

# 代码优化方法论

## Code Modernization

英特尔软件部

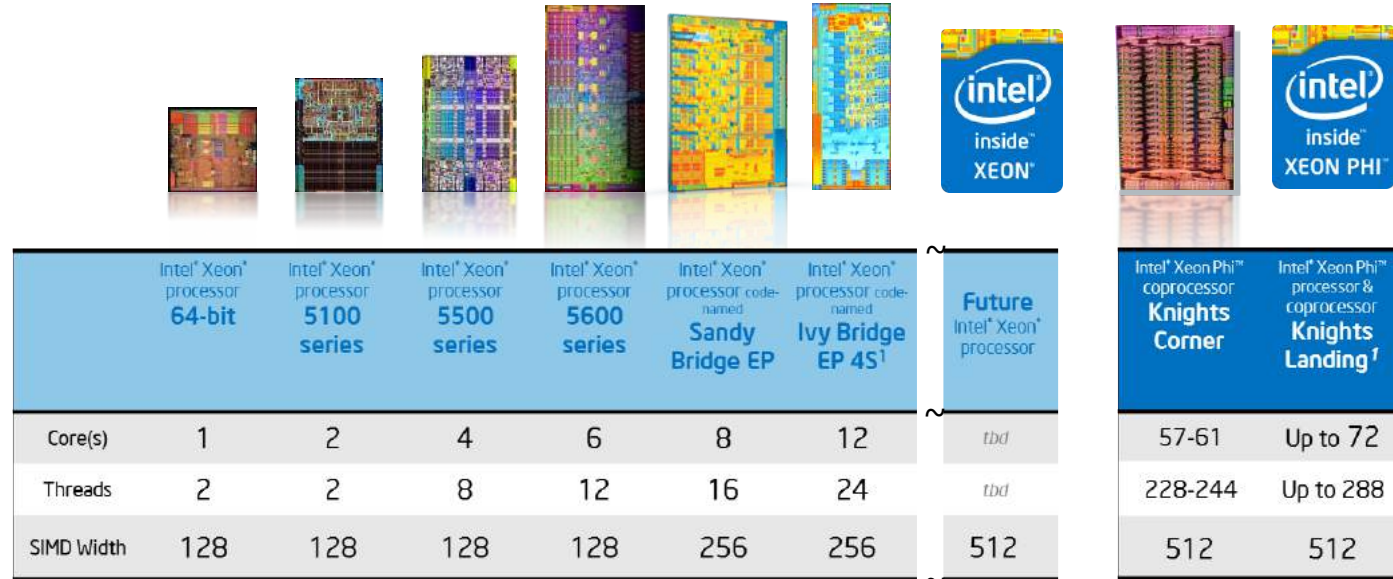
顾彤

# Agenda

- **Overview**
- **Intel Software tools**
- **Code Modernization**
- **Summary**

# Hardware Development : Parallel is the Path Forward

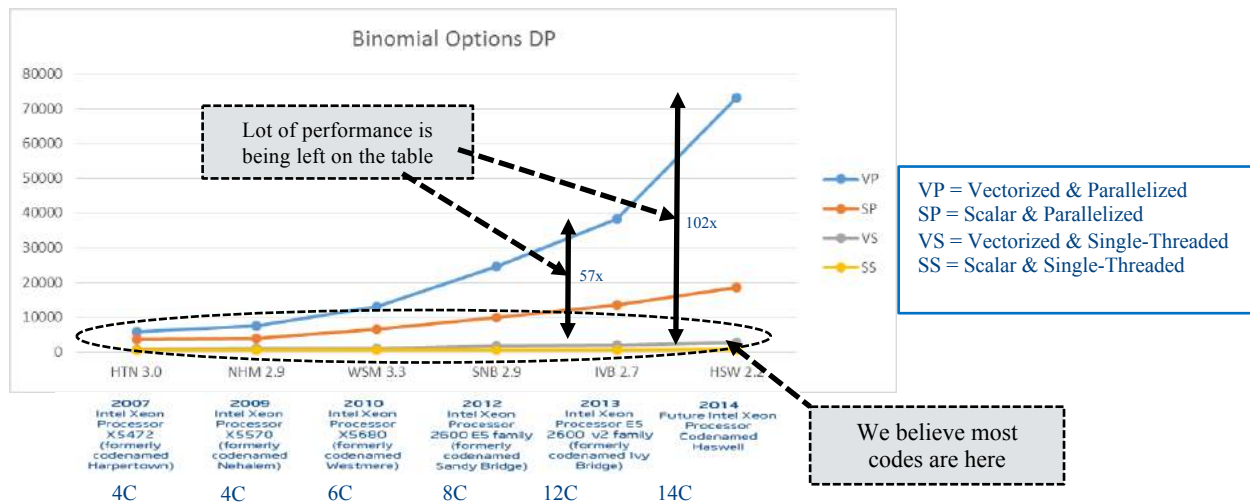
Intel® Xeon® and Intel® Xeon Phi™ Product Families are both going parallel



More cores → More Threads → Wider vectors

# How can I achieve high performance

How to get benefit from Exascale with your code in the future?



**Modernization of your code is the solution**

# Intel SW tools for HPC

- Intel® Parallel Studio XE

- Design, build, verify, and tune  
为程序设计, 代码构建, 校验和调试提供全方位支持
- C++, C, Fortran\*, Python\* and Java\*  
支持C++, C, Fortran and Java\* 等编程环境
- Standards-driven parallel models: OpenMP\*, MPI, and TBB  
支持标准的并行模型

- Highlights from 2017 edition

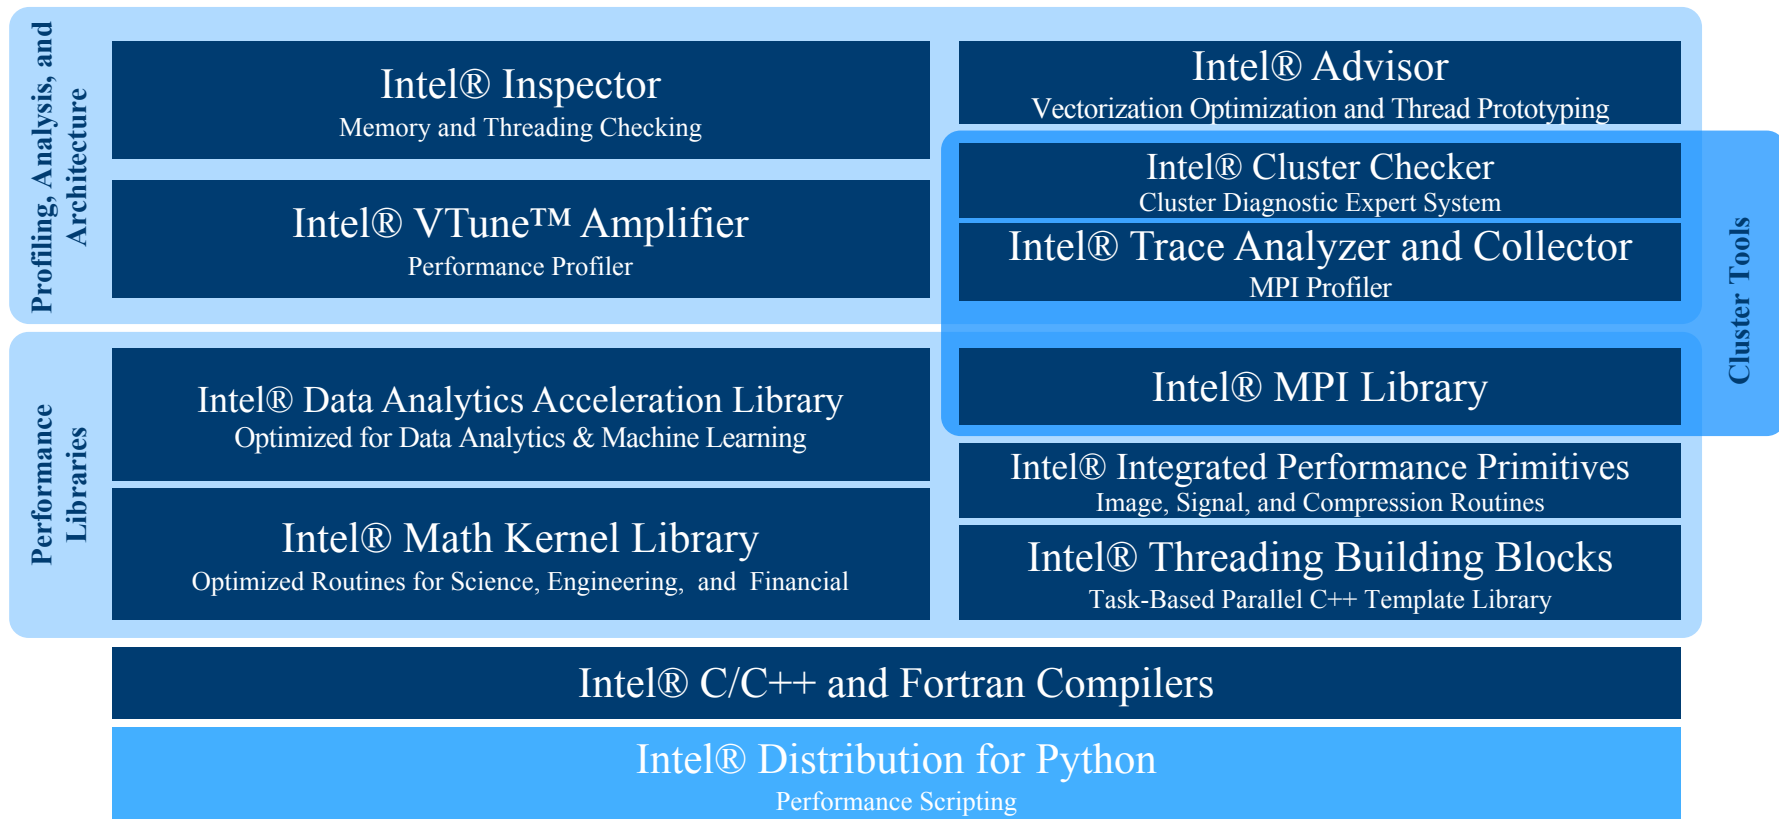
- **Faster Python application performance** using Intel® Distribution for Python and Intel® VTune™ Amplifier XE.  
使用Intel集成的 Python和 Intel® VTune™ Amplifier XE 调优工具, 提升Python 应用的性能
- **Faster deep learning on Intel® architecture** using Intel® Math Kernel Library and Intel® Data Analytics Acceleration Library  
Intel MKL/DAAL类库提升基于IA平台的深度学习系统性能
- **Quickly assess application performance** using snapshot features of Intel® VTune™ Amplifier XE and Intel® Trace Analyzer and Collector  
使用Intel vTune快照功能和追踪收集分析工具快速评价应用性能
- **Scale to next-generation platforms** including the latest Intel® Xeon Phi™ processor. Optimizations for Intel® AVX-512, high bandwidth memory, and explicit vectorization for compiler and analysis tools.  
包括Xeon Phi在内的IA新一代平台的部署, Intel编译器和解析工具帮助开发者实现AVX-512, 高带宽内存和向量化的优化

AVX-512  
PERFORMANCE  
MPI PYTHON  
DATA ANALYTICS XEON  
XEON PHI  
MACHINE LEARNING  
VECTORIZATION  
THREADING



<http://intel.ly/perf-tools>

# Intel® Parallel Studio XE



## Optimization Notice

Copyright © 2016, Intel Corporation. All rights reserved.  
\*Other names and brands may be claimed as the property of others.



# What's Inside

## Intel® Parallel Studio XE 2017



		Composer Edition	Professional Edition	Cluster Edition
Build	Intel® C++ Compiler Intel® Fortran Compiler Intel® Distribution for Python* Intel® Math Kernel Library – fast math library Intel® Integrated Performance Primitives – image, signal & data processing Intel® Threading Building Blocks – threading library Intel® Data Analytics Acceleration Library – machine learning & analytics	✓ ✓ ✓ ✓ ✓ ✓ ✓	✓ ✓ ✓ ✓ ✓ ✓ ✓	✓ ✓ ✓ ✓ ✓ ✓ ✓
Analysis	Intel® VTune™ Amplifier XE – performance profiler Intel® Advisor – vectorization optimization and thread prototyping Intel® Inspector – memory and thread debugging		✓ ✓ ✓	✓ ✓ ✓
SCALF	Intel® MPI Library – message passing interface library Intel® Trace Analyzer and Collector – MPI Tuning and Analysis Intel® Cluster Checker – cluster diagnostic expert system			✓ ✓ ✓
	Rogue Wave IMSL* Library – Fortran numerical analysis	Bundle or Add-on	Add-on	Add-on

Additional configurations including, floating and academic, are available at: <http://intel.ly/perf-tools>

### Optimization Notice

# Educating about Intel SW Tools

## ■ Series Webinars

- Expert talks about the new features
- Attend live or watch after the fact

<https://software.intel.com/events/hpc-webinars>

## ■ High-Performance Programming Books

- Knights-Landing-specific details, programming advice, and real-world examples.
- Intel® Xeon Phi™ Processor High Performance Programming

<http://lotsofcores.com>

### What's New In Intel® Parallel Studio XE 2017 (Online)

Sep 13, 2016 (9:00am - 10:00am PST)

This webinar will go over the latest features of the new release of Intel® Parallel Studio XE 2017.



Add

### Parallel Programming and Optimization for Intel® Architecture (Workshop)

São Paulo, Brazil

Sep 21, 2016 (11:00am - 5:00pm EBT)

Offered by UNESP in partnership with Intel software Brazil, aims to a practical approach to parallel programming on Intel® Xeon® and Intel® Xeon Phi™ based systems



Add

### Code for Speed with High Bandwidth Memory on Intel® Xeon Phi™ Processors (Online)

Oct 11, 2016 (9:00am - 10:00am PST)

Cover methods for users to analyze suitable memory mode and "memkind" library interface, a user-extensible heap manager built on top of jemalloc.



Add

### Vectorization, the "Other" Parallelism You Need

(Online)

Oct 18, 2016 (9:00am - 10:00am PST)

This session demonstrates how process of identifying and modifying code to take advantage of the vector hardware will boost application performance.



Add

### Roofline analysis: A new way to visualize performance optimization tradeoffs (Online)

Oct 25, 2016 (9:00am - 10:00am PST)


Join us in the webinar to see a demonstration and an introduction to how to use roofline analysis to make your own code more efficient.



Add

*"I believe you will find this book is an invaluable reference to help develop your own Unfair Advantage."*

James A.  
Manager  
Sandia National Laboratories





# Optimization Methodology

## Performance profiling

- Hotspots and Micro-arch analysis using Intel® VTune™ Amplifier
- Application behavior analysis and characteristic extraction

## Optimization on Single Node

- Basic optimization
- Improve memory access efficiency
- SIMD optimization – Vectorization
- Achieve high scalability – Parallelism

## Extend to Multi nodes and hybrid platform

- Specific optimization for hybrid platform
- Extend the optimization to multi nodes by improving scalability

# Create Faster HPC and Cloud Software

## What's New in intel® Parallel Studio XE 2018 Beta

### Modernize Code for Performance, Portability and Scalability on the Latest Intel® Platforms

- Use fast **Intel® AVX-512** instructions on **Intel® Xeon®** and **Xeon Phi™** processors.
- **Intel® Advisor** - Roofline finds high impact, but under optimized loops
- **Intel® Distribution for Python\*** - Accelerate applications with high performance Python
- Stay up-to-date with the latest standards and IDE:
  - Full **C++14**, initial **C++2017** draft
  - Full **Fortran\* 2008**, initial **Fortran 2015** draft
  - **OpenMP\* 5.0** draft, **Microsoft Visual Studio\* 2017**
- Accelerate MPI applications with **Intel® Omni-Path** Architecture

### Flexibility for Your Needs

- **Application Snapshot** - Quick answers: Does my hybrid code need optimization?
- **Intel® VTune™ Amplifier** – Profile private clouds with Docker\* and Mesos\* containers, Java\* daemons

And much more\*...

**Register for Beta at: <http://intel.ly/intel-parallel-studio-xe-2018-beta>**

#### Optimization Notice

Copyright © 2016, Intel Corporation. All rights reserved.  
\*Other names and brands may be claimed as the property of others.

\* See Release Notes for the full list with further updates and new features.



# Performance Profiling

## *Hotspots and Mircro-arch characteristic extraction*

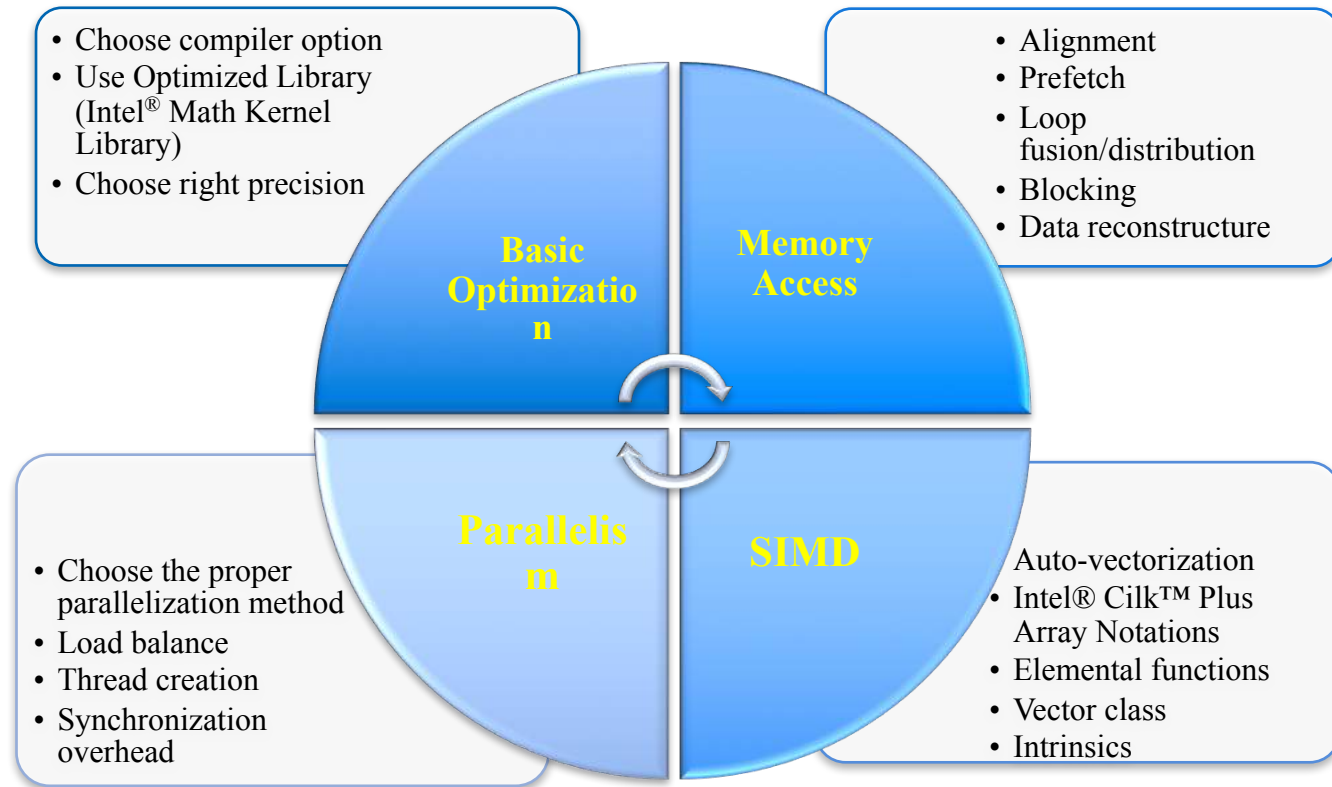
- **Define hotspots**
- **CPI**
- **Cache Hit Ratio**
- **TLB Miss Ratio**
- **Vectorization**
- **Read/Write Bandwidth**

**Intel Tools**

## Application behavior analysis

- **Multi-thread behavior**
  - *Concurrency*
  - *Locks and Waits*
- **MPI Behavior**
  - *Load Balance*
  - *MPI communication*
- **IO analysis**

# Optimization on Single Node



## Optimization Notice

# Extend to Multi Nodes

Good Performance on Single  
Node

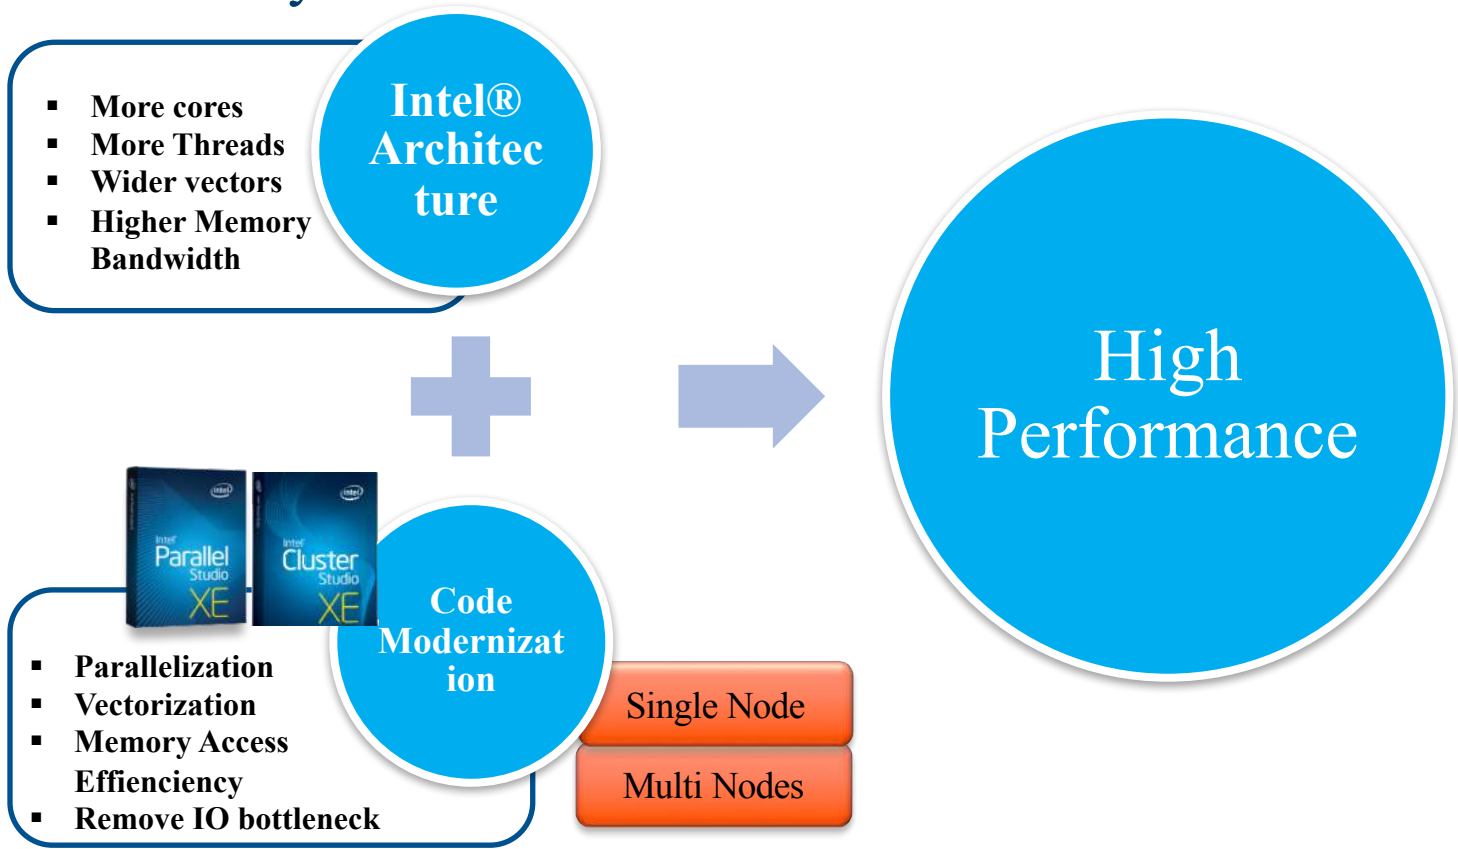


Good Performance on Multi  
Nodes

More Optimization Needed

- Reduce MPI communication
  - Adding multi-thread
  - Reduce data transfer frequency
- Modify serial IO to parallel IO
- Check load balance
  - Modify the load split methods or granularity
  - Async communication

# Summary



## Optimization Notice

