, but instead can express a program's algorithms in terms of operations on collections of data. Ct's deterministic semantics avoid race conditions and deadlocks by design, improving reliability and maintainability, and can be used for both rapid prototyping and production-stable codes.

Ct manages the complexity of mapping the high-level description of the program's operations onto an efficient implementation by employing JIT compilation techniques. Its internal JIT compiler dynamically optimizes a program to whatever hardware is used for execution, automatically emitting vectorized and multi-threaded code appropriate for that hardware's microarchitecture. With Ct's JIT compiler it becomes possible to efficiently execute the program on multiple computing platforms (e.g. Intel® SSE, Intel® AVX) without recompiling the application. Ct's JIT compiler is also the key to support upcoming execution environments without the need to recompile a program: updating the Ct library alone, without recompilation of the application, will suffice to enable future platforms.

In this tutorial, we introduce to the participants the programming model and the execution environment of Intel® Ct Technology. We provide an in-depth guide to the basic building blocks of the Ct language: scalar types, dense and sparse vector data types, vector operations, elemental functions, and control flow. We present how Ct fits into an application's control flow and can be used to express different levels of abstraction. Based on real-world scientific codes and other examples, we then show how to construct data-parallel algorithms from these basic building blocks. We demonstrate how to smoothly move an existing sequential code base to a parallel code base. In addition, we illustrate how to make use of external libraries such as the Intel® Math Kernel Library. We close the tutorial with a live demonstration of

performance and scalability analysis as well as performance optimization of Ct applications.

##### Topics Covered:

The topics covered will include the following items:

- Introduction to Ct and its goals
- Getting started with Ct
- The Ct Language
  - Data types
  - Operators
  - Control flow constructs
- Example programs
- Ct Dynamic Engine and JIT compiler
- Porting applications to Ct
- Utilizing external libraries
- VTune performance analysis and performance tuning (demo)

##### Workshop 1:

### GPUs and Scientific Applications (GPUScA)

*Eduard Mehofer, Markus Schordan, Dan Quinlan and Beniamino Di Martino*

**Clubraum, Saturday, September 11**

8:45	<b>Opening &amp; Welcome</b>
9:00	
9:00	<b>Keynote Address: Vivek Sarkar</b> Towards a Portable Execution Model for Extreme Scale
10:00	<b>Multicore Systems</b>
10:30	<b>Applications with strong algorithmic aspects</b>  Improving the GPU-based Collision Check Procedure for Distributed Crowd Simulations <i>Guillermo Viguera, Juan M. Orduña, Miguel Lozano, Jose M. Cecilia, and Jose M. Garcia</i>  Fast GPU perspective grid construction and triangle tracing for Exhaustive Ray Tracing of Highly Coherent Rays <i>Lancelot Perrotte and Guillaume Saupin</i>  Solving Planted Motif Problem on GPU <i>Naga Shailaja Dasari, Ranjan Desh, and Zubair M</i>
12:00	<b>Applications with strong domain aspects</b>  Scalability of Color-Based Segmentation of Football Players over GPUs <i>Miguel Angel Montañes, Enrique F. Torres, Jesus Martinez, and J. Elias Herrero</i>  Fluid Simulation With CUDA Using the Lattice Boltzmann Method <i>Andreas Monitzer</i>  A Framework for GPU Accelerated Deformable Object Modeling <i>Aria Shahingohar and Roy Eagleson</i>
15:00	<b>Parallel programming technology</b>  ViennaCL - A High Level Linear Algebra Library for GPUs and Multi-Core CPUs <i>Karl Rupp, Florian Rudolf, and Josef Weinbub</i>  Dynamic Work Scheduling for GPU Systems <i>Miguel Angel Lastras-Montañó, Maged M. Michael, and J. Alan Bivens</i>
16:30	
16:30	<b>Closing Remarks</b>
16:45	

## Workshop 2: Parallel Architectures and Bioinspired Algorithms (WPABA)

Jose L. Risco-Martín, Francisco Fernández and Juan Lanchares

Museumszimmer, Saturday, September 11

8:30	Welcome and introduction
8:40	
8:40	A Parallel Memetic Algorithm for Workload Distribution in Dynamic Multi-Agents Systems
9:00	<i>David Millá and J. Ignacio Hidalgo</i>
9:00	Communication-focussed approach for real-time neural simulation
9:25	<i>Paul Fox and Simon Moore</i>
9:30	Effective Mutation Operator for Nurse Scheduling by Cooperative GA and Its Parallel Processing
9:55	<i>Makoto Ohki</i>
10:30	Hybridizing Memetic Algorithms and Particle Filters for Visual Tracking on GPU
10:55	<i>Raul Cabido, Antonio Sanz and Juan José Pantrigo Fernández</i>
11:00	P System Simulations on Massively Parallel Architectures
11:25	<i>José María Cecilia, José Manuel García, Ginés D. Guerrero, Miguel A. Martínez del Amor, Mario J. Pérez-Jiménez and Manuel Ujaldon</i>
11:30	GPU-Accelerated Genetic Algorithms
11:55	<i>Rajvi Shah, P J Narayanan and Kishore Kothapalli</i>

## Workshop Dinner

7:00	For a casual dinner on Saturday we have booked a table at the <b>Wirtshaus Biergarten Zattl</b> (Freyung 6, 1010 Vienna). Workshop and tutorial participants are welcome to join. Note, the dinner is not included in the fees. You have to pay for your own consumption.
open end	

## Program of Sunday, September 12

On Sunday we have a workshop in the morning and a tutorial in the afternoon:

### Workshop 3: Programming Models for Emerging Architectures (PMEA)

Xavier Martorell, Rosa M. Badia, Marc González and Alejandro Duran

Sitzungssaal, Sunday, September 12

8:30	<b>Keynote: Calin Cascaval</b>
9:30	Power Programming
9:30	Mapping a Class of Applications in Heterogeneous Multithreading Architectures
10:00	<i>J. Lucena, O. Plata, N. Guil</i>
10:30	Enjing - a JIT Backend for CUDA Devices
	<i>M. Bentsen, B. Vinter</i>
	Combining Processor Virtualization and Component-Based Engineering in C for Many-Core Heterogeneous Embedded MPSoCs
	<i>E. Rohou, A. Ornstein, A. Özcan, M. Cornero</i>
	QoS-ocMPI: QoS-aware on-chip Message Passing Library for NoC-based Many-Core MPSoCs
	<i>J. Joven, F. Angiolini, D. Castells-Rufas, G. De Micheli, J. Carrabina-Bordoll</i>
12:00	

### Tutorial 3: Automatic Parallelization Techniques and the Cetus Source-to-Source Compiler Infrastructure

Rudi Eigenmann and Sam Midkiff (Purdue University)

Sitzungssaal, Sunday, September 12, 1:30 pm – 5:00 pm

The increased importance of parallelism has made parallelizing compiler technology, and easy-to-learn and use compiler infrastructures implementing this technology, increasingly important to researchers, developers and students in the fields of computer architecture, compilers and high performance applications. This tutorial will cover basic parallelizing compiler technology, including dependence analysis, dependence breaking transformations, optimizing transformations, and limits on compiler technology. We will then describe the Cetus source-to-source restructuring compiler infrastructure for C programs which is already used by a substantial number of research projects around the world. Cetus is a freely available, open source community compiler developed with support from the National Science Foundation. Its main distinction from related infrastructure efforts is its focus on high-level source-to-source translation for C programs and abstract internal representation. These features have already proven to enable highly efficient design and implementation of new compilation techniques. The tutorial aims to reach a wider audience and provide guidance for the use of the resource and its advanced optimization techniques. These techniques include new symbolic analysis methods, such as range analysis, automatic parallelization for multicores, and optimizations for heterogeneous multicores.

#### Topics Covered:

The first part of the tutorial will consist of an introduction to parallelizing compiler technology. Topics include:

- Dependence analysis and its relation to parallelization;
- Dependence eliminating transformations (i.e. privatization and expansion, forward substitution)
- Program optimizations, including vectorization and parallelization, parallelization of fully and partially parallel loops;
- Issues raised by optimizing explicitly parallel programs;
- Limits on compiler technology, and techniques such as cloning and profiling to overcome these.

The second part of the tutorial introduces Cetus, a source-to-source restructuring compiler infrastructure for C programs. Cetus is a community resource developed in support by the National Science Foundation. The infrastructure is available at [cetus.ecn.purdue.edu](http://cetus.ecn.purdue.edu). The tutorial will cover:

- Introduction to Cetus' capabilities;
- Internal abstract program representation;
- Optimization and analysis passes currently available in Cetus, including symbolic analysis methods, such as range analysis, automatic parallelization for multicores, and optimizations for heterogeneous multicores;
- Roadmap of ongoing and future enhancements.

