

## Einladung

zum Informatik-Kolloquium des  
AB Programmiersprachen und Übersetzer am  
**Donnerstag, den 8. August 2013, um 11:00 Uhr s.t.**  
in der Bibliothek E185.1, Argentinierstr. 8, 4. Stock (Mitte)

Es spricht

**Dr. Michael Gschwind**

IBM Corp., Poughkeepsie, NY, USA

über

### **Blue Gene/Q: Design for Sustained Multi-Petaflop Computing**

The Blue Gene/Q system represents the third generation of optimized high-performance computing Blue Gene solution servers and provides a platform for continued growth in HPC performance and capability. Blue Gene/Q started with a new design of the hardware platform, while retaining and significantly expanding an established, trusted and successful software environment.

To deliver a system that enables users to fully exploit the promise of high-performance computing for both traditional HPC applications and new commercial application areas, the Blue Gene/Q system architecture combines hardware and software innovations to overcome traditional bottlenecks, most famously the memory and power walls which have become emblematic of modern computing systems. At the same time, to deliver a platform for sustainable petascale computing, and beyond to exascale, we had to address a new set of “walls” with the many innovations described below: a scalability wall, a communication wall, and a reliability wall.

The new Blue Gene/Q system increases overall system performance with a new node architecture: Each node offers more thread-level-parallelism with a coherent SMP node consisting of eighteen 64-bit PowerPC cores with 4-way simultaneous multithreading. Each core provides for better exploitation of data-level parallelism with a new 4-way quad-vector processing unit (QPU). The memory subsystem integrates memory speculation support which can be used to implement both Transactional Memory and Speculative Execution programming models.

The compute nodes are connected in a five dimensional torus configuration using 10 point-to-point links, and a total network bandwidth of 44 GB/s per node. The on-chip messaging unit provides an optimized interface between the network routing logic and the memory subsystem, with enough bandwidth to keep all the links busy. It also offloads communication protocol processing by implementing collective broadcast and reduction operations, including integer and floating point sum, min and max.

Built on the Blue Gene hardware design is an efficient software stack that builds on several generations of Blue Gene software interfaces, while extending these capabilities and adding new functions to support new hardware capabilities. The hardware functions were designed with a focus on providing efficient primitives upon which to build the rich software environment.

To ensure reliable operation of a petascale system, reliability has to be a pervasive design consideration. At the architecture level, new QPX store-and-indicate instructions support the detection of programming errors. To ensure reliable operation in the presence of transient faults, we conducted exhaustive single event upset simulations based on fault injection into the simulated design. The operating system was structured to use firmware in a small on-chip boot eDRAM to avoid silent system hangs.

Together, the hardware and software innovations pioneered in Blue Gene/Q give application developers a platform and framework to develop and deploy sustained petascale computing applications. These petascale applications will allow its users to make new scientific discoveries and gain new business insights, which will be the true measure of the success of the new Blue Gene/Q systems.

**Biographie:** Dr. Michael Gschwind is a Senior Technical Staff Member and Senior Manager of System Architecture in the IBM Systems and Technology Group. In his dual role as a technical leader and manager, he is responsible for leading the architecture evolution of IBM's mainframe System z and Power systems and manages the architecture teams for both System z and Power brands. Previously, Dr. Gschwind served as Blue Gene Floating Point Chief Architect and Lead. Dr. Gschwind served as IFU and IDU lead and chief microarchitect for the Komal core which is the foundation for Power7 systems and as architecture lead for Productive, Easy-to-use, Reliable Computing Systems (PERCS) for DARPA's High Productivity Computing Systems (HPCS) initiative where he served as lead architect for IBM's vector/scalar SIMD architecture (VSX). Dr. Gschwind helped initiate the Cell project, served as one of its lead architects for the Cell definition, developed the first Cell compiler and served as technical lead for the software team. He also helped create the Xbox360 chip. Dr. Gschwind received his PhD from Technische Universität Wien. Dr. Gschwind is an IEEE Fellow, an IBM Master Inventor and a member of the IBM Academy of Technology and was named as an industry-leading "IT Innovator and Influencer" by InformationWeek in 2006. ([http://researcher.watson.ibm.com/researcher/view\\_person\\_subpage.php?id=3796](http://researcher.watson.ibm.com/researcher/view_person_subpage.php?id=3796))

Zu diesem Vortrag lädt der *Arbeitsbereich für Programmiersprachen und Übersetzer am Institut für Computersprachen* herzlich ein.

Tea: 10:30 Uhr in der Bibliothek E185.1, Argentinierstr. 8, 4. Stock (Mitte).