

## Intel Tuning for Juwels and Jureca

FZ-Jülich, May 24, 2024 Heinrich.Bockhorst@Intel.com



iclassethal def soll(((Seconstation)) def soll(((Seconstation)))

### Agenda

- oneAPI Initiative
- Intel<sup>®</sup> Compiler
- Application Performance Snapshot (APS)
- VTune Profiler
- Advisor
- Intel<sup>®</sup> MPI

### Cross-Architecture Programming for Accelerated Compute, Freedom of Choice for Hardware ONEAPI: Industry Initiative & Intel Products

One Intel Software & Architecture group Intel Architecture, Graphics & Software November 2020





### Free Download of all packages!



performing Python libraries

Data Scientists & Al Developers

#### Get started quickly

Code Samples, Quick-start Guides, Webinars, Training https://software.intel.com/oneapi https://devcloud.intel.com/oneapi/get\_started/



### New Cloud Access cloud.intel.com

- Select (free access)
  - "Training Catalog"
- "Hardware Catalog" shows more options only for paying customers!
- Download Samples from:

https://github.com/oneapi-src/oneAPI-samples



### New Cloud Access cloud.intel.com

- Free Access to new Intel CPU/GPU hardware
- Training material on Jupyter notebooks.

AI		
Al Kit XOBoost Predictive Modeling Learn predictive modeling with decision trees using Intel <sup>®</sup> Al Analytics Toolit	Heterogeneous Programming Using Data Parallel Extension for Numba® for Al and HPC Data Parallel Extension for Numba accelerates Python® code on Intel® XPUs Exunch C	Machine Learning Using oneAPI Intel® AI Analytics Toolkit accelerates data science and analytics with Python®
C++ SYCL		
Essentials of SYCL Learn to write performant and portable code using oneAPI and SYCL C++	Performance, Portability and Productivity Learn to write performant and portable HPC code for multiple platforms with one API and SYCL C++	Introduction to GPU Optimization Learn GPU optimization techniques using SYCL
₩ Launch 🖻	😇 Launch 🖾	≓ Launch II
Migrate from CUDA* to C++ with SYCL* Optimize apps from traditional CUDA environments S Launch C		
Gen AI Essentials		
Text-to-Image with Stable Diffusion A Creative Playground for Artists, Writers, and Engineers	Image-to-Image Generation with Stable Diffusion Perfect for artists and engineers who want to see their images transform in creative and unexpected ways.	Simple LLM Inference: Playing with Language Models A hands-on experience on language models and text generation, no technical background needed.
👺 Launch 🖾	₩ Launch 🕑	≓ Launch 🗹



## Intel<sup>®</sup> Compilers

### Intel<sup>®</sup> Compilers Going Forward

New underlying back-end Compilation Technology based on LLVM

New compiler technology available today in Intel® oneAPI Base & HPC Toolkit for DPC++, C++ and Fortran

Existing Intel proprietary "ILO" (ICC, IFORT) Compilation Technology compilers provided alongside new compilers

CHOICE! Continuity!

BUT Offload (DPC++ or OpenMP TARGET) supported only with new LLVM-based compilers

All Intel compilers are available on Juwels/Jureca: \$ module load Intel

### Intel<sup>®</sup> Compilers

Intel Compiler	Target	OpenMP Support	OpenMP Offload Support	Included in oneAPI Toolkit
Intel® C++ Compiler Classic, IL0 <i>icc/icpc/icl</i>	CPU	Yes	No	HPC
Intel® Fortran Compiler Classic, ILO <i>ifort</i>	CPU	Yes	No	HPC
Intel® Fortran Compiler, LLVM <i>ifx</i>	CPU, GPU	Yes	Yes	HPC
Intel® oneAPI DPC++/C++ Compiler, LLVM <i>dpcpp</i>	CPU, GPU, FPGA*	Yes	Yes	Base
Intel <sup>®</sup> oneAPI DPC++/C++ Compiler, LLVM <i>icx/icpx</i>	CPU GPU*	Yes	Yes	Base

#### Compiler Binary Compatible and Linkable!

tinyurl.com/oneapi-standalone-components

### Common optimization options

	Linux* icx (icc)
Disable optimization	-00
Optimize for speed (no code size increase)	-01
Optimize for speed (default)	-02
High-level loop optimization	-03
Create symbols for debugging	-g
Multi-file inter-procedural optimization	-ipo
Profile guided optimization (multi-step build)	-fprofile-generate (-prof-gen) -fprofile-use (-prof-use)
Optimize for speed across the entire program ("prototype switch")	-fast same as "-ipo -O3 -static -fp-model fast" (-ipo -O3 -no-prec-div –static -fp-model fast=2 -xHost)
OpenMP support	-fiopenmp (-qopenmp)

### SIMD: <u>Single Instruction</u>, <u>Multiple Data</u>

#### for (i=0; i<n; i++) z[i] = x[i] + y[i];</pre>



### Basic Vectorization Switches I

- Linux\*, OS X\*: -x<feature>
  - Might enable Intel processor specific optimizations
  - Processor-check added to "main" routine: Application errors in case SIMD feature missing or non-Intel processor with appropriate/informative message
  - Example: -xCORE-AVX512 (Juwels Xeon SKL)
- Linux\*, OS X\*: -ax<features>
  - Multiple code paths: baseline and optimized/processor-specific
  - Multiple SIMD features/paths possible, e.g.: -axSSE2, CORE-AVX512
  - Baseline code path defaults to -xSSE2

### **Basic Vectorization Switches II**

- Special switch for icc, Linux\*, OS X\*: -xHost
- Compiler checks SIMD features of current host processor (where built on) and makes use of latest SIMD feature available
- Code only executes on processors with same SIMD feature or later as on build host
- For AMD please check for the GNU flags e.g. -march=native



## **LLVM-BASED INTEL COMPILERS**

### What is ICX?

- Close collaboration with Clang\*/LLVM\* community
- ICX is Clang front-end (FE), LLVM infrastructure
  - PLUS Intel proprietary optimizations and code generation
- Clang FE pulled down frequently from open source, kept current
  - Always up to date in ICX
  - We contribute! Pushing enhancements to both Clang and LLVM
- Enhancements working with community better vectorization, opt-report, for example

#### tinyurl.com/blog-on-icx

### Major Changes Overview <u>tinyurl.com/icc-to-icx-migration-guide</u>

- LLVM is a different compilation technology. EXPECT differences
- Options:
  - icx -qnextgen-diag option to get a list of supported and unsupported options
- Use-fiopenmpor-fiopenmp-simd for OpenMP
- C/C++ Pragmas a lot of Intel proprietary ones not supported
  - enable Wunknown-pragmas to warn on unsupported pragmas
- INTEL\_LLVM\_COMPILER is defined instead of \_\_INTEL\_COMPILER

### Please switch to icx/icpx Compiler!

- Deprecation planed for 2024
- Check the user guide for supported flags:

https://www.intel.com/content/www/us/en/docs/dpcpp-cppcompiler/developer-guide-reference/2024-I/overview.html

- Check results and compare with icc/icpc results:
  - -fp-model=fast is the default
  - -fp-model=precise might help to reproduce previous results

### Build your own compiler (only for experts)

- Most of the features are included in the public llvm Intel version. You may test and contribute to the development.
- Interested? Check out: <u>https://intel.github.io/Ilvm-docs/GetStartedGuide.html</u>
- Build a clang compiler with latest features ahead of icpx and icx
- Some features may be missing in the public version
- Can also configure and build a CUDA/AMD backend compiler for offload to NVIDIA/AMD cards
- NVIDIA backend is also available for the oneAPI icpx version: <u>https://developer.codeplay.com/products/oneapi/nvidia/2024.1.0/guides/get-started-guide-nvidia.html</u>



## Which tool should I use?

### Performance Analysis Tools for Diagnosis



### Before dive to a particular tool..

- How to assess easily any potential in performance tuning?
- What to use on big scale not be overwhelmed with huge trace size, post processing time and collection overhead?
- Which tool should I use first?
- Answer: try Application Performance Snapshot (APS)
- Look for VTune module if available: \$ module load VTune

### APS Usage

Setup Environment

\$ source <path\_to\_vtune>/vtune\_vars.sh # or load module

#### **Run Application**

- \$ aps <application and args>
- MPI: \$ mpirun < mpi options > aps < application and args >

Generate Report on Result Folder

\$ aps -report <result folder>



Generate CL reports with detailed MPI statistics on Result Folder

\$ aps-report -<option> <result folder>

Rank> Rank	Volume (MB)	Volume(%)	Transfers
0023> 0024	84.35	1.56	13477
0025> 0026			
0024> 0025			
0021> 0022			
0022> 0023			
[filtered out 16	lines]		
0012> 0011	69.60		
0020> 0019			
0026> 0025	68.78		
0025> 0024			
0022> 0021	68.38		
[filtered out 17			
0016> 0015	58.81		
0028> 0027			
0007> 0008			
0030> 0031	54.74		
0006> 0007			
[filtered out 110			
TOTAL			
1 417.6	4 67	0.09	

### Application Performance Snapshot (APS)

Data in One Place: MPI+OpenMP+Memory Floating Point

#### Quick & Easy Performance Overview

- Does the app need performance tuning?
- MPI & non-MPI Apps<sup>+</sup>
- Distributed MPI with or without threading
- Shared memory applications

#### Popular MPI Implementations Supported

- Intel<sup>®</sup> MPI Library
- MPICH & Cray MPI

**Richer Metrics on Computation Efficiency** 

- CPU (processor stalls, memory access)
- FPU (vectorization metrics)



<sup>+</sup>MPI supported only on Linux<sup>\*</sup>

#### APS Command Line Reports – Advanced MPI statistics

 Data Transfers for Rankto-Rank Communication
 aps-report –x <result>

And many others – check • aps-report -help

Rank> Rank	Volume(MB)	Transfers	
0023> 0024	84.35	1.56	13477
0025> 0026	84.35	1.56	13477
0024> 0025	84.15	1.56	13477
0021> 0022	83.84	1.55	13477
0022> 0023	83.43	1.54	13477
[filtered out	16 lines]		
0012> 0011	69.60	1.29	13477
0020> 0019	69.29	1.28	13477
0026> 0025	68.78	1.27	13477
0025> 0024	68.38	1.27	13477
0022> 0021	68.38	1.27	13477
[filtered out	17 linesl	1.0	20111
0016> 0015	58.81	1.09	13477
0028> 0027	57.69	1.05	13477
	56 09	1.07	12477
	50.90	1.03	10477
	54.74	1.01	13477
0006> 0007	54.44	1.01	13477
[filtered out	1108 linesj		
======================================			
TOTAL	5403.22	100.00	1415619
AVG	4.67	0.09	1224

### **APS** potential issues

 If drivers are not available and for other hardware like AMD, you may use only mpi or/and openMP analysis:

#### \$ aps --collection-mode=mpi[,openmp]

If drivers are needed, you have to add this line to your batch job:

#### **#SBATCH** --disable-perfparanoid



## Intel<sup>®</sup> VTune<sup>TM</sup> Profiler

### Analyze & Tune Application Performance

#### Intel<sup>®</sup> VTune<sup>™</sup> Profiler

Advanced Hotspots + (2) INTEL VTUNE AMPLIFIER 2019										
Analysis Configuration Collection Log Summary Bottom-up Caller/Callee Top-down Tree Platform										
Grouping: Function / Call Stack										
		CPU Time 🔻					Context Swit	( ^		
Function / Call Stack	Effective Time	Spin Time	Overhead Time	Wait Time	Inactive Time	Preemption				
updateBusinessAccount	7.915s		0s	0s	0s	0.055s	934	ĺ		
main\$omp\$parallel_for@269	7.915s		0s	0s	0s	0.055s	934	]		
Kmp_invoke_microtas	7.915s		0s	0s	0s	0.042s	815			
updateBusinessAccount	0s		0s	0s	0s	0.013s	119			
updateCustomerAccount	7.766s		0s	0s	0s	0.052s	1,111			
kmpc_atomic_fixed8_add	2.772s		0s	0s						
kmpc_critical	Os		2.021s	0s	0s	0.014s	262	¥		
< ×	<						>			
0: <b>+ - r</b> r	5s 5.2s	5.4s 5.6s	5.8s	6s 6.	2s	✓ Thread	~	1		
OMP Worker Thread #2 (TI	and data sector of	المراقية المراجع	h	ا جاب ا	11.1.1 ^	Runi	ning			
OMP Worker Thread #3 (TI		and a second		<del>ة بن كارومية بدار</del> • • • • •	1	Contex	t Switches			
	ا ماليندية المريقة، منذ كمت اليد	And a list product and the list of the list.	الالبيد بليرائيك ورعد	بالبير بالخمام	A 110.1	Pre	emption			
rtmtest_openmp (TD: 12752)	وانتفاقح هعنيونيا البرجانة	بالمعادي فريني فالبر فالت	ابيلل توأسيله حايي	بنينالية هيبهاي	ملال، بها	Syı	nchronization			
OMP Worker Thread #1 (TI	CPU Time									
CPU Time	يتماننه محصفتان المتاري	والجريقاط فتستحيين سيارقه				🗹 📥 Spin	and Overh			
	<				>	CPU	_CLK_UNH			
FILTER 🕥 100.0% 🦹 Any Proc 🗸 Any Thread 🗸 Any Moc 🗸 Any L 🗸   User function Show inl 🗸 Function 🗸										

https://software.intel.com/content/www/us/en/develop/tools/vtu \_\_\_\_\_\_ne-profiler/get-started.html

- Accurately profile C, C++, Fortran\*, Python\*, Go\*, Java\*, or any mix
- Optimize CPU, threading, memory, cache, storage & more
- Take advantage of <u>Priority Support</u>
  - Connects customers to Intel engineers for confidential inquiries (paid versions)
- A more accessible user interface provides a simplified profiling workflow
- Smarter, faster Application Performance Snapshot: Analyze CPU utilization of physical cores, pause/resume, more... (Linux\*)

### Start a new Project

💯 Intel VTune Profiler	_	- D X	
Project Navigator	三 音 译 ▷ 占 🕩 🗅 💿 Welcome × Configure Analysis ×		
ttt	💯 Configure Analysis 📆	INTEL VTUNE PROFILER	
xtest	WHERE	HOW	- USE GUI
	Local Host 👻	Performance Snapshot 👻	• Or
	Launch Application -	Performance Snapshot ALGORITHM MICROARCHITECTURE	Command -Line
	Specify and configure your analysis target: an application or a script to execute.		
	Application:	Hotspots Anomaly Microarchitecture Detection Exploration	
	C:\Users\hbockhor\Documents\lssues\2020\OSC_04816644\test_intel\source_cor	(preview)	
	Application parameters:	Memory Access	
	Ose application directory as working directory	PARALLELISM I/O	
	Advanced		Cat Carara and Lina
		Threading HPC Input and Output	Get Command-Line
		Characterization (preview)	
		ACCELERATORS PLATFORM AHALYSES	
		GPU Offload GPU System Throttling (preview) Compute/Media Overview (preview)	
		Hotspots	



## INTEL<sup>®</sup> ADVISOR

### Intel<sup>®</sup> Advisor – Vectorization Advisor

#### Get breakthrough vectorization performance

- 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

- The data and guidance you need:
  - Compiler diagnostics + Performance Data + SIMD efficiency
  - Detect problems & recommend fixes
  - Loop-Carried Dependency Analysis
  - Memory Access Patterns Analysis

Elapsed time: 70.29s 🧩 🙆 Vectorized 🙆 Not Vectorized 🖉																
FIL	TER: All Modules 🔻 All Sources 🔻	Loc	ops And Functio	ns 🔻 All Threads 🔻								INT	ELA	IVISOR 20	18	Optimize for
P	Summary 😂 Survey & Roofline 🔌 R	lefinem	ent Reports													AVX-512
8	E Exaction Coll Sites and Leave		@ Vector	Calf Time	Total	T	FLOPS	$\gg$	Why No	Vectoriz	ed Loops		$\gg$	Trip 🚿	1	with/without
ŎĘ	- Function Call Sites and Loops		Issues	Ser Time*	Time	туре	GFLOPS	AI	Vectorization?	Vector	Efficiency	Gain	VL	Counts		
I.	⊕ [loop in S252 at loops90.f:1172]	-	9 1 Possible	3.129s 7.0%	3.129	Vectorized	0.191	0.115	🖬 1 vectorizat	AVX2	17%	1.36x	4; 8	99; 6; 1; 1		accessto
	[loop in S2101 at loops90.f:1749]	✓	2 Possible	2.765s 6.2%	2.765s	Scalar	0.1421	0.067	vectorizatio					12		AVX-512
	∃⊡ [loop in s442_\$omp\$parallel_for		💡 1 Ineffecti	1.492s 3.4%	1.492s	Vectorized+	0.5861	0.165		AVX2	14%	1.09x	8	30; 1; 3		
	f _svml_sinf8_19			1.108s 2.5%	1.108s	Vector Funct	3.9111	0.156		AVX2						hardware
	⊕ [loop in S353 at loops90.f:2381]		9 1 Possible	0.989s 2.2%	0.989s	Vectorized (	2.0231	0.134		AVX2	27%	2.16x	8	6; 4; 1	~	
	< >	<												>		

Part of oneAPI Base Toolkit

software.intel.com/advisor



#### Critical Data Made Easy Loop Trip Counts

#### Knowing the time spent in a loop is not enough!



### What is a Roofline Chart?

- A Roofline Chart plots application performance against hardware limitations.
  Performance (GFLOPS)
  Q \* × © | Use Single-Threaded Roofs
  - Where are the bottlenecks?
  - How much performance is being left on the table?
  - Which bottlenecks can be addressed, and which should be addressed?
  - What's the most likely cause?
  - What are the next steps?



Roofline first proposed by University of California at Berkeley: <u>Roofline: An Insightful Visual Performance Model for Multicore Architectures</u>, 2009 Cache-aware variant proposed by University of Lisbon: <u>Cache-Aware Roofline Model: Upgrading the Loft</u>, 2013

### Advisor Resources

Intel<sup>®</sup> Advisor

- Product page overview, features, FAQs...
- What's New?
- Training materials <u>Cookbooks</u>, <u>User Guide</u>, <u>Tutorials</u>
- Support Forum
- Online Service Center Secure Priority Support

#### Additional Analysis Tools

- <u>Intel® VTune™ Profiler</u> performance profiler
- Intel® Inspector memory and thread checker/ debugger
- Intel® Trace Analyzer and Collector MPI Analyzer and Profiler

#### **Additional Development Products**

Intel<sup>®</sup> oneAPI Toolkits



### Online Resources

- Intel<sup>®</sup> MPI Library product page.
  - www.intel.com/go/mpi
- Intel<sup>®</sup> Clusters and HPC Technology forums
  - <u>http://software.intel.com/en-us/forums/intel-clusters-and-hpc-technology</u>

### Intel Modules installed on Juwels

- Compiler:
  - : check available: default: APS: check available:

default:

- VTune + APS:
- Advisor: check available default:
- Intel MPI: check available: default:
- Intel MKL: check available: default:

\$ module spider Intel \$ module load Intel \$ module spider vtune \$ module load VTune \$ module spider advisor \$ module load Advisor

\$ module spider intelMPI \$ module load IntelMPI

\$ module spider mkl \$ module load imkl

### How to start?

- Compile with minimal options and run with APS (will provide tuning tips)
- Compile with -xhost and and check timing and APS report
- Optional! Compile with –xhost and –no-vec disables vectorization. Compare with previous timing
- Use: VTune Profiler: \$ module load VTune/<version>
- Use: Advisor: \$ module load Advisor/<version>
- Google for Intel related topics  $\rightarrow$  Intel Developer Zone etc.
- For APS/VTune add to your batch job: #SBATCH --disable-perfparanoid
- Please set thread affinity e.g.: \$ export KMP\_AFFINITY=scatter,verbose This can speed up OMP programs up to 10X!
- Any questions: Heinrich.Bockhorst@Intel.com

### Notices & Disclaimers

Intel technologies may require enabled hardware, software or service activation. Learn more at intel.com or from the OEM or retailer.

Your costs and results may vary.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

**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. <a href="https://software.intel.com/en-us/articles/optimization-notice">https://software.intel.com/en-us/articles/optimization-notice</a>

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. See backup for configuration details. For more complete information about performance and benchmark results, visit <u>www.intel.com/benchmarks</u>.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See configuration disclosure for details. No product or component can be absolutely secure.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.



### Intel<sup>®</sup> oneAPI Math Kernel Library (oneMKL)

#### What's Inside Intel<sup>®</sup> oneAPI Math Kernel Library (oneMKL)



What's New for Intel® oneAPI Math Kernel Library(oneMKL) 2021.2-2022.0

- Introduced GPU support for the following new functionality:
  - BLAS Batch & copy for unified shared memory(USM) & buffer APIs
  - Vector Statistics RNG multinomial, PoissonV, hypergeometric, negative binomial and binomial distributions.
  - BLAS Added SYCL support for in-place and out of place matrix copy/transposition
  - LAPACK Enabled C/Fortran OpenMP offload support for select functions.
  - Sparse BLAS Added support for variance matrix-matrix multiplication operations.
- General performance optimizations
- For detailed information please refer to the oneMKL <u>Release Notes</u>

### **Basic Vectorization Switches III**

- Special switch in addition to CORE-AVX512: -qopt-zmm-usage=[keyword]
  - [keyword] = [high | low] ; Note: "low" is the default
  - Why choosing a defensive vectorization level?

Frequency drops in vectorized parts. Frequency does not immediately increase after the vectorized loop. Too many small vectorized loops will decrease the performance for the serial part.

### Next steps

- Toolkits are free but maybe too large ( > 10 GB). For this workshop you may download to your laptop: VTune, Advisor, Inspector
- Standalone tools download: <u>https://software.intel.com/content/www/us/en/develop/articles</u> /oneapi-standalone-components.html

#