

The Advanced Computing Hub at BSC: Improving fusion codes following the principles of EUROfusion standard software

X. Sáez¹, M. Garcia-Gasulla¹, J. Vinyals¹, C. Morales¹,
D. Vicente¹, A. Soba^{1,2}, I. Gasilova¹, M. J. Mantsinen^{1,3}

¹ Barcelona Supercomputing Center, Spain

² National Scientific and Technical Research Council (CONICET), Argentina

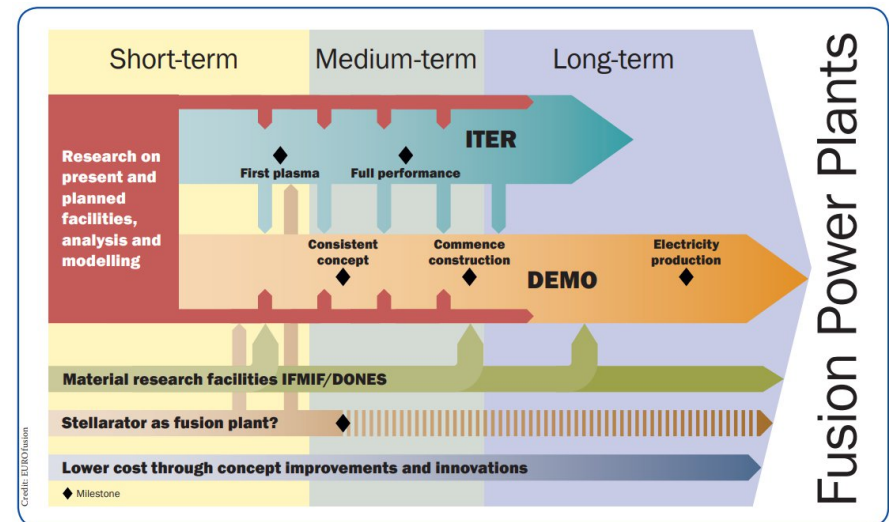
³ ICREA, Barcelona, Spain

3rd Fusion HPC Workshop

15/12/2022

European Research Roadmap

- Provides a **clear and structured way toward commercial fusion energy**
- Is the **basis of the EUROfusion and Fusion for Energy programmes** for the realisation of Fusion Energy
- **Includes DEMO** design studies that are being conducted in Europe
- The DEMO project follows ITER with the goal of designing a commercial fusion power plant



Motivation

- **The step from ITER to DEMO is challenging.** DEMO design is complex in terms of the amount of systems needed to produce and control the plasma
- **Experimental data from ITER and IFMIF-DONES are essential, but not sufficient to design DEMO** with confidence in an unexplored environment to predict plasma and materials performance
- There is a need to create a high-quality suite of research codes (**EUROfusion-standard software**) to model data from existing EUROfusion facilities and to reliably extrapolate to future devices

EUROfusion ACHs

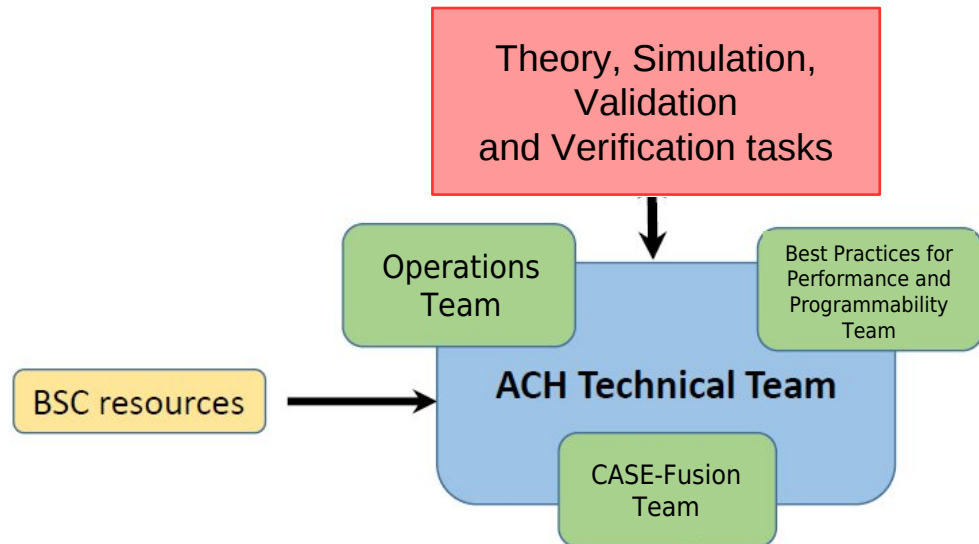
- The goal is to bring together fusion physicists, materials scientists and engineers with computer scientists within the same organisational framework to make use of high-performance computers (HPC) and accelerate the development of fusion energy
- EUROfusion has initiated coordination among theory and advanced simulation creating:
 - **Theory, Simulation, Validation and Verification tasks (TSVV)** that perform fundamental research and channel science
 - **Advanced Computing Hubs (ACHs)** that provide advanced simulation expertise and knowledge to TSVVs



BSC-ACH

- Barcelona Supercomputing Center (BSC) has been awarded one of these ACH that will offer advanced simulation to TSVVs
- **BSC-ACH provides expert support to users regarding High-Performance Computing (HPC):**

- *scalable algorithms*
- *code parallelization*
- *performance optimization*
- *code refactoring*
- *GPU-enabling*
- ...



Codes assigned to BSC-ACH


Code	Work required	Code Coordinator	Groups in charge	Start
ERO 2.0	GPU porting	Dmitry Matveev, Juri Romazanov	CASE-Fusion, Best performance...	2022
SPICE2	Implementing a parallel Poisson Solver and a Parallel electric field calculations.	Michael Komm	CASE-Fusion, Operations	2022
KNOSOS	Optimization.	José Luis Velasco	Operations	2022
STELLA	Optimization.	Michael Barnes	Best performance...	2022
SOLPS	<ul style="list-style-type: none"> • Check the correctness of the OpenMP multi-threading • Improve compiler options and/or job submission script • Consider vectorization of the code profitable for use on Cray platforms 	David Coster	Operations, Best performance...	2023
STELLA	Optimization.	Michael Barnes	Best performance...	2022
XTOR-K	GPU porting	Hinrich Lutjens	CASE-Fusion, Operations	2023
BIT1	GPU porting	David Tskhakaya	CASE-Fusion, Operations	2023
GENE-X	<ul style="list-style-type: none"> • Implementing the ability to access the unstructured computational grid in arbitrary order • Testing the performance of different reordering strategies 	Philipp Uibl	CASE-Fusion, Operations	2023
JOREK	Reduce memory consumption and improve performance	Matthias Hoelzl	CASE-Fusion, Operations	2023

ERO2

- ERO2.0 is a code for modelling plasma-wall interaction and global material migration in fusion devices. The migration is simulated by following 3D trajectories of Monte-Carlo test particles
- **Language:** C++
- **Parallelization:** MPI + OpenMP

ACH tasks

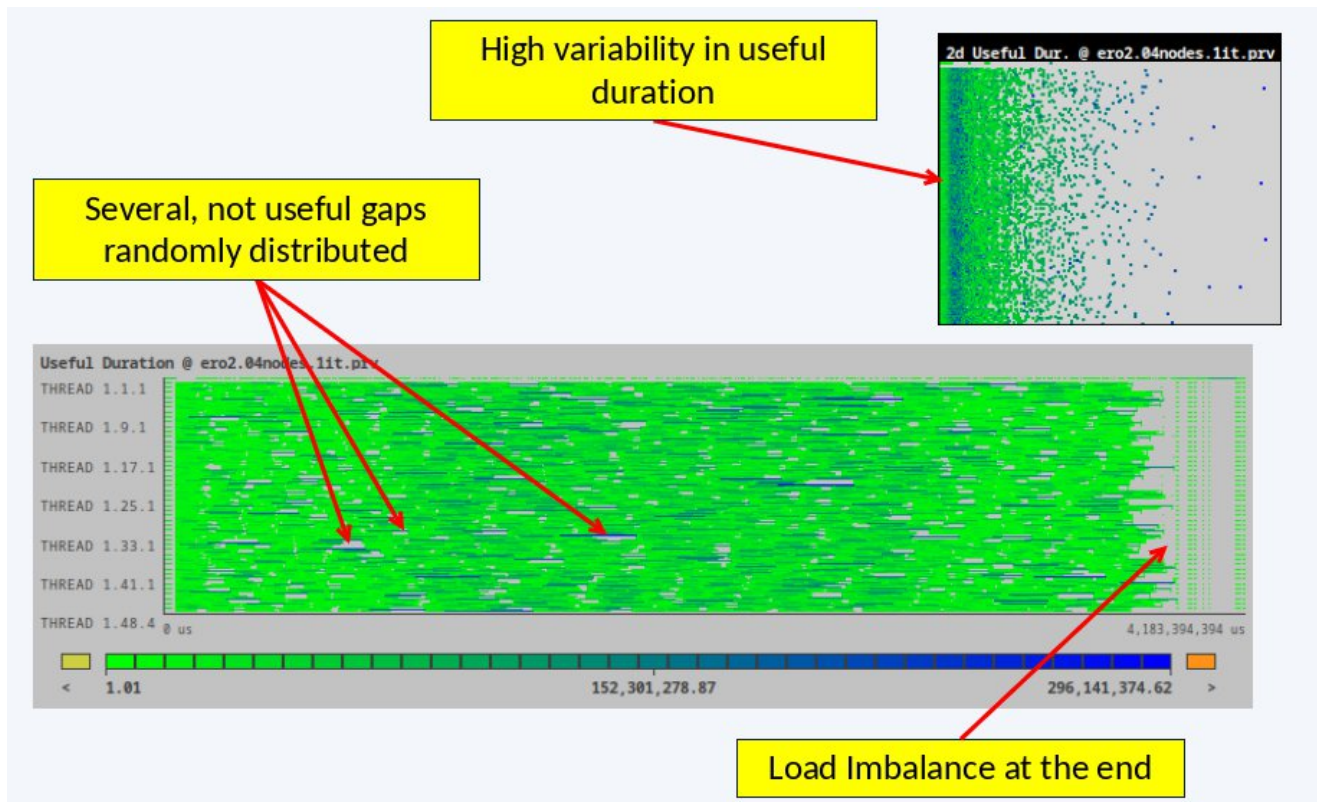
- Performance analysis
- Porting to GPUs



ERO2


Tracing

- BSC Tools: Extrae and Paraver (<http://tools.bsc.es>)




ERO2


4 nodes




8 nodes



16 nodes



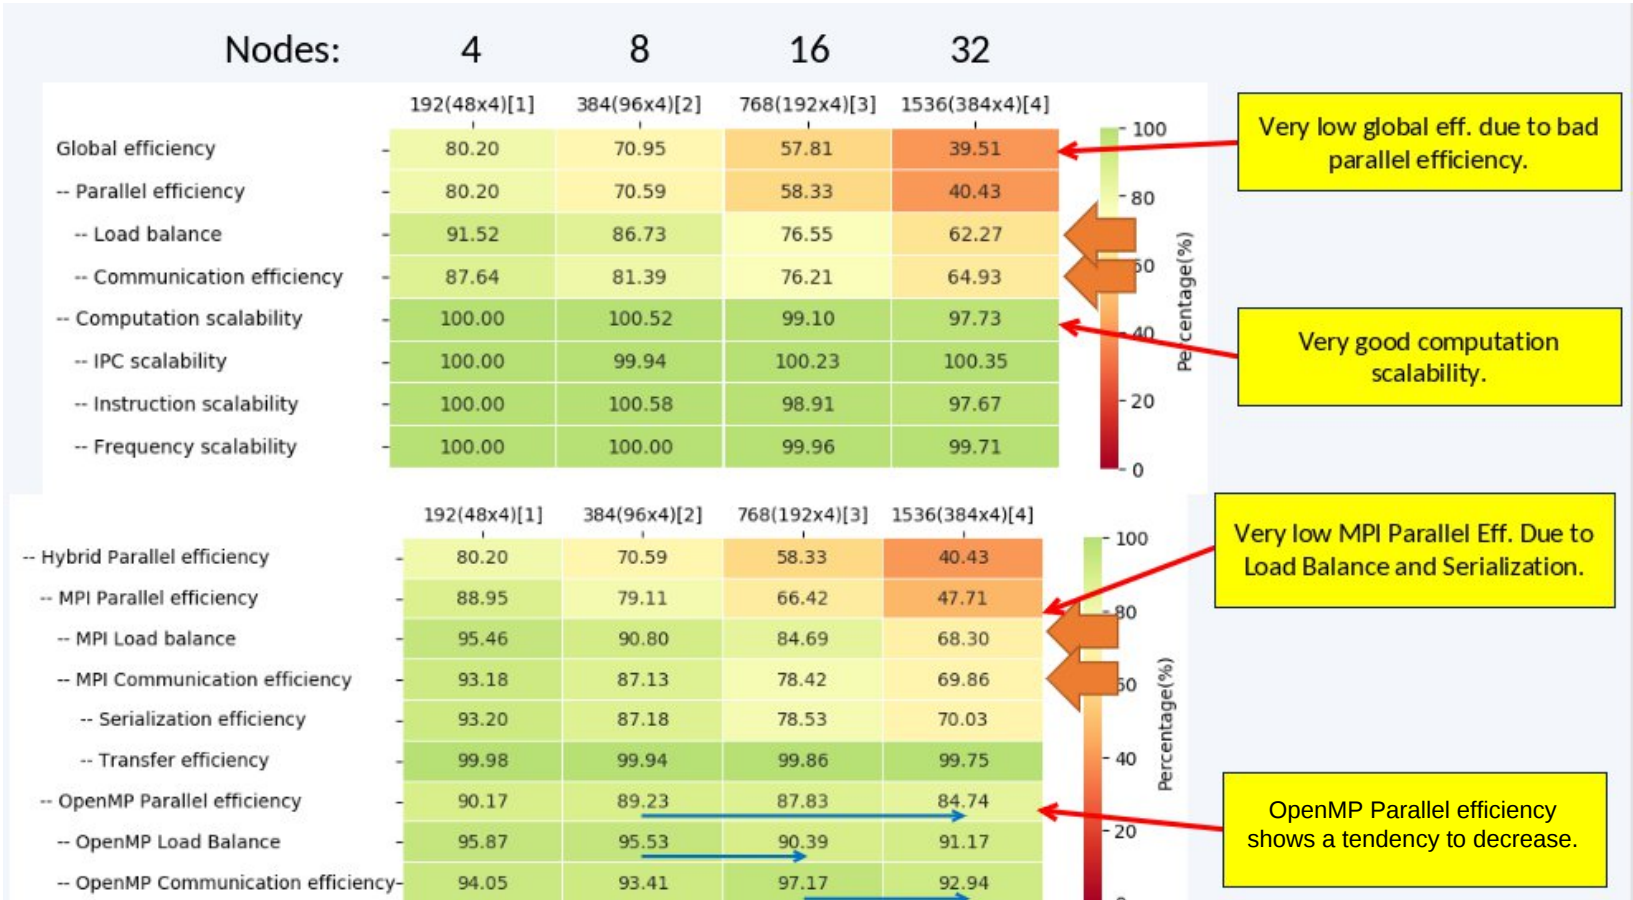
32 nodes



- Gathering phase does not scale.
- Computation seems to scale.
- Load imbalance increases.

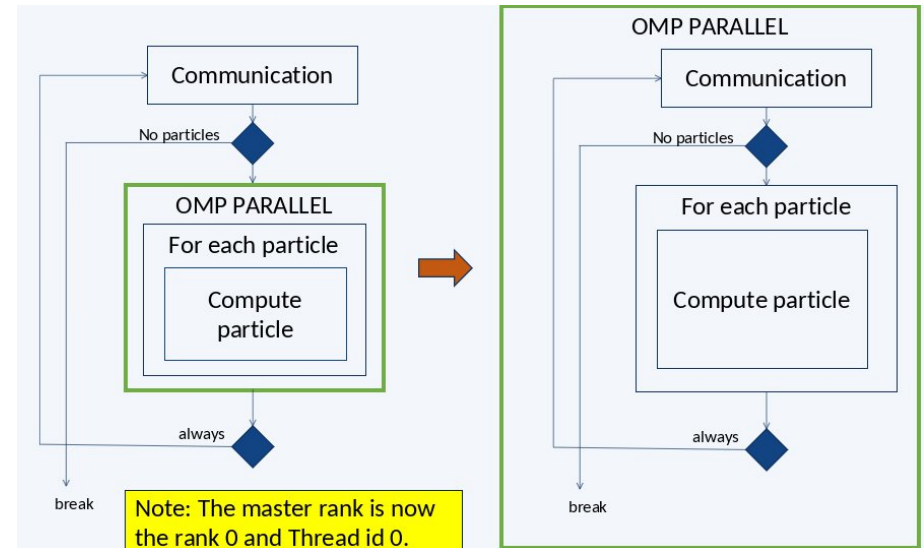
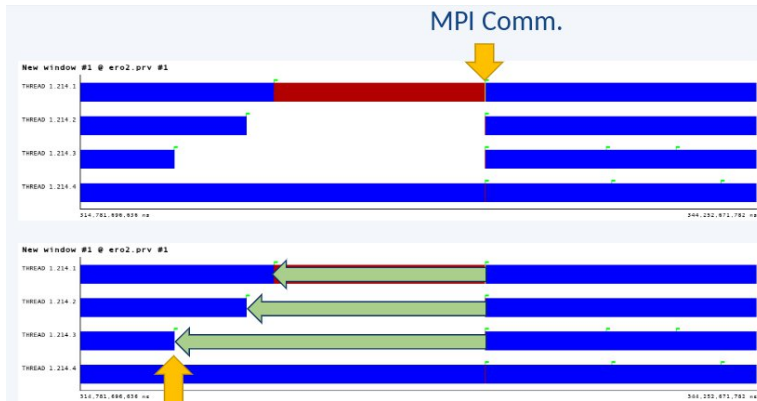
ERO2

Global efficiency metrics



ERO2

Optimization



- Include MPI communication inside the OpenMP parallelization, so new work can be started when threads start idle

ERO2

Porting to GPU

- OpenACC
 - The compiler is responsible for generating kernel and managing data transfers

- Find top time-consuming routine ⇒ octree distance search

Intel VTune

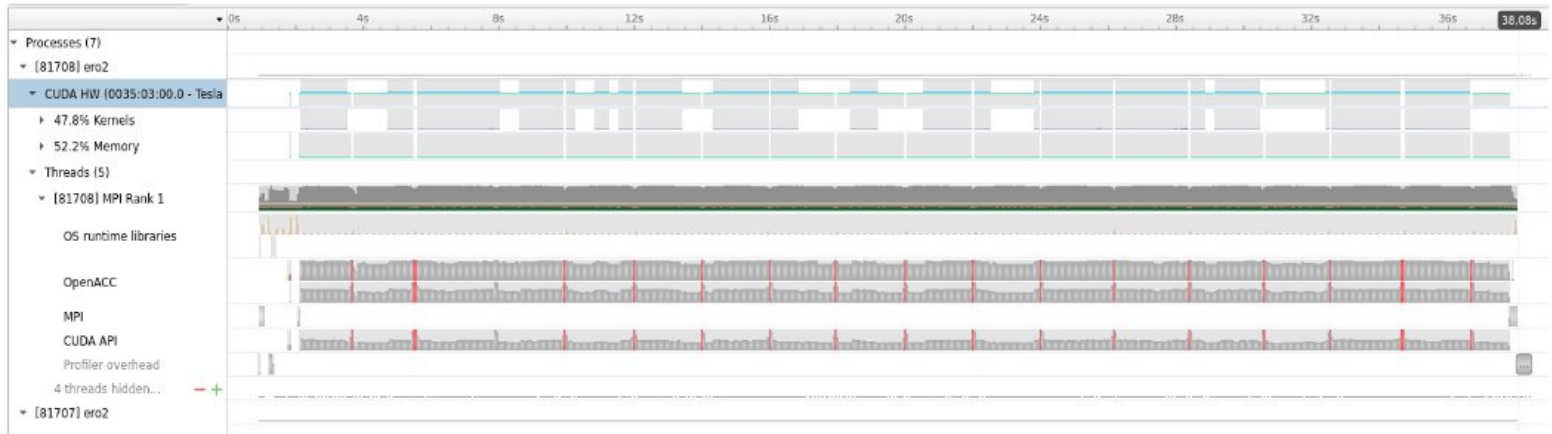
(<https://www.intel.com>)

Symbol Name	Self, %	Total, %
[Broken backtraces]		100.00
0x7fff7e9652f4		99.92
0x7fff7e965100		99.91
main		99.91
ero2::Ero2Simulation::run(int, char**)		99.91
ero2::Ero2Simulation::transportLoop()		99.45
_kmpc_fork_call		99.45
runNewTeam		99.45
launchTeam		99.45
launchInternal		99.45
hxiEmulateHostThreadLaunch		99.45
targetFuncHostTrampoline_1		99.45
ero2::Ero2Simulation::transportParticleLoop(ero2::Particle&, ero2::Plasma&, ero2::DensityManager&, ero2::Sheath&, ero2::Particle*&)		99.44
ero2::Ero2Simulation::transportParticleStep(ero2::Particle&, ero2::Plasma&, ero2::DensityManager&, ero2::Sheath&, ero2::ParticleStepData&, ero2::Particle*&)	0.00	99.33
g3d::Octree::getDistance(g3d::Vector const&, g3d::Vector, double&, g3d::Polygon const*&) const	0.00	69.91
g3d::Octree::getDistanceInNodeSq(g3d::Vector const&, double&, g3d::Polygon const*&, std::set<std::pair<unsigned long, unsigned long>, std::less<std::pair<unsigned long, unsigned long>>> const&) const	0.01	58.75

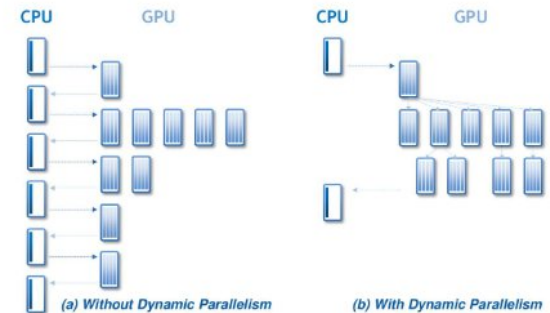

- Issues:
 - C++-style loops are not supported
 - Polymorphism is not supported
 - Routine calls within parallel loops can be problematic
 - ...
- However, it is not as ideal as it was initially expected
 - OpenACC has lacks with C++
 - The feedback from error messages is not very clear

ERO2

- NVIDIA Nsight (<https://developer.nvidia.com/nsight-visual-studio-edition>)



- Next steps
 - OpenACC + CUDA interoperability
 - CUDA Dynamic Parallelism:
Recursive search into Octree



SPICE2

- Particle-In-Cell (PIC) code dedicated to performing simulations of particles in a fixed magnetic and self-consistent electric field for the study of plasma deposition near castellated plasma facing components
- **Language:** Fortran
- **Parallelization:** MPI

ACH tasks

- Implementation of 2D parallel Poisson solver that the number of cores in simulations can be increased to at least 128
- Implementation of a parallel routine for E-field calculation
- General improvement of the code

SPICE2

Profiling

- Intel Tools : VTune and Advisor
(<https://www.intel.com/content/www/us/en/developer/tools/oneapi/vtune-profiler.html>)



Top Time-Consuming Loops

Loop

- loop in [umfdl_lhsolve](#)
- loop in [umfdl_uhsolve](#)
- loop in [umfdl_lhsolve](#)
- loop in [leapfrog_](#)
- loop in [mainloop_](#)


Function	Module	CPU Time
weight_particles_	spice-2.14-debug.bin	204.805s
umfdl_lhsolve	spice-2.14-debug.bin	68.761s
leapfrog_	spice-2.14-debug.bin	58.497s
calc_e_field_general_	spice-2.14-debug.bin	47.722s
psolver_direct2_periodic_	spice-2.14-debug.bin	45.066s
{Others}	N/A*	131.634s

- By modifying the most consuming parts of the code, we achieve small improvements


SPICE2

Tracing

- BSC Tools: Extrae and Paraver (<http://tools.bsc.es>)






- All processes write to the same file at the beginning of the simulation
- Each process waits $6 * \text{their_rank}$ seconds to start writing



- We updated the code to only do one writing on the process with rank 0 and the others processes wait on a Barrier

SPICE2

Optimizing

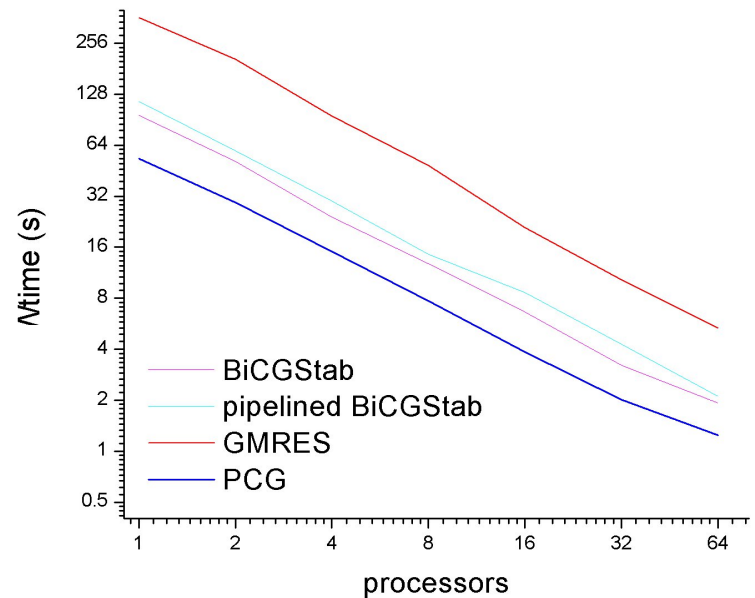


- The potentials are sent to all the processes from rank 0 with Point-to-Point (P2P) calls
- The density vector is calculated on process 0 and sent to all the processes with P2P
- Code updated to send the potentials with the MPI_Broadcast collective
- Code updated using MPI_Reduce collective to compute the density vector

SPICE2

Implementation (Poisson solver)

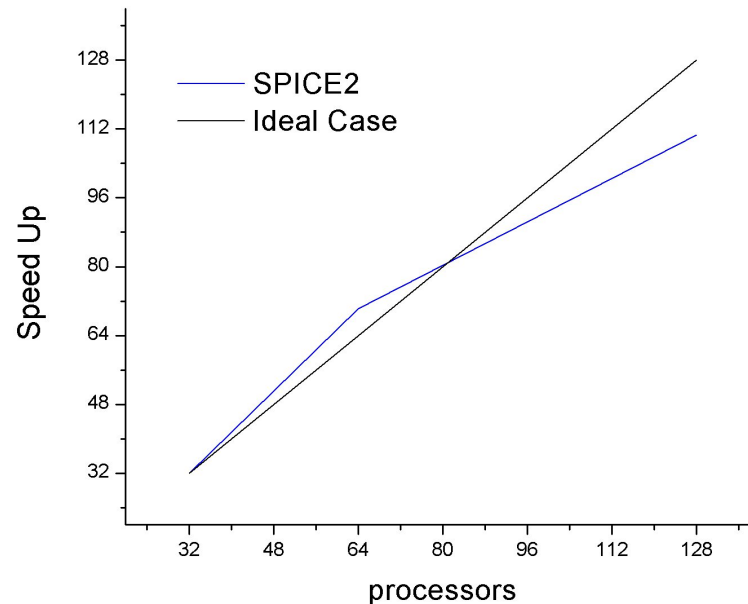
- **Current implementation:** Direct solver based on LU decomposition using the UMFPACK and BLAS libraries
- **Limits:** good scaling up to 32 cores (memory limit of the system) and a grid limit of ~ 4000 cells in one dimension
- **Solution:** use a KSP linear solver based on PETSc library (<https://petsc.org>). Several options of solvers were explored
- The best time is obtained for Conjugate Gradient with Jacobi preconditioner solver



SPICE2

- **Results:**

- Speed up of the whole code with the new solver included
- Good performance up to 128 processors
- Loss of scalability for more than 256 processor, but the code can now run over 512 processors, solving domains of 10000x4000 cells (Marconi)



KNOSOS

- Code that calculates neoclassical transport in low-collisionality plasmas of 3D magnetic confinement devices by solving the radially local drift-kinetic and quasineutrality equation
- **Language:** Fortran
- **Parallelization:** MPI

ACH tasks

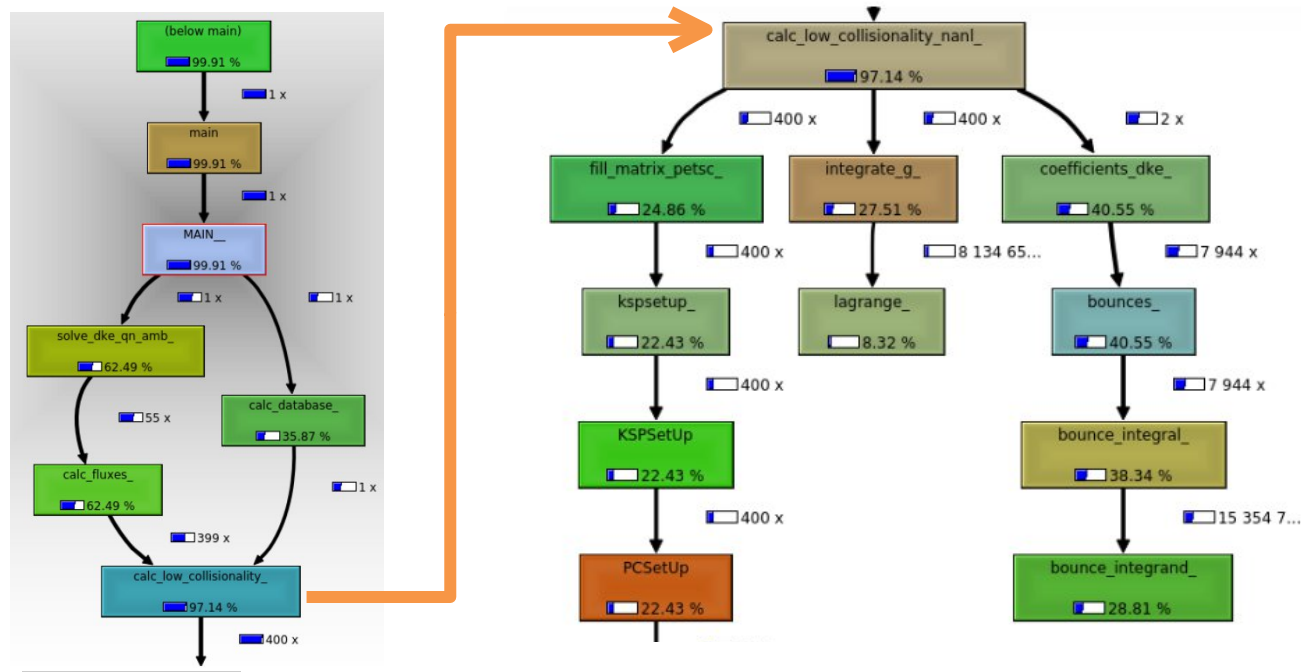
- Overall performance assessment
- Optimize application

KNOSOS

Profiling

- There is a high load imbalance between MPI ranks, which seems to be produced by different number of iterations in subroutines ⇒ feedback from developers
- Discover that parallelization relies on launching “independent” calculations ⇒ **improving sequential version would translate to an overall improvement in the MPI version**


- Firstly, identify the most time consuming parts of the code with Valgrind (<https://valgrind.org>)



KNOSOS

Optimization

- Introduce a new parallelization level using OpenMP in `coefficients_dke` routine
- It was found that different instances of the BOUNCES subroutine could be performed in parallel without producing data conflicts between iterations
- **Result:** Execution time was reduced from 15 sec (sequential) to 10 sec (using 4 OpenMP threads)
- We repeat the process with `integrate_g` routine. However, the improvement was negligible




```
!$OMP PARALLEL DO FIRSTPRIVATE(Q)
DO ipoint=1,npoint
  iw =i_w(ipoint)
  ila=i_l(ipoint)
  CALL BOUNCES(iw,z1(iw),t1(iw),B1(iw),hBpp1(iw),vd1(:,iw), &
    &          zb(iw),tb(iw),Bb(iw),hBppb(iw),vdb(:,iw), &
    &          z2(iw),t2(iw),B2(iw),hBpp2(iw),vd2(:,iw), &
    & 1./lambda(ila),ipoint.EQ.1,nq,Q,&
    & zlw(ipoint),tlw(ipoint),zrw(ipoint),trw(ipoint))
  BI1(ipoint)=Q(1)
  BI2(ipoint)=Q(2)
  BI3(ipoint)=Q(3)
```

Conclusions

- The Advanced Computing Hubs have an ambitious mission to accomplish in a small time window
- The creation of EUROfusion-standard software needs close collaboration between developers and computer science professionals
- We have shown some examples of our work, with different degrees of success to build trustable software capable to run efficiently in HPC systems
- In some cases, the best solution does not accomplish the objectives or the wishes of the developers

Fusion Group



Thank you



fusion.bsc.es



[@Fusion_BSC](https://twitter.com/Fusion_BSC)



fusion@bsc.es

xavier.saez@bsc.es