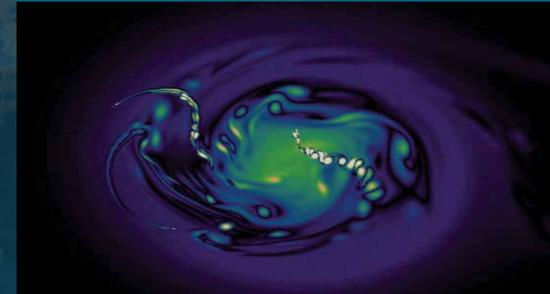


Evaluating the Intel Skylake Xeon Processor for HPC Workloads



PRESENTED BY

Clayton M. Hughes, Sandia National Laboratories

Simon D. Hammond, Sandia National Laboratories

Courtenay T. Vaughan, Sandia National Laboratories

HPBench 2018
July 16 – 20, 2018



Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

SAND2018-7482 C



Scientific computing relies on strong server-class processors

- Wide availability of GPUs, many-core processors, and special-purpose accelerators and functional units
- Majority of calculations still take place on commodity server-class processors
- Many applications still have large regions of serial code
- Necessitates the need for powerful cores

```
lea ecx,dword ptr [esp+0x4]
and esp,0xf0
push dword ptr [ecx-0x4]
push ebp
mov ebp,esp
push edi
push ebx
```

Two classes of computing platforms for U.S. Department of Energy

- Advanced Technology Systems (ATS)
- Capacity Technology Systems (CTS)





Evaluate the impact of the significant hardware changes on Intel's leading Skylake-based server platform using microbenchmarks

- Memory bandwidth
- Cache bandwidth
- Floating point performance

Project the impact that these changes will have on real applications of interest to the scientific community using a selection of mini-applications

- Memory bandwidth
- Indirect memory accesses from cache and main memory
- Throughput
- Vectorization and SMT



Skylake Microarchitecture



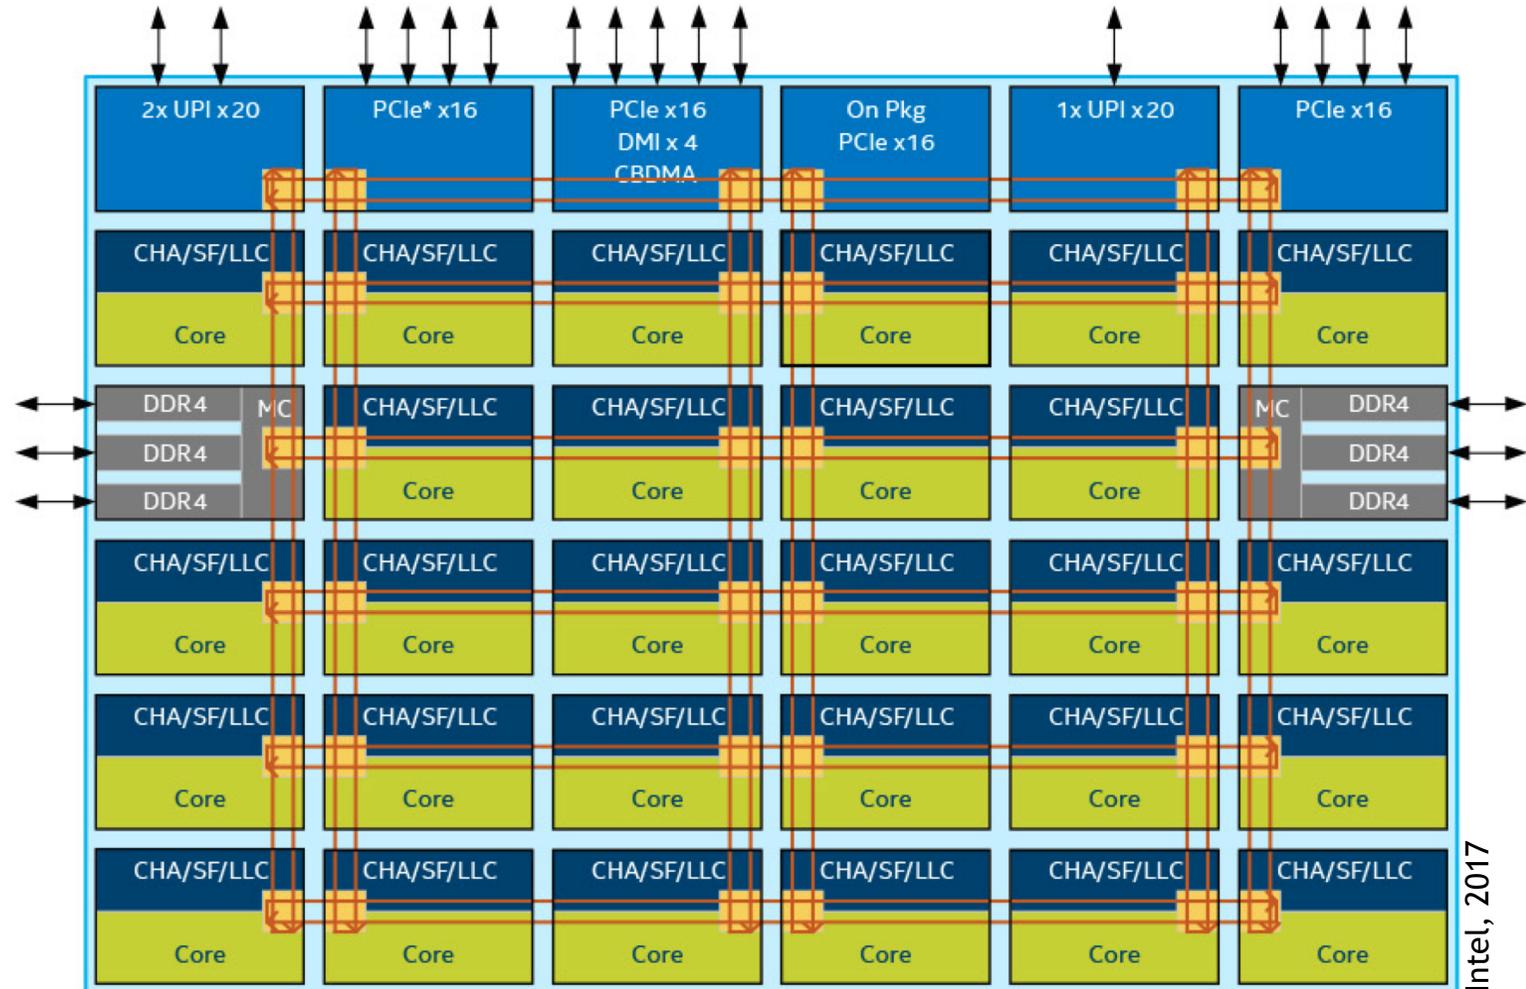
6th generation Core microarchitecture

14nm+ process technology

Mesh-based interconnect

6 memory channels

- Should help some memory bandwidth-bound codes





Core and cache subsystem redesigned to support greater locality and reduce L3 contention

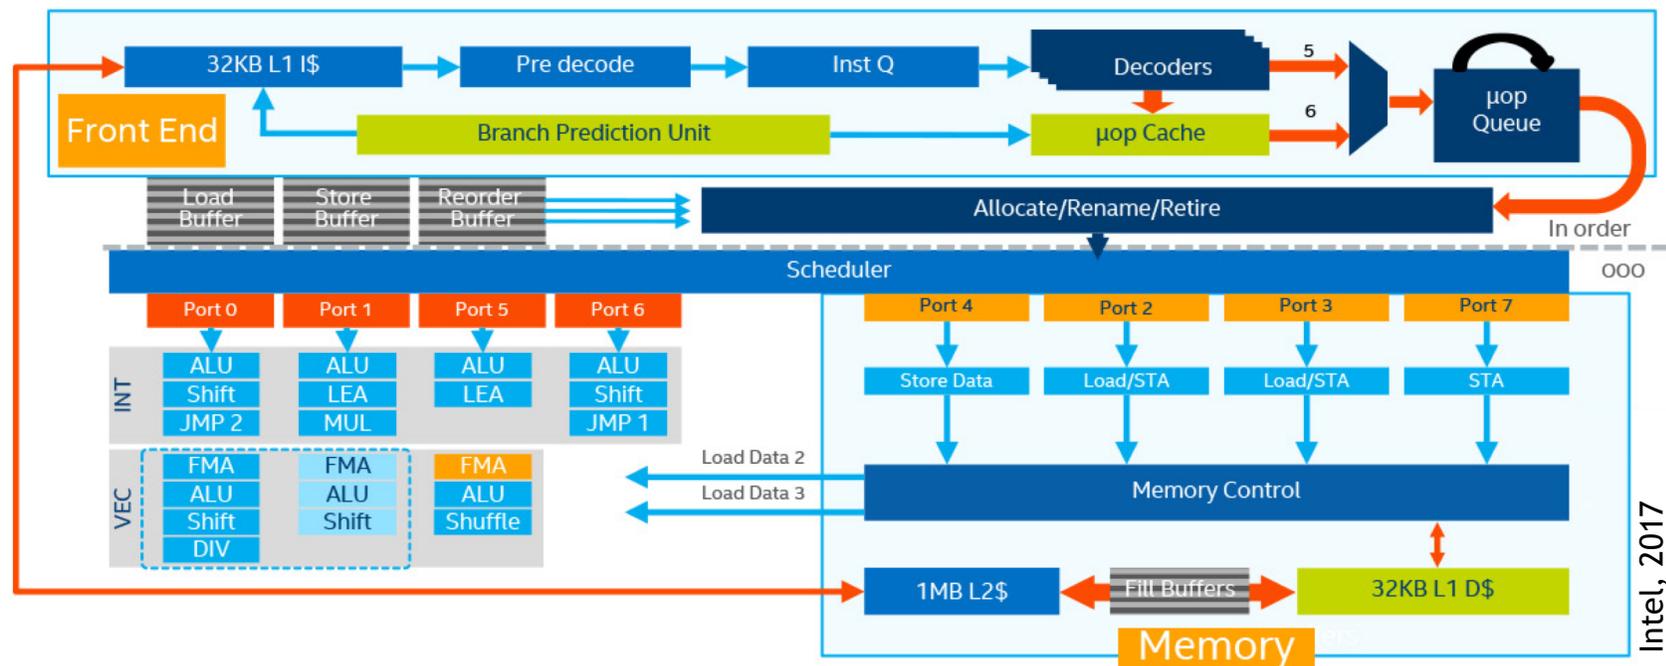
Double L1 bandwidth

Native support for large pages

	Haswell	Skylake
Out-of-Order Window	192	224
In-Flight Mem. Ops, LD+ST	72+42	72+56
Scheduler Entries	60	97
Allocation Queue	56	64
Registers, INT+FP	168+168	180+168
L1 BW, LD+ST (B/cyc)	64+32	128+64
L2 Unified TLB	2k+2M: 1024	2k+2M: 1536 1G: 16

AVX-512

- Masks
- Three vector ports (0,1,5)
 - 1x 512b
 - 2x 256b that can be merged



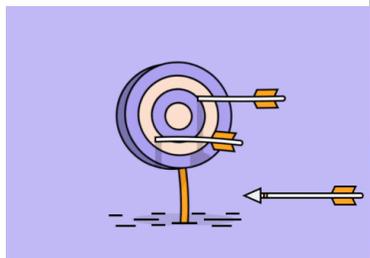
Results averaged over 10 runs with random nodes chosen for each trial

Two Intel-based platforms

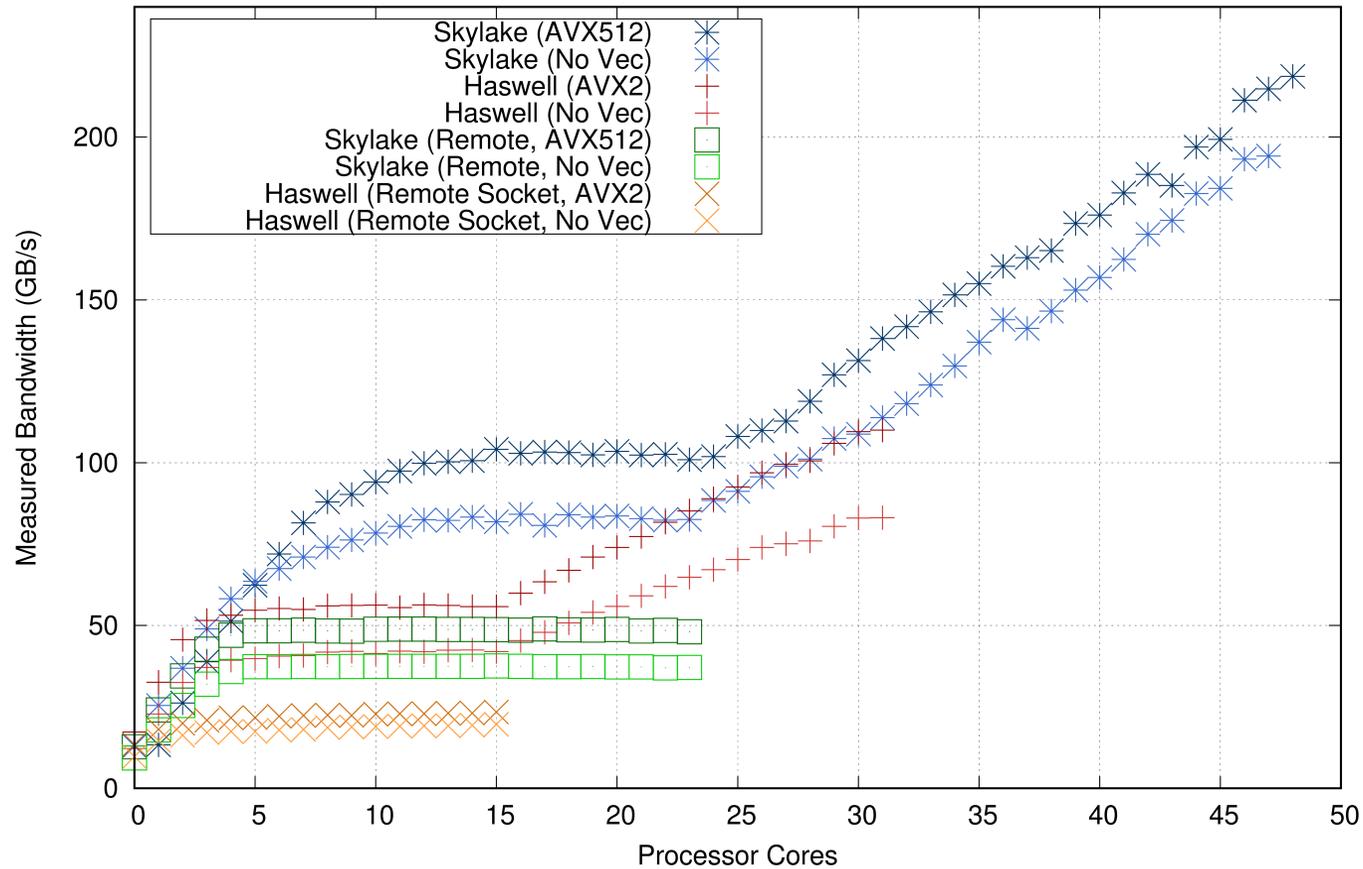
- Shepard (Haswell)
 - Dual-socket Xeon E5-2697v3
 - 2.3GHz
 - 16 cores with dual SMT
 - 32KiB L1/256KiB L2/40MiB distributed L3
 - 128GB 2133MT/s DDR4
- Blake (Skylake)
 - Dual-socket Xeon Platinum 8160
 - 2.1GHz
 - 24 cores with dual SMT
 - 32KiB L1/1MiB L2/33MiB distributed L3
 - 192GB 2666MT/s DDR4

ICC 18.1.0

- GCC 4.9.3 compatibility
- MKL 18.1
- OpenMPI 2.1.2



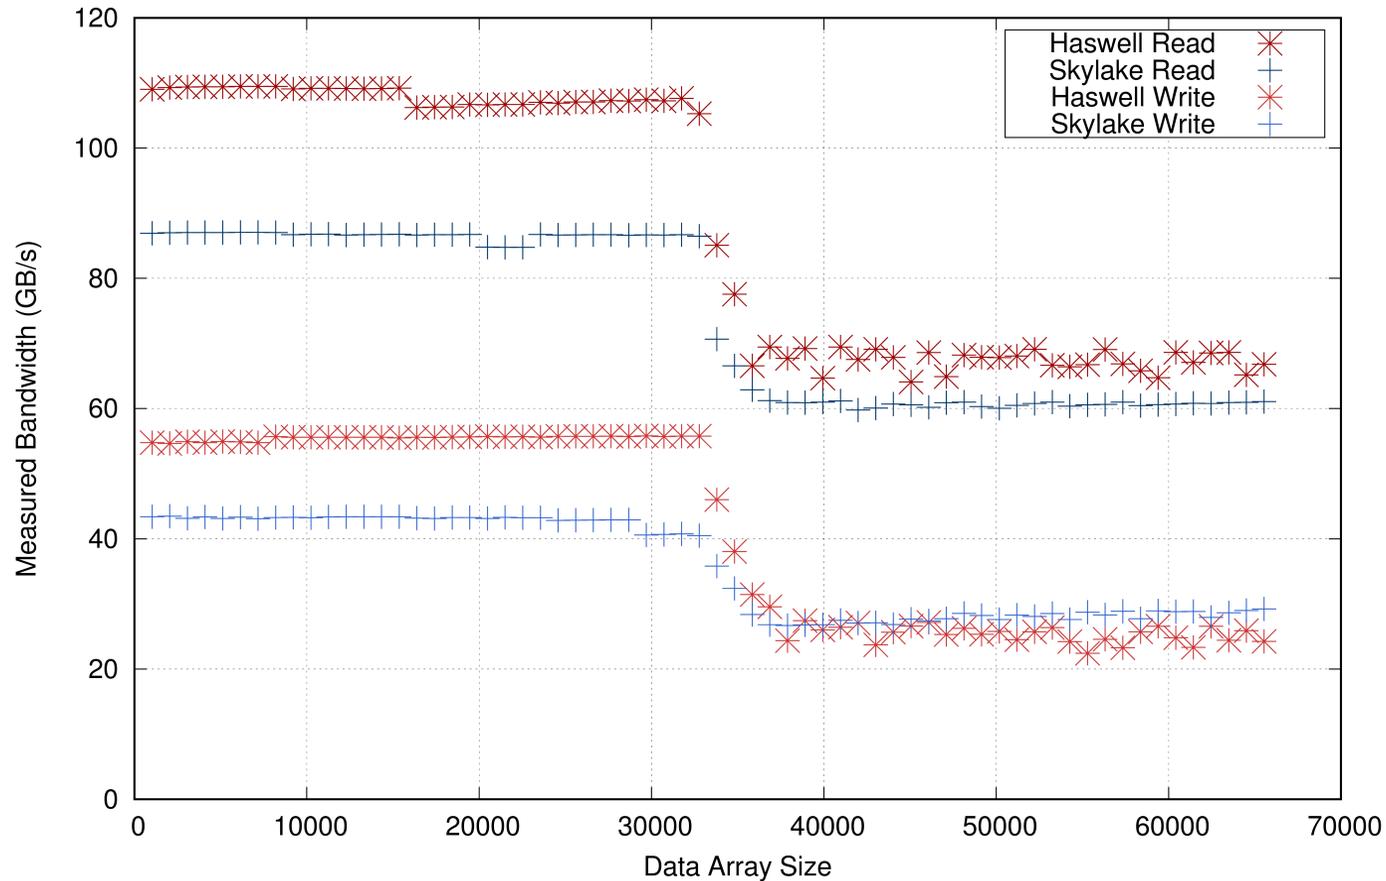
Results – Memory Bandwidth



Vectorization improves memory bandwidth regardless of architecture

UPI links on SKX improve remote socket bandwidth (2.2x higher)

Results – Cache Bandwidth (Small Array)

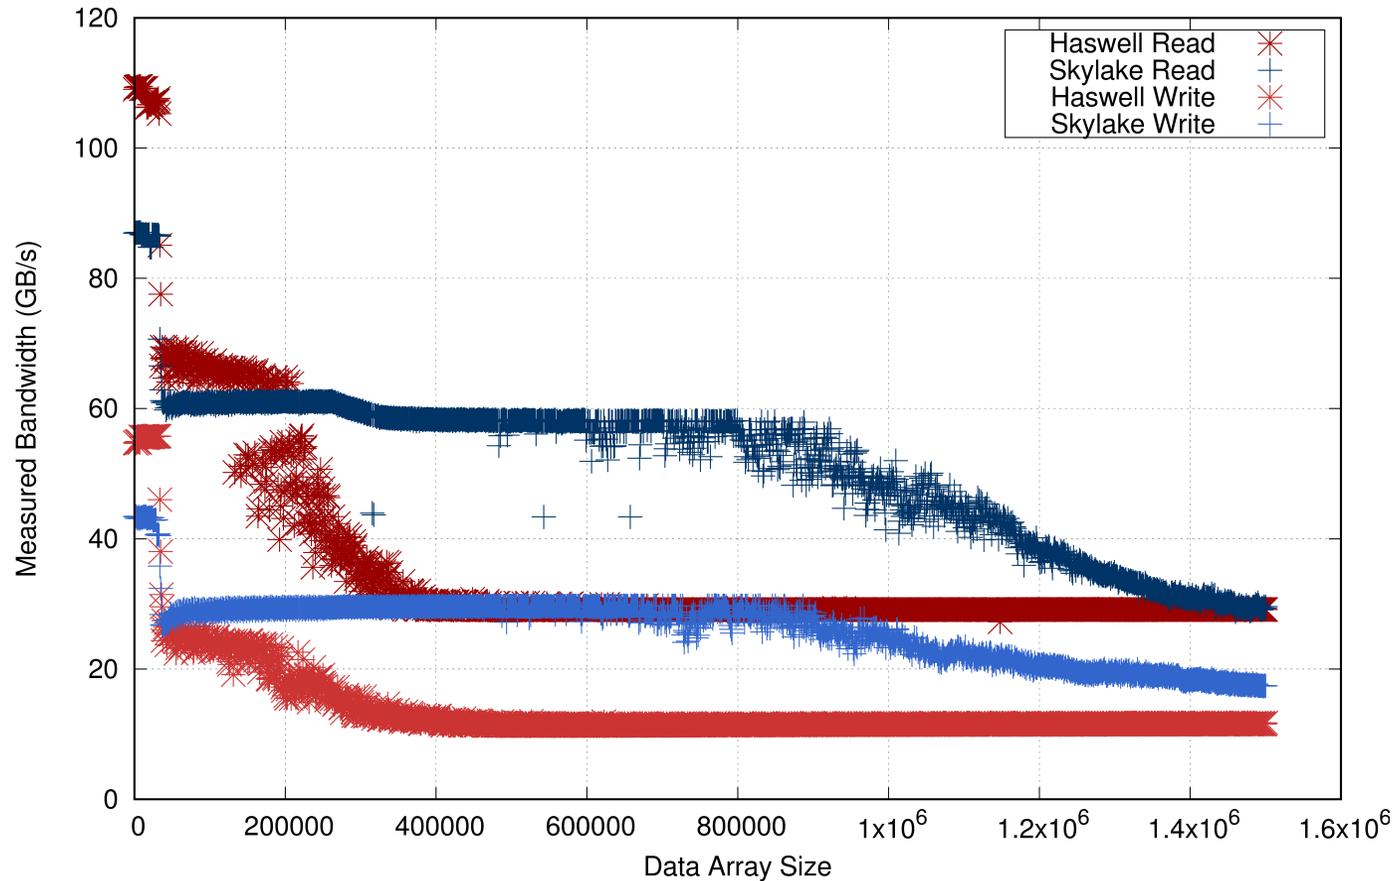


L1 bandwidth is lower on SKX

- Can be partially attributed to lower clock frequency

L2 bandwidth is lower for reads and slightly higher for writes

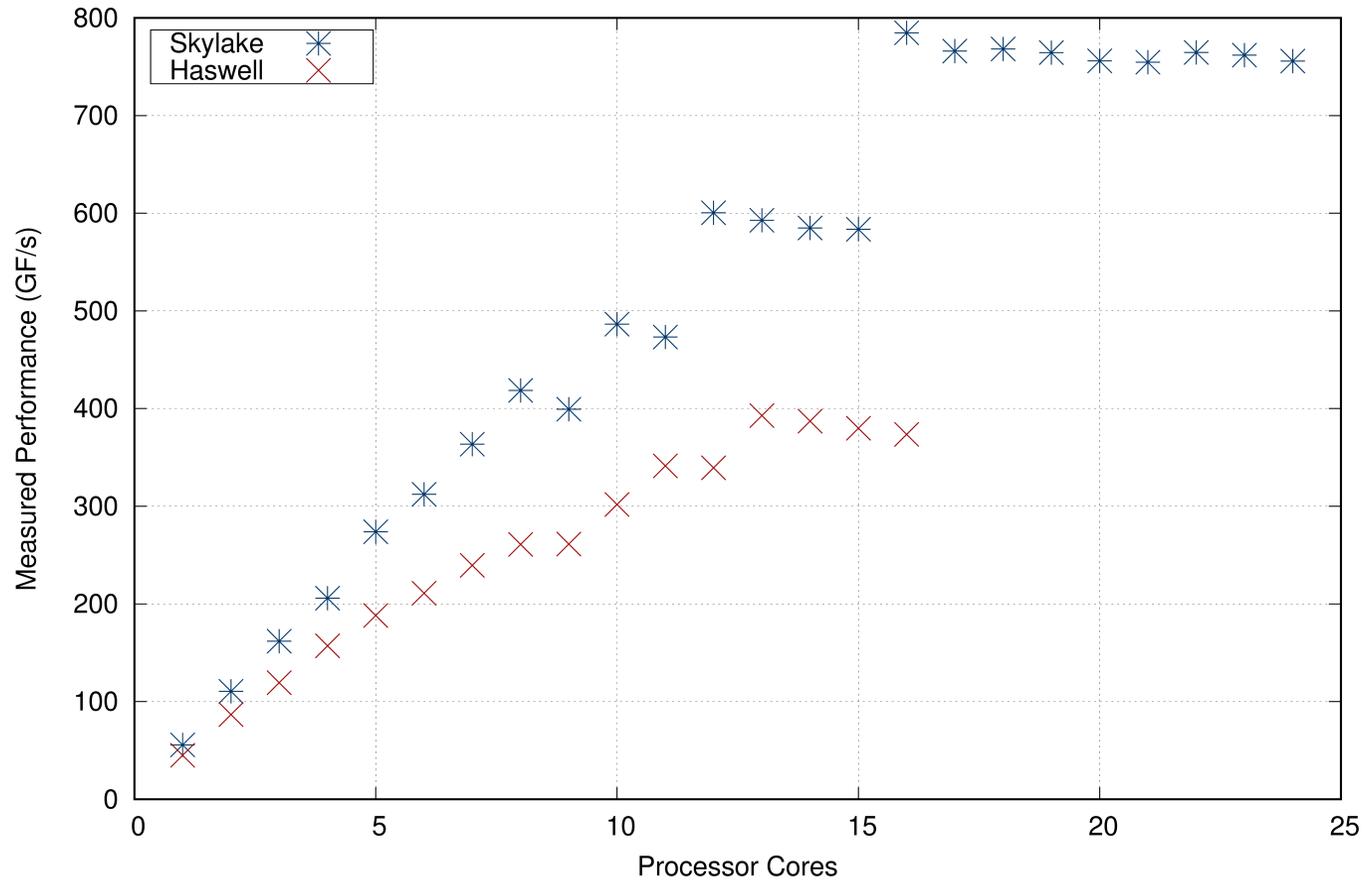
Results – Cache Bandwidth (Large Array)



Moving to a macro view, the larger L2 size on SKX clearly improves the bandwidth for a much greater range of array sizes

Some variation in Haswell despite setting core affinity

Results – Floating Point Arithmetic



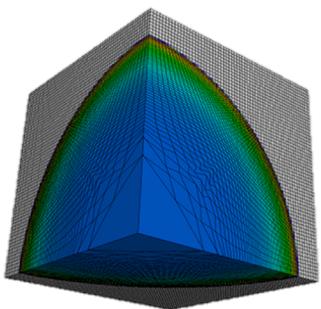
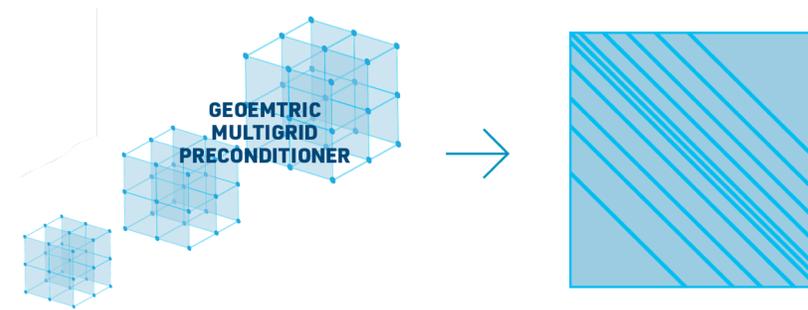
Doubling of vector size does result in doubling of FLOP rate

- 784GF/s at 16 SKX cores vs. 392FG/s at 13 HSW cores

Additional cores do not affect the FLOP rate, most likely due to thermal throttling

Mini-Applications and Benchmarks

High Performance Conjugate Gradient (HPCG)

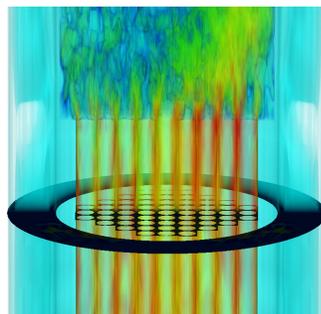


LULESH

- Hydrodynamics over unstructured meshes

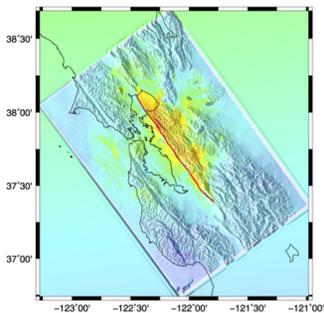
XSBench

- Monte Carlo transport



SW4Lite

- 3D modeling of seismic activity





Kernel	Skylake (AVX512)	Skylake (NoVec)	Haswell (AVX2)	Haswell (NoVec)
DDOT	20.05	30.50	9.87	11.41
WAXBY	16.70	16.88	9.53	9.35
SpMV	18.56	17.95	10.22	10.20
Multi-Grid	18.29	17.94	10.01	9.89
Solve (Total)	18.33	18.04	10.03	9.95

HPCG kernels are considered to be memory bandwidth-bound

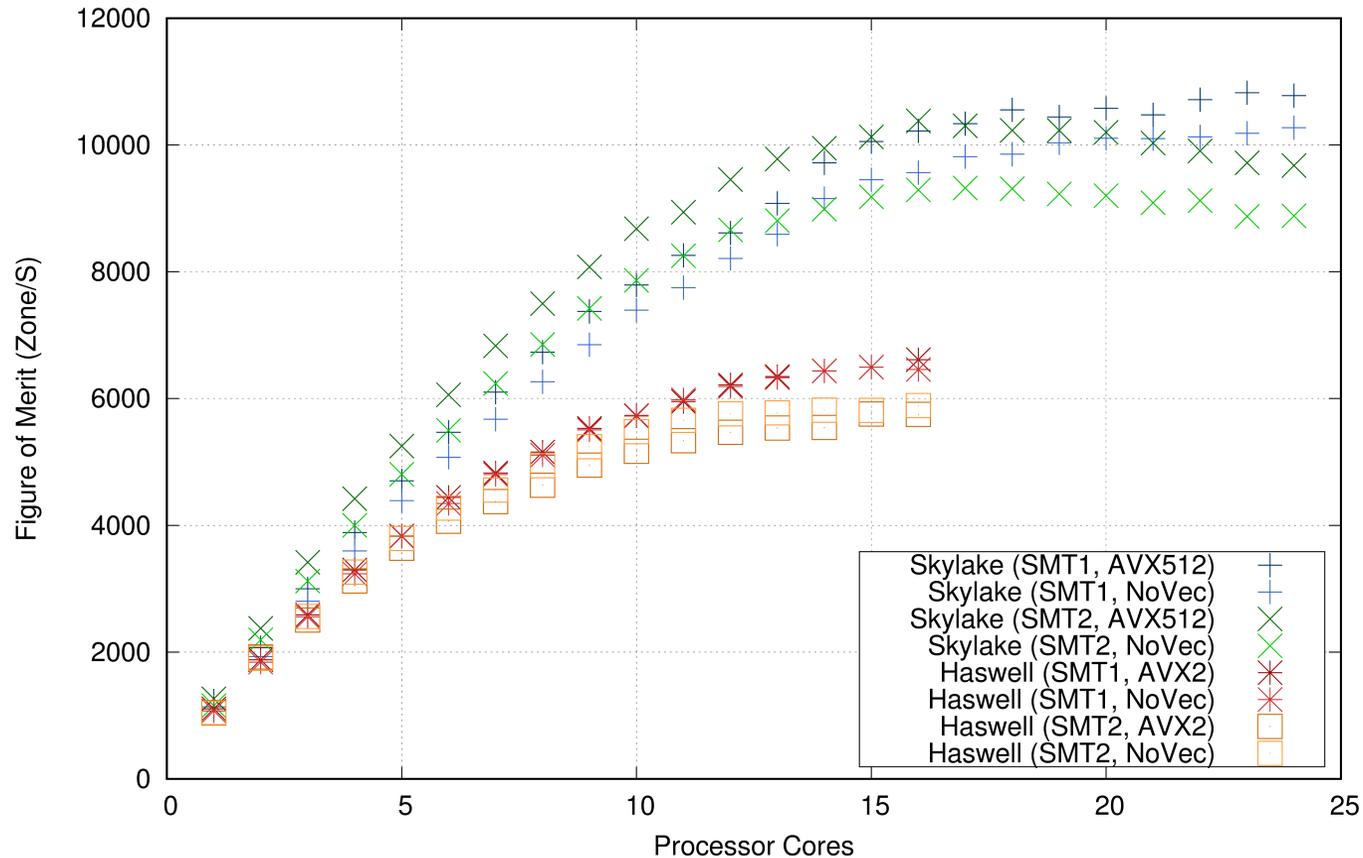
- Would expect ~50% performance improvement due to additional memory channels but see ~80%

Additional performance gains can be attributed to general processor enhancements

- Larger OoO window, scheduler, and additional entries in the store queue

Vectorization makes almost no difference in default implementation

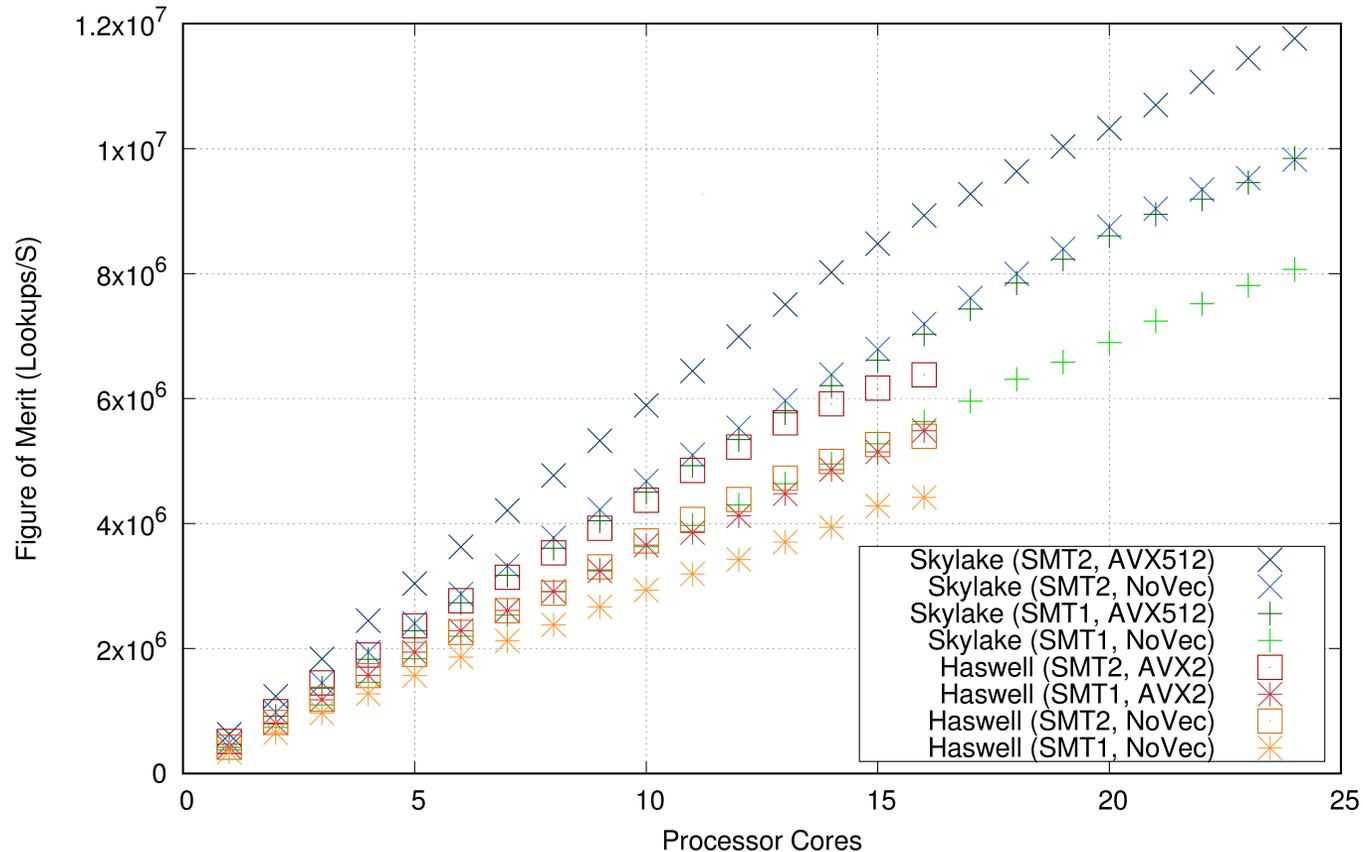
Results – LULESH



Only 4-8% difference in vectorized/non-vectorized code

SMT on SKX good from 1-16 threads (1-12% improvement)

SMT on HSW always underperforms

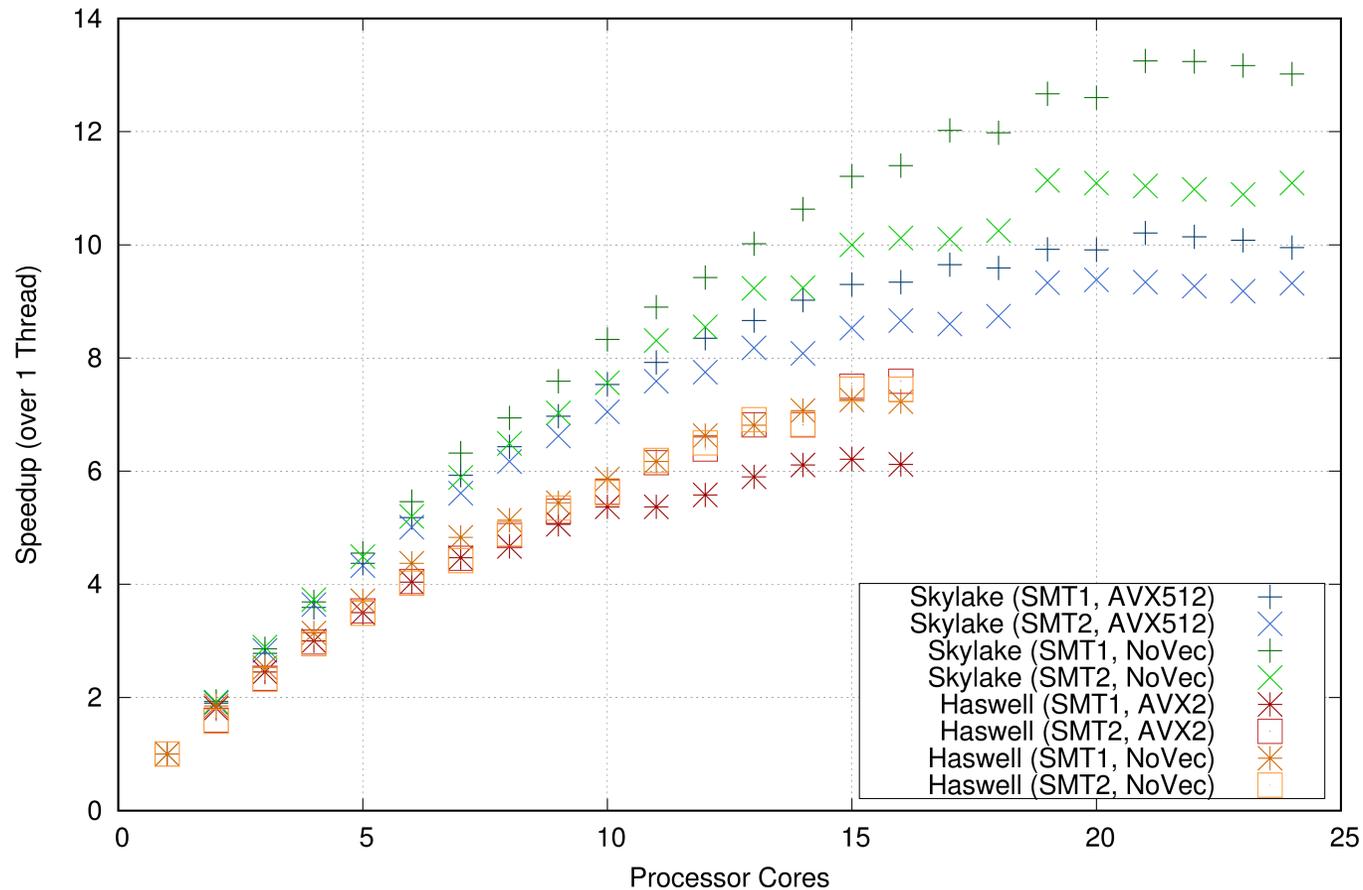


SMT improves performances on both SKX (~20%) and HSW (~16%)

- Able to hide memory latencies with additional lookups

Vectorization improves performance on SKX by 22% and HSW by 24%

- Scatter/gather instructions help here where accesses aren't necessarily in the cache



Vectorization improves performance significantly on both systems

- 39-45% fast on SKX and 45-47% faster on HSW

SMT hinders performance on both systems

- 11% slower on SKX and 30% slower on HSW

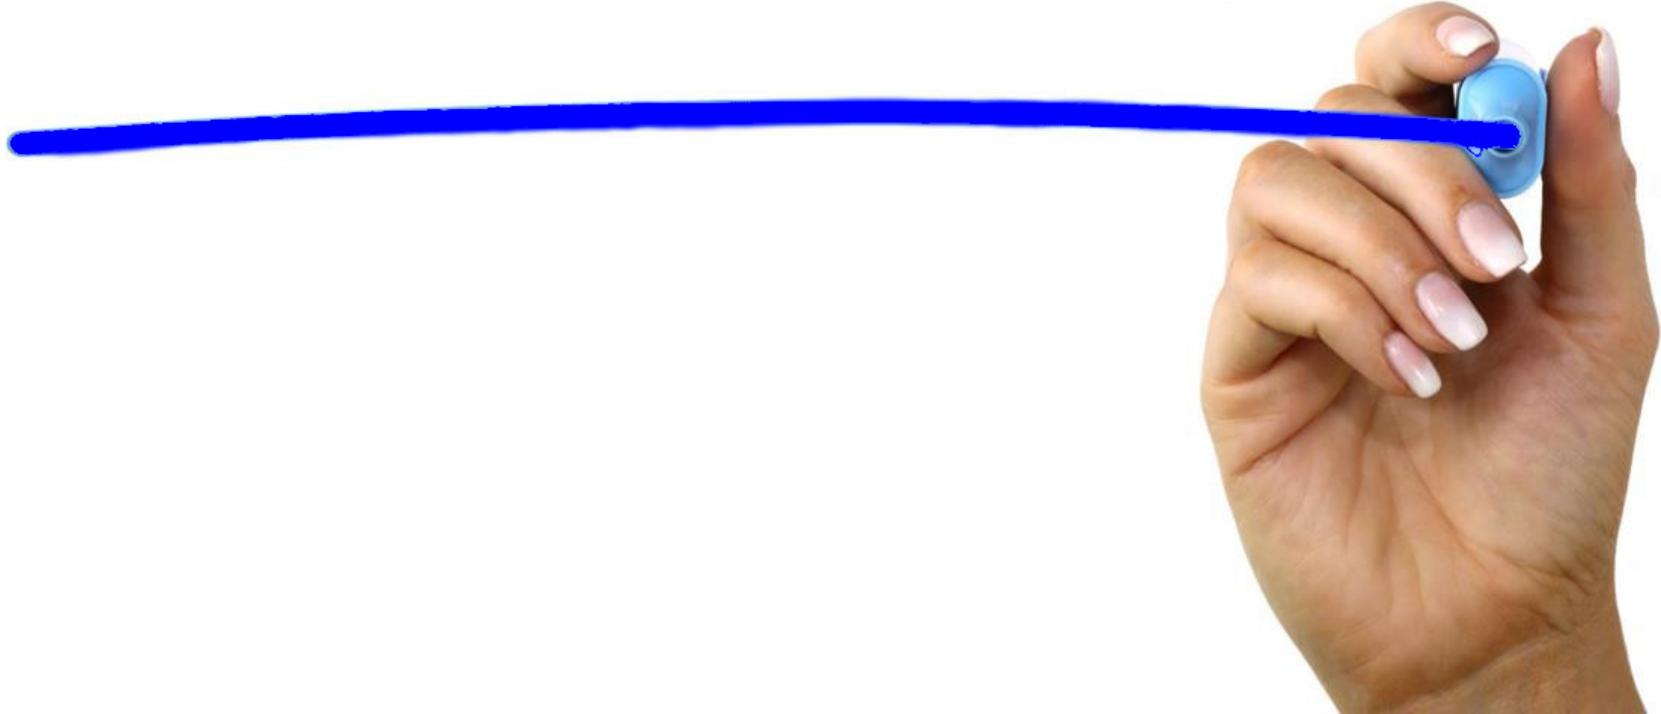


Skylake's redesigned core architecture provides a host of improvements

- Redesigned cache and 6 memory channels
 - STREAM shows nearly 2x improvement over previous generation → 223.8GB/s vs 112.6GB/s
 - HPCG shows a 0.8x improvement
 - LULESH shows a gain of 1.6x
- Wider vector units
 - DGEMM had a 2x performance improvement over previous generation
 - SW4Lite isn't quite as impressive but still shows a 0.83x improvement
- 2D mesh
 - XSBench shows a 1.85x improvement on Skylake over Haswell

The changes made in the core while minimizing power increases are impressive and HPC workloads should benefit greatly from them

QUESTIONS





Slide Left Blank