Inspira Crea Transforma

UNIVERSIDAD EAFIT®

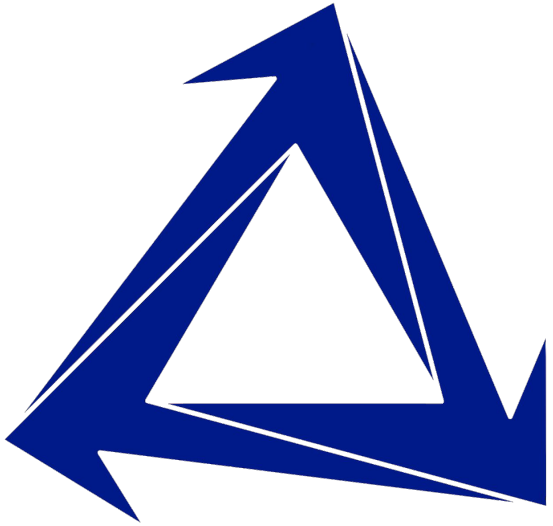

APOLO

Centro de computación científica

Universidad EAFIT

Optimizando su código serial y paralelo con algunas de las herramientas de desarrollo de Intel

Juan David Pineda-Cárdenas
<jpineda2@eafit.edu.co>

Ciclo de Conferencias
Centro de Computación Científica APOLO
Universidad EAFIT

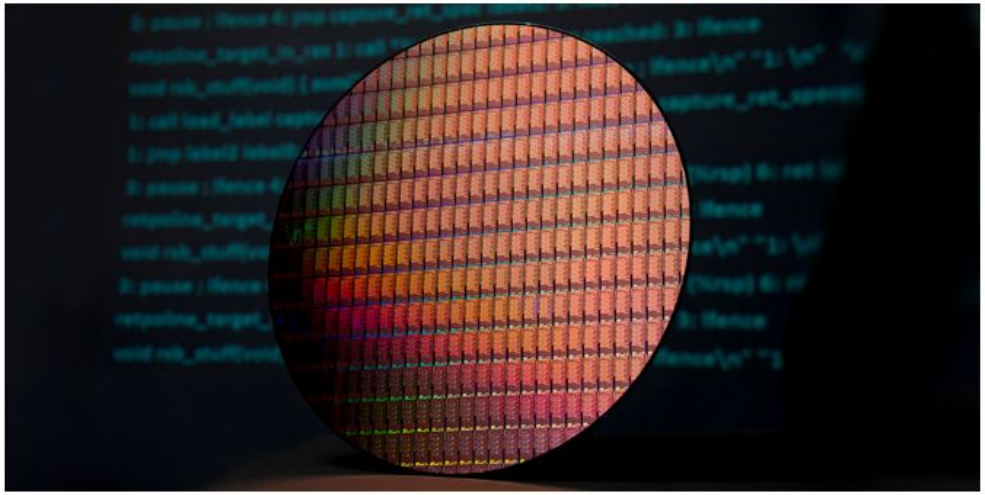

ADVANCING SECURITY AT THE SILICON LEVEL

Hardware-based Protection Coming to Data Center and PC Products Later this Year

By Brian Krzanich

In addressing the vulnerabilities reported by Google Project Zero earlier this year, Intel and the technology industry have faced a significant challenge. Thousands of people across the industry have worked tirelessly to make sure we delivered on our collective priority: protecting customers and their data. I am humbled and thankful for the commitment and effort shown by so many people around the globe. And, I am reassured that when the need is great, companies – and even competitors – will work together to address that need.

But there is still work to do. The security landscape is constantly evolving and we know that there will always be new threats. This was the impetus for the Security-First Pledge I penned in January. Intel has a long history of focusing on security, and now, more than ever, we are committed to the

Tomado de:
https://newsroom.intel.com/editorials/advancing-security-silicon-level/

Avance de la seguridad a nivel del silicio

"Today, I want to provide several updates that show continued progress to fulfill that pledge. First, ***we have now released microcode updates for 100 percent of Intel products launched in the past five years that require protection against the side-channel method vulnerabilities discovered by Google****. As part of this, I want to recognize and express my appreciation to all of the industry partners who worked closely with us to develop and test these updates, and make sure they were ready for production."*

By Brian Krzanich - CEO Intel Corp.

Avance de la seguridad a nivel del silicio

*"While **Variant 1 will continue to be addressed via software mitigations, we are making changes to our hardware design to further address the other two. We have redesigned parts of the processor to introduce new levels of protection through partitioning that will protect against both Variants 2 and 3**. Think of this partitioning as additional "protective walls" between applications and user privilege levels to create an obstacle for bad actors."*

By Brian Krzanich - CEO Intel Corp.

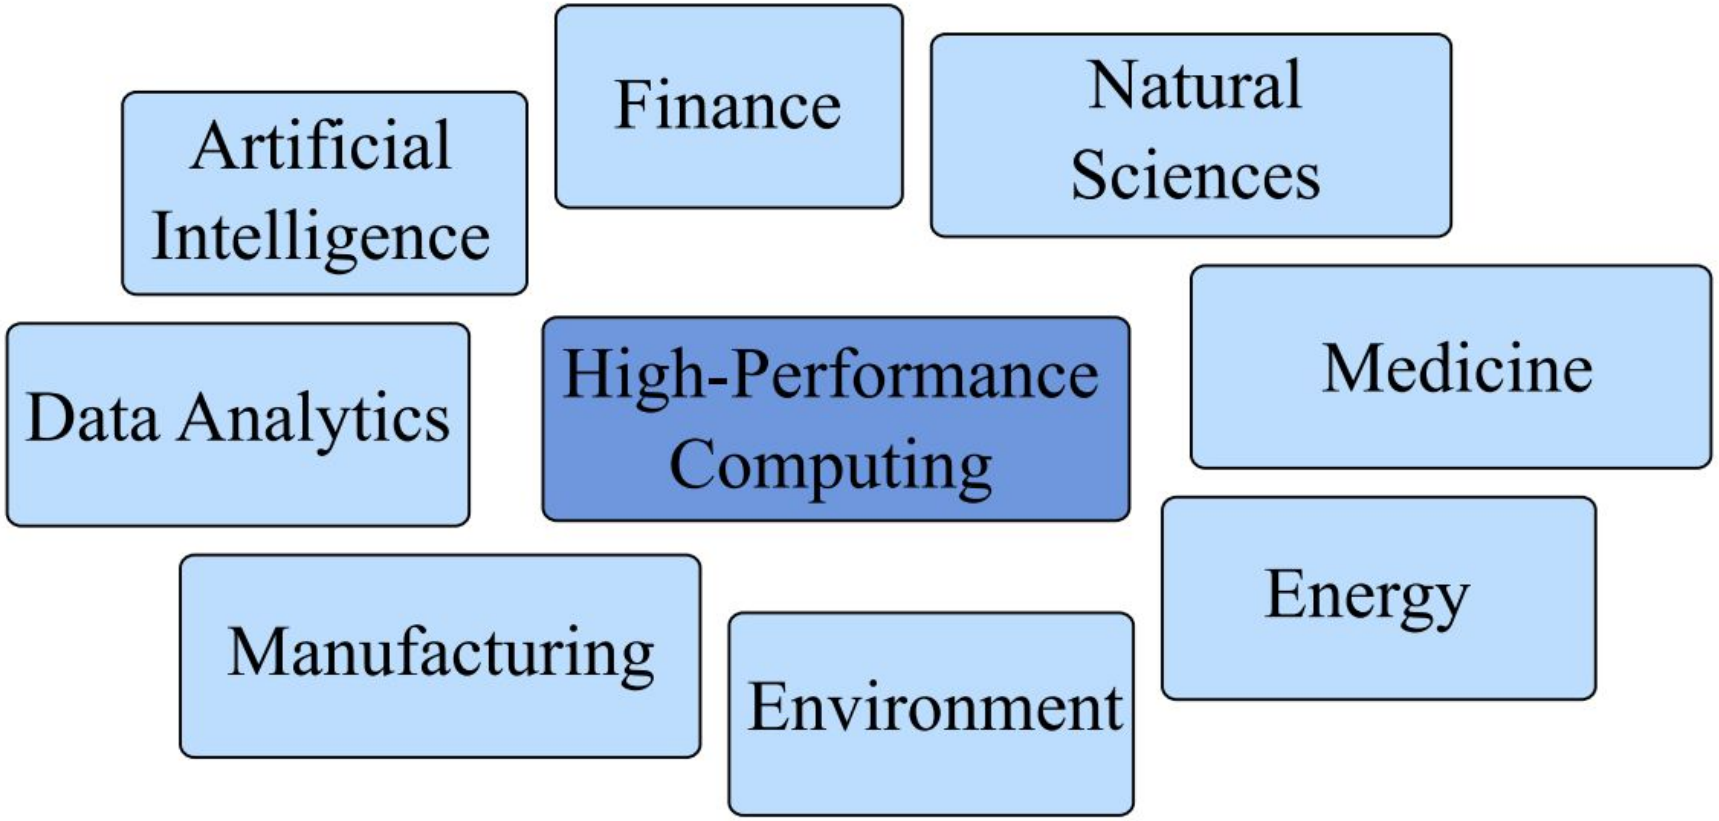

Tomado de: Fundamentals of Parallelism on Intel Architecture. Colfax Research. Intel. Coursera. https://es.coursera.org/learn/parallelism-ia

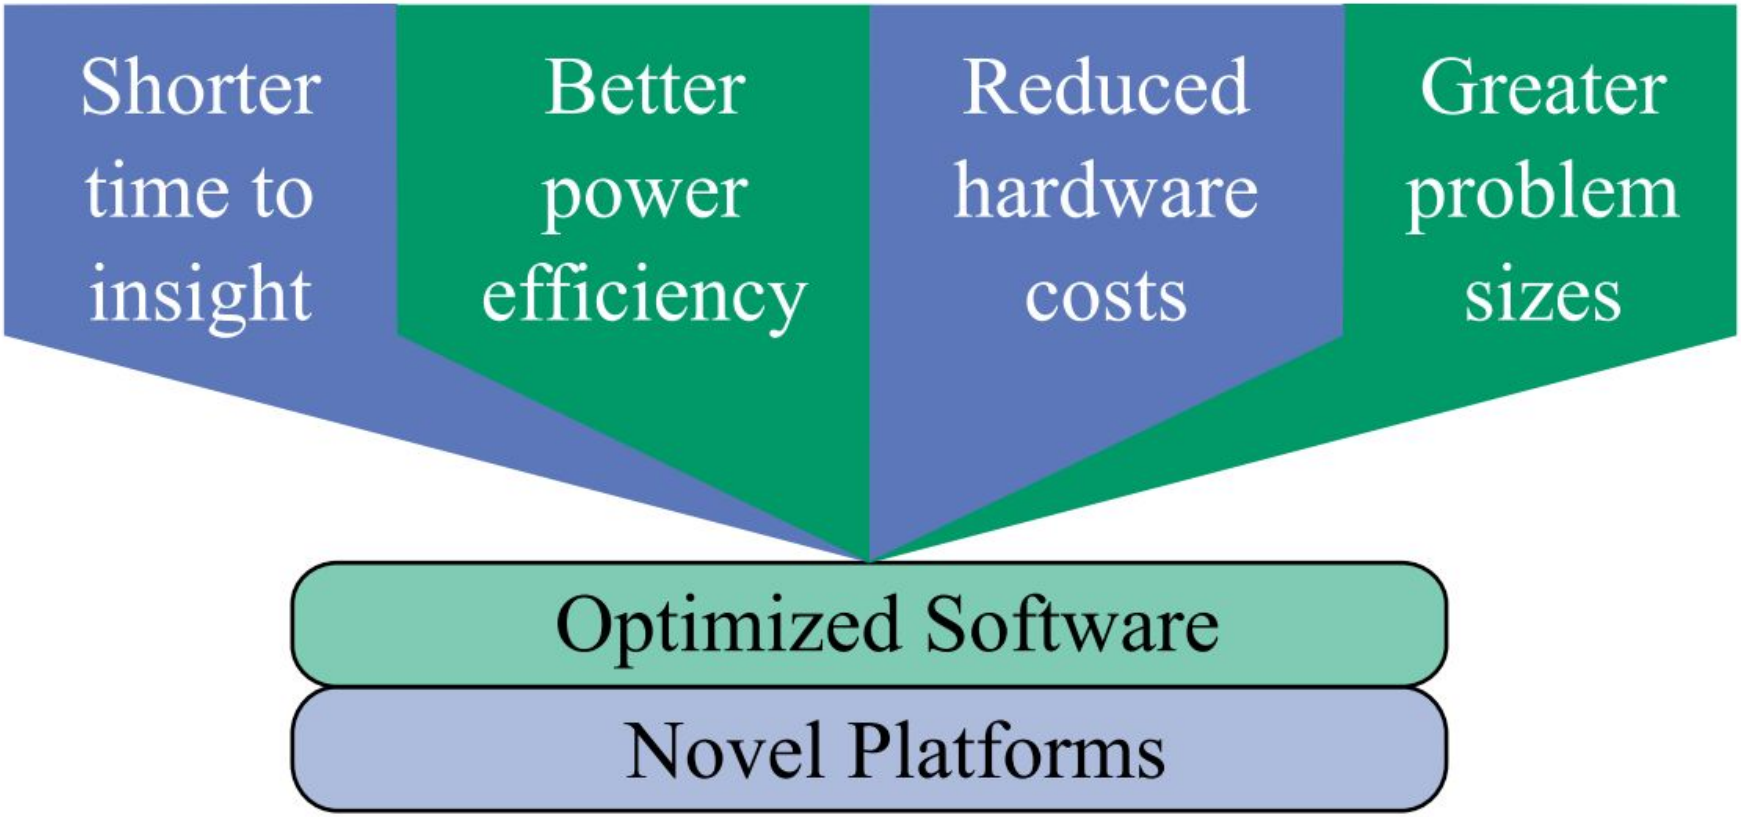

Tomado de: Fundamentals of Parallelism on Intel Architecture. Colfax Research. Intel. Coursera. https://es.coursera.org/learn/parallelism-ia

Capas de Programación Paralela

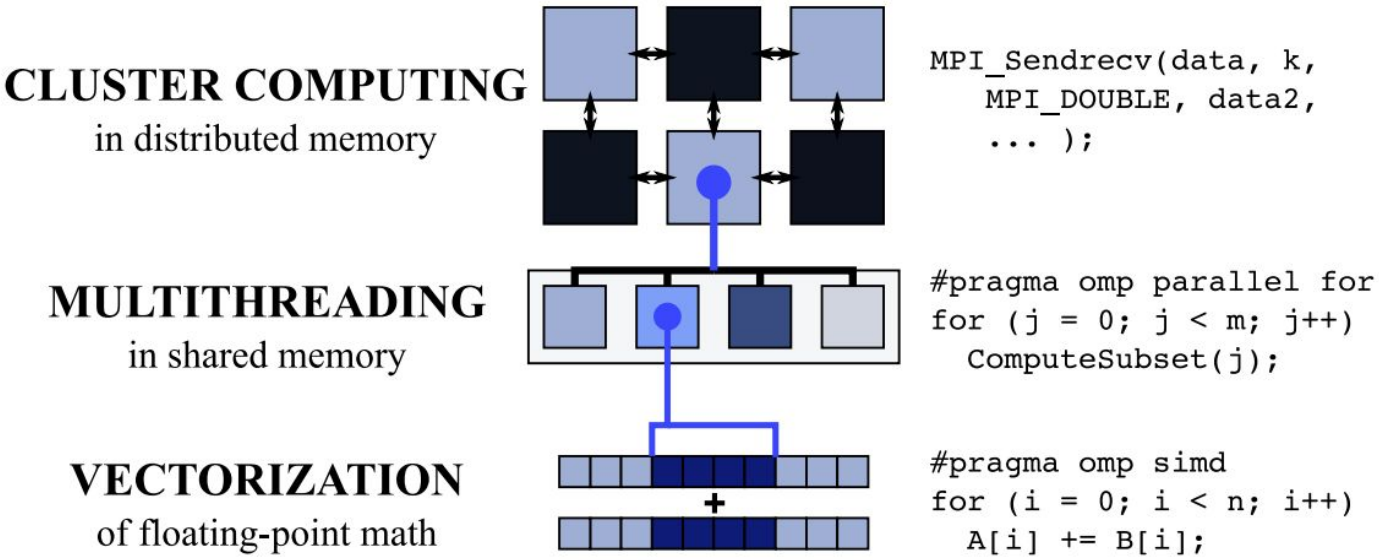

Tomado de: Fundamentals of Parallelism on Intel Architecture. Colfax Research. Intel. Coursera. https://es.coursera.org/learn/parallelism-ia

Paralelismo: Núcleos y vectores

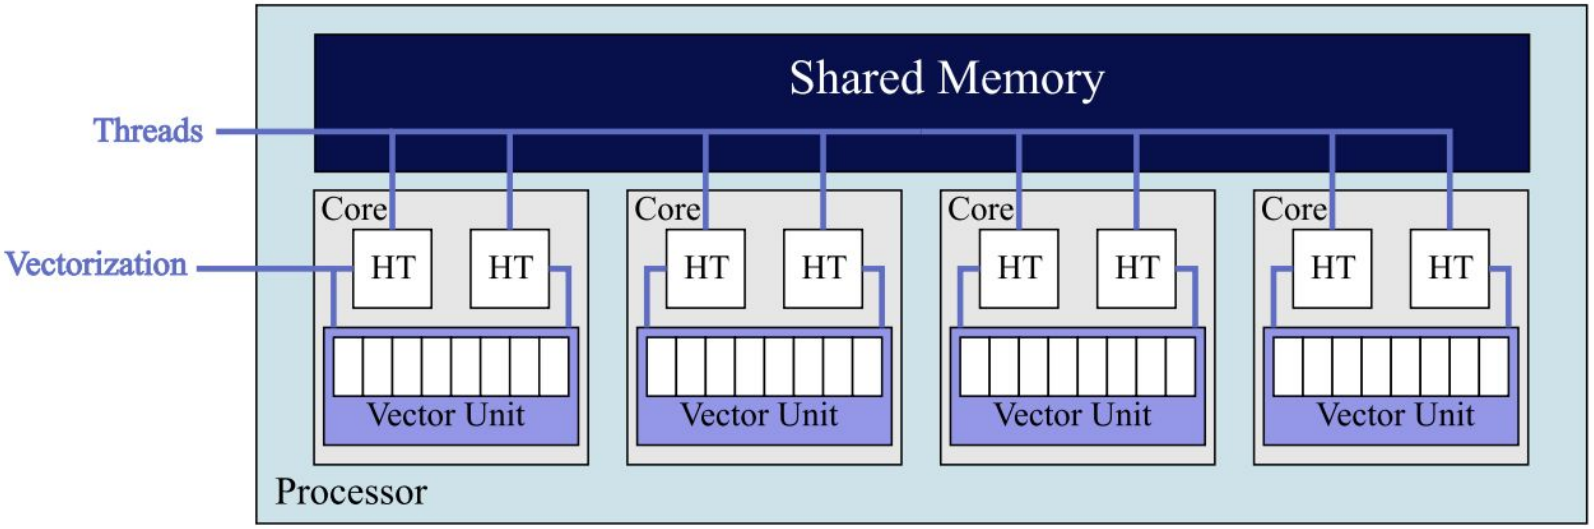

Unbounded growth opportunity, but **not automatic**

Tomado de: Fundamentals of Parallelism on Intel Architecture.
https://es.coursera.org/learn/parallelism-ia

Plataformas de Cómputo de Intel

General-Purpose Processors

Intel® Xeon®
Intel® Core™
Intel® Atom™, ...

Specialized Processors

Intel® Xeon Phi™ processors and coprocessors

Computing Accelerators

Intel® VCA (x86)
Intel® Nervana™ Platform
Intel® DLIA™ (FPGAs)

Network Interconnects

Intel® Omni-Path™ Architecture

Tomado de: Fundamentals of Parallelism on Intel Architecture.
https://es.coursera.org/learn/parallelism-ia

Intel Xeon

- 1-, 2-, 4 CPUs por Servidor
- Propósito General
- Altamente Paralelos (44 núcleos*)
- Rico en recursos (+ cores + cache)
- "*Forgiving Performance*"
- ~ 1.0 TFLOP/s en DP* (Teórico)
- ~ 154 GB/s de ancho de banda* (Medido)

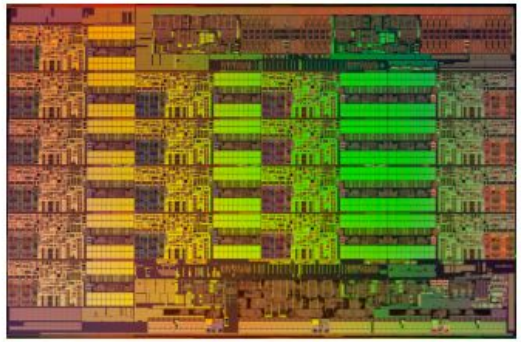

*2-way Intel Xeon, Broadwell architecture (2016), top of the line. (e.g. E5-2699 V4)

Tomado de: Fundamentals of Parallelism on Intel Architecture.
https://es.coursera.org/learn/parallelism-ia

Intel Xeon Phi (2a. Generación)

- Booteable PCIe add-in card
- Especializado para HPC
- Altamente paralelo (72 núcleos*)
- Balanceado para cómputo
- "*Less forgiving than Xeon*"
- ~ 3.0 TFLOP/s en DP* (Teórico)
- ~ 490 GB/s de ancho de banda* (Medido)

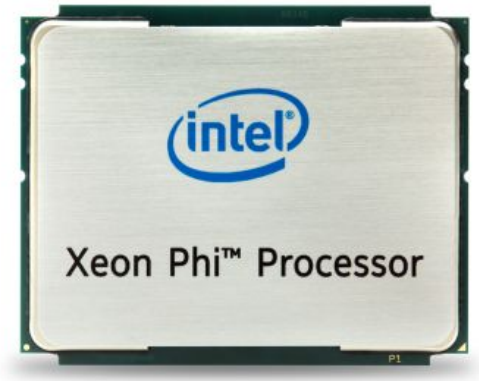

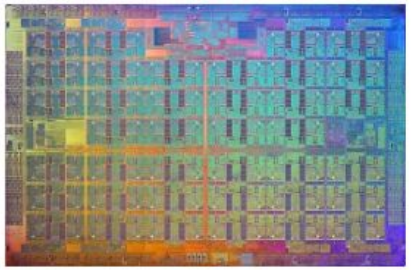

* Intel Xeon Phi, Knights Landing architecture (2016), top of the line. (e.g. 7290P)

Tomado de: Fundamentals of Parallelism on Intel Architecture.
https://es.coursera.org/learn/parallelism-ia

Un código para todas las plataformas

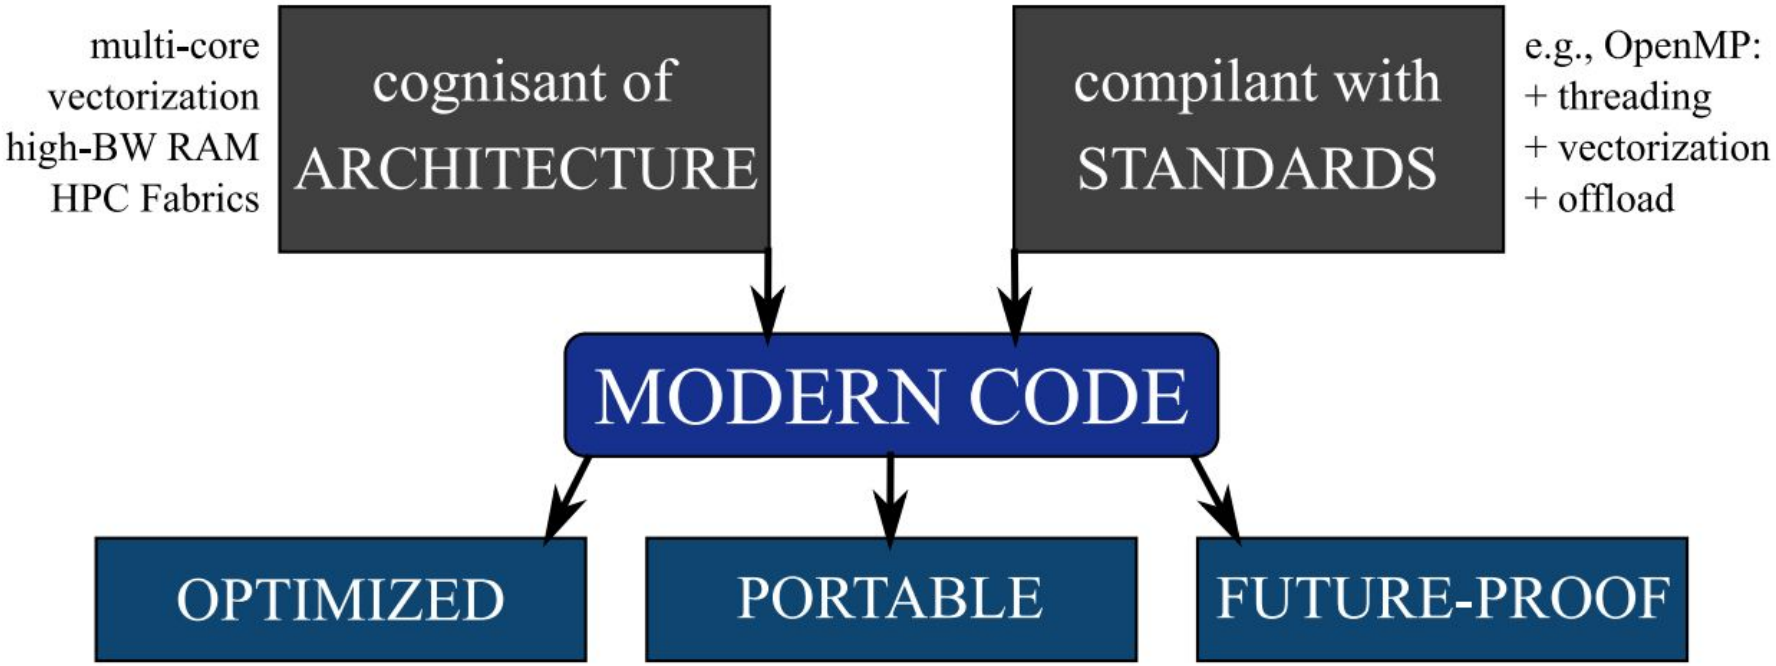

Tomado de: Fundamentals of Parallelism on Intel Architecture.
https://es.coursera.org/learn/parallelism-ia

Computación en Ciencia e Ingeniería

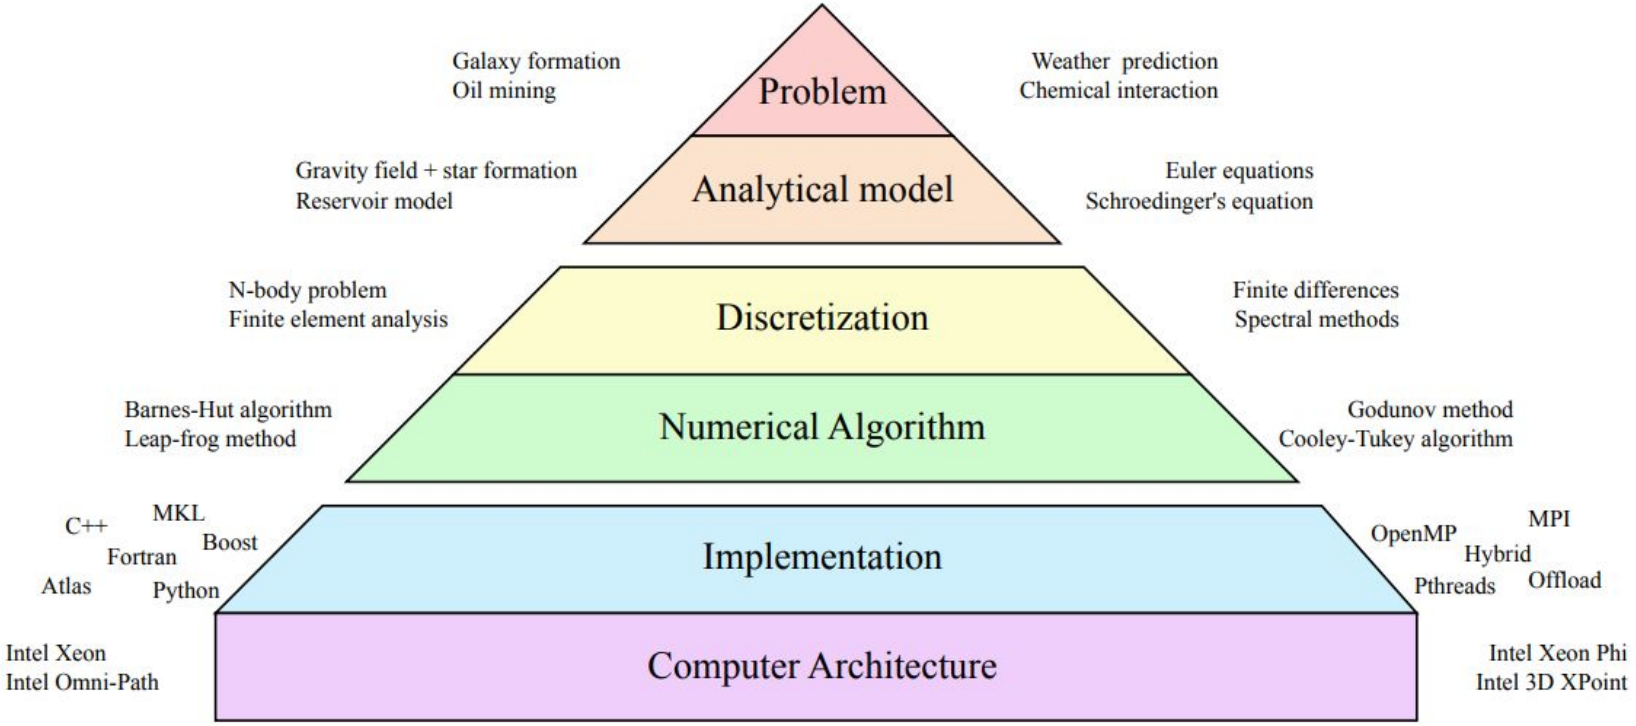

Tomado de: Fundamentals of Parallelism on Intel Architecture.
https://es.coursera.org/learn/parallelism-ia

Áreas de Optimización

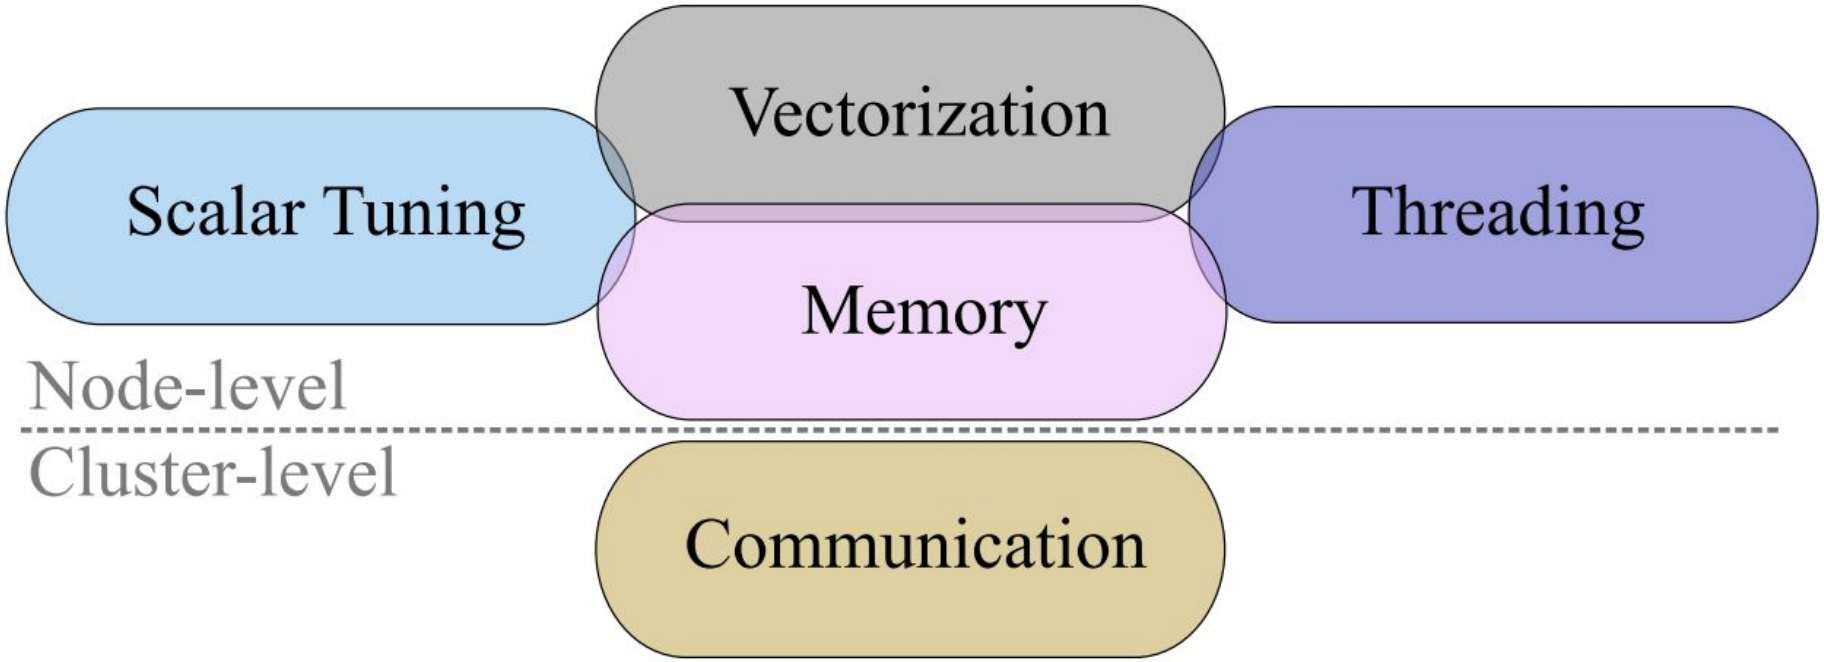

Tomado de: Fundamentals of Parallelism on Intel Architecture.
https://es.coursera.org/learn/parallelism-ia

Experiencia en Modernización de Código Común

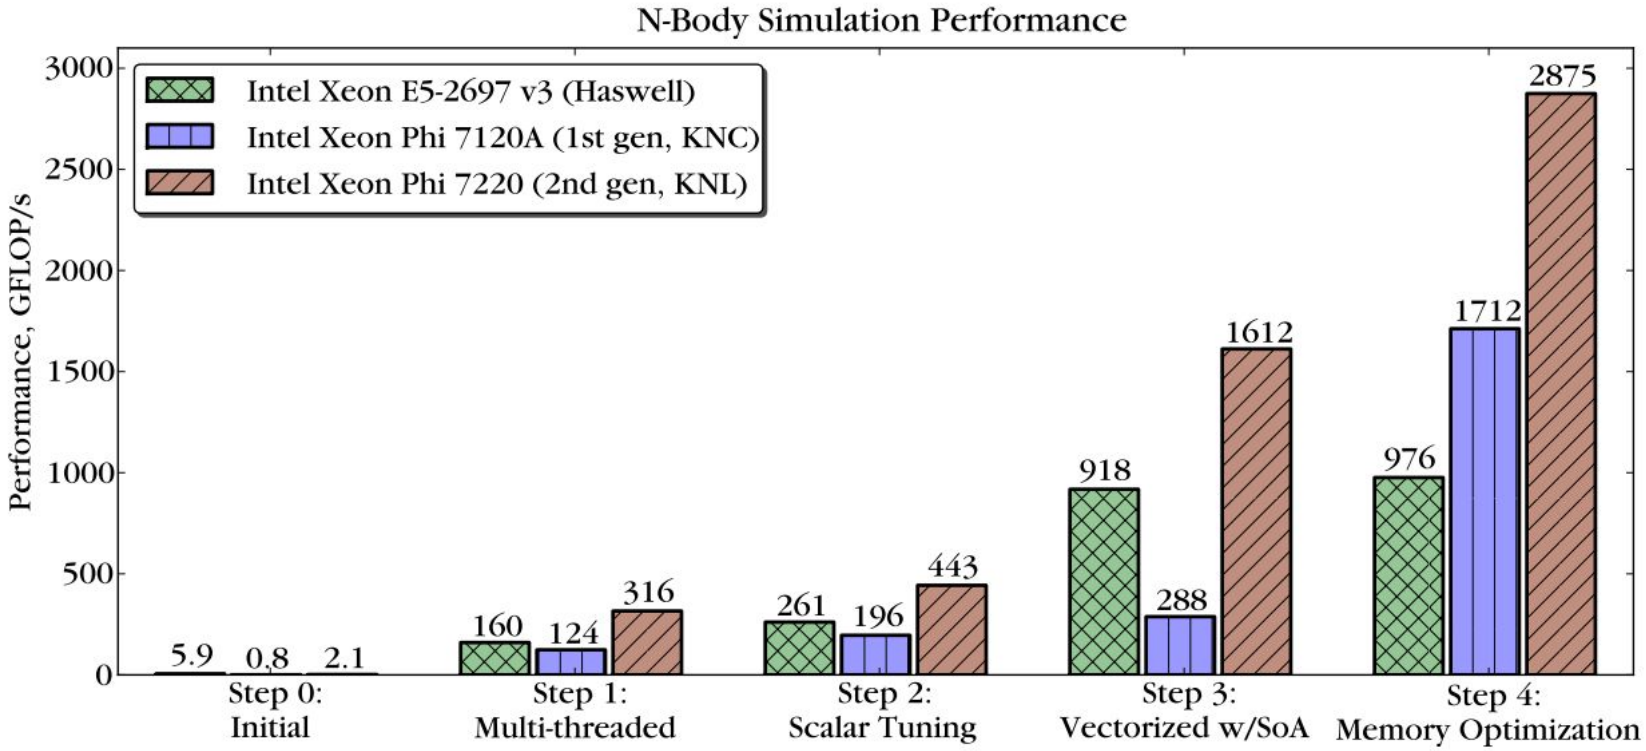

Tomado de: Fundamentals of Parallelism on Intel Architecture.
https://es.coursera.org/learn/parallelism-ia

Multihilos con OpenMP

Cores implement MIMD (Multiple Instruction Multiple Data) arch

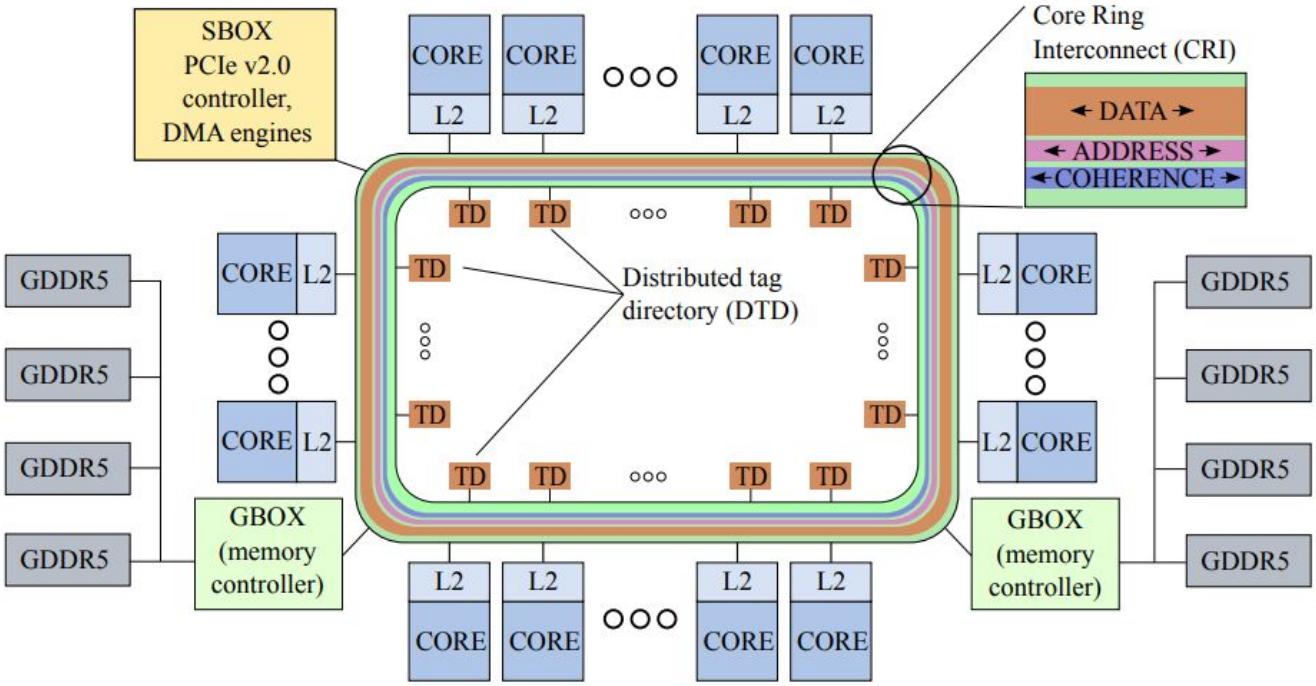

Tomado de: Fundamentals of Parallelism on Intel Architecture.
https://es.coursera.org/learn/parallelism-ia

Vectorización con Compiladores Intel

Vectors – form Single Instruction Multiple Data (SIMD) architecture

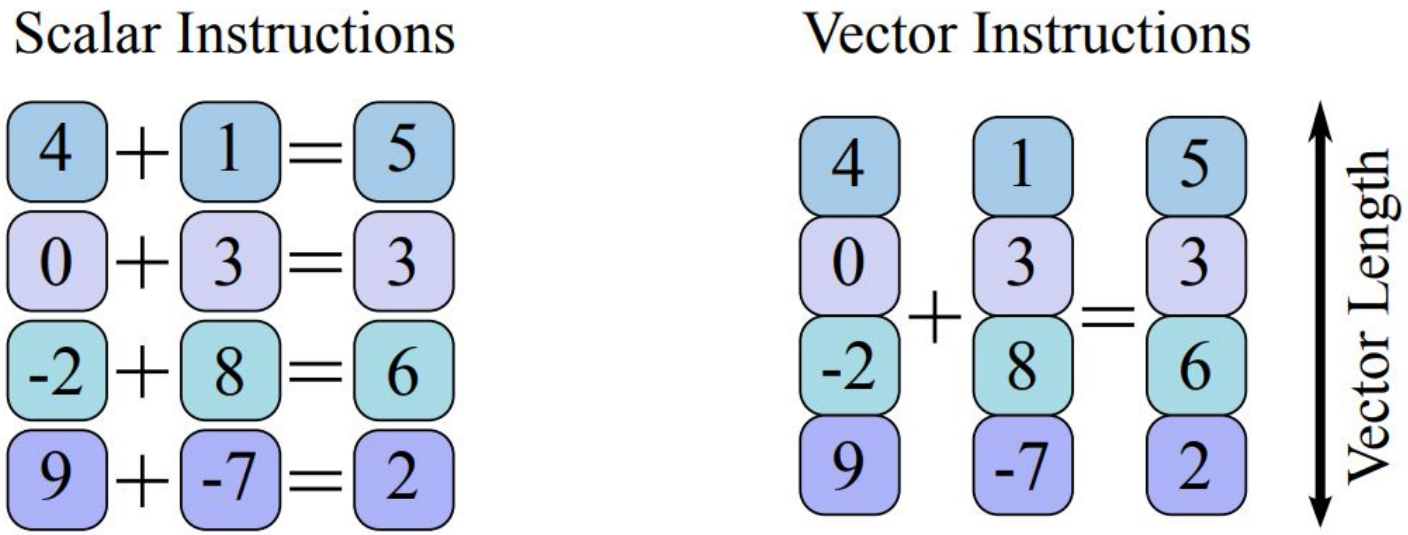

Tomado de: Fundamentals of Parallelism on Intel Architecture.
https://es.coursera.org/learn/parallelism-ia

Tráfico de memoria

Caches facilitate data re-use

RAM is optimized for streaming

Tiling

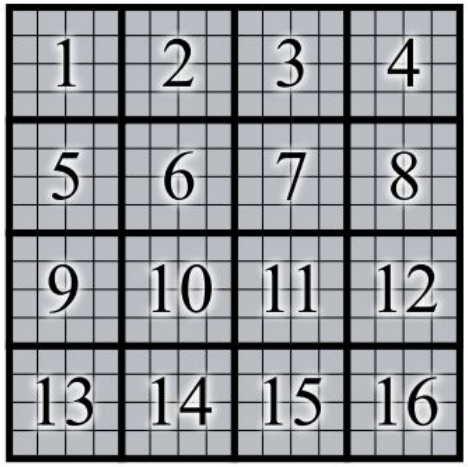

Cache-Oblivious Recursion

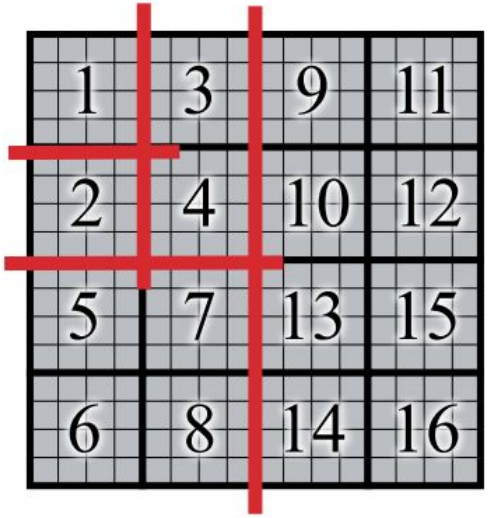

Tomado de: Fundamentals of Parallelism on Intel Architecture.
https://es.coursera.org/learn/parallelism-ia

Clusters y MPI

Clusters form distributed-memory systems with network interconnects

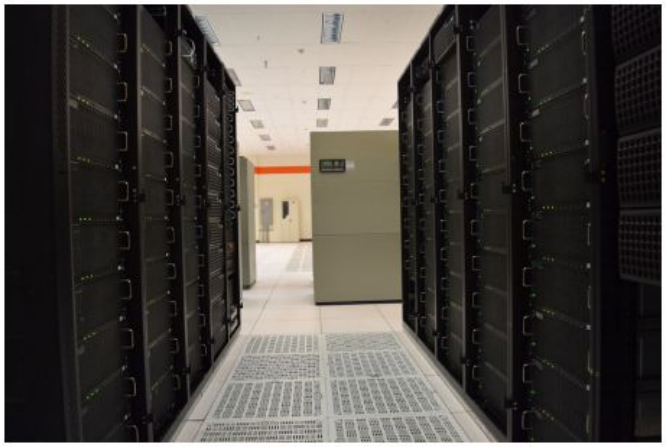

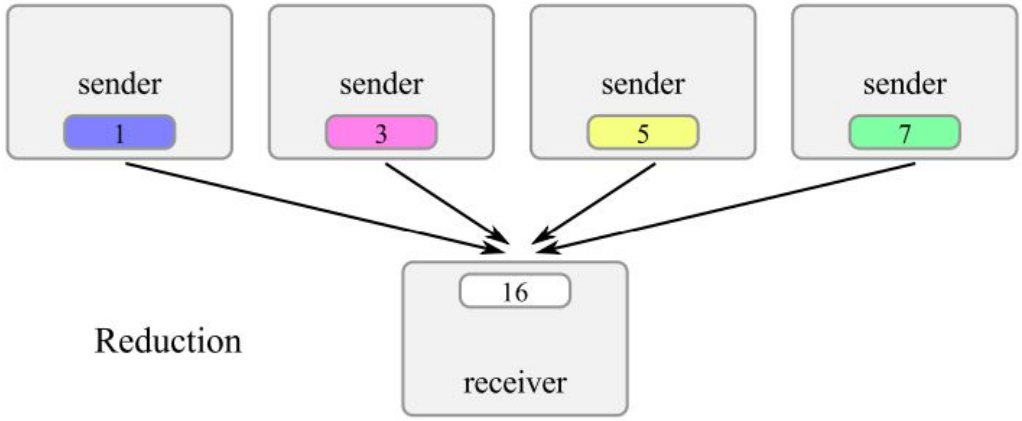

Tomado de: Fundamentals of Parallelism on Intel Architecture.
https://es.coursera.org/learn/parallelism-ia

Intel en Apolo

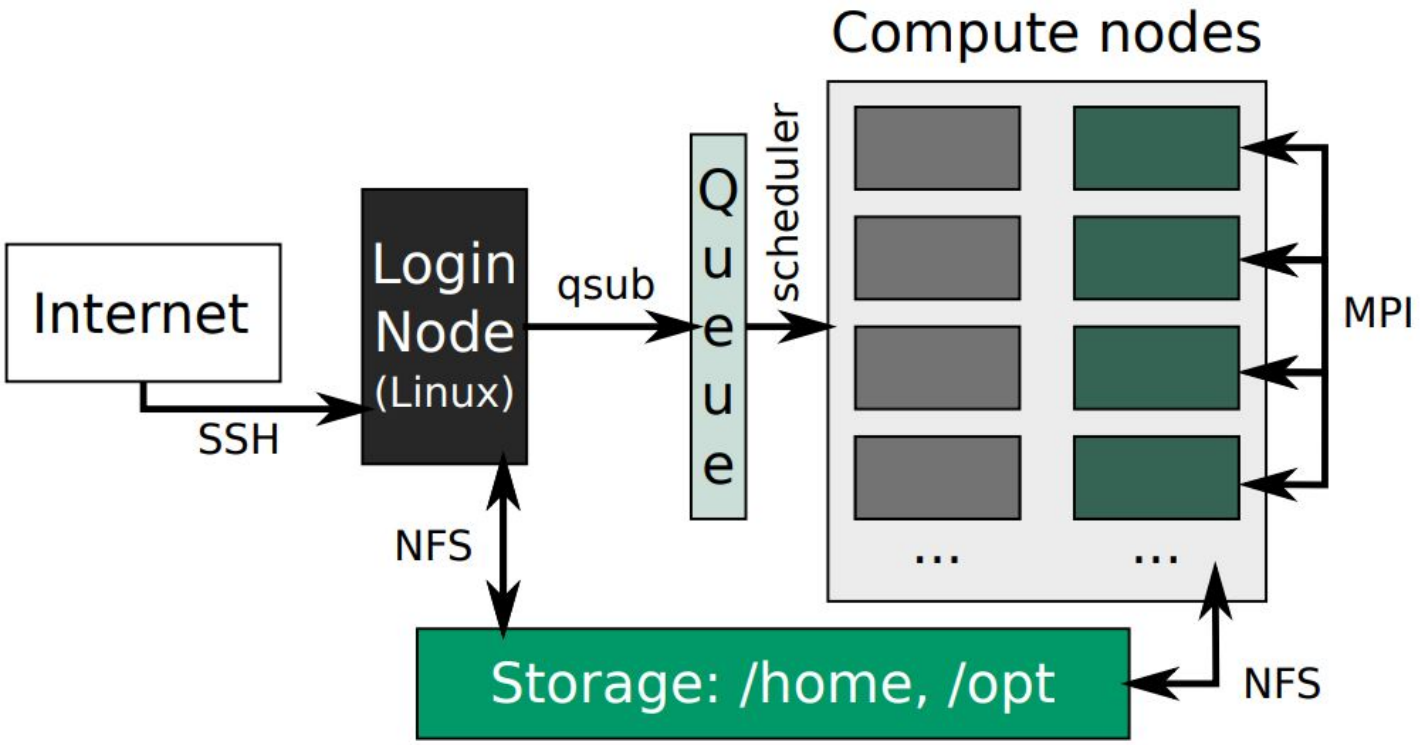

Tomado de: Fundamentals of Parallelism on Intel Architecture.
https://es.coursera.org/learn/parallelism-ia

Opciones generales para optimización

- -O1
  - Optimizar tamaño del código, no vectorización
- -O2 - Por defecto
  - Inlining, vectorización
  - Optimizaciones básicas de ciclos
- -O3
  - Aplicaciones intensivas en ciclos
  - Incluye todas las de O2
  - Transformaciones de ciclos más agresiva

Tomado de: Optimization for Intel C++ and Fortran Compilers
https://software.intel.com/en-us/articles/intel-c-compiler-introduction

Optimización para Multi-Núcleo y Multi-procesamiento

- -parallel
  - Optimización Auto-paralelismo
- Lenguajes soportados para paralelismo explícito
  - OpenMP* 4.0 y la mayor parte de 4.5
    - Disponible para C/C++ y Fortran
    - Paralelismo basado en Pragmas/Directivas
    - Reducciones
- -xHost
  - Le dice al compilador que genere instrucciones para el conjunto de instrucciones más alto que esté disponible

Tomado de: Optimization for Intel C++ and Fortran Compilers
https://software.intel.com/en-us/articles/intel-c-compiler-introduction

Opciones específicas por procesador

- -m<target>
  - Le dice al compilador cuales características puede usar, incluyendo cuál conjunto de instrucciones puede generar según <target>
    - Targets:
      - SSE
      - SSE2 (default)
      - SSE3
      - SSSE3
      - SSE4.1
      - SSE4.2
      - AVX
      - AVX2
      - AVX512
      - IA32

Tomado de: Optimization for Intel C++ and Fortran Compilers
https://software.intel.com/en-us/articles/intel-c-compiler-introduction

Opciones específicas por procesador

- -x<target>
  - Le dice al compilador cuales características puede usar, incluyendo cuál conjunto de instrucciones y optimizaciones puede generar según <target>
    - Targets:
      - COMMON-AVX512
      - MIC-AVX512
      - CORE-AVX512
      - CORE-AVX2
      - CORE-AVX-I
      - AVX
      - SSE4.2
      - SSE4.1
      - ATOM_SSE4.2
      - ATOM_SSSE4.2
      - ATOM_SSSE3
      - ...

Tomado de: Optimization for Intel C++ and Fortran Compilers
https://software.intel.com/en-us/articles/intel-c-compiler-introduction

Math Kernel Library

- Acelera rutinas matemáticas, incrementa el rendimiento de las aplicaciones y reduce el tiempo de desarrollo
- Librería lista para usar que incluye:
  - Algebra Lineal
  - Transformadas rápidas de Fourier
  - Redes neuronales profundas
  - Estadística vectorial y ajuste de datos
  - Matemática vectorial
- Usa la API de C y Fortran para Compatibilidad con BLAS, LAPACK y FFTW. No se requieren cambios de código

¿Dónde Obtenerlo?

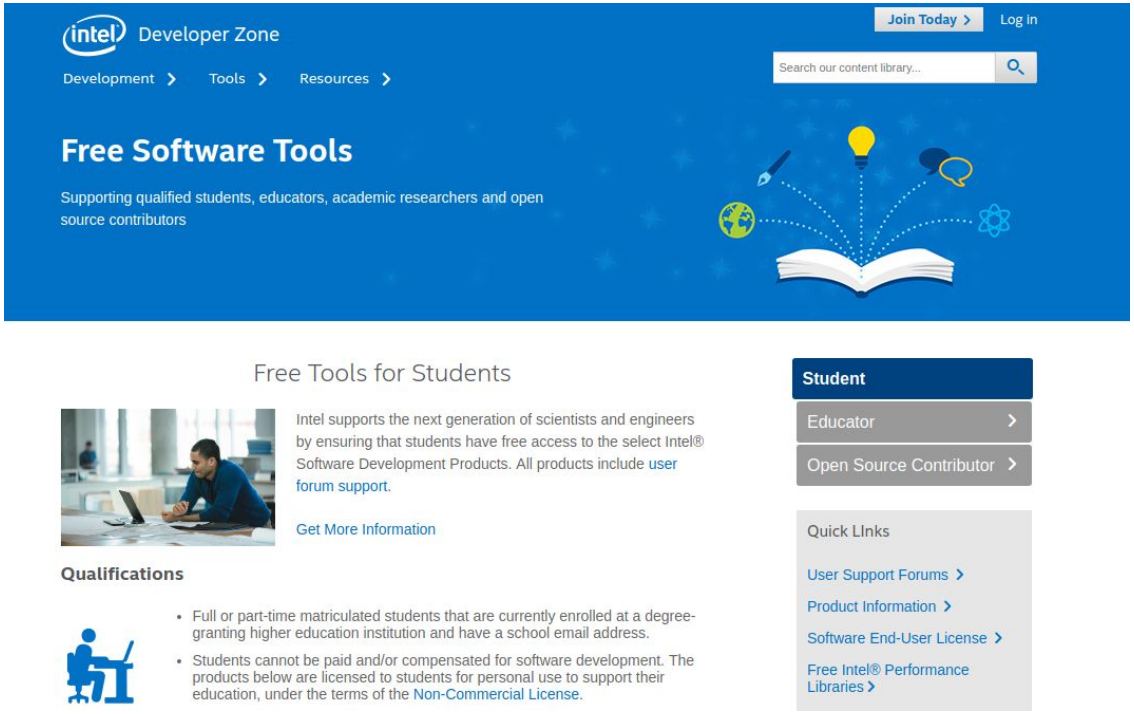

- https://software.intel.com/en-us/qualify-for-free-software/student
- https://software.intel.com/en-us/qualify-for-free-software/educator

Recurso: Introducción a OpenMP

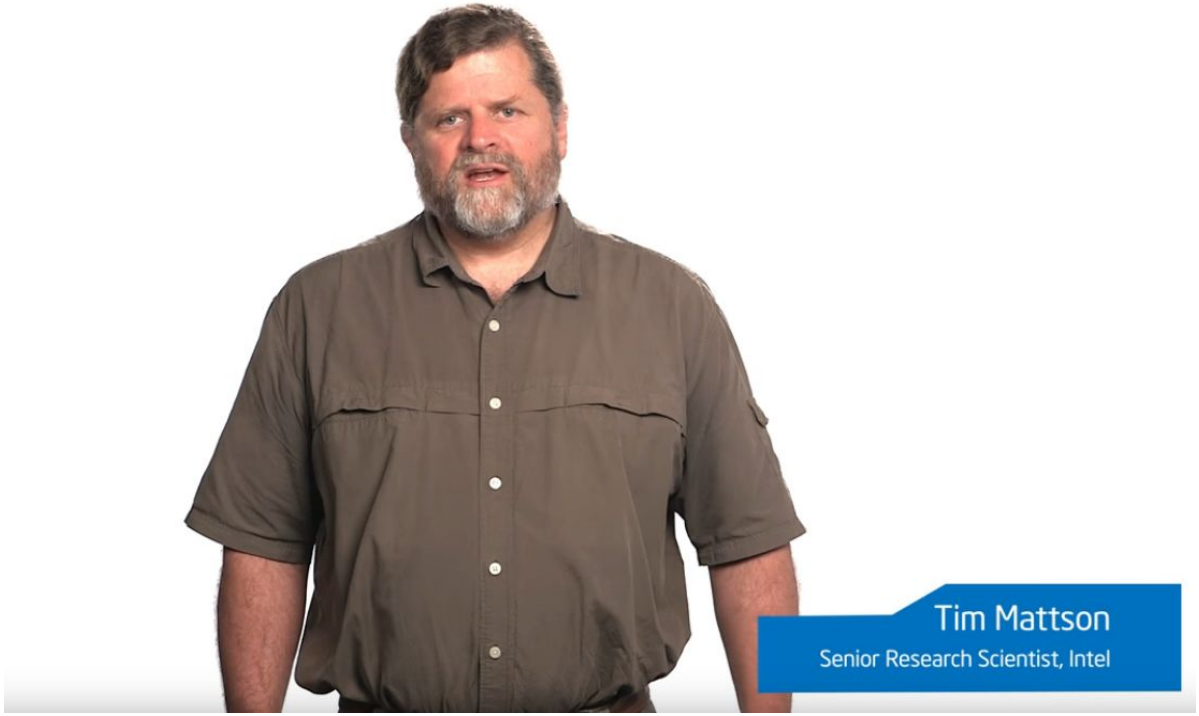

https://youtu.be/nE-xN4Bf8XI

Python Intel Distribution

https://software.intel.com/en-us/distribution-for-python

Fundamentals of Parallelism on Intel Architecture

by Intel

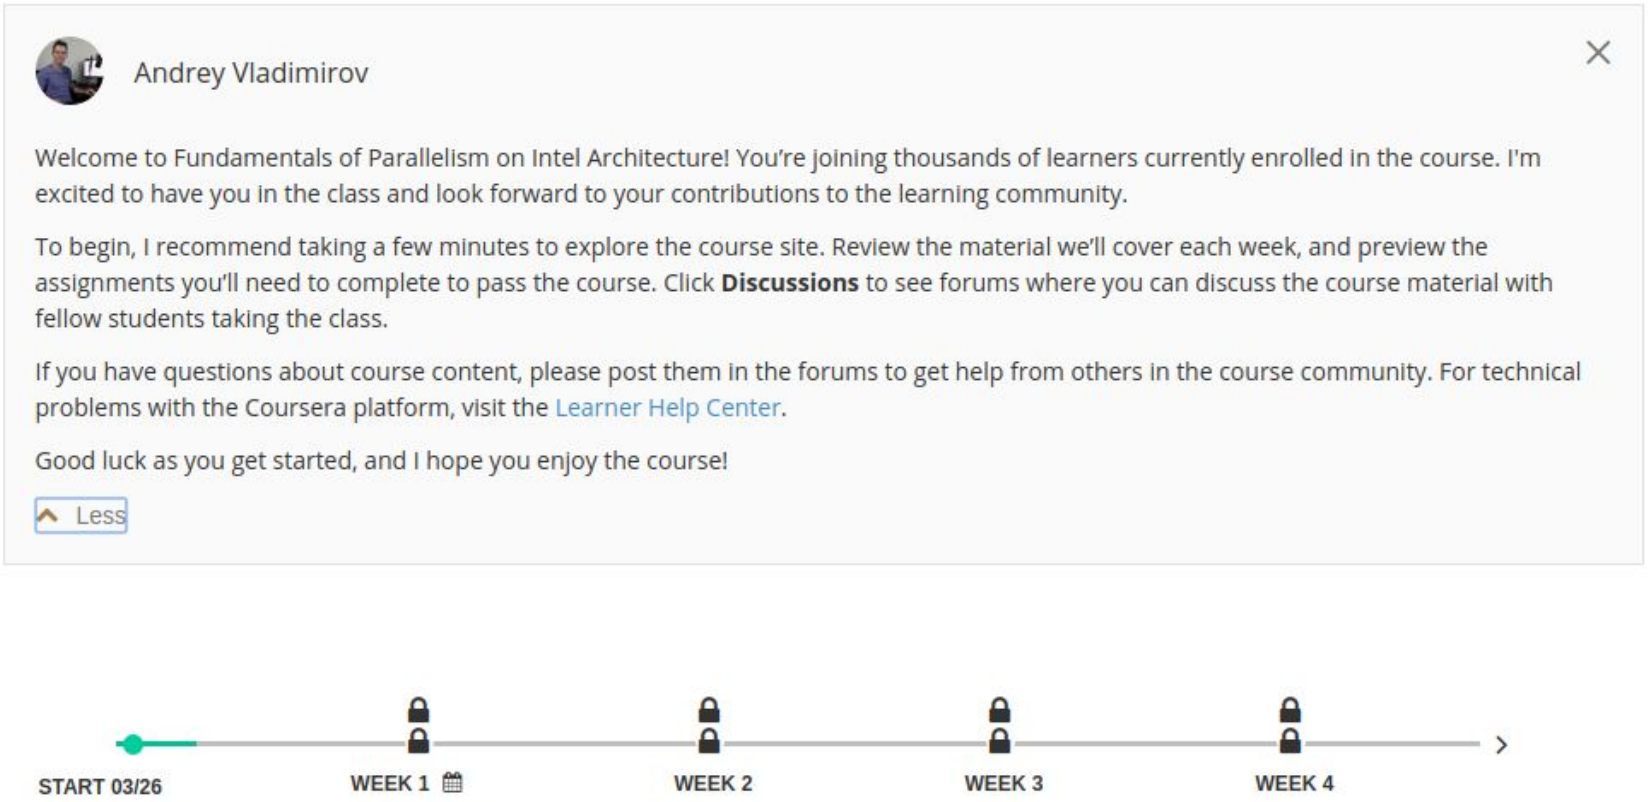

http://coursera.org/learn/parallelism-ia/

Referencias y Recursos

- Compiler options: https://software.intel.com/en-us/node/677967
- Math Kernel Library: https://software.intel.com/en-us/mkl
- Fundamentals of Parallelism on Intel Architecture: https://www.coursera.org/learn/parallelism-ia/home/welcome
- Intel Software Development Products Samples and Tutorials: https://software.intel.com/en-us/product-code-samples?field_software_product_tid[]=20813
- Performance essentials using OpenMP* 4.0 vectorization with C/C++ https://software.intel.com/en-us/videos/performance-essentials-using-openmp-40-vectorization
- Vectorization Essentials: https://software.intel.com/en-us/articles/vectorization-essential
- Intel C++ Compiler Introduction: https://software.intel.com/en-us/articles/intel-c-compiler-introduction
- Tutorial: Using Auto Vectorization: https://software.intel.com/en-us/intel-cplusplus-compiler-auto-vectorization-tutorial
- Explor our documentation: https://software.intel.com/en-us/documentation?field_software_product_tid%5B0%5D=20813&field_software_product_tid%5B1%5D=20813&field_article_type_tid%5B0%5D=20783&field_article_type_tid%5B1%5D=20783
- Intel C++ Compiler in Intel Parallel Studio XE Support: https://software.intel.com/en-us/c-compilers/ipsxe-support/training