

# POP training at IT4Innovations and results from PERMON analysis

David Horák

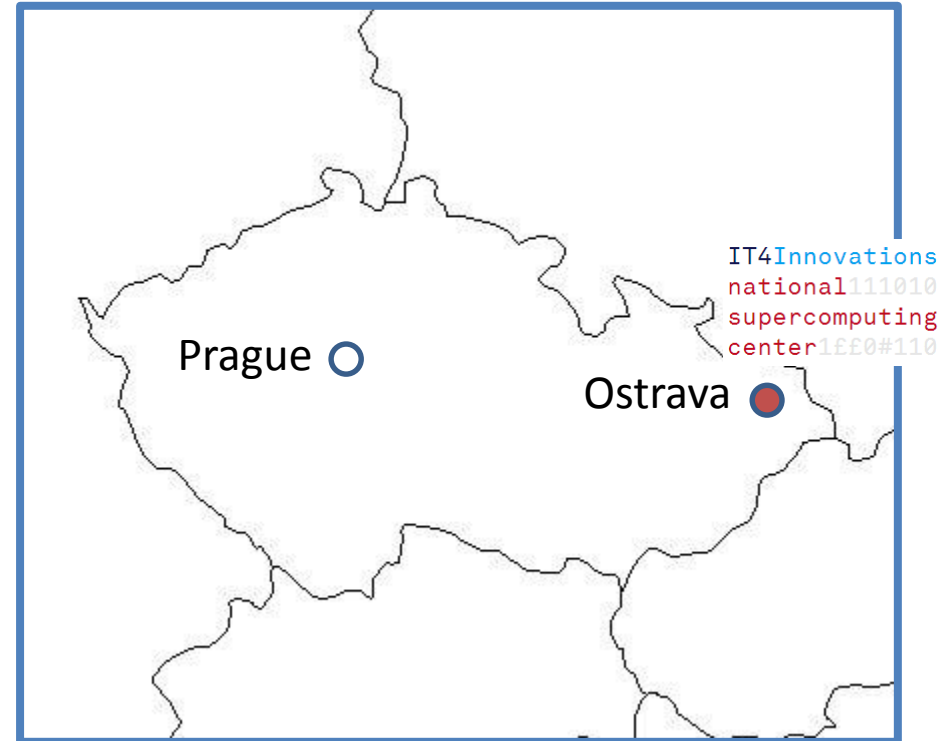
Ondřej Jakl

Radim Sojka

IT4Innovations

National Supercomputing Center

Ostrava, Czech Republic



# Introducing IT4Innovations

One of the six infrastructure centres built 2007-2015 in the Czech Republic using EU structural and investment funds

- established 6/2011
- five partners from academia
  - VŠB-Technical University of Ostrava (coordinator)
  - Brno University of Technology
  - Institute of Geonics, Czech Academy of Sciences
  - University of Ostrava
  - Silesian University in Opava
- 200 FTEs
  - 90% dedicated to the research in IT and HPC
- [www.it4i.cz](http://www.it4i.cz)

# IT4Innovations infrastructure milestones


- **2013:** Anselm supercomputer
- **2014:** IT4Innovations building
- **2015:** Salomon supercomputer



IT4Innovations  
national  
supercomputing  
center

# Salomon – main features

- In 6/2015 No. 40 in TOP500
  - the largest MIC installation in Europe
  - currently (11/2016) No. 67
- 1008 nodes, 24192 cores, 129 TB RAM
- R<sub>peak</sub> 2011 TFlop/s, R<sub>max</sub> ~1500 TFlop/s
- SGI system, Centos 6.6
- 576 regular compute nodes
  - 2x Intel Xeon E5-2680v3 (Haswell)
- 432 accelerated nodes
  - 2x Intel Xeon E5-2680v3 (Haswell)
  - 2x Intel Xeon Phi 7120P
- InfiniBand FDR, SGI Enhanced Hypercube
- SMP/NUMA SGI UV 2000 server, visualization server, tape library,...



# IT4Innovations & PRACE



- National partner of PRACE
- Provider of Tier-1 systems for PRACE
  - contribution with 10% of computational resources of IT4I
- Since 4IP, PRACE provides on-demand training services for/with Centres of Excellence (CoEs)
  - 6 on-demand events organized with CoEs until now
  - among them ...
- A joint event of IT4Innovations and POP- CoE (Performance Optimisation and Productivity Centre of Excellence)


# Performance Optimization and Productivity Tutorial @ IT4Innovations

- 14-15 December 2016
- VŠB - Technical University Ostrava, IT4Innovations building
- 21 participants, mostly from 6 developer teams
- 6 trainers (matching the number of teams)
  - 4 from Germany, 2 from Spain
- Agenda:
  - (1<sup>st</sup> day) performance analysis tools and the methodologies used to create POP performance audit; social event
  - (2<sup>nd</sup> day) hands-on sessions on performance analysis with guidance of POP performance experts



# POP Tutorial – participants

- **6 developer teams**
  - qualified based on a Call addressing application developers interested in a performance audit of their codes
    - 3 from IT4Innovations
    - 1 from Brno University of Technology
    - 1 from Institute of Geonics, Czech Academy of Sci.
    - 1 from Czestochowa University of Technology (Poland)
- **Preparatory course Practical Parallel Performance Analysis on Salomon**
  - conducted at IT4Innovations in October 2016 by POP experts



# POP Tutorial – codes

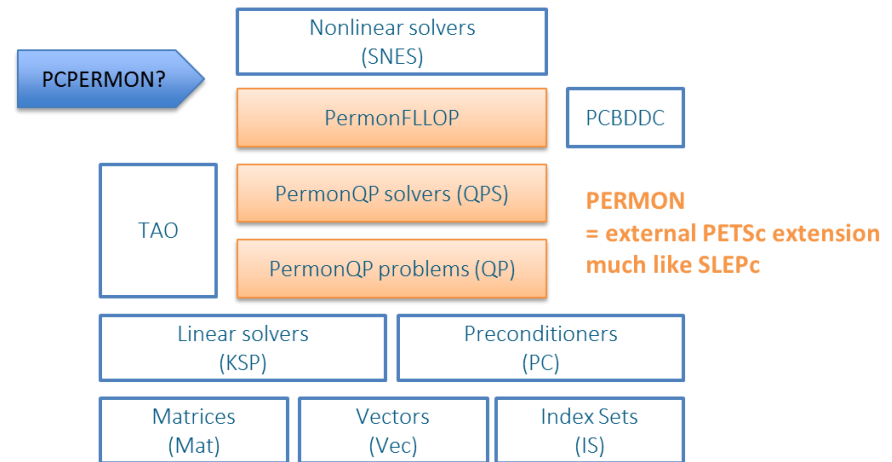
- PERMON, IT4Innovations, contact Václav Hapla  
toolbox based on FEM and domain decomposition methods of the FETI type
- BEM4I, IT4Innovations, contact Michal Merta  
library of parallel BEM based solvers
- Blender, IT4Innovations, contact Milan Jaroš  
modified kernel of the Blender Cycles 3D rendering engine to support the HPC environments
- GEM, Institute of Geonics, contact Jiří Starý  
modelling of thermo-mechanical processes, solution of large-scale problems especially in geosciences
- k-Wave, Brno University of Technology, contact Jiří Jaroš  
simulation of ultrasound waves moving through soft tissues
- MPDATA, Czestochowa University of Technology, contact Krzysztof Rojek  
stencil computations



# PERMON – brief info

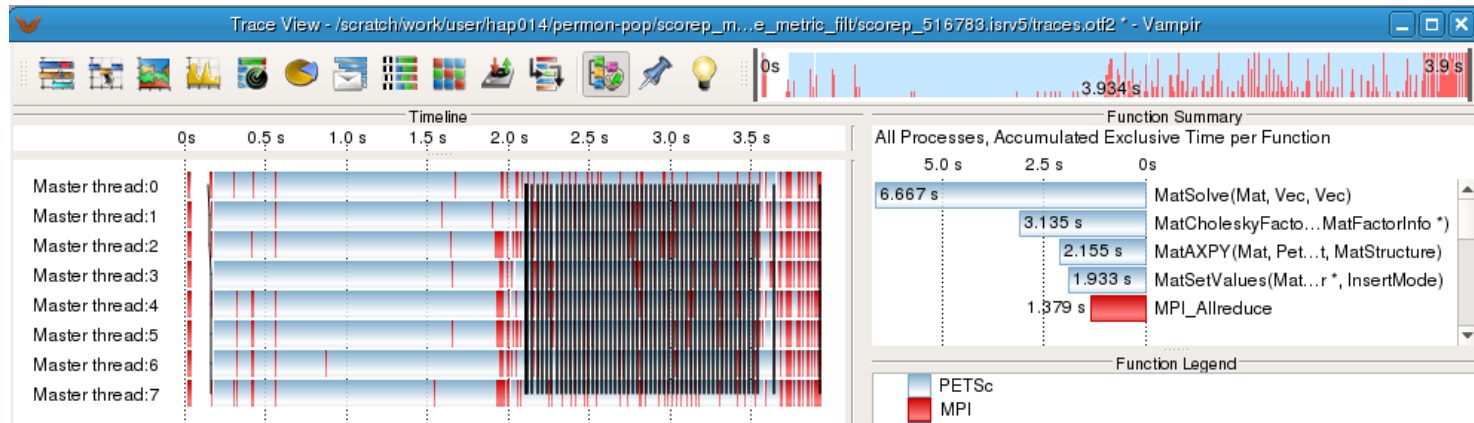


- **Parallel, Efficient, Robust, Modular, Object-oriented, Numerical**
- **Based on/extending PETSc – pure MPI implementation**
- **PermonQP**
  - generic quadratic programming solution framework
  - QP problems, transforms, solvers
  - published under FreeBSD license
- **PermonFLOP**
  - FETI DDM implementation
  - extends PermonQP
- **PermonSVM**
  - machine learning package, Support Vector Machines approach
  - uses PermonQP
- **For more info and codes visit <http://permon.it4i.cz>**



# PERMON analysis (1)


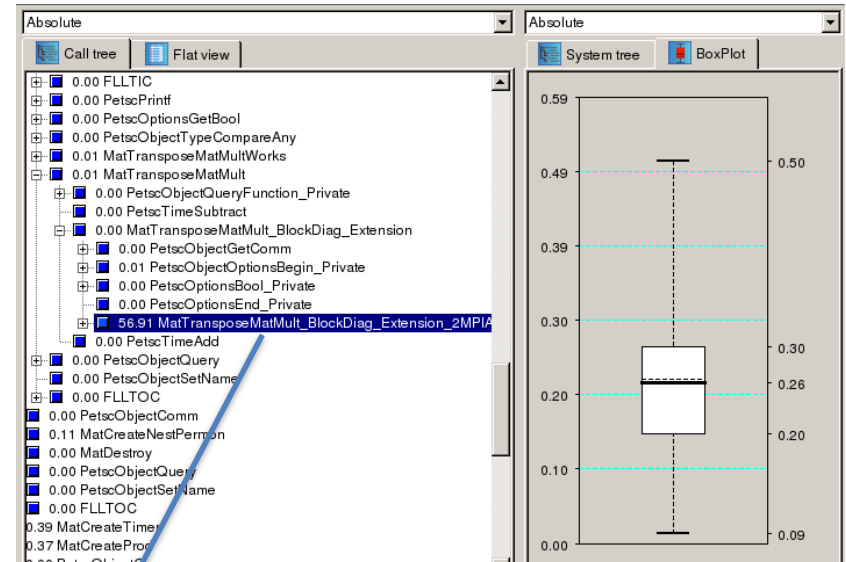
- Most time spent in PETSc library or in direct solver libraries
  - User Library Wrapping (libwrap) employed, to make PETSc calls visible (on-going collaboration with Ronny Brendel)



- PERMON targets for massively parallel computations
  - some bottlenecks may occur in runs on more than thousand processing elements (cores)

# PERMON analysis (2)

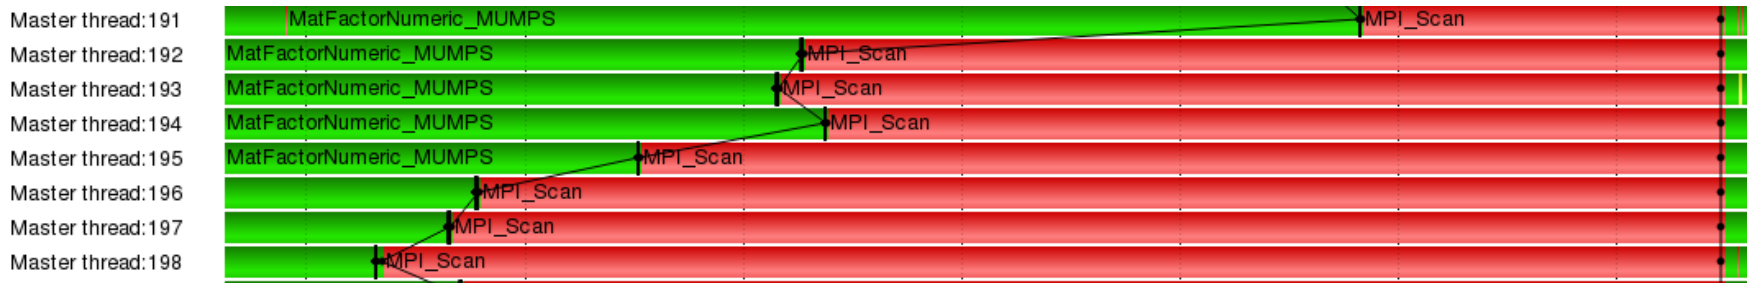
- Load imbalance detected via Score-P and Cube in function for Matrix-Matrix multiplication
  - with 216 MPI processes, the ratio of maximum to minimum runtime for all processes is 5.4
  - the ratio increases with increasing number of processes
    - ratio 96 detected on a larger benchmark
  - most time spent in MPI\_Scan



# PERMON analysis (3)

- **Vampir analysis/visualization**

- synchronization of processes by MPI\_Scan -> the imbalance is caused in previous function MatFactorNumeric from the external MUMPS library



- **Possible mathematical solution**

- decompose problem into more subdomains -> more matrix factorizations with lower computational complexity performed by MUMPS -> better load balancing between processes

# Lessons learned & Conclusions

- **Event unique in high concentration of developer teams**
  - not every developer (team) does want to undergo a third-party performance analysis and optimization
- **2 teams currently continue collaboration with POP**
  - k-Wave, PERMON
  - it takes some time to find a mode of a beneficial collaboration
- **PERMON – POP plans for the future: optimization of codes for SVM in the ExCAPE project**
- **PRACE support appreciated**

# Lessons learned & Conclusions

- **Event unique in high concentration of developer teams**
  - not every developer (team) does want to undergo a third-party performance analysis and optimization
- **2 teams currently continue collaboration with POP**
  - k-Wave, PERMON
  - it takes some time to find a mode of a beneficial collaboration
- **PERMON – POP plans for the future: optimization of codes for SVM in the ExCAPE project**
- **PRACE support appreciated**
  
- **Professional optimization – a standard of the future?**



Thank you for your attention.

Questions?



# IT4Innovations National Supercomputing Center

- Provides the research community with high-performance computer systems and corresponding services (training, etc.)
  - including the most powerful scientific supercomputer in the CR
- One of the three pillars of the Czech ICT/e-infrastructure, next to
  - CESNET – the National Research and Education Network (NREN)
  - CERIT-SC – the National Grid Infrastructure (NGI)

# Research at IT4Innovations

- RP 1 – IT for disaster and traffic management
- RP 2 – Numerical modelling for engineering
- RP 3 – HPC libraries and supercomputing for industry
- RP 4 – Modelling for nanotechnologies
- RP 5 – IT for knowledge management
- RP 6 – Soft computing methods with supercomputer applications
- RP 7 – Multimedia information recognition and presentation
- RP 8 – Secure and safe architectures, networks and protocols