



Meeting of HPC ACHs

Alejandro Soba

HPC ACHs | 15-16 November 2023



This work has been carried out within the framework of the EUROfusion Consortium, funded by the European Union via the Euratom Research and Training Programme (Grant Agreement No 101052200 — EUROfusion). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them.

SPICE2-BIT1



- **SPICE** (SheathParticle In Cell) package includes two codes: SPICE2 (2D3V) and SPICE3 (3D3V).
- PIC code for simulations of particles in a fixed magnetic and self-consistent electric field
- Written in Fortran 90, outputs in the Matlab MAT binary format.
- Parallelization implemented using domain decomposition principles and message passing interface (MPI).
- All internal routines are parallel except for the Poisson solver. The Poisson solver is serial (taking 3% of the overall calculation time). It operates with global matrices of potential and charge density.
- Best solver are SPARCE, based in UMFPACK.



Work required into ACH

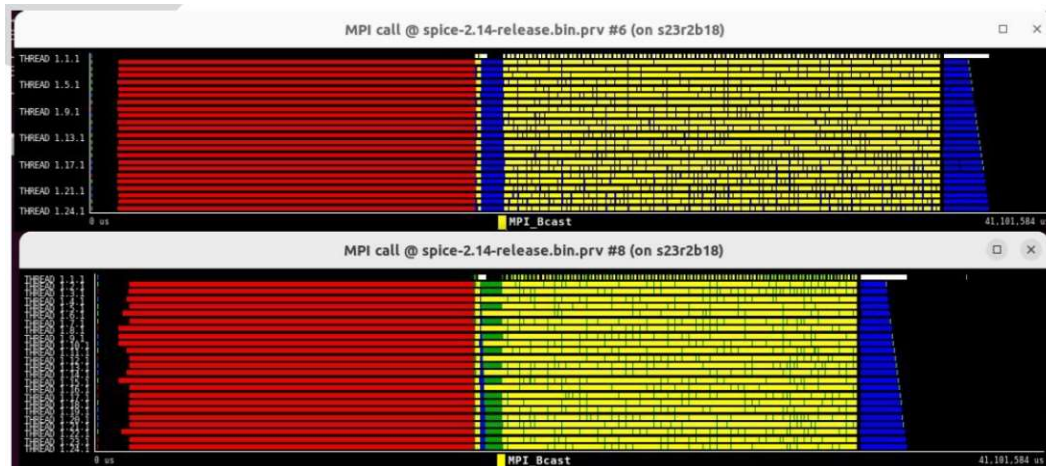
- Implementation of 2D parallel Poisson solver with good scaling and speed so that the number of cores in simulations can be increased to at least 128 (current practical limit is around 32) and the grid size can be increased (UMFPACK has a limit of ~4000 cells in one dimension);
- Implementation of parallel routine for E-field calculation (later task).

SPICE2-BIT1

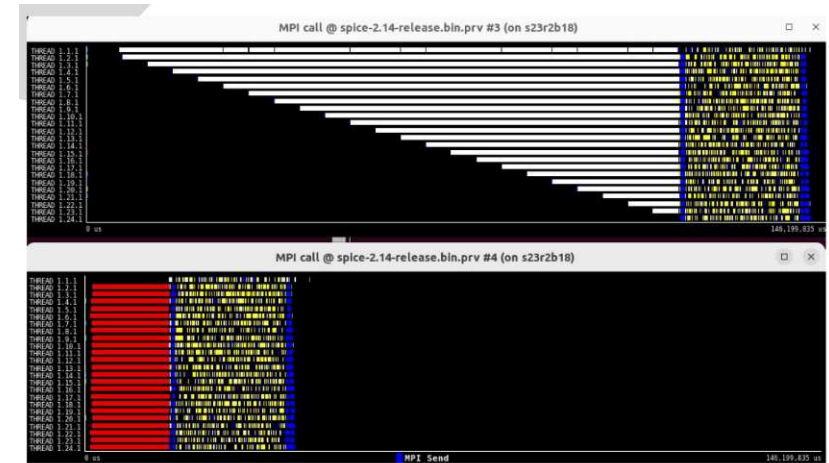


SPICE2

Firstly, the Operation team does an exhaustive performance analysis. Several optimizations were included within the code.



Paraver's traces showing the time saved when collective communicators are used.




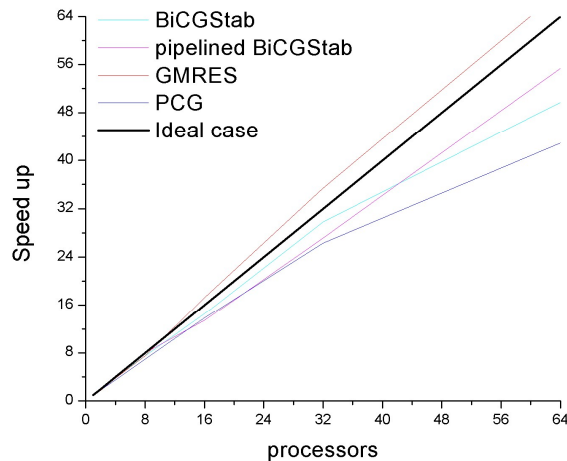
Paraver's traces showing the barrier and the waiting time saved.

SPICE2-BIT1



SPICE2

- We included PETSC solver, based in KSP linear solver context and we explored several options
- From all the solvers tested, we identify 4 of them that possess the better CPU time and scalability. The best is KSP-1




Scalability of selected KSP Petsc solver analyzed (left); Wall time consumed by each KSP solver (right).

SPICE2-BIT1



We included several options in the input files to choose variables for PETSC.

SPICE2 is able to run in big domains up to 128 processors with good scalability

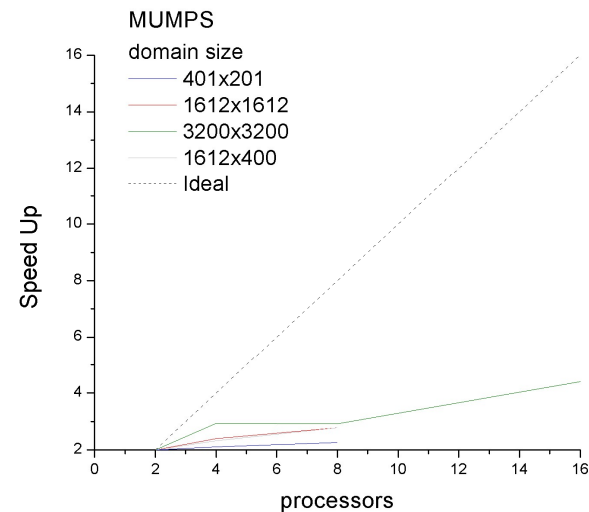
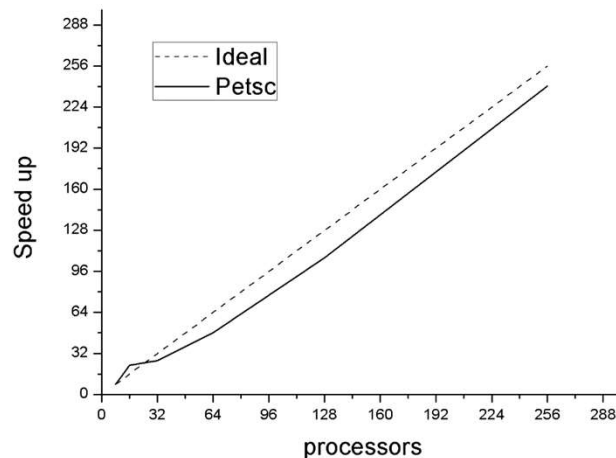


SPICE2-BIT1



SPICE2

- The developers don't like that the solver were slower than the serial one.
- We implement other version in PETSC and in MUMPS




Petsc (Left) and MUMPS (right) scalability in MN4.

SPICE2-BIT1



SPICE2

- We implemented also the Field parallel calculation.
- The results are identical using the serial solver (PS=2) and the Petsc solver (PS=10)





SPICE2

- During this work we identified several issues in the memory management.
- Bad distribution of the domain and several arrays that consume memory until making it starving.
- A compromise exists: It is not possible to run SPICE2 faster than the serial version using UMFPACK, therefore it is not possible to run the code in bigger domains



SPICE2 :: next steps to improve the code.


- **Common repository** → we need to share some space to work together.
- **Portability** → change the use of libraries to system installed libraries.
- **Optimization** → change the use of arrays with time dimension. Use of a disk file or use the information directly in the code.
- **Scalability** → use domain division in all the code.
- **Acceleration** → Introduce a solver using GPUs.

SPICE2-BIT1



BIT1

- BIT1 is an electrostatic 1D3V PIC direct Monte Carlo code for plasma simulations used to edge plasma simulations.
- The first Analysis of performance was done by the Operation group.
- MareNostrum 4



	fdistr	avq_1	Ep_p	arrj	
THREAD 1.1.1	60,658.67 us	43,183.12 us	52,475.89 us	30,969.13 us	
THREAD 1.2.1	77,836.46 us	44,308.57 us	41,937.91 us	20,584.50 us	
THREAD 1.3.1	78,952.92 us	48,371.88 us	38,093.62 us	22,121.51 us	
THREAD 1.4.1	68,081.10 us	48,877.08 us	54,023.53 us	16,610.81 us	
THREAD 1.5.1	68,518.85 us	44,560.45 us	49,924.16 us	24,788.26 us	
THREAD 1.6.1	60,758.92 us	47,693.65 us	57,255.93 us	21,979.01 us	
THREAD 1.7.1	78,732.22 us	44,816.49 us	40,915.55 us	23,169.62 us	
THREAD 1.8.1	80,969.60 us	45,571.57 us	37,646.81 us	23,105.34 us	
THREAD 1.9.1	73,012.10 us	42,974.30 us	47,105.68 us	24,509.65 us	
THREAD 1.10.1	82,986.60 us	44,241.62 us	34,074.57 us	28,295.65 us	
THREAD 1.11.1	84,000.39 us	45,306.71 us	31,600.37 us	25,503.90 us	
THREAD 1.12.1	95,141.59 us	45,654.59 us	26,300.40 us	19,277.69 us	
	909,649.41 us	545,560.03 us	511,354.41 us	280,915.08 us	7
je	75,804.12 us	45,463.34 us	42,612.87 us	23,409.59 us	
im	95,141.59 us	48,877.08 us	57,255.93 us	30,969.13 us	
im	60,658.67 us	42,974.30 us	26,300.40 us	16,610.81 us	
v	9,653.43 us	1,839.72 us	9,238.68 us	3,692.90 us	
ax	0.80	0.93	0.74	0.76	

SPICE2-BIT1



Next steps, porting to GPU. But...

- BIT1 is a one-dimensional PIC code, which means that it treats a big number of particles in big domains, and for each of them a relatively low number of calculi must be performed.
- Optimum for MPI or OPEMMP.
- Not so good to use in a GPU.
- Starting using OPENACC directives in CTE-Power
- First tests were bad.

Then...

- We need to test in Leonardo or MN5
- We ask to developer a version avoiding spurious communication and diagnostics.

CTE-POWER 9

2 login node and 52 compute nodes, each of them:

- **2 x IBM Power9 8335-GTH @ 2.4GHz (3.0GHz on turbo, 20 cores and 4 threads/core, total 160 threads per node)**
- **512GB of main memory distributed in 16 dimms x 32GB @ 2666MHz**
- **2 x SSD 1.9TB as local storage**
- **2 x 3.2TB NVME**
- **4 x GPU NVIDIA V100 (Volta) with 16GB HBM2.**
- **Single Port Mellanox EDR**
- **GPFS via one fiber link 10 GBit**



Conclusions

- We did the required work with SPICE2, but the results are not the expected by the developers.
- We identified several points to improve the performance of the SPICE2 code. Of course, it is difficult to be faster than the serial sparse solver
- A version of BIT1 for work in GPU will need a considerable reformulation of the engineering of the code to reach improvements.