MANY-CORE EXPERIENCE WITH HEP SOFTWARE AT CERN OPENLAB

SVERRE JARP, ALFIO LAZZARO, JULIEN LEDUC, ANDRZEJ NOWAK

CERN openlab regularly benchmarks and analyzes major representatives of the x86 processor family. Even though the processor is the heart of the platform, performance is becoming increasingly influenced by the platform itself (as well as its firmware) and the software that runs on it.

A set of carefully selected benchmarks serves as the foundation for a reference measurement. The programs used are representative and cover major corners of HEP processing: General C++ workloads, Simulation and Data Analysis.

Inside a modern computing platform, parallelism is expressed on many levels . These multiplicative dimensions of performance grow from generation to generation of hardware. A gain in one of them often is a gain in all - ensuring a continued application of Moore's Law.

THE HARDWARE

Dual-socket Intel Xeon E5-2680 @ 2.7 Ghz
16 cores / 32 threads: "Sandy Bridge-EP"

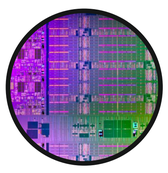

Quad-socket Intel Xeon E7-4870 @ 2.4 Ghz
40 cores / 80 threads: "Westmere-EX"

THE ENVIRONMENT

Benchmarking of modern processors is an increasingly complex affair. One has to control (at least) the following features: processor frequency, overclocking via Turbo mode, the number of physical cores in use, the use of logical cores via Simultaneous Multi-Threading (SMT), the cache sizes available, the memory configuration installed, as well as the power configuration if throughput per watt is to be measured. In latest tests even the OS version makes a difference in platform power consumed and performance limits.

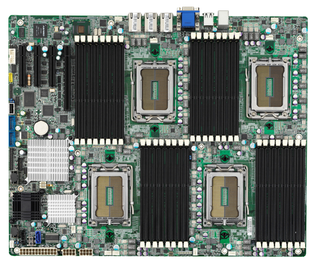

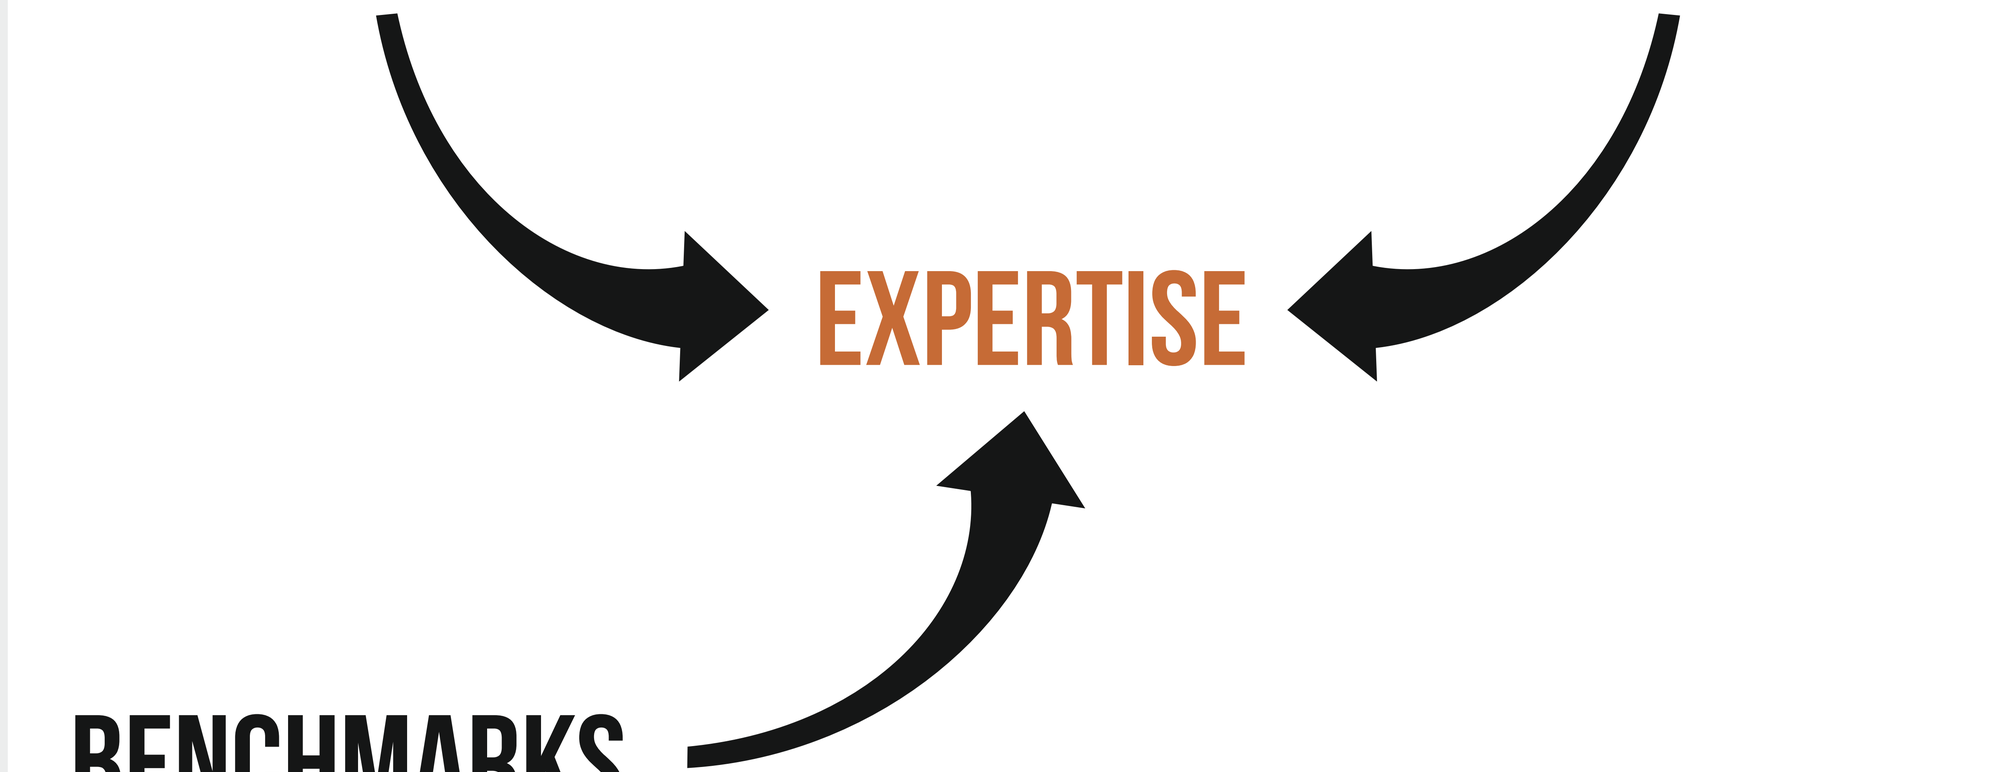

HEPSPEC06

The SPEC CPU2006 benchmark from the SPEC Corporation is one of the important performance benchmark suites in the IT industry. It can be used to measure both individual CPU performance and the throughput rate of servers. SPEC CPU2006 is designed to stress a system's processor, the caches and the memory subsystem. It is based on real user applications, and the source code is commercially available. A High Energy Physics (HEP) working group demonstrated good correlation between the SPEC results and High Energy Physics (HEP) applications when using the C++ subset of the tests from the SPEC CPU2006 benchmark suite. As a result, and because SPEC2006 is readily available, the HEP community has decided to use the C++ subset of SPEC2006, dubbed "HEPSPEC06", rather than internal benchmarks.

Multi-threaded Geant4 prototype

Geant4 is one of the principal toolkits used in Large Hadron Collider (LHC) simulation. Its primary purpose is to simulate the passage of particles through matter. This type of simulation is a CPU-intensive part of a bigger overall framework used to process the events coming from the detectors. It is representative to an extent of real life workloads and can constitute a substantial portion of the CPU time of the Worldwide LHC Computing Grid. Since HEP has always been blessed with parallelism inherent in the processing model, it is natural to try to utilize modern multi-core systems by converging towards multi-threaded event processing. The Geant4 prototype discussed here is one of the key steps in that direction. Parallelized with pthreads and ParTopC, it delivers predictable scalability on a wide range of platforms.

Maximum Likelihood Fit

The HEP community makes large use of many complex data analysis techniques, like maximum likelihood fits, neural networks, and boosted decision trees. These techniques are employed for a better discrimination between interesting events with respect to the total events collected by the physics detectors, in order to discover possible new physics phenomena. This benchmark, developed by CERN openlab, is based on an unbinned maximum likelihood data analysis application. It represents a prototype of the RooFit package (a package inside the ROOT software framework developed at CERN), generally used in the HEP for maximum likelihood fits. This prototype makes use of an optimized algorithm with respect to the algorithm used in RooFit for likelihood evaluation, with OpenMP based parallelization and vectorization applied.

RESULTS

Maximum Likelihood Fit

"Westmere-EP"
vs.
"Sandy Bridge-EP"

(higher is better)

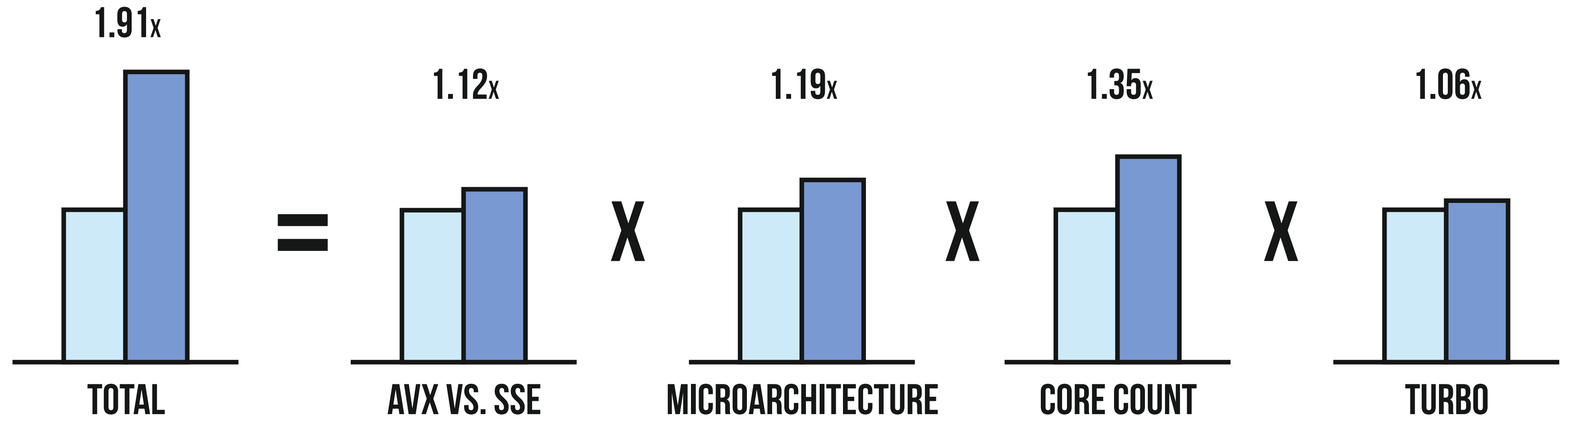

HEPSPEC06: Sandy Bridge-EP features

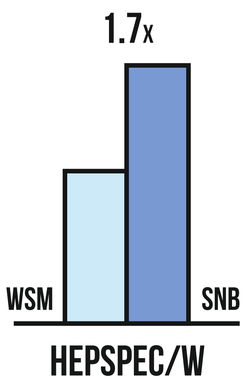

	SMT off	SMT on
SLC6 vs. SLC5	+9%	+6.5%

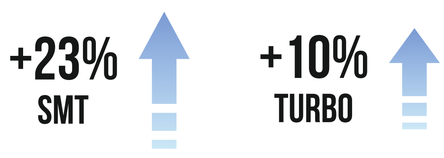

Multi-threaded Geant4 prototype on Westmere-EX
(40 cores, 80 threads)

	1 core	20 cores	40 cores	80 cores
Scaling	1x	20.2x	40.2x	49.3x
SMT				+23%

COMPUTING IN HIGH ENERGY AND NUCLEAR PHYSICS 2012