Evaluation of power efficiency and performance of multi-core platforms using HEP workloads

Paweł Szostek
pawel.szostek@cern.ch

CERN openlab Open Day, June 10th 2015, Geneva, Switzerland

Hardware

Haswell-EP is the latest dual-socket computing platform from Intel. In our tests we compare it with the preceding families: Sandy Bridge-EP and Ivy Bridge-EP. We put particular focus on the power efficiency and Cluser-on-die - new memory management feature.

SKU	Platform	Cores	Frequency	L3 cache	Power	Feature size
E5-2690	Sandy Bridge-EP	8	2.9GHz	20MB	135W	32nm
E5-2695v2	Ivy Bridge-EP	12	2.4GHz	30MB	115W	22nm
E5-2683v3	Haswell-EP	14	2.0GHz	35MB	120W	22nm
E5-2698v3	Haswell-EP	16	2.3GHz	40MB	135W	22nm
E5-2699v3	Haswell-EP	18	2.3GHz	45MB	145W	22nm

Benchmarks

A set of selected benchmarks representing major areas of HEP computing is used for a reference measurement of system and processor performance.
In modern CPUs complexity of benchmarking is constantly growing, as more and more factors contribute to overall

HEP-SPEC 06

HEP-SPEC06 has been chosen by a work group affiliated by HEPiX as a standard HEP benchmark. It is a set of real-life applications available on a commercial basis. Since it has proven to **correlate significantly with HEP applications**, the community decided to adopt it as a reference benchmarks. It is widely use to define pledges at data centres worlwide. Since t, in order to leverage all cores available on a CPU, we run a separate process on each core and sum up the obtained scores.

ParFullCMS / Geant 4

Geant4 is the major toolkit used in the simulation of the detectors in Large Hadron Collider (LHC). It is used to simulate the passage of particles through matter. It is a principal representative of the workloads run in the Worlwide LHC Computing Grid and **constitutes a considerble portion of its CPU time.**
ParFullCMS is a benchmark built on the top of the Geant4 library. It is a stand-alone application implementing the full geometry of the CMS detector with full physics simulation, but a simplified setup for data collection, and operating a uniform magnetic-field.

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 implements partially the Maximum Likelihood Fit algorithm. **It applies vectorization and OpenMP-based parallelization.**

Results

It is very important to know how well a specific hardware platform can scale with different workloads. In addition, we always try to understand how new features and technologies influence total system performance.

In our tests we stress different platform subsystems that contribute to the overal system performance, such as Hyper-Threading and Turbo Boost technologies, available instruction sets, vectorization, Non-Uniform Memory Access, cache hierarchy etc.

When testing Haswell-EP we put particular focus on Cluster-on-die, which is a feature allowing to split the shared memory in one socket into two independent nodes.

Obtained results prove significant improvements in power efficiency, vector efficiency and single core performance.

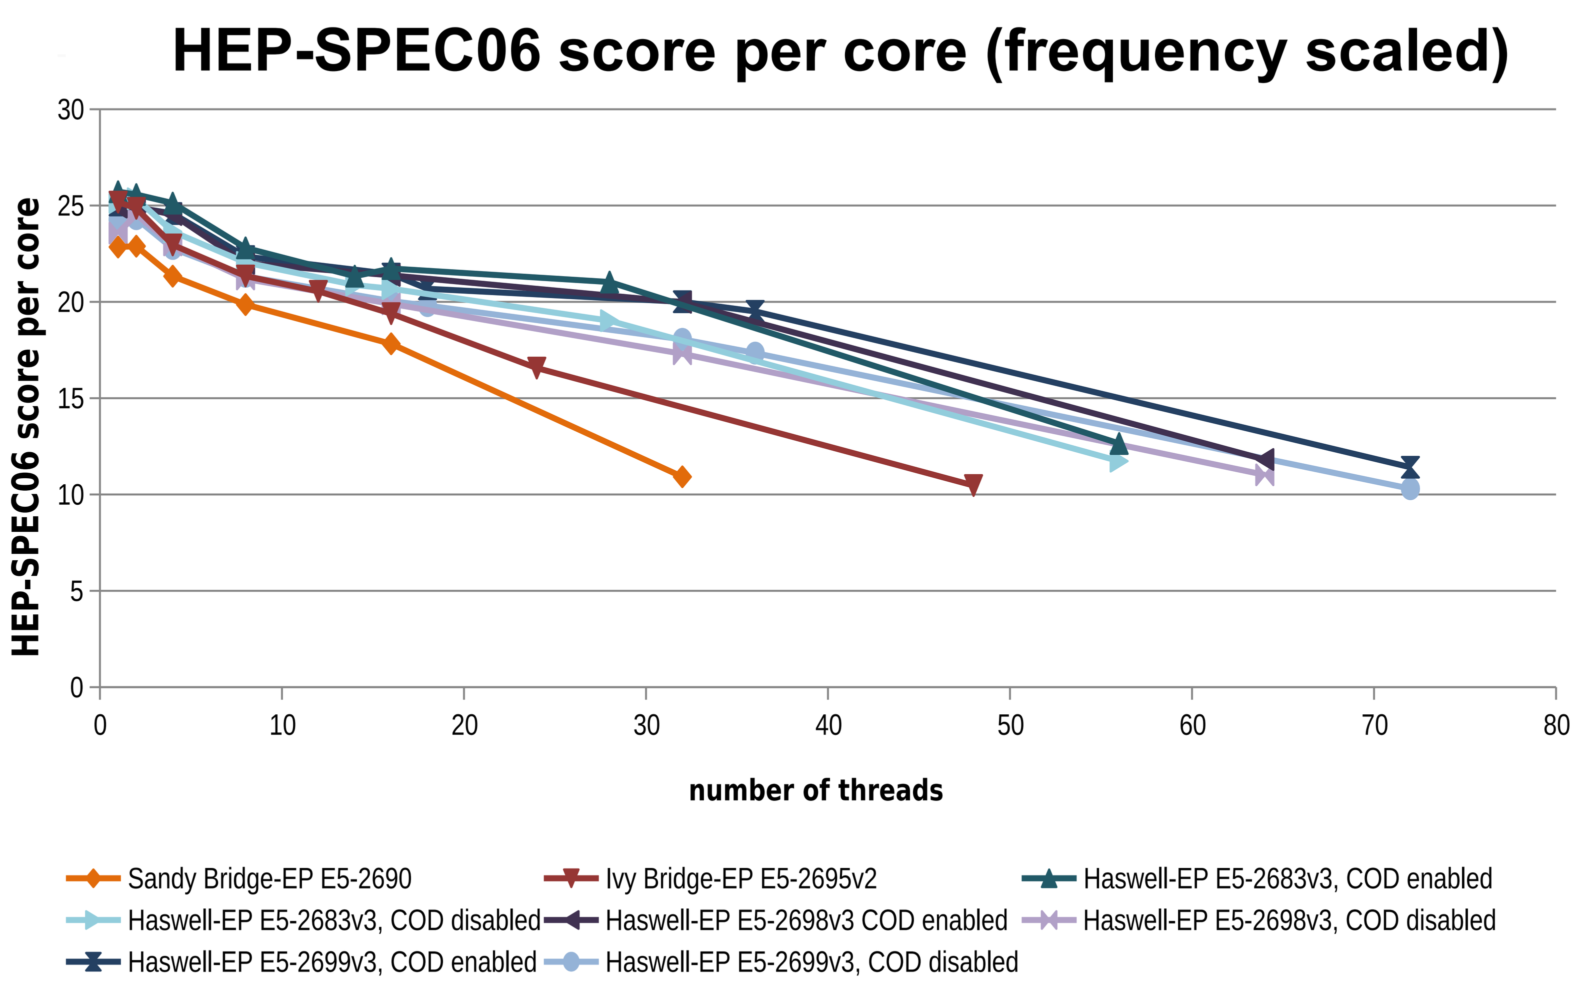

Figure above presents HEP-SPEC06 results scaled down to a common frequency of 2.7GHz and divided by the number of workload instances run at a time. These results can be interpreted as the input of every loaded core to the total score obtained with the given number of parallel instances. **At full load Haswell cores provide from 10% to 20% better performance than Sandy Bridge-EP.**

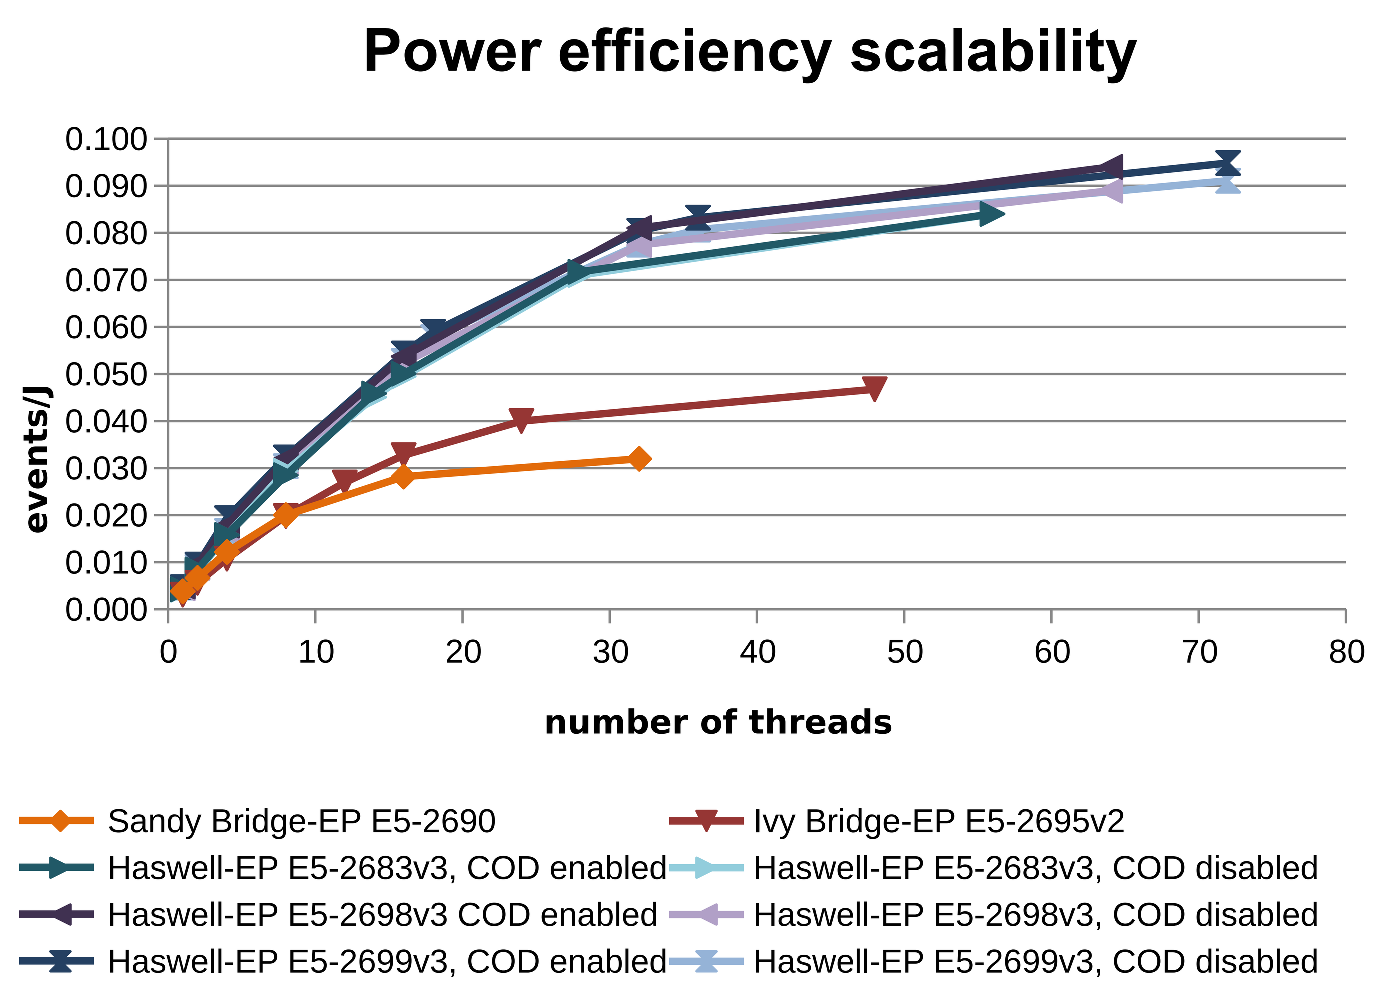

This figure shows the power efficiency scalability for the tested platforms, i.e. how many events can be processed per joule with a given number of threads running in parallel. This measurement proves **significant improvement in power efficiency** in Haswell-EP platform. While Ivy Bridge-EP was able to process 46% more events than Sandy Bridge-EP per the same amount of energy, **Haswell-EP is capable of processing up to 102% more events than Ivy Bridge-EP and 196% more than Sandy Bridge-EP per unit of energy.**

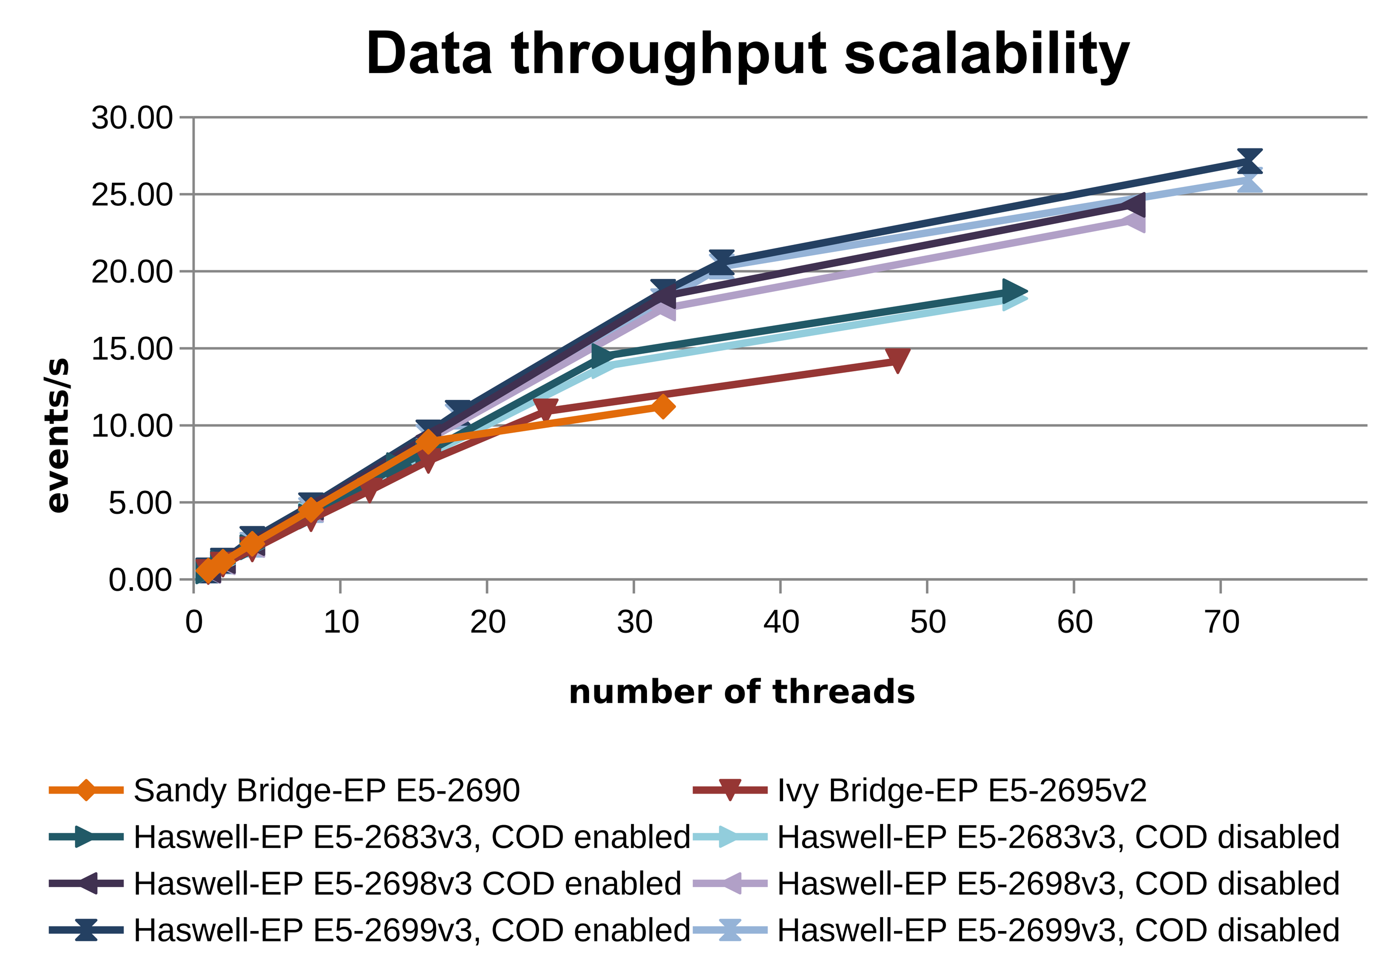

The plot shows throughput scalability of the tested platforms, i.e. how many events can be processed in a unit of time if a certain number of threads is used. **The 18 cores Haswell-EP part (E5-2699v3) can process 92% and 142% more events per second than Sandy Bridge and Ivy Bridge, while it has 50% and 100% more cores respectively.** Cluster-on-die provides from 2% to 4% of additional performance.

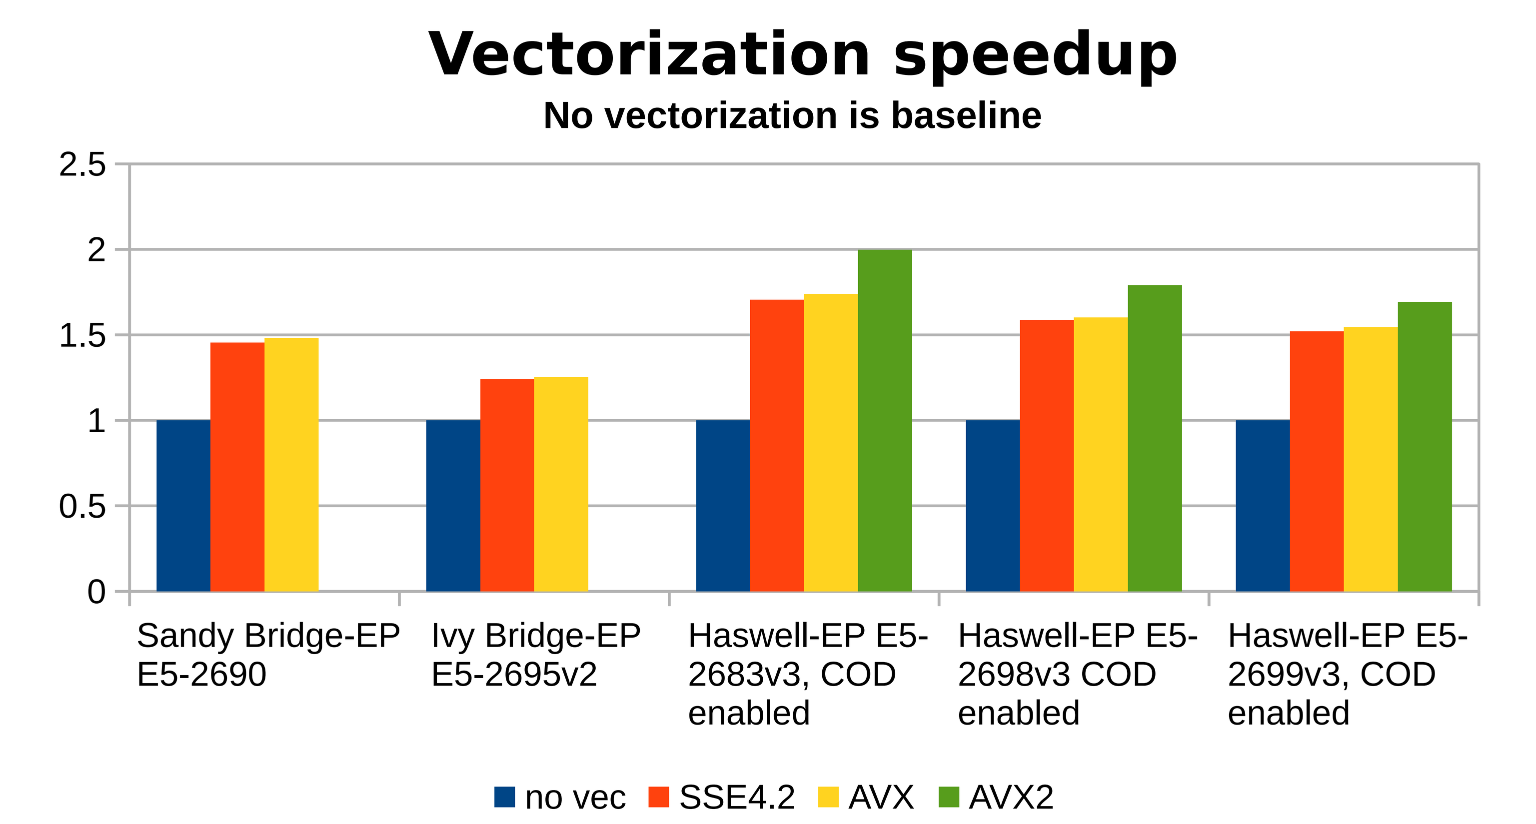

This plot shows vectorization speedup obtained with various instruction sets for VIFit. For each platform non vectorized variant is the baseline. Haswell-EP servers provide up to 2x speedup with AVX2 and 1.7x with AVX instruction sets. **Haswell-EP parts provide significantly better vectorization efficiency than their predecessors.**