

6th VI-HPS Tuning Workshop hosted by SARA in Amsterdam

Markus Geimer Jülich Supercomputing Centre m.geimer@fz-juelich.de 26-28 May 2010











- Presenters/Guides
 - Christian Siebert & Markus Geimer (GRS/FZ Jülich JSC)
 - Andreas Knüpfer & Matthias Jurenz (TU Dresden ZIH)
 - David Cronk (UTK ICL)
 - Yury Oleynik (TU Munich)
 - Wyatt Spear (U Oregon)
- Thanks
 - Host & Systems: SARA
 - ► John Donners, Wim Rijks, Lilit Axner, ...
 - VI-HPS/POINT partners
 - ► RWTH Aachen, UTK-ICL, UIUC-NCSA, U Oregon PRL, ...
 - You
 - ► Your Name Here



We'd like to know a little about you, your application(s), and your expectations and desires from this workshop

- What programming paradigms do you use in your app(s)?
 - only MPI, only OpenMP, mixed-mode/hybrid OpenMP/MPI, ...
 - Fortran, C, C++, mixed-language, ...
- What platforms/systems *must* your app(s) run well on?
 - SGI Altix, IBM Power/AIX, BlueGene/P, Linux cluster™, ...
- Who's already familiar with *serial* performance analysis?
 - Which tools have you used?
 - ► time, print/printf, prof/gprof, SpeedShop, VTune, PTU, ...
- Who's already familiar with *parallel* performance analysis?
 - Which tools have you used?
 - ► time, print/printf, prof/gprof, mpiP/ompP, ITAC/ITT, MPInside, ...



- The workshop concentrates on *hands-on* use of correctness and performance analysis tools with your own application(s):
- Who has prepared their app(s) to be analysed and tuned?
 - A small yet representative test case should build
 - and run (correctly) to completion within a few minutes
 - ... on a relatively small number of processors/cores
- Who has prepared to analyse and improve scalability?
 - One or more larger test cases should build
 - and run (correctly) to completion within less than an hour
 - ... on larger numbers of processors/cores (in batch mode)
 - Strong scaling = fixed total problem size
 - speed-up expected with increasing numbers of processors
 - Weak scaling = constant problem size per process/thread
 - time not expected to change for increasing numbers or processors

It happens



- You may experience problems with the HPC systems
 - They're continually being pushed to the limit (and beyond)
 - It's probably not your fault when they break
- You may discover bugs in your application(s)
 - That's presumably why you're here
 - ... and we're here to help
 - Think of identified bugs as opportunities for improvement, to make your particular application fitter and more competitive, as HPC systems continue to get larger and more complex
- You may encounter defects/deficiencies with the tools
 - We want to learn from you what should be improved
 - It's also nice to hear when they work well and are helpful



- **Goal**: Improve the quality and accelerate the development process of complex simulation codes running on highly-parallel computer systems
- Funded by Helmholtz Association of German Research Centres



- Activities
 - Development and integration of HPC programming tools
 - Correctness checking & performance analysis
 - Training workshops
 - Service
 - Support email lists
 - Application engagement
 - Academic workshops

www.vi-hps.org

Partners



RNTHAACH





Jülich Supercomputing Centre



- RWTH Aachen University
 - Centre for Computing & Communication
- Technical University of Dresden
 - Centre for Information Services & HPC



University of Tennessee (Knoxville)Innovative Computing Laboratory





- Technical University of Munich
 - Chair for Computer Architecture
- University of Stuttgart
 - High Performance Computing Centre









Universität Stuttgart





- Marmot
 - Free MPI correctness checking tool
- PAPI
 - Free library interfacing to hardware performance counters
- Periscope
 - Prototype automatic analysis tool using an on-line distributed search for performance inefficiencies
- Scalasca
 - Open-source toolset for analysing the performance behaviour of parallel applications to automatically identify inefficiencies
- Vampir/VampirTrace
 - Commercial tool for graphical trace visualization & analysis, and open-source event tracing library

[Tuning Workshop Live-DVD contains latest tools releases]

Technologies and their integration





VI-HPS component technologies



- Key tool components also provided as open-source
 - Program/library instrumentation
 - ► OPARI, POMP
 - MPI library/tool integration
 - UniMCI
 - Scalable I/O
 - ► SIONIib
 - Libraries & tools for handling (and converting) traces
 - ► EPILOG, EARL, PEARL, OTF, Œ
 - Analysis algebra & hierarchical/topological presentation
 - ► CUBE





- VI-HPS collaborates with the POINT project in the USA
 - Petascale Productivity from Open, Integrated Tools
 - Funded by US NSF SDCI, Software Improvement & Support
 - University of Oregon, University of Tennessee, UIUC NCSA, and Pittsburgh Supercomputing Center
 - www.nic.uoregon.edu/point



- Goals
 - Give an overview of the programming tools suite
 - Explain the functionality of individual tools
 - Teach how to use the tools effectively
 - Offer hands-on experience and expert assistance using tools
 - Receive feedback from users to guide future development
- For best results, bring & analyse/tune your own code(s)!
- VI-HPS Tuning Workshop series
 - Aachen (3/08), Dresden (10/08), Jülich (2/09), Bremen (9/09), Garching (3/10)
- Joint POINT/VI-HPS Tutorial series
 - SC (11/08), ICCS (5/09), SC (11/09)
- Training with individual tools & platforms (e.g., BlueGene)



- Wednesday 26th May
 - 08:30 (registration & notebook computer set-up)
 - 09:00 Welcome & Introduction to VI-HPS [Geimer, JSC]
 - Virtual Institute High Productivity Supercomputing
 - Building and running the tutorial exercise NPB3.3-MPI/BT
 - 09:45 Marmot correctness checking tool [Knüpfer, TUD-ZIH]
 - Hands-on tutorial exercise with Marmot
 - 10:30 (break)
 - 11:00 Scalasca performance analysis toolset [Geimer, JSC]
 - Hands-on tutorial exercise with Scalasca
 - Case studies using Scalasca
 - 12:00 (lunch)



- Wednesday 26th May
 - 13:00 Introduction to the Vampir toolset [Knüpfer, TUD-ZIH]
 - Hands-on tutorial exercises with Vampir/VampirTrace
 - Case studies using Vampir
 - 14:00 Hardware counter analysis with PAPI [Cronk, UTK-ICL]
 - 14:30 (break)
 - 15:00 Introduction to Periscope toolset [Oleynik, TUM]
 - Hands-on tutorial exercises with Periscope
 - 15:45 TAU toolset for profile analysis [Spear, UOregon]
 - 16:30 Review and preparation for rest of workshop
 - Prepare participants' own codes for analysis on SARA systems
 - Further exercises with tools
 - 17:00 (adjourn)



- Thursday 27th May
 - 09:00 Tools coaching with participants' applications
 - Assistance using Marmot, Periscope, TAU, Scalasca & Vampir
 - 12:00 (lunch)
 - 13:00 Tools coaching & possible additional presentations
 - 17:00 Review of day and schedule for tomorrow
 - 17:30 (adjourn for socializing and drinks)
- Friday 28th May
 - 09:00 Tools coaching & possible additional presentations
 - 12:00 (lunch)
 - 13:00 Review of workshop and participants' experiences
 - 14:00 (adjourn or continue to work to 16:30)



- Bootable Linux installation on DVD (or USB stick)
- Includes everything needed to try out our parallel tools on an x86-architecture notebook computer
 - GCC compiler suite (with OpenMP support), OpenMPI library
 - POINT tools: PAPI, PerfSuite, TAU
 - VI-HPS tools: Marmot, Periscope, Scalasca, VT/Vampir*
 - Other tools: BUPC, dyninst, Eclipse/PTP, PPW, TotalView*
 - * time/capability-limited evaluation licences provided for commercial products
 - Manuals/User Guides
 - Tutorial exercises and examples
- Prepared by U. Oregon Performance Research Laboratory
 - Sameer Shende & Alan Morris

VI-HPS

% module ...

- list
- avail
- avail package(s)

print currently loaded packages# print all available packages# print versions of *package(s)* available

- load package(s)
- unload package(s)

configure access to package(s)

- # remove access to package(s)
- **swap** *package1 package2*# replace *package1* with *package2*
- whatis package(s)
- help package(s)
- show package(s)

print short description of package(s)
print longer description of package(s)
print settings done for package(s)



6th VI-HPS Tuning Workshop hosted by SARA in Amsterdam

Markus Geimer Jülich Supercomputing Centre m.geimer@fz-juelich.de 26-28 May 2010











Review



- You've survived a long day with lots of hands-on work
- We've presented and you've used a variety of tools from POINT & VI-HPS for parallel performance engineering
 - MARMOT, Periscope, Scalasca & Vampir toolsets provide complementary functionality and increasing interoperability
 - often installed together under the UNITE module configured for a *uniform integrated tool environment*
- The tools are actively supported by VI-HPS & developers
 - development teams welcome feedback, including requests
 - bug-reports when appropriate also help us improve our tools
 - VI-HPS consultancy service offered to German public HPC institutions (e.g., Gauss Centre for Supercomputing)
- Take time to complete our workshop evaluation survey
- ... and discuss your ideas/experience with us right here!



- The "real" hands-on work starts tomorrow
 - using *your* application codes on the HPC computer system (IBM Power6: Huygens@SARA)
- Ensure your application codes build and run to completion with appropriate datasets
 - initial configuration should ideally run in less than 15 minutes with around 16 processes/threads
 - to facilitate rapid turnaround and quick experimentation
 - Iarger/longer scalability configurations are also interesting
 - turnaround may be limited due to busyness of batch queues
- Compare your application performance on other systems
 - requires your own account on the other systems
 - the tools are already installed, ask for details if necessary



- Thursday 27th May
 - 09:00 Recap of tools usage and local facilities set-up
 - 09:30 Tools coaching with participants' applications
 - ► Assistance using Marmot, Periscope, TAU, Scalasca & Vampir
 - 12:00 (lunch)
 - 13:00 Tools coaching & possible additional presentations
 - 17:00 Review of day and schedule for tomorrow
 - 17:30 (adjourn for socializing and drinks)
- Friday 28th May
 - 09:00 Tools coaching & possible additional presentations
 - 12:00 (lunch)
 - 13:00 Review of workshop and participants' experiences
 - 14:00 (adjourn or continue to work to 16:30)