



MULTICORE WORLD 2013

19 - 20 February 2013, Wellington, New Zealand

PROGRAMME

Multicore World: a destination conference to discuss all things multicore at peer-peer level

All modern computers are now multicore and increasingly so; we need the programming and application expertise to harness this current transformation. This is a permanent change that affects all computing.

Multicore World is a content intensive event that will allow you to network with other professionals from industry, academia, science, engineering and software communities, business and government. You will take the pulse of what is happening in the multicore processing and parallel programming ecosystem in a unique destination conference.

Sessions about the leading edge of multicore and parallel processing

Learn about all aspects of multicore technology and business. Sessions and panel discussions will cover the latest developments in software and hardware, applications and use cases, and the businesses behind these trends. **Confirmed keynote guest speakers:**

- **How to Grow the Economy by 10% per Year.**
Tuesday 19 - 9:00 am
Professor Ian Foster - Arthur Holly Compton Distinguished Service Professor of Computer Science at the University of Chicago & Argonne Distinguished Fellow at Argonne National Laboratory. USA
- **Bare-Metal Multicore Performance in a General-Purpose OS.**
Tuesday 19 - 11:00 am
Paul McKenney - IBM Distinguished Engineer & Linux CTO, IBM Academy of Technology. USA
- **The (Massive) Opportunities of a Multicore World.**
Tuesday 19 - 3:30 pm
Professor Barbara Chapman - OpenMP Architecture Review Board & University of Houston. USA
- **Massive Parallel Processing in the Varnish Software.**
Wednesday 20 - 9:00 am
Poul-Henning Kamp - Chief Architect, Varnish Software & Author (of a lot of) FreeBSD. Denmark
- **Transactional memory hardware, algorithms, and C++ language features.**
Wednesday 20 - 11:00 am
Mark Moir - Oracle, Consulting Member of Technical Staff. USA - New Zealand
- **The Revolution in Your Pocket: Invisible Computers with Dissociative Identity Disorder.**
Wednesday 20 - 4:00 pm
Tim Mattson - Intel Principal Engineer, Khronos OpenCL Group. USA

Multicore World 2013 Activities

Monday 18 February

- Workshop: "Introduction to Parallel Programming with OpenMP". Tim Mattson (Intel)
- Speakers' Welcome Cocktail

Tuesday 19 February - Conference Day 1	Wednesday 20 February - Conference Day 2
<ul style="list-style-type: none"> • Full day sessions and Panel • Conference Dinner - Wellington Town Hall 	<ul style="list-style-type: none"> • Full day sessions and Panel • Conference Closure - Drinks and Nibbles - Wellington Town Hall

Thursday 21 and Friday 22 February

- Follow-up business meetings and networking - NZTE

Multicore World 2013 is an initiative of

Open Parallel
Software for Multicore and Parallel Computing




MULTICORE WORLD 2013

SCHEDULE

Tuesday 19 February	Wednesday 20 February
8:50 - 9:00 Conference Opening	8:55 - 9:00 2 nd day opening
9:00 - 9:45 KEYNOTE - Prof Ian Foster (Argonne Labs / University of Chicago)	9:00 - 9:45 KEYNOTE - Poul-Henning Kamp (FreeBSD / Varnish)
9:55 - 10:05 Joe Gamman (CCANZ)	9:55 - 10:05 Vikas Singh (University of Auckland)
10:10 - 10:25 Nicolas Erdody (Open Parallel)	10:10 - 10:25 David Eyers (University of Otago)
<i>10:30 - 11:00 Morning tea</i>	<i>10:30 - 11:00 Morning tea</i>
11:00 - 11:45 KEYNOTE - Paul McKenney (IBM)	11:00 - 11:45 KEYNOTE - Mark Moir (Oracle Labs)
11:55 - 12:25 Simon Spacey (Imperial College / University of Waikato)	11:55 - 12:25 Richard O'Keefe (University of Otago)
<i>12:30 - 1:30 Lunch</i>	<i>12:30 - 1:30 Lunch</i>
1:30 - 1:45 Andrew Ensor (AUT)	1:30 - 1:55 Dave Fellows & Ivan Towlson (GreenButton)
1:50 - 2:15 TN Chan (Compucon)	2:00 - 2:15 Christian Rolf (Corvid / Open Parallel)
2:20 - 3:00 Tim Mattson (Intel)	2:20 - 3:00 Panel - Killer App (Poul-Henning Kamp, Tim Mattson, Barbara Chapman, Paul McKenney)
<i>3:00 - 3:30 Afternoon tea</i>	<i>3:00 - 3:30 Afternoon tea</i>
3:30 - 4:15 KEYNOTE - Prof Barbara Chapman (University of Houston)	3:30 - 3:55 James McPherson (Oracle)
4:25 - 5:00 Panel - Winners (Ian Foster, Paul McKenney, Mark Moir, Poul-Henning Kamp)	4:00 - 4:45 KEYNOTE - Tim Mattson (Intel)
	4:45 - 5:00 Conference Closure
<i>5:00 - 6:30 Free</i>	<i>5:00 - 6:30 Cocktail</i>
<i>6:30 - 10:00 Open Parallel - Conference Dinner</i>	

Open Parallel
Software for Multicore and Parallel Computing



MULTICORE WORLD 2013

19 - 20 February 2013, Wellington, New Zealand

PROGRAMME

(as 4 February 2013)

Monday 18 February, 12:30 - 5:00 pm - Majestic Centre Level 15, New Zealand Trade & Enterprise (NZTE)

Workshop: "An introduction to Parallel Programming with OpenMP"

Tim Mattson, PhD, Principle Engineer and Parallel Computing Evangelist, Intel Corporation, USA

Outline

- Introduction to parallel programming
- An introduction to OpenMP
- Creating threads
- Basic synchronization
- Parallel loops (introduction to worksharing)
- The rest of worksharing and synchronization
- Data environment
- OpenMP tasks
- The OpenMP memory model
- A survey of parallel programming models

Hands on workshop: Bring your **own laptop** with an OpenMP compiler (anything from Apple with X-code or any Linux-laptop with a modern GNU compiler, Windows laptops with Microsoft Visual Studio, or any laptop/OS with an Intel Compiler. Windows plus cigwin and a modern GNU compiler works as well. Note that we will not have time to help you install your programming environment, please do this in advance.)



Tuesday 19 February - Conference Day 1

Wellington Town Hall

8:50 - 9:00 - *Conference Opening*

9:00 - 9:45 - **KEYNOTE**

How to grow the economy by 10% per year

Professor Ian Foster, Arthur Holly Compton Distinguished Service Professor of Computer Science at the University of Chicago & Argonne Distinguished Fellow at Argonne National Laboratory, USA

Abstract - Economic growth is driven by the evolving interplay between innovation and automation: the former providing new products and services and the latter enabling cost-competitive and timely production and delivery. Information technology plays an increasingly central role in this process - a role that is only going to accelerate in the next decade as a result of advances in cloud and multicore computing. I discuss the challenges and opportunities inherent in a hyper-connected, hyper-automated world, with a particular emphasis on what they mean for the wonderful country of New Zealand.

Bio - Ian Foster is the Arthur Holly Compton Distinguished Service Professor of Computer Science at the University of Chicago and an Argonne Distinguished Fellow at Argonne National Laboratory. He is also the Director of the Computation Institute, a joint unit of Argonne and the University. His research is concerned with the acceleration of discovery in a networked world. Foster was a leader in the development and promulgation of concepts and methods that underpin grid computing. These methods allow computing to be delivered reliably and securely on demand, as a service, and permit the formation and operation of virtual

organizations linking people and resources worldwide. These results, and the associated Globus open source software, have helped advance discovery in such areas as high energy physics, environmental science, and biomedicine. Grid computing methods have also proved influential outside the world of science, contributing to the emergence of cloud computing. His new Globus Online project seeks to outsource complex and time-consuming research management processes to software-as-a-service providers; the goal here is to make the discovery potential of massive data, exponentially faster computers, and deep interdisciplinary collaboration accessible to every researcher, not just a select few "big science" projects.

Dr. Foster is a fellow of the American Association for the Advancement of Science, the Association for Computing Machinery, and the British Computer Society. Awards include the British Computer Society's Lovelace Medal, honorary doctorates from the University of Canterbury, New Zealand, and CINVESTAV, Mexico, and the IEEE Tsutomu Kanai award.

9:55 - 10:05

Big data and Networked Infrastructure

Joe Gamman - Education and Development Manager, Cement & Concrete Association of New Zealand (CCANZ)

Abstract - Infrastructure assets are designed to typically last decades with many existing assets approaching 100+ years of operation with no foreseeable decommission inviting the general description 'perpetual asset'. However, due to the large finance requirements infrastructure investments are 'lumpy' over a generational timescale and can favour short term construction cost minimisation over longer term lifetime maintenance cost. Maintenance on existing assets can also be deferred seemingly indefinitely - unfortunately, this can't go on forever and in 2009 the American Society of Civil Engineers for instance, rated their national infrastructure "D" and estimated a required 5-year investment of \$2.2 trillion.

Considering the lifetime maintenance costs can approach the initial construction cost there is a push in research toward 'structural health monitoring' - a system that allows active management of infrastructure assets whereby the asset itself can be queried for status and maintenance and inspection can be prioritised rather than the traditional 'brute force' approach of manual inspections on a regular basis. Currently, a number of individual pieces of infrastructure around the world have had sensors included in their build but asset owners are still struggling to understand how to use this data as anything other than a piece of research.

In this talk, we discuss the aspirations of the recently funded New Zealand 'Networked Infrastructure' project. In brief, the project is investigating the feasibility of designating the Christchurch rebuild (a 30-year, \$30bn central city construction project) as the world's demonstration site for sensors in the built environment. We believe that this is a unique opportunity for the world to field test the expectations of a massively networked built environment that is likely to be the default standard in 2050.

Over and above the challenges in sensor development, communications and construction engineering, the opportunity is very much one of understanding real-time large data sets and converting a tsunami of data into key pieces of asset management advice for users. We hope that by bringing the world to New Zealand in this niche area of big data, we can develop in parallel an ecosystem of technology and software development creating niche products for the world leveraging the credibility of large multi-national companies demonstrating their technical prowess. I will describe the aims of the project, the project design and our hopes for 2013. We also believe the talk will be of particular interest to established big data players and large enterprise management software companies.

Bio - Joe describes himself as a 'bi-lingual' technical and commercial manager. Joe's PhD and early career were focused on developing clean low-carbon distributed electricity devices. However, an understanding of a technology is limited without a similar understanding of the economic and policy frameworks that also exist: innovation doesn't occur in a vacuum, but an innovation can certainly become stranded if it fails to address, or even acknowledge non-technical systems and incentives. While Joe's background is in clean energy research and management and most recently the construction sector, his interests tend toward investigating problems and opportunities at the boundaries of standard disciplines. Information technology is transforming every sector and as the saying goes: if your business hasn't been transformed by the internet yet, it will be.

How this plays out in the construction space is by no means settled. Built environment information management tools are the new New Thing but are plagued by walled gardens, closed protocols and vague vendor promises of value. The idea of using the Christchurch re-build to launch a self-sustaining ecosystem of demonstration and innovation in the built environment is ambitious to say the least. In his talk Joe will be highlighting some of the major themes that are emerging and highlight how New Zealand is placed to take advantage of the opportunity.

10:10 - 10:25

Open Parallel: you are not alone

Nicolas Erdody - Director & Founder, Open Parallel

Abstract - An overview of Open Parallel's projects, from the integration of Intel's Threading Building Blocks (TBB) with Facebook's HipHop, to SKA (Square Kilometre Array radio-telescope) middleware scaling. The talk will present capabilities and track record plus current projects including functional

programming in relation to Big Data and the Financial Industry

Bio - Nicolás Erdödy is a serial high tech entrepreneur and former maths lecturer who founded and implemented 15 knowledge-intensive start ups in different countries and industries -from a venture specialised in algorithms for images' super-resolution to the 1st online academy in LatinAmerica (yes, a 1999 ".dot com"). He spent decades establishing and leading highly specialised teams, and designing, integrating and managing systems for new ventures - creating them from scratch. Nicolás is also Director of Erdödy Consultancy Ltd, a professional services firm that specialises in Strategy, Innovation and Corporate Entrepreneurship

Following the vision that New Zealand could become a global hub in entrepreneurship based in multicore technologies, Nicolás has been organising a yearly friends gathering (aka conference in multicore and parallel computing) since 2010, when he founded Open Parallel. Previously he founded and was CEO of World45, the first multicore software company established in New Zealand. Prior to that, Nicolás established dGV, a boutique venture capital firm specialised in early stage high tech ventures with strong IP position. He was dGV's Venture Capital Manager and CEO of two of the invested companies and Director of all of them.

Nicolás holds a Master of Entrepreneurship from the University of Otago, New Zealand; a Research Diploma for New Technologies applied to Education from INRP (Uruguay - France) and forgot long time ago the FORTRAN learned between Mathematics and Hydraulics at the School of Engineering of Universidad de la Republica, Montevideo, Uruguay. He knows how to ask for a beer in five human languages.

10:30 - 11:00 - Morning Tea

11:00 - 11:45 - KEYNOTE

Bare-Metal Multicore Performance in a General-Purpose OS

Paul McKenney, IBM Distinguished Engineer & Linux CTO, IBM Academy of Technology. USA

Abstract - A constant refrain over the decades from database, high-performance computing (HPC), and real-time developers has been: "Can't you just get the kernel out of the way?". Recent developments in the Linux kernel are paving the way to just that ideal: Linux is there whenever you need it, but if you follow a few simple rules, it is completely out of your way when you don't need it.

This adaptive-idle approach will provide bare-metal multicore performance and scalability to databases as well as to HPC and real-time applications. However, it is at the same time able to improve energy efficiency for upcoming asymmetric multicore systems, allowing these systems to better support workloads with extreme peak-to-mean utilization ratios. This talk will describe how this feat is accomplished and how it may best be used.

Bio - Paul E. McKenney has been coding for almost four decades, more than half of that on parallel hardware, where his work has earned him a reputation among some as a flaming heretic. Over the past decade, Paul has been an IBM Distinguished Engineer at the IBM Linux Technology Center, where he maintains the RCU implementation within the Linux kernel, dealing with a variety of workloads presenting highly entertaining performance, scalability, real-time response, and energy-efficiency challenges.

Prior to that, he worked on the DYNIX/ptx kernel at Sequent, and prior to that on packet-radio and Internet protocols (but long before it was polite to mention Internet at cocktail parties), system administration, business applications, and real-time systems. His hobbies include what passes for running at his age along with the usual house-wife-and-kids habit.

11:55 - 12:25

Remote Procedure Call (RPC) considered harmful

Simon Spacey, University of Waikato, New Zealand

Abstract - In the 1980's architectures consisting of two computational components communicating through a loosely-coupled network were considered advanced. For these architectures the Remote Procedure Call (RPC) communication paradigm is efficient and many authors concerned with the complexities of computational partitioning on distributed architectures quickly adopted RPC as their standard program communication paradigm which led to RPC's firm embedding in the Client-Server, Custom Instruction and Shared Memory implementation libraries that we take for granted today.

Unfortunately though, RPC is not efficient for modern Distributed and High-Performance Computing (DHPC) architectures which invariably include more than two computational components or tightly-coupled busses. In this presentation I explain the critical problems that make RPC inefficient for modern DHPC systems, introduce a possible solution to these problems called the Write-Only Architecture (WOA) and provide results and formal bounds showing the performance improvements deliverable for both homogeneous and heterogeneous DHPC systems over current RPC based implementations.

Bio - Simon Spacey (PhD - Imperial College, London) is an ex-missile scientist with decades of professional IT and Computer Science experience. Simon is a Senior Lecturer of Computer Science at the University of Waikato and provides strategic consulting to organisations around the globe. His research interests centre around Computational Optimisation and he currently lectures topics in Software Engineering, Computer Systems and DHPC Architectures

12:30 - 1:30 - Lunch - Civic Suites 1&2

1:30 - 1:45

Exascale Computing in the SKA

Andrew Ensor - Auckland University of Technology, New Zealand

Abstract - The Square Kilometre Array Project will require exascale computing capabilities and unprecedented levels of parallelism. The project is entering its pre-construction phase, where a range of potential computing technologies will be investigated, including highly multicore CPU, GPU, heterogeneous, and custom systems. Each will have its performance and power consumption modelled, and it is anticipated new approaches and digital platforms may need to be developed to find feasible solutions to the demands of the SKA.

This talk will give a brief background of the SKA and its processing requirements, and outline some of its computing developments that are hoped will be centred in New Zealand.

Bio - Dr Andrew Ensor received a BSc(Hons) from the University of Auckland and a PhD from the University of California at Berkeley in 1995. After three years at the Università di Siena in Italy he returned to New Zealand as a Senior Research Lecturer at AUT University in Computer Science and Mathematics.

His research and industry experience includes distributed and mobile systems, GPU computing, algorithms and concurrency. Andrew is the computing lead for the NZ SKA Open Consortium, a partnership of New Zealand universities and industries that will be tackling some of the computing challenges of the SKA.

1:50 - 2:15

Heterogeneous Parallel Computing with Kepler and CUDA5

TN Chan - Compucon New Zealand

Abstract - This paper takes a system integration engineering perspective for industry technology transfer purposes. The context is high performance computing focussed on heterogeneous and CUDA, and it is more hardware architecture than software application based. Themes discussed include heterogeneous versus homogeneous approach, high level versus low level compilation, the latest advances of CUDA eco-system, and the differences between digital content and concept creation.

Bio - TN Chan was born and educated in Hong Kong with a Mechanical Engineering honours degree from the University of Hong Kong. He migrated to Wellington in 1985 and worked for Electricorp New Zealand till he moved to Auckland to start up Compucon New Zealand in 1992. He is a Chartered Electrical Engineer, professional member of the Institution of Engineering & Technology, and professional member of the Institution of Professional Engineers of New Zealand. He designed the introduction of the first locally built dual processor RAID server in New Zealand in 1998. He is an Industry Supervisor of the University of Auckland since 2002. His current roles include quality assurance management, project management, technology transfer, and maintaining Compucon New Zealand as a centre of excellence in serious computing.

2:20 - 3:00

Programming heterogeneous computers: a look at OpenCL, CUDA, OpenACC, and the OpenMP accelerator directives

Tim Mattson, Intel, USA

Abstract - I will discuss the evolution of heterogeneous platforms and then survey programming models to use with them. This would be a look at OpenCL, CUDA, OpenACC, and the OpenMP accelerator directives

Bio - Tim Mattson is a parallel programmer (Ph.D. Chemistry, UCSC, 1985). Tim has been with Intel since 1993 where he has worked with brilliant people on great projects such as: (1) the first TFLOP computer (ASCI Red), (2) the OpenMP API for shared memory programming, (3) the OpenCL programming language for heterogeneous platforms, (4) Intel's first TFLOP chip (the 80 core research chip), and (5) Intel's 48 core, SCC research processor.

Tim has published extensively including the books *Patterns for Parallel Programming* (with B. Sanders and B. Massingill, Addison

Wesley, 2004), *An Introduction to Concurrency in Programming Languages* (with M. Sottile and C. Rasmussen, CRC Press, 2009), and the *OpenCL Programming Guide* (with A Munshi, B. Gaster, J. Fung, and D. Ginsburg, Addison Wesley, 2011).

3:00 - 3:30 - Afternoon Tea

3:30 - 4:15 - KEYNOTE

The (massive) opportunities of a Multicore World

Professor Barbara Chapman, OpenMP Architecture Review Board & University of Houston. USA

Abstract - Multicore technology is everywhere! Multicore processors are increasing the computational power available to high-end technical calculations in supercomputers and to embedded systems alike. They are being used to reduce database response times in social networking services and to create high-end medical devices as well as to validate astrophysical theory. Yet to take advantage of the parallelism of such platforms, suitable programming standards are essential.

OpenMP has established itself as a programming model with productivity benefits that offers a portable means to express shared memory parallel computations. It is being extended to increase the range of applications and systems that it can support. The Multicore Association is creating new standards for embedded system development that intend to facilitate the use of multicore platforms in this area too. In this presentation, we discuss the status of these efforts and their application in different industries, and how the broad adoption of these standards might be a game changer.

Bio - Barbara Chapman is Professor of Computer Science at the University of Houston in the United States where she specialises on Compiler Technology, Parallel Programming Languages, Tool Support for Application Development, Parallel Computing and High-Performance Computing. Among many activities, Professor Chapman is part of

- OpenMP Architecture Review Board
- CEO of cOMPunity, Inc.
- Member of International Exascale Project (IESP)
- Member of Open64 Steering Group
- Member of Executive Council, Educational Alliance for a Parallel Future (EAPF)

Previously, Prof. Chapman was director at the European Centre for Parallel Computing at Vienna, Austria. Her PhD is from Queen's University, Belfast, Northern Ireland

4:25 - 5:00 - PANEL

Do we have winners in a Multicore World?

Panellists: **Prof Ian Foster, Paul McKenney, Mark Moir, Poul-Henning Kamp.**

Moderator: **Nicolas Erdody**

In the next 10-15 years huge opportunities will emerge while translating and transforming sequential programming ('traditional' legacy code). We will create new software that takes full advantage of thousands of cores in the new generation of chips. How would a Multicore World look like? Keynote speakers will discuss with the Audience which organisations and / or industries are better positioned today to rapidly take advantage of this paradigm shift.

5:00 - 6:30 - Free

6:30 - 10:00pm - Open Parallel Conference Dinner - Civic Suites 1&2



Wednesday 20 February - Conference Day 2

Wellington Town Hall

8:55 - 9:00 - 2nd day Opening

9:00 - 9:45 - KEYNOTE

Massive Parallel Processing in the Varnish software

Poul-Henning Kamp, Chief Architect, Varnish Software & Author (of a lot of) FreeBSD. Denmark.

Abstract - The Varnish HTTP accelerator was written to show what modern MPP hardware with a modern UNIX kernel is capable of, if you stop programming like it was still the 1970-ies -as 90% of programmers worldwide still do. The answer is north of 1 million webpages per second, per machine.

Poul-Henning will talk about what Varnish taught a UNIX kernel programmer with 25 years under his belt about multiprogramming, and will share his tricks and suggestions for writing good MPP programs (that could dramatically change how you write software from now)

Bio - Poul-Henning Kamp has done weird things with computers for more than 30 years. Amongst the weird things are a IBM S/34 disassembler written in RPG-II, a PC defragmenter before Peter Norton wrote his, a LOT of FreeBSD, (including the "md5crypt" password scrambler), DARPA research on encrypted disks, microsecond timing for Air Traffic Control, and most recently the Varnish HTTP accelerator, which is used by more than half a million websites, including some of the biggest in the world. Poul-Henning is self-employed, lives in Denmark and looks forward to see the other side of the planet.

9:55 - 10:05

An Eclipse plugin for object-oriented parallel programming with Pyjama

Vikas Singh - Postgraduate student and Research Assistant, The University of Auckland, New Zealand

Abstract - Developing parallel applications is notoriously difficult, but is even more complex for desktop applications. The added difficulties are primarily because of their interactive nature, where performance is largely perceived by its users. Desktop applications are typically developed with graphical toolkits that in turn have limitations in regards to multi-threading.

This paper presents our latest object-oriented and GUI-aware tool that assists programmers in developing parallel user-interactive applications. More specifically, we present PJPlugin: an Eclipse plugin which aids in the development of parallel and concurrent programs using Java and the Pyjama compiler-runtime system. This paper provides an overview of PJ- Plugin and the Pyjama system, thus demonstrating the crucial role the plugin plays in a wider adoption of parallel programming.

Bio - Vikas is a Master of Engineering student and research assistant in the Department of Electrical and Computer Engineering, at the University of Auckland, New Zealand. His research includes Pyjama, a compiler-runtime system, and PJPlugin, the associated Eclipse plug-in. Pyjama research focuses on supporting OpenMP-like directives with GUI-extensions for Java. PJPlugin aims to increase programmer's productivity and add to ease of parallelisation and concurrency, while using Pyjama. In the past, Vikas has worked in software R&D in companies including Samsung and Sharp. He was technical lead at Samsung and led the teams to develop messaging, social and network applications for mobile devices. He contributed to product commercialisation and framework development for Android mobile devices and feature phones.

10:10 - 10:25

Multicore scores and resource optimisation within the Galaxy Project

David Eyers - University of Otago, Dunedin, New Zealand

Abstract - The Galaxy Project provides web-based access to bioinformatics experiments. By providing a consistent, accessible interface that wraps stand-alone analysis software, it allows scientists to focus on their actual work, rather than needing to become highly skilled in computing methods.

However, the user-friendly interface that Galaxy provides hides some key inefficiencies in the operation of the analysis tools it orchestrates. This paper describes our work in bringing resource monitoring information, including a multicore score, to the attention of users and tool developers, so that they can focus optimisation efforts on the least efficient parts of their scientific workflows.

Bio - Before joining the University of Otago Computer Science Department, David worked as a senior research associate at the University of Cambridge, from where he was awarded his PhD. He has undergraduate degrees in Computer Engineering and Maths from UNSW in Sydney, Australia.

David's research interests are in distributed systems, particularly regarding security enforcement and data dissemination mechanisms within wide-area applications. This has become particularly relevant within cloud computing: large-scale public services, such as electronic health record repositories, must manage sensitive data in a secure manner.

David's desire to examine green computing, and to become more involved with distribution at the level of CPU cores, has been met by him joining the Otago Systems Research Group, and its existing research projects in those fields .

10:30 - 11:00 - Morning Tea

11:00 - 11:45 - KEYNOTE

Transactional memory hardware, algorithms, and C++ language features

Mark Moir - Oracle Labs, USA - New Zealand

Abstract - Transactional memory (TM) aims to make it significantly easier to develop multicore programs that are scalable, efficient, and correct.

This talk will introduce TM, and discuss recent efforts towards standardization of language features in C++ for exploiting it. We will also touch on hardware transactional memory, its use to simplify and improve tricky concurrent algorithms, and its relationship to transactional language features such as those being proposed for C++

Bio - Mark Moir received the B.Sc.(Hons.) degree in Computer Science from Victoria University of Wellington, New Zealand in 1988, and the Ph.D. degree in Computer Science from the University of North Carolina at Chapel Hill, USA in 1996. He was an assistant professor in the Department of Computer Science at the University of Pittsburgh from 1996 until 2000. He then joined Sun Labs and subsequently formed the Scalable Synchronization Research Group, of which he is the Principal Investigator (the group is now part of Oracle Labs). He was named a Sun Microsystems Distinguished Engineer in 2009.

Dr. Moir's main research interests concern practical and theoretical aspects of concurrent, distributed, and real-time systems, particularly hardware and software support for programming constructs that facilitate scalable synchronization in shared memory multiprocessors.

11:55 - 12:25

Multicore COBOL: Three approaches

Richard O'Keefe - University of Otago, New Zealand

Abstract - COBOL is a commercially important historic programming language. This survey of ways to adapt it to a multicore world shows that support for parallelism can be added below a language, in it, or above it, and COBOL is a good example. It turns out that concurrent COBOL has been a commercially successful reality for 43 years.

Bio - I have a BSc(hons) in Statistics and MSc(hons) in Underwater Acoustics from Auckland, and a PhD in Artificial Intelligence from Edinburgh. I worked at Quintus Computer Systems, then wanted to get back into academia and went to RMIT to be part of the logic programming scene in Melbourne. I then wanted to be back in New Zealand and so came to Dunedin. I would sum up my research interests by saying - Everything connects! My biggest problem is that just about everything is interesting. My biggest concern is that most programs don't work, hence my interest in Software Engineering, and my liking for Ada. How can we apply neat ideas like logic programming, functional programming, meta-programming, formal methods, natural language processing, and so on to help build programs that work?

12:30 - 1:30 - Lunch - Civic Suites 1&2

1:30 - 1:55

Parallel Computing in the Cloud

Dave Fellows & Ivan Towlson - GreenButton, New Zealand

Abstract - One of the advantages of parallel computing is the ability to divide work across a large number of low-cost machines instead of having to invest in very expensive high-end hardware. But as processing loads grow, the cost of all those low-end machines still adds up. Cloud environments like Windows Azure and Amazon Web Services solve that problem by allowing you to pay only for the computing resources you need, only when you need them. In this talk, we share our experiences of running a diverse range of workloads in the cloud, the challenges we've encountered and what cloud

providers are doing to make the cloud a more hospitable environment for parallel computing.

Bio - Dave is the Chief Technology Officer at GreenButton, a platform for bursting workloads to the cloud. He has many years of experience designing massively scalable applications and is a regular speaker at cloud computing conferences. Ivan is the Chief Architect at GreenButton and has worked on distributed applications across a wide range of industries.

2:00 - 2:15

AutoCloud: Scalable Client-Side Replication

Christian Rolf - Corvid & Open Parallel, New Zealand

Abstract - In this paper we introduce AutoCloud; an architecture for offloading web servers by automatically constructing a cloud of replicating clients, henceforth known as mirrors. The design is similar to BitTorrent, where each client has a copy of the data, and can pass it on to all other clients. The inherent scalability of this approach is undeniable as it is the gold standard for distributing large files. We have simplified the BitTorrent protocol to increase the speed, as the data provided by web servers usually consists of many small files. We achieve maximal scalability by ensuring high compatibility. Replication in AutoCloud relies solely on JavaScript and the WebSockets interface, which is about to be standardized in HTML5. Most up-to-date browsers, like Chrome and Firefox, already support WebSockets.

The principle behind AutoCloud is to tag each element of a webpage as static, semi-static, or dynamic. Static elements typically include images, especially logos, which rarely change. Elements that are tagged as semi-static are given a timestamp until which they are valid. For many webpages, instant updates are desirable but not necessary, and a time-lag in the order of minutes is acceptable. Dynamic elements include breaking news and chat sessions, where any delay is undesirable. Our preliminary results show that server load can be significantly reduced, while increasing the total throughput to the clients. Basically, our approach allows servers to spend their cycles delivering dynamic data faster and cheaper, while concurrently reducing bandwidth consumption.

We have three focuses of our future research. Firstly, to order mirrors based on bandwidth and physical distance, allowing faster response times, and reducing the load of backbone networks. Secondly, we're working on peer-exchange between mirrors to further offload servers. Thirdly, we are looking at replicating the entire browser cache, rather than only the content that is currently open.

Bio - It was gaming that sucked Christian into the computer world, at the age of six. After a long and arduous journey, he emerged with a PhD (Lund, Sweden) in cloud computing for operations research. He's currently doing a start-up that provides optimality at the click of a button. His focus is on generating perfect rosters for medical staff, saving hospitals worldwide millions per year. He also provides cloud computing expertise for Open Parallel. This includes developing strategies for business cooperation, planning of large-scale government projects, and processing of big data.

2:20 - 3:00 - PANEL

Entrepreneurship and Multicore

Panellists: **Paul McKenney, Prof Barbara Chapman, Poul-Henning Kamp, Tim Mattson.**

Moderator: **Nicolas Erdody**

It is said that nowadays is 10x cheaper than 15 years ago to start an internet based business... entrance barriers are also lower. Would entrepreneurs take advantage of multicore hardware? Is multicore software expertise a "killer app" to differentiate from competitors? Or thanks to the tremendous processing power new business models will emerge based on multicore technologies?

3:00 - 3:30 - Afternoon Tea

3:30 - 3:55

(Re) Designed for High Performance: Solaris and Multi-Core Systems

James C. McPherson - Oracle, Australia

Abstract - Solaris 11 provided a rare opportunity to redesign and reengineer many components of the Solaris kernel. Designed to efficiently scale to the largest available systems (currently 256 cores and 64 TB of memory per system on SPARC, a little less on x64), significant work was done to remove obsolete

tuning concepts and ensure that the OS can handle the large datasets (both in-memory and on disk) which are now common- place. The continuing growth in virtualisation plays to Solaris' strengths: with the built-in hypervisor in the multi-threaded multi-core SPARC T- series we can provide hardware partitions even within a single processor socket. Solaris also features support for Zones (somewhat similar to the BSD Jails concept) which provide soft partitioning and allow presentation of an environment which mimics earlier releases of Solaris. This allows software limited to running on those older releases to obtain some of the performance and observability benefits of the host operating system.

Many years of development and support experience have given the Solaris engineering division an acute awareness that new features must include the capacity to observe them. While we have extensive testing, benchmarking and workload simulation capabilities to help bring a new release to the market, building in tools to help customers and support teams diagnose problems with their real-world usage are essential. The Solaris 11 release extended the work done with DTrace in Solaris 10, providing more probe points than ever before. This paper will describe some of the changes made to several parts of the operating system in Solaris 11, and the motivations behind those changes.

Bio - James joined Sun Microsystems in November 1999, working in the Support Services organisation. After several years doing technical escalation management, bug fixing in the storage software area and training colleagues in analytical troubleshooting, he moved to the Solaris development organisation. Working first in the fibre channel and multipathing group, he then moved on to the serial attached SCSI (SAS) multipathing project. From January 2009 until November 2011 he was the Gatekeeper for the Solaris kernel. Following the successful launch of Solaris 11 he moved to the Solaris Modernization team where he is working on redesigning some basic system administration tools to work more usefully in the massive multicore environments which are now standard.

4:00 - 4:45 - KEYNOTE

The revolution in your pocket:

Invisible computers with Dissociative Identity Disorder

Tim Mattson - Intel Principal Engineer, Khronos OpenCL Group. USA

Abstract - Predicting the future is hard; especially when it hasn't happened yet. This is especially true in the computer industry. If you pay close attention to hardware trends and emerging software standards, however, you can sketch out the high level details of where we are going.

We will discuss this future in this talk and explore the consequences of the fact that: (1) computers are becoming invisible and (2) they have developed dissociative identity disorder. This is great for hardware designers. Software developers, however, will have their work cut out for them as they adapt to this brave new world.

Bio - Tim Mattson is a parallel programmer (Ph.D. Chemistry, UCSC, 1985). Tim has been with Intel since 1993 where he has worked with brilliant people on great projects such as: (1) the first TFLOP computer (ASCI Red), (2) the OpenMP API for shared memory programming, (3) the OpenCL programming language for heterogeneous platforms, (4) Intel's first TFLOP chip (the 80 core research chip), and (5) Intel's 48 core, SCC research processor.

Tim has published extensively including the books *Patterns for Parallel Programming* (with B. Sanders and B. Massingill, Addison Wesley, 2004), *An Introduction to Concurrency in Programming Languages* (with M. Sottile and C. Rasmussen, CRC Press, 2009), and the *OpenCL Programming Guide* (with A Munshi, B. Gaster, J. Fung, and D. Ginsburg, Addison Wesley, 2011).

4:45 - 5:00 - Conference Closure

5:00 - 6:30 - Cocktail - Civic Suites 1&2



MULTICORE WORLD 2013

Team Open Parallel @MulticoreWorld

Lev Lafayette (MC) - President Linux Users Victoria (Melbourne, Australia); Systems Admin & Project Manager at VPAC (Victorian Partnership of Advanced Computing)

Andrew McMillan (photos), **Erica Hoehn** (registration), **Peter Kerr** (journo), **Christian Rolf**, **Lenz Gschwendtner**, **Beau Johnson**, **Nicolas Erdody**. Very special thanks to **Pete Salerno** (Chicago) and **Kevin Black** (Dunedin)

Programme Committee

<u>PC Chair</u> - Prof. Ian Foster	Alistair Rendell (Professor, Deputy Dean Research School of Computer Science, Australian National University)
Jan-Philip Weiss (Jun.-Prof. Head of Computing Lab Hardware-aware Numerics, KIT-Germany)	Mario Nemirovsky (Research Professor, Barcelona Supercomputing Centre, Spain)
Tim Cornwell (Head of Computing, SKA Organisation, UK)	Ariel Hendel (Senior Technical Director, Broadcom. USA)
Akash Deshpande (Distinguished Engineer, Cisco. USA)	Martin McKendry (SVP Seven Networks, USA)
Paul McKenney (IBM, Linaro. USA)	Mark Moir (Oracle Labs, USA/NZ)
Zhiyi Huang (University of Otago, NZ)	John Michalakes (Snr. Scientist NREL, USA)
Dr. Ing. Hector Cancela (Dean, School of Engineering, UdelaR, Uruguay)	Suren Byna (Research Scientist, Lawrence Berkeley National Lab, USA)
Lindsay Groves (Victoria University of Wellington, NZ)	Nick Jones (Director NeSI, NZ)
François Marier (Mozilla, NZ)	David Eyers (University of Otago, NZ)
Esteban Mocskos (Universidad de Buenos Aires, Argentina)	Andrew McMillan (Debian Project)
Nathan Torkington (PM Perl 6, Former Chair OSCON, NZ)	Sergio Nesmachnow (Universidad de la República, Uruguay)

Sponsors

Open Parallel

Software for Multicore and Parallel Computing



 Scoop Media Cartel

Supporters



4 February 2013 - DISCLAIMER

NOTE: Titles and abstracts are a work in progress and are an indication of the level of the conference: Guest and Keynote Speakers will present the latest innovation in the field and their updated thinking by the time of the conference. More speakers could be added plus specialised panels and events for audience participation. Final version will be released on Monday 18 February 2013. Multicore World organisation does not assume any liability if for any circumstance a Speaker cannot be present at the conference. Regular updates appear in www.MulticoreWorld.com