What's new@intel

P. Thierry

Principal Engineer, Intel Corp

philippe.thierry@intel.com

CPU trend

Memory update

Software

Characterization

… in 30 mn

RTC Workshop, Paris 2016

10 000 feet view

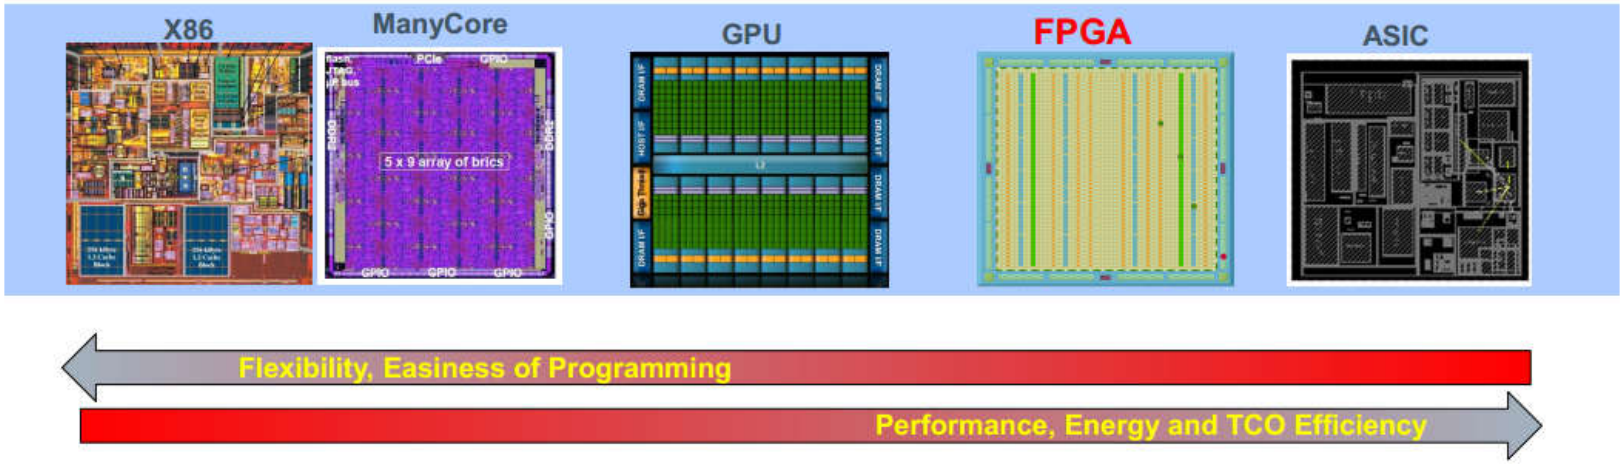

CPU : Range of few TF/s and <200 GB/s per node. Large memory footprint. Standard programing model (PM)

MIC : bootable. hundreds of threads. 3TF/s DP , 500GB/s HBM , Standard PM

FPGA : Discrete and with Xeon CPU. 1.5 TF/s A10 (now)

GenGraphics: Single socket CPU. 45W. 1.5 GF/s SP . DDR4 and eDRAM. Omp and Ocl

ASIC : > 50 Tops/s , HBM low level PM for now

+ SSD/ NVM + Parallel File System + HPC Software stack, Compilers, Math Libraries

10 000 feet view

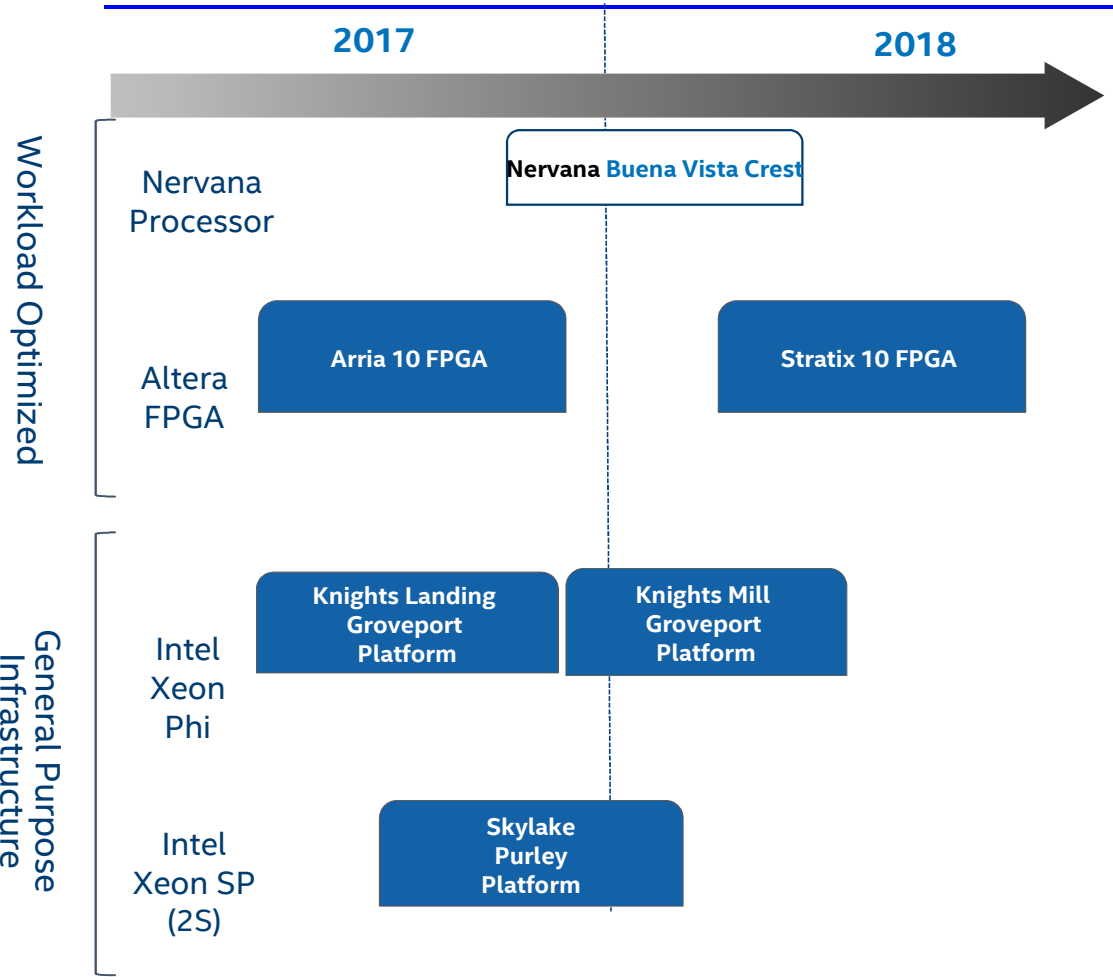

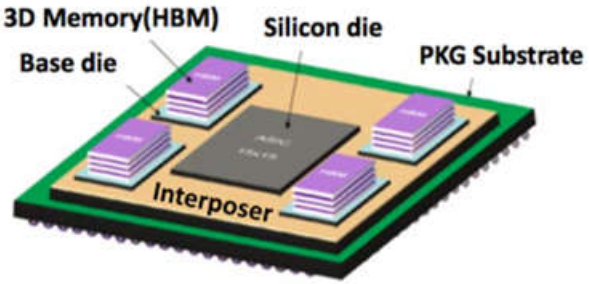

Nervana Engine
https://www.nervanasys.com/technology/engine/

- Patented ***FlexPoint*** precision for maximum Tput and high accuracy
- Over Tera b/s of inter/Intra connectivity for optimal scaling
- Standard PCIe Gen3x16

Storage Evolution

Storage class memory

On package memory

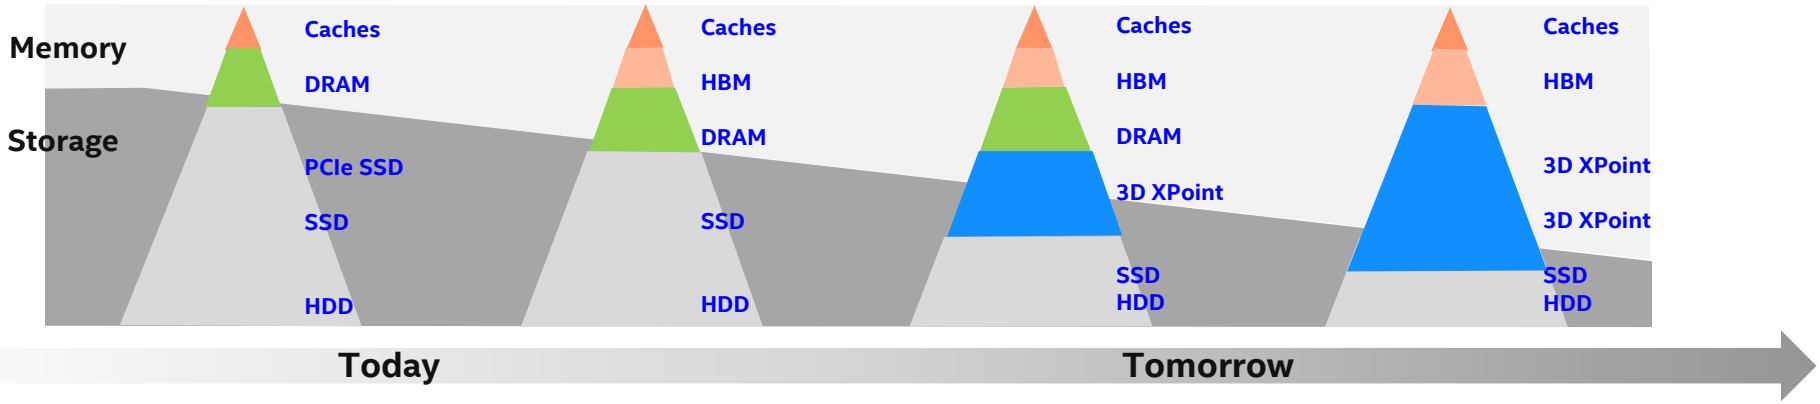

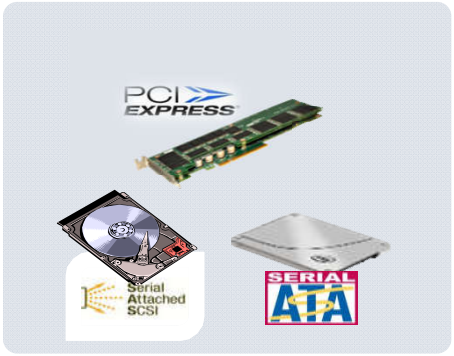

NAND based Intel P3700 (Fultondale) for NVMe

3D XPoint

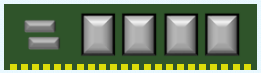

3D XPoint "DIMM"

3D Xpoint : DIMM form factor

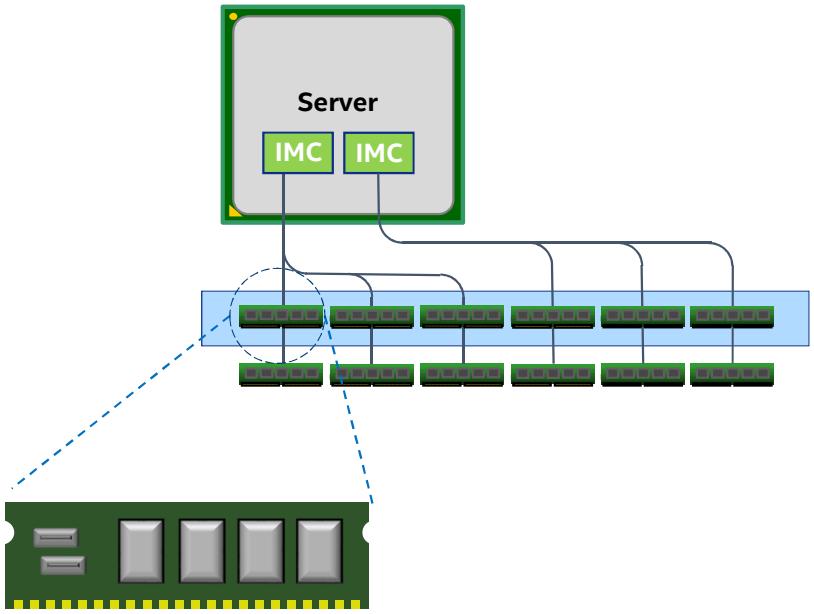

- DDR4 electrical & physical
- Close to DRAM latency

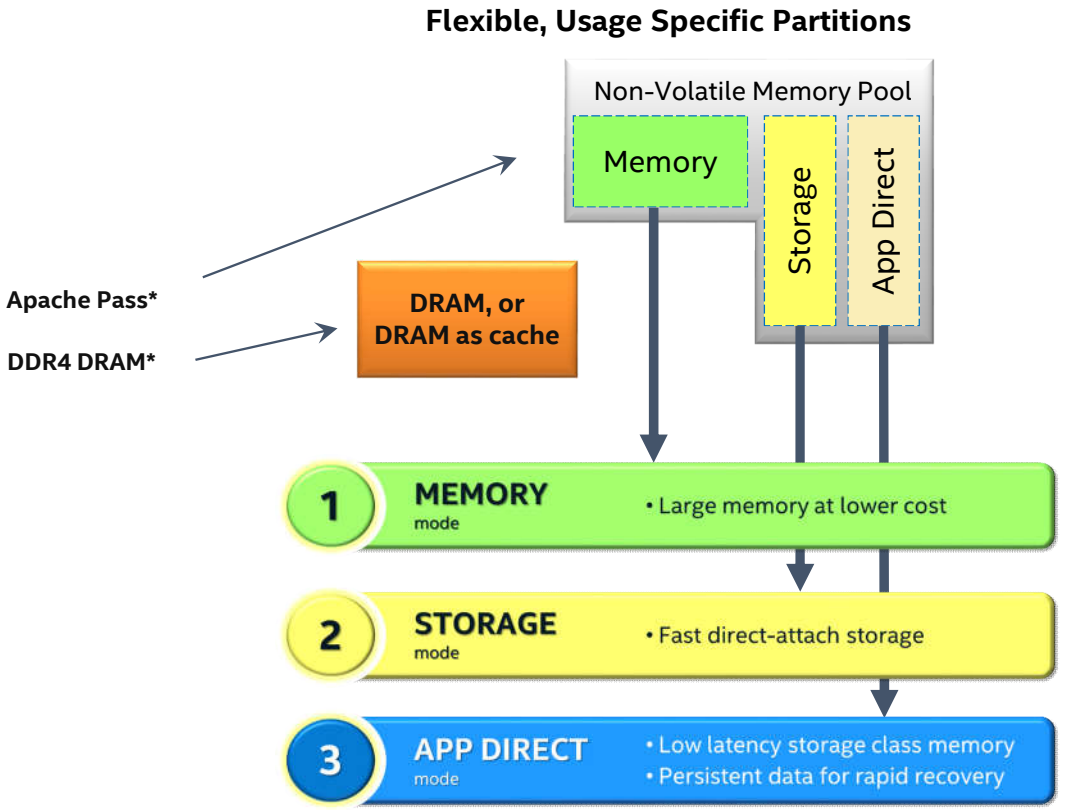

Compiler, mkl, ipp, MPI , Openmp, Opencl

vtune, advisor, itac, inspector ..

CPU trend

Memory update

Software

Characterization

… in 5 mn

MPI Performance Snapshot

Your application is OpenMP bound.
High OpenMP imbalance has been identified.
Use Intel VTune Amplifier for further analysis.

Application: /nfs/inn/home/yshchyok/p/svn/testing/ts/results/2015.09.23 12.31.09/itac_testspec/vt_key_default_test_c_icc15_n2_itac_it_mps/test
Number of ranks: 4
Used statistics: app_stat_4p28t.txt, stats_4p28t.txt
Creation date: 2015-09-28 14:58:48

Wallclock time

1.78 sec

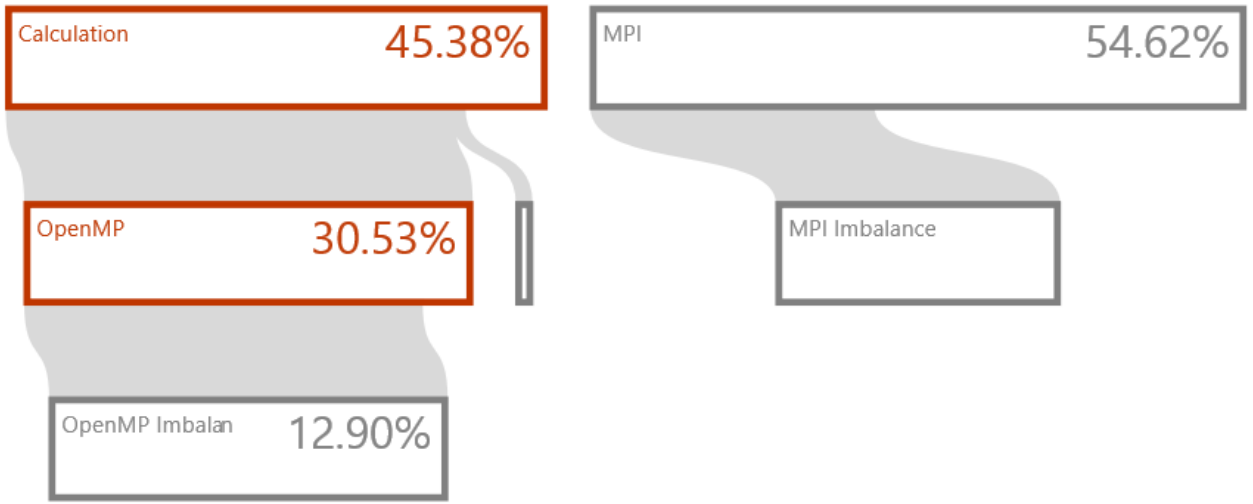

TOP 5 MPI functions

Func	%
Wait	71.98
Barrier	20.92
Init	3.98
Send	2.04
Recv	0.93

Memory usage

PEAK 25.29 MB
MEAN 20.95 MB

Per-process memory usage affects the application scalability.

Cycles Per Instruction Rate

4
1.67
0
max: 2.23
min: 1.10

This could be caused by such issues as memory stalls, instruction starvation, branch misprediction or long latency instructions.

Please use Intel® VTune™ Amplifier XE to identify the cause of this bottleneck.
High values are usually bad. The CPI value may be *too high*.

GFLOPS

20.67

I/O operations

I/O wait: **0.00 sec** 0.00%

PEAK 0.00%
MEAN 0.00%

This is the time the application spends waiting for an I/O operation to complete. High percentage of I/O wait time indicates that your application actively reads data from the storage device. This application does not spend much time on I/O operations.

Memory Bound Coefficient

1
0.18
0
max: 0.23
min: 0.14

It indicates that the application doesn't spend much time waiting for data. High values are usually bad. The application is *not Memory Bound*.

Free download: http://www.intel.com/performance-snapshot. Also included with Intel® Parallel Studio Cluster Edition.

Storage performance snapshot

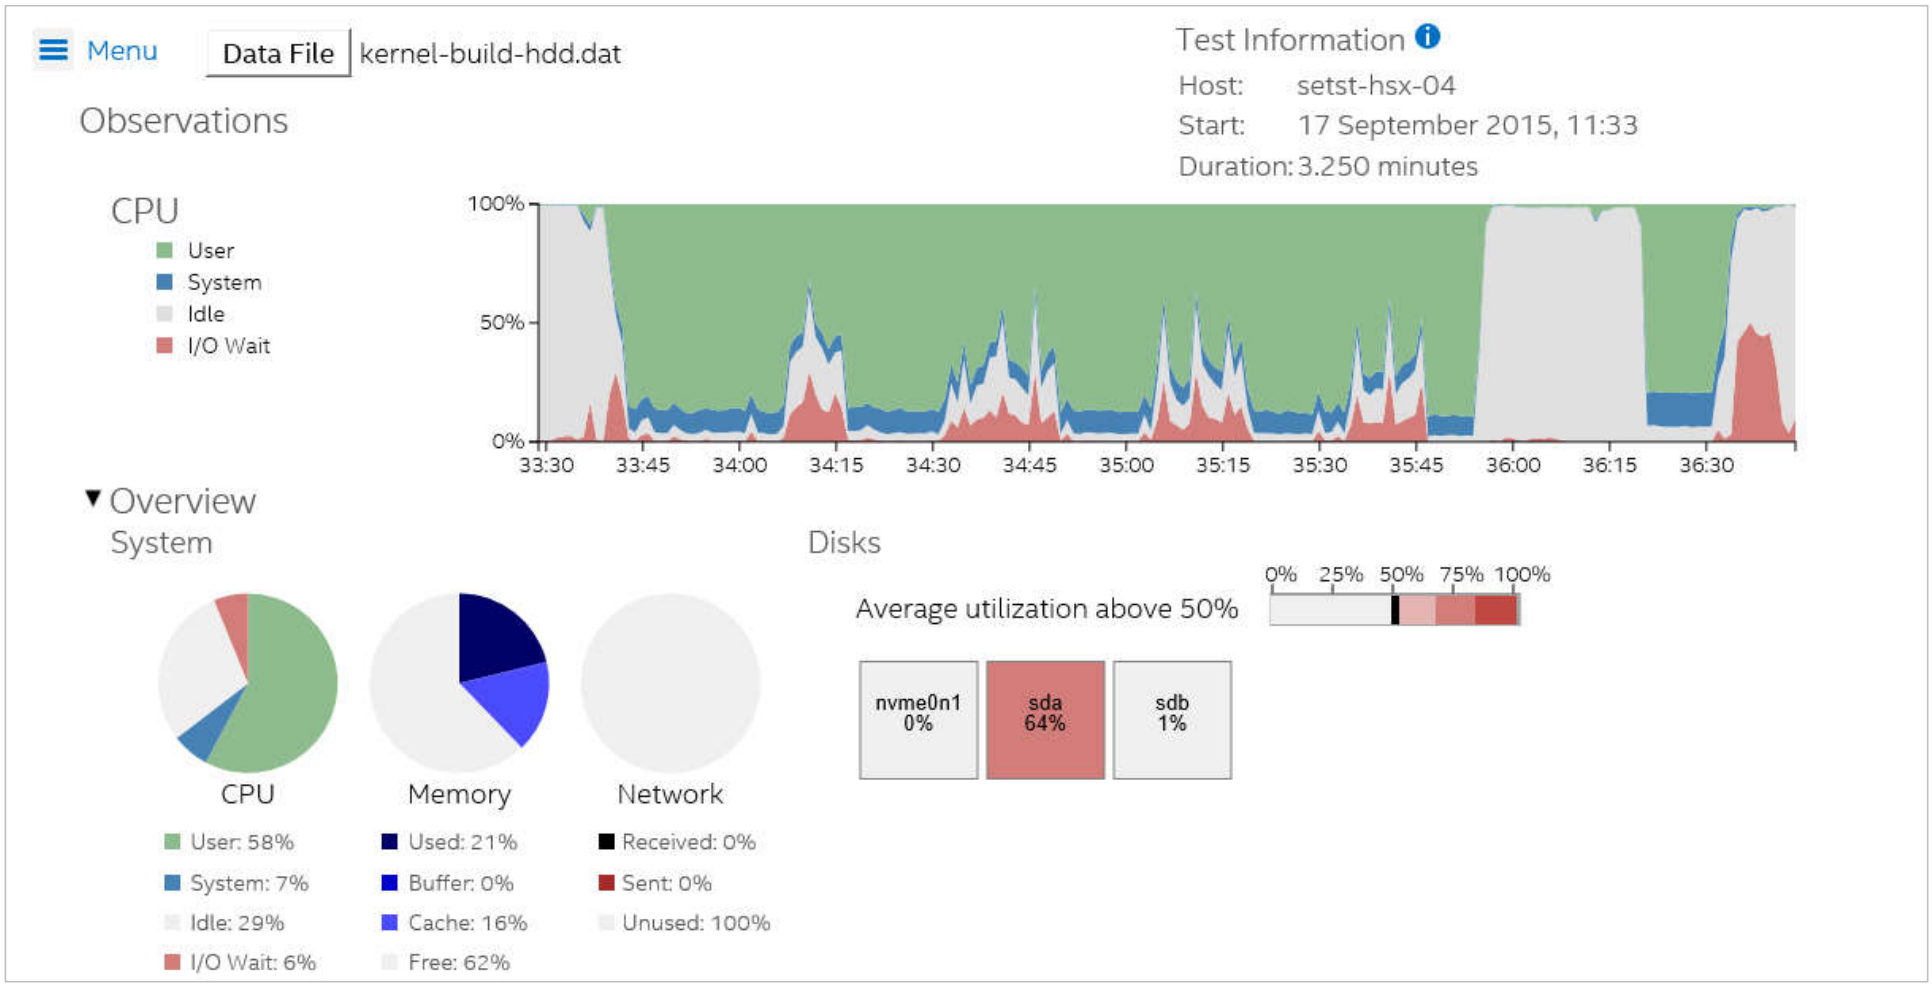

Intel® Advisor - Vectorization Advisor

The data and guidance you need:

- Compiler diagnostics + Performance Data + SIMD efficiency
- Detect problems & recommend fixes
- Loop-Carried Dependency Analysis
- Memory Access Patterns Analysis

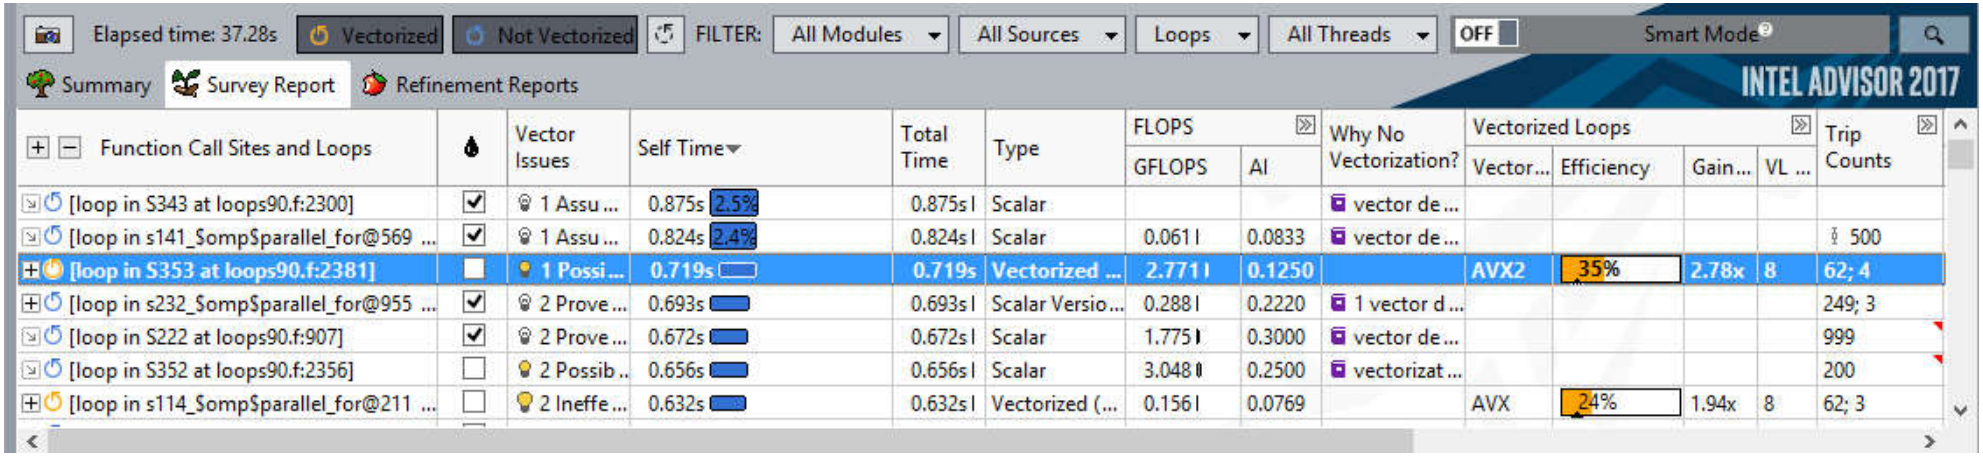

Function Call Sites and Loops	Vector Issues	Self Time	Total Time	Type	FLOPS		Why No Vectorization?	Vectorized Loops				Trip Counts
					GFLOPS	AI		Vector...	Efficiency	Gain...	VL ...	
[loop in S343 at loops90.f:2300]	1 Assu...	0.875s 2.5%	0.875s	Scalar			vector de...					
[loop in s141_ompparallel_for@569 ...	1 Assu...	0.824s 2.4%	0.824s	Scalar	0.061	0.0833	vector de...					500
[loop in S353 at loops90.f:2381]	1 Possi...	0.719s	0.719s	Vectorized ...	2.771	0.1250		AVX2	35%	2.78x	8	62; 4
[loop in s232_ompparallel_for@955 ...	2 Prove...	0.693s	0.693s	Scalar Versio...	0.288	0.2220	1 vector d...					249; 3
[loop in S222 at loops90.f:907]	2 Prove...	0.672s	0.672s	Scalar	1.775	0.3000	vector de...					999
[loop in S352 at loops90.f:2356]	2 Possib...	0.656s	0.656s	Scalar	3.048	0.2500	vectorizat...					200
[loop in s114_ompparallel_for@211 ...	2 Ineffe...	0.632s	0.632s	Vectorized (...	0.156	0.0769		AVX	24%	1.94x	8	62; 3

Roofline Model. High level characterization

R.M gives the max achievable performance on a given platform

$$GFlop/s(AI) = min \begin{cases} p_f \\ AI \times p_b \end{cases} or \ \min \begin{cases} xGEMM \\ AI \times StreamBW \end{cases}$$

See where your application stands
And what you can expect

AI: arithmetic intensity
p_f: peak FP
p_b: peak bandwidth

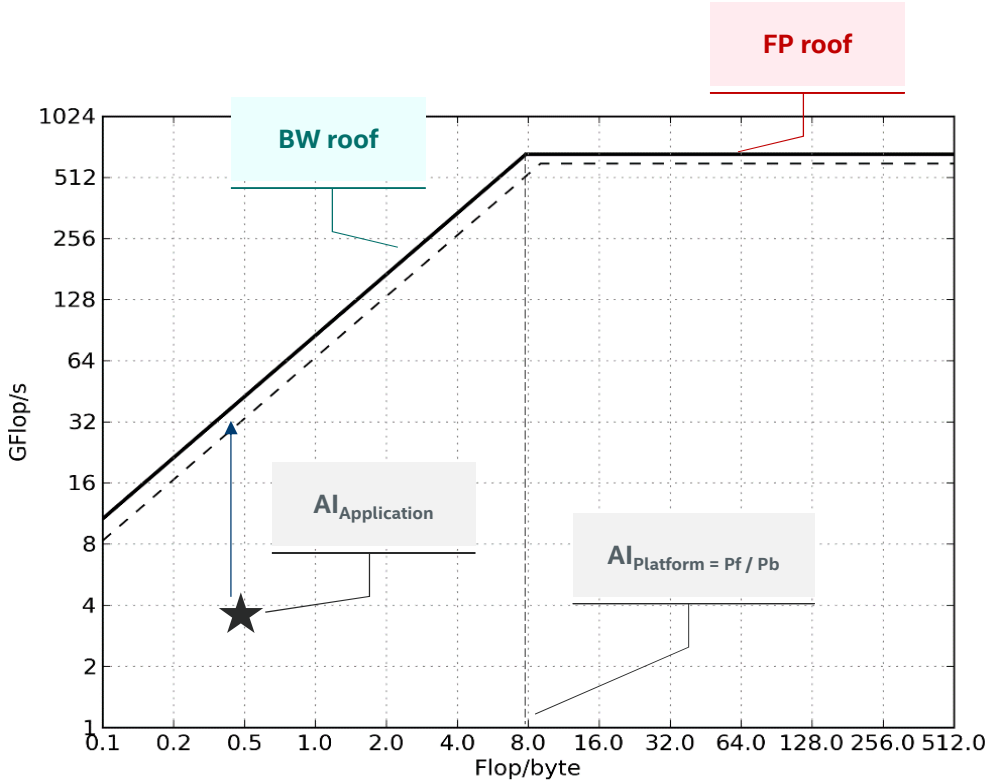

D. Lazowska, J. Zahorjan, G. Graham, K. Sevcik,"*Quantitative System Performance*" (1984)

S. Williams, "*Roofline: An Insightful Visual Performance Model for Floating-Point Programs and Multicore Architectures*". (2009)

Performance Characterization

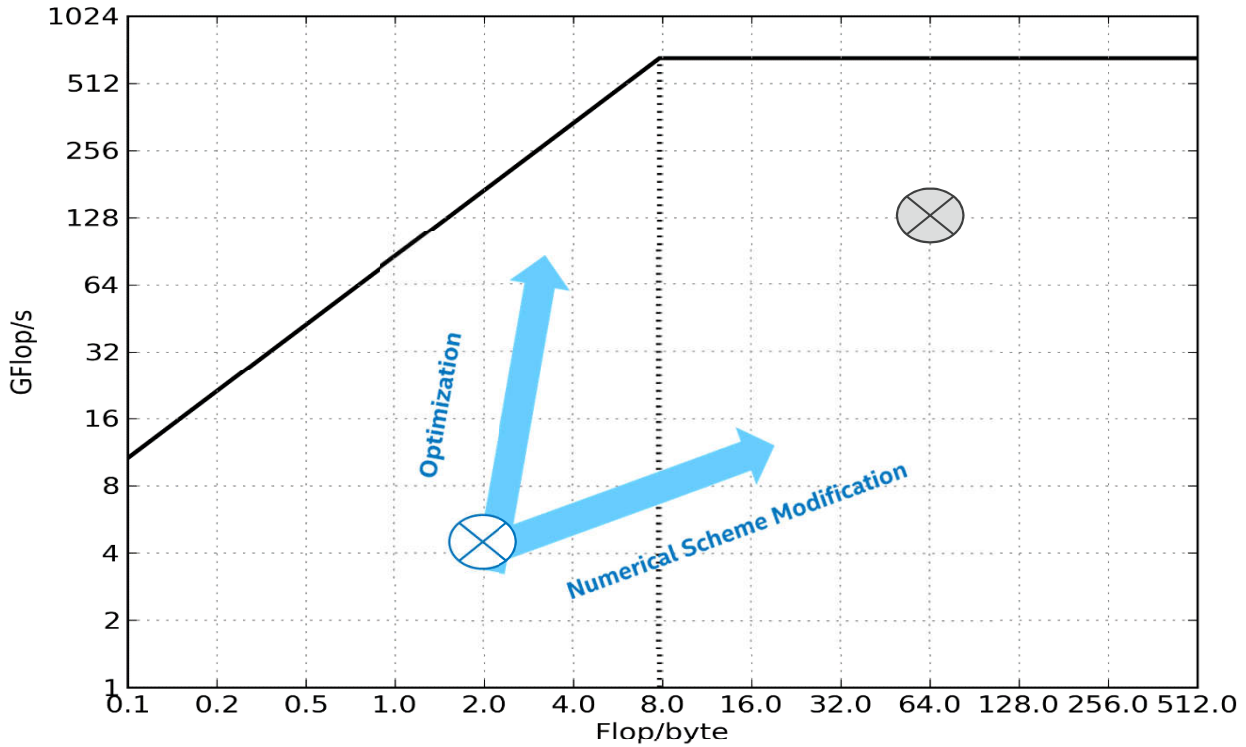

Temporal roofline: Application phases identification

An application with 2 phases in the extremities on E5-2697 v2

- ✓ Bandwidth bound: Stream
- ✓ CPU bound: HPL

The scalar AI and flops are averaged and are not representative of the application evolution.

Temporal roofline identifies these phases distinctively.

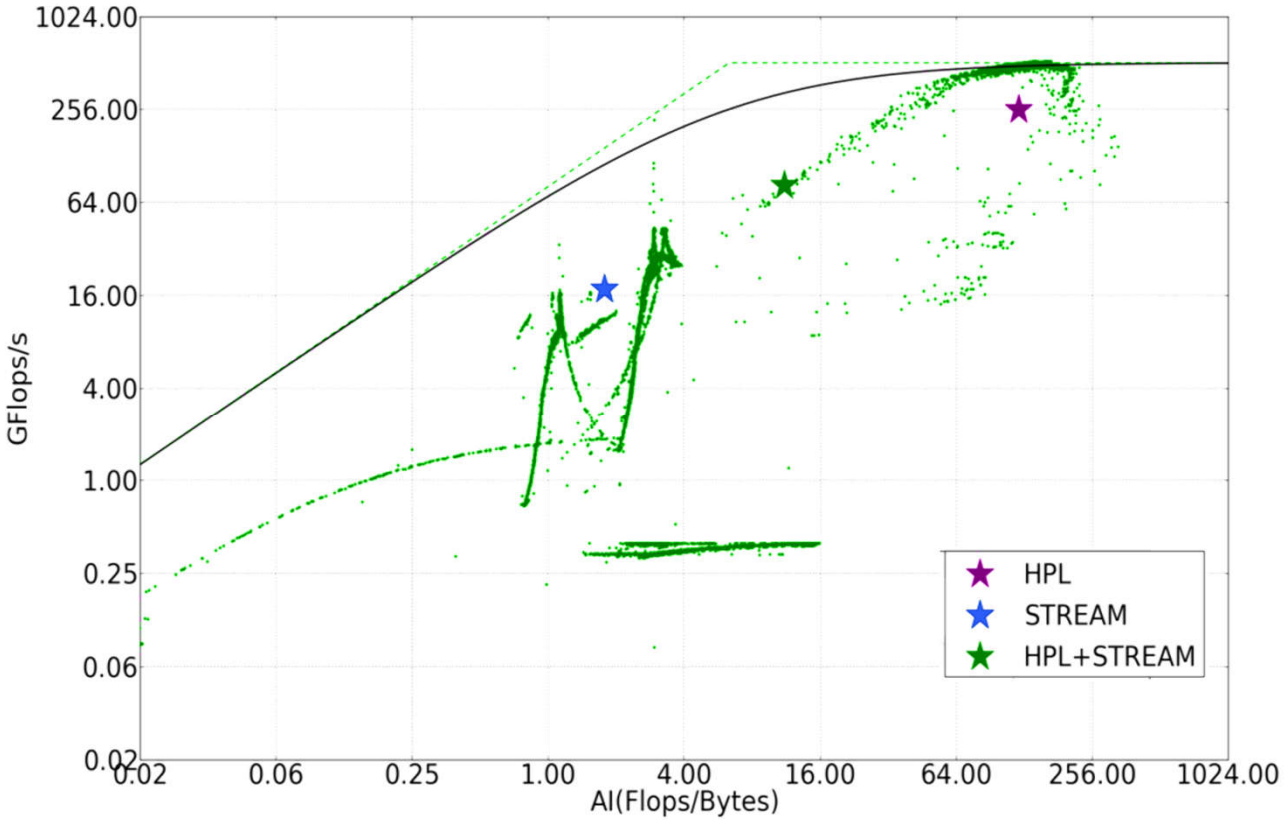

Courtesy of A. Mrabet et al. 2015

Cache aware roofline

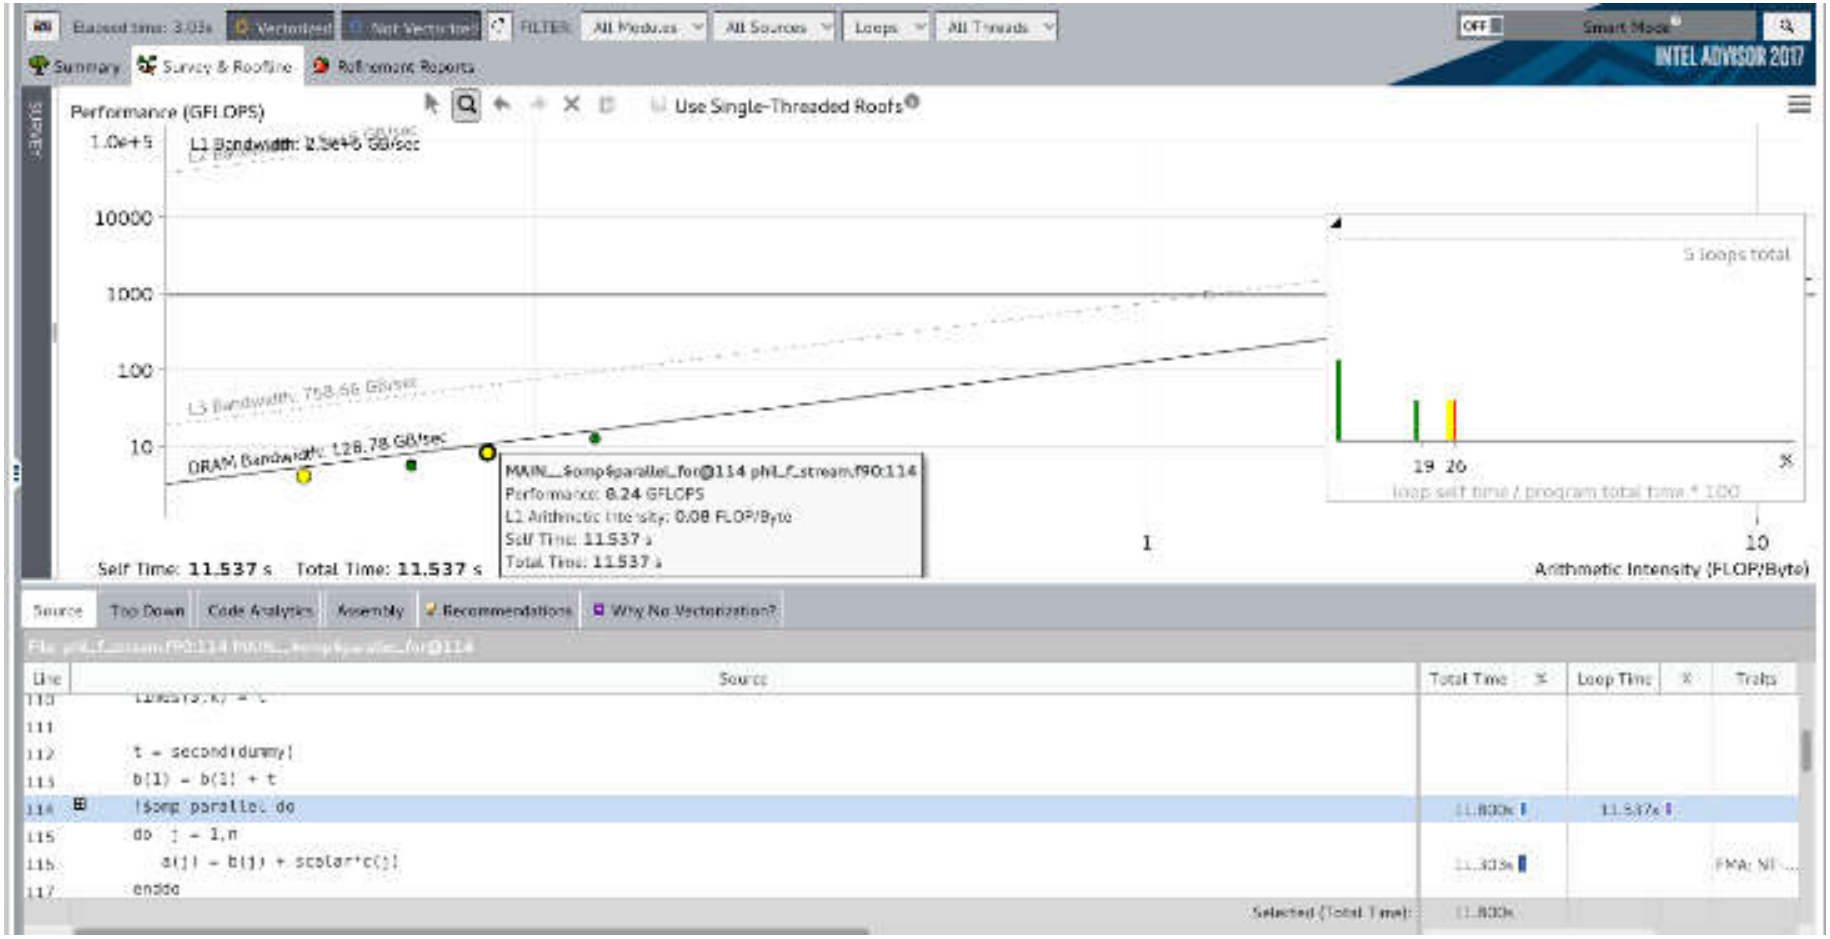

Cache aware roofline

- Know where you are vs peaks
- “what to expect”
- Per block view for the whole apps.
- Linked to source and assembly

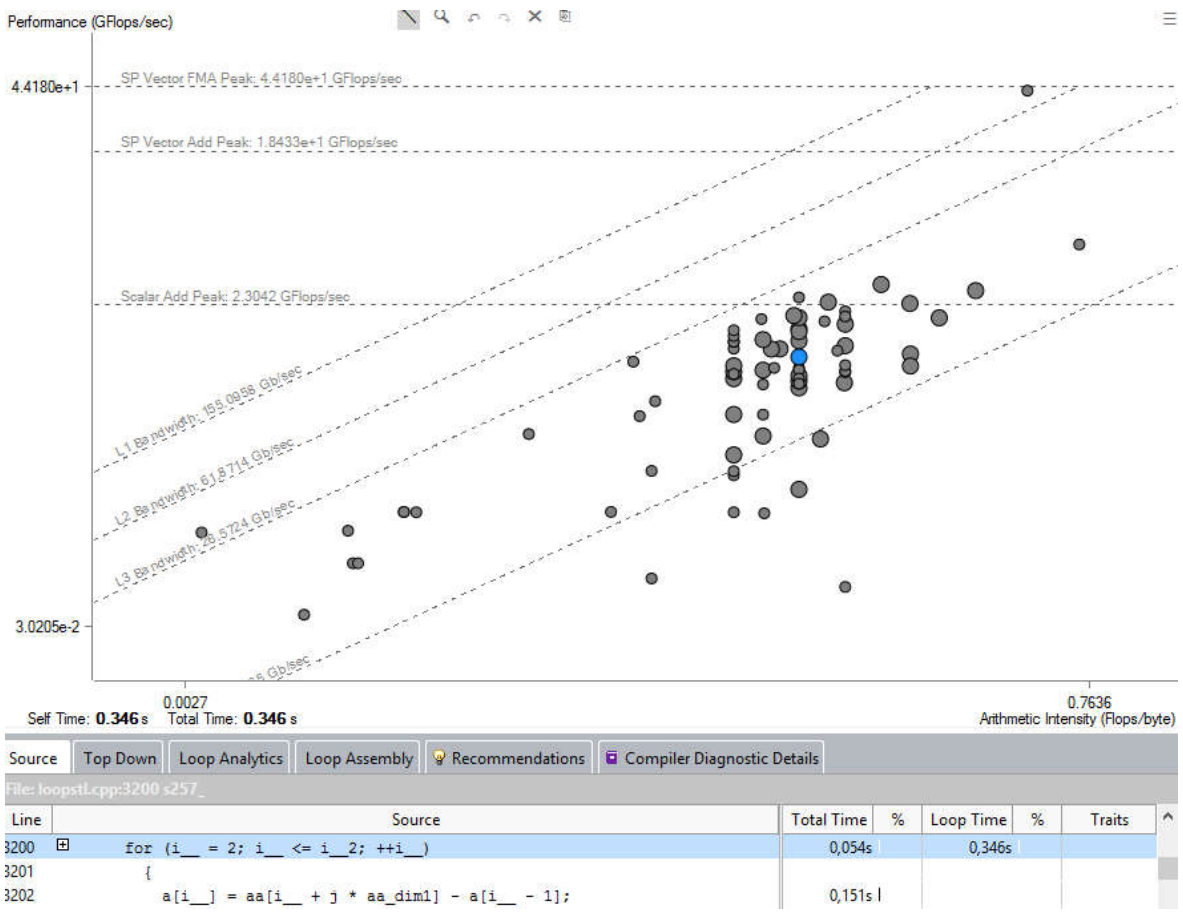

How to collect Flops and Bytes (AI definition)

$$AI = \frac{Flop}{Byte}$$

SDE
- ✓ Possible for future architectures
- ✓ Average over execution time

Hardware counters
- ✓ Not always available
- ✓ Vtune, PCM, LIKWID, PAPI

By hands
- ✓ Not always possible

Hardware counters
- ✓ Vtune, PCM, LIKWID, PAPI

By hands
- ✓ Not always possible

How to define memory demand without cache impact or latency ?

DRAM demand *: how many DRAM transactions a workload 'wants' to do.*
versus DRAM BW : *the no. of DRAM transactions completed per unit time*

Programming model

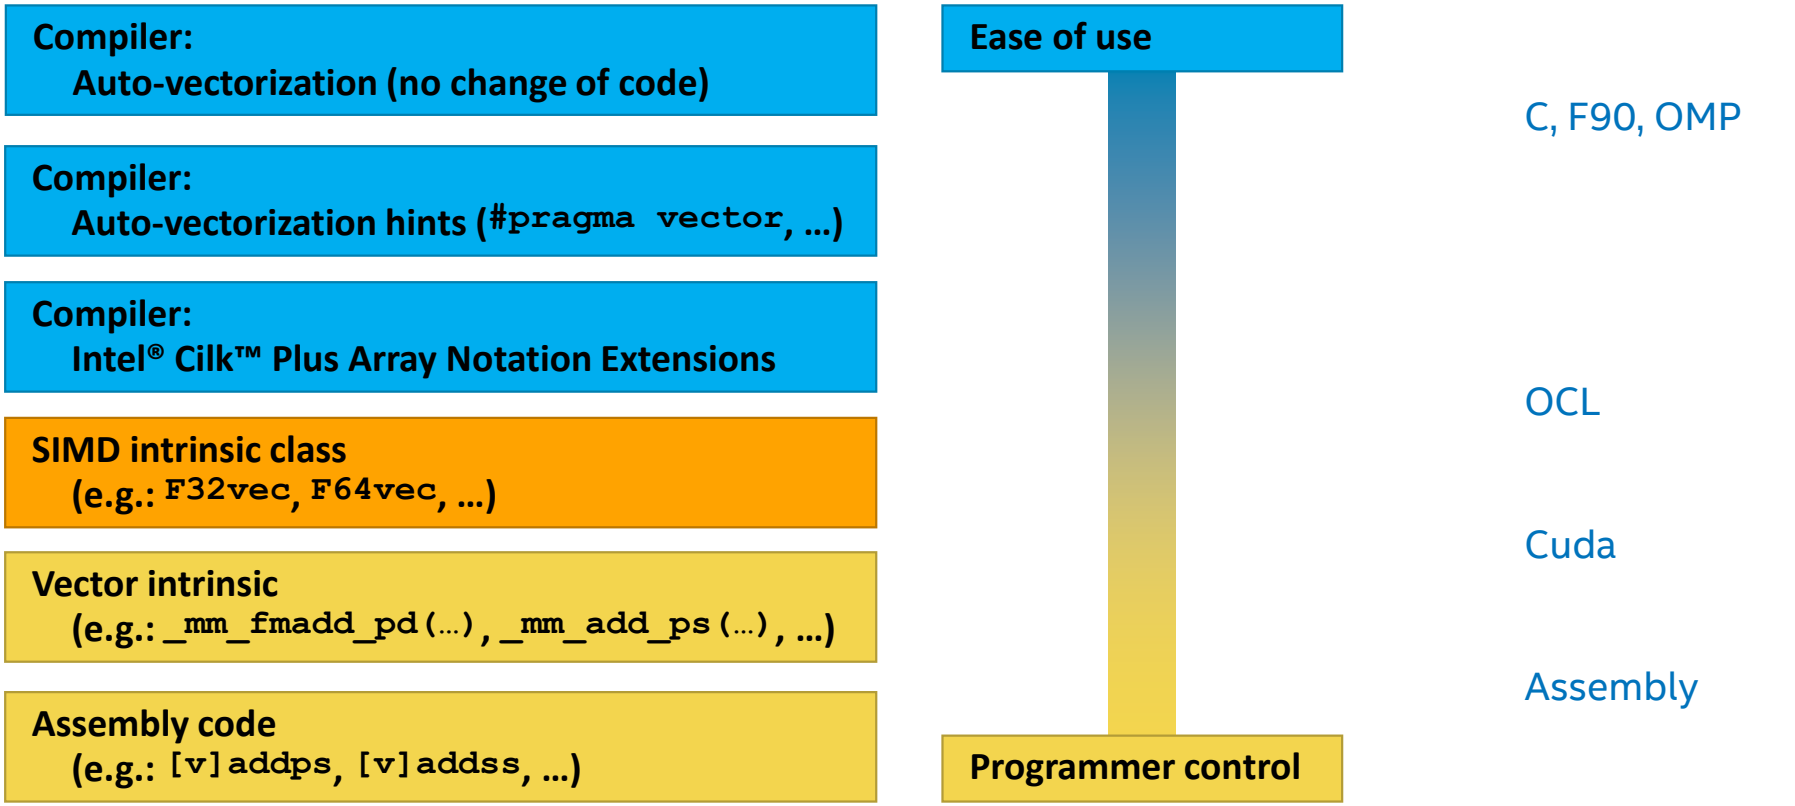

Questions?


intel
experience
what's inside™

Legal Notices and Disclaimers

Legal Disclaimer & Optimization Notice

Optimization Notice

Intel’s compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804