The POP Centre of Excellence - Improving Parallel Codes

Craig Lucas & Phil Tooley (NAG)

EU H2020 Centre of Excellence (CoE)

1 December 2018 – 30 November 2021

Grant Agreement No 824080

POP CoE

- A **Centre of Excellence**
  - On **Performance Optimisation and Productivity**
  - Promoting **best practices in parallel programming**

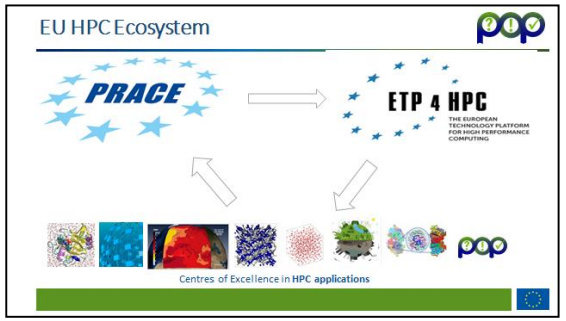

- Providing **FREE** Services
  - Precise understanding of application and system behaviour
  - Suggestion/support on how to refactor code in the most productive way
  - Or confirmation that all is good!
  - Approaching 200 services so far.
- Horizontal
  - Transversal across application areas, platforms, scales
- For (EU) **academic AND industrial** codes and users!

POP Team

- **Who**
  - BSC, ES
  - HLRS, DE
  - IT4I, CZ
  - JSC, DE
  - NAG, UK
  - RWTH Aachen, DE
  - TERATEC, FR
  - UVSQ, FR
- **A team with**
  - Excellence in performance tools and tuning
  - Excellence in programming models and practices
  - R & D background in real academic and industrial use cases
- **Developing open-source tools**
  - Extrae, Paraver & Dimemas
  - Score-P, Cube & Scalasca
  - MAQAO

Motivation

- **Why?**
  - HPC machines and codes getting ever more complex
    - ⇨ Difficult to build quantitative picture of application behaviour
    - ⇨ Often unclear what is best approach to improve performance
  - Important to maximize efficiency (performance, power) of compute intensive applications and productivity of the development efforts
- **What?**
  - Parallel programs, mainly MPI/OpenMP
  - Although also CUDA, OpenCL, OpenACC, Python, …

Services provided by POP

- **Parallel Application Performance Assessment**
  - Identifies performance issues of customer code (at customer site)
  - Possibly restricted to a “region of interest”
  - If needed, identifies the root causes of the issues found and qualifies and quantifies approaches to address them (recommendations)
  - A report is produced
  - Medium effort (1-3 months)

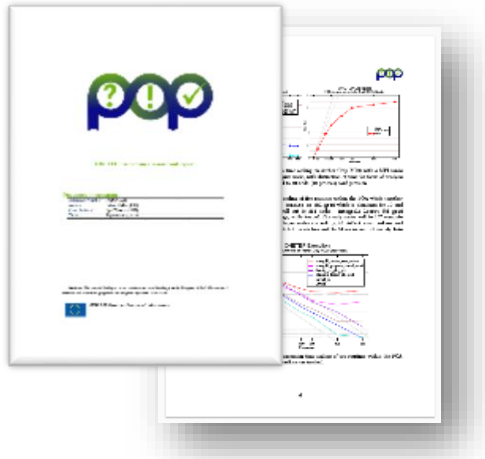

- **Proof-of-Concept**
  - Follow-up service
  - Experiments and mock-up tests for customer codes
  - Kernel extraction, parallelisation, mini-apps experiments to show effect of proposed optimisations
  - Larger effort (3-6 months)

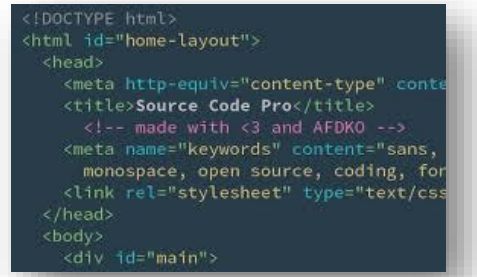

Note: Effort shared between our experts and customer!

POP Metrics

- The following metrics are used in a POP Performance Assessments:
- A numerical measure between 0.0 (bad) and 1.0 (perfect)
- Global Efficiency (GE): GE = PE * CompE
  - Parallel Efficiency (PE): PE = LB * CommE
    - **Load Balance** Efficiency (LB): LB = avg(CT)/max(CT)
    - **Communication** Efficiency (CommE): CommE = SerE * TE
      - Serialization Efficiency (SerE):
        SerE = max (CT / TT on ideal network)
      - Transfer Efficiency (TE): TE = TT on ideal network / TT
  - (Serial) **Computation** Efficiency (CompE)
    - Computed out of IPC Scaling and Instruction Scaling
    - For strong scaling: ideal scaling -> efficiency of 1.0

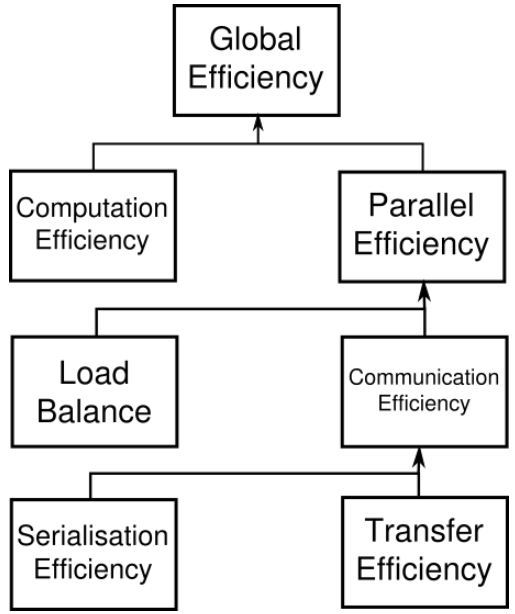

CT = Computational time
TT = Total time

POP Metrics – Parallel Efficiency 1

- Consists of **Load Balance** Efficiency (LB) which reflects how well the distribution of work to threads/processes is done in the application:

$$LB = \frac{\text{average computational time}}{\text{max computational time}}$$

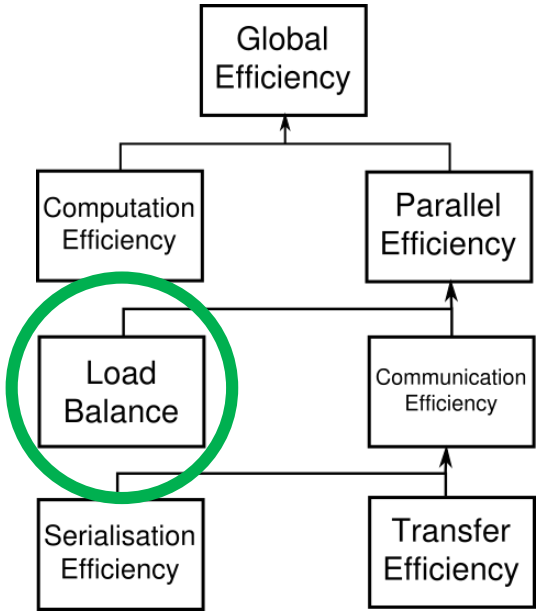

Example 1: good load balance (LB = 100%)

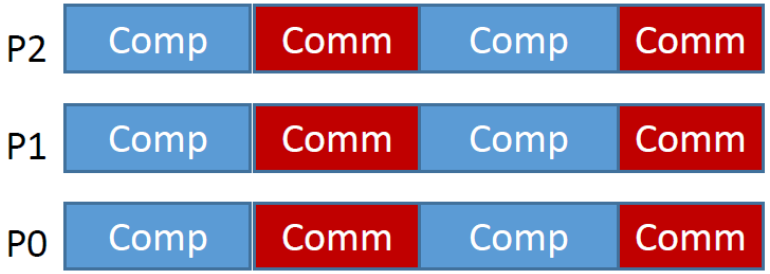

Example 2: bad load balance (LB = 77%)

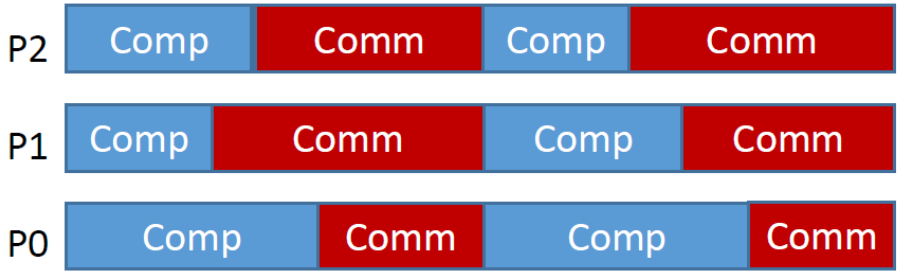

CT = Computational time
TT = Total time

POP Metrics – Parallel Efficiency 2

- And **Communication Efficiency**: CommE = SerE * TE
  - **Serialization Efficiency** describes loss of efficiency due to dependencies between processes, waiting in MPI, say:
    SerE = max (CT / TT on ideal network)
  - **Transfer Efficiency** describes loss of efficiency due to actual data transfer:
    TE = TT on ideal network / TT

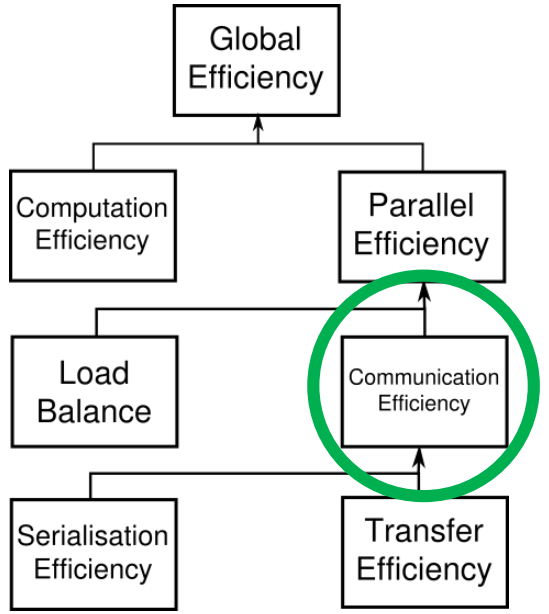

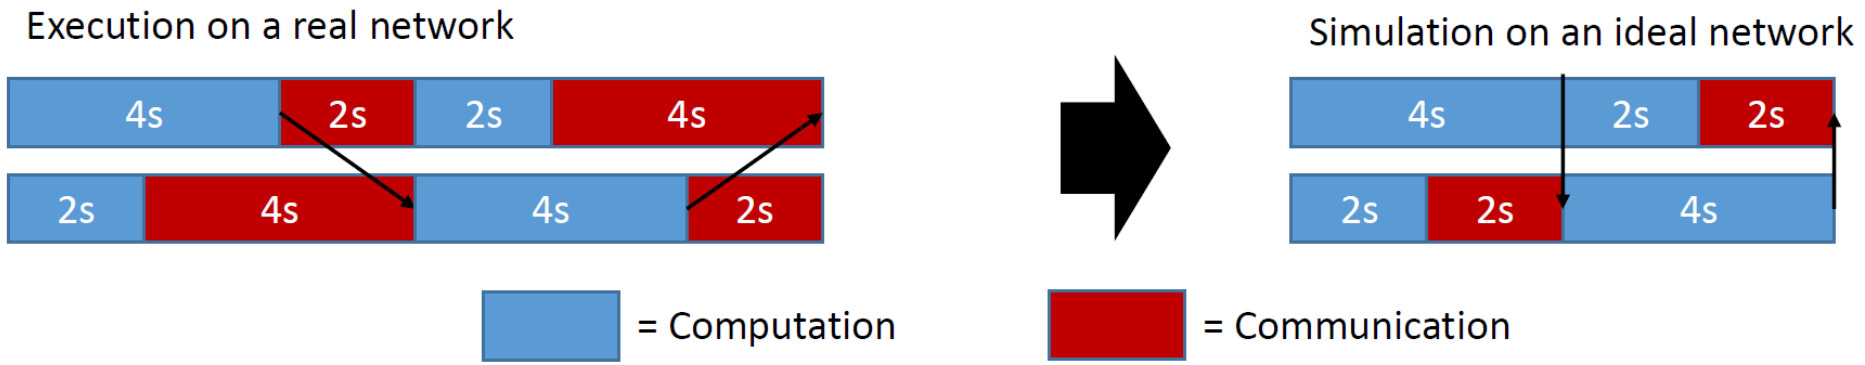

CT = Computational time
TT = Total time

$$SerE = \frac{6}{8} = 75\%$$

$$TE = \frac{8}{12} = 66.6\%$$

POP Metrics – Comp. Efficiency

- (Serial) **Computation Efficiency** (CompE)
  - Computed out of IPC Scaling and Instruction Scaling
  - Describes how well the computational load of an application scales with the number of processes.
- **Instruction Scaling** compares the total number of instructions executed for a different number of threads/processes.
- With more processes more instructions may be executed, e.g. some extra computation for the domain decomposition is needed.
- **IPC Scaling** compares how many instructions per cycle are executed for a different number of threads/processes.
- The same number of instructions is computed but the computation takes more time, say. This can happen e.g. due to shared recourses like memory channels.

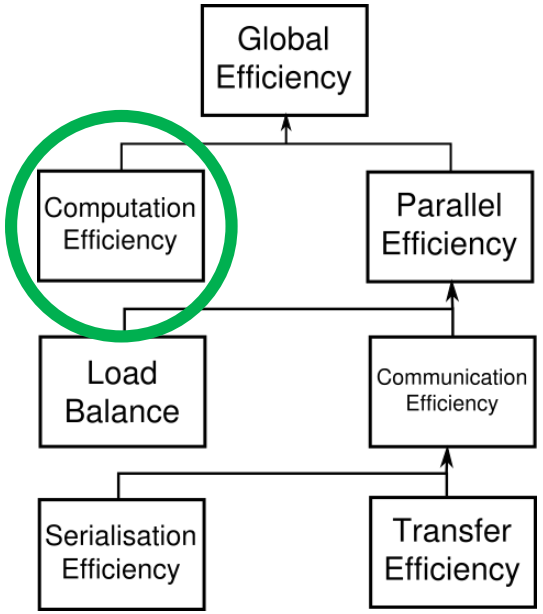

CT = Computational time
TT = Total time

POP Metrics

- Simply presented as:

			2	4	8	16
Global efficiency			0.98	0.90	0.78	0.59
X	**Computation Efficiency**		1.00	0.96	0.90	0.76
	Parallel Efficiency		0.98	0.94	0.90	0.85
	X	**Load Balance**	0.99	0.97	0.91	0.92
		Communication Efficiency	0.98	0.97	0.98	0.92
		X **Serialization efficiency**	0.99	0.98	0.99	0.94
		Transfer Efficiency	0.99	0.99	0.99	0.98

X	**IPC Scaling**		1.00	0.99	0.96	0.84
	Instruction Scaling		1.00	0.97	0.94	0.91

POP and CompBioMed

- POP working with CompBioMed on flagship code HemeLB. (C++ & MPI)
- Analysis helped in development of a new version of HemeLB including a new load-balancing library from E-CAM

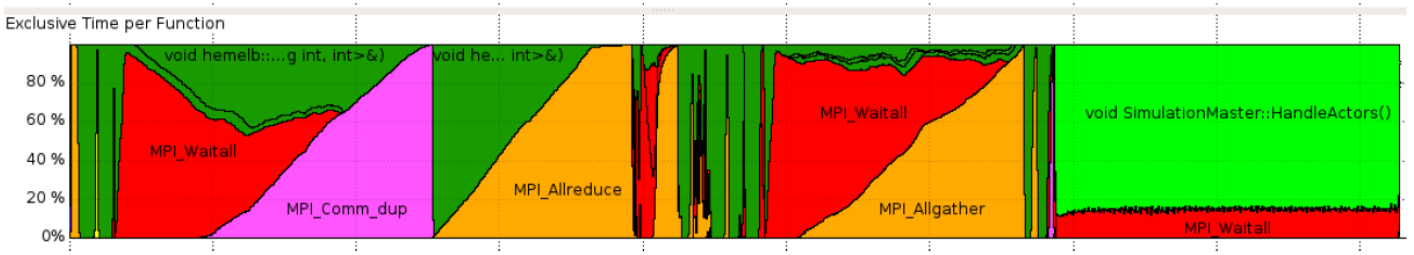

- Studies show excellent communication efficiency and strong scaling (up to 239,615 MPI processes Blue Waters Cray XE.)
- Observations also include load balancing affected by one (weak) process per node.
- Would benefit from reduction in memory footprint as some cores unused on node.

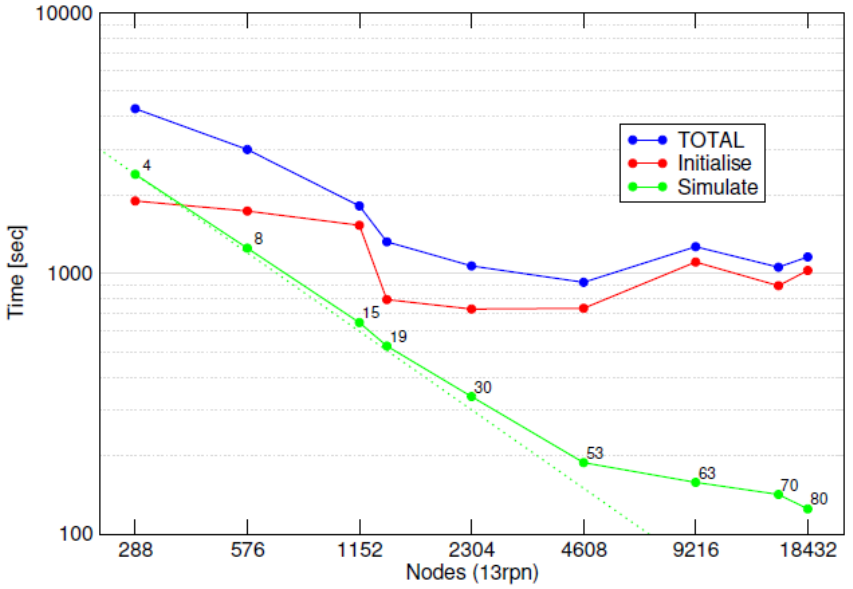

[See summary and reports]

Return on Investment Examples

Application Savings after POP Proof-of-Concept

- POP PoC resulted in 72% faster-time-to-solution
- Production runs on ARCHER (UK national academic supercomputer)
- Improved code saves €15.58 per run
- Yearly savings of around €56,000 (from monthly usage data)

Application Savings after POP Performance Assesment

- Cost for customer implementing POP recommendations: €2,000
- Achieved improvement of 62%
- €20,000 yearly operating cost
- Resulted in yearly saving of €12,400 in compute costs ⇨ ROI of 620%

POP Website

- www.pop-coe.eu
- All the information you need to access POP services, and…

- Blogs
- Newsletter: subscribe and archive
- Information on our webinar series, also…

YouTube

- See our YouTube Channel

  https://www.youtube.com/pophpc

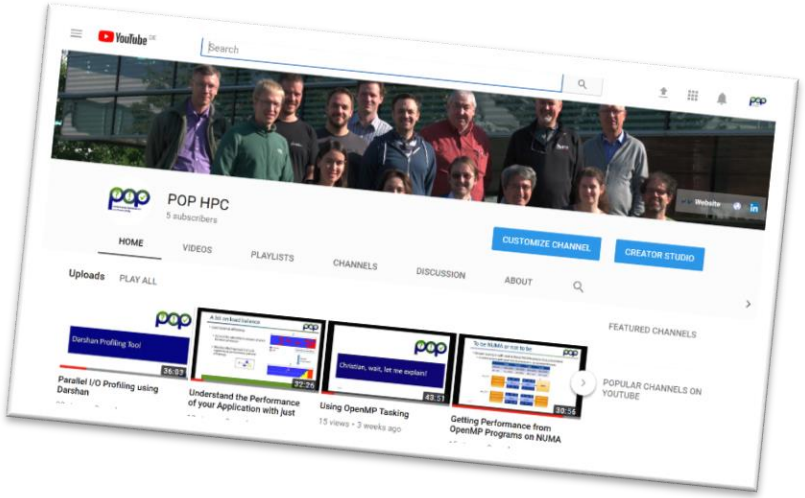

- Recordings available:
  - How to Improve the Performance of Parallel Codes
  - Getting Performance from OpenMP Programs on NUMA Architectures
  - Understand the Performance of your Application with just Three Numbers
  - Using OpenMP Tasking
  - Parallel I/O Profiling Using Darshan
  - The impact of sequential performance on parallel codes
  - Large scale Application Execution Performance Assessment

The POP Process

- **When?**

  POP runs to November 2021

- **How?**
  - Fill in small questionnaire describing application and needs https://pop-coe.eu/request-service-form
  - Questions? Ask pop@bsc.es

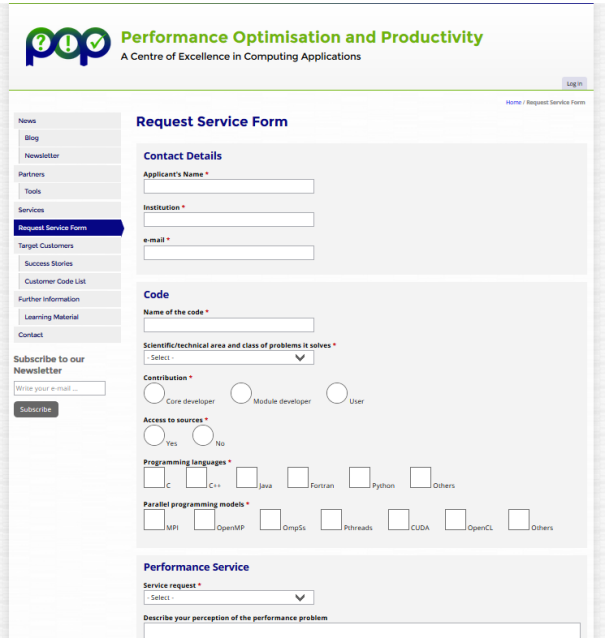

  - Install tools on your production machine (local, PRACE, …)
  - Collaboratively: Gather data → Analysis → Report

Performance Optimisation and Productivity

A Centre of Excellence in HPC

Contact:

https://www.pop-coe.eu

mailto:pop@bsc.es

@POP_HPC

NAG is hiring for POP!

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 676553 and 824080.