
INTEL®
INNOVATION
DAY

КОМПОНЕНТЫ ЭФФЕКТИВНОГО КОДА – ВЫСОКООПТИМИЗИРОВАННЫЕ БИБЛИОТЕКИ INTEL

Татьяна Бышева,
старший инженер по разработке программного обеспечения

Agenda

- Introduction
- Intel® Integrated Performance Primitives (IPP) overview
- Intel® Math Kernel Library (MKL) overview
- Intel® Data Analytics Acceleration Library (DAAL) overview
- Summary

INTRODUCTION

Introduction

Take advantage of powerful and award-winning performance libraries that optimize your code and shorten development time. These libraries are offered for free as part of Intel's mission to support innovation and impressive performance on Intel® architecture.

Intel® Integrated Performance Primitives

Gain a competitive performance advantage with this library that offers image, signal, compression, and cryptography functions for multiple operating systems and platforms.

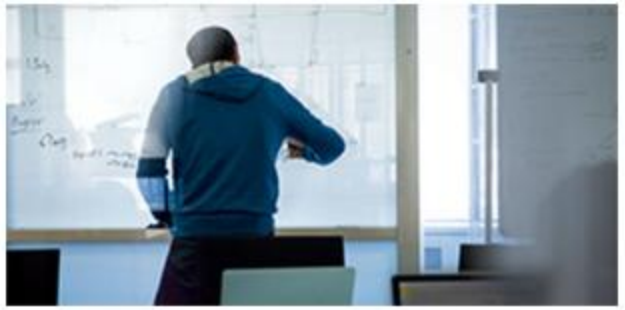

Intel® Data Analytics Acceleration Library

Boost machine learning and big-data analytics with this industry-leading, easy-to-use performance library. Features include highly tuned functions for analytics performance across the spectrum of Intel® architecture devices.

Intel® Math Kernel Library

This popular, fast math library for Intel® and other compatible processors features highly optimized, threaded, and vectorized functions to maximize performance on each processor family.

and more…

INTEL® INTEGRATED PERFORMANCE PRIMITIVES

Intel® IPP: Your Building Blocks for Image, Signal & Data Processing Applications

What is Intel® IPP?

Intel IPP provides developers with ready-to-use, processor- optimized functions to accelerate ***Image & Signal processing, Data Compression & Cryptography computation tasks***

Why should you use Intel® IPP?

- High Performance
- Easy to use API's
- Faster Time To Market (TTM)
- Production Ready
- Cross-Platform API
- Small footprint

How to get Intel® IPP?

- Intel® Parallel Studio XE
- Intel® System Studio
- Free Tools Program
- IPP Crypto open source
- YUM, APT-GET and Conda packages

Optimized for

Supports

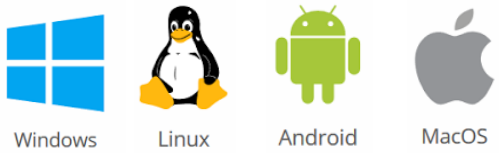

Addresses

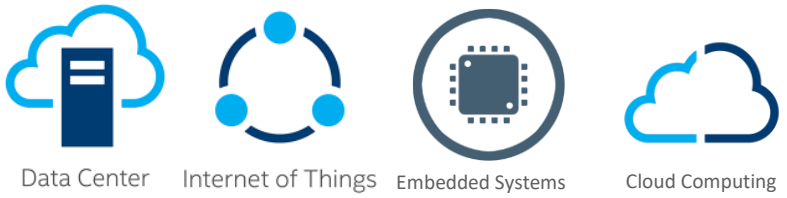

Image Processing Uses

- Medical Imaging
- Computer Vision
- Digital Surveillance
- ADAS
- Automated Sorting
- Biometric Identification
- Visual Search

Signal Processing Uses

- Games (sophisticated audio content or effects)
- Echo cancellation
- Telecommunications
- Energy

Data Compression & Cryptography Uses

- Data centers
- Enterprise data management
- ID verification
- Smart Cards/wallets
- Electronic Signature
- Information security/cybersecurity

What's Inside Intel® Integrated Performance Primitives

High Performance , Easy-to-Use & Production Ready APIs

Image Processing

Computer Vision

Color Conversion

Image Domain

Signal Processing

Vector Math

Signal Domain

Data Compression

Cryptography

String Processing

Data Domain

Intel® Architecture Platforms

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

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

OpenCV Optimization

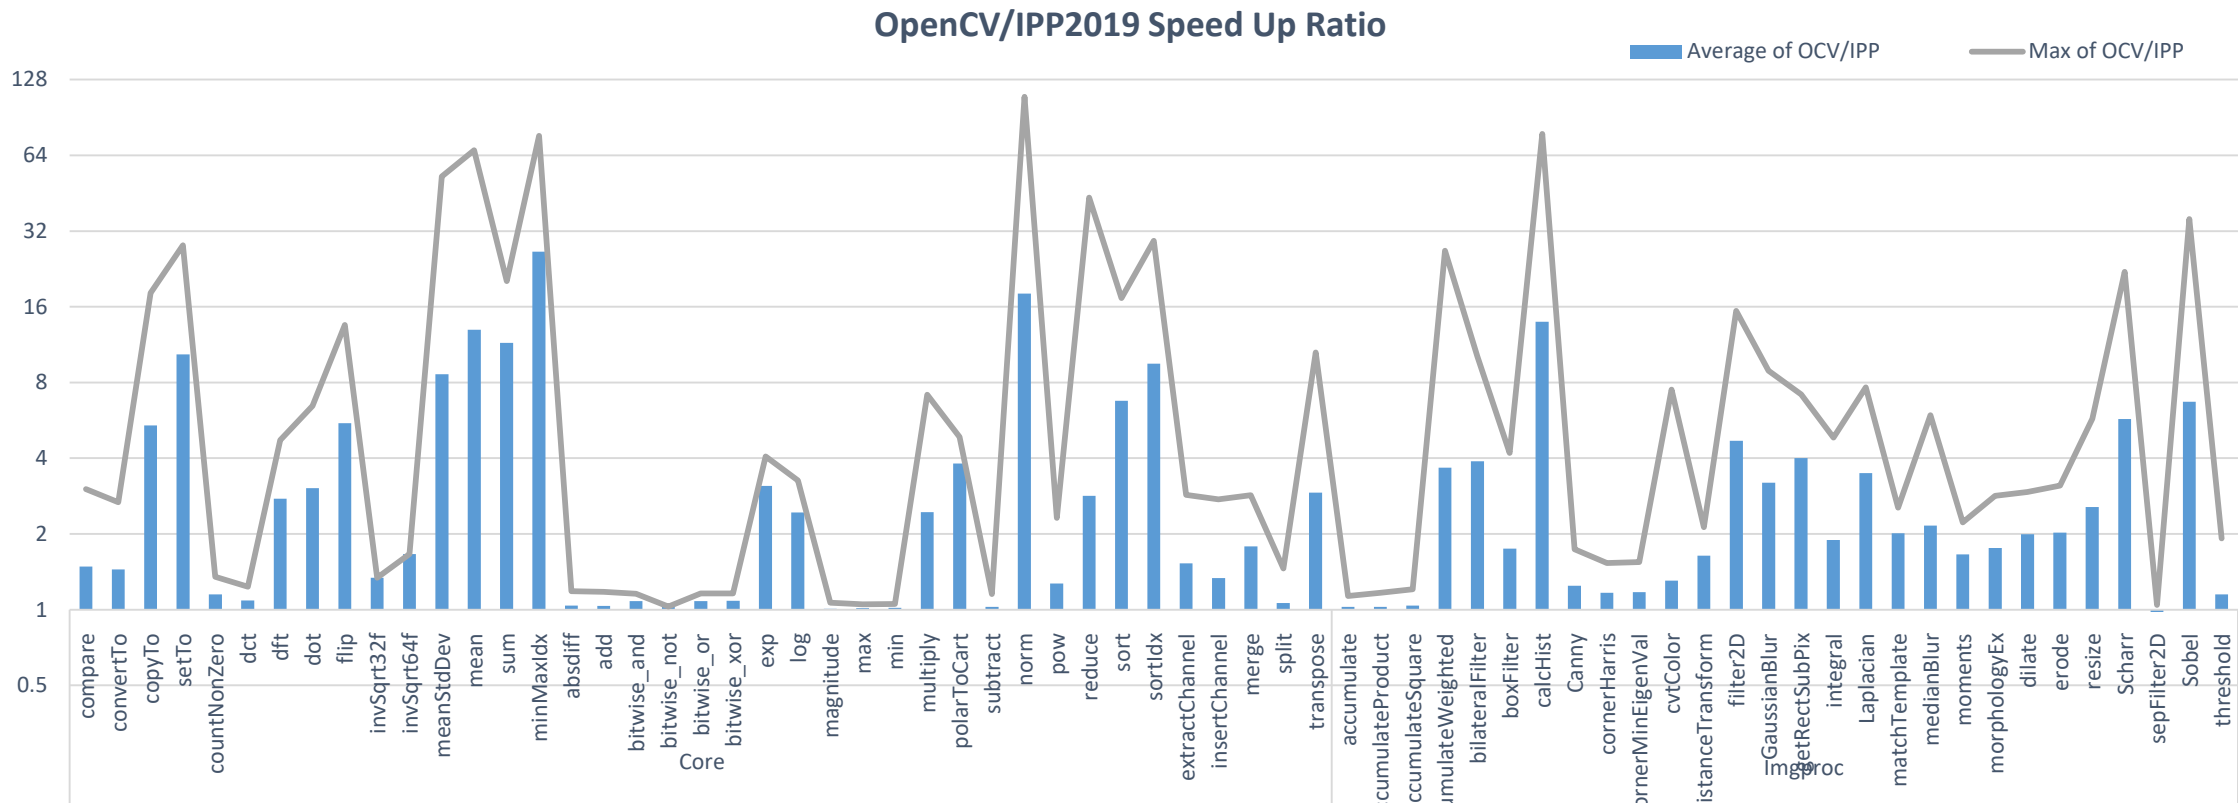

Performance results are based on testing as of **July 18, 2018** and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure.

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, see Performance Benchmark Test Disclosure.

Testing by Intel as of July 18, 2018. Configuration: Intel® Xeon® Platinum 8168, 2.7 GHz, 2x24 cores; Intel® C++ compiler 18.0; 94 GB, L3=33 MB; Red Hat Enterprise Linux Server 7.2

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.

For more complete information about compiler optimizations, see our Optimization Notice.

Performance Improvement for Data Decompression

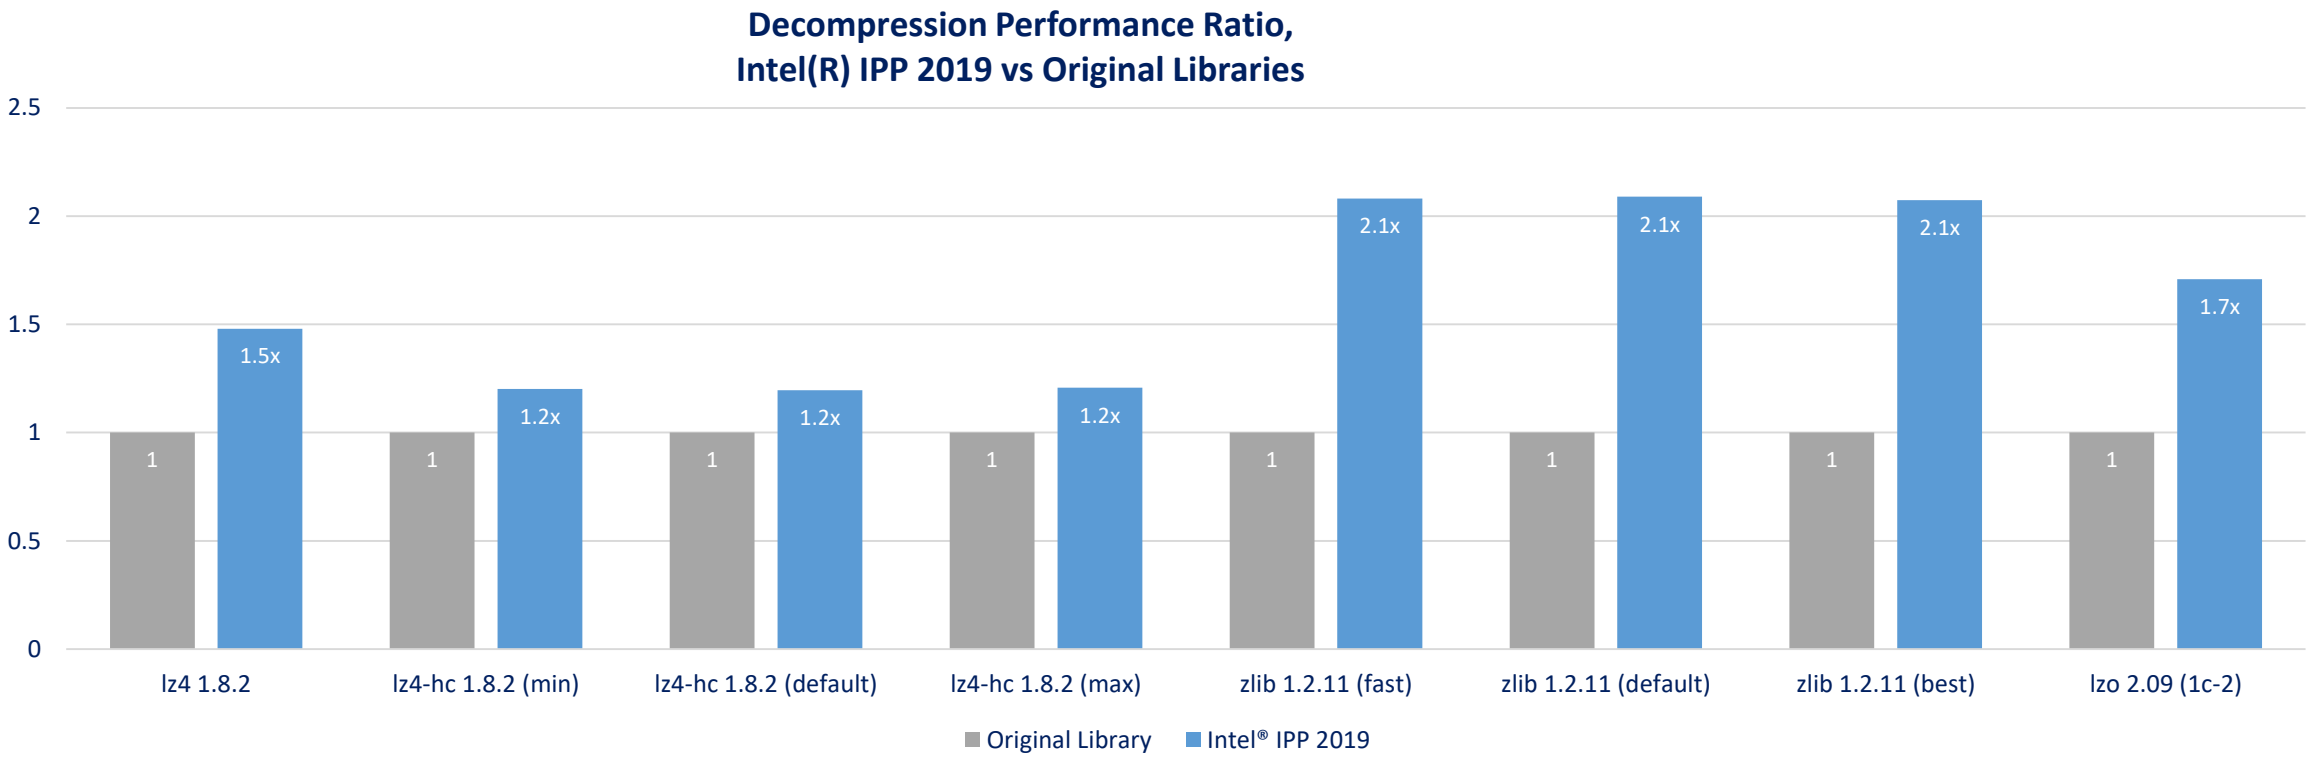

Performance results are based on testing as of **August 15, 2018** and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure.

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, see Performance Benchmark Test Disclosure.
Testing by Intel as of August 15, 2018. Configuration: Intel® Core™ i5-7600 CPU @3.50GHz, 4 cores, hyper-threading off; Cache: L1=32KB, L2=256KB, L3=6MB; Memory: 64GB; OS: RH EL Server 7.2

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.
For more complete information about compiler optimizations, see our Optimization Notice.

Performance Improvement for Data Compression/Decompression

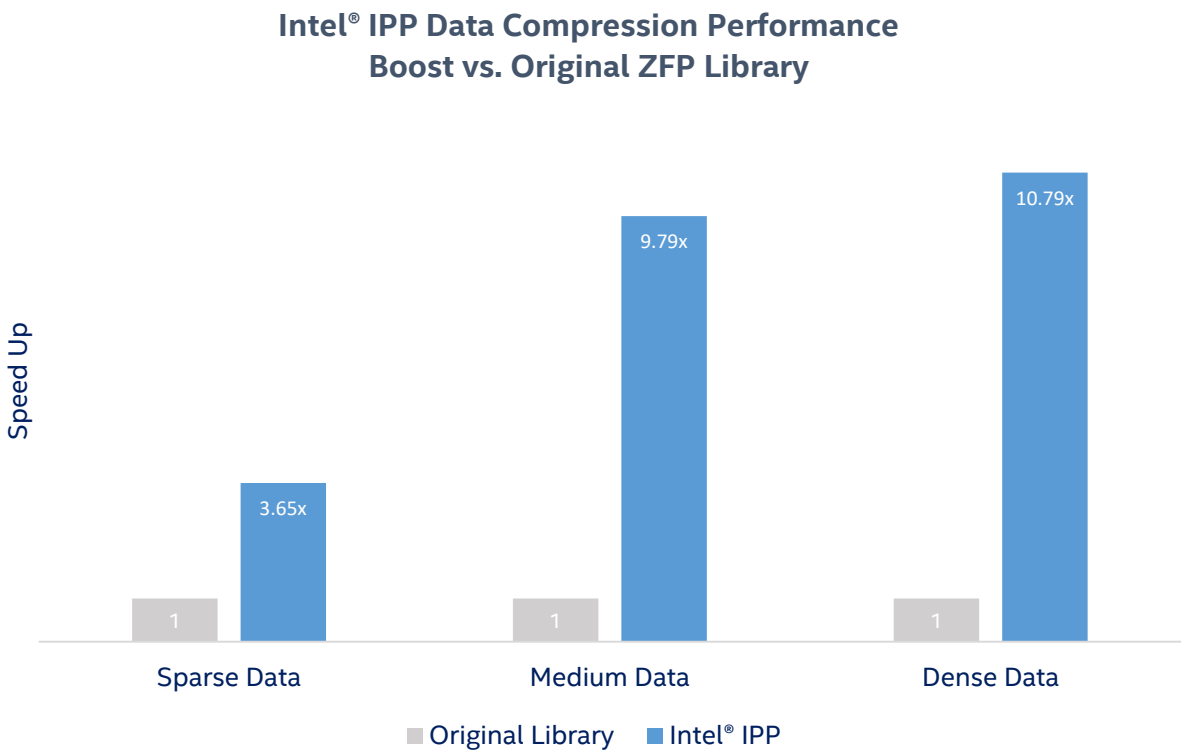

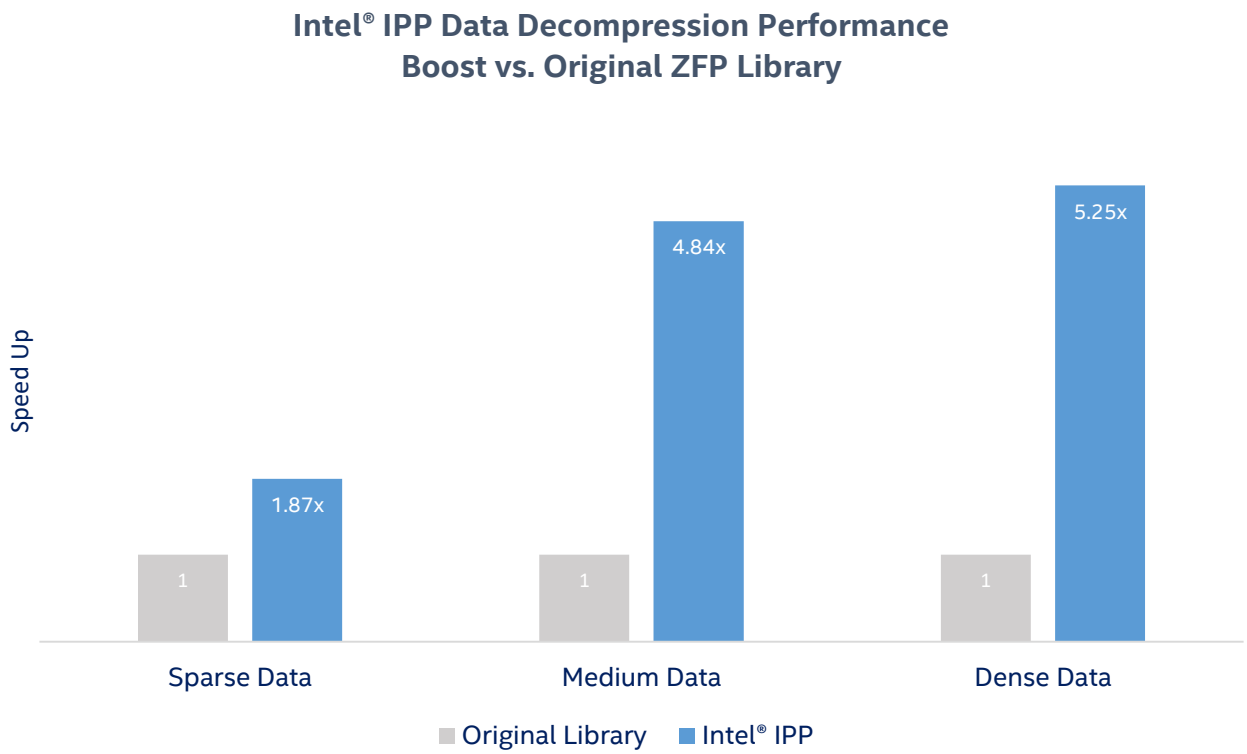

Performance results are based on testing as of **August 6, 2018** and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure.

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, see Performance Benchmark Test Disclosure.
Testing by Intel as of August 6, 2018. Configuration: Intel® Xeon® Platinum 8168, 2.7 GHz, 2x24 cores; Intel® C++ compiler 18.0; 94 GB, L3=33 MB; Red Hat Enterprise Linux Server 7.2

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.
For more complete information about compiler optimizations, see our Optimization Notice.

INTEL® MATH KERNEL LIBRARY

Faster, Scalable Code with Intel® Math Kernel Library

- Speeds computations for scientific, engineering, financial and machine learning applications by providing highly optimized, threaded, and vectorized math functions
- Provides key functionality for dense and sparse linear algebra (BLAS, LAPACK, PARDISO), FFTs, vector math, summary statistics, deep learning, splines and more
- Dispatches optimized code for each processor automatically without the need to branch code
- Optimized for single core vectorization and cache utilization
- Automatic parallelism for multi-core and many-core
- Scales from core to clusters
- Available at no cost and royalty free
- Great performance with minimal effort!

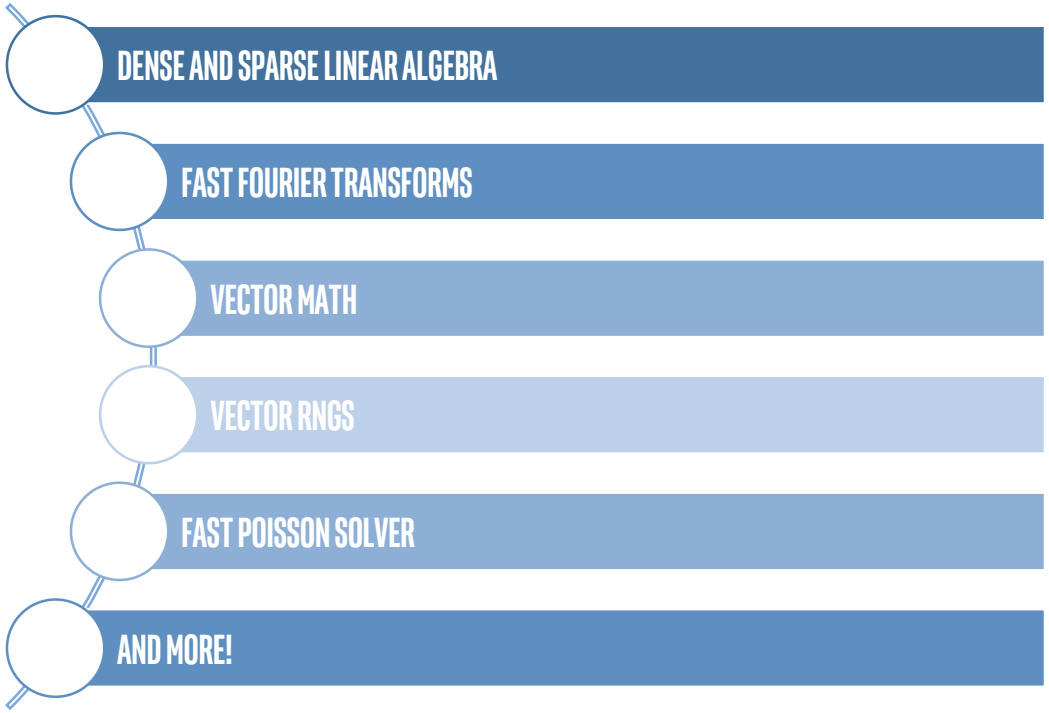

Available as standalone or as a part of Intel® Parallel Studio XE and Intel® System Studio

Intel® Architecture Platforms

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

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

What's Inside Intel® MKL

LINEAR ALGEBRA

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

FFTS

- Multidimensional
- FFTW interfaces
- Cluster FFT

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

AND MORE

- Splines
- Interpolation
- Trust Region
- Fast Poisson Solver

Intel® Math Kernel Library 2019 Gold vs Competitors on Intel® Xeon® Processor

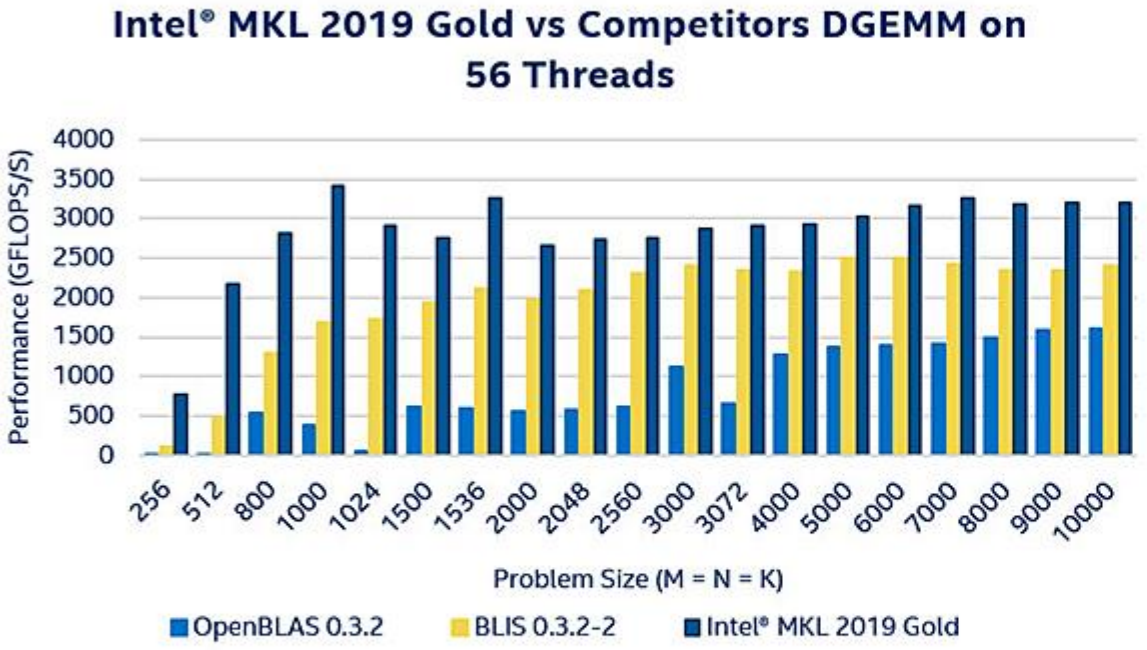

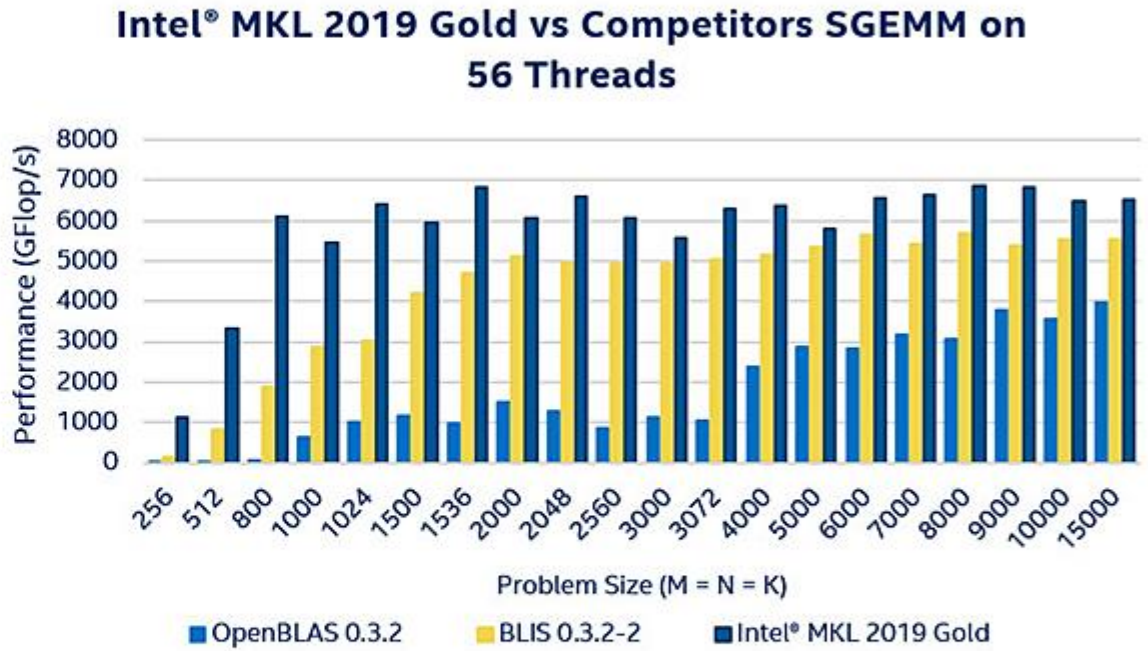

Performance results are based on testing as of **July 9, 2018** and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure.
Software and workloads used in performance tests may have been optimized for performance only on 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, see Performance Benchmark Test Disclosure.
Testing by Intel as of July 9, 2018. Configuration: Intel® Xeon® Platinum 8180 H0 205W 2x28@2.5GHz 192GB DDR4-2666

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.
For more complete information about compiler optimizations, see our Optimization Notice.

More cases

INTEL® DATA ANALYTICS ACCELERATION LIBRARY

Intel® Data Analytic Acceleration Library (Intel® DAAL)

Easy to use

- Highly optimized functions for classical machine learning and analytics performance
- Optimizes data ingestion together with algorithmic computation for highest analytics throughput
- Includes Python*, C++, and Java* APIs and connectors to popular data sources including Spark* and Hadoop*

Easy to get

- Free and open source community-supported
- Github: https://github.com/intel/daal
- Distributions via YUM, APT-GET, PIPY and Conda, Maven repositories
- Static and dynamic library
- Windows*, Linux*, OS X*

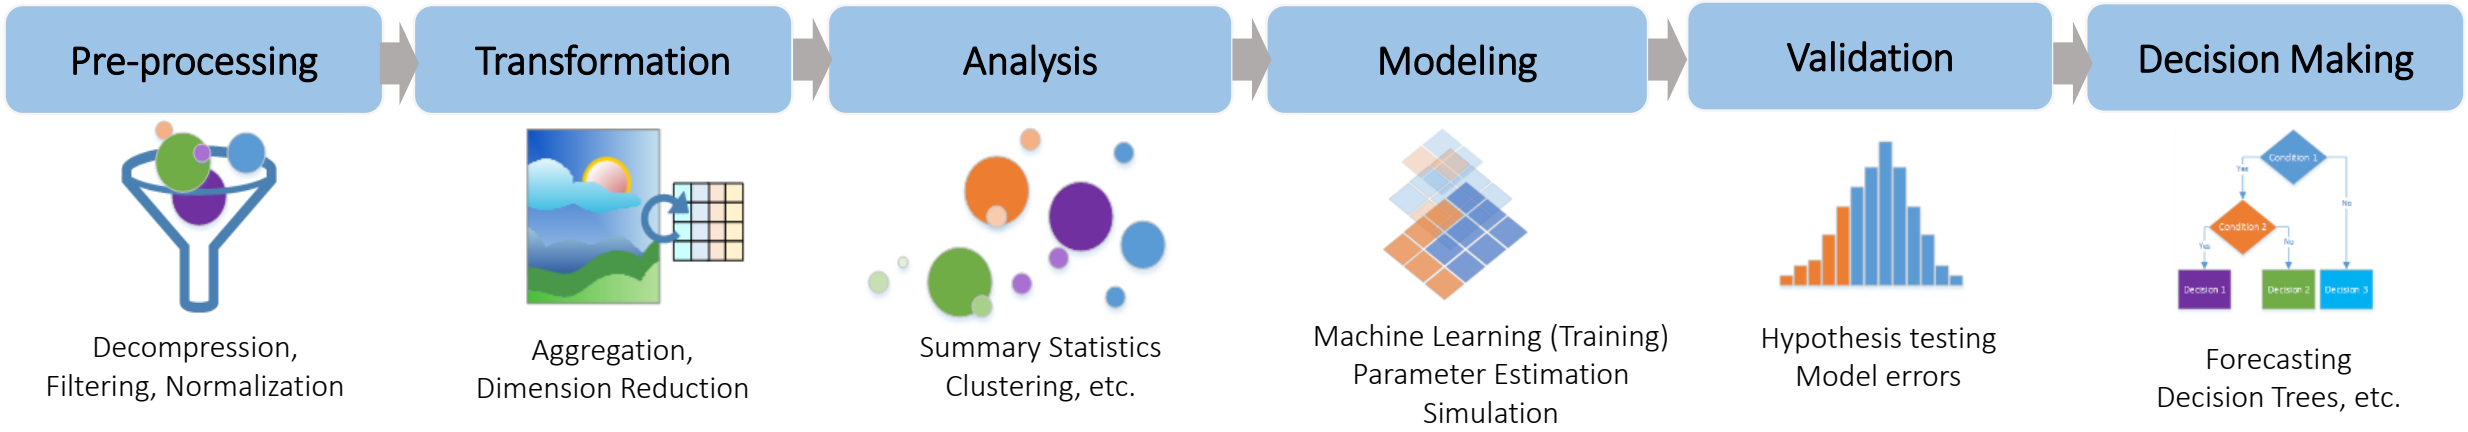

Learn More: software.intel.com/daal

Data Transformation and Analysis in Intel® DAAL

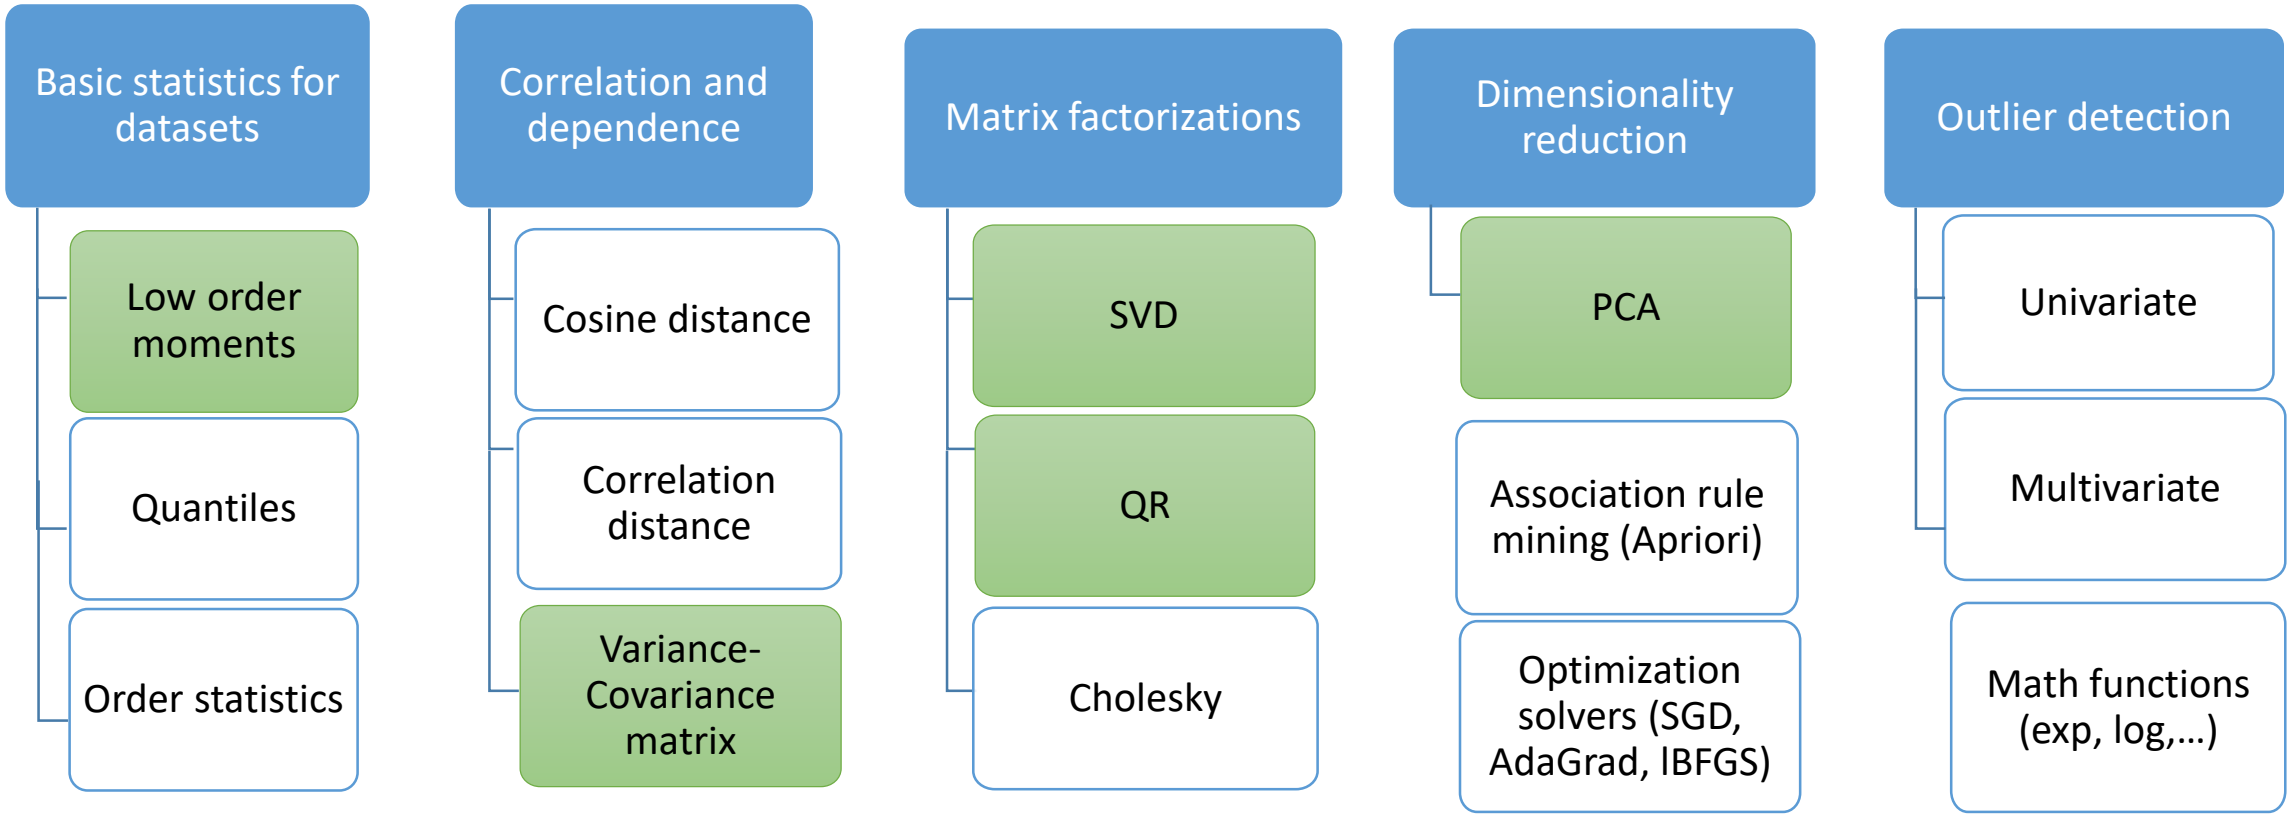

Algorithms supporting batch processing

Algorithms supporting batch, online and/or distributed processing

Intel® DAAL Algorithms - Machine Learning in Intel® DAAL

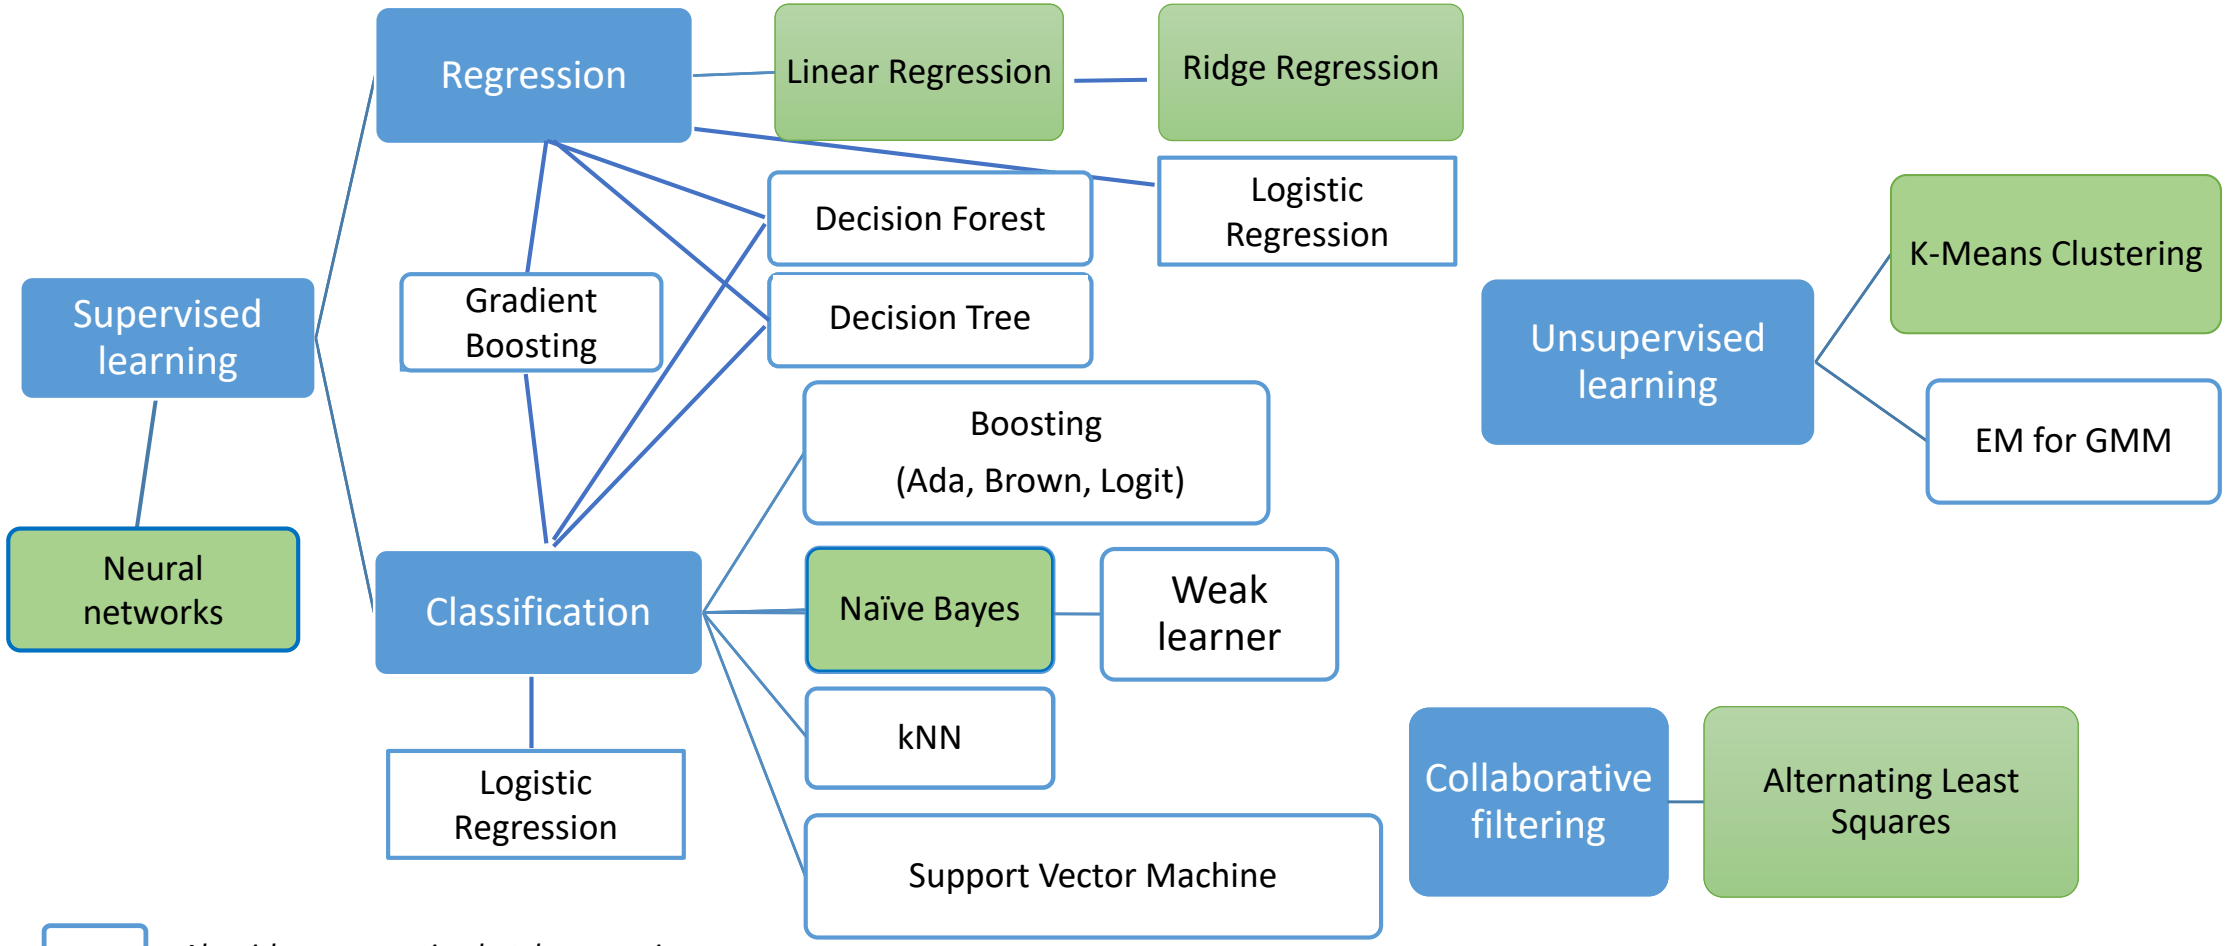

Intel® DAAL performance – Scikit-learn optimization

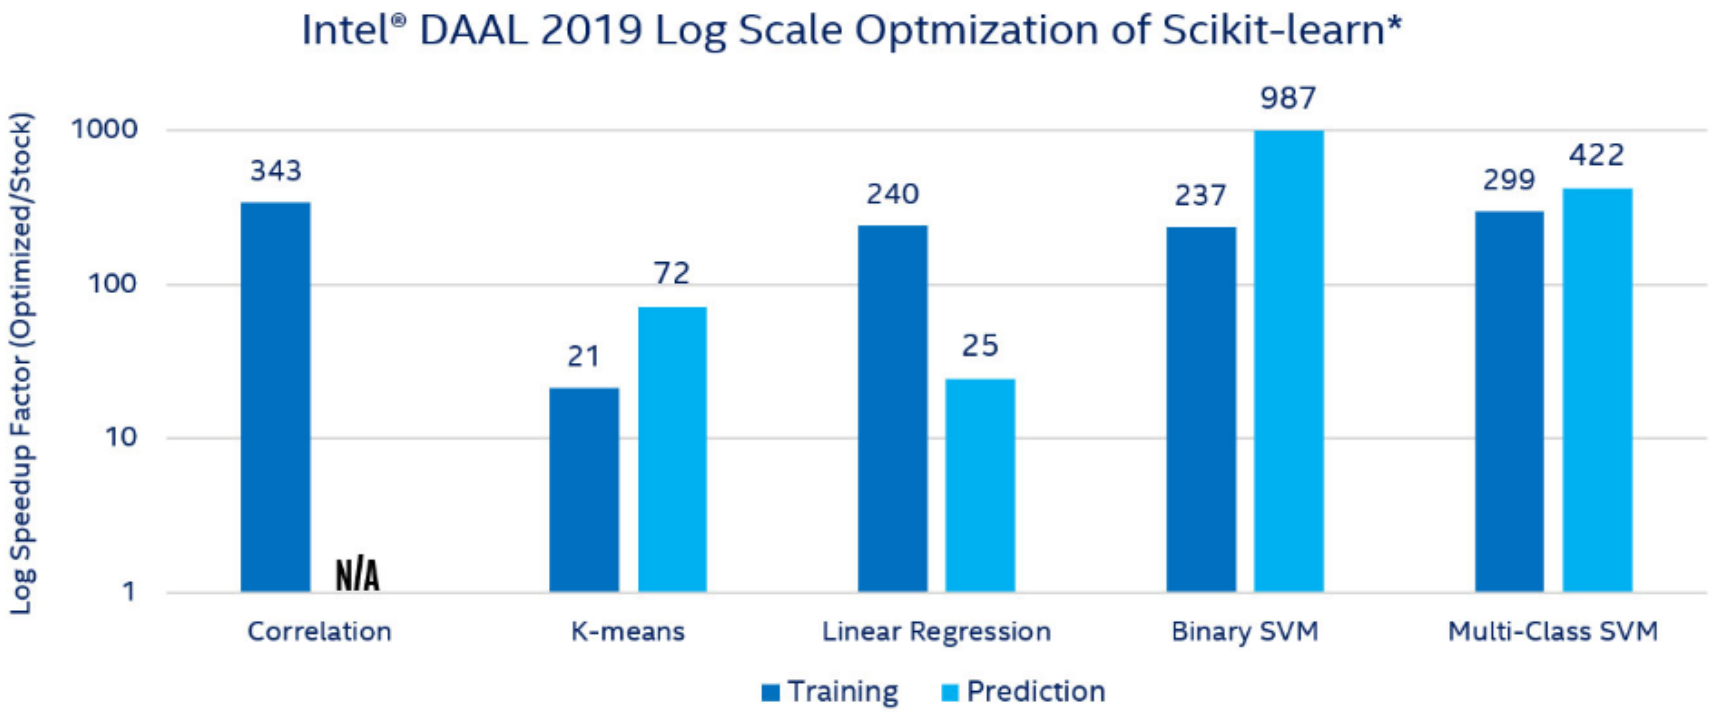

Performance results are based on testing as of July 9, 2018 and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure.

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, see Performance Benchmark Test Disclosure.

Testing by Intel as of July 9, 2018. Configuration: Intel® Xeon® Gold 6140 CPU, 2x18@2.30GHz, 256GB, 16x16gb DDR4-2666, Intel® Data Analytics Acceleration Library (Intel® DAAL 2019), Optimized: Scikit-learn*_intel 0.19.1, Numpy*_intel 1.14.3 Stock: Scikit-learn* 0.19.2, Numpy* 1.15.0, CentOS Linux 7.3.1611

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.
For more complete information about compiler optimizations, see our Optimization Notice.

SUMMARY

Summary

- Boost application performance with minimal effort
- Feature set is robust and growing
- Provide scaling from the core, to multicore, to manycore, and to clusters
- Automatic dispatching matches the executed code to the underlying processor
- Future processor optimizations included as well before processor ship

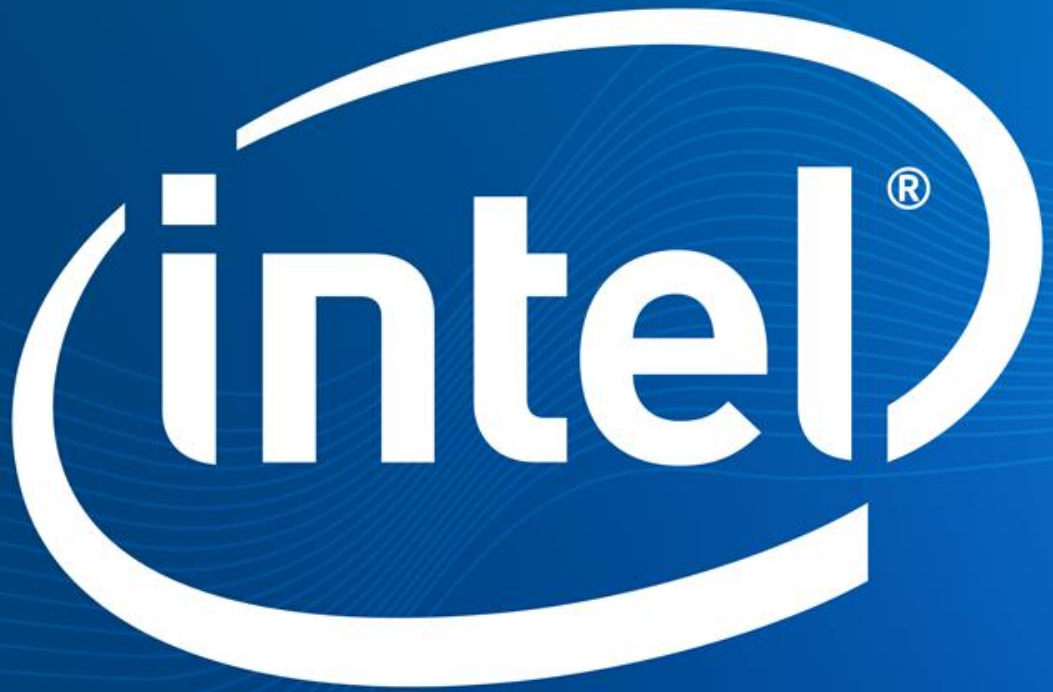
intel

BACKUP

Intel® IPP Benefits to Applications

Cloud and Server application
- Web image processing(resize, filtering, etc.)
- Web data compression and transferring, data encryption/decryption

Medical Images
- CT, MRI signal processing
- Medical image processing

Storage
- Storage data compression
- Storage data encryption/decryption

Print Imaging
- Image enhancement and correction
- Data compression

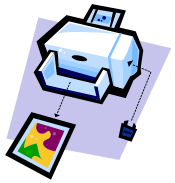

Digital Surveillance
- Computer vision
- Image recognition

Signal Processing
- Seismic data analysis, radar and sonar signal processing.

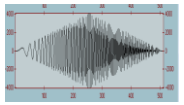

Machine Vision
- Image filtering, segmentation
- Edge detection, pattern recognition

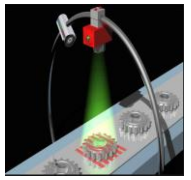

In-Vehicle Infotainment
- Image and audio data processing

Biometric Identification
- Biometric image and signal processing

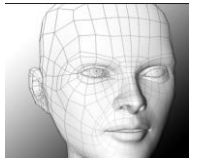

Visual Search
- Examining image content(color, shape, texture…)

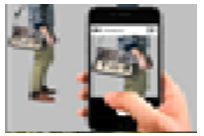

Communication
- Wireless communication single processing
- CRC and MIMO functions for communication.

And More
- Digital media, security, mobile……

Gets Good Performance with Intel® IPP

In popular apps like WeChat, QQ*, and QQ Album* the volume of newly generated images reach about 100 petabytes. Some users may try to upload illegal images (e.g., porn). The system has to run a check on each image to try to block them. Imagine trying to search through 100 petabytes of data.*

IPP filter function (ipp_filter2D) took 9ms to perform the operation when compared to 143ms with openCV. ***The IPP filter2D is 15x faster*** *than the OpenCV* plain code.*

Tencent 腾讯

Tencent doubled the speed of its image filter System

JD.com business has grown rapidly, from offering approximately 1.5 million SKUs in 2011 to approximately 25.7 million in 2013. Today, JD.com must handle petabytes of data, which takes an efficient, robust, distributed file system.

JD.com ***speeds up its image processing 17x – handling 300,000 images in 162 seconds instead of 2800 seconds.***

JD.com sped image processing with Intel® IPP

JD.京东 JD.COM

[More Case Studies]

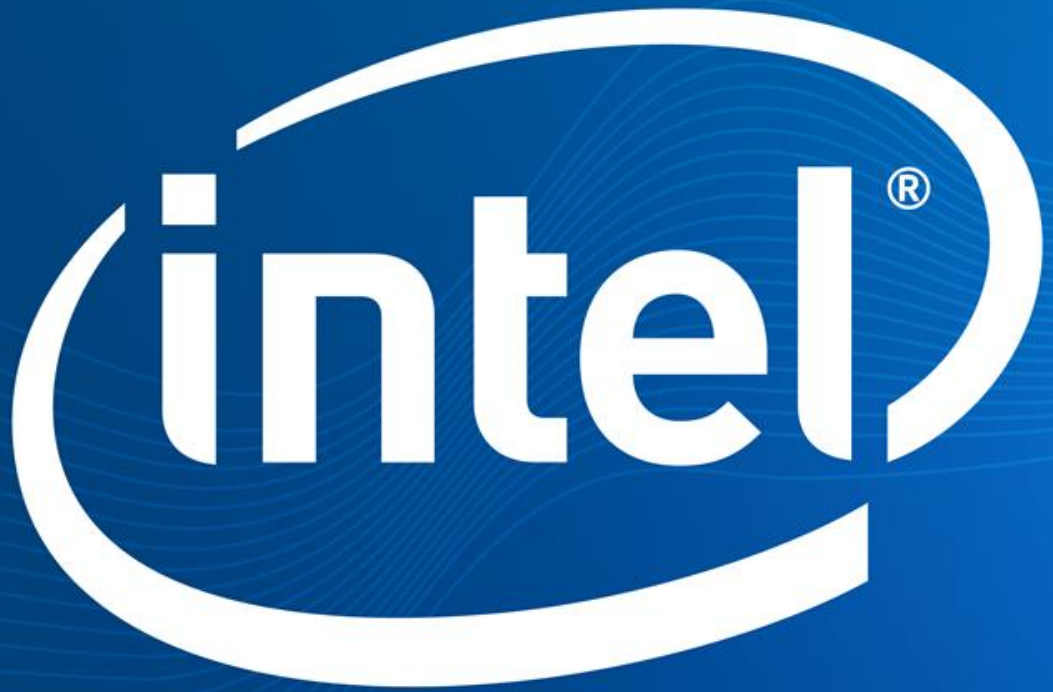
intel


INTEL® INNOVATION DAY
В РИТМЕ ТЕХНОЛОГИЙ