

Conference Sessions

Monday, April 24, 2006

8:30

Registration & Exhibits Open

Bayshore Foyer / Sierra & Cascade • 8:30–9:00

9:00

General Interest

Welcome

Mark K. Smith, Gelato Central Operations
Siskiyou & Donner Pass • 9:00–9:20

Welcome, introduction, and overview of Gelato Federation activities.

9:20

Keynote

Itanium: Its Rationale and Potential from an HP Labs Perspective

William S. Worley, Secure64 and Itanium Solutions Alliance
Siskiyou & Donner Pass • 9:20–10:20

The Intel/HP Itanium architecture definition effort started with the results of an HP Labs research program, called PA Wide Word internally, conducted from January 1990 to December 1993. Concepts and conclusions formulated during this research program established technical principles for a fundamental advance in processor architecture and led to the Intel/HP partnership. Less noticed in published accounts is the fact that many capabilities Intel and HP jointly innovated in the Itanium architecture were specifically designed to enable construction of secure systems.

Non-security objectives have led modern general-purpose operating systems to continue to rely upon a more than 40-year-old, CPU-only, hardware protection model. This limited hardware protection model simply is incapable of supporting the levels of remote-attack security required in today's massively complex systems, in today's online world. As a result, we find vulnerable servers surrounded by vulnerable external protective appliances. All require periodic patching and re-testing. It's not clear the good guys are winning.

Intel's Itanium 2 systems now offer the means for building "inherently secure" systems. Inherently secure means that the software controlling the hardware platform has specific, strong security properties. Without an inherently secure foundation, the current trends of virtualizing servers and consolidating network protective appliances magnify,

rather than mitigate, security risks. Secure64's inherently secure hardware platform control software fully utilizes the capabilities of the Itanium architecture to provide such a foundation. This offers substantial benefits for information systems and infrastructures, and can establish Itanium hardware platforms as the winners both for secure consolidation and for secure virtualization.

10:20

Topics for Enterprise

BP: An Enterprise Itanium Use Case Study

Keith Gray, BP

Siskiyou & Donner Pass • 10:20–10:45

I will discuss the use of high-performance computing at BP, including the seismic imaging research we support, the implications on the computing solutions deployed, and the directions of future efforts.

11:00

Topics for Enterprise

Oracle: An Enterprise Itanium Use Case Study

Brian Hirano, Oracle

Donner Pass • 11:00–12:15

Oracle's 4-way Itanium 2 TPC-C benchmarks, announced in November of 2002, were the culmination of a two-year project involving engineers from Intel, HP, and Oracle. Since that time, multiple groups in Oracle and Intel have continued to work closely on multiple versions of Oracle and Linux-based Itanium platforms to ensure performance and stability for enterprise solutions. This talk discusses the initial performance work and the evolution of Oracle's and Intel's focuses, and presents some of the current areas Oracle and Intel are jointly investigating.

Tools & Tuning

Compiling for the Intel Itanium 2 Processor Code Named Montecito

Martyn Corden, Intel

Siskiyou • 11:00–12:15

New compiler options and strategies are presented for taking advantage of the new features of Intel Itanium 2 processors, code named Montecito, including dual cores and hyperthreading. Other new features of the 9.1 version of the Intel C++ and FORTRAN compilers for Linux will be covered as time permits.

2 Jane Smith Exercise

Advanced Topics

Mathematical Modeling to Formally Prove Correctness

John R. Harrison, Intel
Cedar • 11:00–12:15

Formal verification attempts to establish the correctness of a computer artifact (hardware, software, microcode, protocol, etc.) by rigorous modeling and mathematical proof, rather than merely by testing or simulation. Formal verification in the hardware industry is widely practiced, and increasingly seen as necessary. We can perhaps identify at least three reasons:

- Hardware is designed in a more modular way than most software, with refinement an important design method. Constraints of interconnect layering and timing means that one cannot really design “spaghetti hardware.”
- More proofs in the hardware domain can be largely automated, reducing the need for intensive interaction by a human expert with the mechanical theorem-proving system.
- The potential consequences of a hardware error are greater, since such errors often cannot be patched or worked around, and may in extremis necessitate a hardware replacement.

It is not surprising that a considerable amount of effort has been in the floating-point domain. Floating-point algorithms have proven themselves difficult to get right. Yet in marked contrast to some other targets for formal verification, it is not hard to come up with widely accepted formal specifications of how floating-point operations should behave. In fact, many operations are specified almost completely by the IEEE Standard governing binary floating-point arithmetic. However, in some other respects, floating-point operations present a difficult challenge for formal verification. We will describe some of our work in formally verifying algorithms for operations such as division, square root, and transcendental functions for the Intel Itanium architecture.

Research

Preparing for the First Beam at the LHC

Lawrence Pinsky, University of Houston
San Carlos • 11:00–11:45

The Large Hadron Collider (LHC) at CERN, the European Laboratory for Particle Research in Geneva, Switzerland, is expecting to have the first beam next year. This is the culmination of more than a decade construction project, including the development of the supporting software and computing models. ALICE is one of the four major detectors that is being prepared for physics at the LHC, and the University of Houston is a member of the US contingent of institutions involved in that experiment. Along with the Ohio Supercomputer Center and the facility at NERSC

(LBL), the University of Houston Itanium cluster has been participating in the increasingly severe sequence of “data challenges” that are being wrapped up now in preparation for the actual turn-on of the LHC. The ALICE computing model is necessarily dependent upon a grid-based model that will include many different platforms, Itanium among them. The data challenges have provided a good venue to compare the relative attributes of the various platforms in running the kind of simulations and analysis codes that are relevant to particle physics applications. An overview of these results will be presented along with a summary of the overall ALICE computing plans and the status of its deployment.

11:45

Research

Computing Optimal Equilibrium Strategies for Network Economies

Alejandro Jofré, University of Chile
San Carlos • 11:45–12:30

Models for regulating, planning, and operating industries working on networks such as energy, transportation, and telecommunication are key ingredients today of which to take advantage. These models correspond to large stochastic optimization/equilibrium problems, which are very difficult to solve. In this talk, we will show three new distributed algorithms/strategies to compute a solution and its implementation on Itanium 2 clusters. This family of models and/or solutions are currently used by several companies and institutions participating in these industries.

12:15

Lunch & Exhibits Open

Sierra & Cascade • 12:15–13:30

13:30

General Interest

Basic Itanium Architecture

Cameron McNairy, Intel
Donner Pass • 13:30–14:45

The Itanium architecture and the paradigm of explicit parallel instruction computing (EPIC) are often poorly understood. This presentation will cover important aspects of the EPIC paradigm, including software pipelining, register save engine, predication, parallel instruction groups, data and control speculation, and many other mysteries of the Itanium application and system architectures.

13:30

Tools & Tuning**Columbia Application Tuning Case Studies**

Johnny Chang, NASA

Siskiyou • 13:30–14:15

This talk will present several case studies of application performance enhancements on the SGI Altix platform. The enhancements include both explicit (dplace) and implicit (cpubind/cpuset_pin) process-pinning, eliminating memory contention in OpenMP applications, eliminating unaligned memory accesses, and system profiling. These enhancements enabled 2- to 20-fold improvements in application performance.

Advanced Topics**Kernel Optimization for Enterprise Workloads**

Kenneth Chen, Intel

Cedar • 13:30–14:15

Linux has been receiving a great deal of attention in the past few years. The popularity is being propelled by a wide range of adoption of Linux for enterprise computing. Major software vendors have been supporting their products on Linux for many years. As the enterprise software solution stack builds up everyday, it is crucial that Linux kernel development takes this opportunity to ensure that the kernel provides necessary infrastructure for enterprise application to excel. This means developing enterprise focused OS features, improving performance by extending the scalability, as well as improving many other areas.

Adding to the excitement, the Intel Itanium 2 processor is built with many innovative features that push the performance envelope. Featuring massive caches and CPU execution resource, EPIC technology (explicitly parallel instruction computing) provides a variety of optimization opportunities. In this talk, we will highlight kernel optimization work done on Linux-ia64, ranging from several critical low level assemblies to generic kernel components. We will present how the linux-ia64 kernel utilizes Itanium architecture features to extend scalability and performance for enterprise workloads.

Research**Mathematical Libraries and the Implementation of Parallel Solvers for Engineering**

Hugo Daniel Scolnik, University of Buenos Aires

San Carlos • 13:30–14:15

Our research is focused on developing a highly efficient parallelizable solver of huge systems of linear equations that arise from finite element discretizations of complex nonlinear engineering problems. Those problems are nonlinear, require many linearizations, and hence several

days of CPU time on Itanium platforms. Another important application is the reconstruction of tomographic images.

This work includes a comparison of the mathematical libraries like MKL (Linux) and MLIB (HP) from the point of view of the performance on numerical problems using sequential and parallel implementations. The new solver uses BLAS routines at levels 1,2,3 excluding complex data types. The conclusions of our study present the results obtained with several problems.

14:15

Tools & Tuning**HP Caliper: An Update to the Linux IPF Performance Tool**

Curt Wohlgemuth and Steve Williams, HP

Siskiyou • 14:15–15:00

HP Caliper is a sophisticated general-purpose performance analysis tool that takes advantage of the Itanium processor's advanced performance monitoring unit to provide detailed and accurate performance measurements at the application and system level with minimal perturbation to the system's behavior.

Besides an overview of HP Caliper, we will discuss new features, including system-wide profiling and a new graphical user interface based on the rich client platform of Eclipse.

Focus on GCC**The ISP RAS Effort to Improve GCC for Itanium**

Arutyun I. Avetisyan, Russian Academy of Science

San Carlos • 14:15–15:00

Ongoing work at ISP RAS on improving GCC for Itanium processors will be presented. Discussion will cover a past project with HP on improving GCC instruction scheduling and the current effort on implementing a new VLIW-targeted instruction scheduler. Future plans on improving GCC for Itanium and potential collaboration projects will also be presented, including plans for a GCC meeting in Moscow this summer.

Advanced Topics**Suggested Improvements in Itanium and Software**

Clemens C. J. Roothaan, Gelato Honorary Member

Cedar • 14:15–15:00

In general, the Itanium is a major step forward in computer design. Nevertheless, there are still gaps in the instruction repertoire, and the specifications of some instructions could be expanded or modified. There are also some mandates by C++ concerning corner cases, which cannot

4 Jane Smith Exercise

be justified by any mathematical reasoning whatsoever; there is even one IEEE mandate that cannot pass muster. A detailed list of shortcomings and possible remedies will be presented for your consideration.

15:00

General Interest

An Evaluation of High Performance Octave on Itanium

Ashok Krishnamurthy, Ohio Supercomputer Center
Donner Pass • 15:00–15:45

GNU Octave is a MATLAB-style interactive application for performing numerical computations. The Octave language is mostly compatible with MATLAB. MATLAB (and Octave) are being used as an executable specification language to develop synthetic compact applications for the DARPA HPCS program. This work has identified a clear need for a MATLAB-style interpreter that can handle large address spaces, run on multiple processors, and leverage high-performance interconnects.

The Ohio Supercomputer Center (OSC), Ohio State University, and Indiana University have been collaborating on research and software technologies for parallel Octave. We have constructed a version of parallel Octave for the Itanium 2 cluster at OSC. This interpreter has a 64-bit address space for large matrix support and uses the high-bandwidth Myrinet interconnect. This talk will review the software architecture, performance, and scalability of parallel Octave on the OSC Itanium 2 cluster.

Topics for Enterprise

Computational Requirements for Compute Intensive Applications

Christian Tanasescu, SGI
Cedar • 15:00–15:45

We will investigate the computational requirements, performance characteristics, and scalability of some of the most used compute intensive applications in scientific and engineering communities. Our analysis covers HPC application segments like Computational Structural Mechanics (CSM), Computational Fluid Dynamics (CFD), Computational Chemistry and Material Science (CCM), Bioinformatics (BIO), Seismic Processing and Reservoir Simulation (SPR), and Climate/Weather/Ocean Modeling and Simulation (CWO).

The performance of an application results from many factors and not just from peak processing capability of a system. Generally speaking, compute intensive applications exhibit a range of various HPC resource demands. Applications in Computational Chemistry and Computational Structural Mechanics (static problems) reuse data read from memory in a large degree and are

considered cache-friendly. On the other side, applications in Computational Fluid Dynamics, Computational Structural Mechanics (dynamic problems like noise, vibration, harshness), and Reservoir Simulation consume the data read from memory and have to load continuously new data from memory. To keep the FP units busy, these applications require computer architectures with high memory bandwidth, mainly due to the data addressing patterns and heavy IO activities.

We will conclude by analyzing how the Itanium 2 processor features address the computational needs of HPC applications. For large-scale problems, understanding the communication patterns of MPI applications is essential to achieve scalability on hundreds of processors in cluster architecture.

Tools & Tuning

VTune Update

Paul M. Cohen, Intel
Siskiyou • 15:00–15:45

This talk will cover what's new for tuning Intel Itanium 2-based applications, including native Eclipse IDE and NUMA aware support for data collection.

Focus on GCC

GCC IP Issues

Dan Berlin, IBM
San Carlos • 15:00–15:45

This talk will cover a variety of intellectual property issues that come up during working on GCC, including:

- Copyright: Assignments of copyright, and how we deal with issues of contributions of code from other open source/commercial projects.
- Patents: How we deal with them in GCC, and what we require of companies that are going to contribute to GCC.
- General other issues related to intellectual property and GCC.