



21st VI-HPS Tuning Workshop

Leibniz-Rechenzentrum, Garching-bei-München, Germany



18-22 April 2016

<http://www.vi-hps.org/tws/tw21.html>

21st VI-HPS Tuning Workshop (LRZ)

▪ Tools instructors

- Christian Feld (Jülich Supercomputing Centre)
- Michael Firbach & Josef Weidendorfer (Technische Universität München)
- Judit Gimenez & German Llort (Barcelona Supercomputing Center)
- Florent Lebeau (Allinea Software Ltd)
- Joachim Protze (RWTH Aachen University)
- Martin Schulz (Lawrence Livermore National Laboratory)
- Sergei Shudler (Technische Universität Darmstadt)
- Cédric Valensi & Emmanuel Oseret (Université de Versailles Saint-Quentin)
- Bert Wesarg & Johannes Ziegenbalg (Technische Universität Dresden)

▪ Local organisation

- Anupam Karmakar & Volker Weinberg

▪ Sponsor: Gauss Centre for Supercomputing PRACE Advanced Training Centre

Outline

Monday 18 April

- 09:00 Welcome [Anupam Karmakar & Volker Weinberg, LRZ]
 - Introduction to VI-HPS and overview of tools [Michael Gerndt, TUM]
 - SGI *UltraViolet2* computer system and software environment [LRZ]
 - Building and running NPB-MZ-MPI /BT-MZ on SGI UV2 [Christian Feld, JSC]
 - Introduction to parallel performance engineering [Michael Gerndt, TUM]
- 10:30 (break)
- 11:00 **MAQAO performance analysis tools** [Cédric Valensi & Emmanuel Oseret, UVSQ]
 - MAQAO hands-on exercises
- 12:30 (lunch)
- 13:30 Hands-on coaching to apply tools to analyze your own code(s)
- 17:30 Review of day and schedule for remainder of workshop
- 18:00 (adjourn)

- Hands-on exercises part of each presentation to familiarise with tools every morning session
- Hands-on coaching to apply tools to analyse and tune your own codes each afternoon

Outline of rest of week

Tuesday 19 April

- 09:00-10:30 **Allinea performance tools suite** [Florent Lebeau, Allinea]
- 11:00-12:30 **Open|SpeedShop & mpiP performance tools** [Martin Schulz, LLNL]

Wednesday 20 April

- 09:00-10:30 **Score-P instrumentation & measurement** [Christian Feld & Bert Wesarg]
- 11:00-12:30 **CUBE profile explorer** [Christian Feld, JSC]
Extra-P automated performance modelling [Sergei Shudler, TUDarmstadt]

Thursday 21 April

- 09:00-10:30 **Scalasca automated trace analysis** [Christian Feld, JSC]
Vampir interactive trace analysis [Bert Wesarg & Johannes Ziegenbalg, TUDresden]
- 11:00-12:30 **Paraver tracing tools suite** [Judit Gimenez & German Llort, BSC]

Friday 22 April

- 09:00-10:30 **Kcachegrind & Periscope TF** [Josef Weidendorfer & Michael Firbach, TUM]
- 11:00-12:30 **MUST/ARCHER runtime error detection** [Joachim Protze, RWTH]
STAT stack trace analysis [Martin Schulz, LLNL]
- 13:30-14:00 Conclusion & Review

Participant survey

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

- What programming language(s) do you use?
 - Fortran, C, C++, multi-language, ...
- What parallelisation mode(s) do you use?
 - only MPI, only OpenMP, mixed-mode/hybrid MPI+OpenMP, ...
- What platforms/systems *must* your code run well on?
 - SuperMUC, Cray, IBM BlueGene, Linux cluster, ...
- Are you already familiar with *serial* performance analysis? Using which tools?
 - time, print/printf, prof/gprof, VTune, ...
- Are you already familiar with *parallel* performance analysis? Using which tools?
 - time, print/printf, prof/gprof, ITAC, Paraver, Scalasca, TAU, Vampir, ...

Prepare to analyse your own application code(s)

- Ensure that your application code(s) build and run correctly to completion with appropriate datasets
 - initial configuration should ideally run in less than 15 minutes with 1-4 compute nodes
 - to facilitate rapid turnaround and quick experimentation
 - larger/longer scalability configurations are also interesting
 - turnaround may be limited due to busyness of batch queues, but perhaps overnight
- Compare your application performance on other computer systems
 - VI-HPS tools are already installed on many HPC systems
 - if not, ask your system administrator to install them (or install a personal copy yourself)

Disclaimer

Tools will ***not*** automatically make you, your applications or computer systems more productive.

However, they can help you understand ***how*** your parallel code executes and ***when / where*** it's necessary to work on correctness and performance issues.

Evaluation / Feedback

- Please also complete and return the VI-HPS workshop paper form, which provides valuable feedback
 - to tools developers for improving their tools and training material
 - to improve future workshops and training events
 - can be anonymous if desired
- Tools support queries and bug reports are also welcome
 - should be submitted to respective support mailing lists