

Gelato Federation Meeting

Understanding Your Itanium System to Maximize Performance

May 22-25, 2005, San Jose, California, USA

Host: HP at their Palo Alto and Cupertino Campuses

Sunday, May 22, 2005

18:30

Sponsor Hosted Dinner at the Eulipia in Downtown San Jose

Monday, May 23, 2005 - HP Palo Alto Campus

8:30 Bldg. 20 Auditorium **Registration**

8:45 Bldg. 20 Auditorium **Meeting Time**

HP Welcome



Dick Lampman - HP

Richard H. (Dick) Lampman leads HP's central research and development organization, HP Labs. He has responsibility for the research activities in six research labs around the world, with the dual objectives of providing leading-edge technologies to HP's current businesses and creating new technology-enabled opportunities for HP. Lampman is based at the main laboratory in Palo Alto, California.

State of the Gelato Federation

Mark will give an overview and status of the Federation's growth and 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.



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 has worked at several small software companies and brings a diverse background in business operations, start-up management, and technical insight to Gelato. He holds a PhD in Engineering and an MBA from the University of Illinois at Urbana-Champaign.

The Itanium Image: Cultivating an Informed Perspective

The Gelato community is doing a vast amount of outstanding work with Linux on the Itanium platform. Publicizing information about this work will help counteract some of the negative popular press that Itanium has received and in turn, accelerate the platform adoption rate. In this presentation, a panel of marketing and PR experts from HP, Intel, and SGI will talk about their efforts to positively promote Itanium and how the Gelato community can help.

Representatives from HP, Intel, and SGI

Commercial Deployment of a Large Scale Itanium System

We will review the purpose, strategy, and use of the Itanium systems in BP's High Performance Computing Center. Our HP and SGI systems today have over 3200 Intel Itanium CPUs, the storage systems exceed 650 TB. We will also discuss future plans and ideas.



Keith Gray - British Petroleum

Keith Gray is the Manager of high-performance computing at British Petroleum, where he is the architect for technical infrastructure for the company's Upstream Digital Business. He graduated from Virginia Tech with a degree in Geophysics and has spent 17 years within the oil and gas industry.

10:45

Morning Break

11:15 Bldg. 20 Auditorium

Meeting Time

Low-Level Details for High Performance Using Intel's Itanium Compiler

Two keys to obtaining high performance on IPF are managing the memory subsystem and optimally scheduling loops. In the Intel compiler, the High Level Optimizer (HLO) optimizes the memory hierarchy via prefetching, loop: unrolling, interchange, blocking, etc. The Software Pipeliner (SWP) employs modulo scheduling to achieve compact loops by overlapping the execution of multiple iterations. How does the programmer know what optimizations were performed (or not), and why?

The compiler provides an optional set of Optimization Reports for the different phases. This talk will describe the optimizations performed and how to use the Optimization Reports to achieve high performance on your applications.



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.

12:15

Lunch (1.5 hours)

13:45 Bldg. 20 Auditorium

Meeting Time

NASA Applications on the Altix

Update on GCC Improvements for the Itanium

Agenda:

Reviewing GCC Improvement Meeting
Rotating Register Support
Memory Disambiguation
SuperBlock Scheduling
Q&A



Wen-Mei Hwu - University of Illinois at Urbana-Champaign

Wen-mei W. Hwu is the Walter J. ("Jerry") Sanders - Advanced Micro Devices Endowed Chair in Electrical and Computer Engineering in the Coordinated Science Laboratory of the University of Illinois at Urbana-Champaign. From 1997 to 1999, Dr. Hwu served as the chairman of the Computer Engineering Program at the University of Illinois. His research interest is in the area of architecture, implementation, and software for high-performance computer systems. He is the director of the OpenIMPACT project, which has delivered new compiler and computer architecture technologies to the computer industry since 1987. For his contributions to the areas of compiler optimization and computer architecture, he received the 1993 Eta Kappa Nu Outstanding Young Electrical Engineer Award, the 1994 Xerox Award for Faculty Research, the 1994 University Scholar Award of the University of Illinois, the 1997 Eta Kappa Nu-Holmes MacDonald Outstanding Teaching Award, the 1998 ACM SigArch Maurice Wilkes Award, the 1999 ACM Grace Murray Hopper Award, the 2001 Tau Beta Pi Daniel C. Drucker Eminent Faculty Award, and the 2002 ComputerWorld Honors Archive Medal. Prof. Hwu holds four patents and is a fellow of IEEE and ACM. Dr. Hwu serves on the Executive Committee of the MARCO/DARPA Centers for Circuit and System Solutions and Gigascale Systems Research. Dr. Hwu received his PhD degree in Computer Science from the University of California, Berkeley.

15:45

Afternoon Break

16:15 Bldg. 20 Auditorium

Meeting Time

Montecito: A Look Inside the Next Itanium Processor

The Montecito delivers performance, power efficiency, and high reliability while maintaining full compatibility with the Itanium 2 processor. Montecito improves the Itanium 2 processor core with additional cache and a better organized memory hierarchy, multiple-thread and multiple-core capabilities, new RAS+M technologies, and a fundamental design shift in the area of power. The final result is a single die with 4x the contexts, 7x the cache, 70% of the power, 2.6x the power efficiency, and significant RAS+M capabilities compared to the original Itanium 2 processor. This presentation will cover the Montecito micro-architecture including details of the processor's pipeline, memory sub-system, and MT/MC capabilities, plus an overview of power management features (Foxton Technology) and RAS+M features (Silvervale Technology).

October Pre-Meeting Highlights

Brazil—the primordial, tropical paradise, the passion of Carnival, the immensity of the Amazon—truly a country of mythic proportions and also the venue for the next Gelato meeting! The Pontifical Catholic University of Rio Grande do Sul (PUCRS) will host us October 3-5 in Porto Alegre, Rio Grande do Sul, Brazil, with a focus on Linux Itanium research in Latin America and beyond. This will be a short presentation about the location and logistics of the next meeting.



Cameron McNairy - Intel

Cameron McNairy is an Architect for the Montecito program. Previous to Montecito, Cameron was a micro-architect for the Itanium 2 processor, contributing to its design and final validation. He plans to focus on performance, RAS, and system interface issues in the design of future IPF products. He came to the Itanium 2 team soon after its inception from performance work on the first Itanium processor.

Cameron received a BSEE and an MSEE from Brigham Young University. He is a member of the Institute of Electrical and Electronics Engineers.



César De Rose - 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 his doctoral degree 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.

17:30

End of Day 1 Meetings

Tuesday, May 24, 2005 - HP Cupertino Campus

The term "Bundle" is used to denote a parallel session. During parallel sessions, no specific breaks are scheduled. Refreshments will be available throughout the day. It is up to you to choose break times that work best for you.

8:30

Member Poster Session

9:30

Bundle 0

Maple Room

Intel Compiler Switches and Effects on Architecture

This lecture will introduce engineers to the Intel C++ Compiler 9.0 for Linux. It will cover some of the basic switches (O1, O2, O3) as well as some of the more advanced switches (Parallel, OpenMP, PGO, IPO). For each switch, there will be a short explanation on how that switch takes advantage of a specific Intel Itanium Architecture feature.



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, HP, SGI, Cisco, Intel, several universities, as well as all over the world, including and around Korea, China, Brazil, and Europe.

Madrone Room

Large-Scale Shared-Memory Parallelism

For the last eight years, Nasa Ames has been a partner with SGI in pushing the state of the art in shared memory parallelism on single system image machines, in the belief that shared memory provides simpler and faster methods of doing parallel applications. In this talk, we will discuss some of the results and techniques that have been used to take advantage of shared memory machines with hundreds of CPUs, and to achieve high levels of scaling on real world applications.

If time permits, the speaker will veer off and gratuitously rant about any of a number of other possible topics.

Bron Nelson - SGI

Bron Nelson is an SGI Software Engineer assigned fulltime to the Nasa Ames "Columbia" system, an installation of twenty 512 CPUs SGI Altix machines running Linux. Currently ranked #2 on the Top 500 list, the Columbia machines are used by a variety of researchers in aeronautics, cosmology, weather, chemistry, and other fields. Bron works with the researchers to help them achieve high performance on their codes; does trouble shooting and diagnosis when problems arise; and communicates back to R&D about what works, what doesn't, and what future directions would be useful.

Bron has been with SGI for the past 17 years. Prior to that, he worked for Cray Research and LLNL. He has an MS in Computer Science from UCLA, and a BA in Mathematics from UC Berkeley.

10:00

Bundle 1

Maple Room

Intel Compiler Switches and Effects on Architecture (Cont.)

Eric W. Moore - Intel

Madrone Room

Update on DIKU's Work on MySQL

Cache-conscious indexes, such as CSB+-tree, are sensitive to the underlying processor architecture. In this talk, we will focus on how to adapt the CSB+-tree so that it performs well on a range of different processor architectures, including Itanium 2.

Previous work has focused on the impact of node size on the performance of the CSB+-tree. We argue that it is necessary to consider a larger group of parameters in order to adapt CSB+-tree to processor architectures as different as Pentium and Itanium. We identify this group of parameters and study how it impacts the performance of CSB+-tree on Itanium 2. Finally, we propose a systematic method for adapting CSB+-tree to new platforms. This work is a first step towards integrating CSB+-tree in MySQL's heap storage manager.



Eric Jul - University of Copenhagen

Dr. Eric Jul is a Professor of Computer Science at the Dept. of Computer Science at the University of Copenhagen (DIKU) and heads the Distributed Systems Group. He is also the Director of the Danish Center for Grid Computing. Dr. Jul has been doing research in distributed systems for 30 years and was one of the principal developers of the Emerald distributed programming language as well as the main implementer of the Emerald Virtual Machine, which includes full, on-the-fly, fine-grained object mobility.

Performance Comparison Between GCC and icc on Itanium

Itanium requires extra effort with the compiler to extract performance. Traditionally, Itanium compiler teams have focused on profile-based optimization to extract the performance on Itanium machines. Unfortunately, most application developers use either the default optimization level of the compiler or -O2 with some caution. We collected the runtime performance generated by both GCC and icc on a set of applications that have not been tuned by either compiler. Such measurement is more likely to represent the out-of-the-box performance of the Linux-Itanium platform. It also helps us identify areas for performance tuning efforts.

**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 compilers targeted for the Itanium processor. In this project, he helped redesigned 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/SGL in the area of compiler front end, middle end, back end, and linker. During that time, he coauthored several technical publications.

Madrone Room

Work-in-Progress: Itanium Virtual Memory

Some of today's workloads use massive amounts of virtual address space and have very large working sets. For such workloads, handling TLB misses can take a significant amount of time. One solution to the problem is to use larger pages, so more memory is mapped with each TLB entry. However, using larger pages for everything means slower page fault times (as it takes longer to load a page from disk) and possibly wasted memory, where mapped objects are smaller than a page.

The Itanium architecture has two hardware-walked page table formats: a "short format" virtual linear array and a "long format" hash table. Linux currently uses the short format page table as a cache of the Linux page table. We have a patch to allow the long format page table to be used. Our measurements show that for most workloads changing the page table format doesn't affect processing times.

Because the hardware walked page table is not the same as the Linux page table, it's possible to change Linux's page table. We are currently just starting a project to provide an abstract page table interface, and then experiment with software page tables optimized for large address spaces, and for super page implementation.



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.

11:00

Bundle 3

Maple Room

HP Caliper: A Powerful Performance Tool, Now on Linux IPF

For four years, HP Caliper has offered HP-UX Itanium developers a powerful performance analysis tool that leverages the innovative Performance Monitoring Unit present on Itanium processors. Coinciding with the May 2005 Gelato meeting, HP Caliper becomes available to Linux IPF programmers. This talk will provide a quick overview of Caliper's capabilities, examples of its use and teasers for upcoming attractions.



Eric Gouriou - HP

Eric Gouriou is a Lead Engineer in the HP Caliper project. Before joining HP in 1999, Eric worked as a technical consultant for Oracle in France. Eric holds a master's of Computer Science from UCLA and a Diplôme d'Ingénieur from Ecole Centrale Paris.

Madrone Room

Advancing High-Performance Applications on Linux Itanium

The success of Linux on Itanium requires exploring new and improved methods for permitting efficient and scalable interaction among applications, libraries, and the operating system. We are examining existing concurrency mechanisms and profiling tools for the programming language C++; alternate interfaces, algorithms, data structures, and locking mechanisms to improve the efficiency of library and system calls; methods for reducing the number of kernel boundary crossings; and techniques for communicating information between applications and the kernel (currently achieved by system call or signals). The target applications are those that heavily rely on or whose performance is most affected by multi-threading and the interactions at the boundaries between languages, libraries, run-time systems, applications, and the operating system.

Therefore, the following fundamental software areas are being examined: concurrency, network I/O, and memory management, which form the core components of many applications. The progress of the C++ concurrency, profiling with hardware counters, TCP server performance, and application development will be presented.



Peter Buhr - University of Waterloo

Peter Buhr received BSc Hons, MSc, and PhD degrees in computer science from the University of Manitoba in 1976, 1978, 1985, respectively. He is currently an Associate Professor in the School of Computer Science at the University of Waterloo, Canada. He was a Research Scientist with SUN Microsystems Labs in 1993/4. His research interests include concurrency, concurrent profiling/debugging, persistence, and polymorphism. He is the principal designer and primary developer of both μ C++, a C++ high-level thread library, and μ Profiler, a profiling and debugging toolkit for μ C++. Dr. Buhr is a member of the Association of Computing Machinery.

11:30

Lunch (1 Hour)

12:30 Maple Room

Meeting Time

Gelato Federation Strategy Session

We will use this time to identify: 1) general areas of the software infrastructure which are lacking Itanium support, and 2) specific applications/tools within these areas that are needed for you to perform work on Linux-Itanium systems.

The session goal is to produce a short list of applications/tools (both open source and proprietary) that Gelato members can either: 1) port and optimize (open source), or 2) approach owning ISVs about providing Itanium support (proprietary).

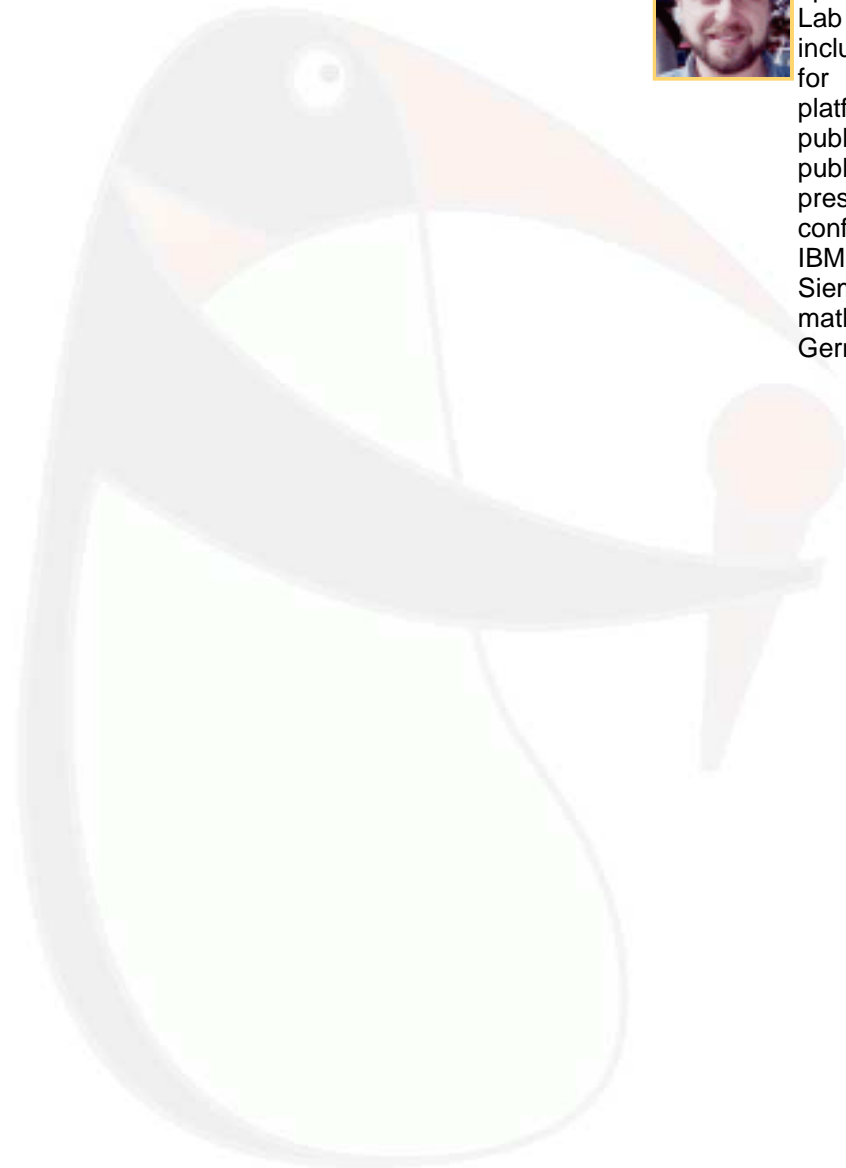
Mark K. Smith - Gelato Central Operations

Roundtable with Senior Itanium Compiler Engineers



Gerolf F. Hoflehner - Intel

Gerolf Hoflehner has been a code generation and optimization expert working in the Itanium Compiler Lab at Intel since 1999. His recent accomplishments include driving analysis and compiler optimizations for large enterprise applications on Itanium-based platforms, which resulted in record-setting TPC-C publications. He co-chaired two EPIC-4 workshops, published papers on compiler optimizations, and presented compiler tutorials and talks at various conferences. Prior to joining Intel, he worked on IBM370/390 and Mips R4K/R10K compilers at Siemens Pyramid. He earned a Diplom degree in mathematics from the Technical University Munich, Germany.



13:45

Afternoon Break

14:00 Maple Room

Meeting Time

Performance Profiling for Fun and Profit

The Itanium Processor Family includes a rich set of performance-monitoring hardware, which supports a series of unique optimization techniques. These include stall accounting, miss address profiling, and instruction address and path profiling. The focus will be on quickly locating performance opportunities in application code, but examples will also include tuning the Linux kernel. Details of the Montecito performance monitor will also be given as well as tips on how to use its advanced monitoring features.



Stéphane Eranian - HP

Stéphane Eranian is a Senior Research Scientist at HP Labs, where he has been working on the porting of Linux to the IA-64 platform since 1998. He has made numerous contributions to the Linux/IA-64 kernel and related user-level programs. He is the main architect of the Linux/IA-64 kernel performance monitoring subsystem (perfmon). He is also the creator of the pfmon tool, which uses this subsystem to collect performance information.

Before joining HP, Stéphane worked on his PhD at Chorus Systems (now Jaluna) in France. He holds a D.E.A. (BSc degree) in Operating Systems from Université PARIS 6, France, and a Doctorate (PhD degree) in Computer Science from Université PARIS 7, France. He is a member of USENIX and co-author of "IA-64 Linux Kernel: Design and Implementation."



Jim Callister - Intel

Jim Callister works for Intel's Fort Collins Design Center in Colorado, USA. In addition to working in firmware and CPU design, Jim has devoted most of his 23-year career to all aspects of CPU performance engineering. He participated in the architecture definition and evaluation of the Itanium architecture, with special emphasis on the performance monitor, and for the past nine years has been a member of the Itanium 2 processor design team. If he remembers correctly, Jim has degrees in Mathematics, Computer Science, and Electrical Engineering.

15:00

Bundle 4

Maple Room

Multithreaded Programming Issues for Application Programs

This presentation will review the often misunderstood ground rules for multi-threaded programming. We will concentrate on a Pthreads environment, though many of the issues are similar elsewhere. We will then look at two obstacles to following those rules. First, the rules themselves are neither completely clear nor completely sufficient. Second, there are occasions when programs that do not strictly follow the rules can greatly outperform those that do. These issues do not currently have completely satisfactory solutions on any platform. We will discuss both ongoing efforts to improve matters and partial solutions, particularly from an Itanium perspective.



Hans Boehm - HP

Hans Boehm is the primary author of a widely used open source garbage collector library. He was involved in the recent revision of the Java memory model, and is now participating in an effort to clearly define the meaning of multithreaded C++ programs. Boehm holds a PhD from Cornell University. He was on the faculty at the University of Washington and at Rice University before joining Xerox PARC, SGI, and finally HP Labs. Boehm has chaired several programming language research conferences and is the past chair of ACM SIGPLAN.

Madrone Room

The Vanilla Project

Everybody in the Gelato community knows the importance of optimizing your software to get the best efficiency from your Itanium processor. However, it appears this knowledge is far from being widespread outside Gelato. Furthermore, access to practical examples for using compiler options and tools related to performance monitoring seem difficult according to users. Enter Gelato Vanilla, whose aim is to spread the word "optimization" to the masses and to provide practical examples. In this presentation, we will present in detail the origin of the Vanilla project, what it is, and what it is not. In addition, we will review current as well as upcoming results.



Matthieu Delahaye - Gelato Central Operations

Although Matthieu Delahaye has worked on the Gelato portal since its creation in 2002, he officially joined Gelato Central Operations as a Software Engineer in August 2004. In addition to maintaining the Gelato portal, Matthieu works on Gelato Coconut, Gelato Vanilla, and other challenging infrastructure and development projects around the Itanium processor. Matthieu made his first kernel hacks while involved in the parisc-linux port effort, and then joined the Debian Project. At the same time, he received an MS in Computer Science from ESIEE, where he subsequently worked for two years in the IT Department.

David Moore - Intel

David Moore has been involved with programming languages and compilers for almost 30 years. He was a contributor to the Pascal language, the author of a Modula-2 compiler for CP/M 80 and other operating systems, part of a team that ported an Ada compiler and environment to UNIX, and a contributor to the National Compiler Infrastructure project. He is currently a member of the C++ Compiler Product Team at Intel with special responsibility for Linux support.

15:30

Bundle 5

Maple Room

Multithreaded Programming Issues for Application Programs (Cont.)

Hans Boehm - HP

Madrone Room

Focus on Grid

Agenda:

Update on GridAsia 2005 Gelato BoF
CERN LCG/EGEE Collaborations Using Itaniums
OurGrid Project



Jon Lao Khee Erng - National Grid Office

Jon Lau is the Assistant Head (Technical) at the National Grid Office (NGO) as well as the Technical Manager of the National Grid Pilot Platform (NGPP). He coordinates the technical issues of the NGPP and virtual grid communities, which span from network and security to middleware software. He developed the first Access Grid (AG) node in Singapore, and is promoting the deployment of more AG sites in Singapore through the demonstration of its benefits of AG-enabling conferences and meetings.

Prior to joining the NGO in January 2003, he was the Director of Engineering at eXage Private Limited, a high-tech spin-off from the Kent Ridge Digital Labs (KRDL), where he led the development team in designing a scalable architecture ready to evolve to meet the needs of eXage's customer.

Jon's technological experience is driven from both hardware interests and software R&D work at both the Information Technology Institute and KRDL. The many projects that Jon has been involved with include WinViz, a data visualization tool, as well as the Expert Advisory System on the Internet (a national project), where he performed the role of a Technical Manager.

Jon holds a bachelor's degree in Computing and a Master's in Technology, both from the National University of Singapore.

16:00

Bundle 6

Maple Room

Focus on Scalability in a Box

Agenda:
Overview
Progress Reports
Collaboration Opportunity Discussion
Action Items



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.

Madrone Room

AAOGS2 Itanium Grid: Practical Experiences in In-Silico Drug Discovery

Ligand-based drug discovery and structure-based drug discovery, including quantitative structure-activity/property relationship study, usually involves a large amount of data processing and numerical operations. An Itanium processor system is one of the economical approaches to get the efficiency of the massive calculations of the computational chemistry. And networking virtual computing based on grid computing is a challenging way to share and/or manage cooperative systems.

BMD has developed an in-house grid based on in-silico drug discovery solutions and has applied this to study structural similarities of proteins and molecules. The performance results of the calculation were documented. The demonstration of the system will be presented. Additionally, some ideas and concrete plans, harnessing large-scale memory to provide expansive opportunities of computer-aided drug discovery, will be suggested.



Jaesik Kwak - Research Institute of Bioinformatics & Molecular Design

Jaesik Kwak is the Director of a co-development project between BMD and IDRTEch, Inc., and is also a developer of the AAOGS Grid System. From 2002-2003, Kwak was an organizer of the Computational Chemistry and Nanomaterial Grid in the Korean National Grid Project.

Kwak received his master's degree at KAIST (Korea Advanced Institute of Science and Technology) in Quantum Chemistry. His current research areas include: development of total solution for computer-aided drug discovery, grid application to quantum chemistry for high throughput processing of drug candidates, high-performance chemical database system, and computational similarity research of chemicals. Kwak has implemented a grid system to share and integrate Itanium servers to perform high throughput processing of chemical data. He is working to improve the system and to modularize its functionalities and architecture to share with other people. He is also working on handling and tuning a large-scale database to use for drug discovery, adopting a large-scale memory model of the Itanium system.

16:30

Bundle 7

Maple Room

Focus on Scalability in a Box (Cont.)

Lee Schermerhorn - HP

Madrone Room

The Gelato Presence

Awareness about the Gelato Federation, our mission, and our members' work is on the rise. This presentation will cover Gelato's presence both inside the Linux Itanium community (via the Gelato portal) and outside the community (via PR initiatives). We will also cover ways in which members can participate and benefit from these outreach efforts.



Nan Holda - Gelato Central Operations

Nan Holda is the Marketing and Web Development Specialist for Gelato. Her main responsibilities include: increasing public awareness about Gelato and its mission, assisting Federation marketing strategies, overseeing portal development, keeping the portal up-to-date regarding Linux Itanium news and Federation Members' activities, and aiding conference exhibition and meeting coordination. Nan formerly worked at SourceGear Corporation as a Technical Support and Marketing Specialist for their version control software. She is also a freelance journalist for several local newspapers, and is a graduate of the University of Illinois at Urbana-Champaign (UIUC).

Madrone Room

Coconut

Gelato Central Operations will open remote-computing resources to Gelato members. The system we will use, COCONUT, stands for "Central Operations Computing Network for User Testing". This talk will address what COCONUT means for Gelato members.

Matthieu Delahaye - Gelato Central Operations

17:00 Depart Cupertino for HP Dinner at the Opera House in Los Gatos

17:30 HP Dinner at the Opera House in Los Gatos

Wednesday, May 25, 2005 - HP Cupertino Campus

The term "Bundle" is used to denote a parallel session. During parallel sessions, no specific breaks are scheduled. Refreshments will be available throughout the day. It is up to you to choose break times that work best for you.

8:30 **Bundle 8**

Maple Room

NUMA Scalability Issues and Locking Techniques

Effective locking is necessary to have satisfactory performance on large Itanium-based NUMA systems. Synchronization of parallel executing streams on NUMA machines is currently realized in the Linux kernel through a variety of mechanisms, which include atomic operations, locking, and memory ordering. However, the various synchronization methods may also be combined in order to increase performance. This talk will present the realization of basic synchronization in Linux on Itanium and then investigates more complex locking schemes.

Current locking mechanisms rely heavily on a simple spinlock implementation that may be fitting for systems of up to 8 CPUs. However, the existing spinlocks cause more and more cacheline bouncing during contention in the higher range of CPUs. The approaches that have so far been made to solve the contention issue are presented and then an implementation for Linux of an approach by, first proposed by Radovic called "Hierarchical Backoff Locks," will be discussed in detail.



Christoph Lameter - SGI

Since 2004, Christoph Lameter has worked at SGI as a Linux Developer for the Linux Core System Group, working on the Linux timer subsystem, page faults, and page tables as well as locking mechanisms.

Christoph's experience includes:

2004: PhD, Fuller Graduate School on Quantum Theory and Reality

2003: Infrant Technologies—scalability issues for embedded systems

2002: FoundryOne Corporation—scalable network infrastructures

1999-2001: SiteROCK Corporation—Open Source Ambassador

1999-2001: Board Member of Linux International

2000-today: Advisory Council Linux Professional Institute

1997: Elected to the Board of Directors of the Debian Project

1996-today: Member of the Debian Project

1994: Master of Divinity, Fuller Graduate School

1986: Master of Computer Science, University of Bremen on compiler technology

Madrone Room

Intel 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.

Eric W. Moore - Intel

9:30

Bundle 9

Maple Room

Scalability on Large SMP Linux Boxes with Linux 2.6.X

The Linux operating system has gained a very rapid acceptance in almost all segments of the software industry. It is used in a wide range of solutions from laptops to servers and large clusters. This operating system has been recognized as a good price-competitive solution for entry-level machines. There is now an increased interest to use it in some large SMP machines with 16+ processors. Some scalability issues were identified with the Linux 2.4 kernel. This discussion will present the state of the new 2.6 kernel and in particular, all improvements in terms of scalability for very large SMP machines.



Patrick Demichel - HP

Patrick Demichel has worked at HP since graduating from the Control Data Corporation (CDC) School in France in 1980. He worked on Real Time Executive (RTE), and then moved to UNIX in 1986 when HP introduced the series 500, 300, and later the 800. For four years, Patrick worked for the HP Workstation Division in Fort Collins, Colorado, on the Itanium project. Its activity was focused on software development and performance. Currently, Patrick works in Grenoble, France, for HP's Europe, Middle East, and Africa (EMEA) division in a group dedicated to HPC business. In particular, he is member of a team focused on developing a program called Bigtux to make standard Linux versions available and tuned for the IA-64 Superdome in a single 64-way partition, accessing some very large memory and very large I/O resources.

Optimizing Scientific Libraries for the Itanium

One of the key operations in many scientific applications is the computation of standard mathematical functions. These include algebraic operations (division, square root), elementary transcendental functions (exp, log, sin etc.), and more exotic transcendental functions (Gamma, Bessel etc.). Both Intel and HP have expended considerable effort in providing a library of accurate and efficient mathematical functions for the Itanium architecture. These can be used in place of the generic ones provided by GCC or other C compilers, even without using our compilers.

The internals of these functions are an interesting case study in how certain features of the Itanium architecture lend themselves to efficient and accurate floating-point computation. This is mostly a consequence of good architecture design—for example the combination of parallelism and extended precision permits the highly-efficient evaluation of long polynomials. But at least one case is a lucky accident: the "frcpa" reciprocal approximation, intended for use in division, can be used in a highly-efficient argument reduction step for the logarithm function.

We will give an overview of some of the highlights of the implementation of these functions, with the emphasis on where special features of the Itanium architecture give it an advantage over others such as IA-32 and Power.

Recommended Background Reading:

Markstein: "IA-64 and Elementary Functions: Speed and Precision," Prentice-Hall 2000.

Harrison, Kubaska, Story and Tang: "The Computation of Transcendental Functions on the IA-64 Architecture," Intel Technology Journal Q4 1999 (http://www.intel.com/technology/itj/q41999/articles/art_5.htm).

Cornea: "Proving the IEEE Correctness of Iterative Floating-Point Square Root, Divide and Remainder Algorithms," Intel Technology Journal Q2 1998 (http://www.intel.com/technology/itj/q21998/articles/art_3.htm).

**John R. Harrison - Intel**

John Harrison has worked in formal verification and automated theorem proving since 1990, when he joined Mike Gordon's "Hardware Verification Group" (HVG) at the University of Cambridge Computer Laboratory. As well as working on the development of the HOL theorem prover, he developed a particular interest in the formalization of real analysis and its application to formal verification of floating-point hardware.

After completing his PhD research in 1995, Harrison spent a very enjoyable year at Åbo Akademi University and Turku Centre for Computer Science (TUCS) in Turku, Finland, where he was a member of Ralph Back's Programming Methods Research Group. Harrison then returned to Cambridge and worked on a formal model of floating-point arithmetic and its application to the verification of some realistic algorithms for transcendental functions. This work attracted the attention of Intel, and in 1998 Harrison joined the company as a Senior Software Engineer specializing in the design and formal verification of mathematical algorithms. He has formally verified and in many cases designed or redesigned numerous algorithms for mathematical functions including division, square root and trigonometric functions.

In his limited spare time over the past 10 years, Harrison has been working on a book giving a comprehensive introduction to automated theorem proving. He hopes that this book will finally reach publication in 2005, and the associated code is already available from his Web page.

Bundle 10**Scalability on Large SMP Linux Boxes with Linux 2.6.X (Cont.)****Patrick Demichel - HP**

Madrone Room

Chasing Java Performance on Itanium

Java's write-once, run-anywhere design makes it one of the most popular software development environments today. A consistent set of APIs across different platforms makes it relatively painless to debug and run a Java application on a new platform. However, a proliferation of JVMs means different performance tuning opportunities from the different providers. In this talk, we will discuss the various JVMs available for Itanium, especially for Linux on Itanium, and do some baseline performance characterization for them based on standard benchmarks. Then we will talk about the differences in performance tuning mechanisms on these and the best ways to get the application to perform well. We will also look at some performance tuning tools available for the platforms to help this process.



Virendra Mehta - HP

Virendra Mehta has worked at HP for 10 years and is currently a Senior Architect in the Java, Compilers, and Tools Lab of the Infrastructure Solutions Division. Previously, he worked at the IBM T. J. Watson Research Center. Mehta has worked with languages and compilers for 15 years. In particular he has worked on static and dynamic compilers for Itanium for a long time. Lately he has been working on extracting Java and J2EE performance on Itanium systems. He started a Java performance project which improved Itanium JVM performance several fold and established Itanium as the performance leader with SPECjbb2000 and SPECjAppServer2002 benchmarks.

Mehta received a BTech in Computer Engineering from IT-BHU, India, and an MS in Computer Science from Stanford University, USA.

10:30

Bundle 11

Maple Room

Scalability on Large SMP Linux Boxes with Linux 2.6.X (Cont.)

Patrick Demichel - HP

Madrone Room

OpenIMPACT Compiler Update

Since the Beijing meeting, work on the OpenIMPACT compiler has concentrated in two areas: C++ support and improving compile time. We will present a status update on our C++ work and some interesting methods we have developed to reduce compile time. We will also present some detailed information on the effects of control flow specializing optimizations on code performance. This has particular relevance to GCC, as the OpenIMPACT team will be porting such optimizations to the open-source compiler.



Robert Kidd - University of Illinois at Urbana-Champaign

Robert Kidd is a Compiler Engineer in the OpenIMPACT team at the University of Illinois at Urbana-Champaign (UIUC), where he is responsible for developing a new front-end library to support future compiler research. He has also worked on back-end enhancements, as well as scripting drivers for GCC compatibility and compiler usability. Robert has a BS in Computer Science with Highest Honors from UIUC.

11:00

Bundle 12

Maple Room

Preview of the Upcoming Open|SpeedShop Performance Debugger

Open|SpeedShop is SGI's next generation Linux performance tool. Initially targeted to support performance analysis of applications running on the SGI Altix platform, Open|SpeedShop is based on the concepts of SGI's IRIX SpeedShop. Open|SpeedShop is being co-funded by the Department of Energy (DOE) and will have the infrastructure and base components released as open source under the GPL and LGPL licenses.

Open|SpeedShop is designed to be modular and easily extendable. It supports the concept of plugins which allow users to create their own performance experiments. Another key feature of the performance tool is its usability. Its user interface is designed for scientists in general, not just computer scientists.

The Open|SpeedShop baseline functionality will include support for: single system image (SSI) machines, for clusters (i.e., multiple OS kernels), exclusive and inclusive user time, program counter (PC) sampling, MPI call tracing, input/output tracing, floating point exception tracing, and CPU hardware performance counter experiments. SGI will also produce an enhanced Pro version by creating advanced Altix specific plugins.

Open|SpeedShop on Linux platforms will enable FORTRAN (77, 90, and 95), Open|C, and C++ programmers to use an advanced performance analysis tool within the C++ open-source environment.



Jack Carter - SGI

Jack Carter has spent the last 20 odd years working on compilers and compiler tools. Much of this has been developing linkers for both native and embedded systems along with extensive work developing object transformation tools used for performance debugging, code coverage, and memory simulation input. Jack's current work is with the SGI Open|SpeedShop project.

Itanium-Related Research Activities at CERN

This presentation will review the Itanium-related activities inside the CERN openlab, where intensive work has been done in the areas of grid-enabling, C++ compiler optimization, high-speed networking, and high-speed data export to labs in other countries. A planned project to perform high-speed streaming of data to tape will also be presented. Furthermore, CERN is currently doing some work in the area of virtualization (based on Xen) and the potential for using virtualization in grids will be discussed.



Sverre Jarp - CERN

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.

Improving GCC Instruction Scheduling for Itanium

The instruction scheduler is one of the weak points of GCC on the IA-64 architecture. This presentation will be devoted to preliminary results of an ongoing project. The project's goal is improving GCC instruction scheduling for Itanium. We address the following scheduler problems: absence of support for IA-64 control and data speculation, weak alias analysis, and primitive basic block probability evaluation. We add support for control speculation using parts of existing framework for interblock motions. As for data speculation, we suggest generating speculative loads without recovery code first, then allowing interblock load motion, and adding support for the recovery code.

To benefit from added speculation support, usage of more accurate alias information in the scheduler is a must. The current scheduler calculates dependencies based on RTL alias analysis. The analysis cannot disambiguate most memory references on Itanium due to lacking of "base + displacement" addressing mode. The solution we suggest is: 1) propagation of name memory tag information from tree-ssa to RTL for pointers, and 2) tracking pointer arithmetic, which is performed for addressing non-pointer variables, in `alias.c`.

The scheduler uses its own primitive algorithm of evaluation of basic block probabilities. This can be improved by using GCC standard probability analysis or profile information, if available.

**Arutyun I. Avetisyan - Russian Academy of Science**

Arutyun Avetisyan is the Deputy Director of the Institute for System Programming (ISP) of the Russian Academy of Sciences (RAS) in Moscow, Russia. Avetisyan is the leader of projects involved with parallel program development systems and cluster management. His research focuses include parallel and distributed programming, cluster technologies, compiler technologies, and software quality. Current compiler-related work includes C++ compiler optimization.

The LRZ Linux-Itanium Experience

LRZ is presently in the process of migrating its complete line of HPC systems to Intel Itanium based architectures. This talk gives an overview of the migration process, first experiences with Itanium in HPC, and the requirements for the top-end system to be installed at LRZ in 2006.



Reinhold Bader - Leibnitz Computing Centre

Since 1999, Reinhold Bader has been a member of the scientific staff at Leibniz Computing Centre (LRZ) in Munich, Germany. He is responsible for benchmarking for the purpose of system procurements, libraries, tools, and documentation on HPC systems, user support, and optimization of codes on the LRZ HPC platforms.

Bader attended a physics diploma course at Ludwig-Maximilian University in Munich 1985-89 and 1990-92, including diploma thesis in theoretical solid state physics, and final exams. The course was interrupted by an 8 month industrial training course completed at Rutherford Appleton Laboratory, England from September 1989 to April 1990. For his PhD studies, Bader worked with Prof. H. Bross at Ludwig-Maximilian University in Munich, Germany from 1993-98, and completed his PhD thesis magna cum laude on the electronic and structural properties of tetrahedral semiconductors.

12:00

Bundle 14

Maple Room

HP's Math and Message Passing Libraries

This talk will describe the development process, and business model used by HP to develop its high-level math library, HP MLIB, and its message passing library, HP-MPI. Features and advantages of these libraries will also be described. In the interactive portion of the talk, I would like to learn from Federation members about their library development and usage patterns. Topics for discussion include open-source alternatives, standards organization participation, and other commercial products.



Steven Rowan - HP

Steve Rowan is an Engineering Manager in HP's High Performance Computing Division. His department is responsible for HP's message passing library, HP-MPI. HP-MPI is supported on Itanium, X86-64, PA-RISC, and Alpha platforms. The product supports numerous interconnects on Linux, HP-UX, and Tru64 Unix. In addition to HP-MPI, the department is also responsible for HP's high-level scientific math library product, HP MLIB. HP MLIB runs on Itanium Linux and HP-UX. In addition to libraries, the department is responsible for HP's Unified Parallel C compiler, the HP-UX FORTRAN compiler, and the HP-UX cluster tool suite, ClusterPack.

Before taking his current position, Steve was responsible for various compiler products and components at HP and Convex Computers. Projects included high-level optimization, FORTRAN, C, C++, and Ada compilers.

Madrone Room

HPCToolkit

HPCToolkit is an open-source, multi/cross-platform tool set designed to improve productivity in identifying and diagnosing performance problems for large multi-image, multi-language programs. It helps the analyst to collect multiple profiles (primarily from hardware performance counters), compute derived metrics from the measured data, produce a hierarchy of program performance views, and hyperlink the data to whatever program sources are available. We will present the design of the current toolkit using Itanium examples, and we will discuss our current research directions.



Robert Fowler - Rice University

Robert J. Fowler is a Senior Research Scientist and Associate Director of the Center for High Performance Software Research at Rice University. His education includes an AB in Physics from Harvard (1971) and MS (1981) and PhD (1985) degrees from the University of Washington. His research interests are in the area of high-performance distributed and parallel computing. Specific interests include compilers and programming environments, architectures, operating systems, performance evaluation, and simulation.

12:30

Lunch and Wrap Up Session