5th VI-HPS Tuning Workshop
hosted by TUM/LRZ/MPG in Garching

Brian Wylie
Jülich Supercomputing Centre
b.wylie@fz-juelich.de

8-10 March 2010

Introduction

- Presenters/Guides
  - David Böhme, Markus Geimer & Brian Wylie (FZ Jülich JSC)
  - Tobias Hilbrich & Thomas William (TU Dresden ZIH/GWT)
    - Jens Doleschal unfortunately unable to attend
  - Michael Gerndt, Ventsislav Petkov & Yury Oleynik (TU Munich)
- Thanks
  - Host: Technische Universität München
  - Sponsor: ParTec Cluster Competence Center GmbH
  - Systems: Leibniz RZ & Max-Planck Institüt RZ Garching
    - Orlando Rivera (LRZ), Andreas Schmidt (RZG), ...
  - VI-HPS/POINT partners
    - RWTH Aachen, UTK-ICL, UIUC-NCSA, U Oregon PRL, ...
  - You
    - *Your Name Here*

Participant survey

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, ...

Warning

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

Virtual Institute – High Productivity Supercomputing

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

Forschungszentrum Jülich

- 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

Productivity tools

- 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

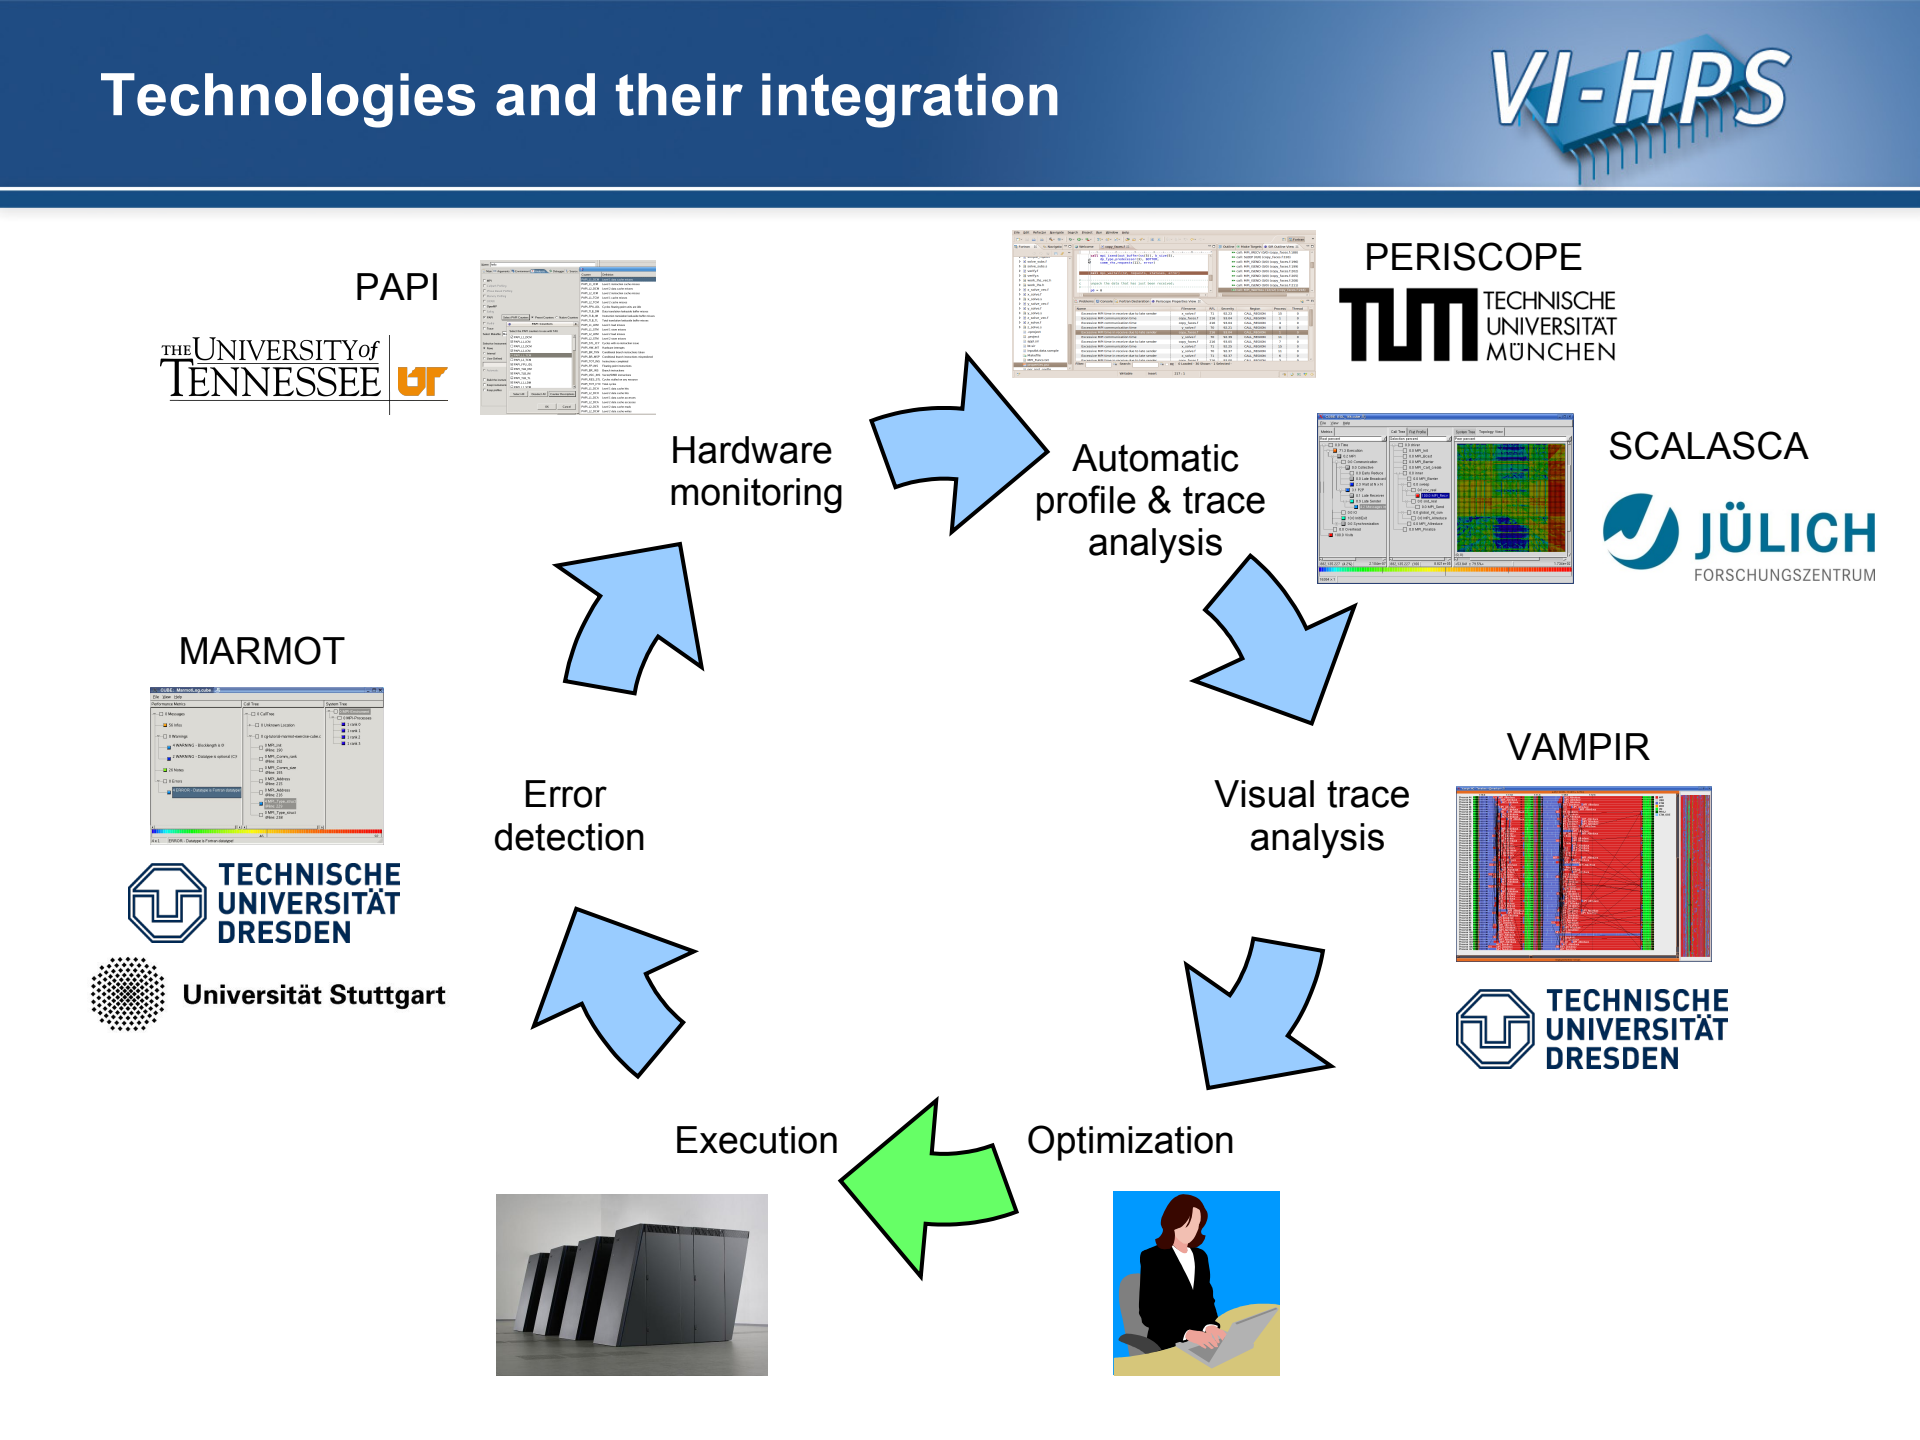

VI-HPS component technologies

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

POINT/VI-HPS collaboration

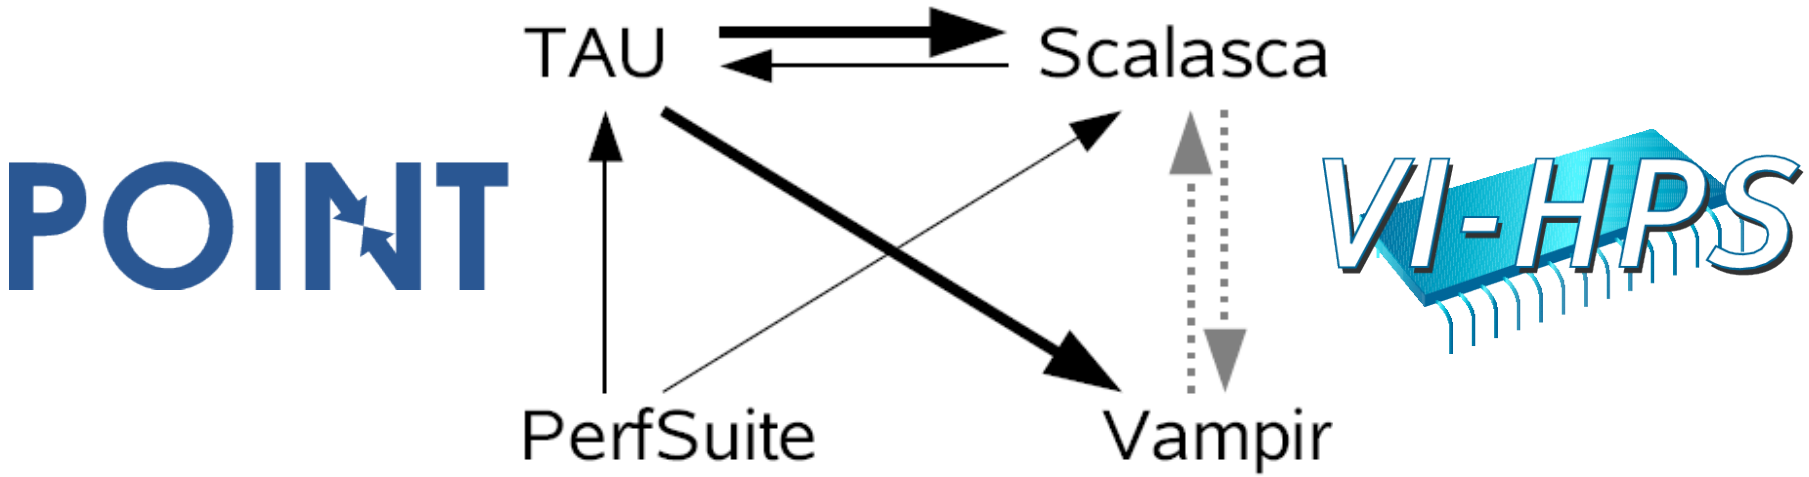

- VI-HPS collaborates with the POINT project in the USA
  - Petascale Productivity from Open, Integrated Tools
  - Funded by US NSF SDCI, Software Improvement & Support
  - University of Oregon, University of Tennessee, UIUC NCSA, and Pittsburgh Supercomputing Center
  - www.nic.uoregon.edu/point

VI-HPS Training & Tuning Workshops

- 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)
- Joint POINT/VI-HPS Tutorial series
  - SC (11/08), ICCS (5/09), SC (11/09)
- Training with individual tools & platforms (e.g., BlueGene)

Upcoming VI-HPS training events

- 6th VI-HPS Tuning Workshop
  - 26-28 May 2010
  - Hosted by SARA in Amsterdam, The Netherlands
    - priority for users of SARA systems, e.g., Huygens (Power6/Linux)
    - places for non-SARA users may become available
- Further events to be determined
  - (one-day) tutorials
    - with guided exercises using Live DVD
  - (multi-day) training workshops
    - with your own applications on real HPC systems

- Check www.vi-hps.org/training for announced events
- Contact us if you might be interested in hosting an event

Outline (tutorial, am)

- Monday 8th March
  - 08:30 (registration & notebook computer set-up)
  - 09:00 Welcome & Introduction to VI-HPS [Wylie, JSC]
    - Virtual Institute – High Productivity Supercomputing
    - Building and running the tutorial exercise NPB3.3-MPI/BT
  - 09:30 Parallel performance engineering [Gerndt, TUM]
    - Introduction to performance analysis techniques and tools
  - 10:15 Marmot correctness checking tool [Hilbrich, TUD-ZIH]
    - Hands-on tutorial exercise with Marmot
  - 10:45 (break)
  - 11:15 Scalasca performance analysis toolset [Geimer, JSC]
    - Hands-on tutorial exercise with Scalasca
    - Case studies using Scalasca
  - 12:30 (lunch)

Outline (tutorial, pm)

- Monday 8th March
  - 13:30 Introduction to the Vampir toolset [William, TUD-ZIH]
    - Hands-on tutorial exercises with Vampir/VampirTrace
    - Case studies using Vampir
  - 14:45 (break)
  - 15:15 Introduction to Periscope toolset [Oleynik/Petkov, TUM]
    - Hands-on tutorial exercises with Periscope
    - PAPI library & utilities
  - 16:30 Review and preparation for rest of workshop
    - Distribution/set-up of workshop accounts for participants
    - Prepare participants' own codes for analysis on LRZ/RZG/JSC HPC systems
    - Further exercises with tools
  - 17:30 (adjourn)

Outline (workshop)

- Tuesday 9th March
  - 09:00 Tools coaching with participants' applications
    - Assistance using Marmot, Periscope, Scalasca & Vampir
  - 12:30 (lunch)
  - 13:30 Tools coaching & possible additional presentations
  - 17:00 Review of day and schedule for tomorrow
  - 17:30 (adjourn for sponsored dinner)
- Wednesday 10th March
  - 09:00 Tools coaching & possible additional presentations
  - 12:00 (lunch)
  - 13:00 Review of workshop and participants' experiences
  - 15:00 (adjourn or continue to work to 17:30)

POINT/VI-HPS Live-DVD

- 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

Basic module commands

% module ...

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

- **load** *package(s)* # configure access to *package(s)*
- **unload** *package(s)* # remove access to *package(s)*
- **swap** *package1 package2* # replace *package1* with *package2*

- **whatis** *package(s)* # print short description of *package(s)*
- **help** *package(s)* # print longer description of *package(s)*
- **show** *package(s)* # print settings done for *package(s)*

UNITE: Uniform Integrated Tool Environment

- Based on flexible and convenient “module” package
  - widely employed on HPC systems
    - local installation determines modules loaded by default
  - selects desired software products and versions
  - dynamically updates user environment configuration
    - e.g., PATH, licenses, default settings
    - applies to current session/shell/job only!
  - provides software description and/or basic help/usage
- but unfortunately in practice
  - unstructured assemblies of modules quickly become awkward
  - module names and configurations vary according to system
- UNITE modules standardize tool access & documentation
  - deployed on JSC, LRZ, RWTH & ZIH production systems
- Latest information/version at http://apps.fz-juelich.de/unite

UNITE@HLRB-II

- UNITE installation on LRZ HLRB-II

```
% module help UNITE
-------------------------------------------------------------------
UNITE: UNIform Integrated Tool Environment
This module initializes the UNITE modules environment

For more information:
- http://apps.fz-juelich.de/unite/
- mailto:vi-hps-support@rz.rwth-aachen.de
-------------------------------------------------------------------

% module load UNITE
UNITE loaded
% module avail 2>&1 | head
------------ /lrz/sys/tools/unite/modulefiles/tools ------------
marmot/2.4.0-sgimpt-intel
scalasca/1.3.0-sgimpt-intel
vampirtrace/5.8-sgimpt-intel-marmot
vampirserver/2.1.1
vampir/5.2.0
------------ /lrz/sys/tools/unite/modulefiles/utils ------------
cube/3.3-intel
papi/3.6.2
```

- Warning: non-UNITE modules are older versions of tools!

5th VI-HPS Tuning Workshop
hosted by TUM/LRZ/RZG in Garching

Brian Wylie
Jülich Supercomputing Centre
b.wylie@fz-juelich.de

8-10 March 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!

Getting prepared

- The “real” hands-on work starts tomorrow
  - using ***your*** application codes on the HPC computer system (SGI Altix: HLRB-II@LRZ)
- 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
  - larger/longer scalability configurations are also interesting
    - turnaround may be limited due to busyness of batch queues
- Compare your application performance on other systems (e.g., AIX: VIP@RZG, BG/P: GENIUS@RZG, JUGENE@JSC)
  - requires your own account on the other systems
  - the tools are already installed, ask for details if necessary

Outline (workshop)

- Tuesday 9th March
  - 09:00 Recap of tools usage and local facilities set-up
  - 09:30 Tools coaching with participants' applications
    - Assistance using Marmot, Periscope, 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 sponsored dinner at 19:30)
- Wednesday 10th March
  - 09:00 Tools coaching & possible additional presentations
  - 11:00 Review of workshop and participants' experiences
  - 12:00 (lunch)
  - 13:00 Tools coaching & possible additional presentations
  - 15:00 (adjourn or continue to work to 17:30)