GridKa School 2019

KIT, Campus North, FTU

INTEL'S HARDWARE & SOFTWARE SOLUTIONS - DIRECTIONS FOR ARTIFICIAL INTELLIGENCE

Edmund Preiss

Business Development Manager, EMEA
Intel Computing Performance and Software Products (CPDP)

Agenda

Intel Xeon Family – Latest Status

Intel Software Development Tools

- Intel Parallel Studio XE (IPS XE) Tool Suites
- News on the IPS XE 2020 Edition
- Upcoming **oneAPI** development tools concept

Intel optimized AI Solutions and Directions

EVOLUTION OF INTEL® XEON® PLATFORM

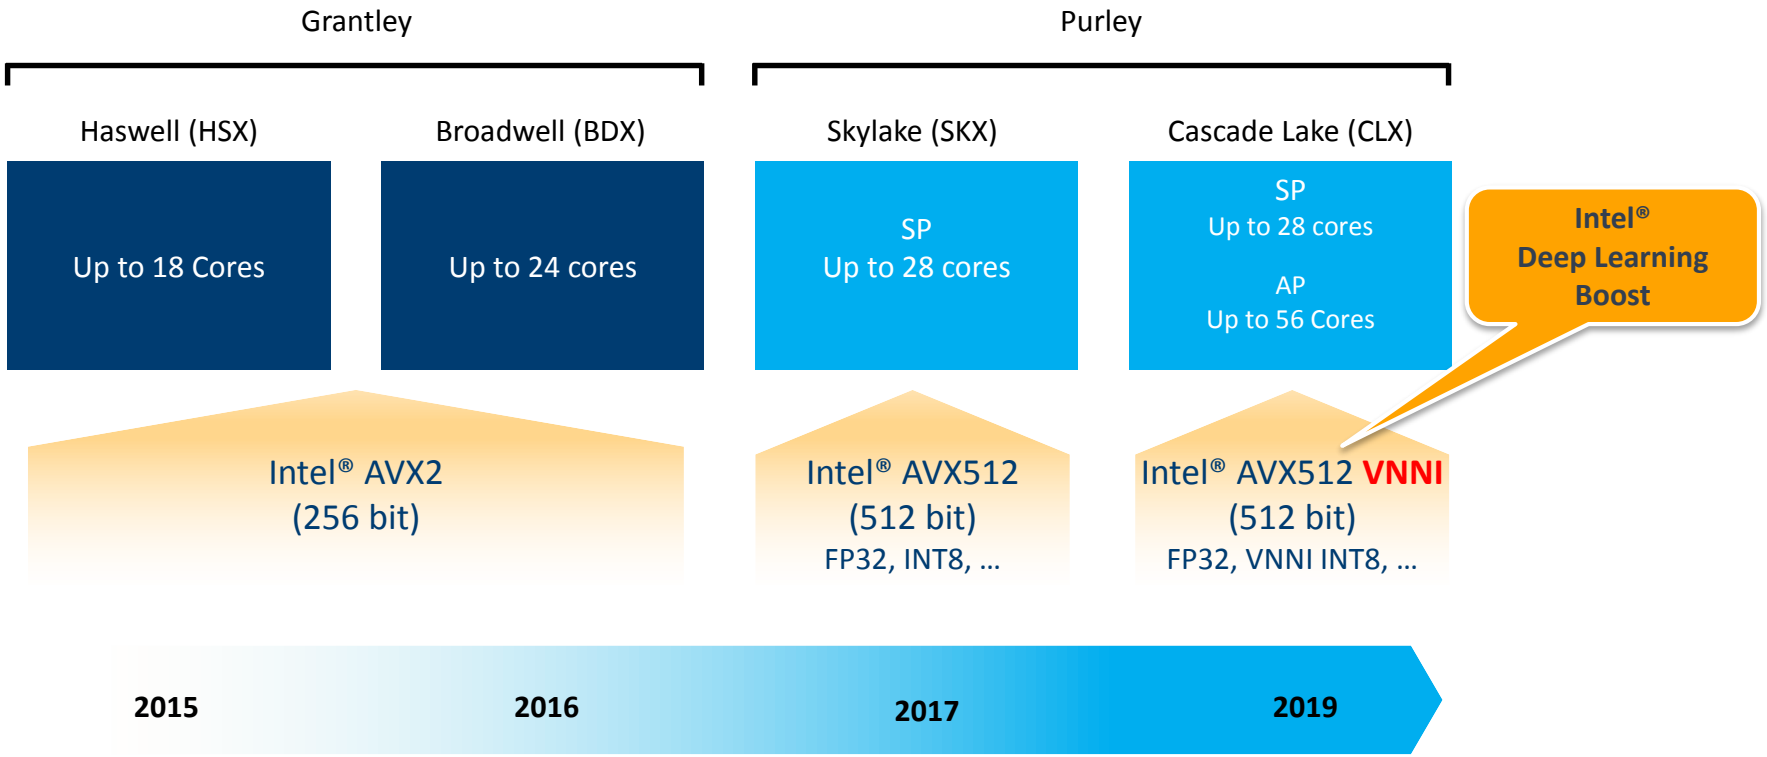

Agenda

Intel Xeon Family Update

Intel Software Development Tools

- Intel Parallel Studio XE (IPS XE) Tool Suites
- News on the IPS XE 2020 Edition
- Upcoming **oneAPI** development tools concept

Intel optimized AI Solutions and Directions

Code Modernization

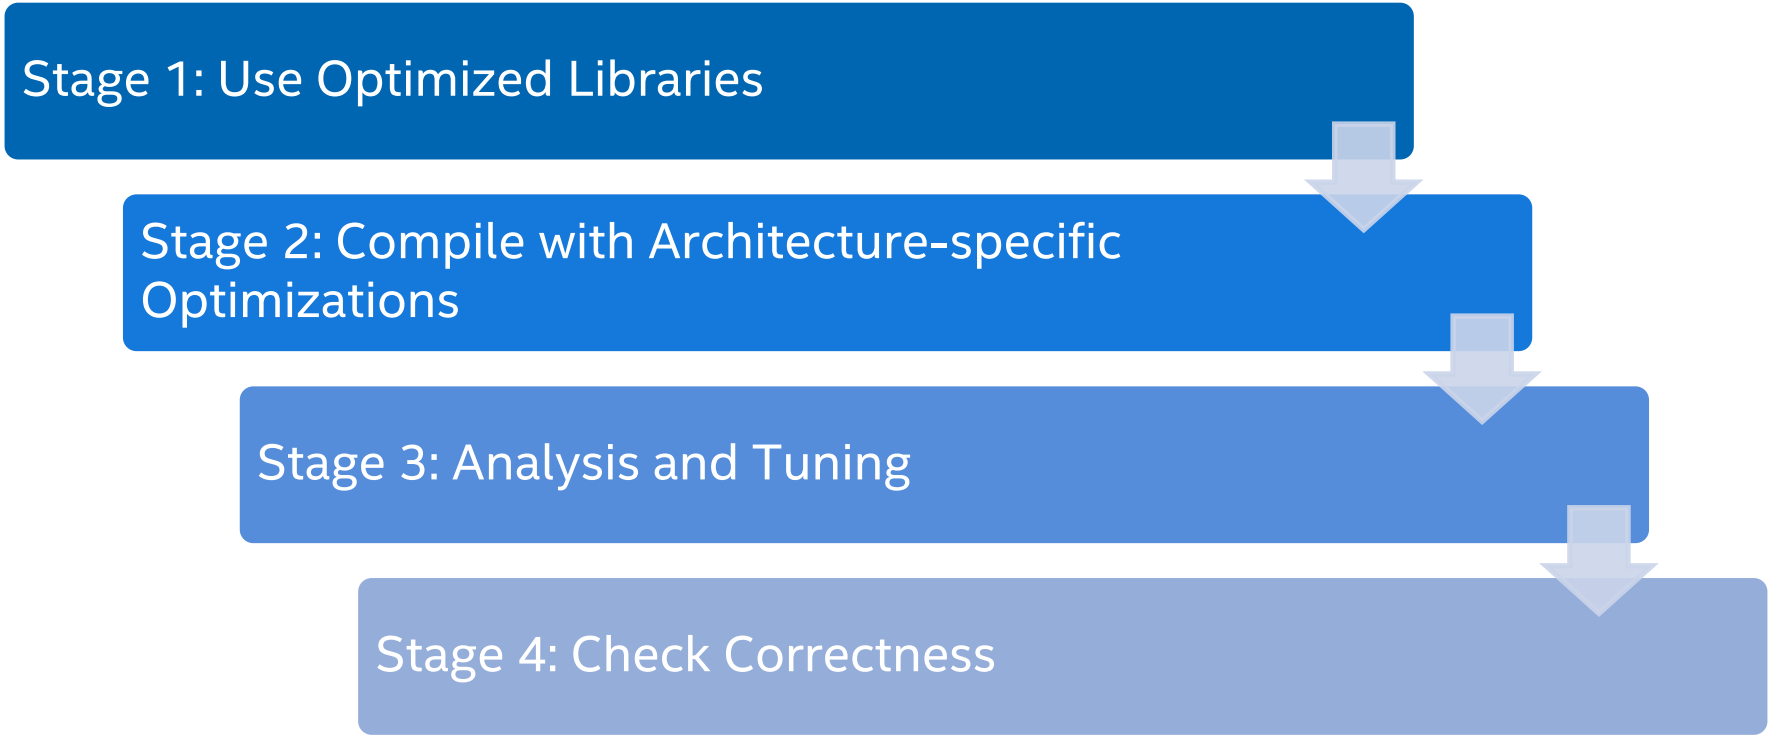

What's Inside Intel® Parallel Studio XE

Comprehensive Software Development Tool Suite

COMPOSER EDITION

BUILD
Compilers & Libraries

- C / C++ Compiler
  Optimizing Compiler
- Fortran Compiler
  Optimizing Compiler
- Intel® Threading Building Blocks
  C++ Threading Library
- Intel® Math Kernel Library
- Intel® Integrated Performance Primitives
  Image, Signal & Data Processing
- Intel® Data Analytics Acceleration Library
- Intel® Distribution for Python*
  High Performance Scripting

PROFESSIONAL EDITION

ANALYZE
Analysis Tools

- Intel® VTune™ Amplifier
  Performance Profiler
- Intel® Inspector
  Memory & Thread Debugger
- Intel® Advisor
  Vectorization Optimization & Thread Prototyping

CLUSTER EDITION

SCALE
Cluster Tools

- Intel® MPI Library
  Message Passing Interface Library
- Intel® Trace Analyzer & Collector
  MPI Tuning & Analysis
- Intel® Cluster Checker
  Cluster Diagnostic Expert System

Intel® Architecture Platforms

Operating System: Windows*, Linux*, MacOS[1]*

More Power for Your Code - software.intel.com/intel-parallel-studio-xe

SELECTED INTEL PARALLEL STUDIO XE HIGHLIGHTS

What's Inside Intel® Math Kernel Library

LINEAR ALGEBRA

- BLAS
- LAPACK
- ScaLAPACK
- Sparse BLAS
- Iterative sparse solvers
- PARDISO*

FFTS

- Multidimensional
- FFTW interfaces

VECTOR RNGS

- Congruential
- Wichmann-Hill
- Mersenne Twister
- Sobol
- Neirderreiter
- Non-deterministic

SUMMARY STATISTICS

- Kurtosis
- Variation coefficient
- Order statistics
- Min/max
- Variance-covariance

VECTOR MATH

- Trigonometric
- Hyperbolic
- Exponential
- Log
- Power
- Root

& MORE

- Splines
- Interpolation
- Trust Region
- Fast Poisson Solver

[1]Available only in Intel® Parallel Studio Composer Edition.

Application Performance Snapshot (VTune Amplifier)

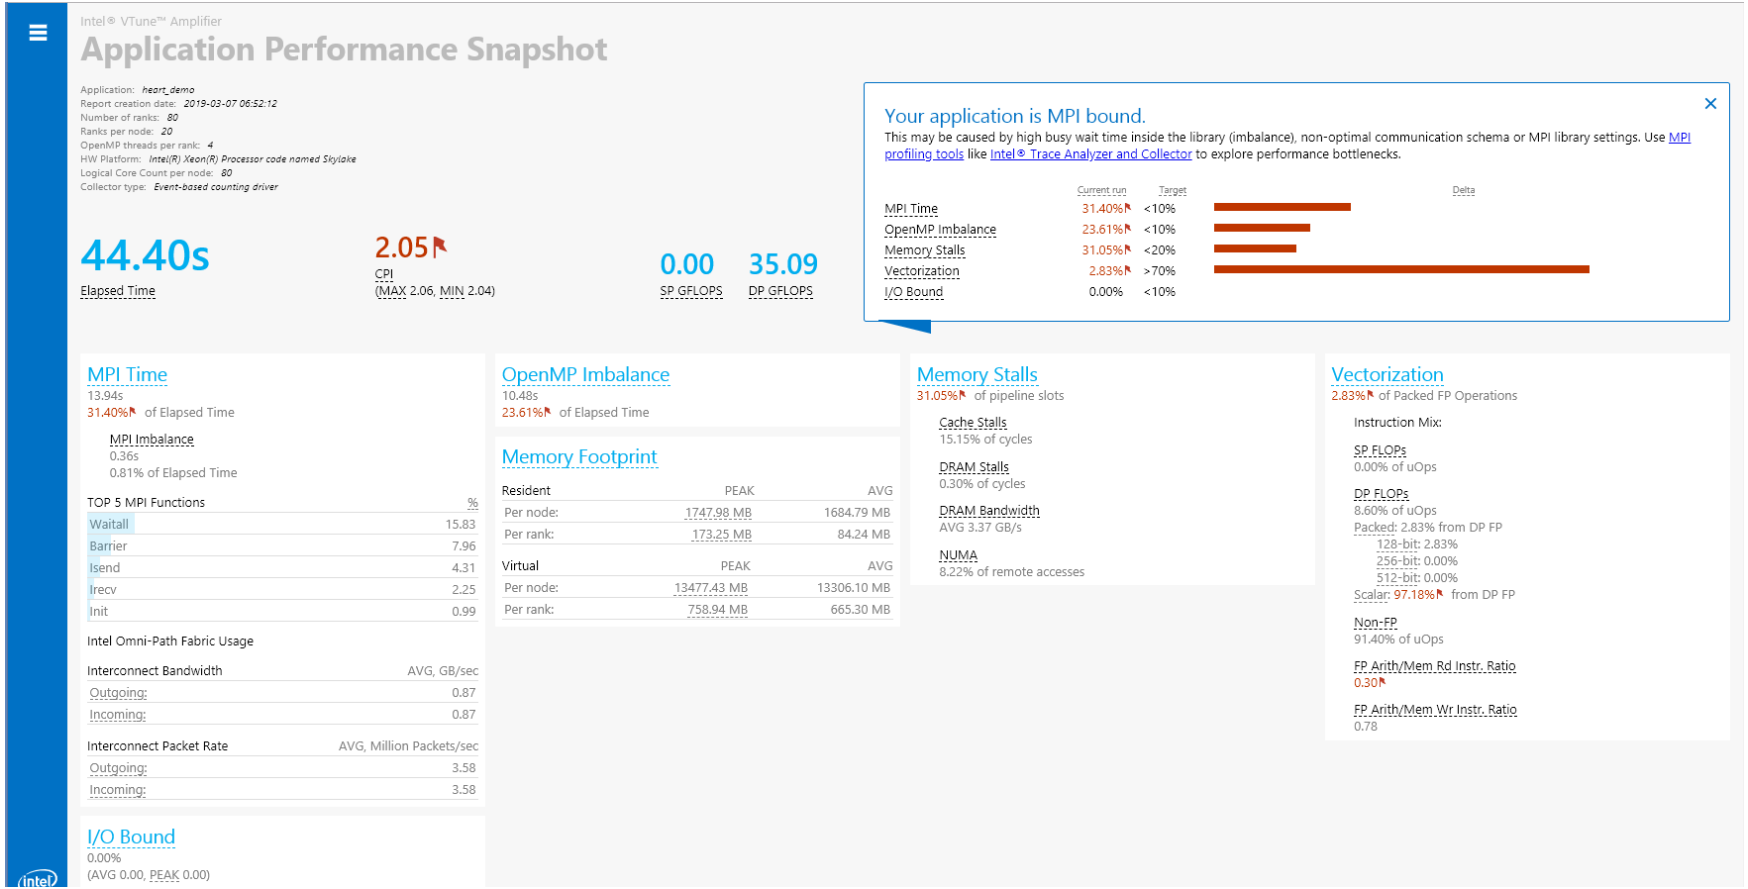

Find Effective Optimization Strategies

Intel Advisor: Cache-aware roofline analysis

Roofs Show Platform Limits

Memory, cache & compute limits

Dots Are Loops

Bigger, red dots take more time so optimization has a bigger impact

Dots farther from a roof have more room for improvement

Higher Dot = Higher GFLOPs/sec

Optimization moves dots up

Algorithmic changes move dots horizontally

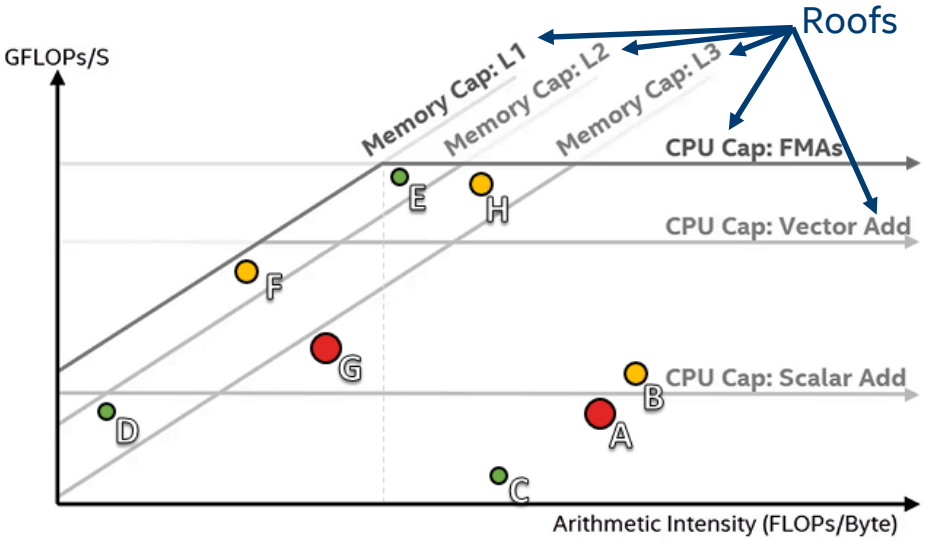

Which loops should we optimize?

- A and G are the best candidates
- B has room to improve, but will have less impact
- E, C, D, and H are poor candidates

Roofline tutorial video

Analyze & Tune Application Performance & Scalability with Intel® VTune™ Amplifier—Performance Profiler

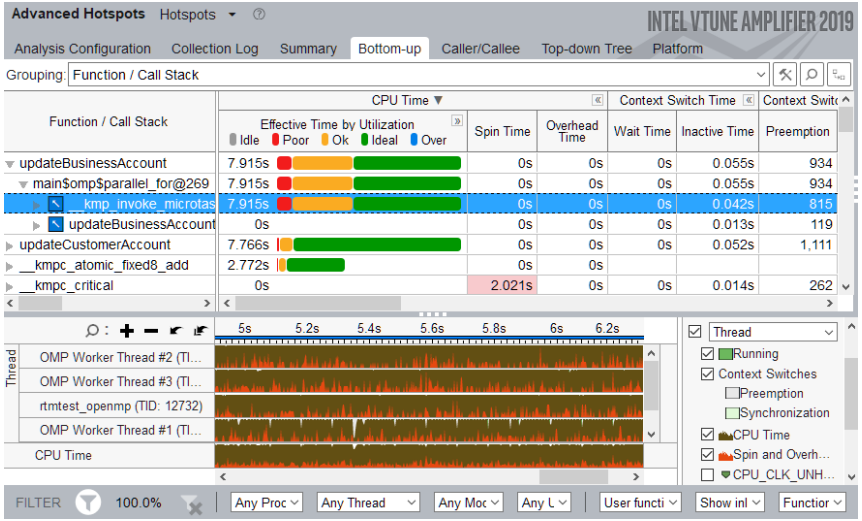

Fast, Scalable Code, Faster

- Accurately profile C, C++, Java*, Python*, Go*, or any mix
- Optimize CPU/GPU, threading, memory, cache, storage & more
- Save time: rich analysis leads to insight

What's New in 2019 Release (Highlights)

- Simplified workflow for easier tuning
- I/O Analysis—Tune SPDK storage & DPDK network performance
- New Platform Profiler—Get insights into platform-level performance, identify memory & storage bottlenecks & imbalances

Optimize Vectorization & Threading with Intel® Advisor

Performance Increases Scale with Each New Hardware Generation

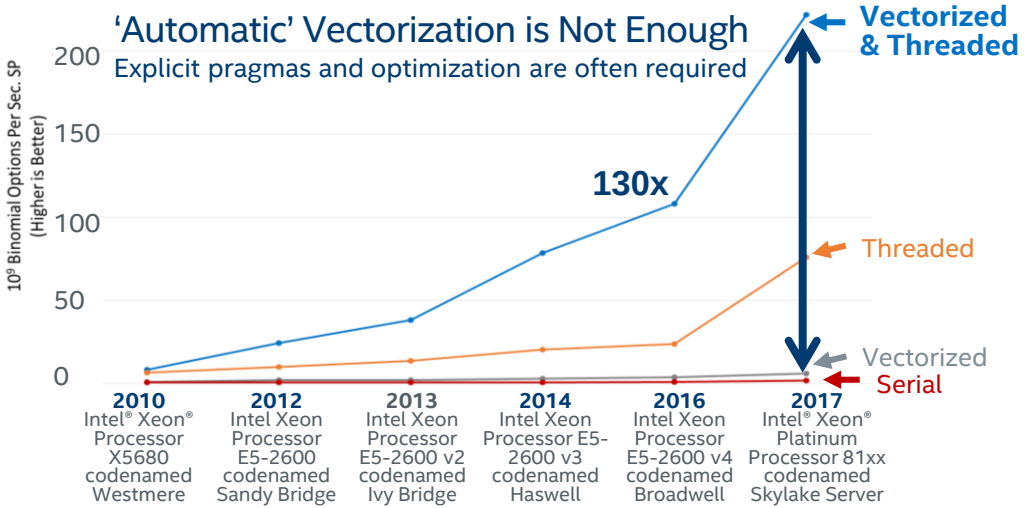

Modern Performant Code

- Vectorized for Intel® Advanced Vector Extensions (Intel® AVX-512 & Intel® AVX)
- Efficient memory access
- Threaded

Capabilities

- Adds & optimizes vectorization
- Analyzes memory patterns
- Quickly prototypes threading

Benchmark: Binomial Options Pricing Model software.intel.com/en-us/articles/binomial-options-pricing-model-code-for-intel-xeon-phi-coprocessor

Performance results are based on testing as of August 2017 and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks. See Vectorize & Thread or Performance Dies Configurations for 2010-2017 Benchmarks in Backup. Testing by Intel as of August 2017.

Learn More: http: intel.ly/advisor

Get Breakthrough Vectorization Performance

Intel® Advisor—Vectorization Advisor

Faster Vectorization Optimization

- Vectorize where it will pay off most
- Quickly ID what is blocking vectorization
- Tips for effective vectorization
- Safely force compiler vectorization
- Optimize memory stride

Data & Guidance You Need

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

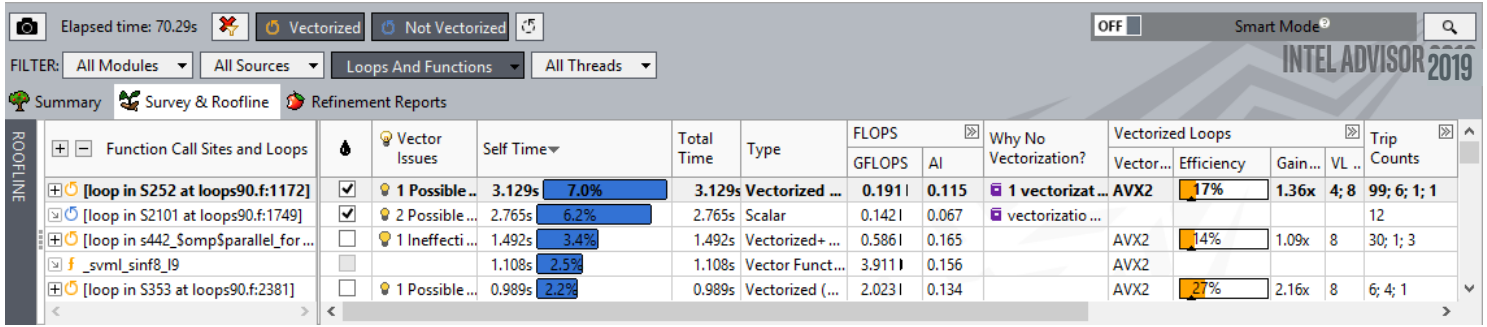

Optimize for Intel® Advanced Vector Extensions 512 (Intel® AVX-512) with or without access to Intel AVX-512 hardware

Agenda

Intel Xeon Family Update

Intel Software Development Tools

- Intel Parallel Studio XE (IPS XE) Tool Suites
- News on the IPS XE 2020 Edition
- Upcoming **oneAPI** development tools concept

Intel optimized AI Solutions and Directions

Key Updates for Intel® Parallel Studio XE 2020

Speed Artificial Intelligence Inferencing - Intel® Compiler and analyzer support for Vector Neural Instructions (VNNI) in Cascade Lake/AP platform

512GB DIMMs with Persistent Memory – Identify, Optimize & Tune Platforms for Intel® Optane™ Persistent Memory with Intel® VTune™ Amplifier

Extended Coarse Grain Profiling – Platform level collection and analysis in Intel® VTune™ Amplifier

Cache Simulation Insights for Vectorization - Roofline analysis for L1, L2 L3, DRAM in Intel® Advisor

Expanded standard support — More Fortran 2018 features & Expanded support of C++17 with initial C++20 support

Latest Processor Support - Intel® Xeon® Scalable Processors (codenamed Cascade Lake / Cascade Lake AP)

New OS Support – Clear Linux & Amazon Linux 2*

* Supported features of tools and libraries may vary by instances and configurations

Intel® VTune™ Profiler Server Architecture

Just launch your browser and go

2020 Preview Feature

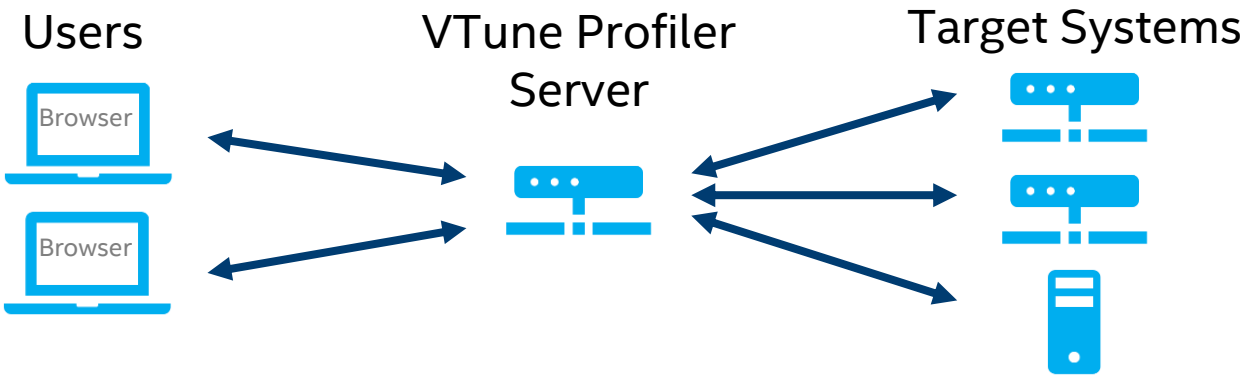

Easier profiling

- **Access with a web browser** – no install required by users
- **Share results** – all results available to all users with server access
- **Profile any system on the network** – server installs collector on the target

What's Coming in Intel® Parallel Studio XE 2020

Coming Q3'2019

- **Intel® Math Kernel Library**
  - Increased AVX512 Optimizations for Complex Vector Math Functions
  - Strided Vector Math API
- **Intel® Data Analytics Acceleration Library**
  - Performance improvements and feature improvements such as
    - Gradient Boosted Trees for large dimensional data sets
    - Extended Z-score support for PCA algorithm
    - XGBoost accelerated with DAAL

Agenda

Intel Xeon Family Update

Intel Software Development Tools

- Intel Parallel Studio XE (IPS XE) Tool Suites
- News on the IPS XE 2020 Edition
- Upcoming **oneAPI** development tools concept

Intel optimized AI Solutions and Directions

DIVERSE WORKLOADS REQUIRE DIVERSE ARCHITECTURES

The future is a **diverse** mix of scalar, vector, matrix, and spatial **architectures** deployed in CPU, GPU, AI, FPGA and other accelerators

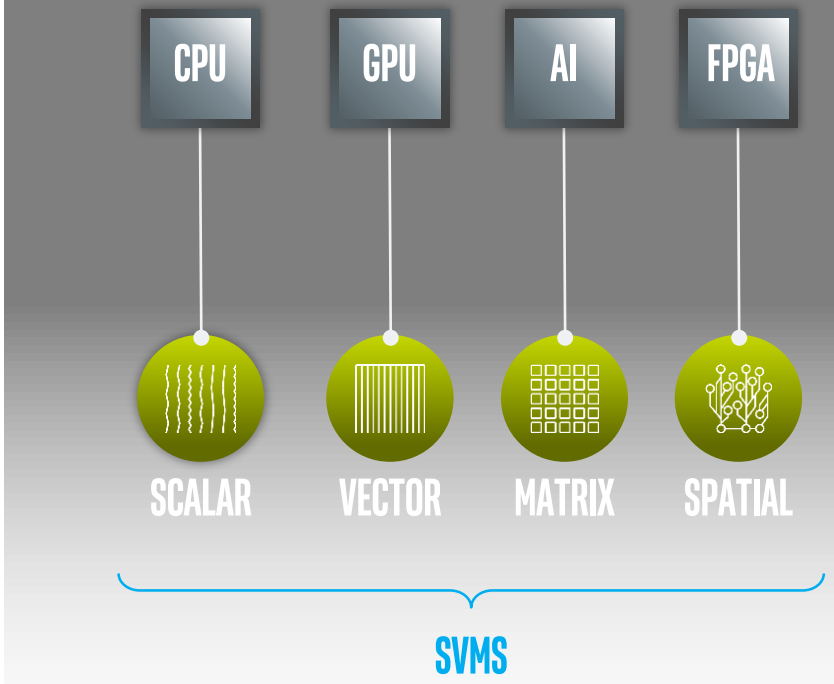

PROGRAMMING CHALLENGE

Diverse set of data-centric hardware

No common programming language or APIs

Inconsistent tool support across platforms

Each platform requires unique software investment

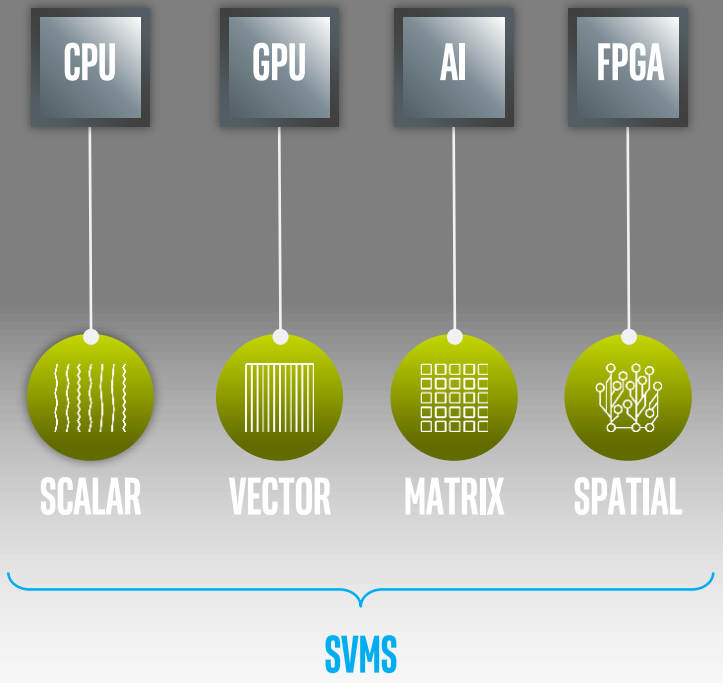

INTEL'S ONE API CORE CONCEPT

Project oneAPI delivers a unified programming model to simplify development across diverse architectures

Common developer experience across Scalar, Vector, Matrix and Spatial (SVMS) architecture

Unified and simplified language and libraries for expressing parallelism

Uncompromised native high-level language performance

Support for CPU, GPU, AI and FPGA

Based on industry standards and open specifications

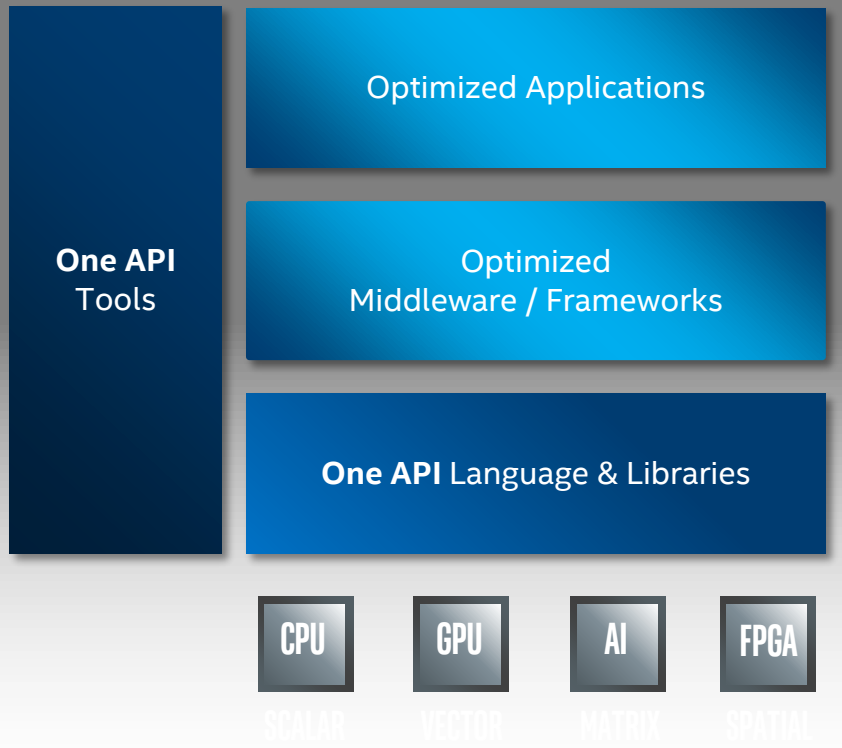

ONE API FOR CROSS-ARCHITECTURE PERFORMANCE

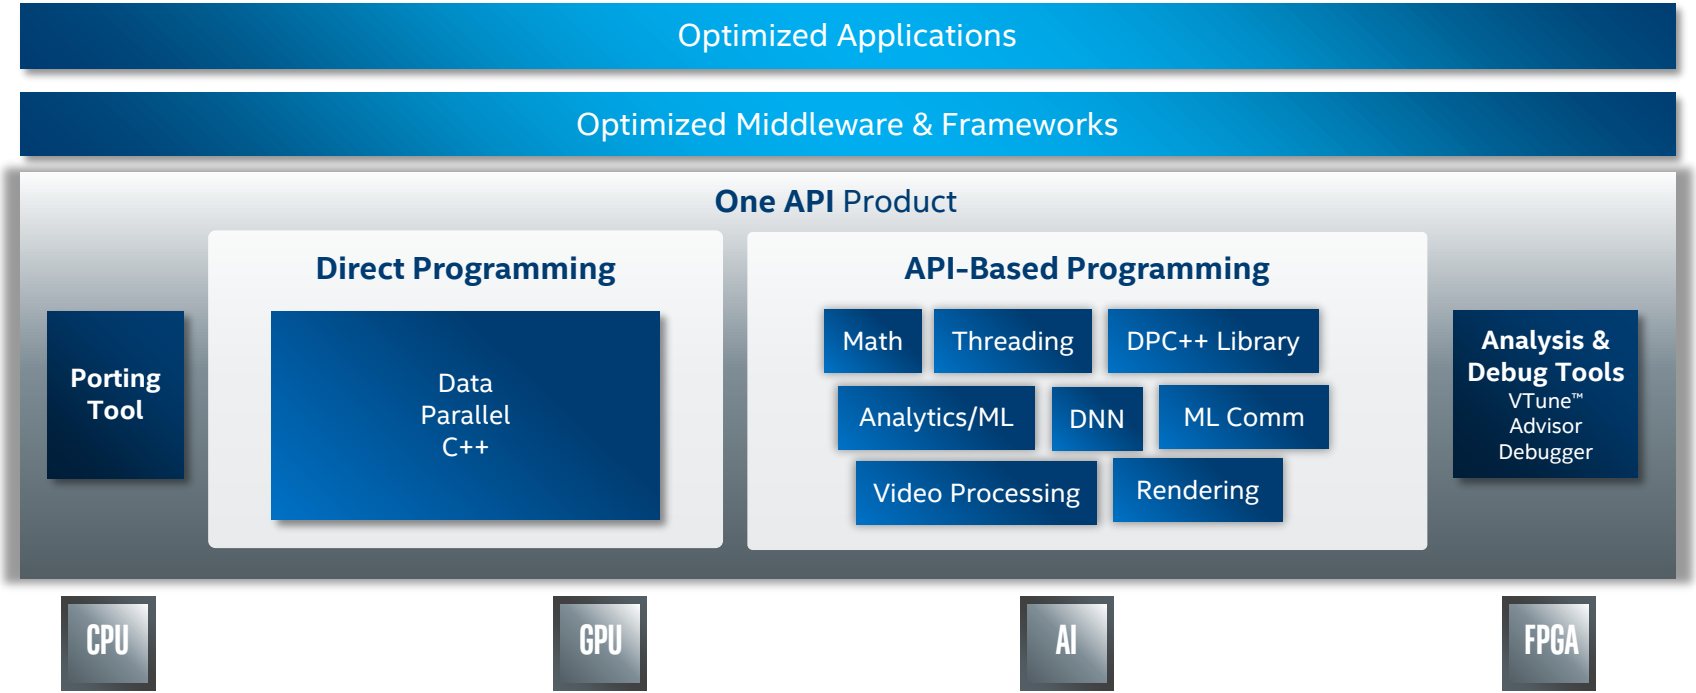

Some capabilities may differ per architecture.

Agenda

Intel Xeon Family Update

Intel Software Development Tools

- Intel Parallel Studio XE (IPS XE) Tool Suites
- News on the IPS XE 2020 Edition
- Upcoming **oneAPI** development tools concept

Intel optimized AI Solutions and Directions

ARTIFICIAL INTELLIGENCE

is the ability of machines to learn from experience, without explicit programming, in order to perform cognitive functions associated with the human mind

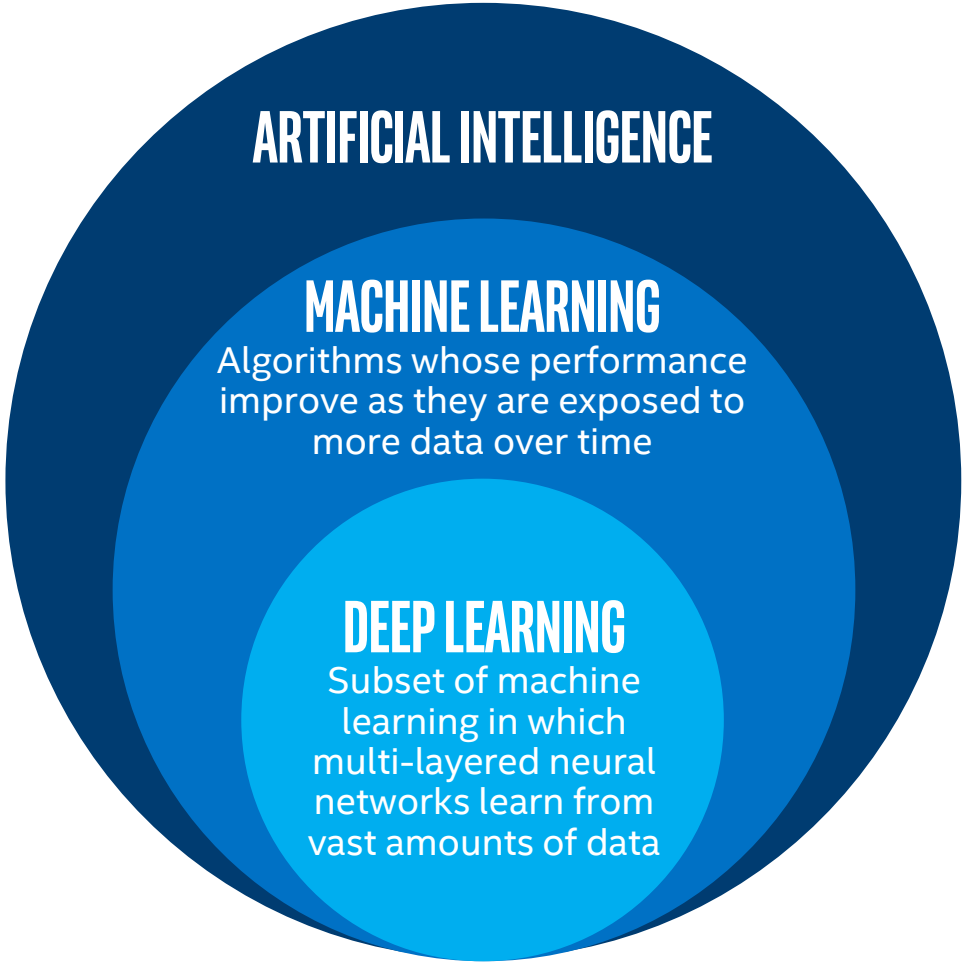

ARTIFICIAL INTELLIGENCE

SOLUTIONS
Solution Architects

AI Solutions Catalog (Public & Internal)

Platforms | Finance | Healthcare | Energy | Industrial | Transport | Retail | Home | More...

TOOLKITS
App Developers

DEEP LEARNING DEPLOYMENT

OpenVINO™ †
Open Visual Inference & Neural Network Optimization toolkit for inference deployment on CPU, processor graphics, FPGA & VPU using TF, Caffe & MXNet**

Intel® Movidius™ SDK
Optimized inference deployment for all Intel® Movidius™ VPUs using TensorFlow & Caffe**

DEEP LEARNING

COMING SOON!

Intel® Deep Learning Studio‡
Open-source tool to compress deep learning development cycle

LIBRARIES
Data Scientists

MACHINE LEARNING LIBRARIES

Python
- Scikit-learn
- Pandas
- NumPy

R
- Cart
- Random Forest
- e1071

Distributed
- MLlib (on Spark)
- Mahout

DEEP LEARNING FRAMEWORKS

COMING SOON!

Now optimized for CPU

TensorFlow* | MXNet* | Caffe* | BigDL/Spark*

Optimizations in progress

Caffe2* | PyTorch* | PaddlePaddle*

FOUNDATION
Library Developers

ANALYTICS, MACHINE & DEEP LEARNING PRIMITIVES

Python
Intel distribution optimized for machine learning

DAAL
Intel® Data Analytics Acceleration Library (for machine learning)

MKL-DNN **clDNN**
Open-source deep neural network functions for CPU, processor graphics

DEEP LEARNING GRAPH COMPILER

Intel® nGraph™ Compiler (Alpha)
Open-sourced compiler for deep learning model computations optimized for multiple devices (CPU, GPU, NNP) using multiple frameworks (TF, MXNet, ONNX)

HARDWARE
IT System Architects

AI FOUNDATION

ATOM | CORE 8th Gen | XEON PLATINUM

Data Center
Edge
Device

DEEP LEARNING ACCELERATORS

NERVANA (COMING 2019) | STRATIX 10 | ARRIA 10 | MOVIDIUS | GNA

NNP L-1000

Inference

† Formerly the Intel® Computer Vision SDK
*Other names and brands may be claimed as the property of others.
All products, computer systems, dates, and figures are preliminary based on current expectations, and are subject to change without notice.

AI.INTEL.COM

Components Comparison : Intel MKL-DNN vs Intel MKL

MKL-DNN (Open Source)

- Convolution
- Pooling
- ReLU
- Inner Product
- Normalization

MKL (Math Kernel Library)

Linear Algebra	Fast Fourier Transforms	Vector Math	Summary Statistics	And More…
• BLAS • LAPACK • ScaLAPACK • Sparse BLAS • Sparse Solvers • Iterative • PARDISO* • Cluster Sparse Solver	• Multidimensional • FFTW interfaces • Cluster FFT	• Trigonometric • Hyperbolic • Exponential • Log • Power • Root • Vector RNGs	• Kurtosis • Variation coefficient • Order statistics • Min/max • Variance-covariance	• Splines • Interpolation • Trust Region • Fast Poisson Solver

Data Transformation & Analysis Algorithms

Intel® Data Analytics Acceleration Library

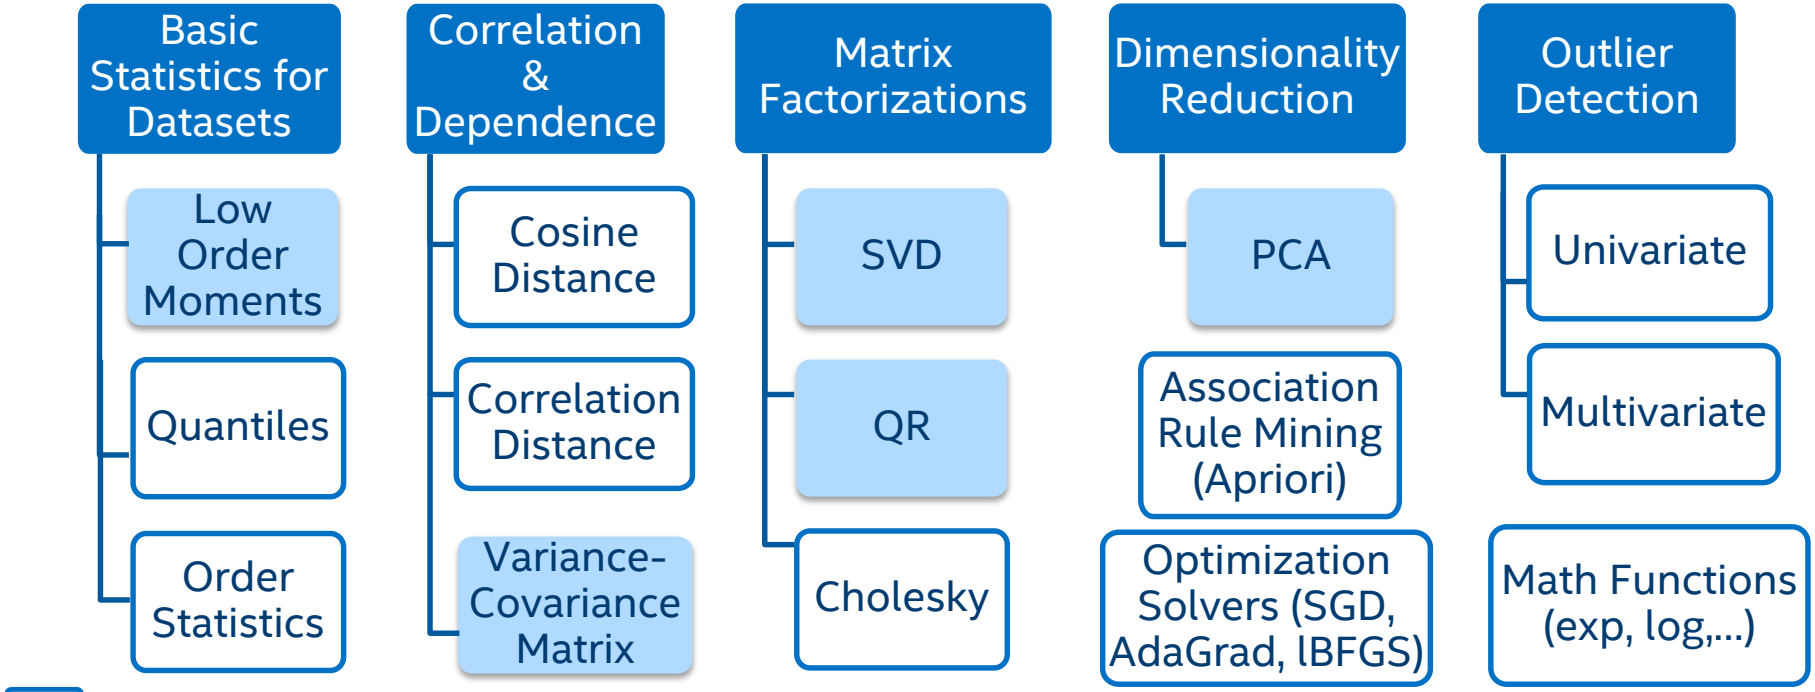

Algorithms supporting batch processing

Algorithms supporting batch, online and/or distributed processing

Machine Learning Algorithms

Intel® Data Analytics Acceleration Library

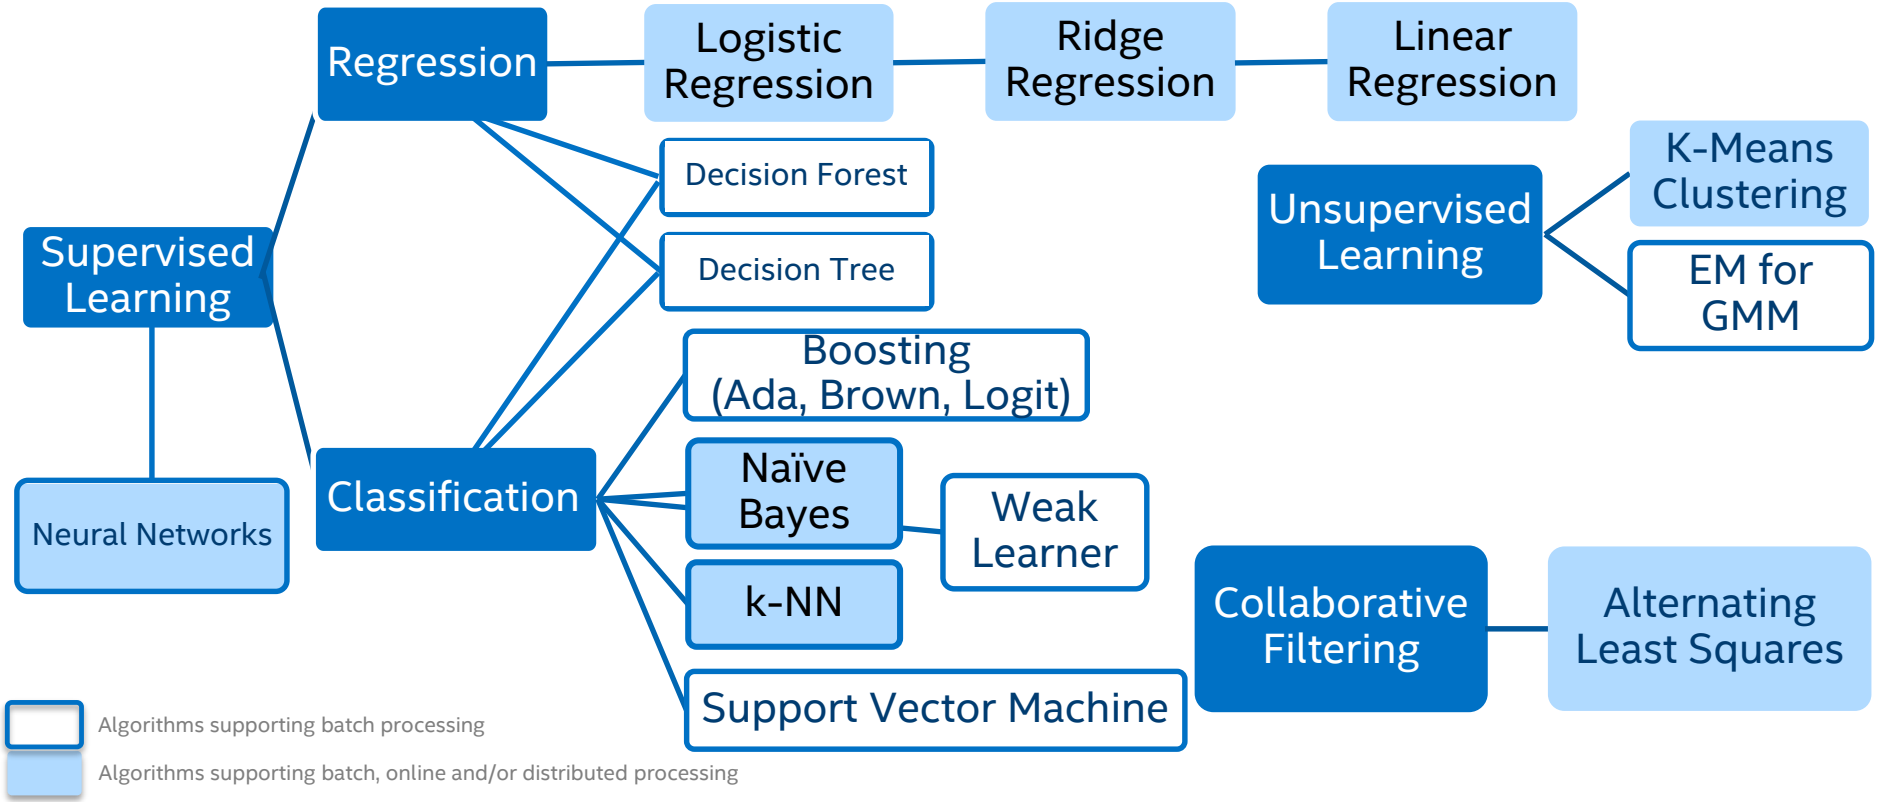

Faster Python* with Intel® Distribution for Python

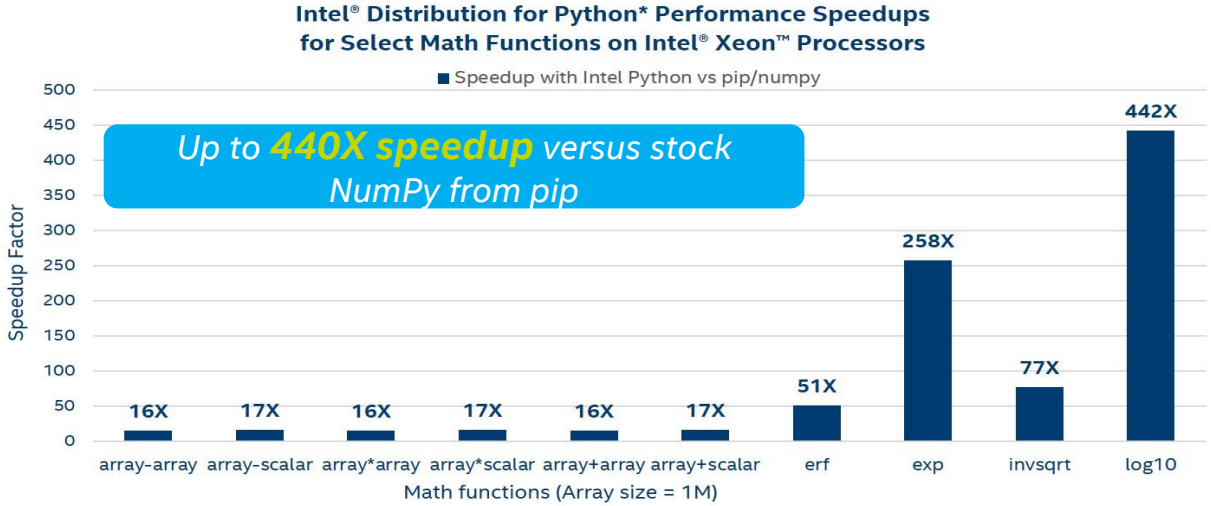

Configuration: Hardware: Intel® Xeon® CPU E5-2699 v4 @ 2.20GHz (2 sockets, 22 cores per socket, 1 thread per core – HT is off), 256GB DDR4 @ 2400MHz. Software: Stock: CentOS Linux release 7.3.1611 (Core), python 3.6.2, pip 9.0.1, numpy 1.13.1, scipy 0.19.1, scikit-learn 0.19.0. Intel® Distribution for Python* 2018 Gold: mkl 2018.0.0 intel_4, daal 2018.0.0.20170814, numpy 1.13.1 py36_intel_15, openmp 2018.0.0 intel_7, scipy 0.19.1 np113py36_intel_11, scikit-learn 0.18.2 np113py36_intel_3

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more information go to http://www.intel.com/performance. Benchmark Source: Intel Corporation.

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. *Other brands and names are the property of their respective owners.

Learn More: software.intel.com/distribution-for-python

Intel-Optimized AI Frameworks

Popular DL Frameworks are now optimized for CPU!

See installation guides at ai.intel.com/framework-optimizations/

More under optimization: Caffe2 PYTORCH PaddlePaddle

SEE ALSO: Machine Learning Libraries for Python (Scikit-learn, Pandas, NumPy), R (Cart, randomForest, e1071), Distributed (MLlib on Spark, Mahout)
**Limited availability today*
Other names and brands may be claimed as the property of others.

INTEL® XEON® PROCESSOR PLATFORM PERFORMANCE

Hardware plus optimized software

INFERENCE THROUGHPUT

198x

Intel® Xeon® Platinum 8180 Processor
higher Intel optimized Caffe GoogleNet v1 with Intel® MKL
inference throughput compared to
Intel® Xeon® Processor E5-2699 v3 with BVLC-Caffe

TRAINING THROUGHPUT

127x

Intel® Xeon® Platinum 8180 Processor
higher Intel Optimized Caffe AlexNet with Intel® MKL
training throughput compared to
Intel® Xeon® Processor E5-2699 v3 with BVLC-Caffe

Inference and training throughput uses FP32 instructions

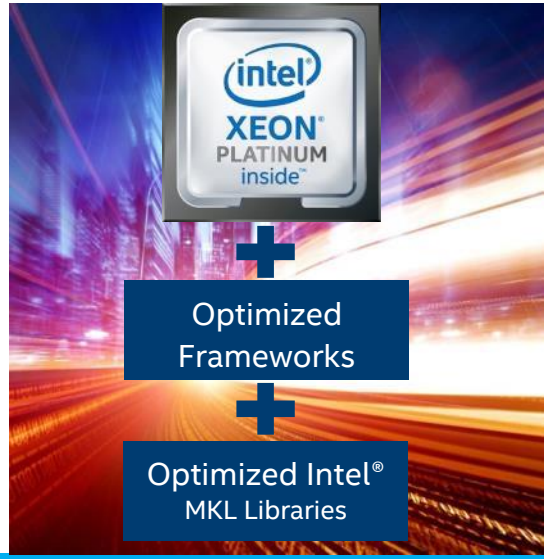

Deliver significant AI performance with hardware and software optimizations on Intel® Xeon® Scalable Family

Up to 191X Intel® Xeon® Platinum 8180 Processor higher Intel optimized Caffe Resnet50 with Intel® MKL inference throughput compared to Intel® Xeon® Processor E5-2699 v3 with BVLC-Caffe
Up to 93X Intel® Xeon® Platinum 8180 Processor higher Intel optimized Caffe Resnet50 with Intel® MKL training throughput compared to Intel® Xeon® Processor E5-2699 v3 with BVLC-Caffe

Use your Intel Architecture based HPC Infrastructure ---> Shorten Deep Learning Training Time

HPC

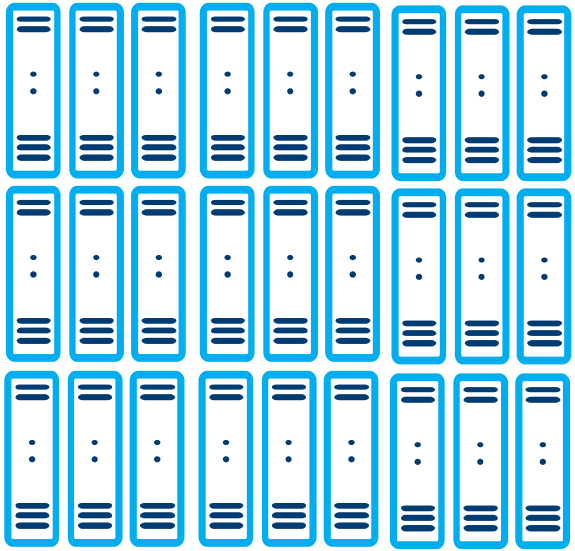

AI

AI (ML & DL) SOFTWARE STACK FOR INTEL® PROCESSORS

Intel MKL

Intel MKL-DNN

Intel Processors

Distributed DL Training

Uber's open source Distributed training framework for TensorFlow

Intel MPI

Intel MLSL

Intel Processors

More details on Intel and AI at GridKa School

Join the Intel hands-on Training : Thursday afternoon at 13:15

Title : Enhance Machine Learning Performance with Intel® Software tools

1st Session:

- Introduction of Intel tools for ML and DL
  - DLBoost, VNNI instructions ; Intel distribution for Python

2nd session:

- Classical ML
  - Numpy and MKL, K-Means, clustering and DAAL4PY Distributed ML algorithms

3rd session:

- Intel MKL for Deep Neural Network Intel optimized Framework and Tensorflow
- Distributed Tensorflow with Horovod

THANK YOU

Legal Disclaimer & Optimization Notice

Performance results are based on testing as of August 2017 to September 2018 and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit www.intel.com/benchmarks.

INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS". NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO THIS INFORMATION INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

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

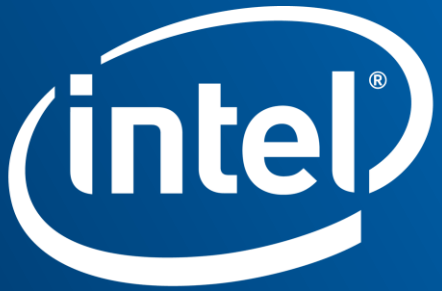

Software

INTEL® OPTANE™ DC PERSISTENT MEMORY

THE BEST OF BOTH WORLDS WITH INTEL® OPTANE™ DC PERSISTENT MEMORY

DRAM ATTRIBUTES

Performance comparable to DRAM at *low latencies*[1]

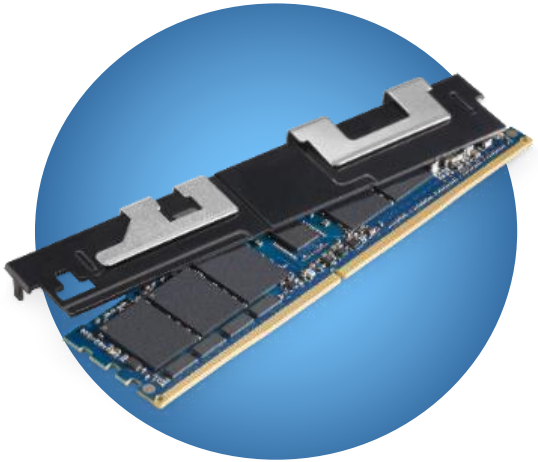

NAND SSD ATTRIBUTES

Data persistence with higher capacity than DRAM[2]

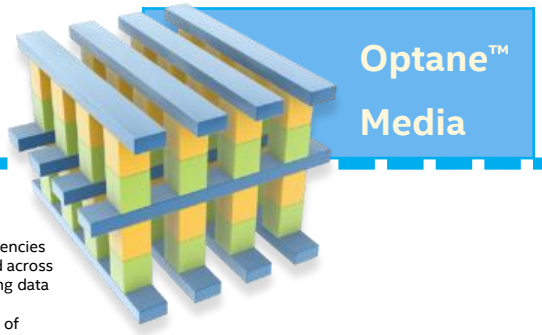

1. "*Fast performance comparable to DRAM*" - Intel persistent memory is expected to perform at latencies near DDR4 DRAM. Benchmarks and proof points forthcoming. "*low latencies*" - Data transferred across the memory bus causes latencies to be orders of magnitude lower when compared to transferring data across PCIe or I/O bus' to NAND/Hard Disk. Benchmarks and proof points forthcoming.
2. Intel persistent memory offers 3 different capacities – 128GB, 256GB, 512GB. Individual DIMMs of DDR4 DRAM max out at 256GB.

Latency Estimates for different Memory and Storage Devices

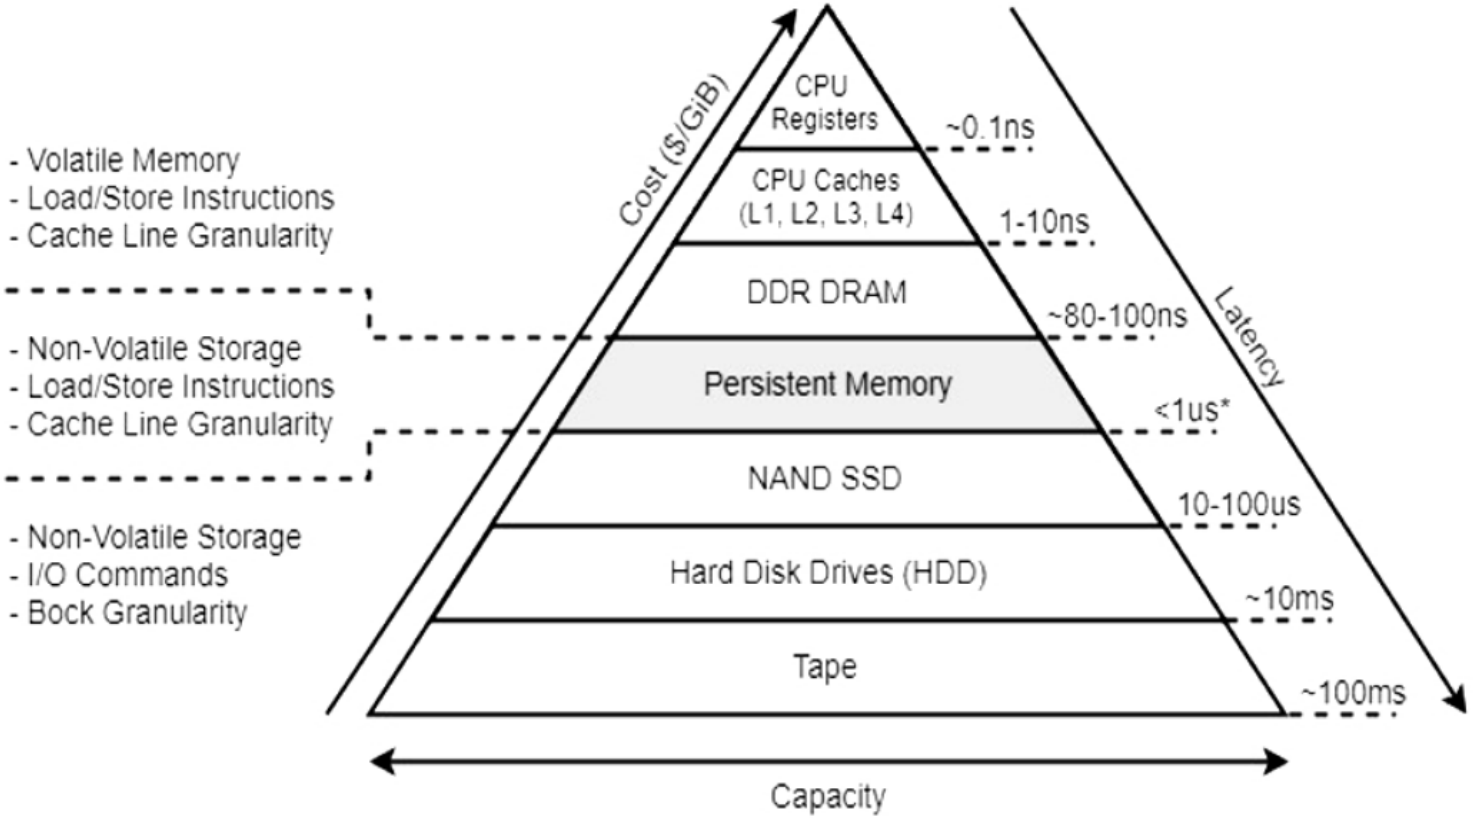

Distributed TensorFlow™ Compare

Distributed Tensorflow with Parameter Server

With Parameter Server →

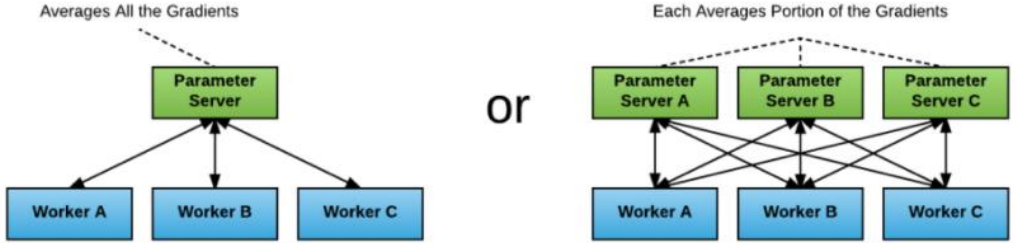

The parameter server model for distributed training jobs can be configured with different ratios of parameter servers to workers, each with different performance profiles.

HOROVOD

Uber's open source Distributed training framework for TensorFlow

No Parameter Server →

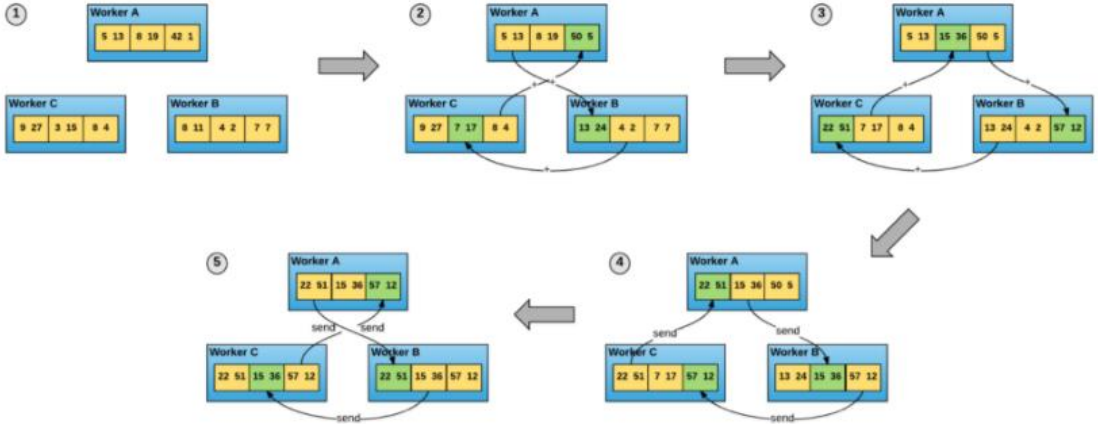

The ring all-reduce algorithm allows worker nodes to average gradients and disperse them to all nodes without the need for a parameter server.

Source: https://eng.uber.com/horovod/

UNSUPERVISED LEARNING EXAMPLE

MACHINE LEARNING

- Regression
- Classification
- **Clustering**
- Decision Trees
- Data Generation

DEEP LEARNING

- Image Processing
- Speech Processing
- Natural Language Processing
- Recommender Systems
- Adversarial Networks
- Reinforcement Learning

MACHINE LEARNING (CLUSTERING)

An 'Unsupervised Learning' Example

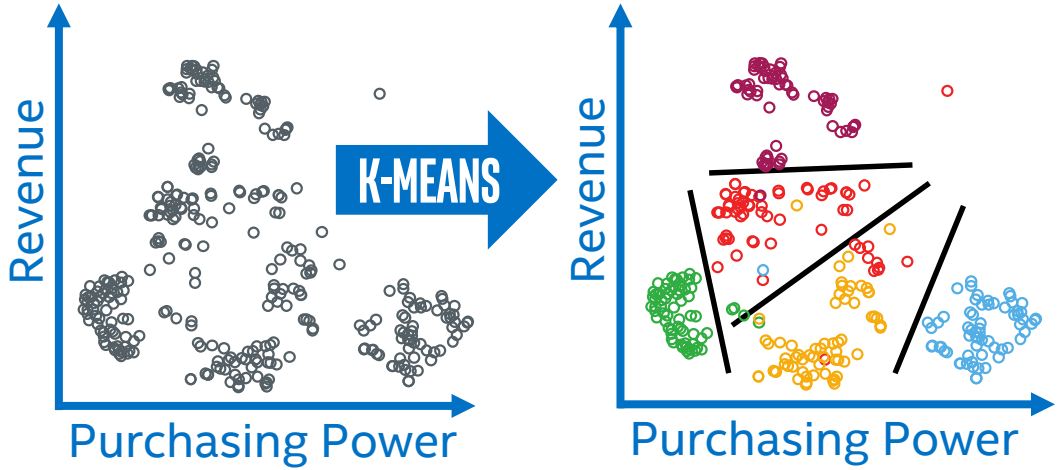

Choose the right AI approach for your challenge

SUPERVISED LEARNING EXAMPLE

MACHINE LEARNING

- Regression
- Classification
- Clustering
- Decision Trees
- Data Generation

DEEP LEARNING

- Image Processing
- Speech Processing
- Natural Language Processing
- Recommender Systems
- Adversarial Networks
- Reinforcement Learning

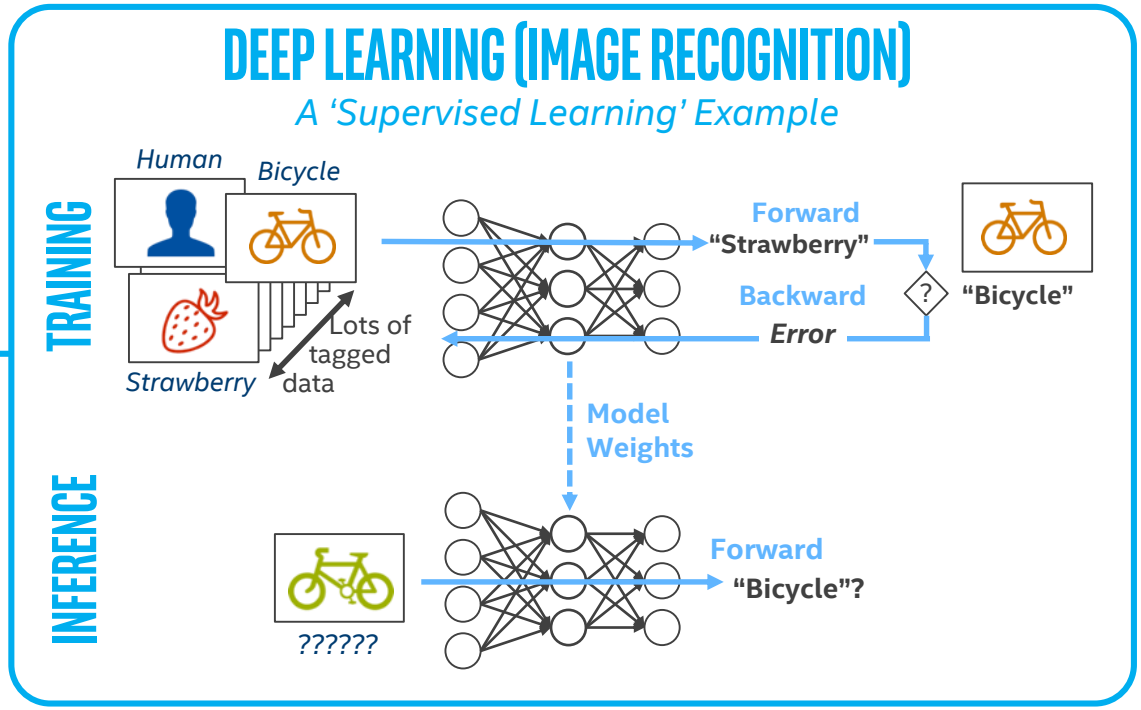

Choose the right AI approach for your challenge

REINFORCEMENT LEARNING EXAMPLE

MACHINE LEARNING

- Regression
- Classification
- Clustering
- Decision Trees
- Data Generation

DEEP LEARNING

- Image Processing
- Speech Processing
- Natural Language Processing
- Recommender Systems
- Adversarial Networks
- Reinforcement Learning

DEEP LEARNING (REINFORCEMENT LEARNING)

A 'Deep Reinforcement Learning' Example

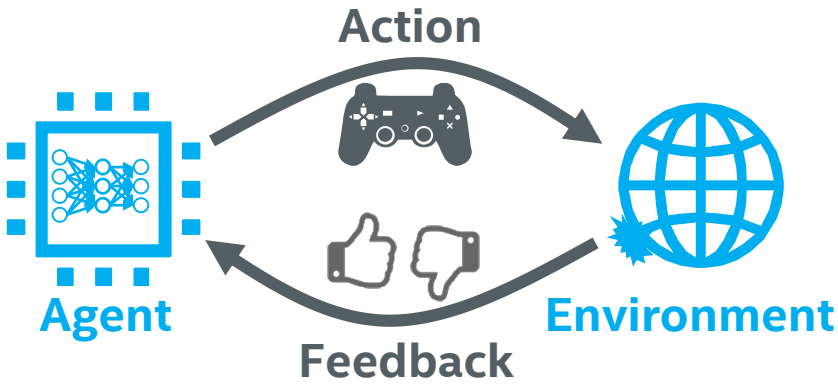

Choose the right AI approach for your challenge

Intel-Optimized Frameworks: How To Get?

Check out our intel.ai for the framework optimizations page

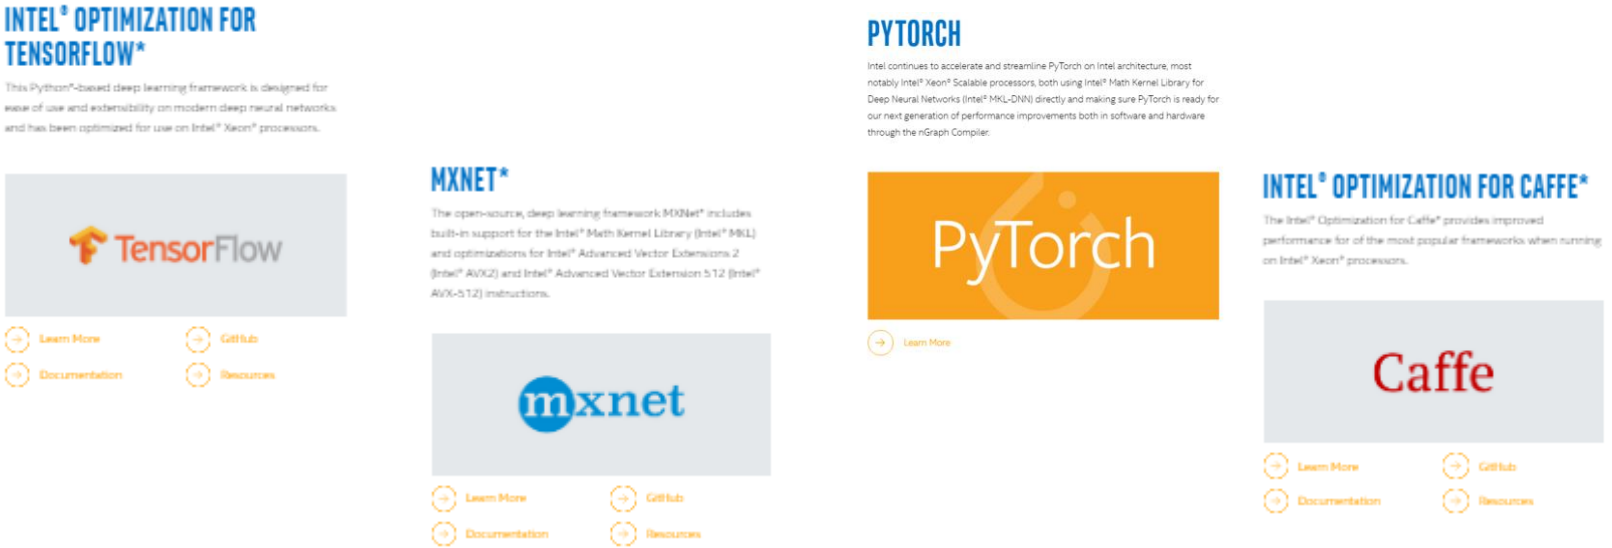

Installation example for optimized Tensorflow / Keras for a Medial Sample application
https://docs.google.com/document/d/1AAsWLSfYBx-pYfvzFQxOYi4Z1zx3TAEB5534s61Lf8w/edit?usp=sharing

Topology Examples : https://www.intel.ai/framework-optimizations