



ABP on HPC

X. Buffat, M. Schenk, G. Iadarola, H. Rafique and G. Rumolo
for the ABP Computing Working Group



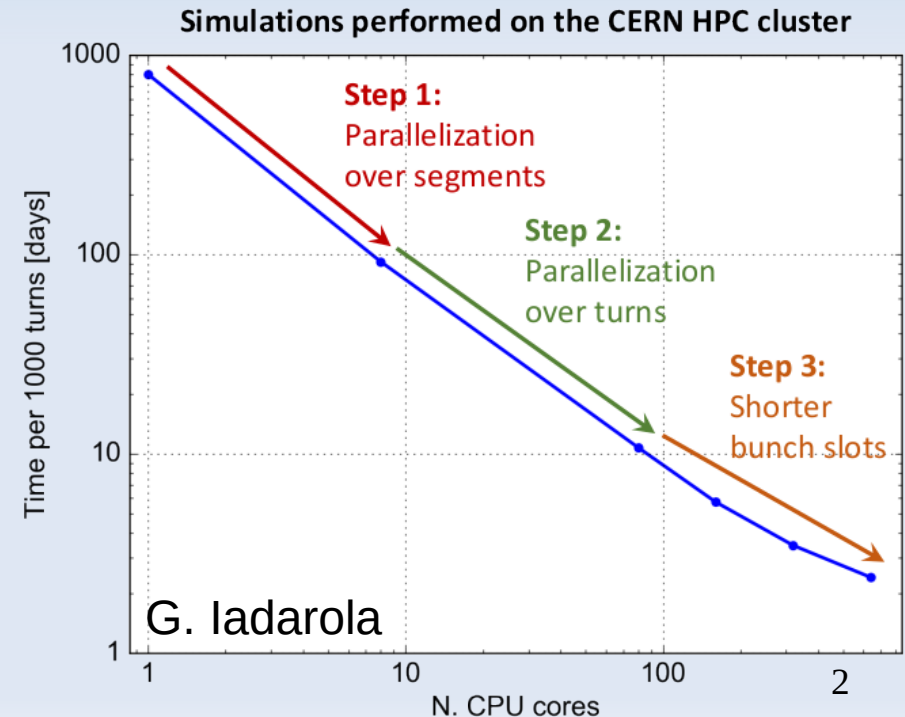
- PyHEADTAIL(-PyECLLOUD): Coherent instabilities
- COMBI : Coherent beam-beam effect
- PyOrbit : Space-charge effects
- Overall feedback



Multibunch PyHEADTAIL - PyECLLOUD



- Watch the result of a simulation of a full SPS batch (288 bunches, modelled with 10^6 macroparticles each) circulating in the LHC, interacting with electron clouds around the machine ($3 \cdot 10^8$ macroparticles) at this [link](#)
- Impressive scaling with various levels of parallelisation
 - Multi-node parallelisation is essential
- Mainly python, using MPI4Py
- Currently ramping up the usage, smooth testing / debugging
 - Could profit from IT expertise on profiling tools



Multibunch PyHEADTAIL

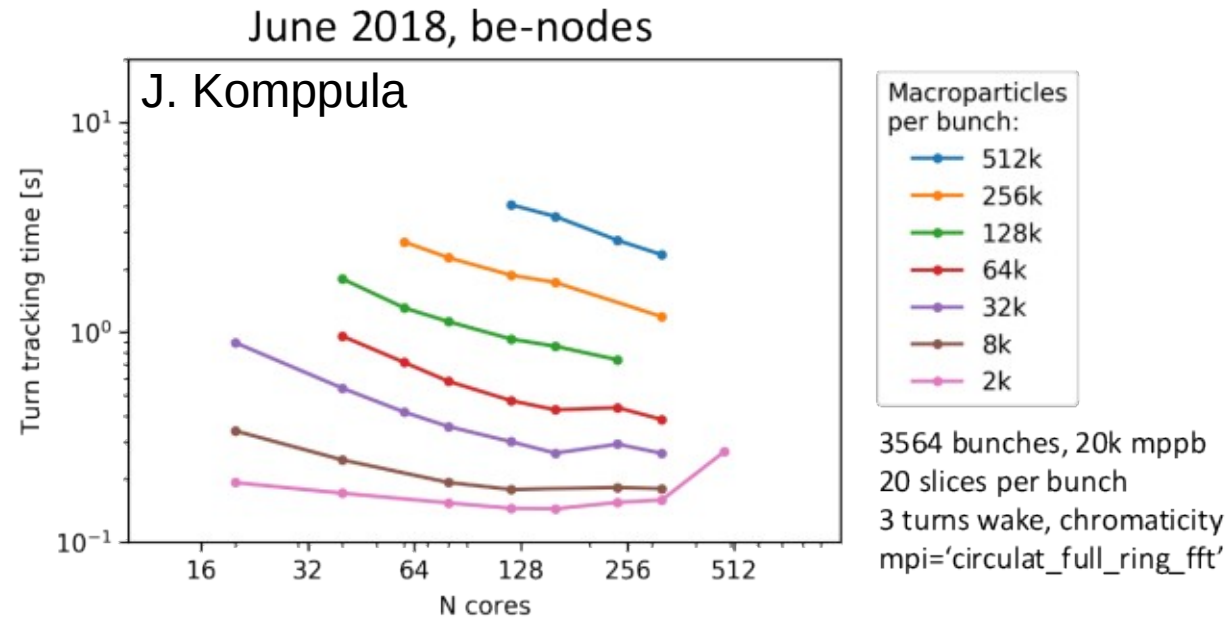
- Used CERN HPC between Sept.'18 and Jan.'19
 - *PyHEADTAIL*: multi-core particle tracking simulations
 - based on Python 2.7, using also *h5py-parallel*, *cython*, *mpi4py*, etc.
 - openMPI 1.8.4 module

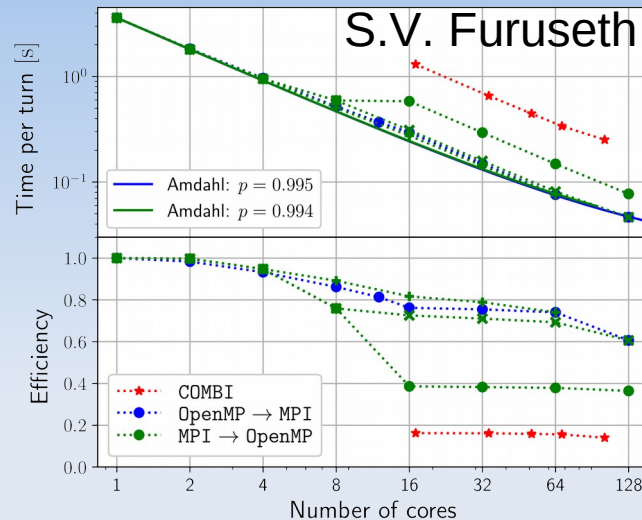

• Experience

- ✓ Efficient job handling: typically jobs start running within minutes, *scancel* with immediate effect, *queue* refreshes fast
- ✓ Usually useful logs for debugging
- ✓ No unexplained aborts once using “correct” submission parameters (# cores, hyper threading) – see issues below
- ✓ Generally a very positive experience

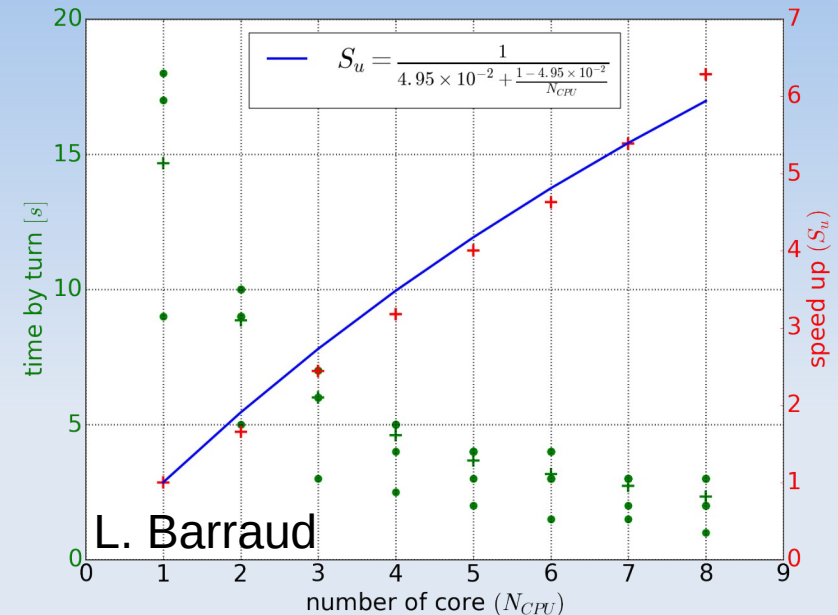
Issues

- # cores: jobs would fail if # cores not integer multiple of # cores per node (depending on queue: BE 20 cores per node, Batch: 16 cores per node)
- Had to disable hyper threading, otherwise inexplicable job aborts at beginning (input from experts in reply to trouble ticket opened by J. Komppula)





(d) 2 beams ($N_{b1} = 8 = N_{b2}$). 1 wakefield calculation and 3 beam-beam calculations, as in Fig. 4c.




- COMBI is used mostly for instability simulations involving beam-beam interactions and/or noise

- 10⁵ to 10⁷ macroparticles per bunch, up to tens of bunches per beam
- Hybrid OpenMP and MPI parallelisation (C and Fortran, using gcc and mvapich2 on HPC) : Some effort when starting, but currently running smoothly
- Multi-node parallelisation is required for multibunch studies, probably ramping up the usage towards end of summer
- Most single/two bunch studies are now ported to HTCondor (still running on HPC when parallel resources fall short on HTCondor)

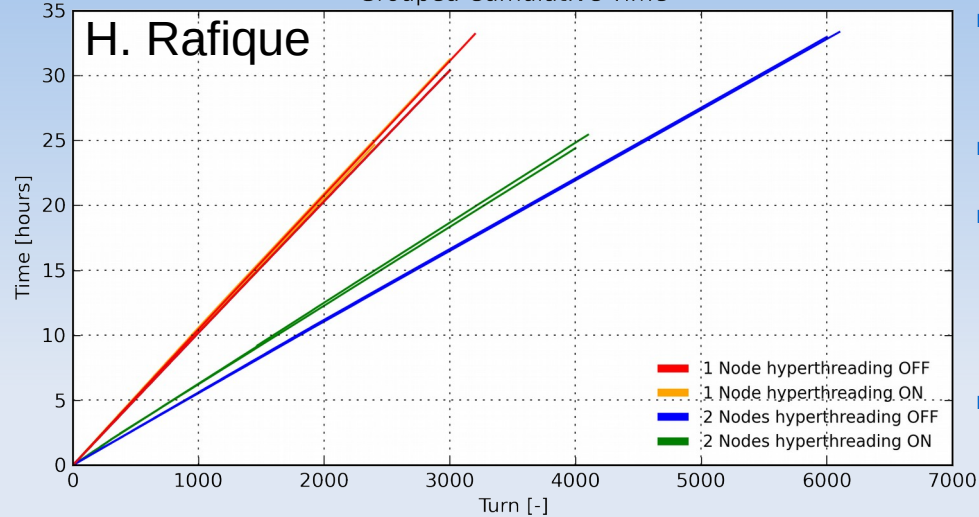


PyORBIT

Space-charge effects



Grouped Cumulative Time




- PyORBIT is used on HPC mainly for working point studies in the PS
- $5 \cdot 10^5$ macroparticles, 2.5D PIC solver
- Good scaling on two nodes, but absence of scaling / detrimental effect with hyperthreading
- Access to a shared system is critical as the building of the code is heavy
 - Currently AFS performs well, but after ?

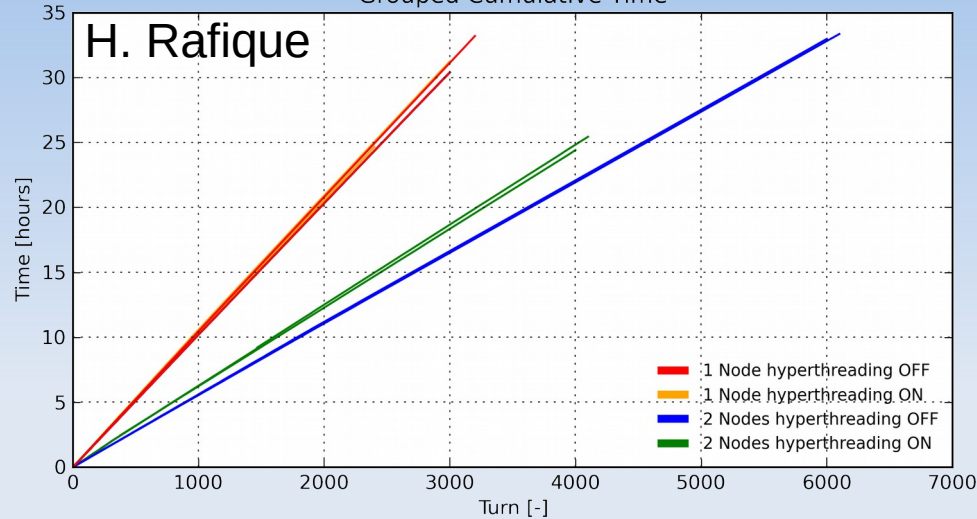


PyORBIT

Space-charge effects

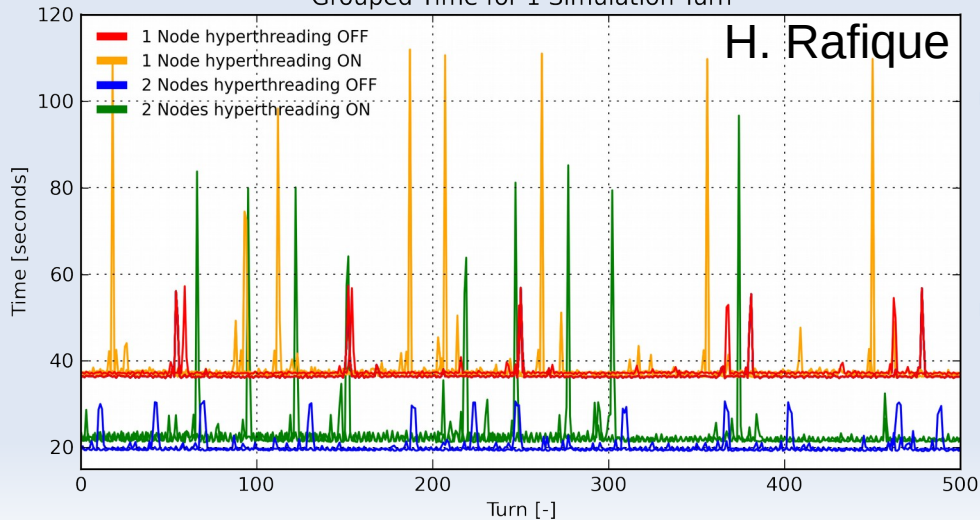


Grouped Cumulative Time



- PyORBIT is used on HPC mainly for working point studies in the PS
- $5 \cdot 10^5$ macroparticles, 2.5D PIC solver
- Good scaling on two nodes, but absence of scaling / detrimental effect with hyperthreading
- Access to a shared system is critical as the building of the code is heavy
 - Currently AFS performs well, but after ?

Grouped Time for 1 Simulation Turn



- The execution does not seem smooth, with identical parts executing significantly slower at times
- This feature is not limiting, but can it be understood ?



Feedback



- Overall the feedback from the users of the HPC in ABP is very positive: The machines are efficient and most of the issues solved rapidly via tickets
- Four potential improvements :
 - Provide expertise with profiling tools
 - Often the output data can be (at least partially) processed on the cluster to avoid large data transfer to other file systems → The frontends are sweating and requesting a an interactive job seems like an overkill (and is pretty inconvenient if the queues are full)
 - Is there an efficient solution, e.g. powerful frontends or dedicated 'post-processing' nodes
 - Hyperthreading does not always behave well (jobs killed, absence/negative scaling)
 - Could jobs running batch-* queues have access to bescratch ?