




ChEESE

Exascale ambition towards D2.4: status

Team EARTH
CINECA



WP2 Tasks



WP2 tasks				
Task	Name	Lead partner	Start month	End month
2.1	Code audits	HLRS	3(27)	15(36)
2.2	Increasing the performance portability of flagship codes for different accelerators	CIN	1	48
2.3	Optimisation for Exascale	CIN	6	48
2.4	Continuous Performance Efficiency Monitoring	BSC	18	48



WP2 Deliverables



D	Deliverable name	Short description	Lead	Type	Level	M
2.1	First code audit report	First code audit round results and the recommendations arising from the analysis	HLRS	R	PU	12
2.2	First report on improving code performance at exascale	Report on porting, optimization and tuning for the selected applications on the available hardware	CIN	R	PU	24
2.3	Second code audit report	Update of D2.1 and further recommendations arising from the analysis	HLRS	R	PU	30
2.4	Final report on improving code performance at exascale and Best Practice Guide	Technical report describing the achieved scalability level and related performances obtained for the selected applications on the available hardware including a final Best Practice guide	CIN	R	PU	48



ToC of D2.4 (updated draft)


- 1. Executive Summary
- 2. Exascale ambition final update
- 3. Flagship codes improvements highlights
 - 3.1 Code 1
 - 3.2 Code 2
- 4 Tasks achievements
 - 4.1 Code audit
 - 4.2 Performance portability
 - 4.3 Performance for exascale
 - 4.4 Performance monitoring
- 5. Cross-cutting activities
 - 5.1 PETSc
 - 5.2 TALP
- 6. Conclusions/Lessons learned
 - 6.1 from mini-app to full-app: outcomes from WP2 and WP3
 - 6.2 ...
- Appendix





ChEESE

Status of Cross-cutting activities



Cross-cutting Tasks

- Main objectives:
 - Investigation on the TALP side
 - Scalability testing (production versions) and resolution of any issues (porting). It will require EuroHPC budget.
- A reasonable goal: to extract lessons learned for D2.4 and metrics for new code releases
- BSC should coordinate this part together with HLRS and CINECA (as WP2 lead)

Task	Codes	WP2 task	Priority	Partners
Performance monitoring at scale	SeisSol, xSHELLS, HySEA, FALL3D, ELMER/ICE	T2.4	High	BSC, CINECA
Deepening PETSc GPU performance	Tandem, LaMEM, pTatin3D	T2.3	High	CINECA



TALP: Phase 1 updated status

Flagship code	Target use case(s)	Scaling parameters	Planned EuroHPC resources/budget	Notes
FALL3D	1- High-res scalability use case (Raikoke 2019 eruption)	Iterations per second	EHPC-DEV-2025D10-040 MN5, Leonardo & LUMI	In August the project will be renewed.
	2- Large ensemble use case (Etna 2018 eruption)	Ensemble throughput, ensemble weak scaling and time to solution		
HySEA	Strong scaling (use case: 2011 Tohoku tsunami)	Runtime	Leonardo	
ExaHyPE	Strong scaling (Use cases: Tafjord landslide and ExaSeis TPV)	Throughput in degrees of freedom per second	LUMI	
xSHELLS	Strong scaling of geodynamo simulation	Time per iteration		Some exploration done with TALP but not access to EuroHPC resources to make a proper measurement.



xSHELLS (courtesy by Nat/Sebastien Valat)


- Not yet access to a large stable machine to make a real scaling test (64 core machine (AMD EPYC Milan CPU))
- used TALP for its efficiency measurement (mostly the MPI Efficiency metrics).
- Idea: to check what is the minimal mesh size parameters over the 3 dimensions managed by xSHELLS so that MPI is efficient enough, either globally or looking for some parts of the code (using TALP APIs)
- test done varying the two main dimensions which impact the MPI communications.



● xSHELLS (courtesy by Nat/Sebastien Valat)

As expected the tool reported:

- a mesh being too small is not efficient.
- efficiency is lost for some specific NR values which are not multiple of the number of MPI ranks (64) because of the non symmetric splitting.
(clear vertical lines at NR=96, 160 and 224).
- Of course with NR=64 (meaning a single shell per rank) is not efficient at all.
- From preliminary TALP reporting you would need at least ~ 4 shells per rank to be efficient (NR $\sim$ 256 on 64 cores)
- Still something to tune at tool level...



Cross-cutting Tasks

Main objectives:

investigation on the PETSc side
potential engagement of
expertise on the library
(particularly regarding GPU
support)

Benchmarking campaign (use of
demonstrators)

Minimum goal to extract lessons
learned for D2.4

CINECA proposes to coordinate this
task

Task	Codes	WP2 task	Priority	Partners
Performance monitoring at scale	SeisSol, xSHELLS, HySEA, FALL3D, ELMER/ICE	T2.4	High	BSC, CINECA
Deepening PETSc GPU performance	Tandem, LaMEM, pTatin3D	T2.3	High	CINECA




PETSc





ChEESE

Status of vertical activities



Vertical Tasks

Vertical tasks address **code-specific needs** that cannot be generalized across multiple applications.

Task	Codes	WP2 task	Priority	Partners
Performance portability campaign using SPECfem-Waves	SPECfem3D	T2.2	Medium	CINECA
Avoiding GPU-CPU memory transfers in full-app by changing memory layout to SoA	ExaHyPE	T2.3	High	
Overlapping of communication and computation	ExaHyPE	T2.3	High	
Performance portability studies with different GPU backends on new hardware	ExaHyPE	T2.2	Medium	CINECA
Tuning I/O performances	ELMER/ICE	T2.3	Medium	
Clang/flang version tailored to AMD hardware	ELMER/ICE	T2.3	Medium	
Performance portability campaign using new hw	ELMER/ICE	T2.2	Medium	CINECA



SPECFEM3D



ExaHyPE



ELMER/ICE




Next steps

- Monitoring overall advancement (**CINECA**):
 - Start drafting D2.4 (**CINECA/WP2 team**)
 - Cross-cutting PETSc & TALP (**CINECA/BSC**)
 - Vertical Tasks (**CINECA/code owners**)
 - WP2@General Assembly (**CINECA/WP2 team**)



Coming next: ChEESSE-Factory



AI generated picture...





ChEESE

Questions/Comments?

