

The Itanium ERA: Education, Research, Application

Gelato Federation Meeting | October 2-5, 2005, Brazil



Sunday, October 2, 2005

18:30 Sponsor-Hosted Kickoff Dinner, Sheraton Porto Alegre Hotel, 22nd Floor

Monday, October 3, 2005

8:00 Bus Departs Hotel for Gelato Meeting

8:30 Registration

9:00
Main Room

Welcome and Itanium Research at PUCRS

PUCRS will welcome Gelato meeting attendees to their Porto Alegre campus and will present their three main Itanium-related projects.



Jorge Audy - Pontifical Catholic University of Rio Grande do Sul

Jorge Audy is a Professor of Computer Science and Vice President of Research and Graduate Studies at the Pontifical Catholic University of Rio Grande do Sul (PUCRS), Brazil. He holds a PhD in Information Systems and is currently involved in projects related to global software development and project management.



César De Rose - Pontifical Catholic University of Rio Grande do Sul

César De Rose is an Associate Professor in the Computer Science Department at the Pontifical Catholic University of Rio Grande do Sul (PUCRS), Porto Alegre, Brazil. His primary research interests are parallel and distributed computing and parallel architectures. He is currently conducting research on a variety of topics applied to clusters and grids, including resource management, resource monitoring, and distributed allocation strategies.

Dr. De Rose received a PhD in Computer Science from the University Karlsruhe, Germany, in 1998. He currently leads the Research Center in High Performance Computing (CPAD - PUCRS/HP) at PUCRS.

10:00
Main Room

State of the Federation

Mark will give an overview and status of the Federation's growth and will report on collective member characteristics. In addition, he will discuss some of the Federation's key areas of interest and recent activity at Gelato Central Operations.

Additionally, Matthieu Delahaye will brief attendees on the Gelato Coconut project. Andy Schuh will preview the 2006 Gelato San Jose conference.



Mark K. Smith - Gelato Central Operations

Mark K. Smith is the Managing Director of the Gelato Federation. He works with Federation members and sponsors around the world, fostering collaborative relationships among members, sponsors, and the general community to advance the Linux-Itanium platform. Mark leads a technical team at University of Illinois and dedicates time to educating the general community about the advantages of the platform. Prior to joining Gelato, he worked in the software industry for 10 years. Mark holds a PhD in Engineering from the University of Illinois.

10:45 **A Word from HP**

Main Room This brief keynote will take a short look back at Gelato's growth and activities, and Itanium adoption past, present, and future.



Steve Geary - HP

Steve Geary is responsible for overseeing HP's open-source and Linux technical strategy as well as development of open-source and Linux technologies and solutions. Steve is also responsible for centralized open-source and Linux functions such as technical support; HP's Open Source Program Office, which hosts internal and external development efforts; HP's Open Source Review Board; and HP's R&D relationship with commercial distributions. Steve is active in various capacities in open source. He is presently HP's board member for the Open Source Development Lab (OSDL), recently resigning his 3 year tenure as OSDL's Data Center Linux Steering Chair. Steve is also responsible for HP's activity with the Linux Standards Base and Eclipse.

11:00 **Itanium Solutions Alliance, An Introduction**

Main Room

11:15 **Itanium Tools**

Main Room Experts will introduce the pfmon tool, qtools, OProfile, PAPI-based tools, Intel VTune, Intel C++ Compiler (icc), and performance libraries.

Audience participation is encouraged.

12:15 **Lunch**

Caryl Malone

Mark Davis - Intel

Eric W. Moore - Intel

Paul M. Cohen - Intel

Stéphane Eranian - HP

Al Stone - HP

Philip Mucci - Royal Institute of Technology

13:30
Main Room

Java on the Itanium

In a world where the habit of using 32-bit computing has ruled the market for a long time, it's a challenge to convince users to switch to 64-bits. Usually arguments against migrating to 64-bit computing take one of two forms: the need to recompile and perhaps re-design and re-implement the Java application; and memory requirements that eventually might result in longer response times due to more exhaustive garbage collection pause times.

This session will refute these common arguments by introducing JRockit as the solution. In this presentation, we will discuss the pros and cons of a 64-bit world and provide an overview of the new Itanium-specific features in JRockit. We will also show how they mitigate the known problems of using 64-bit computing. Focus will be on MME (Mixed Mode Execution), compressed references, and the Deterministic Garbage Collector. The session will also contain a discussion about DPGO (Dynamic Profile Guided Optimizations) and present benchmark results showing JRockit performance on Itanium.

Bundle #0

14:30
Annex 1

Itanium Architecture 101

This lecture will cover the principles of the EPIC architecture that software developers need to know. In this pursuit, the lecture will introduce the registers, instruction formats, predication, and speculation of the Intel Itanium architecture, culminating with the vocabulary and principles of how software pipelining is accomplished.



Eva Andreasson - BEA Systems

Eva Andreasson received an MSc in Computer Science from the Royal Institute of Technology in Stockholm, Sweden, in 2002. She also received the award of *Best Master Thesis Project 2002* for her project on the use of machine learning in JVMs.

Since then, she has been developing the BEA JRockit JVM at BEA Systems. Her focus is on memory management, including performance improvements of the parallel garbage collector and heuristics behind the deterministic garbage collector. Her areas of interest include: algorithms to improve the performance of JRockit, better and more stable behavior through adaptivity, and ideas and solutions around distributed garbage collection. Currently, she is acting development manager and is in charge of developing the upcoming BEA JRockit release.

During her time at BEA Systems, she has made several presentations at conferences, like USENIX JVM-02 and BEA World, and has contributed to two pending patents: one for the use of reinforcement learning in JVMs and one for control theory in VEEs (virtual execution environments).



Eric W. Moore - Intel

Eric Wynne Moore is a Senior Software Engineer working in the Software Products Division at Intel Corporation. In the past, he has worked at Rational, Microsoft, RealAudio, Digital, Compaq, and Keane. His specialties include operating systems, compilers, high-performance computing, and performance tuning. In the last couple of years, Moore has trained more than 500 engineers in performance optimization, including engineers at PNNL, CIA, FBI, IBM, Dell, Hewlett Packard, SGI, Cisco, Intel, several universities, as well as all over the world, including and around Korea, China, Brazil, and Europe.

14:30
Annex 2**Transcendental Function Calculations on Itanium 1 - Determining Optimal Polynomial Expansions**

Chebyshev Expansions of a real function of a real argument are formulated as linear transformations of the MacLaurin Expansion of that function. The transformation matrices connecting the Chebyshev and MacLaurin expansions are independent of the function at hand; however they depend explicitly on two parameters, namely the highest order of the polynomials used in the expansion, n , and the maximum absolute value of the argument, $d(L)$, where the index L labels appropriate different choices for that maximum. A significant simplification of the MacLaurin-Chebyshev transformation algebra is obtained by adopting an alternative definition of the Chebyshev Polynomials which is known, but has rarely been used. After choosing the ranges of n and L , and the magnitudes $d(L)$ for a given function, straightforward application of this new algebra yields a Valid Bits Table and, for any particular choice of L , n , and $d(L)$, the Coefficients of the Chebyshev Expansion. In addition, an efficient, quadratically convergent iterative procedure for calculating a Remez Expansion from a Chebyshev Expansion, with or without additional constraints, will also be presented.

**Clemens C. J. Roothaan - Secure64; University of Chicago**

Clemens Roothaan is Professor Emeritus of Physics and Chemistry at the University of Chicago. In the 1950's, he published detailed algorithms to solve quantum mechanical movements of electrons in molecules and atoms. Today, most computer programs in this area are based on his method. After his retirement from the University in 1988, Roothaan started to work for HP Labs in Palo Alto, California. He has worked on the Itanium design team since 1990. Currently, Roothaan is working on a large software suite of scientific tools for function evaluation.

15:30

Afternoon Break**Bundle #1**16:00
Annex 1**Compiler Optimizations for Transaction Processing Workloads on Itanium Linux Systems**

This talk describes a repertoire of well-known and new compiler optimizations that help produce excellent server application performance and investigates their performance contributions. These optimizations combined produce a 40% speed-up in on-line transaction processing performance and have been implemented in the Intel C/C++ Itanium compiler.

**Mark Davis - Intel**

Mark Davis is a Senior Principal Engineer at Intel. He serves as an architect of Intel's Itanium Compiler Lab, providing high-quality, high-performance compilers for enterprise-class Itanium platforms. He has also been co-manager of the Itanium Compiler Development team, and co-manager of the Itanium code generator. Mark has specialized in compiler optimizations, performance analysis, and architecture design in his career. At Digital/Compaq, he worked on the GEM compiler for Alpha, and was technical lead when GEM was targeted to Itanium. As compiler lead at Stellar/Stardent, he helped design the world's first graphics supercomputer and delivered high performance compilers for it. While at Intermetrics, he was a language designer in DoD's competition for the Ada language and then was technical lead for an Ada compiler; later he helped design and became technical director of a highly-optimizing PL/I compiler developed for IBM.

Dr. Davis received his PhD in Computer Science from Harvard.

16:00
Annex 2**Focus on Grid (Session #1)**

This section intends to discuss the grid efforts of Gelato members and how these efforts relate to Linux on Itanium. There will be an open plenary discussion and three short presentations. Dr. De Rose will talk about the GerpavGrid Project and how it utilizes the power of computational grids in the public administration of the City of Porto Alegre. Dr. Jarp will bring the latest updates on the CERN grid. Dr. Cirne will present the OurGrid project and how it intends to reduce the "scientific digital divide."

**Walfredo Cirne - Universidade Federal de Campina Grande**

Since July 1995, Dr. Walfredo Cirne has been a Professor at UFCG, acting as a researcher in distributed and parallel computing, advisor of MS and PhD students, and teacher of graduate and undergraduate classes. Previous to 1997, Dr. Cirne worked on computer networks and machine learning. Since then, his research has focused on grid computing. Currently, he coordinates one of the greatest research projects in computational grids in Brazil, OurGrid, a project developed in cooperation with HP that aims to provide a complete grid solution for bag-of-tasks applications, and includes the search of good out-of-the-box performance for grid middleware running on Itanium.

Dr. Cirne holds a PhD in Computer Science from the University of California San Diego, and has published dozens of papers in the main international publications and computer science conferences.

16:45
Annex 1**Bundle #2****CERN Loops - Can We Extrapolate Total Application Performance?**

CERN programs often have very flat execution profiles, which implies that the execution time is spread over many routines/methods. Consequently, compiler optimization must be applied to the whole program and not just a few inner loops. In this talk we, nevertheless, discuss the value of extracting some of the most solicited routines (relatively speaking) and using them to gauge overall performance. An initial set of ten C++ routines have been extracted from three CERN packages (ROOT, GEANT4, and CLHEP). One main advantage is, of course, that the routines compile and execute in seconds, allowing lots of testing of different platforms, compilers, and compiler options. The speaker will review the initial selection and show results with GCC and icc on both the Xeon and the Itanium platforms.

**Sverre Jarp - European Organization for Nuclear Research**

Sverre Jarp is the Chief Technology Officer in CERN's openlab for DataGrid Application, which is a joint collaboration with industry in order to assess leading-edge information technology for the Large Hadron Collider's Computing Grid in 2007.

He has been working in computing at CERN, the European Organization for Nuclear Research, for over 28 years and has held various managerial and technical positions promoting advanced but cost-effective computing solutions for the Laboratory. In 2001-02, he spent a sabbatical year at the HP Labs (Palo Alto, USA) working on software for the Itanium Processor Family. His current field of interest is, in particular, compiler optimization.

S.Jarp holds a degree in Theoretical Physics from the Norwegian University of Science and Technology in Trondheim.

16:45
Annex 2**Focus on Grid (Session #2)****Walfredo Cirne - Universidade Federal de Campina Grande**

17:30 **End of Monday's Main Meeting - Board Bus for Hotel**

19:30 **GCC Improvement (Session #1)**

Hotel

Three presentations are planned for this evening:

19:30 - 20:00 - Diego Novillo - Red Hat

In this talk, Diego will provide a general architectural overview of GCC. He will also describe the development process: how to participate and contribute to its development. Finally, he will talk about some of the major challenges and opportunities for improvement.

20:00 - 20:30 - Shin-Ming Liu - HP

The HP GCC project team is made up of three full-time engineers. The primary objective of the project is to improve the customer experience on the Linux-Itanium platform. The current focus of the team is to fix GCC defects posted in the Bugzilla database, submit patches to the FSF tree, and regularly monitor the quality of the major support line of GCC and report or fix any regressions. The project also helps coordinate the GCC performance enhancement activities among various organizations.

20:30 - 21:00 - Bob Kidd - UIUC

Superblock formation creates a chain of single predecessor basic blocks which can then be specialized by later optimizations. Moving the Superblock formation pass to early in the Tree-SSA stages of GCC opens the door for application of optimizations in a more powerful and maintainable intermediate language. Bob will discuss the details of his patch, present results, and future plans.



Shin-Ming Liu - HP

Shin-Ming Liu is the Project Manager for High-Level Optimization and GCC of the Itanium C/C++ Compiler Section of the Java, Compiler, and Tools Lab at HP in Cupertino, California. Liu led the development effort for the high-level optimization and code generator project in compiler targeted for the Itanium processor. In this project, he helped redesigned the high-level optimization into a highly-robust, scalable, and efficient component by rearchitecting the infrastructure, from which many new techniques were developed. Many highly-recognized programming analysis methods were adopted as well. Liu led the reinvention of compiler development methodology by focusing on modulization, memory footprint control, canonical internal representation, and automatic error detection. Before joining HP, he worked at MIPS/SGI in the area of compiler front end, middle end, back end, and linker. During that time, he co-authored several technical publications.

Tuesday, October 4, 2005

8:30 **Bus Departs Hotel for Gelato Meeting**

9:00 **Poster Session**

Bundle #310:00
Annex 1**Bugs! Getting Them Stomped**

Eric Raymond invented what he called, Linus's Law: "Given enough eyeballs, all bugs are shallow." What he meant by that was that in the open-source world, where bug reporter, bug fixer, and core developer share a common view of the system, bug reports are of better quality, and fixes are easier to find.

Unfortunately, even though much of what we work on is open-source, it's often hard to work out how to report a problem and get it fixed. Consequently, many of us carry along sets of patches and workarounds for many months, maybe even years, until "something happens."

In this talk, I'll be attempting to elucidate how to interact with the open-source community so that the problems *you* have are fixed up-stream. The essence is communication.

**Peter Chubb - University of New South Wales**

Peter Chubb is a Senior Research Engineer at National ICT Australia and a Research Officer at UNSW. He completed his PhD under Associate Professor John Lions in 1989. Peter worked at Softway Pty Ltd as a consultant and software engineer doing UNIX kernel, security, and embedded work. He joined Gelato@UNSW at its inception in 2002.

Peter started using UNIX in 1979 and has never used Microsoft operating systems for more than a few moments. His home life includes wife Lucy, who also works at Gelato@UNSW, and two small daughters. Peter's hobbies include music (he runs a recorder consort), aquaria (3 tanks at present, no room for more), and fine wines.

10:00
Annex 2**GCC Improvement (Session #2)**

10:00 - 10:30 - Canqun Yang - NUDT

Canqun will give a brief introduction of the CCRG group and their improvements for GCC on Itanium 2, which focus on function inlining. Performance results produced by their improved GCC will be given as well as a look at their ongoing work.

10:30 - 11:30 - Next Steps, Action Items

- Review and update IA-64 project list.
- Continue discussion on alias analysis improvements and determine concrete ways to help.

**Shin-Ming Liu - HP**

Shin-Ming Liu is the Project Manager for High-Level Optimization and GCC of the Itanium C/C++ Compiler Section of the Java, Compiler, and Tools Lab at HP in Cupertino, California. Liu led the development effort for the high-level optimization and code generator project in compiler targeted for the Itanium processor. In this project, he helped redesigned the high-level optimization into a highly-robust, scalable, and efficient component by rearchitecting the infrastructure, from which many new techniques were developed. Many highly-recognized programming analysis methods were adopted as well. Liu led the reinvention of compiler development methodology by focusing on modularization, memory footprint control, canonical internal representation, and automatic error detection. Before joining HP, he worked at MIPS/SGI in the area of compiler front end, middle end, back end, and linker. During that time, he co-authored several technical publications.

Bundle #410:45
Annex 1**New Intel Itanium Architecture Specific Features for Intel VTune**

Intel's VTune Performance Analyzer is a robust enterprise grade solution even with large executables (100MB+) that other products are unable to profile. This talk will cover specific features being added to better support supercomputer systems based on Intel Itanium 2 processors.

Productive Eclipse 3.1 Integrated Design Environment

VTune Performance Analyzer 8.1 for Linux makes application performance tuning easier with an improved graphical user interface which will be available for the first time on Intel Itanium 2 processors this includes wizards to simplify configuration and quickly get to application hotspots without learning about the tool. No recompiles or changes to your build script are required to use VTune Analyzer.

Intel Itanium 2 Processor Features

Specific topics I expect to cover include new lower overhead ways to collect data on multi-processor computers, event editing to allow more experienced users to tailor event-based sampling for specific purposes, opcode matching an advanced collection technique specific to Intel Itanium family processors. Compiler optimization report integration allowing users to easily view complex reports. Multi-user callgraph allows several users to do tuning on a single Itanium 2 computer simultaneously.

**Paul M. Cohen - Intel**

Paul Cohen is a Performance Tools Product Line Marketing Manager at Intel. He is responsible for Intel tools targeted at improving the performance of customer applications. His current focus is on improving usability of VTune Performance Analyzer, making it a robust enterprise-grade solution able to deal with extremely large executables (100MB+) that other products are unable to profile. In addition, he is working on integration of the VTune Analyzer with Intel C and FORTRAN compilers under Eclipse with the ability to provide a close connection between Intel compiler optimization reports and performance bottlenecks represented in the VTune Analyzer.

10:45
Annex 2**GCC Improvement (Session #3)****Shin-Ming Liu - HP**11:30
Main Room**Community Discussion Time**

We will use this time to:

- Solicit feedback on the proposed change in Gelato governance as proposed to the gelato-all list prior to the meeting.
- Report back on the progress of the Gelato Vanilla project (Matthieu Delahaye).
- Continue our discussion started in the May '05 meeting about applications (proprietary and open-source) that should be ported to or optimized for Itanium.

**Mark K. Smith - Gelato Central Operations**

Mark K. Smith is the Managing Director of the Gelato Federation. He works with Federation members and sponsors around the world, fostering collaborative relationships among members, sponsors, and the general community to advance the Linux-Itanium platform. Mark leads a technical team at University of Illinois and dedicates time to educating the general community about the advantages of the platform. Prior to joining Gelato, he worked in the software industry for 10 years. Mark holds a PhD in Engineering from the University of Illinois.

12:30

Lunch

Bundle #5

14:00

Annex 1

Itanium Research at the University of Split

The Department of Physics at the University of Split has been enrolled in research and education in natural as well as other sciences (from technical and computer to human sciences). This talk will review Itanium-related research being conducted at the University of Split, including:

- Simulating proton-proton collisions for the CMS detector at the Large Hadron Collider (LHC) at CERN near Geneva, Switzerland.
- Studying the behavior of hot and dense nuclear matter as part of the FOPI heavy-ion experiments at the GSI (Gesellschaft für Schwerionenforschung) in Darmstadt, Germany.
- Developing software for the CROATEA project (Cosmic Ray Observatory at the Eastern Adriatic). Because we are in the early stages of the project, our current focus is the development of MF simulations using a computer cluster for simulation of atmospheric showers, telescope system simulation, design and implementation of the software system.
- Developing our own software:
 - SPLIT4.0 - a program for TMH modeling in proteins
 - GARLIC - a molecular visualization program
- Developing our own modules (currently working on optimization of MD algorithms).
- Preparing to run a non-linear stability analysis of the envelopes of giant planets, using fluid-dynamics code written in Ada95.

All this gives us experience in designing/enhancing/porting/tuning our own or third-party applications between 32- and 64-bit platforms.



Mile Dzelalija - University of Split

Mile Dzelalija is a Professor of Physics at the University of Split. He received his PhD in 1995 at the University of Zagreb, Croatia, with the thesis "Entropy in Au + Au Reactions at Relativistic Energies" after researching at the GSI, Darmstadt, Germany. Recently, he has been very active in quality assurance in education at the university and national level.

Dzelalija's main research interests include: simulations and data analysis of heavy-ion reactions (FOPI and CBM experiments at the GSI, Darmstadt), and Higgs boson and supersymmetry particles observability in high-energy experiments (future CMS experiment at the LHC at CERN in Geneva), especially in determining some global properties of reaction systems and new particles. He is also interested in experimental and theoretical research activities in Biomechanics, Environmental Physics, and Physics in Conservation of Fine Arts. In 2004 as part of the UNESCO (United Nations Educational, Scientific and Cultural Organization) project sponsored by HP, the Department of Physics of the Faculty of Natural Sciences, Mathematics, and Education at the University of Split received two Itanium machines. The focus of the project is to design, enhance, port, and tune the Faculty's own or third-party applications to the Itanium architecture in order to achieve better results regarding the time and precision of the results. The Faculty intends to integrate their computer systems into this grid project.

Focus on Scalability (Session #1)

The focus on scalability session will contain a series of short presentations of work in progress at several organizations to measure and improve the scalability of Linux. Although the focus is on scalability of Linux on IA-64 platforms, all platforms will benefit from much of this work. In addition to the presentations, we hope to have time to discuss the presented work and arrange for on-going collaboration in the areas presented.

Currently scheduled presentations include:

"Adapting Linux Scheduling Domains to Multi-level NUMA Hierarchies"
Avelino F. Zorzo - PUCRS

A multiprocessor computer Linux kernel uses several process queues from which processes are selected to be scheduled. Because a process queue can become overloaded while other queues can be empty (or "underloaded"), Linux has a load balancing algorithm to balance all process queues in the system. Currently, Linux builds a memory access level hierarchy, which does not represent correctly the actual number of memory machine access levels. In our project, we have implemented a new strategy to build the correct hierarchy based on information provided in the SLIT table. We intend to present some results we have already produced based on this new strategy.

"Work in Progress on Linux Scalability" Peter Chubb - UNSW

Peter will be presenting results of file system scalability measurements and other efforts underway at UNSW.

"HP's Linux Scalability and Performance Effort" Lee Schermerhorn - HP

Lee will present an overview of various performance and scalability issues that the HP Linux Scalability and Performance project has investigated or is currently working. The presentation will include information regarding some of the Linux performance measurement/instrumentation tools being used at HP to investigate these issues. Areas of investigation include: AIM7 scaling on midrange systems [16cpu, ~140 file systems, ...], locking bottlenecks [page up to date lock, global inode lock, ...], LVM2/DM scaling vs raw SCSI performance, etc.



Lee Schermerhorn - HP

As a member of the Linux Performance and Scalability team within HP's Linux and Open Source Lab (LOSL), Lee Schermerhorn works on performance engineering for Linux, primarily on HP Integrity (Itanium) platforms, with emphasis on NUMA scheduling/affinity and (storage) IO performance.

SLOT 1

Bundle #6

14:45
Annex 1

Performance Comparison of Data Reordering Algorithms

Several performance improvements for parallel finite element edge-based sparse matrix-vector multiplication algorithms on unstructured grids are presented and tested. Edge data structures for tetrahedral meshes and triangular interface elements are treated, focusing on nodal and edges renumbering strategies for improving processor and memory hierarchy use. Benchmark computations on Intel Itanium 2 processors are performed. The results show performance improvements in CPU time ranging from 2 to 3.



Alvaro Coutinho - UFRJ

Alvaro Coutinho is the Director of the Center for Parallel Computing and a Professor in the Department of Civil Engineering at the Alberto Luiz Coimbra Institute for Graduate Studies and Research in Engineering (COPPE) at the University of Rio de Janeiro (UFRJ), and is also a Research Fellow with the Brazilian National Scientific Research Council. Coutinho's areas of current research include: new algorithms for improving processor efficiency in unstructured grid parallel computations, a computational environment for petroleum systems modeling, and the GRAD-GIGA project, a computational grid for high-performance computing.

Coutinho received MSc and DSc degrees in Civil Engineering from UFRJ and has published 62 journal papers and 250 conference papers.

SLOT 2

14:45
Annex 2

Focus on Scalability (Session #2)

Lee Schermerhorn - HP

15:30

Afternoon Break

Bundle #7

16:00

Main Room

Paravirtualization without Pain

Virtual systems are useful for many purposes. One of my favorites is to allow the development of operating systems without having to reboot. But also, for server consolidation, resource isolation, etc.

However, a fully virtual system usually runs much more slowly than the bare metal, because page faults, and privileged operations, cause expensive traps into the host system. A common approach is to modify the guest system to call directly into the host without trapping, allowing much better performance.

Unfortunately, paravirtualization (as this approach is called) leads to fairly invasive changes to the guest operating system. And because the changes are usually specific to a particular virtual machine, they have to be redone for each virtual machine.

We propose a technique called "afterburning" that automatically paravirtualizes the guest operating system. With a bit of linker/loader magic, the resulting binary can run either on a virtual or a real machine. Performance is very similar to that of a manually paravirtualized system, but the amount of change to the operating system is much less.

**Peter Chubb - University of New South Wales**

Peter Chubb is a Senior Research Engineer at National ICT Australia and a Research Officer at UNSW. He completed his PhD under Associate Professor John Lions in 1989. Peter worked at Softway Pty Ltd as a consultant and software engineer doing UNIX kernel, security, and embedded work. He joined Gelato@UNSW at its inception in 2002.

Peter started using UNIX in 1979 and has never used Microsoft operating systems for more than a few moments. His home life includes wife Lucy, who also works at Gelato@UNSW, and two small daughters. Peter's hobbies include music (he runs a recorder consort), aquaria (3 tanks at present, no room for more), and fine wines.

Itanium Research at the University of Buenos Aires

Topics presented will cover: new fast and robust solvers for very large linear systems implemented on Itanium 2 computers, and applications to computational mechanics and image reconstruction problems.

**Hugo Daniel Scolnik - UBA**

Hugo Daniel Scolnik is a Professor in the Computer Sciences Department (that he founded in 1984) at the School of Sciences of the University of Buenos Aires (UBA) where he teaches Cryptography, Numerical Analysis, and Optimization. For his Gelato-related work, Scolnik co-directed a Gelato-sponsored project comparing 64- and 32-bit architectures from the point of view of their performance for scientific programming. Scolnik is also currently directing three of his five graduate students on Gelato-related theses.

Beyond his work at UBA, Scolnik was an international consultant for United Nations agencies, HP, and Hitachi. He has been a Visiting Professor in several countries. He represents Argentina on the International Federation for Information Processing (IFIP) Technical Committee 7 (TC7). He has published papers on Optimization, Numerical Analysis, Automata Theory, Artificial Intelligence, Robotics, and Mathematical Modeling, and has refereed several journals. In 2003, Scolnik won the Konex Award for the best trajectory in Science and Technology for the 1993-2003 decade in the area of Informatics.

Scolnik received a Licenciado en Ciencias Matemáticas at the University of Buenos Aires in 1964, and a PhD in Mathematics from the University of Zurich, Switzerland, in 1970.

16:45

Main Room

Xen-Virtualized Machines

Xen is rapidly becoming the de facto standard for open-source virtualization, with capabilities and performance matching or exceeding leading industry products. Paravirtualization techniques, efficient inter-domain virtual I/O mechanisms, clever migration, and support for multiple architectures (including VT and Pacifica hardware) have contributed to a large broad base of developers and has piqued industry interest.

Xen/ia64 is the first non-x86 architecture supported by Xen. It is still a work-in-progress, but the core hypervisor component utilizes code and/or experience from Xen, Linux/ia64, and the HP vBlades research project. Many interesting strategies are employed to ensure correctness, optimize performance, and leverage the many rapidly developing layers of tools provided by Xen.

We will provide a brief overview of virtualization in general, Xen specifically, and the current status of Xen/ia64. Then we will spend the remaining time discussing some interesting details about the inner workings of Xen on Itanium.



Dan Magenheimer - HP

Dan Magenheimer is a senior scientist working for HP Labs, Fort Collins, Colorado, USA. Dan joined HP in 1982 as a member of the processor architecture team that developed PA-RISC; he wrote the first PA-RISC simulator, remote debugger, object-code emulator (for the 16-bit HP3000), integer multiplication algorithm, and linker.

From 1985 until 2001, he managed various software teams in HP's software, server, and storage divisions. Returning to HP Labs in 2001, Dan joined a team investigating security and virtualization on the Itanium platform; this team developed vBlades, the first Itanium virtual machine monitor. When the Xen open-source virtual machine monitor was announced in 2003, Dan commenced a port of Xen to Itanium (Xen/ia64), utilizing the lessons learned in vBlades and also directly leveraging Linux/ia64 code. Dan is currently the maintainer of Xen/ia64 and is working with a multi-company, worldwide team of Itanium experts to help deliver the first open-source virtual machine monitor for Itanium capable of running multiple SMP guests and supporting migration.

Dan has a BA in Computer Science from the University of California and a MSEE from Stanford University. He is a member of the ACM and IEEE.

17:30

End of Tuesday's Main Meeting - Board Bus for Hotel

Relax at hotel until social and dinner.

19:30

Board Bus at Hotel for Sponsor-Hosted Social Event and Dinner

Location: Churrascaria Galpão Crioulo, Parque Maurício Sirotsky Sobrinho, s/n°.

Wednesday, October 5, 2005

9:00

Bus Departs Hotel for Gelato Meeting

Bundle #9

SLOT 1

9:30
Annex 1

"Easily" Obtain Low-Level Details of the Intel Compiler with the Intel VTune Performance Analyzer

The software pipeliner (SWP) employs modulo scheduling to achieve fast, compact loops by overlapping the execution of multiple iterations. How does the programmer know if software pipelining occurred (or not), and why? The compiler provides an optional set of optimization reports. Intel VTune Performance Analyzer 8.0 allows one to easily obtain this information from the compiler. This talk will describe how the Intel VTune Performance Analyzer displays the optimization reports and how to use these reports to achieve high performance on your applications.



Eric W. Moore - Intel

Eric Wynne Moore is a Senior Software Engineer working in the Software Products Division at Intel Corporation. In the past, he has worked at Rational, Microsoft, RealAudio, Digital, Compaq, and Keane. His specialties include operating systems, compilers, high-performance computing, and performance tuning. In the last couple of years, Moore has trained more than 500 engineers in performance optimization, including engineers at PNNL, CIA, FBI, IBM, Dell, Hewlett Packard, SGI, Cisco, Intel, several universities, as well as all over the world, including and around Korea, China, Brazil, and Europe.

SLOT 2

9:30
Annex 2

Cluster OpenMP

Cluster OpenMP* is software that virtualizes a cluster of machines to appear as a large SMP resource. This presentation will provide an overview of Cluster OpenMP and describe application characteristics that help predict scalability and performance when using the software. There will be a summary on migrating OpenMP applications to use Cluster OpenMP, as well as performance data from a 16-node Intel Itanium 2 processor cluster running Cluster OpenMP.

*Other names and brands may be claimed as the property of others.



Brock Taylor - Intel

Brock Taylor is part of the Advanced Computing Center at Intel working as an HPC Systems Engineer. He manages research and development clusters and works on multiple HPC storage and applications projects. Previously, Brock was a BIOS engineer responsible for initial power-on of Intel enterprise processors and chipsets, and he also spent a short time as a product BIOS engineer for Itanium workstations. Brock has a BS in Computer Engineering from Rose-Hulman Institute of Technology, Indiana, USA, and an MSc in High Performance Computing from Trinity College, Dublin, Ireland.

Bundle #10

SLOT 1

10:30
Annex 1

Optimizing Loops

An exposition of loop optimization methodology using the Itanium processor performance monitoring unit and the Intel VTune Analyzer. The discussion will focus on issues that are extremely difficult for compilers to resolve automatically. The importance of data prefetching and issues that can cause the default compiler prefetch algorithms to fail, and versioning of loops to assist the compiler in generating optimal instructions sequences are discussed. The SMG2000 code will be used for the purposes of illustration.



David Levinthal - Intel

David Levinthal has been involved in the Itanium processor development since the first version of the simulator based SDK. He is currently focusing on HPC performance analysis and optimization as a member of Intel's VTune Analyzer team. He joined Intel in 1995 working for the Supercomputing Systems Division and then the Microprocessor Software Labs. He has been the Lead Software Support Engineer on the Itanium Processor Family since 1997. Prior to joining Intel, he was a Professor of Physics at Florida State University. He has received the DOE OJI award, the NSF PYI award, and a Sloan Foundation Fellowship.

Proposal for Enhanced Open-Source Compiler Performance

In contrast to superscalar RISC processors which implement an out-of-order instruction execution model, the Itanium Processor Family (IPF) shifts the responsibility on the compiler to expose and exploit available instruction-level parallelism through aggressive code transformation and scheduling techniques. As such, Itanium compilers exert a big influence on delivered application performance. Although GCC performs very well in optimizing scalar codes, it lags behind commercial compiler offerings in optimization techniques employed for floating point intensive or loop intensive codes. While good progress is being made in implementing TREE_SSA based optimizations for GCC, an alternative, highly-leveraged, back-end compile path has the potential for delivering increased performance for loop-intensive codes in the short-term, and could be specifically targeted for the Montecito processor release.

The Open64 compiler is a 64-bit, open-source optimizing compiler that has been derived from the proven SGI MIPSpro production compiler. It has been ported to Itanium and was open sourced in 2000. Intel helped promote the use of this compiler to stimulate Itanium compiler research and christened it as the Open Research Compiler (ORC). ORC integrates the GCC front end, a robust middle end, and an optimized Itanium code generator. It could serve as the basis for a high performance open-source compiler for IPF in a relatively short time, and it can also be leveraged to quickly bring-up an alternative IPF-tuned backend compilation path for GCC that complements the evolving TREE-SSA-based GCC compiler optimization strategy.

In this talk, we advocate this complementary approach to implementing an IPF-tuned alternate GCC back-end to deliver the performance potential of the Montecito processor for the Gelato community and Linux/IPF GCC development community at large. To be clear, the long-term interests of GCC and its broad Linux user base are best served with a single high performance backend. In this talk we'll review a possible evolutionary path to integrating the mainstream GCC optimization path with the alternative backend with the support of the GCC developer community, drawing on the strengths of each.

Bundle #11**Interconnects and Latencies**

Scalability and performance of HPC applications running on clusters are often limited by the interconnects between nodes. This talk will review the advantages of employing high-performance interconnects in cluster computing environments. It will include an overview of RDMA and compare the bandwidth and latencies of various interconnects, including InfiniBand.

**Shin-Ming Liu - HP**

Shin-Ming Liu is the Project Manager for High-Level Optimization and GCC of the Itanium C/C++ Compiler Section of the Java, Compiler, and Tools Lab at HP in Cupertino, California. Liu led the development effort for the high-level optimization and code generator project in compiler targeted for the Itanium processor. In this project, he helped redesigned the high-level optimization into a highly-robust, scalable, and efficient component by rearchitecting the infrastructure, from which many new techniques were developed. Many highly-recognized programming analysis methods were adopted as well. Liu led the reinvention of compiler development methodology by focusing on modulization, memory footprint control, canonical internal representation, and automatic error detection. Before joining HP, he worked at MIPS/SGI in the area of compiler front end, middle end, back end, and linker. During that time, he co-authored several technical publications.

Brock Taylor - Intel

11:30
Annex 2

Transcendental Function Calculations on Itanium 2 - The Vector Math Library (VML) of Standard Functions

A Vector Math Library (VML) of 56 standard functions was implemented for Itanium. The functions are classified in six groups, as follows:

- 1 Reciprocal, Division, Square Root, Reciprocal Square Root;
- 2 Exponential, Logarithm, Power;
- 3 Trigonometric Functions;
- 4 Inverse Trigonometric Functions;
- 5 Hyperbolic Functions;
- 6 Inverse Hyperbolic Functions.

In all cases, the input consists of a vector of arguments (or two such vectors), and the output is a vector of results. The inner loops of these codes are "modulo scheduled" and, with one exception, floating point execution is 100% saturated. There is no penalty for unusual arguments or results: the inner loop code handles all corner cases, thus avoiding costly traps and external exception handling. Accuracy of results for single precision calculations is at worst 0.5002 ulps; for double precision calculations, 0.501 ulps. Speed of calculation ranges from 3.5 cycles/argument for the single precision reciprocal $y=1/x$, to 19.5 cycles/argument for the double precision power function $z=x^y$. Details of special coding techniques and strategies will be provided.



Clemens C. J. Roothaan - Secure64; University of Chicago

Clemens Roothaan is Professor Emeritus of Physics and Chemistry at the University of Chicago. In the 1950's, he published detailed algorithms to solve quantum mechanical movements of electrons in molecules and atoms. Today, most computer programs in this area are based on his method. After his retirement from the University in 1988, Roothaan started to work for HP Labs in Palo Alto, California. He has worked on the Itanium design team since 1990. Currently, Roothaan is working on a large software suite of scientific tools for function evaluation.

12:30

Lunch

13:30

End of Wednesday's Main Meeting - Board Bus for Hotel