


Sandy Bridge EP evaluation



**CERN openlab Minor review meeting
Feb 28 2012
Sverre Jarpe, CERN openlab**

Based on technical contributions by A.Lazzaro, J.Leduc, A.Nowak
Slide master also from A.Nowak


Overview

1. Benchmarking complexity
2. Intel's tick-tock model
3. Sandy Bridge details
4. Architecture/Micro-architecture
5. SNB CPU models
6. SNB Results
 - a. HEPSPEC
 - b. HEPSPEC/W
 - c. MT Geant4
 - d. Mlfit
7. Conclusion

Benchmarking: A complex affair

- **At least the following elements need to be controlled:**
 - **Hardware:**
 - Processor generation
 - Socket count
 - Core count
 - CPU frequency
 - Turbo boost
 - SMT
 - Cache sizes
 - Memory size and type
 - Power configuration
 - **Software:**
 - Operating System version
 - Compiler version and flags

Intel's tick-tock model



SNB in some detail

- **Advanced Vector eXtensions (AVX)**
 - 256 bit registers which can hold 4 doubles/8 floats
 - AVX instruction set
- **More execution units (2 * LD, for instance)**
- **Enhanced Hyper-threading and Turbo-boost technology**
- **Larger on-die L3 cache**
- **Integrated PCI Express 3.0 I/O**

Architecture vs microarchitecture

- **Architecture**

- New register format
 - e.g. 256-bit AVX registers
- New (ternary) instructions
 - e.g. vdivpd ymm1, ymm2, ymm3

The way it works

- **Microarchitecture**

- Lots of design decisions for a given processor
 - Number of execution units (and their width)
 - Data paths (and width)
 - Cache sizes
 - Etc.

The speed at which it works

CPU models

- Long list of models to choose from.
 - Some variants:

Number	Core count	Frequency (GHz)	TDP (W)
E5-2630L	6/12	2.3	60
E5-2650L	8/16	1.8	70
E5-2680	8/16	2.7	130
E5-2690	8/16	2.9	135


SNB results

- **System tested:**
 - Beta-level white box; Dual-socket server.
 - E5-2680 @ 2.7 GHz, 8 cores, 130W TDP
 - 32 GB memory (1333 MHz)
 - C1 stepping
- **Benchmarks used:**
 - HEPspec
 - HEPspec/W
 - MT-Geant4
 - MLfit

HEPSPEC

- **Throughput test from SPEC 2006**

- All the C++ jobs (INT as well as FP); As many copies as cores
- SLC 5.7/gcc 4.1.2/64-bit-mode/Turbo off/SMT on
- Compared to 6-core Westmere-EP X5670 (@2.93 GHz)
 - Frequency-scaled



Using only the “real” cores:

Speed-up per core:	1.2x
Core count:	1.33x
<hr/>	
Total:	1.6x

SMT gain (for both): 1.23x

Energy efficiency

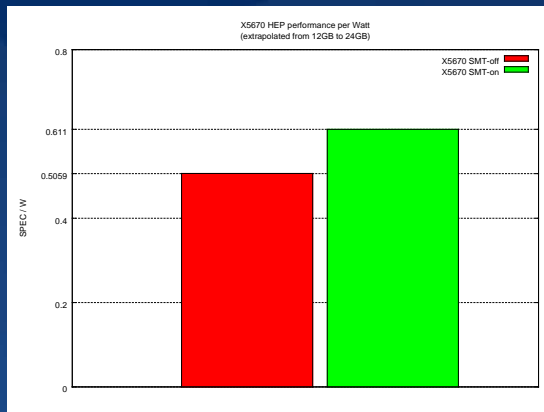
- **For CERN and most W-LCG sites, energy efficiency is paramount**
 - Our centres have (more or less) a fixed amount of electric energy
 - Ideally, we would like to double the throughput/watt from generation to generation
 - This was relatively easy when core count increased geometrically:
 - $1 \rightarrow 2 \rightarrow 4$
 - Recently, however, it has been increasing arithmetically:
 - $4 \text{ (NHM)} \rightarrow 6 \text{ (WSM: 1.5x)} \rightarrow 8 \text{ (SNB: 1.33x)}$

HEPSPEC/Watt

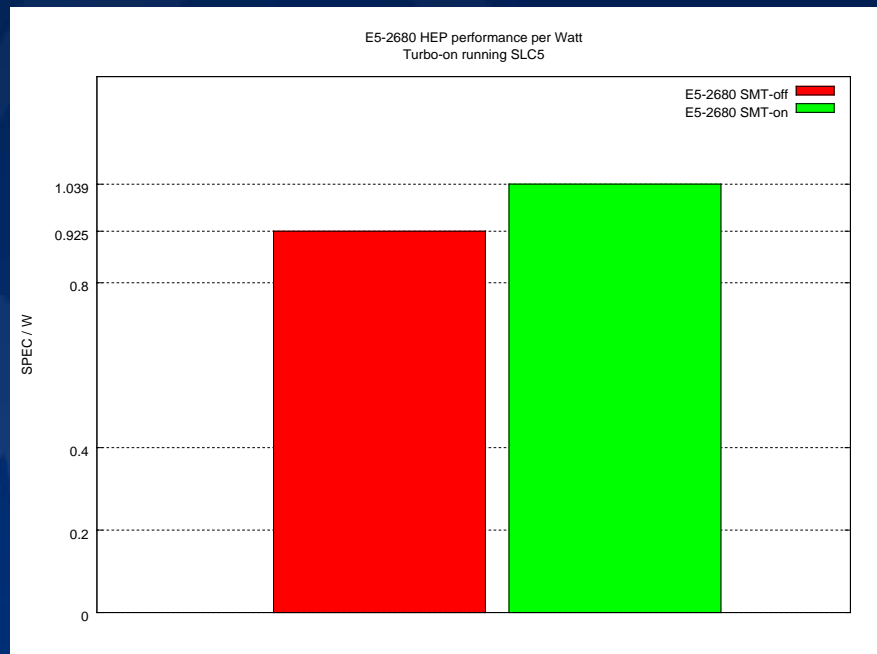
- **Great news: Bigger jump than foreseen in energy efficiency!**
 - Now reaching 1 HEPspec/W which is **1.7x** compared to WSM-X5670
 - SNB options: SLC 5.7, 64-bit mode, SMT on, Turbo on
 - WSM options: SLC 5.4

SNB

Bigger is better!



WSM

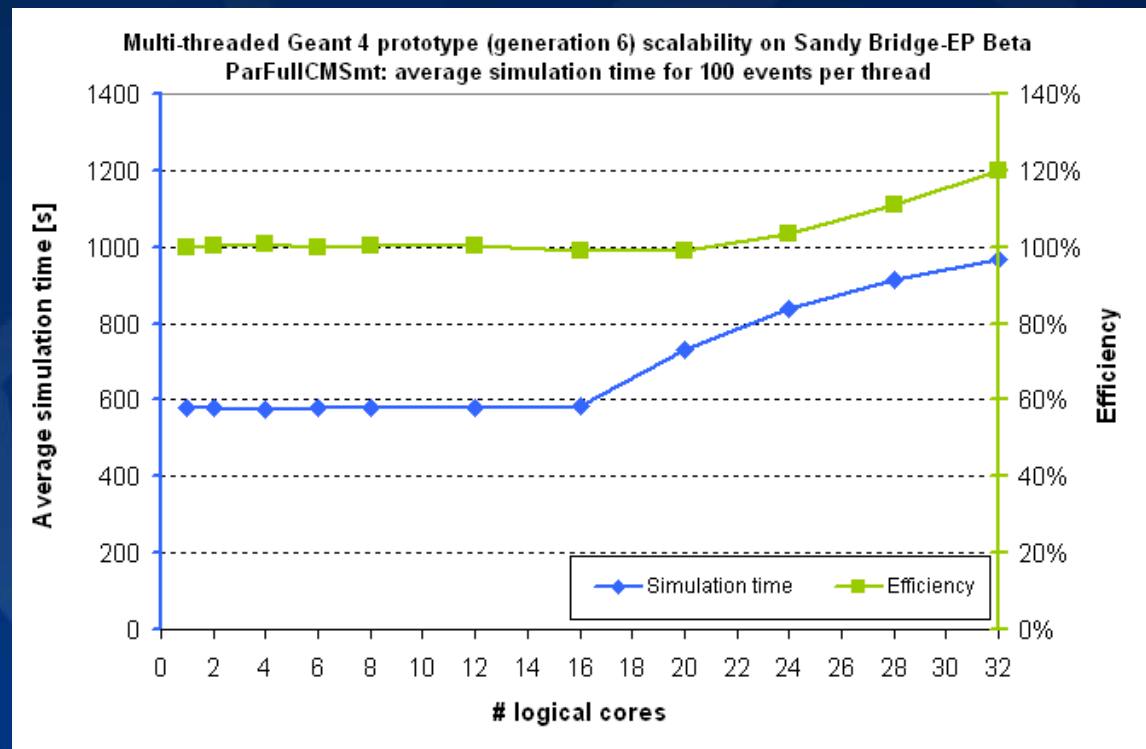


STOP PRESS: With SLC 6 (gcc 4.4.6) we further lower the power consumption by 5% and increase the HEPspec results by 3%: 1.083x in total !

MT Geant4

SLC 5.7, gcc 4.3.3,
pinning of threads

- Our favourite benchmark for testing weak scaling:
 - Speed-up compared to Westmere (L5640@2.26GHz):
 - Both servers with Turbo-off, SMT-on (WSM frequency-adjusted): 1.46x
 - SMT increase: 1.25x



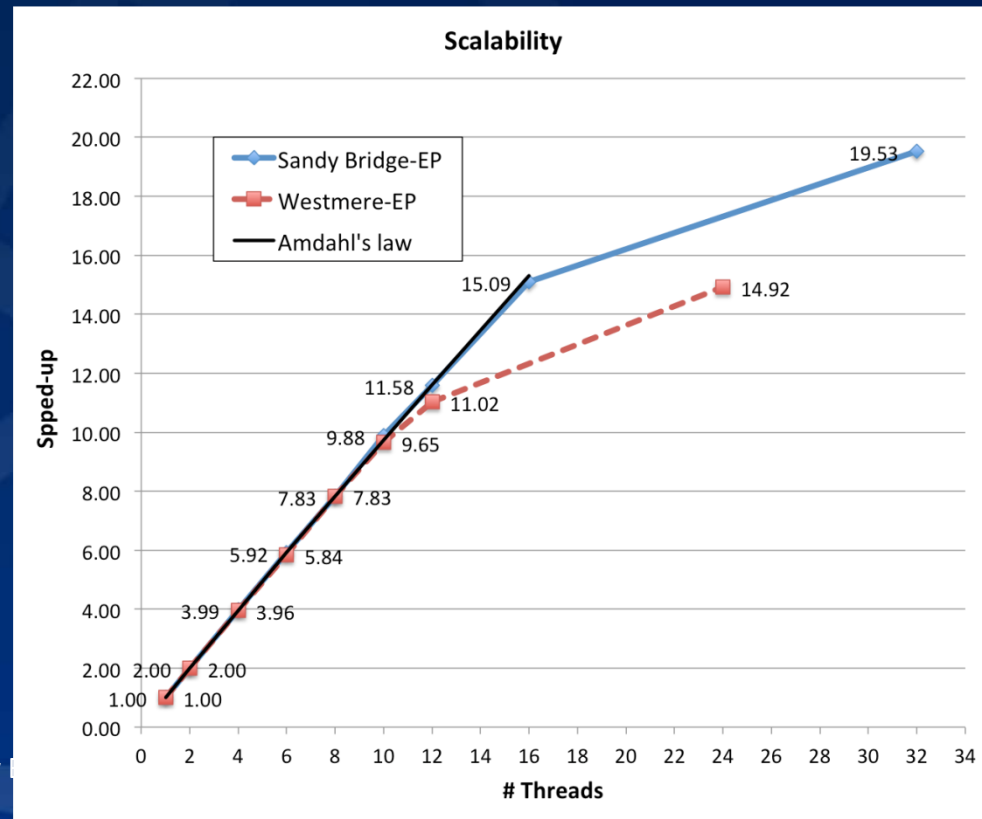
MLFit

SLC 6.2, icc 12.1.0,
pinning of threads

- Our favourite benchmark for testing strong scaling:
 - Single core (Turbo off, using SSE): 1.19x
 - Single core, moving to AVX: 1.12x
 - All the “real” cores w/SSE: $(1.33 * 1.19)$ 1.59x
 - All the “real” cores & AVX: $(1.59 * 1.12)$ 1.78x

1.33x

SNB SMT speed-up:
1.29x



Conclusion

- **Sandy Bridge EP confirms Intel's desire to improve both absolute performance and performance per watt**
- **CERN and W-LCG will appreciate both**
 - In particular, the HEPSPEC/W value
- **The full openlab evaluation report will be published at launch time (as usual)**
 - The Westmere-EP (X5670) report is available since April 2010