

### Lesson learnt from 15+yrs of VI-HPS & POP CoE training

Brian Wylie Jülich Supercomputing Centre



25th Anniversary



### Agenda

- Virtual Institute High Productivity Supercomputing
  - VI-HPS Tuning Workshops
- EU HPC Applications Centre of Excellence Performance Optimisation & Productivity
  - Workshops by&for under-represented groups in HPC
- Hybrid
  - Durham University Performance Analysis Workshop Series
- Etc
  - Conference tutorials, seasonal school contributions & bespoke training workshops

### **Virtual Institute – High Productivity Supercomputing**

- Goal: Improve the quality and accelerate the development process of complex simulation codes running on highly-parallel computer systems
- Start-up funding (2006–2011)

by Helmholtz Association of German Research Centres

- Activities
  - Development and integration of HPC application tools
  - Primarily correctness checking & performance analysis
  - Academic workshops: e.g. ProTools@SC
  - Tools training via conference tutorials and multi-day "bring-your-own-code" Tuning Workshops
    - Face-to-face & side-by-side hands-on coaching now successfully migrated also to virtual/on-line events

### http://www.vi-hps.org



#### **VI-HPS Tuning Workshops**



Multi-day (3-5) hands-on workshops

- hosted by HPC computing centres
- participants bring their own code to analyse applying variety of tools
  - ideally working together in code-teams
- multiple member's tools plus invited guests
  - partially guided by host preferences
- originally always in-person, transitioned to virtual events during pandemic, potentially hybrid in future
- primarily within Europe, previously often co-funded by PRACE (Advanced) Training Centres
  - also Japan & Americas (via JLESC partners)

### **VI-HPS Tuning Workshops evolution**

- Initial workshop hosted by RWTH Aachen in March 2008
  - Format defined by Dieter an Mey and Brian Wylie
  - Short presentations/demonstrations of tools, followed by hands-on application of tools
  - Opportunity to try out and compare use of multiple tools (from VI-HPS and external guests)
  - Feedback from participants (in)directly to tools developers
- 32 workshops followed similar though evolving format
  - Introduction of a representative example code as basis for demonstration of all tools
  - Emphasis on participants bringing their own application code(s) for analysis
  - Brian Wylie as Training Coordinator for VI-HPS
  - Increasing conflict of interest with similar/common workshops organised by POP CoE
- Transitioned to virtual/online events mandated by COVID19 pandemic
  - Cedric Valensi (UVSQ) took over Training Coordinator role for VI-HPS
  - Worked surprisingly well, but lack interaction with and between participants
- Total of 43 workshops to date; 3 more already scheduled in 2024H1

### **VI-HPS-TW** example codes evolution

- Jacobis, SGEMV, ASC SMG2000, ASC LULESH, NAS NPB BT
  - MPI (OpenMP)
- NAS NPB BT-MZ
  - MPI+OpenMP
- UK-MAC TeaLeaf
  - MPI+OpenMP+CUDA
- UK-MAC CloverLeaf
  - MPI+OpenACC(+CUDA)
- ...
- Highly portable benchmark codes, without significant performance inefficiencies
  - generally strong scaling with minimal file I/O, but other execution configurations possible
- Configurable to run from a notebook to a supercomputer: O(4) to O(10000)
  - internal report of (kernel) execution time and validation of correct execution

### Lessons: optimal execution configurations

#### Not too long, but also not too short

- 15sec to 3min typically acceptable/appropriate to allow multi-run experimentation
- comparison of uninstrumented reference with various instrumented versions and measurement configurations

#### Not too large, but also not too small

- Resources are constrained, both on notebook computers and reserved HPC partitions
- Notebook computers typically have only a few (physical) cores
  - virtual cores are generally unhelpful for computationally-intensive scientific codes
- A single compute node can be reasonable (if dedicated for exclusive use)
  - nowadays have multiple processor sockets with many cores (and distinct NUMA domains)
- Several compute nodes is better, but two often sufficient
  - shows distinction of intra-node and inter-node performance aspects
  - more representative of intended HPC usage
- Access to larger numbers of compute nodes also valuable for execution scaling analyses
  - often can be deferred to overnight or other breaktimes

### Lessons: example source bundles

#### Valuable augmentations

- Jobscripts for execution (with & without tools) on different computer systems
  - typical for usage of HPC systems, and better for reference
- Annotation of Makefile(s) with instrumentation options
  - PREP compiler/linker preposition can be uncommented or defined on Make command
  - distinct naming or directory to separate uninstrumented reference executable from instrumented version(s)
- Source code instrumentation annotations can be transparently incorporated
  - enabled on demand via instrumenter flag when desired
- Auxiliary files such as for instrumentation or measurement filtering can be provided
  - suitable wildcard patterns allow single file to be used with multiple compilers
- Single source-file example codes are too trivial to be useful
- Makefile/source-file complexity and compile/build times need to be managable
- Fortran example codes uncover tools issues that developers miss
  - C/C++ are always tested much more
    - C++ example codes present much more complexity (and variability)

### Lessons: hybrid MPI+threading+offload

- Provides flexibility at the cost of additional complexity
  - Single example can be customised for use with different audiences having minimal knowledge of each parallelisation paradigm
  - Can readily adapt example execution runtime according to available compute resources by varying numbers of processes/threads
  - Opens additional investigation of associated execution efficiency changes
- Inherent complexity generally not too much of an issue
  - no source code (language) familiarity or changes are required to do exercises with tools
- While MPI and threading can be addressed as orthogonal parallelisation aspects, combinations of threading or offload schemes are not
  - OpenMP+Pthreads (plus inherent language concurrency)
  - OpenMP/OpenACC+ROCm/HIP/CUDA

#### **Lessons: virtual events**

#### Remote participation quality varies

- plagued by unreliable and/or slow network connections
- additional technological & management demands of multiple screens
  - one dedicated for interactive compute session, others for presentation, communication, etc.
  - only partially alleviated by multiple windows and browser tabs
- time-zone deltas reduce productive time
- Attractive for many potential participants (worldwide)
  - Avoids financial expense of travel/accommodation (also for instructors)
  - Reduces time commitment for participation
    - can flexibly join only the convenient sessions
  - Anonymous/stealth monitoring vs. active participation
- Fundamentally a poor alternative, but can still be valuable
  - Reaches participants for which it's the only (or most viable) option
  - Relatively little cost for hosting organisation and instructors
    - though typically more instructors and assistants required to monitor and manage breakout rooms

### Lessons: added-value/benefits of in-person training

- Direct interaction between instructors and participants
  - clear indications of progress/difficulties and need for assistance
  - open feedback in breaks
- Interaction and information-sharing between participants
  - assisting of neighbours (potential code-team colleagues)
  - identification of common interests and issues
- Dedicated attention
  - fewer distractions in classroom as from (shared/home) office environment

### Accumulated knowledge

- Instruction session recordings are valuable for everyone
  - (partially) supports those who miss the live session
  - allow replay at individual pace and when convenient (possibly months later)
- Combined virtual+presence training is the worst of both options for everyone
  - both physical classroom & streaming resources required
  - twice the work for hosts and instructors to manage both concurrently
  - common schedule is awkward/inconvenient for both kinds of participants
    - in-person training most productive in several consecutive (full) days
    - virtual training most effective when half-day sessions spread over several weeks

### We are HPC!

WOMEN IN HIGH PERFORMANCE COMPUTING

- Representatives (and senior instructrices)
  - Judit Giménez (BSC) & Heike Jagode (UTK-ICL)
- Instructrices & training organisers
  - Marta García-Gasulla (BSC), Anja Gerbes (TUDresden), Radita Liem (RWTH), Katharina Haus (JSC), Anara Kozhokanova (RWTH), Sandra Mendez (BSC), Christina Mühlbach (TUDresden), Anke Visser (JSC)

Other

- Ayesha Afzal (RRZE), Stefanie Brink (LLNL), Carla Guillen (LRZ), Carmen Navarrete (LRZ), Anastasiia Shamakina (HLRS), Isabel Thärigen (RWTH), Sandra Wienke (RWTH)
- (No women in VI-HPS from Linaro, TUDarmstadt, UOregon, UVSQ)

### Inclusivity advocacy

- Established 2022 (at ISC22 F2F)
  - Topic co-chairs: Judit Giménez (BSC) & Brian Wylie (JSC)
- Covers diversity & equality aspects (DEI/IDEA/JEDI/etc)
  - "woman" & "female" are shorthand for members of all under-represented groups in HPC: women<sup>+</sup>
  - EU/EuroHPC focus on improved representation of EU13 nationals
    - Baltic / V4 (Vysegrad) / Balkan / SouthEastern European countries
    - Ukraine, Moldova?
- Specifics
  - Increasing pool of VI-HPS tools instructors for training events
  - Increasing visibility of female instructors in tutorials, workshops, etc.

## EU13/Widening: under-represented nations



- (Atlantic)
- Baltic
- Vysegrad/V4
- (Austro)Balkan
- SE Europe

### Status

#### **2022**

- (2nd) virtual POP/VI-HPS tuning workshop organized (primarily) by and for women
  - well attended (80% women) and received
- several instructrices and assistants for (regular) VI-HPS tuning workshops; 1 instructrix at conference tutorials (SC)

#### **2023**

- No VI-HPS tuning workshops (virtual or in-person)
  - One planned with no female instructors postponed to Jan 2024
- No instructrices at VI-HPS conference tutorials (ISC & SC)
- No women as co-chairs of ProTools workshop (but well represented on program committee)

#### **2024**

- Perhaps a third POP/VI-HPS tuning workshop organized (primarily) by and for women, ideally in-person
- Perhaps some/all women co-chairs for ProTools@SC workshop proposal

• ...







25th Anniversary

# POP3 training

Brian Wylie, Jülich Supercomputing Centre

HORIZON-EUROHPC-JU-2023-COE



1 January 2024–31 December 2026

Grant Agreement No 101143931

## Ecosystem



- EuroHPC (members+affiliates)
  - Commited: CASTIEL2 (CoEs & NCCs?), EPICURE
  - Optional?: HANAMI, HPC Spectra, HPC Train



3

- EC: Horizon Europe, Digital Europe, DG Connect, etc (EuroHPC+)
- National: ACCESS (USA), ExCALIBUR (UK), CSCS/EPFL/ETHZ (CH)
- Regional: ASEAN, CARLA
- Global: VI-HPS, PRACE?
- Project/Partnership: ETP4HPC, JLESC
- Others...

2024/02/12



### Activities



- Training of POP performance analysts
  - Dedicated tools & methodology training mostly for new POP analysts, probably also for affiliates (EPICURE advanced application support, CoEs?)
  - Regular fortnightly VC sessions for all analysts (Wednesday afternoons)
- VI-HPS Tuning Workshops
- Conference (hands-on) tutorials: ISC, SC?
- Contributions to seasonal schools: ASEAN, CARLA, ESM/TerrSys, IHPCSS
- Bespoke training events
  - Training by&for *women*<sup>+</sup> & other under-represented groups [once per year?]
  - Training workshop series for RSEs
    - Durham/UKRSE, deRSE/HiRSE/NHR/HPC.NRW?, others?
  - Training for projects: DEEP, EUPEX/EuPilot, HPC4fRG, TIME-X

4



## **VI-HPS Tuning Workshops**





2024

- 01.29-02.01: **TW43**, CALMIP, Toulouse/F
  - *Turpan* (Arm64+A100)
    - MAQAO, Scalasca, TAU, Verificarlo
- 02.26-03.01: TW44, RWTHAachen+TUDresden/D
  - CLAIX-2023 & Barnard (SapphireRapids)
    - PIKA/JMS, LIKWID, MAQAO, MUST, Scalasca, TAU, Vampir, lo2s
- 06.10-06.13: **TW45**, LRZ/TUM, Garching/D
  - CooLMUC-2 (Haswell)
    - Caliper, MAQAO, Scalasca, Vampir, etc.

#### TBC

- 09.??: TW??, IT4I/VŠB, Ostrava/CZ
  - Karolina (EPYC+A100)
- 11.??: TW??, NCSA or UTK-ICL, USA (hybrid?)
  - Delta (EPYC+A100)



## Other tools training 2024H1



- Performance Analysis & Tools @ BSC
  - March 13-14
  - *MareNostrumX* (SapphireRapids + H100)
- Durham University Performance Analysis Workshop Series
  - April 4 to May 9 (Thursdays)
  - DiRAC, COSMA/DINE (AMD Rome)
    - Linaro PR, Intel APS/VTune, MAQAO, Scalasca, Reframe
- EPCC/Archer2 Scalasca/analyst training
  - April 29 May 3 (2-3 days)
  - Archer2 Cray/HPE AMD CPU+GPU
- ISC (May 12)
  - POP analysis methodology & Score-P/etc hands-on tutorials proposed

## Other tools training 2024H2?



- July: JLESC/HANAMI (in conjunction with IHPCSS [HPC Spectra])
  - *Fugaku*, R-CCS/Kobe or Tokyo?
- Q?: EPICURE advanced application support
  - Leonardo, CINECA (primarily GPU/Booster, perhaps also CPU/Cluster)
  - (LUMI, CSC in 2023; MareNostrum5, BSC in 2025; JUPITER, JSC in 2026?)
- VI-HPS Correctness & Debugging Workshop
  - France? Only MUST/Archer (RWTH) participation from POP?
- Q3?: by&for *women*<sup>+</sup> and other under-represented groups in HPC
  - EuroHPC priority EU-13/Widening: Baltic/V4(Vysegrad)/Balkan/SE European
  - virtual again, hybrid or in person this time (where)?
  - organisers/hosts? IT4I propose combining with analyst training in September

# Workshops by&for *women*<sup>+</sup>



- 42 VI-HPS Tuning Workshops (tunathons) organised since 2008
  - provide free instruction and coaching in applying VI-HPS (performance) tools
  - target HPC application developers/analysts working on their own parallel codes
    - mostly around (Western) Europe, but also in Chile, Japan & USA
  - over 650 participants, around 80% male (some with no women at all)
  - over 40 instructors, more than 90% male (though 1 prodigious instructrix)
- POP CoE women inspired to improve participation of women<sup>+</sup>
  - identical workshops organized by & specifically for women<sup>+</sup>
    - and also members of other groups under-represented in (European) HPC
  - 8 instructrices & assistants delivered 2 virtual workshops for over 35 women<sup>+</sup>
    - plus several men from waiting list accepted to fill available capacity
  - planning in progress for first workshop to be held in person



## Assessment (preliminary)



- *Women*<sup>+</sup> were very enthusiastic (and curious) workshop participants
  - less inhibited asking questions and seeking assistance
  - more open to sharing and collaborating with other participants
  - few had any experience of other training events by or for women<sup>+</sup>
- But caution warranted from limited early results
  - experienced instructors perceived little difference from mixed-gender events
  - impact on overall workshop participation rates still unclear
  - potential danger of establishing segregated training norm
- However, teaching confidence of our less experienced instructrices boosted immensely
  - keen to join all-gender instructor teams as well
    - including high-profile tutorials at SC & ISC international conferences

## Skills certification (badging)



- Certification of skills [HPC Spectra]
  HPC Certification Forum
  - Performance analyst
    - Capable using one of the POP toolsets for parallel performance measurement & analysis
      - Produced quality performance results from a (familiar) application
      - Distinct for each toolset?
    - Competent applying POP performance assessment methodology
      - Produced quality performance assessment report
  - Instructor/Instructrix
    - Junior: Assisted with (hands-on) tools instruction
    - Senior: Lead (hands-on) tools & methodology instruction
- Certified skills may be recognized by EUMaster4HPC, HPC TRAIN, etc.

12

## Recordings



- Much more valuable when archived than livestreamed
- Generally linked from individual training event pages
  - such as those of each VI-HPS Tuning Workshop
    - though tend to become stale quickly
- Archives distributed
  - Mostly done by host/recorder (with instructor permission)
    - Sometimes behind paywalls, particularly for conference tutorials
  - POP training collection (needing attention)
  - PRACE
    - Unclear whether still maintained
  - YouTube
    - VI-HPS, POP and collections of other random stuff (with advertising)
  - Zenodo/etc
    - Responsibility of individual instructors, counts toward publication record





### **Performance Optimisation and Productivity 3**

A Centre of Excellence in HPC



This project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 101143931. The JU receives support from the European Union's Horizon Europe research and innovation programme and Spain, Germany, France, Portugal and the Czech Republic.

