

Speed issues in shower MC

David Grellscheid, IPPP Durham

Summary of blackboard talk

- In Herwig++, Shower/MPI/Hadronization/Decays can process events at ~60 ev/sec. on 3 GHz == ~50M clock ticks per event.
- Fixed init time before the run: ~1500 events' worth / ~30s.
- Share of runtime: Shower/MPI 50%, Hadr. 27%, Decays 23%
- Algorithms are *very* sequential, no good handle to parallelize within the event
- But: trivial to parallelize across events, easy to keep CPUs busy. No need to do internal threading etc.; other CPUs are busy with their own event streams anyway
- Outsourcing these steps to GPU not worth the effort: Architecturally, 32 threads in a warp need to follow identical instruction stream. Code branches lead to threads running 0-ops when the branch does not apply to them; 32 MC events *will* go down 32 different branches => parallelisation factor of 32 is lost.
- However, GPUs for Matrix Elements make a lot of sense! Needs more investigation to link up with serial downstream processing.

Conclusions

- Time spent on parallel implementations after the MEs is not useful, unless we rethink theory algorithms completely with parallel implementations in mind.
- Time spent on profiling and increasing serial execution speed is more promising. Several low hanging fruit still wait to be discovered. Recently: 20% speedup from 3-line code change. Relevant for situations where event generators are run stand-alone with fast MEs and fast downstream processing.