

Exploitation of modern CPU
architectures and its impact on the
computing model of HEP
experiments

Goals of the Workshop

- Review of current industrial technologies
 - Hardware
 - Software
- Trend for the future
 - Short term
 - Next decade
- Review of current activities in HEP

Presentations

- Reviews (hardware, software, activities)
 - Vincenzo: Multi-Core
 - Leone: GPUs
- Specific topics
 - Alfio: “final” physics analysis
 - Gerri: I/O (root) and analysis of large datasets
 - Karen: physics analysis on GPUs

Goals of the Workshop

Exploitation of modern CPU means parallelization

- New Evaluate impact on
 - SuperB computing model
 - Software Architecture
- Identify the most affected areas
- Propose a short term R&D program

Issues for discussion R&D

- Optimization of scalar code
 - It will not go faster because of technology
 - We need to go back to trading accuracy for speed
- Vectorization (in core parallelism)
 - Compiler level auto-vectorization is not enough
 - We may need to rethink algorithms to exploit it

Challenge here is to achieve a significant speed-up while keeping the code flexible and readable/maintainable by a large group of people

Issues for discussion R&D

- High grain parallelism (inside algorithms)
 - Many paradigms around (OMP, TBB, Cilk++, Cl, std::thread,...)
 - We need to evaluate the various approaches and see which best fit our environment
 - Try with existing algorithms (pattern recognition and tracking for instance)

Issues for discussion R&D

- Framework level parallelism
 - Beyond the well established event by event parallelism
 - Algorithm pipelines, parallel I/O
 - Needs to establish the actual gain of different scheduling strategies
 - Leverage Babar's software and data to build "emulators"
- Keep an eye on emerging technology
 - GPGPU: established for specialized use-case, can be used for general computing?
 - Experiment with (HL)T algorithms