High core count and low power systems evaluation

Paweł Szostek, Andrzej Nowak, Georgios Bitzes
{pawel.szostek, andrzej.nowak, georgios.bitzes}@cern.ch

CERN openlab Board of Sponsors, May 7th 2014, Geneva, Switzerland

Hardware

CERN openlab regularly benchmarks and analyzes major representatives of the x86 processor family to see how well they can fit .
In modern CPUs complexity of benchmarking is constantly growing, as more and more factors contribute to overall performance.

Quad-socket Intel Xeon E7-4870 v2 @ 2.4 GHz
60 cores / 120 threads
37.5MB L3 Cache per socket
"Ivy Bridge" microarchitecture
Up to 2.8GHz with Turbo Boost
130W TDP

Dual-socket Intel Xeon E5-2695 v2 @ 2.4GHz
24 cores / 48 threads
30MB L3 Cache per socket
"Ivy Bridge" microarchitecture
Up to 3.2GHz with Turbo Boost
115W TDP

Single-socket Intel Atom C2730 @ 1.8 GHz
8 cores / 8 threads
4MB L2 cache
"Silvermont" microarchitecture
Up to 2GHz with Turbo Boost
12W TDP

Single-socket Intel Xeon E3-1285 v3 @ 3.1GHz
4 cores / 8 threads
8MB L3 Cache
"Haswell" microarchitecture
Up to 3.9GHz with Turbo Boost
65W TDP

Benchmarks

A set of selected benchmarks representing major areas of HEP computing is used for a reference measurement.

HEPSPEC 06

The SPEC CPU2006 benchmark from the SPEC Corporationis 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. It consists of a set of single-threaded workloads. Thus, in order to leverage all cores available on a CPU, we run a separate process on each core and sum up the obtained scores.

Multi-threaded Geant4 prototype

Geant4 is one of the principal toolkits used in 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.

VIFit

The HEP community makes large use of many complex data analysis techniques. 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 and enhanced by Vincenzo Innocente, is based on an unbinned maximum likelihood data analysis application. It represents a prototype of the RooFit package (a package inside the ROOT framework), generally used in the HEP for maximum likelihood fits. It applies vectorization and OpenMP-based parallelization.

Results

It is very important to know how well a specific hardware can scale with different workloads in both weak and strong manner. To this end, we apply both black-box and white-box methodologies. With the latter we stress different technologies and subsystems that contribute to the overal performance, such as HyperThreading and Turbo technologies, available instruction sets, vectorization, Non-Uniform Memory Access, cache hierarchy, hardware prefetching etc.

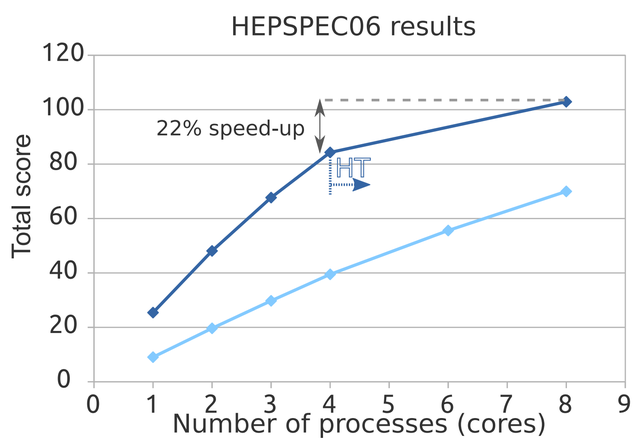

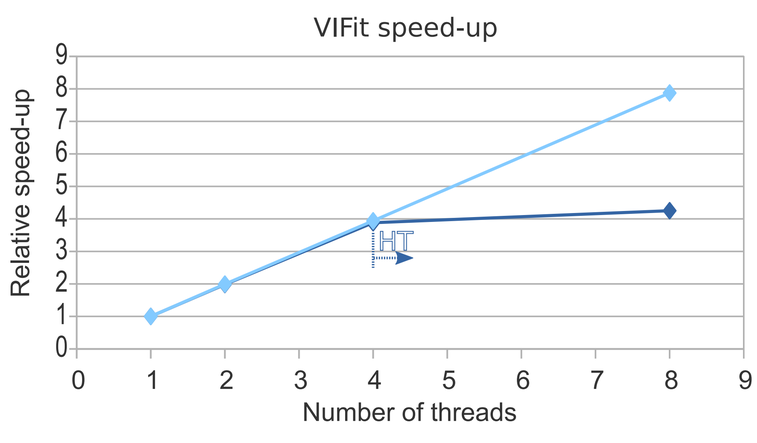

Intel “Silvermont” C2730 — Intel “Haswell” E3-1285 v3

Intel Silvermont shows very good scaling capabilities - there is almost no performance loss when scaling VIFit to the maximum number of cores. These scores are even more promising when taking into account its TDP - 12W. This results in a top-class power efficiency number equal to 3.7 HS06 points per Watt. The same methodology yields 1.8 HS06 per Watt for E3-1285 v3 "Haswell"
However, the tested "Haswell" platform obtained 22% of additional speed-up in HEPSPEC06 with HyperThreading technology enabled. This boost is very handy when only a part of the cores is loaded i.e. when executing non-parallelized regions of code

Last Level Cache statistics and execution times in VIFit on E3-1285v3 "Haswell"

	With prefetching	Without prefetching
Runtime (sec.)	144	290
LLC references (millions)	2,145	4,982
LLC misses (millions)	588	4,335
LLC miss ratio	27%	87%

An important factor affecting this benchmark has been Hardware Prefetching – we have observed speed-ups of 2x simply by enabling BIOS hardware prefetching options. This application is memory bandwidth intensive and as a result prefetching helps reduce last level cache misses. This imbalance indicates that there is certainly a room for improvement in the compiler.

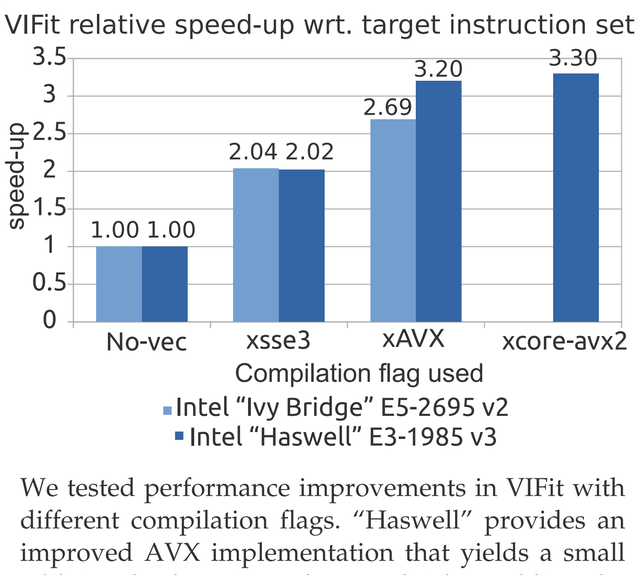

We tested performance improvements in VIFit with different compilation flags. “Haswell” provides an improved AVX implementation that yields a small additional advantage when explicitly enabling the AVX2 flag in the Intel C Compiler v14.0. The small AVX-AVX2 difference comes from the lack of major opportunities for FMA usage in the code.

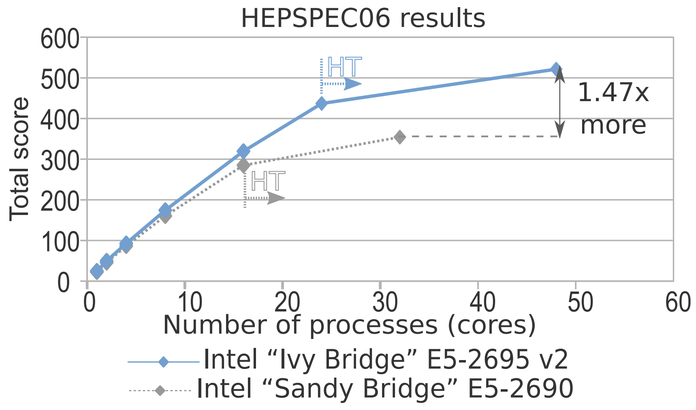

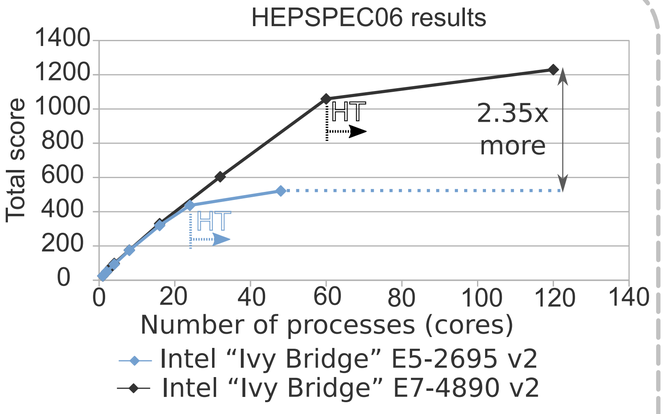

Figure on the left shows scalability with an increasing number of processes of HESPEC06 results of E5-2695 "Ivy Bridge" compared to its predecessor "Sandy Bridge". In the leftmost part of the plot, when the two systems have enough physical cores to accommodate the workload, they scale linearly, with 12% advantage of “Ivy Bridge” HEPSPEC06 performance per core. When we compare results for the maximum number of threads, “Sandy Bridge” has a slight advantage of 2% per core. Nevertheless, when comparing the maximum total performance, “Ivy Bridge” has an advantage of 47%.

Plot on the right shows a comparison of a dual-socket to a quad-socket "Ivy Bridge" platforms. Their scaling in non-HT and HT area respectively is comparable. With E7-4890 we got 2.35x better results, being slightly less than the core-count ratio 120:48 (=2.5).