

Prototype for high performance (full and fast) simulation

Update of the project presented by R.Brun last year

J.Apostolakis, R.Brun, F.Carminati, A.Gheata

Jan 18, 2012

Objectives of the project

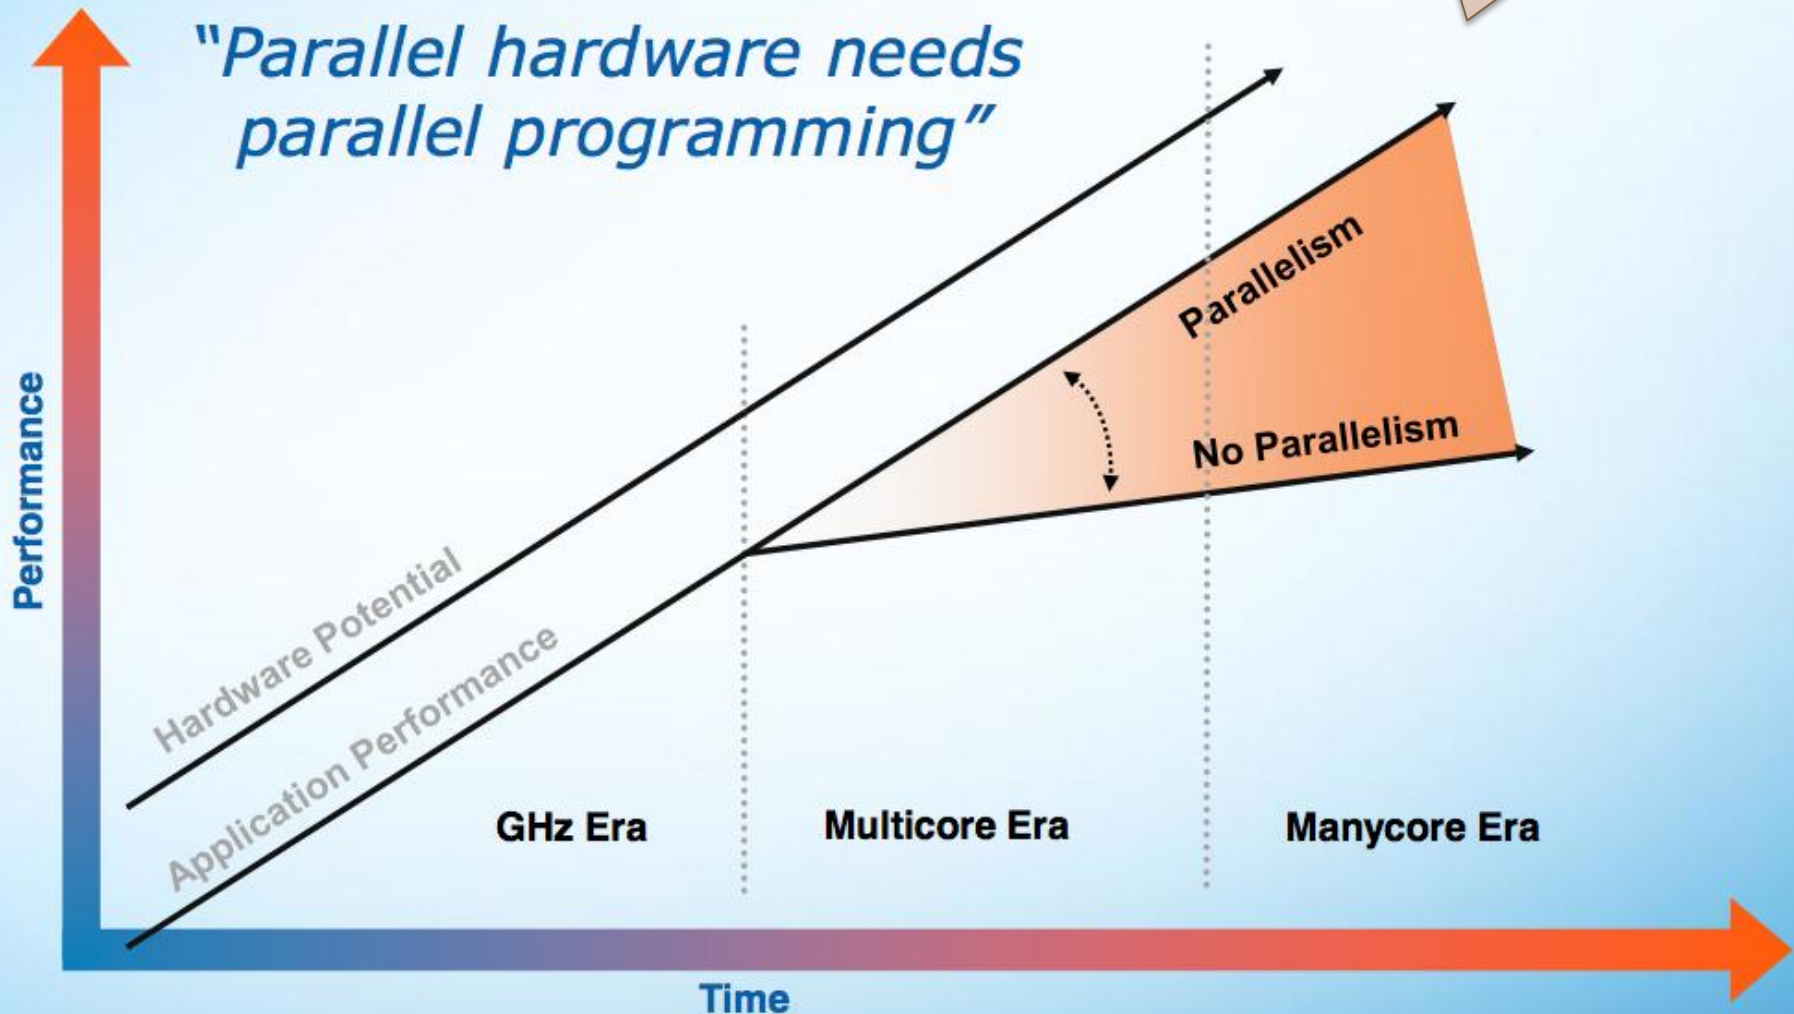
- Develop a simulation aimed at achieving maximum performance on parallel architectures
- Integrated fast and full simulation in the same framework
- Study parallelism at different granularity levels
 - Multi-processors, multi-cores, vectorisation, GPUs...
 - Explore the full “phase space” of parallelism in simulation
- Collaborate with OpenLab
- Use the large experience that is available in the SFT group on these items
- Share findings with the SFT group
- Discuss closely with experiments



Parallelism


Motivation: Performance

From a recent talk by Intel

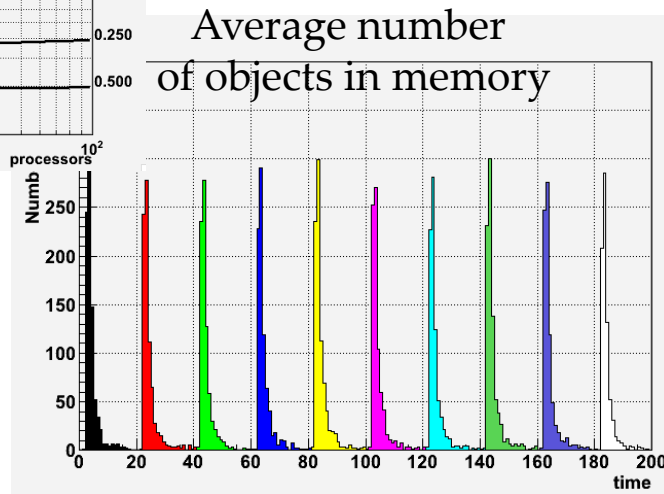


Preliminary findings

Using “real” experiment geometries

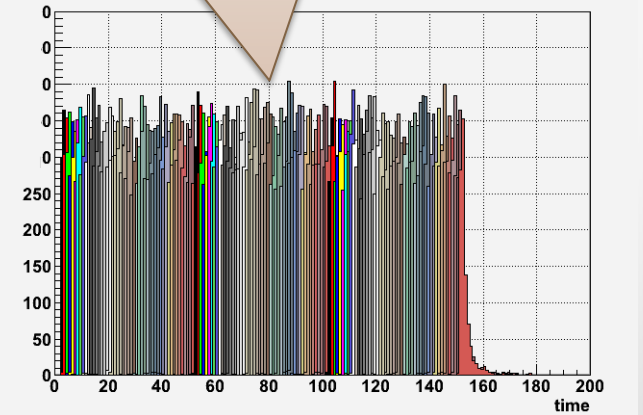


Exploring the handling of "baskets" data structures and workload splitting to fill best the processing pipeline in multi-threaded mode (high-level parallelism)



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

Lower level parallelism (within single task e.g. physics processes, geometry, ...) for optimization on CPU/GPU's is next step



Plan ahead (no timing yet)

- Continue exploring all concurrency opportunities
- Develop “virtual transporter” to include a full and fast option
- Introduce embryonic physics processes (em) to simulate shower development
- Evaluate the prototype on parallel architectures
- Evaluate different “parallel” languages (OpenMP, CUDA, OpenCL...)
- Cooperate with experiments
 - For instance with ATLAS ISF (Integrated Simulation Framework) to put together the fast and full MC.