Application Performance Evaluation and Optimisation Strategy in EoCoE

Paul Gibbon

EoCoE Objectives

*Help accelerate transition to **low-carbon electricity***

Meteorology for energy: Very short term forecasting for solar and wind energy – extreme events

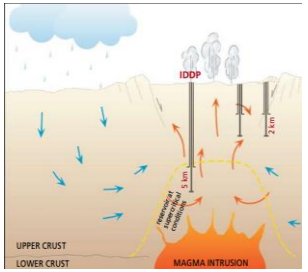

Water for energy: Geothermal and hydro-power – management of resources, usage strategy, influence of climate change.

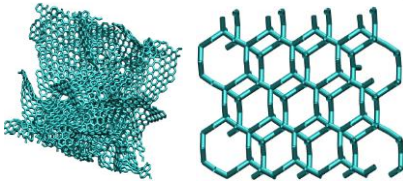

Materials for energy: Photovoltaic cells, batteries and super-capacitors

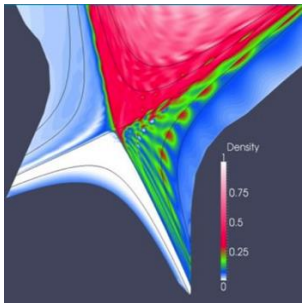

Fusion for energy: Coupling kinetic and fluid codes, mesh alignment with equilibrium configurations, core-edge interaction

Transversal Basis: the 5th element

EoCoE co-design objectives:

- Advanced programming methods and tools (FZJ)
- System tools for high performance (CEA)
- Linear algebra (INRIA)
- Numerical methods – applied mathematics (INRIA)

EoCoE-POP workshops

- **28 participants:** 17 scientists + 11 HPC experts
- **Training 1 day:**
  - JUBE by Sebastian Luehrs
  - Scalasca by Michael Knobloch
  - Paraver by Judit Gimenez
- **Hands-on benchmarking 2 days:**
  - Tools integration in application
  - Performance issue detection
- **Follow up:**
  - *Code teams:* application developers + **HPC mentor**
  - *Continual automated monitoring*: JUBE + performance tools
  - *Deep re-factoring:* math library/ algorithm/ accelerator experts

EoCoE tool development: JUBE

Integration of each code into the EoCoE **automatic metric extraction process** to monitor program improvements on different platforms and optimize code exchangeability

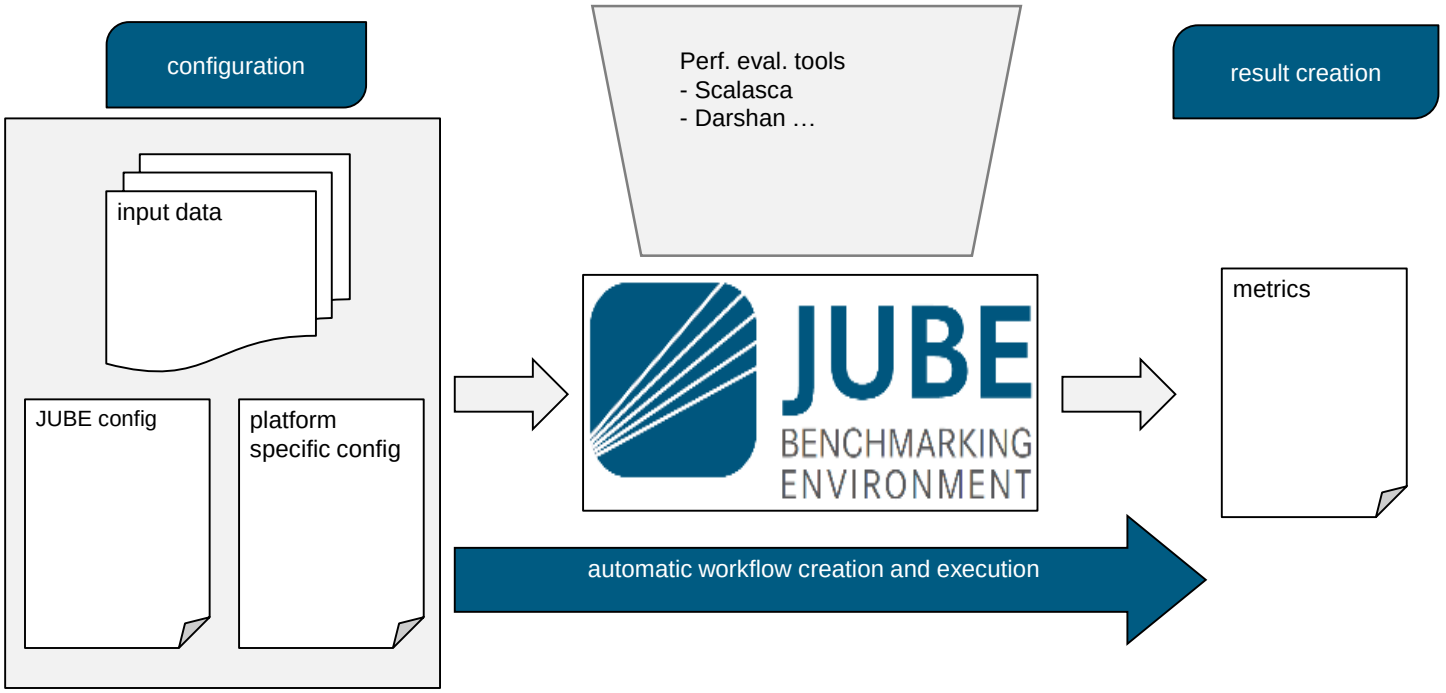

Automatic metric generation

	Metric name	fayalite.json	fayalite_io.json
Global	Total Time (s)	25	19
	Time IO (s)	0.32	0.13
	Time MPI (s)	23.20	15.03
	Memory vs Compute Bound	1.00	1.06
	Load Imbalance (%)	55.56	43.88
IO	IO Volume (MB)	197.73	197.73
	Calls (nb)	3113003	27519
	Throughput (MB/s)	616.35	1571.13
	Individual IO Access (kB)	0.06	7.70
MPI	P2P Calls (nb)	15352	15352
	P2P Calls (s)	0.34	0.33
	P2P Calls Message Size (kB)	0	0
	Collective Calls (nb)	8062	8062
	Collective Calls (s)	21.94	13.88
	Coll. Calls Message Size (kB)	2357	2357
	Synchro / Wait MPI (s)	21.79	13.72
	Ratio Synchro / Wait MPI (%)	92.74	90.32
Node	Time OpenMP (s)	N.A.	N.A.
	Ratio OpenMP (%)	N.A.	N.A.
	Synchro / Wait OpenMP (s)	N.A.	N.A.
	Ratio Synchro / Wait OpenMP (%)	N.A.	N.A.
Mem	Memory Footprint	148336kB	122632kB
	Cache Usage Intensity	0.97	0.97
Core	IPC	0.66	0.66
	Runtime without vectorisation (s)	25	20
	Vectorisation efficiency	1.00	1.05
	Runtime without FMA (s)	25	20
	FMA efficiency	1.00	1.05

Example: CP2K buffered I/O improvement

Work in EoCoE:

- Definition of specific metrics in EoCoE
- JUBE enhancements to include performance tools in work-flow

Result/Status:

- Metrics automatically obtained by developers for chosen tool
- Direct and continuous monitoring of code changes and optimisations
- Regression tests and verifications
- Workflow has been successfully applied to **21 codes** in **3 workshops**

Code progress charts

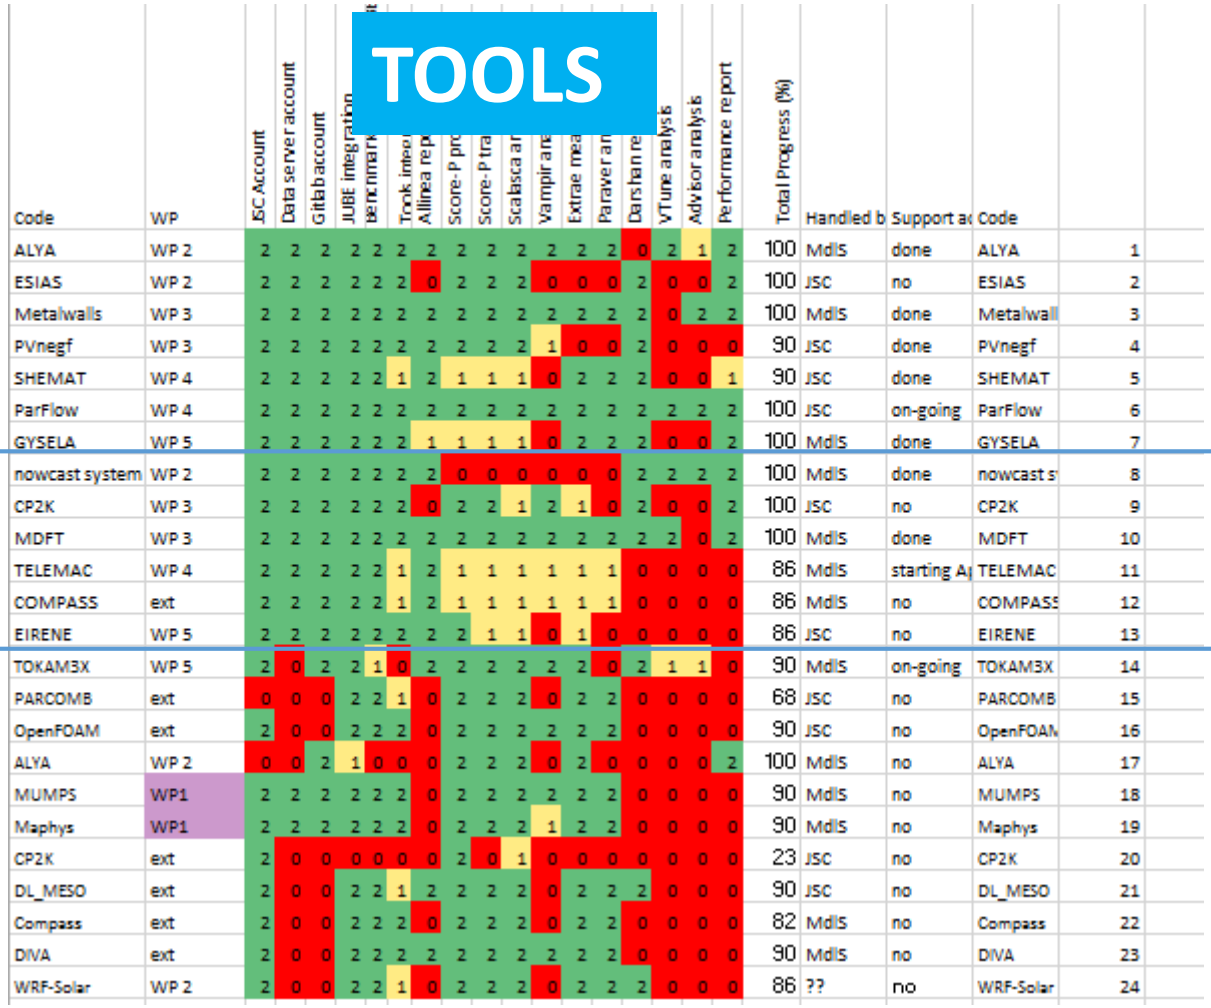

CODES

Workshop 1

Workshop 2

Workshop 3

t.b.c.

Application monitoring with JUBE

ESIAS - EoCoE Code Diary

Code developers	Jonas Berndt (Hendrik Elbern, Charlotte Hoppe)
WP1 HPC experts	Sebastian Lührs
POP HPC experts	Brian Wylie
Location of code on EoCoE site	https://gitlab.eocoe.eu/EOCOE/Esias

CODE TEAM

Workflow element	end of workshop 1	29/02/2016	31/03/2016	30/06/2016	30/09/2016	31/12/2016
JUBE integration	1	2	2	2		
Benchmark definition in JUBE	1	1	1	2		
Tools integrated in JUBE	0	0	1	2		
Allinea report	0	0	0	0		
Score-P profile	2	2	2	2		
Score-P trace	1	1	1	2		
Scalasca analysis	1	1	1	2		
Vampir analysis	0	0	0	0		

Legend	
0	not started
1	in progress
2	established

Platform: JUQUEEN

Metric number	Global level	Unit prop. MH	Tool(s) used	end of workshop 1	2/29/2016	3/31/2016	6/30/2016
Glob.1	Total time elapsed	[s]	time	1,550	838	838	672
Glob.2	Total time IO	[s]	Darshan			79	87
Glob.3	Total time MPI	[s]	Scalasca	398	398	398	355
Glob.4	Memory vs compute bound	1 <= ratio <= 2					

BENCHMARK HISTORY

EoCoE-POP synergy

- Standardisation of performance metrics (subset of POP metric set)
- Integration of performance metrics in JUBE
- EoCoE teams and guests introduced to perf. analysis tools
- Wide set of issues identified and strategies to resolve them proposed quickly
- Follow-up activity taps into 'off-line' POP expertise
- Several EoCoE apps requested POP audits after workshop

Extreme event nowcasting

- Code: **ESIAS**
- PIs: Henrick Elbern, Jonas Bernd, IEK-8, FZ-Jülich
- HPC mentors: Sebastian Lührs, Morris Riedel JSC
- Topic: Meteorological Big-Data analytics

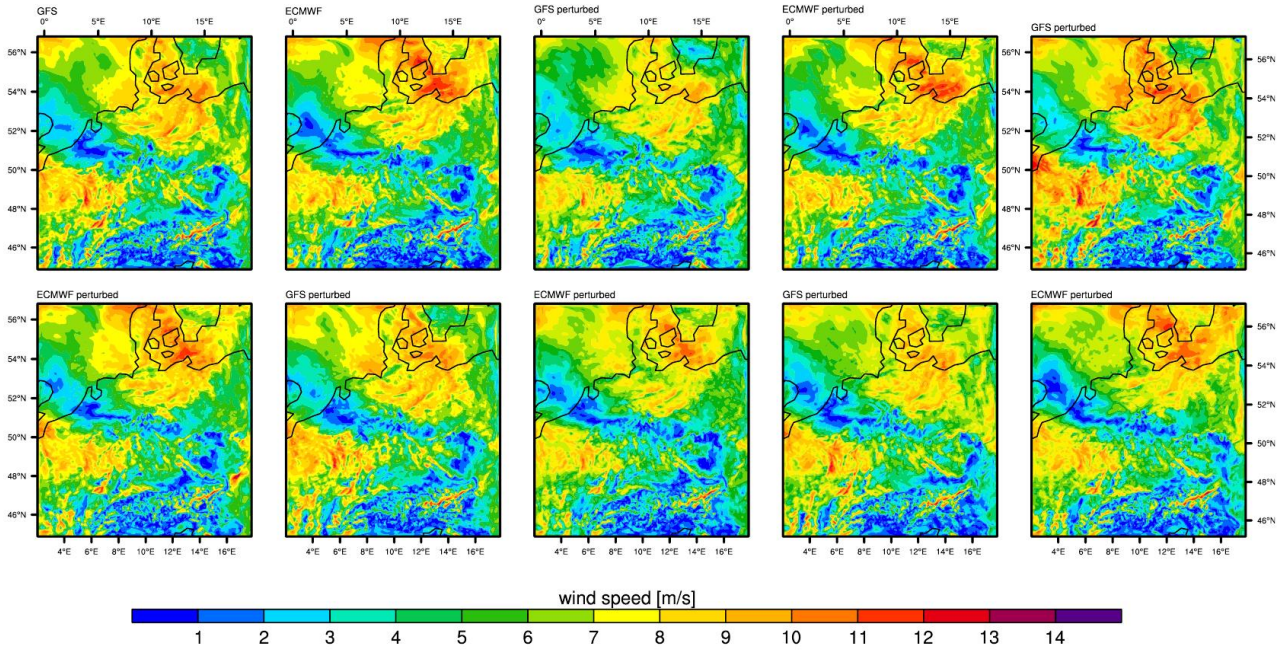

ALYA wind farm modelling

- PI: Herbert Owen, BSC
- Code: **ALYA**
- HPC mentor: Yacine Ould-Rouis, MdlS Saclay; Luc Giraud, INRIA
- Topic: Wind farm modelling

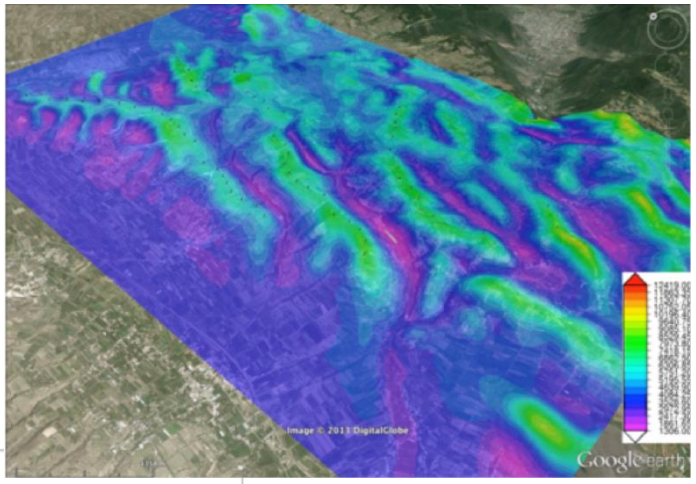

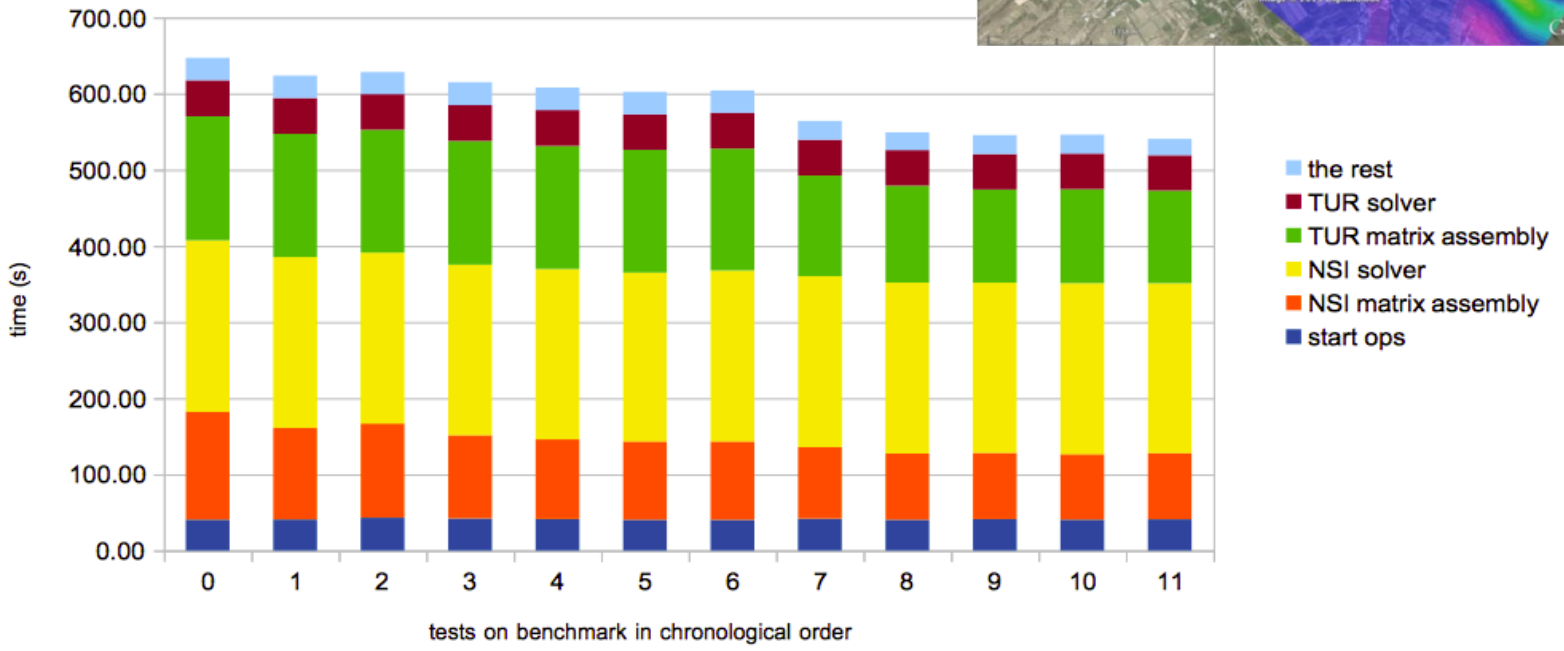

Photovoltaics (WP3: Materials)

- Code: **PVnegf**
- PIs: Urs Aeberhard, Uwe Rau, (IEK-5, FZ-Jülich)
- HPC mentors: Edoardo Di Napoli, Thomas Breuer JSC
- Topic: Multiscale solar cell modelling

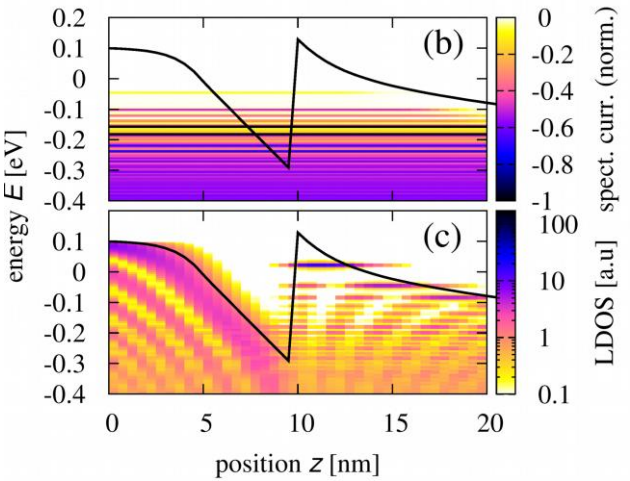

Initial state: No parallelism at all!

Actions two-fold:

i) Implementation of OpenMP

ii) Re-design core algorithm by synthetic examples

Results:

i) roughly 8-fold speed-up ⇨ already in production

ii) *experimental 1D version* scales to 0.5M cores of JUQUEEN

Heat and fluid flow in porous media

Code: **SHEMAT**
PI: Jan Niederau, (RWTH Aachen)
HPC-mentors: S. Lührs, R. Halver

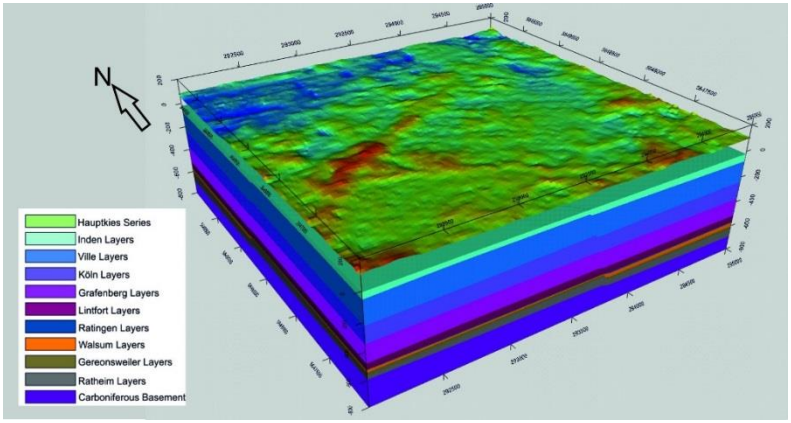

Initial state:

- Slow input file handling

Action:

- Implement parallel I/O with HDF5

⇨ Request to EoCoE I/O experts in WP1, Task 3

Result:

- *50x speed-up* in parsing of large, 200MB inputs
- Now provides option for distributed data management which is essential for Exascale readiness

Plasma turbulence modelling

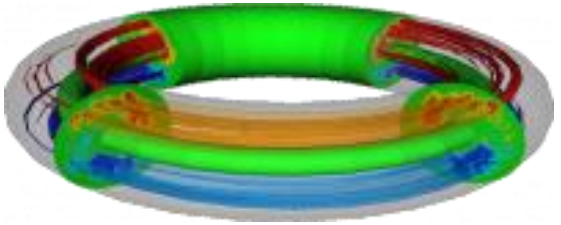

Code: **Gysela**
PI: Virginie Grandgirard
HPC mentors: Guillame Latu, Julien Bigot

Initial state:

- room for improvement in *load balancing*

Actions:

- Hyper-threading + fine-task management
- use StarPU software for dynamic scheduling

Result:

- first gain accomplished with room for more

- Hyper-threading & fine-task management ⇒ lower execution time by improving load balancing

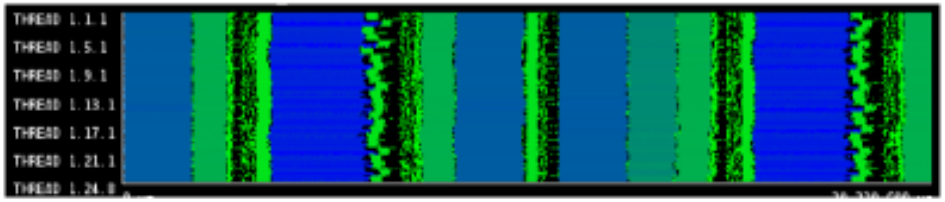

Still idle time…

EoCoE application progress

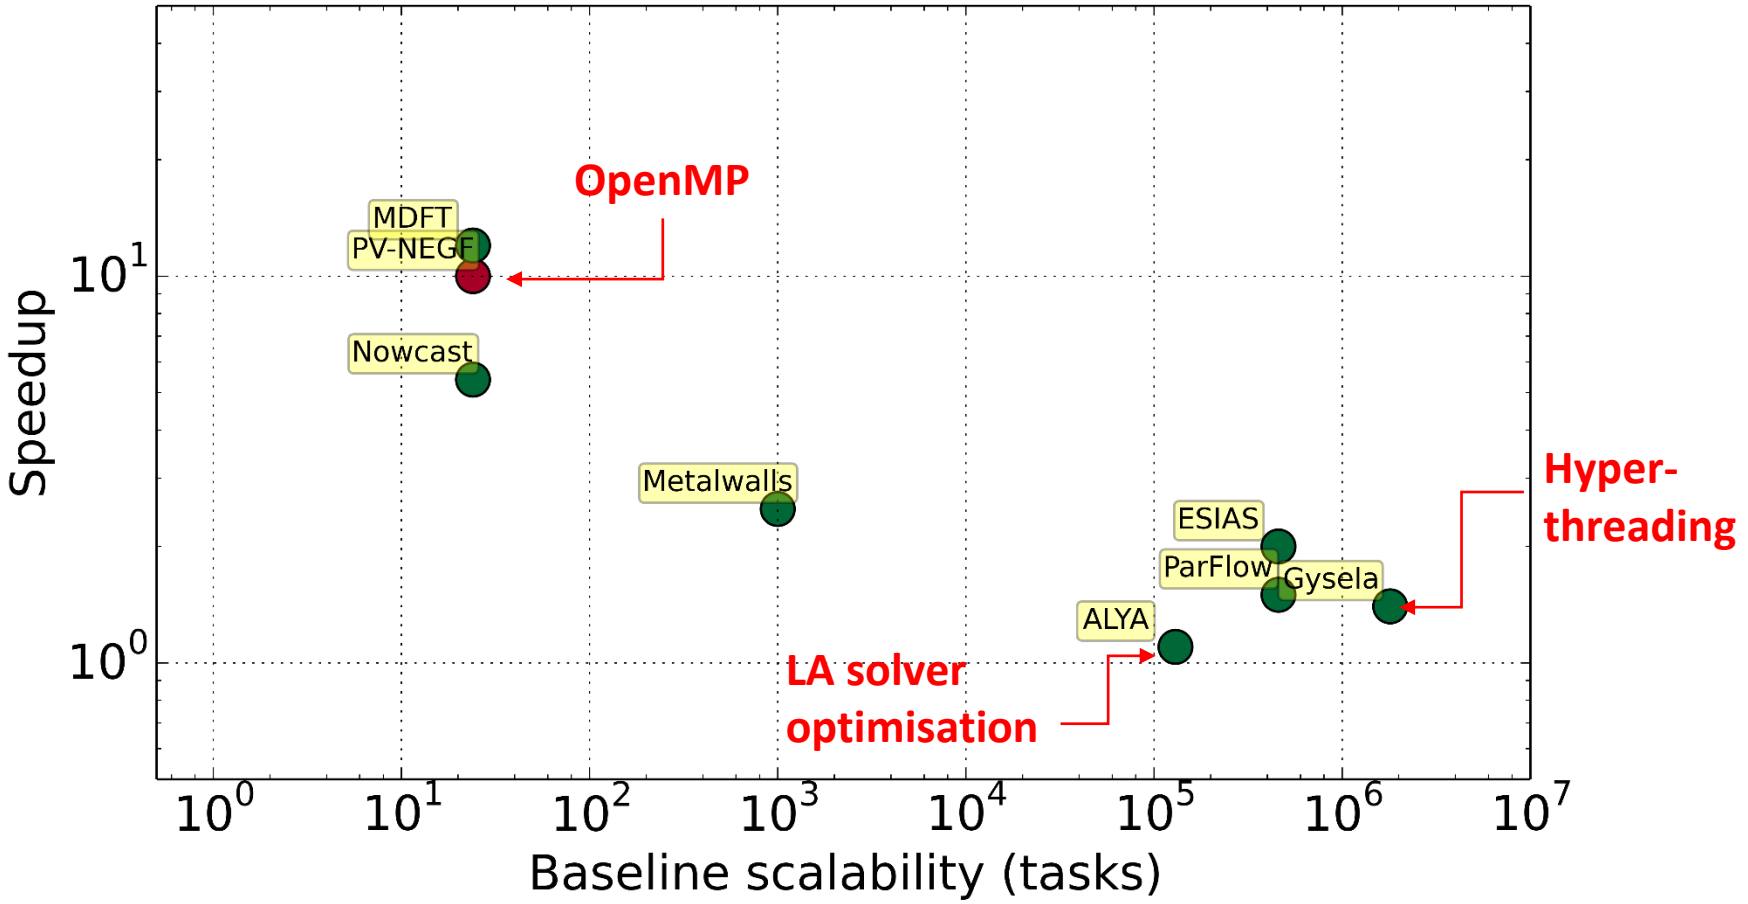

Turning Exascale into benefit

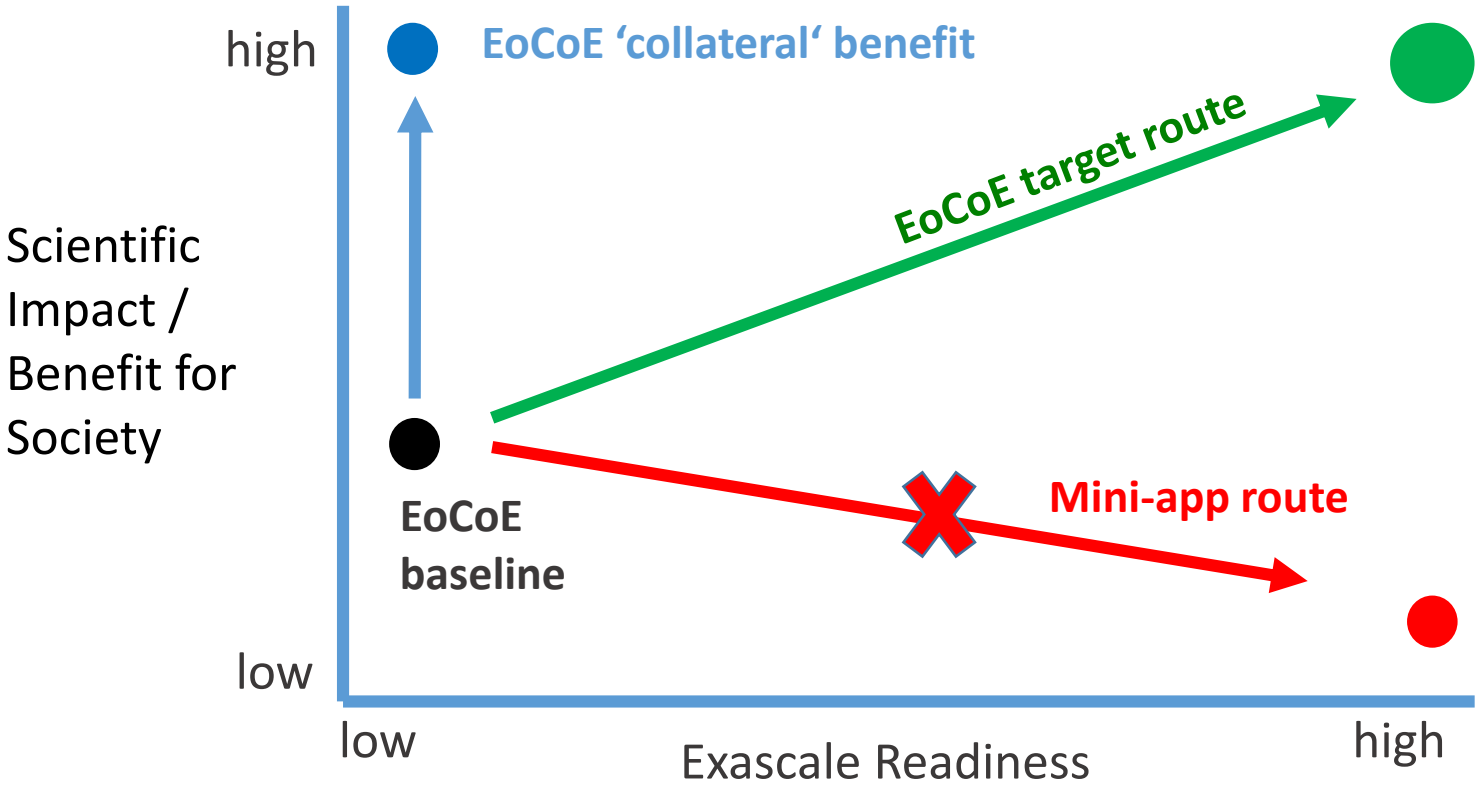

EoCoE exa-readiness pathway

Determine Baseline	Optimisation Strategies	Explore Hardware Architectures
Automatic workflow: • code profiling • extracting metrics • identify bottlenecks	Code teams: • HPC mentor + developer(s) • *address* bottlenecks • automatic workflow • continual code monitoring	- Use automatic workflow to explore next-gen HW - Compare and evaluate - Iterate development process on a continual basis

Acknowledgements

- EoCoE WP1: Matthieu Haefele, Yacine Ould-Rouis, Matthieu Lobel, Abel Marin-Lafléche, Daniel Rohe, Sebastian Lührs, Wendy Sharples, Thomas Breuer, Rene Halver, Benedikt Steinbusch

- POP: Judit Gimenez, Brian Wylie, Ilya Zukhov, Michael Knobloch, Gemàn Llort …

- 20+ EoCoE members from WP2-5

Thank you!

Any questions?