

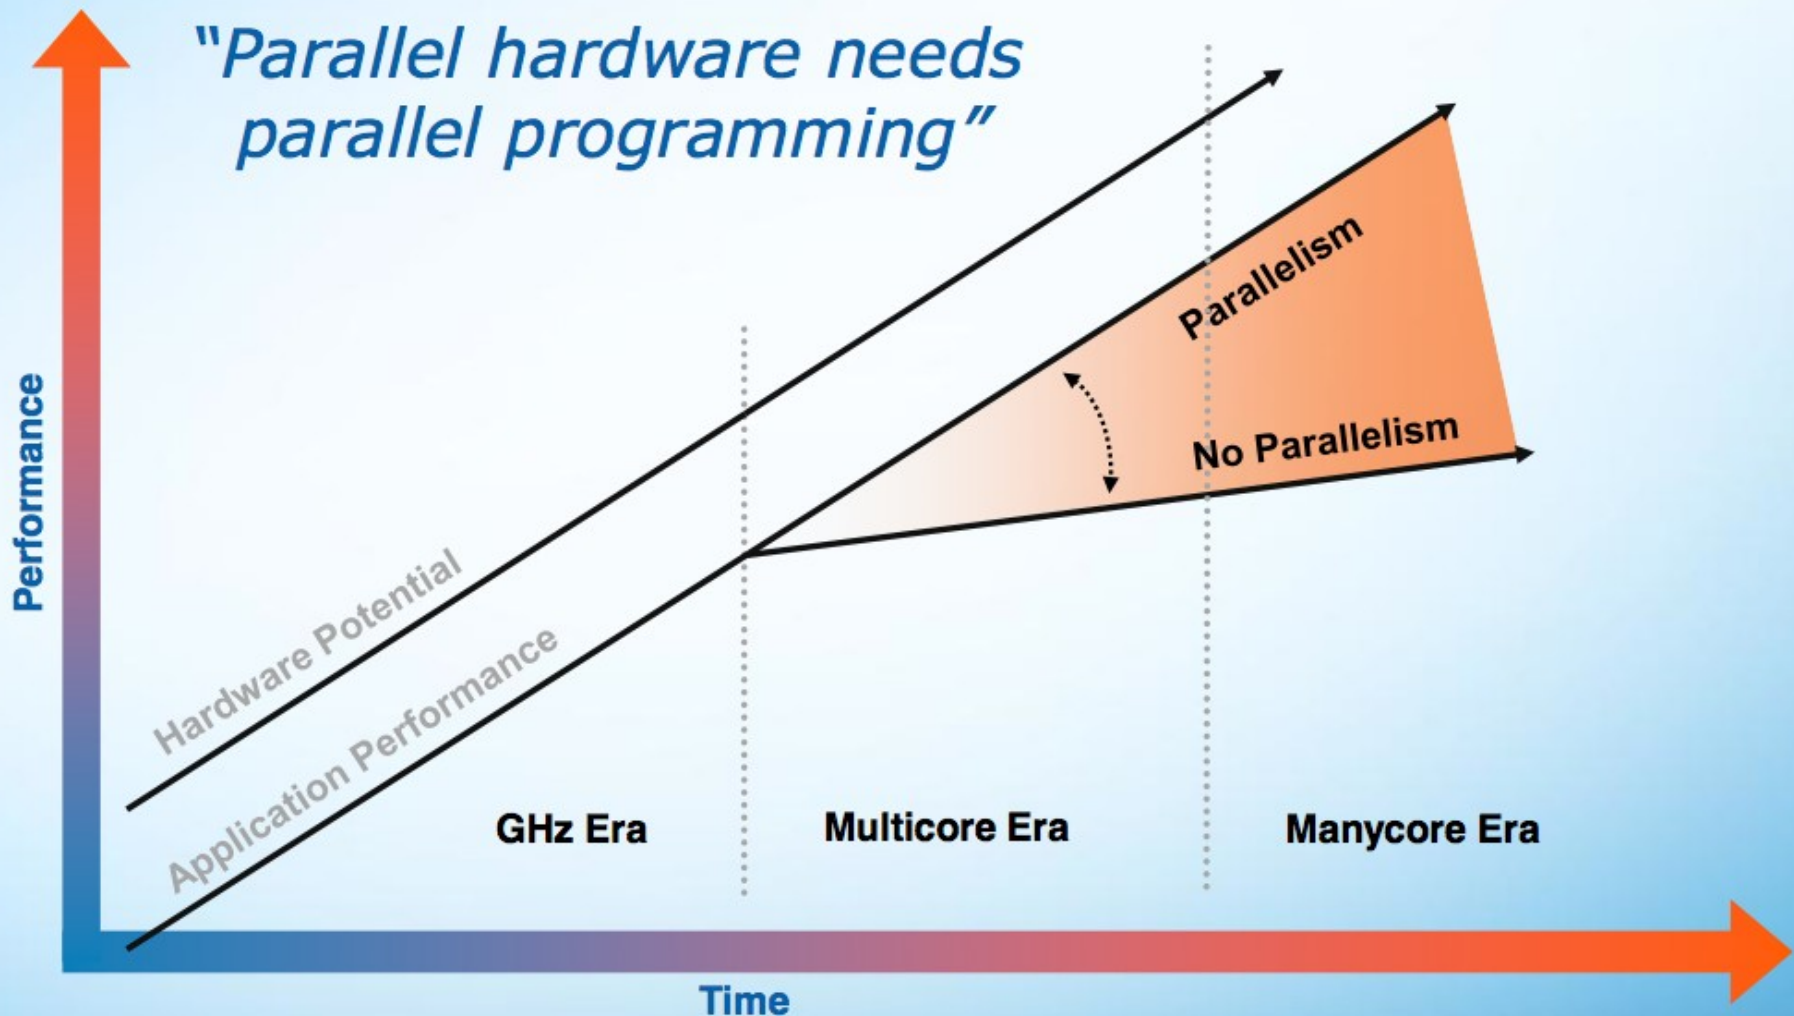
The optimisation of ALICE code

Federico Carminati
January 19, 2012

parallelism

Motivation: Performance

From a recent talk by Intel



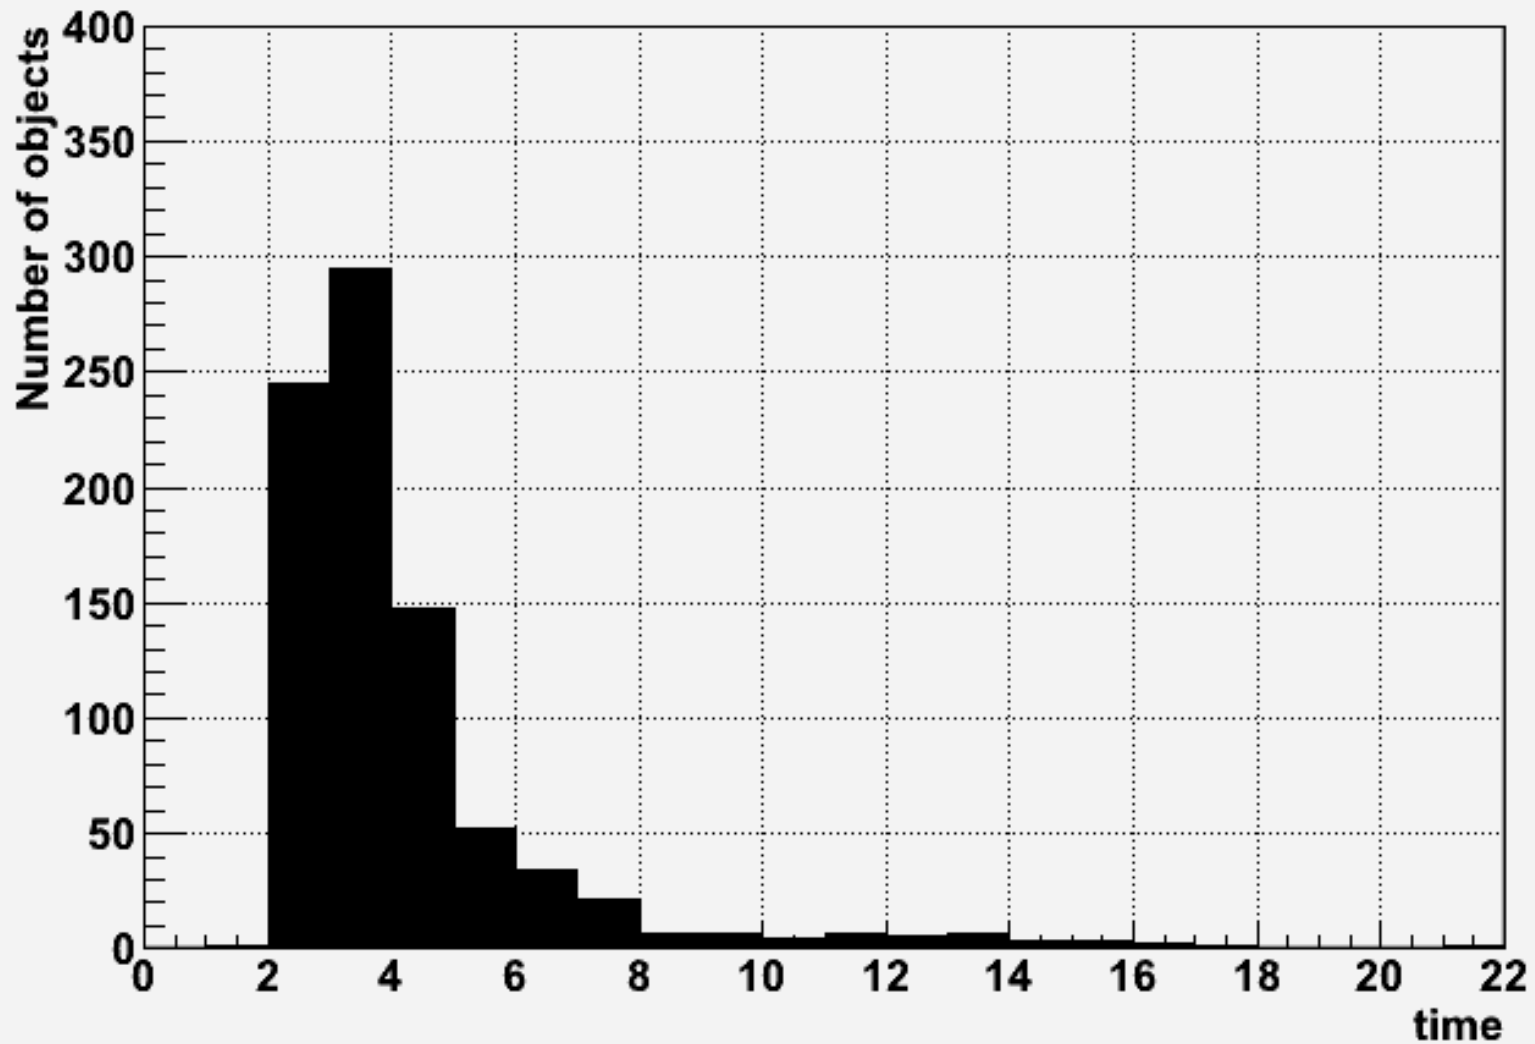
Why it is so difficult?

- No clear kernel
- C++ code generation / optimisation not well understood
- Most of the technology is coming out now
 - Lack of standards
 - Technological risk
- Non professional coders
- Fast evolving code
- No control on hardware acquisition

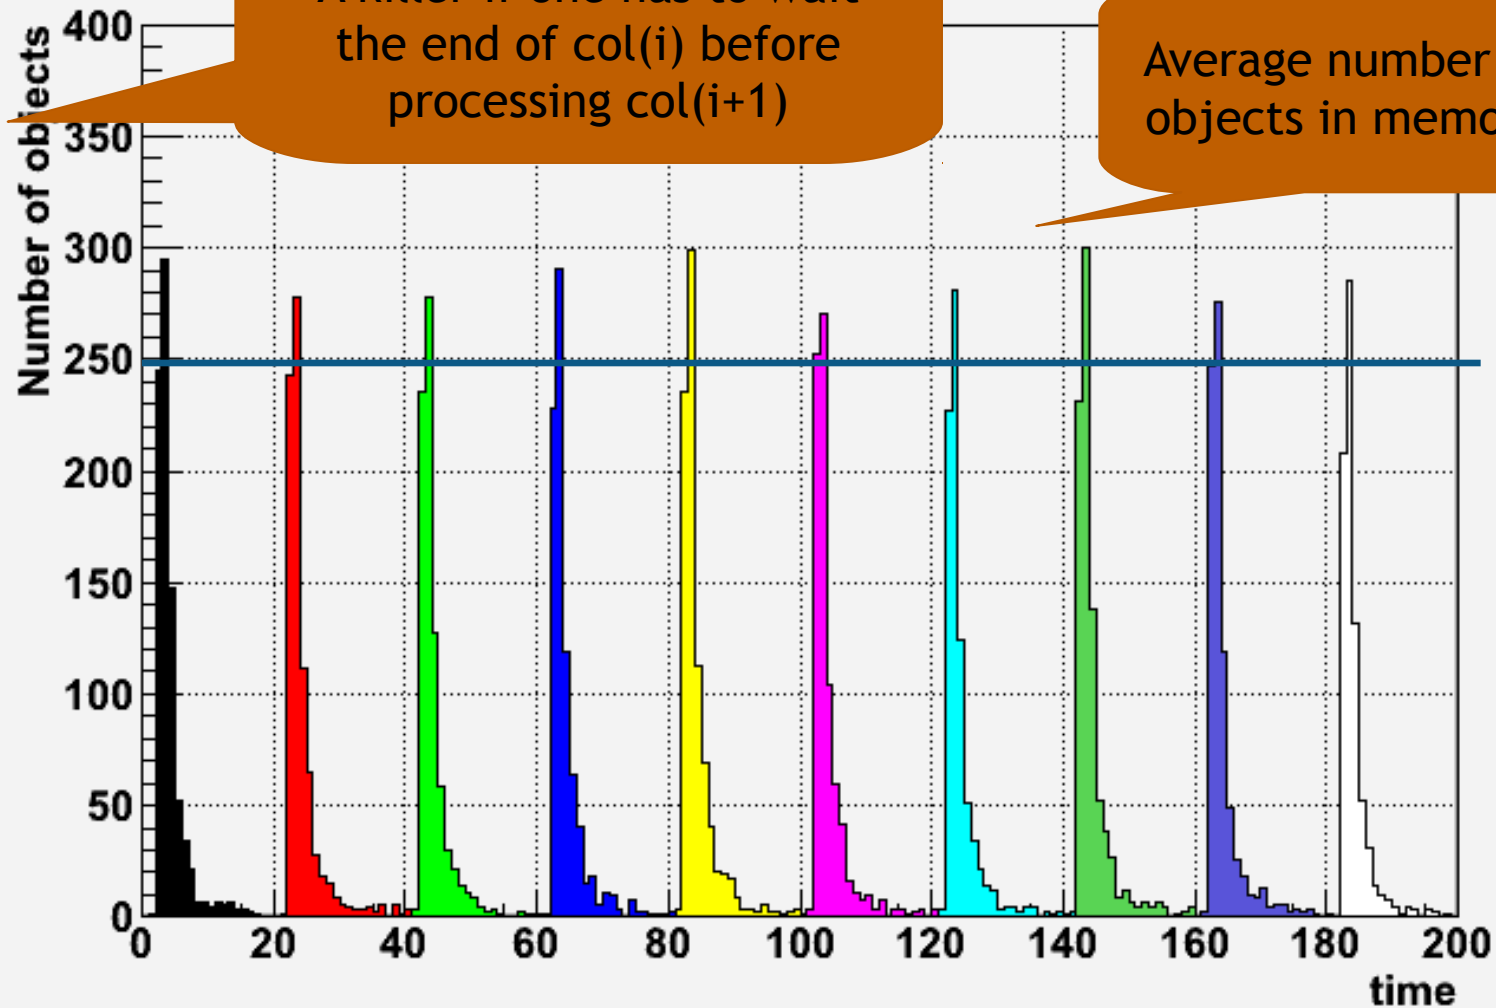
ALICE strategy (unauthorised)

- Use the LSD-1 to essentially re-write AliRoot
- Use the LSD-2 to expand the parallelism to the Grid
 - Hopefully the major thrust will be on MiddleWare
- Refactor the code in order to expose the maximum of parallelism present at each level
- Keep the code in C++ (no CUDA, OpenCL etc.)
- Explore the possible use of #pragma's (OpenMP, OpenACC)
- Experiment on all hardware at hand (OpenLab, but not only)

tails, tails, tails



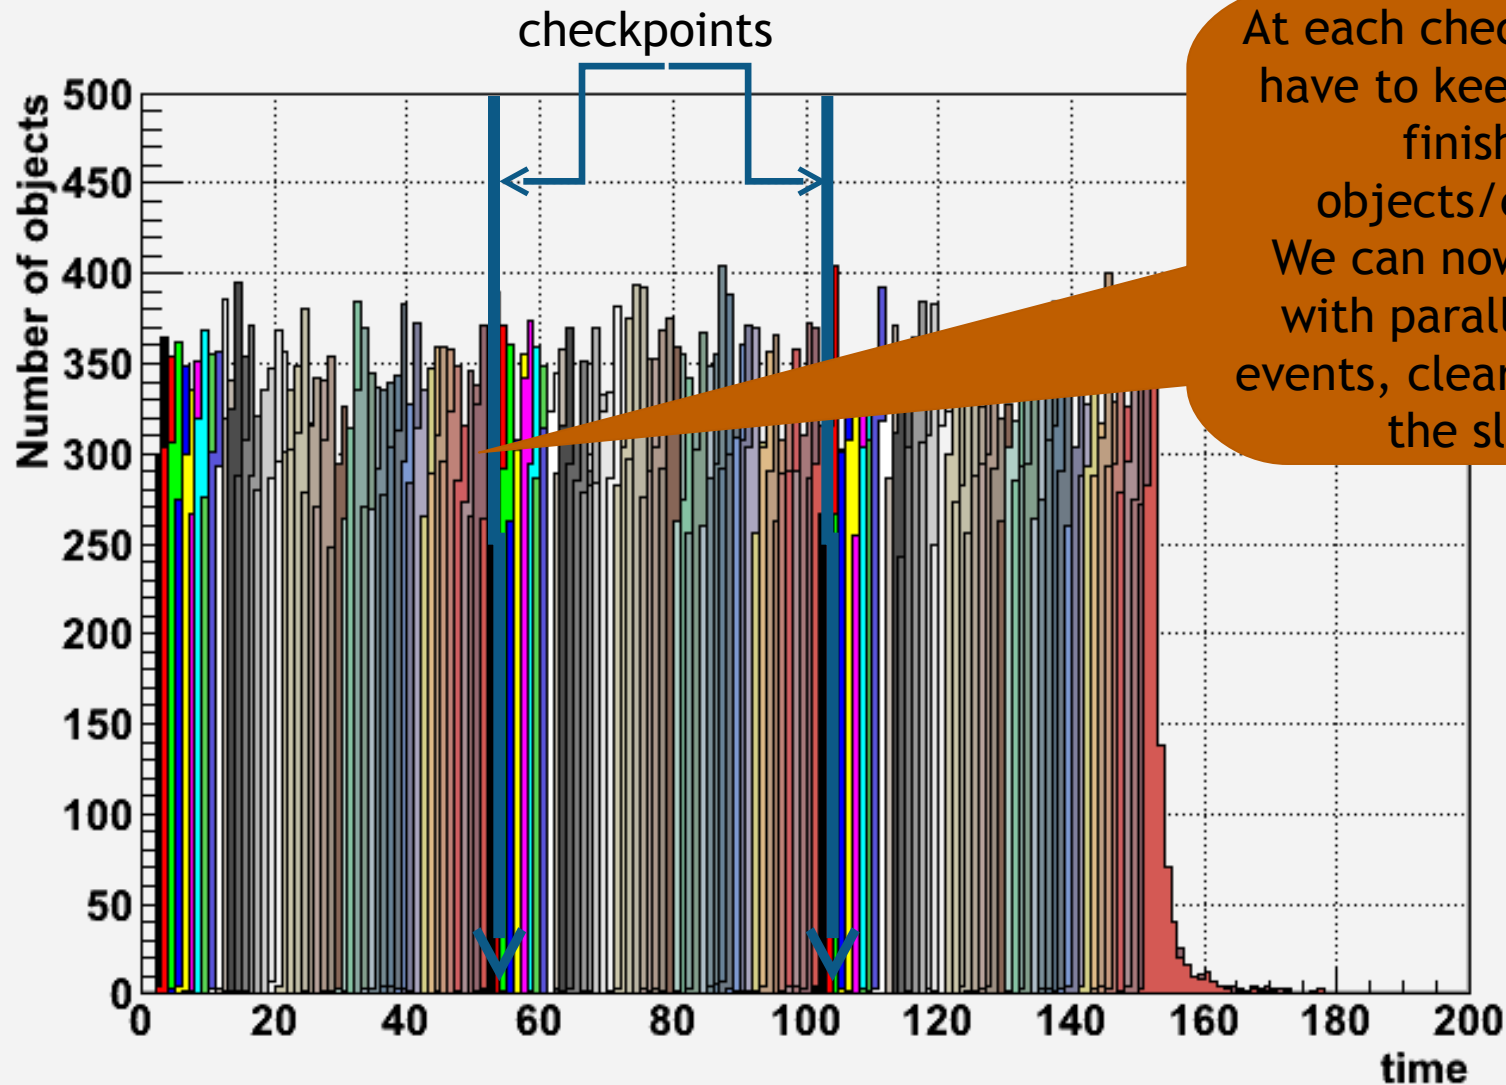
Tails again



A killer if one has to wait the end of $\text{col}(i)$ before processing $\text{col}(i+1)$

Average number of objects in memory

A better better solution



At each checkpoint we have to keep the non finished objects/events. We can now digitize with parallelism on events, clear and reuse the slots.

With CC-IN2P3

- Meeting ALICE + CC held in January
- ALICE proposes collaboration
 - We provide infrastructure
 - multi-core, many-core, clusters of GPU, MICs... will be defined soon
 - They test and give feedback
- According to the outcome, ALICE will push to work on such parallel resources
- Interesting performance test for us
- We will give access to system monitoring tools
- Use of VM (ex-CAPRI)

Another solution proposed

- Use of Grid'5000
- Full access to machines (root)
 - Deploy different OS, kernel, VM etc.
- Useful for thorough comparative tests
- ALICE must write an official request to access those machines

Conclusion

- ALICE parallel computing kick-off meeting in May
- CCIN2P3 will have new manycore machines
- Tests will begin in a few months