



AI MEETS HPC IN A MULTI-ARCHITECTURE COMPUTING FUTURE:
DEVELOP AND OPTIMIZE SOFTWARE
FOR HETEROGENEOUS COMPUTING SYSTEMS

Ralph de Wargny - Intel Software

CODING

КОДИРОВАНИЕ

CODE DRIVES OUR LIVES


AGRICULTURE




ENERGY



EDUCATION



GOVERNMENT



FINANCE



HEALTH




INDUSTRIAL



MEDIA



RETAIL



SMART HOME



TELECOM



TRANSPORT




CODING WHY IS IT IMPORTANT?

MAKE OUR CUSTOMERS SUCCESSFUL!



PERFORMANCE GAP



Optimization Notice

Copyright © 2019, Intel Corporation. All rights reserved.
*Other names and brands may be claimed as the property of others.



SCALAR



VECTOR



MATRIX



SPATIAL

TODAY CPU CENTRIC



SCALAR



POWER THROUGH PERFORMANCE BOTTLENECKS

Deliver faster, scalable, and portable parallel code
For high-performance computing (HPC),
Enterprise, Cloud and AI.



INTEL® PARALLEL STUDIO XE 2019

- A comprehensive tool suite that simplifies creation and modernization of code on Intel® Xeon® Scalable & Core™ processors.
- It helps developers accelerate workloads with advanced vectorization, multi-threading, multi-node, and memory optimization techniques.

Who needs this product?

- OEMs/ISVs
- C/C++, Fortran, & Python* developers
- Developers, domain specialists of enterprise, data center/cloud, HPC & AI applications

Top 2019 New Features

- **Improve application performance** using OpenMP*—latest SIMD-enabled hardware & Intel® AVX 512.
- **Speed machine learning with new high performance Python***
- **Use a new, intuitive user interface** in Intel® VTune™ Amplifier. Preview a platform profiler.
- **Visualize parallelism with rapid visual prototyping** in Intel® Advisor's Flow Graph Analyzer
- **Extend HPC solutions on the path to Exascale** with next gen Intel® MPI Library
- **Supports industry standards & IDEs.**

Free 30-Day Trial: software.intel.com/intel-parallel-studio-xe

WHAT'S INSIDE INTEL® PARALLEL STUDIO XE

COMPOSER EDITION

BUILD

Compilers & Libraries

C / C++,
Fortran
Compilers

Intel® Math Kernel Library

Intel® Data Analytics
Acceleration Library

Intel Threading Building Blocks
C++ Threading

Intel® Integrated Performance Primitives
Image, Signal & Data Processing

Intel® Distribution for Python*
High Performance Python

PROFESSIONAL EDITION

ANALYZE

Analysis Tools

Intel® VTune™ Amplifier
Performance Profiler

Intel® Inspector
Memory, Thread &
Persistence Debugger

Intel® Advisor
Vectorization Optimization
Thread Prototyping
& Flow Graph Analysis

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



















Operating System: Windows*, Linux*, MacOS1*

Intel® Architecture Platforms



INTEL® PARALLEL STUDIO XE:

HIGH PERFORMANCE, SCALABLE SOFTWARE ACROSS MULTIPLE INDUSTRIES

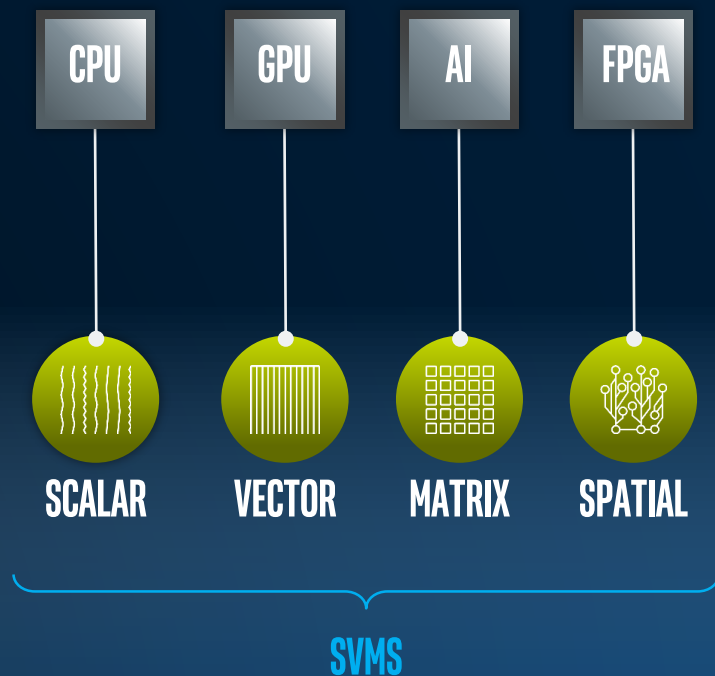
Energy		10X		
EDA		11X		
Science & Research	 4X	 3X <small>*THE REAL SCIENCE</small>	 8X	 35%
Manufacturing	 1.4X	 4X		
Government		 25X		
Computer Software	 2.5X	 1.25X	 1.3X	
IT	 5X		 2X	
Healthcare		 MASSACHUSETTS GENERAL HOSPITAL	20X	
Digital Media				
Telecommunications		 2.5X		

View case studies details [online](#).

Software & workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark & MobileMark, are measured using specific computer systems, components, software, operations & functions. Any change to any of those factors may cause the results to vary. You should consult other information & 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>. 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.

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



INTEL® AI HARDWARE TODAY



DEVICE



INTELLIGENT EDGE



MULTI-CLOUD

OPTIMIZED FRAMEWORKS & SOFTWARE

CPU



GPU



FPGA



ASIC



WORKLOAD BREADTH

AI SPECIALIZATION

Multi-Purpose
Foundation for AI

Data-Parallel Media,
Graphics, HPC & AI

Multi-Function & Real-time
Deep Learning Inference

Deep Learning
Inference

Deep Learning
Training

Media & Vision
DL Inference at
the Edge


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 ONEAPI CORE CONCEPT

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

Common developer experience across Scalar, Vector, Matrix and Spatial architectures (CPU, GPU, AI and FPGA)

Uncompromised native high-level language performance

Based on industry standards and open specifications

oneAPI
Tools

Optimized Applications

Optimized
Middleware / Frameworks

oneAPI Language & Libraries

CPU

SCALAR

GPU

VECTOR


AI

MATRIX

FPGA

SPATIAL

LOW LEVEL INTERFACE, LEVERAGED BY OTHERS



Millions



Thousands

oneAPI Language & Libraries

CPU

GPU

AI

FPGA

ONEAPI FOR CROSS-ARCHITECTURE PERFORMANCE

Optimized Applications

Optimized Middleware & Frameworks

oneAPI Product

Direct Programming

Data Parallel C++

API-Based Programming

Libraries

Analysis &
Debug Tools

CPU

SCALAR

GPU

VECTOR

AI

MATRIX

FPGA

SPATIAL

Some capabilities may differ per architecture.

[Optimization Notice](#)

Copyright © 2019, Intel Corporation. All rights reserved.

*Other names and brands may be claimed as the property of others.

TODAY CPU CENTRIC



CPU



SCALAR



INTEL'S ONEAPI CORE CONCEPT

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

Common developer experience across Scalar, Vector, Matrix and Spatial architectures (CPU, GPU, AI and FPGA)

Uncompromised native high-level language performance

Based on industry standards and open specifications



БОЛЬШОЕ СПАСИБО !

NOTICES & DISCLAIMERS

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice.


Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at [intel.com](https://www.intel.com), or from the OEM or retailer.

Copyright © 2019, Intel Corporation. All rights reserved. Intel, the Intel logo, Xeon, and VTune, are trademarks of Intel Corporation or its subsidiaries in the U.S. and other countries.

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



SUMMARY

Diverse workloads for data-centric computing are driving the need for diverse compute architectures including **CPUs, GPUs, FPGAs, and AI accelerators**

OneAPI unifies and simplifies programming of Intel CPUs and accelerators, delivering developer productivity and **full native language performance**

OneAPI is based on industry standards and **open** specifications to encourage ecosystem collaboration and innovation

CPU



SCALAR



CPU



SCALAR

GPU



VECTOR

AI



MATRIX

FPGA



SPATIAL

