Bonjour,

J'ai le plaisir de vous inviter à la soutenance de mon Habilitation à Diriger des Recherches intitulée ***«Coopération entre le compilateur et le support exécutif pour le parallélisme multi-paradigme en calcul haute performance»*** et au pot qui suivra.

Elle aura lieu le mercredi 8 juillet à 14h00 dans l'Amphi B (Bâtiment Descartes) sur le campus de Versailles de l'UVSQ (45 avenue des Etats-Unis, 78035 Versailles Cedex).
Le jury sera composé de :

- **Jean-François MEHAUT**, Professeur des Universités, à l'*Université Grenoble 1/Laboratoire d'Informatique de Grenoble (LIG) - UMR n°5217 - Montbonnot-Saint-Martin* – Rapporteur
- **Mathias MÜLLER**, Professeur des Universités, à *l'Université d'Aachen - Aachen (Allemagne)* – Rapporteur
- **David PADUA**, Professeur des Universités, à *l'Université d'Illinois - Urbana (Illinois - Etats-Unis)* – Rapporteur
- **William JALBY**, Professeur des Universités, à l'*Université de Versailles Saint-Quentin-en-Yvelines/Laboratoire Parallélisme Réseaux Systèmes Modélisation (PRISM) - Versailles* – Tuteur
- **Allen MALONY**, Professeur des Universités, à *l'Université d'Oregon/Department of Computer and Information Science - Eugene (Oregon - Etats-Unis)* – Examinateur
- **Jim ANG**, Directeur Technique, au *Laboratoire National de Sandia - Albuquerque (Nouveau Mexique - Etats-Unis)* – Invité
- **Pascale ROSSE-LAURENT**, Architecte Expert, à *Bull - Echirolles* – Invitée

Résumé :

The current evolution of high-hand hardware architecture leads to interesting problems for computer-science research and industry. Since 2010, supercomputers reached the Petaflops threshold allowing a program to run at the speed of 10^15 floating-point operations per second. But during the last five years, various computer-architecture designs arose to prepare the next generation of clusters.
Indeed, the Exascale era (10^18 floating-point operations per second) is predicted to appear by the end of this decade or, at least, early after the 2020 horizon. Even if this goal is still years away, the time to prepare the software environment (and update the applications) is really short. One way for those scientific applications to reach the Exascale milestone is to extend the parallelism from MPI to MPI+X. For this purpose, the whole software environment should evolve too. Indeed, the underlying runtime systems of each programming model have to be aware of each other to deal with resource allocation (cores and different memories). Because mixing parallel programming can be a tough trial, the whole toolchain should help this transition from the compiler to the runtime. This presentation brings together the research I conducted during the last 7 years about the possible evolutions for the software stack (compilers and runtime systems) and makes the following contributions: (i) design of a unified MPI+OpenMP runtime lowering the overhead and exposing additional features e.g., taxonomy for thread placement, (ii) resource management for both cores and memory tested on different architectures (CPUs and GPGPUs), and (iii) compiler support for parallel programming models including interaction with the runtime systems for data placement and debugging tools for various parallel paradigms.

Patrick Carribault