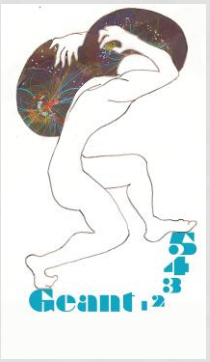

PARTICLE TRANSPORT REFLECTING ON THE NEXT STEP

R.BRUN, F.CARMINATI, A.GHEATA

A STEP FORWARD… (FUNCTIONALITY)

- The LHC (and not only) experiments have expressed the requirement of a continuum spectrum from very detailed to fast simulation
- The advantages of a "common approach" do not have to be repeated
  - Common approach between experiments
  - Common approach between full and fast simulation
- We are very aware of the difficulties along the path
- However before Geant3 even a common detailed simulation was deemed impossible

TECHNOLOGY PROJECTIONS

- Speed of computer no more driven by clock speed but by available parallelism
- Free-lunch era is over
- Parallalelisation is the only way to get the speed we need

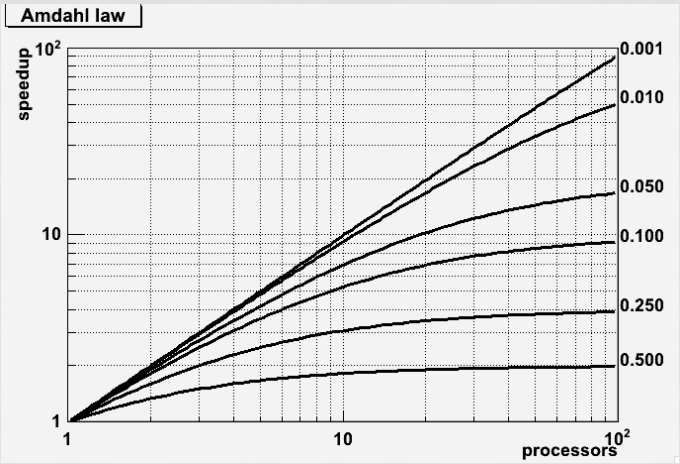

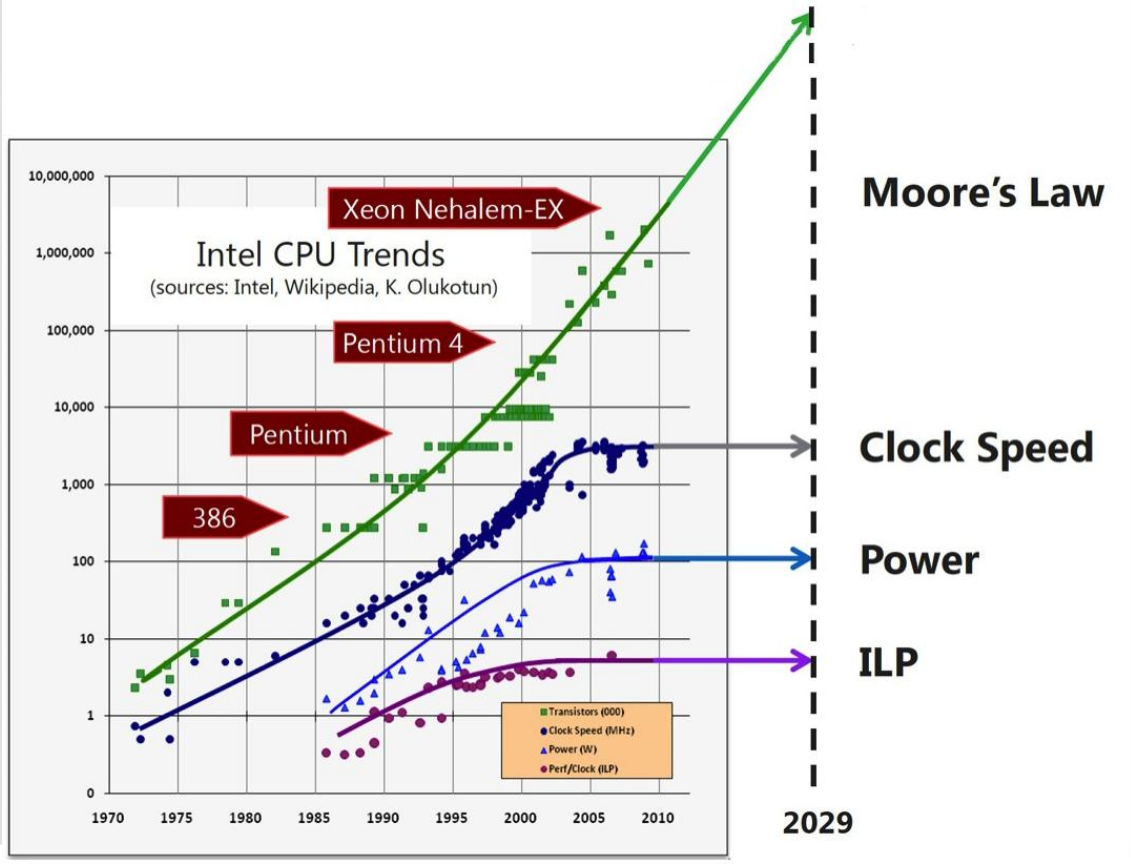

DOING NOTHING WILL ENLARGE THE GAP

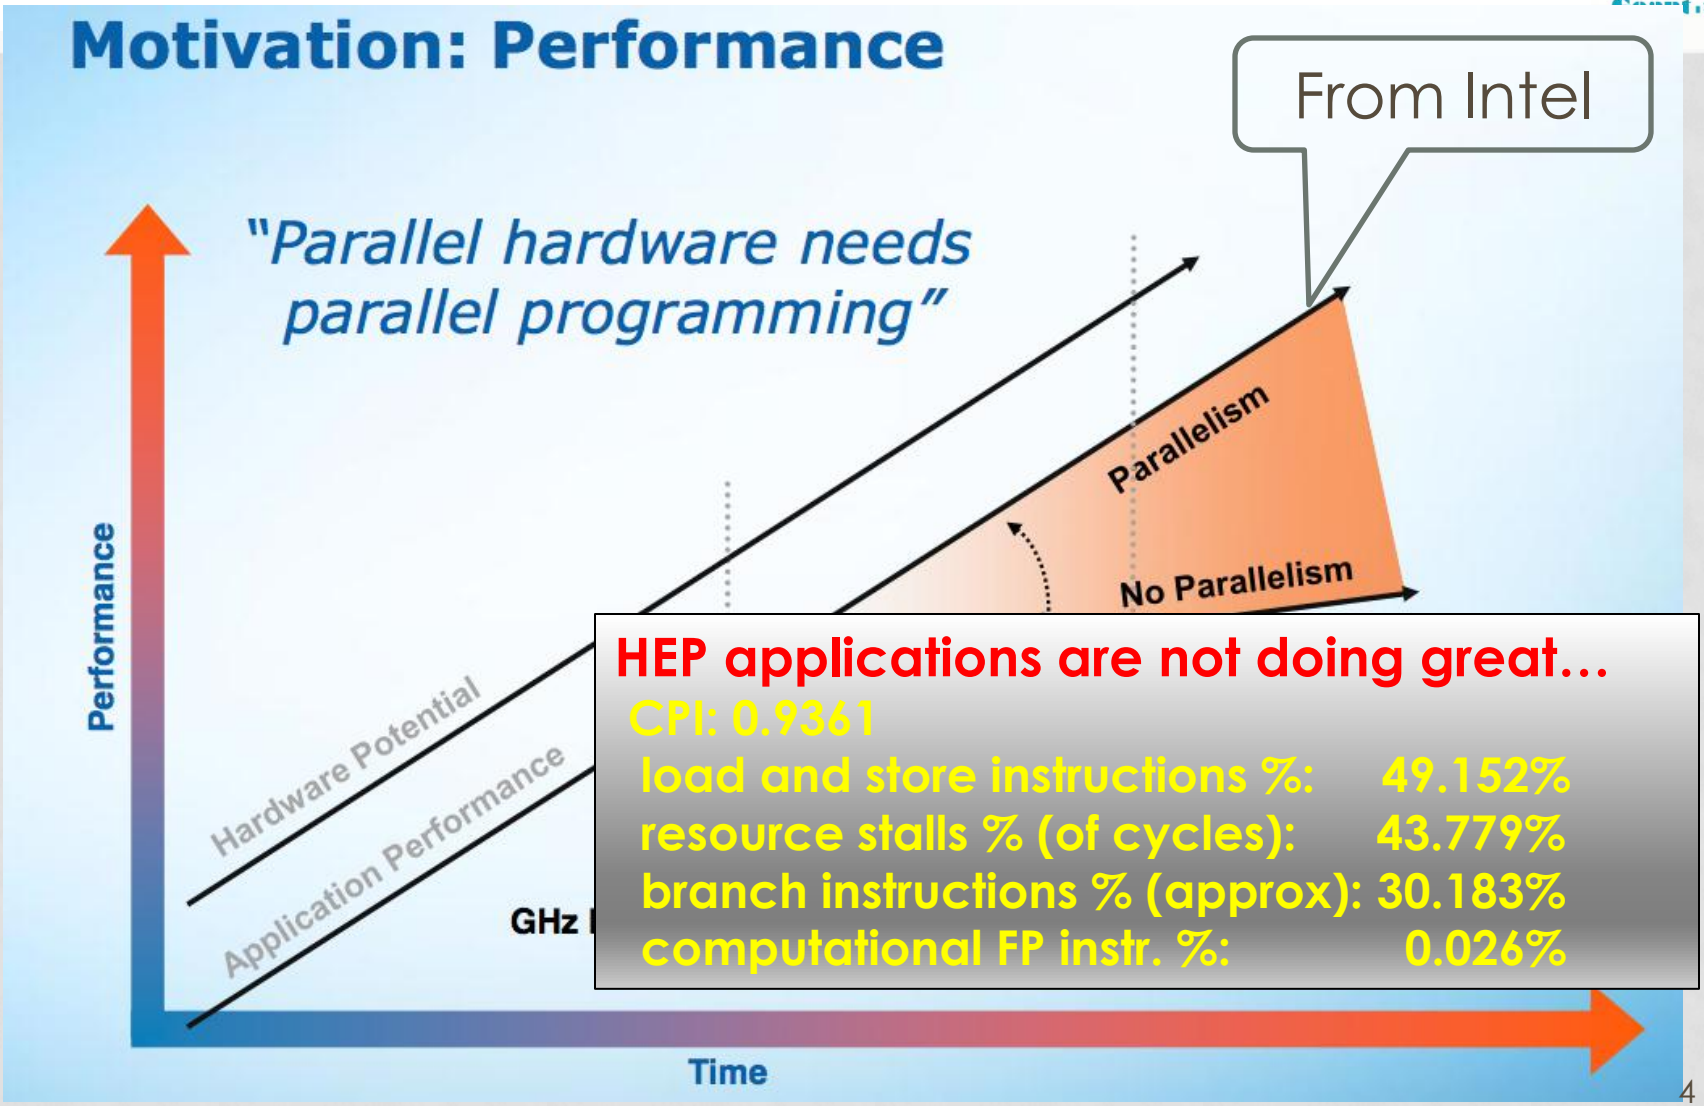

A STEP FORWARD… (SPEED)

- We have been running Monte Carlo simulations for decades
  - A very large experience has been gained, reflected in the quality of the physics, the complexity of the setups and the precision of the transport
  - State of the art full simulation in a sequential approach…
- A factor 100 is probably required for simulation
- We think this is within reach nowhere you talk about speed factors required. Both ATLAS and CMS want to see a factor 100.
- I would say that a factor 10 is reachable with the combinations fast/full along the lines of ATLAS today
- and possible a factor 5 to 10 coming from speed-up on parallel systems (vectorization + MIC-like)
- After a serious investigation, it becomes clear that a rather complete rethinking of the code becomes necessary to exploit the new hardware

THE BIG PICTURE

- The idea is to develop a new framework integrating various levels of fast and detailed simulation
- Keep services, geometry, I/O and scoring the same (as far as possible)
  - The model is the one of ROOT VMC, but substantially extended
- Develop the framework for maximum efficiency on parallel architectures to continue riding the "Moore's wave"

VIRTUALISATION!

- The experience of Root VMC

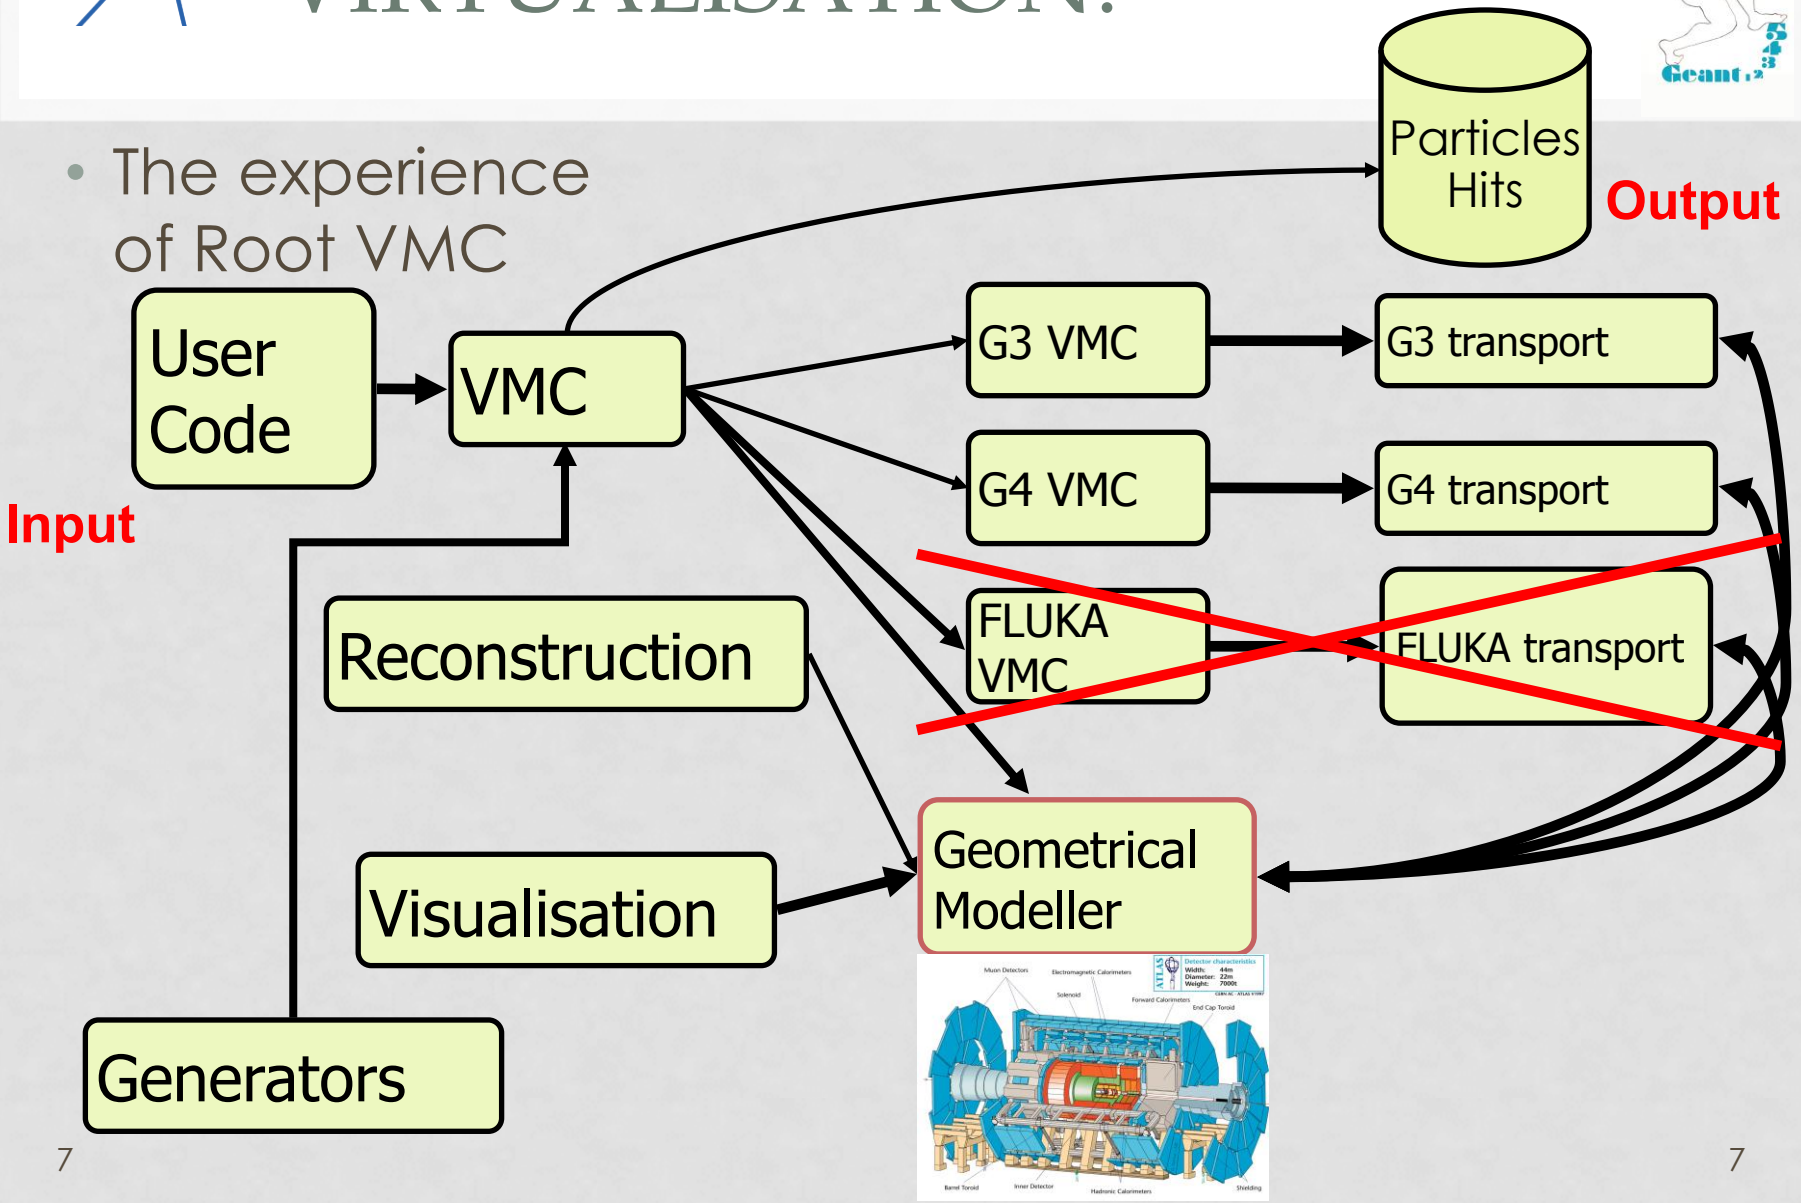

7

GLOBAL VIEW

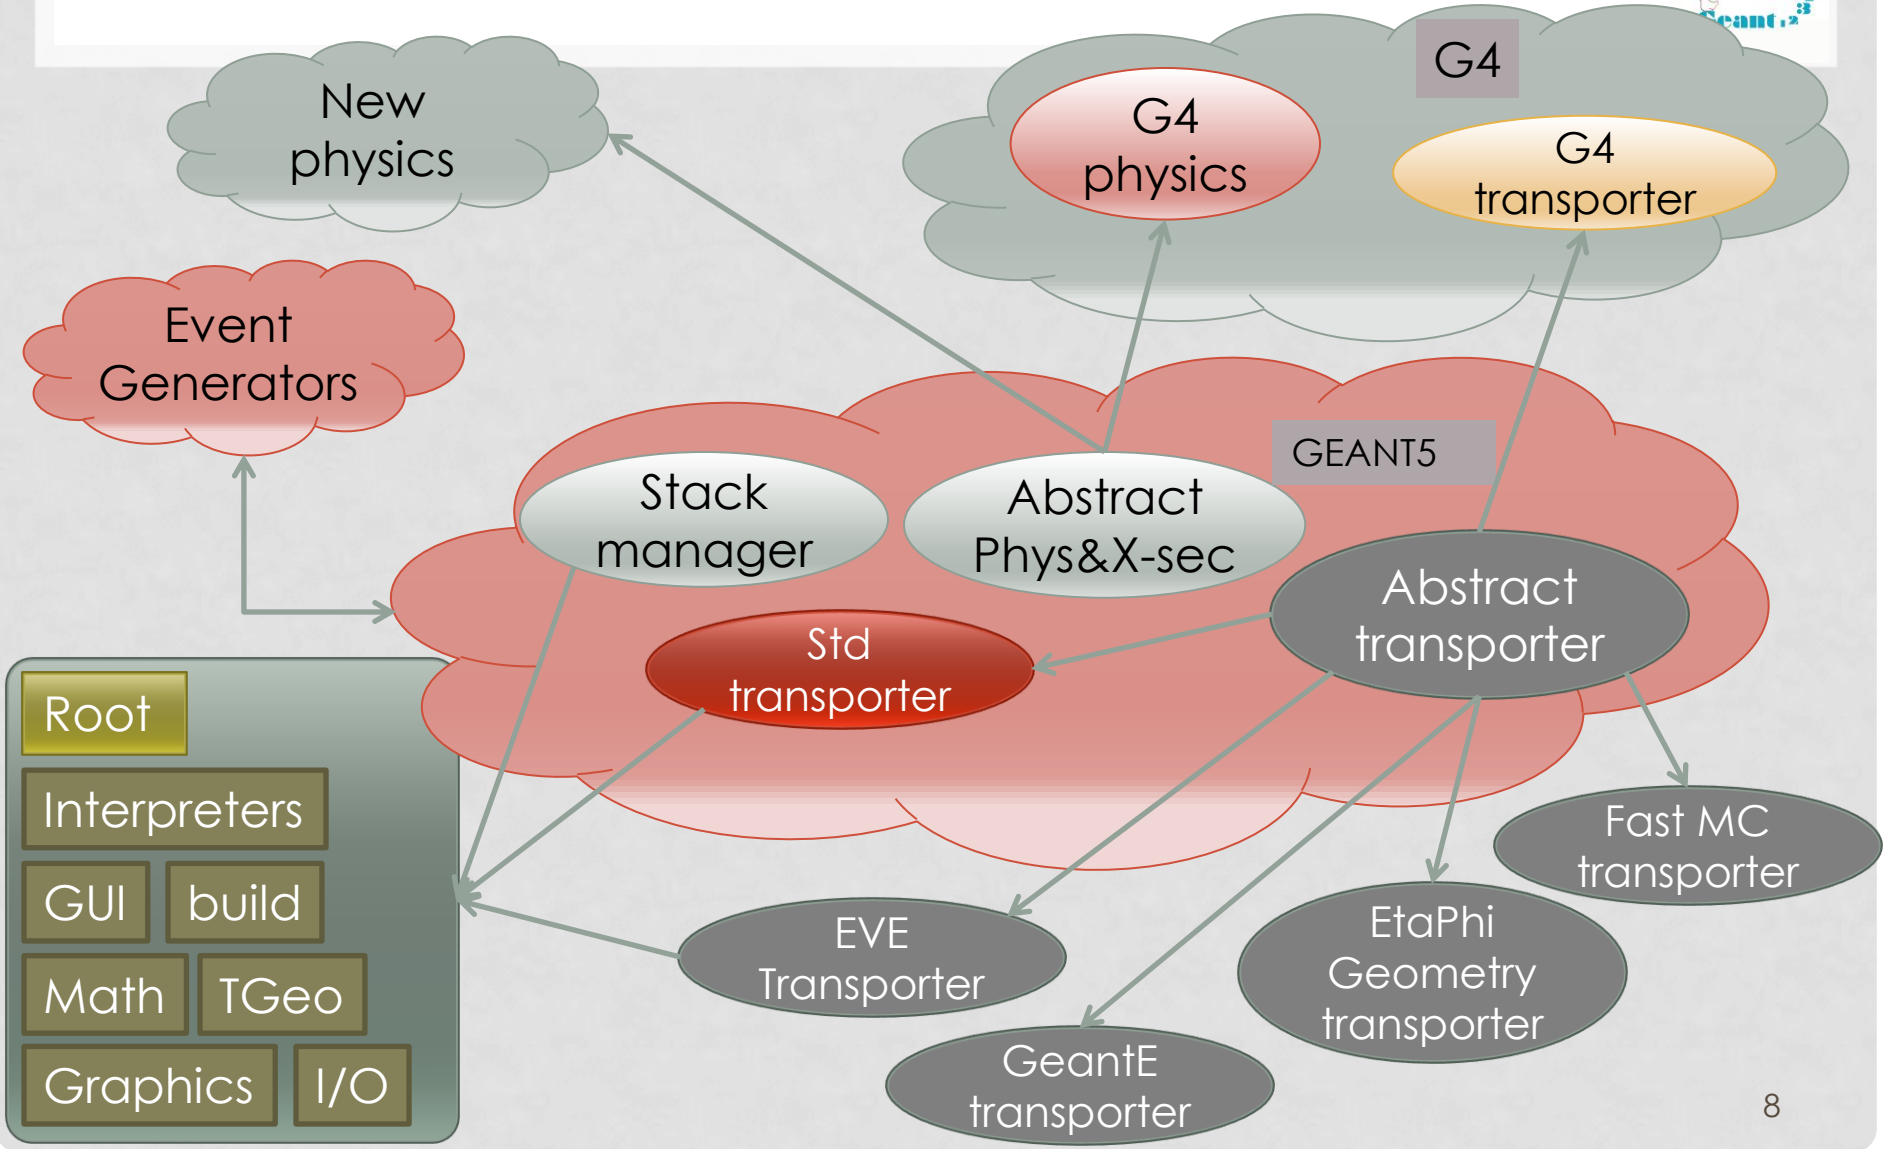
New physics
G4
G4 physics
G4 transporter
Event Generators
GEANT5
Stack manager
Abstract Phys&X-sec
Abstract transporter
Std transporter
Root
Interpreters
GUI
build
Math
TGeo
Graphics
I/O
Fast MC transporter
EVE Transporter
EtaPhi Geometry transporter
GeantE transporter

WHY A NEW EFFORT?

- The obvious "other" place to do this would be G4
- However
  - At the current level of understanding very different approaches have to be tried – impossible to do with a large code
  - We want to build into the mode fast simulation "ab initio"
  - Geant4 has to continue to provide a stable service to the experiments
- G4 MT is a very successful development, but it only addresses coarse level paralellism
- We want to explore all possible levels of parallelism

EMBARRASSING PARALLELISM ?

- Data parallelism on GRID was a savior for HEP, but…
  - Resources get short when one needs to simulate x10 the size of LHC data and uses just a tiny fraction of the CPU power
  - Fast Monte Carlo is a getaway, but cannot help in many performance studies
- Event and track level parallelism: share the code and most of the RO data structures
  - Already a step forward, but does not make jobs more efficient…
- There is an additional need to merge the outputs
  - The process may take longer across different machines than the simulation itself

GOOD NEWS: HEP TRANSPORT IS MOSTLY LOCAL !

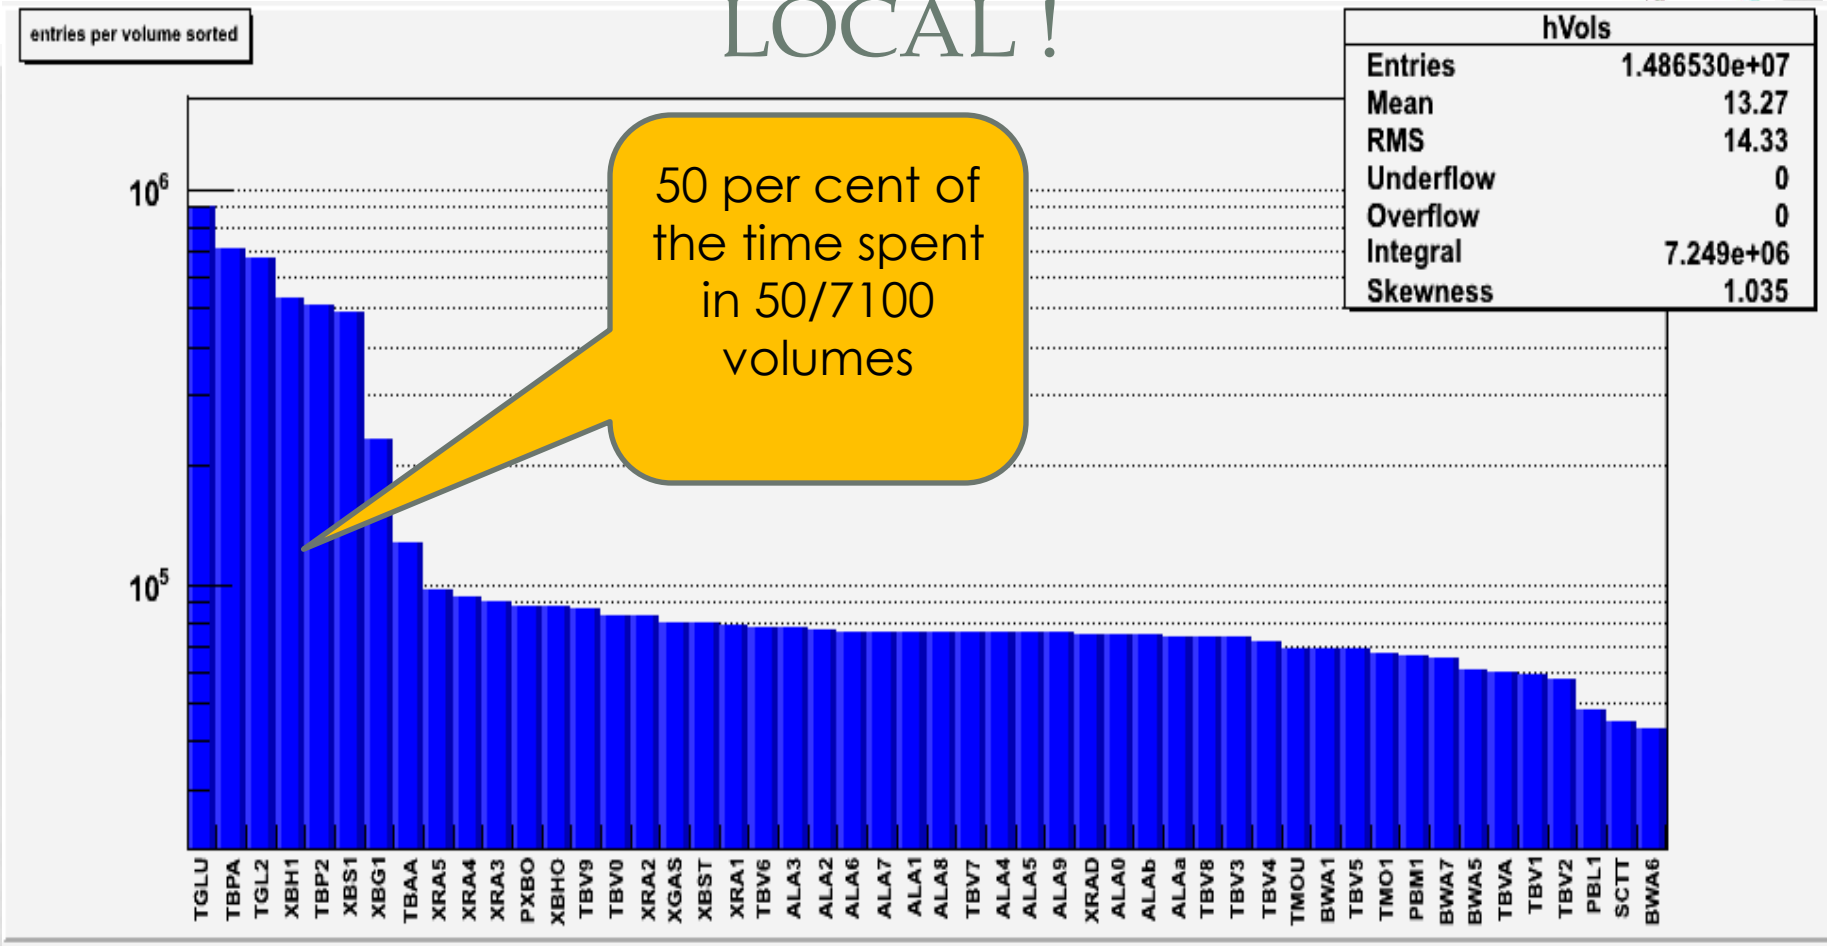

ATLAS volumes sorted by transport time. The same behavior is observed for most HEP geometries.

BAD NEWS: TAILS… TAILS…

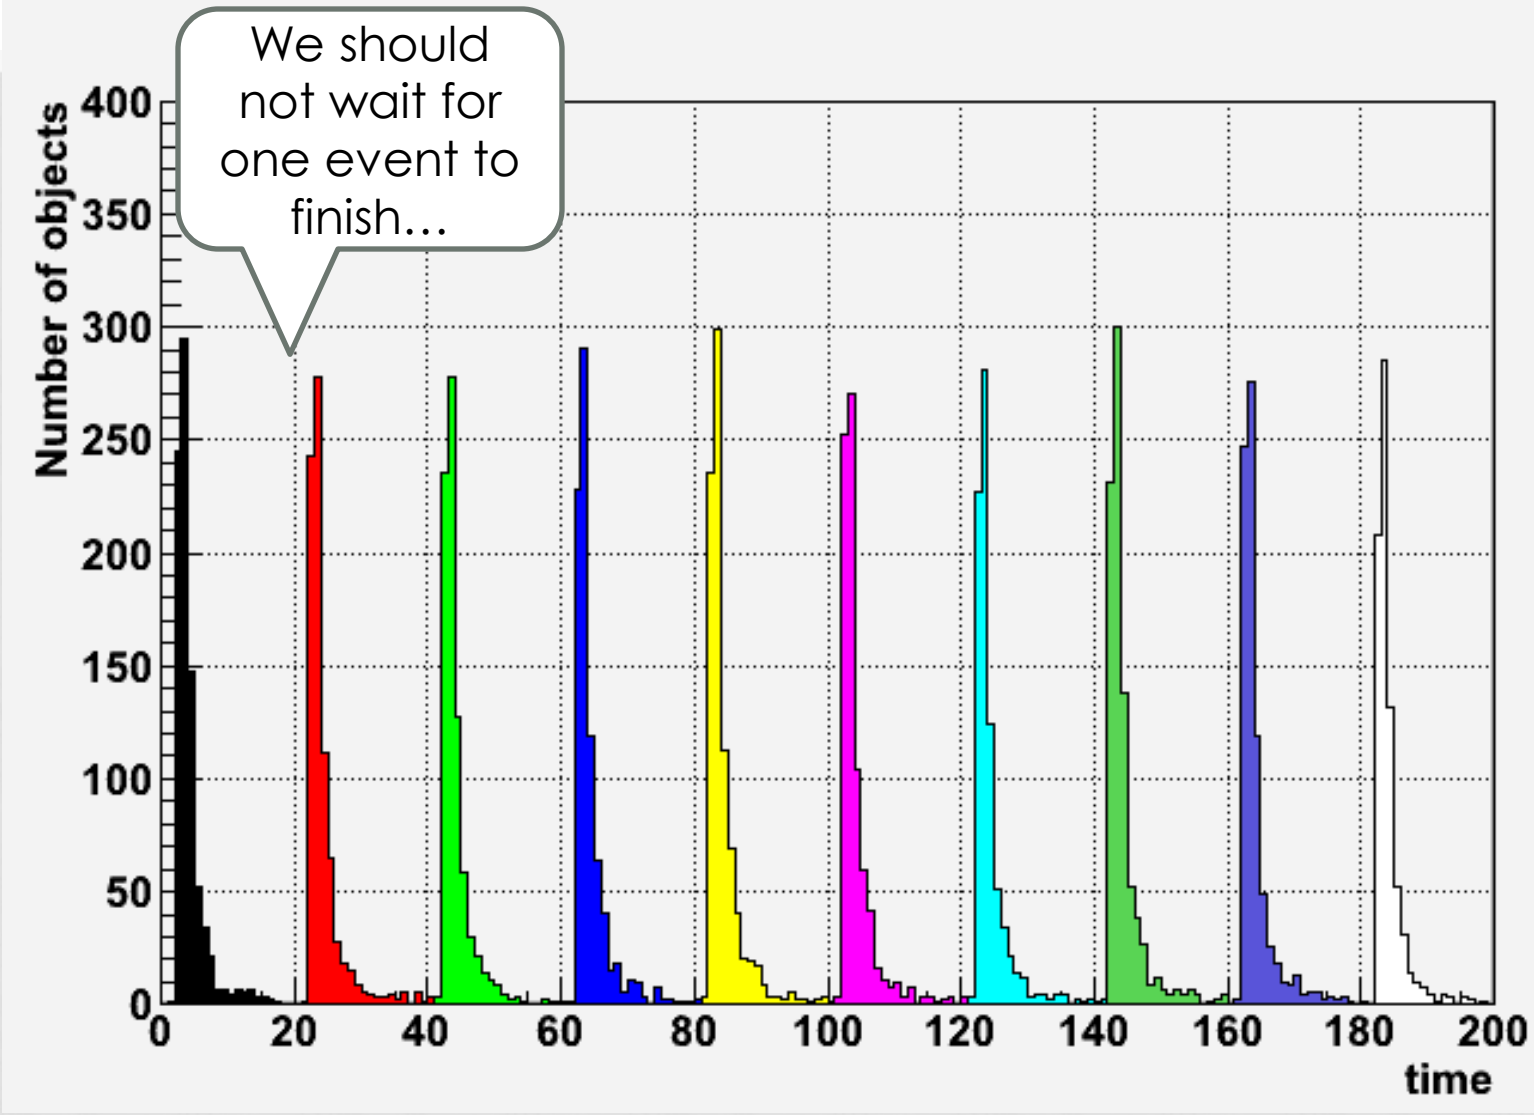

VOLUME-ORIENTED TRANSPORT MODEL

- In our model all particles traversing a given geometry volume are transported together as a vector until the volume is empty
  - **Same volume -> local (vs. global) geometry navigation, same material and same cross sections**
  - Load balancing: distribute all particles from a volume type into smaller work units called **baskets**, give a basket to a transport thread at a time
- Particles exiting a volume are distributed to baskets of the neighbor volumes until exiting the setup or disappearing
  - Like a Champaign cascade, but lower glasses can also fill top ones…
  - No direct communication between threads to avoid synchronization issues
- Only toy physics for the moment

THE TRANSPORT PROTOTYPE

More events better to cut event tails and fill better the pipeline !

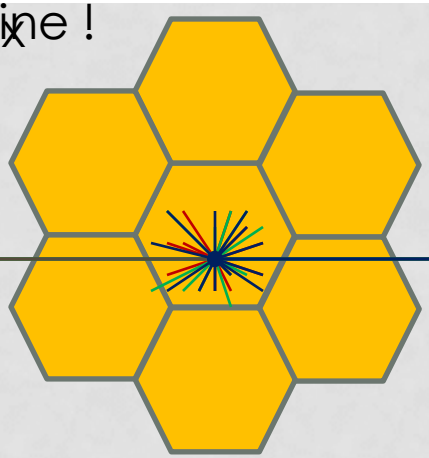

Associate a set of "baskets" to each geometrical logical volume

A FIRST APPROACH

Each thread transports its basket of tracks to the boundaries of the current volume
Move crossing tracks to a buffer, then picks-up the next basket from the queue

Work queue

Transport threads pick-up baskets from the work queue

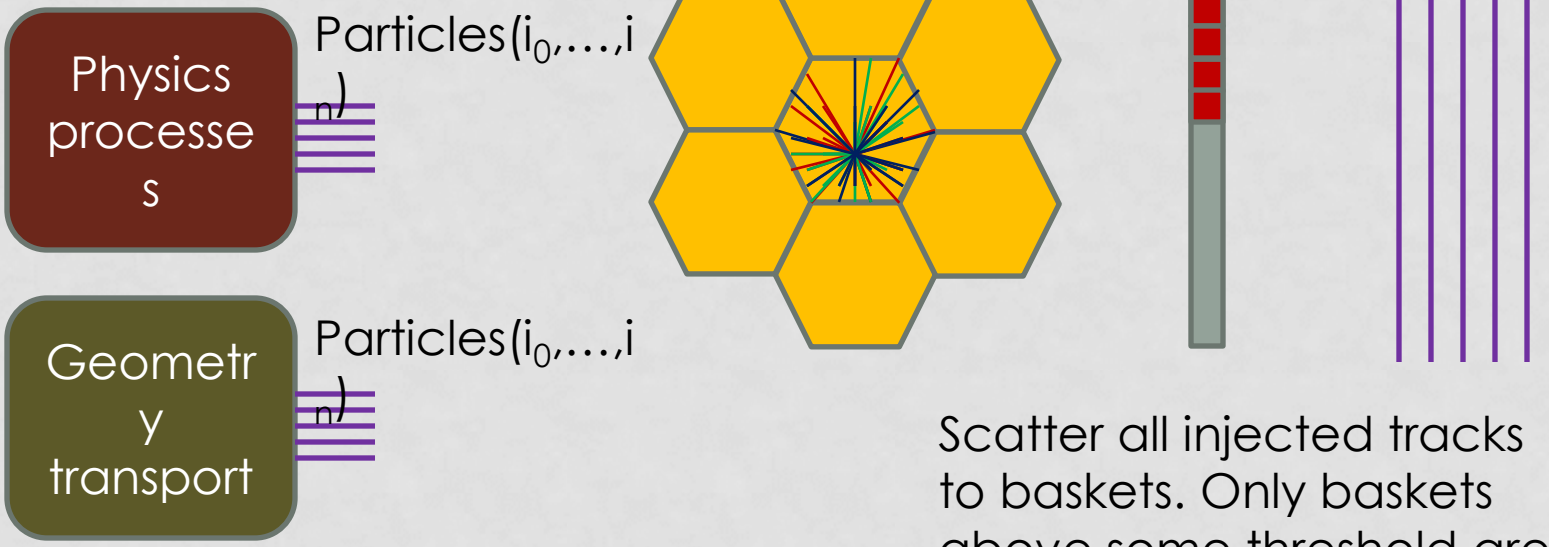

Scatter all injected tracks to baskets. Only baskets above some threshold are transported.

Physics processes and geometry transport called with vectors of particles

FIRST VERSION REQUIRED SYNCHRONIZATION…

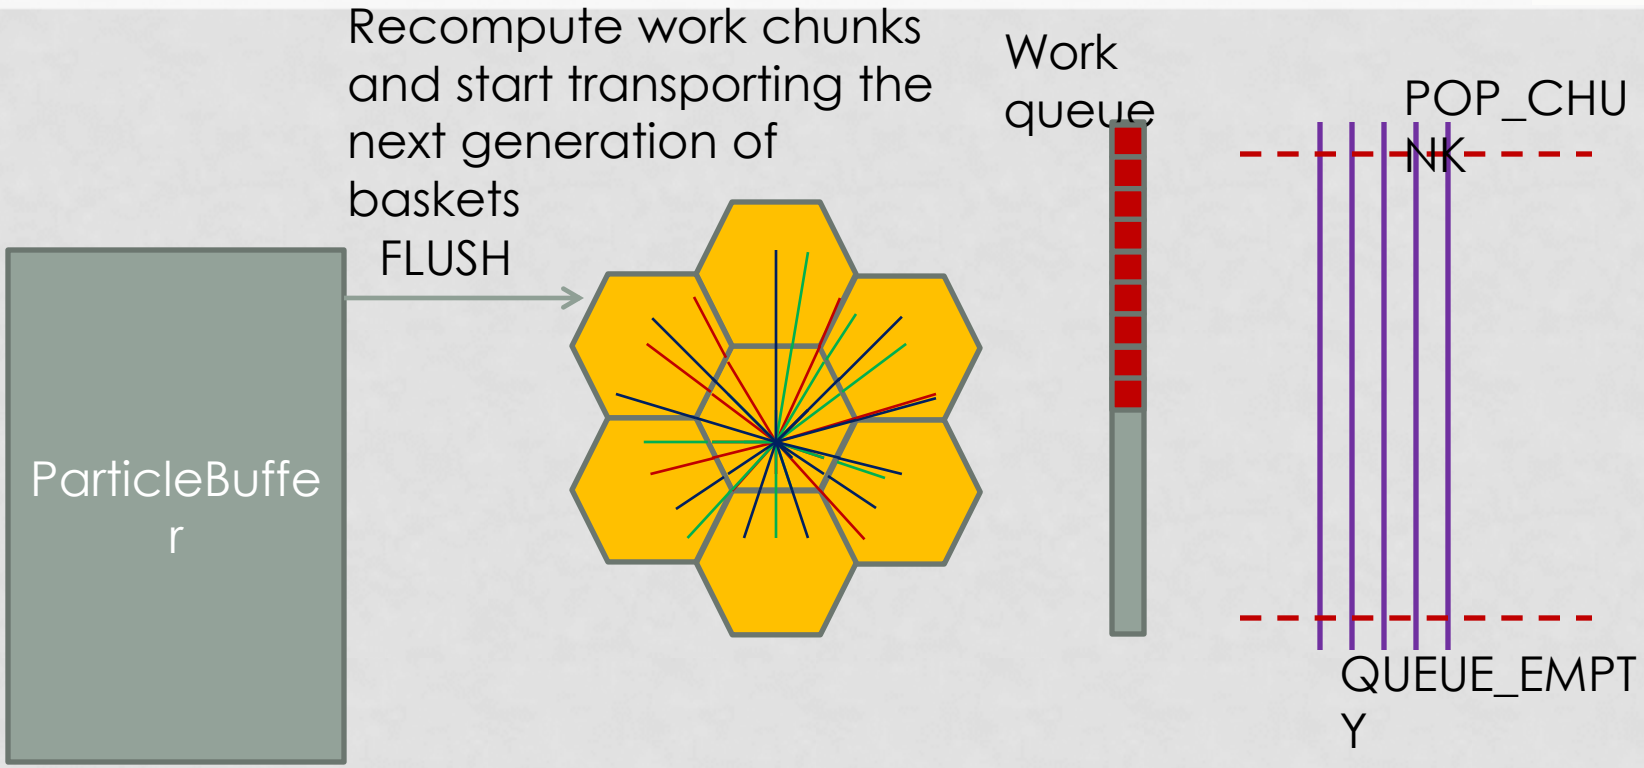

Synchronization point: flush transported particle buffer and sort baskets according content

Generation = Pop work chunks until the queue is empty

EVOLUTION OF POPULATIONS

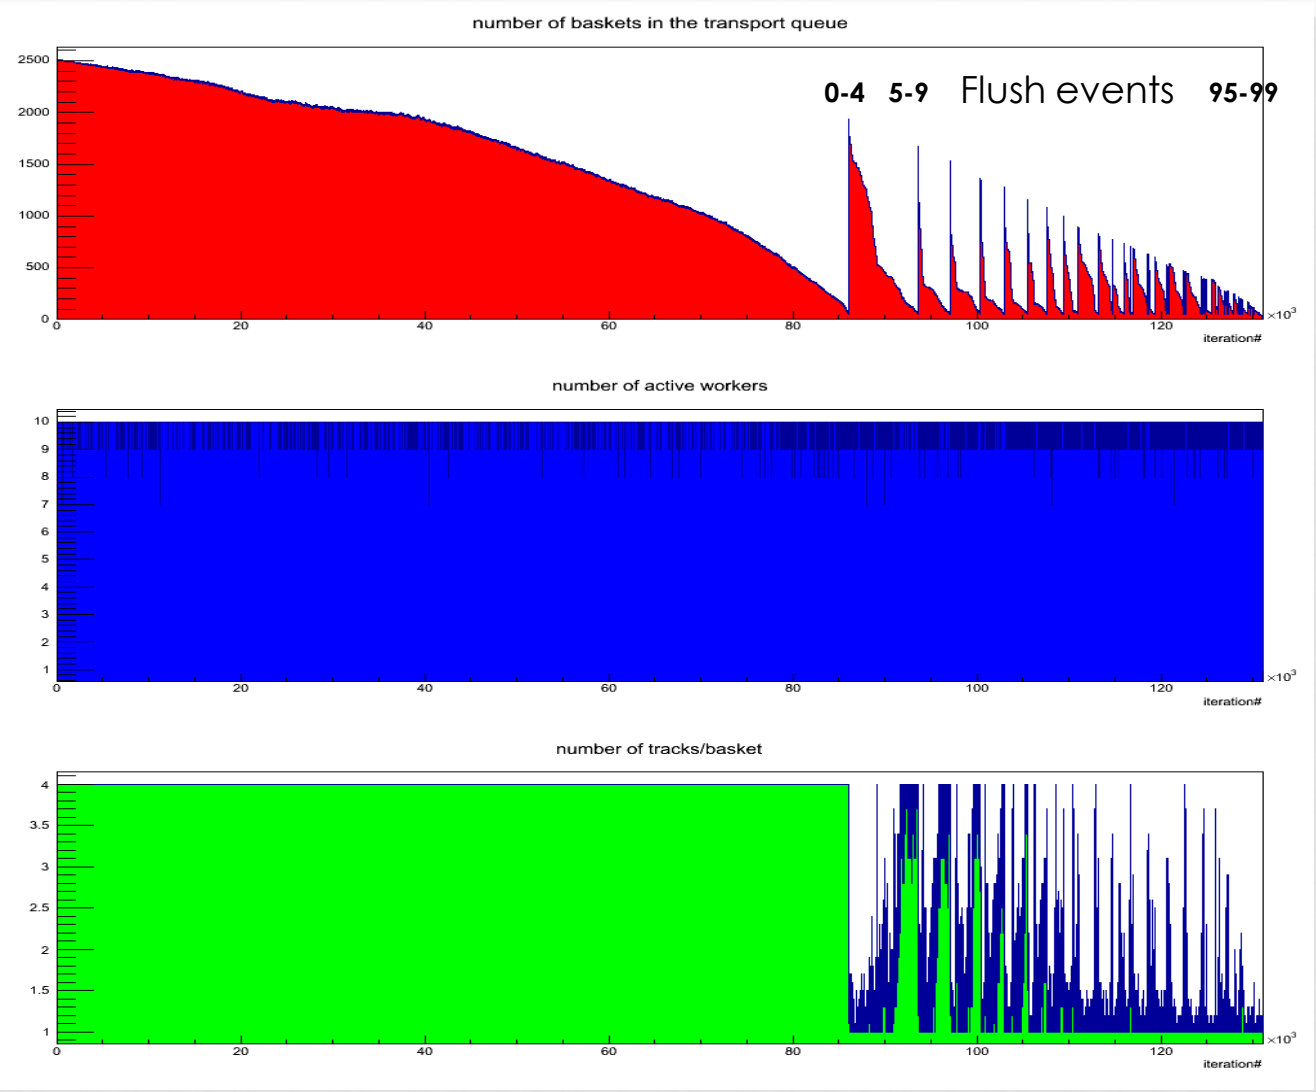

LESSONS LEARNED

- A model requiring synchronization stages is not appropriate
  - Very large Amdahl effects, increasing in the track depletion stage
- Balancing basket populations in time is not trivial
  - Events need to be injected in the system to compensate the basket inefficiencies
    - Will cumulate hits and make memory grow
  - Hits from early introduced events need to be evacuated to the digitization and I/O threads
- We need a good estimate of the percentage of work that can be done with "efficient" vectors
  - A model including realistic physics, digitization and I/O will be needed

CONCURRENCY IN THE FIRST APPROACH

- Ideally all workers should be in running state during the processing phase, so the distribution should peak for the number of workers
- Synchronization becomes critical during the particle depletion regime, when particle baskets are non-optimally filled and garbage collections more often

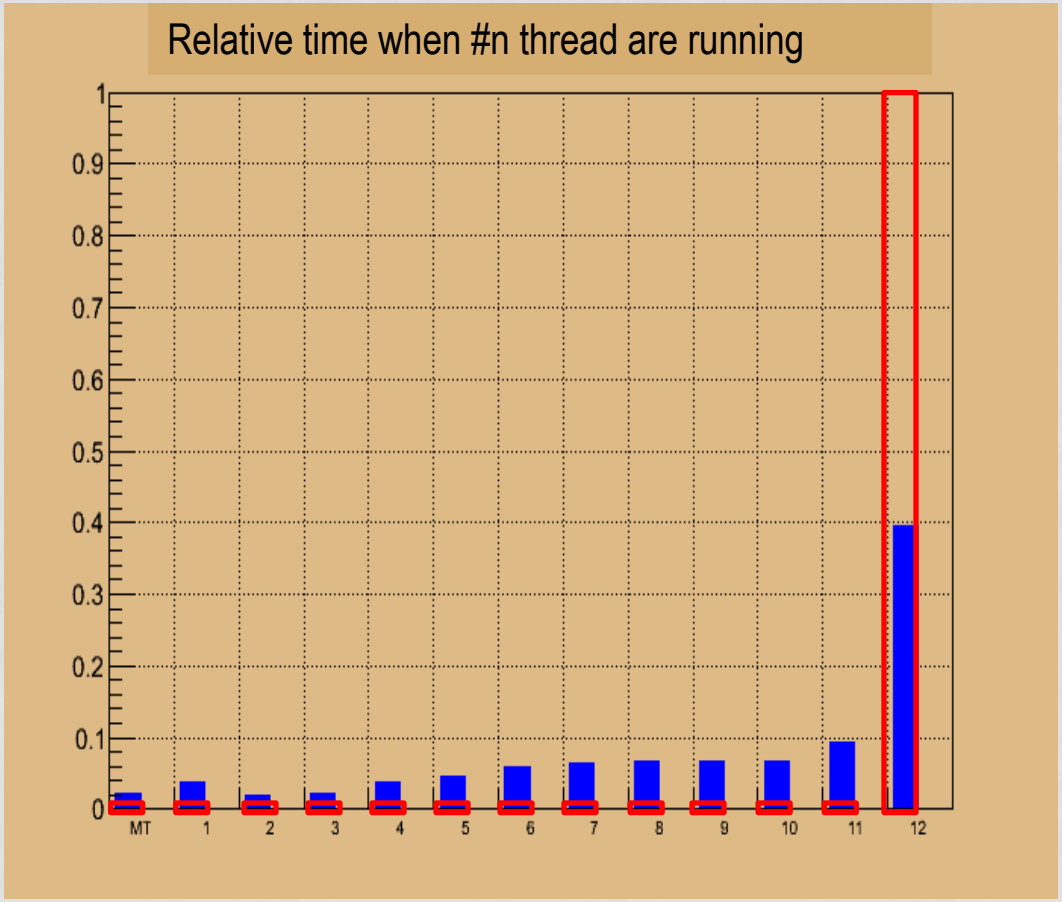

PRELIMINARY BENCHMARKS

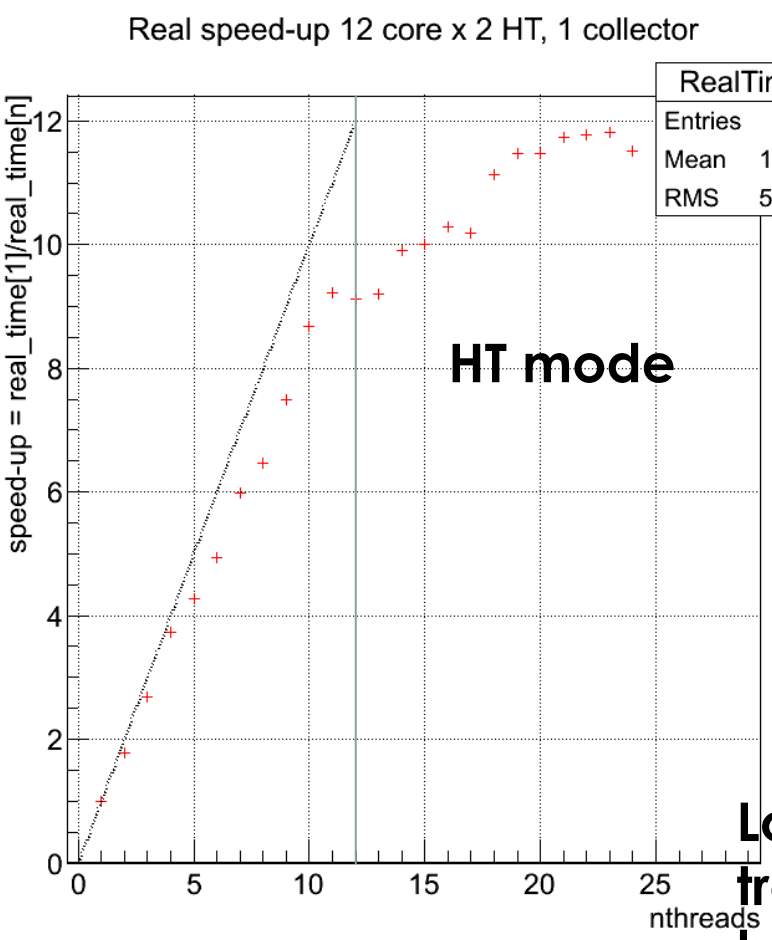

Benchmarking 10+1 threads on a 12 core Xeon

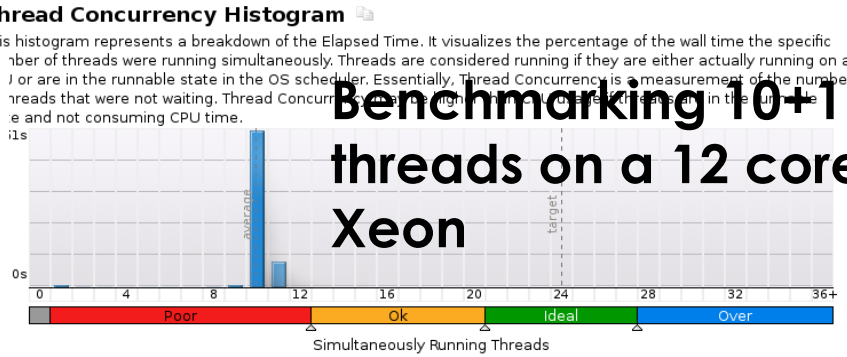

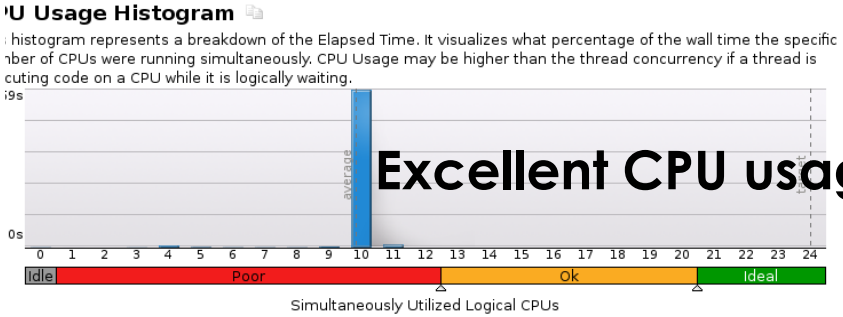

Excellent CPU usage

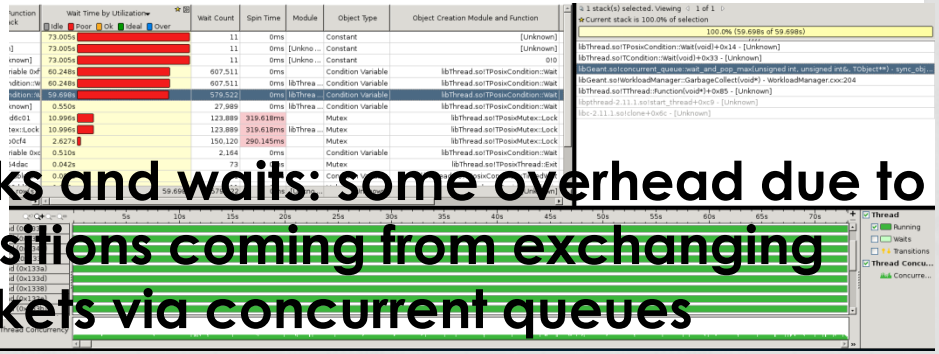

Locks and waits: some overhead due to transitions coming from exchanging baskets via concurrent queues

FUTURE PLANS

- Continue the investigation of parallel event transport
- Develop more realistic physics models
- Integrate fast simulation
- Aiming at a working prototype in 2015
- The activity will start in earnest in September 2012
- As with other large successful projects (ROOT, G4), this will be an international collaboration
  - We count on the HEP community (LHC, ILC, CLIC, FAIR, …) for help and feedback

CONCLUSIONS

- The new generation (Geant5) of detector simulation programs will have to
  - Integrated seamlessly fast and detailed simulation at different levels
  - Make efficient use of parallelism at different levels
  - Capitalizing on the large Geant 1-4 experience
- A prototype is being built at CERN, which will require collaboration with the HEP Community at large
- The first results are interesting and our learning curve very steep (!)
- Stay tuned…