SYCL, DPC++, XPUs, oneAPI

a view from Intel

James Reinders, engineer

IWOCL / SYCLcon 2021

SYCL, DPC++, XPUs, oneAPI

a book to teach
SYCL programming,

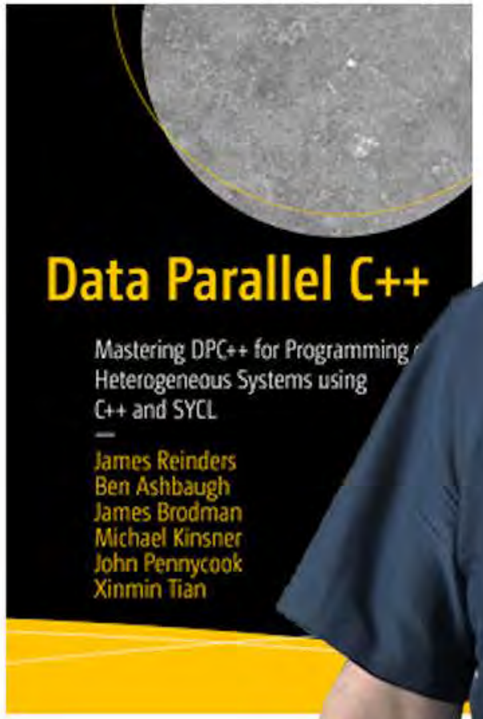

many thanks
to those of
you who helped
with your feedback!

What is “Data Parallel C++” ?

DPC++ is an open-source project to add SYCL to LLVM.

Join in the fun!

Help Out!

Try it out!

Why name “Data Parallel C++”

“SYCL” was taken. ☺

DPC++ name is perfectly descriptive.

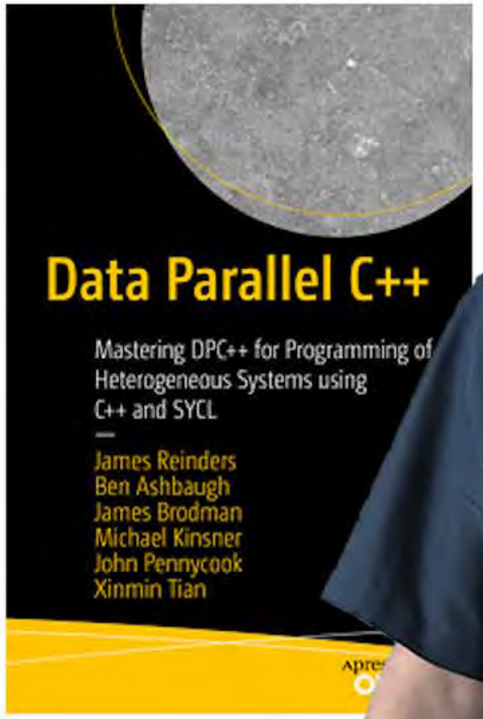

Adding Data Parallelism to C++

"SYCL 2020's primary goal is to achieve closer convergence with ISO C++, furthering our work to bring parallel heterogeneous programming to modern C++ through open standards."

- Michael Wong

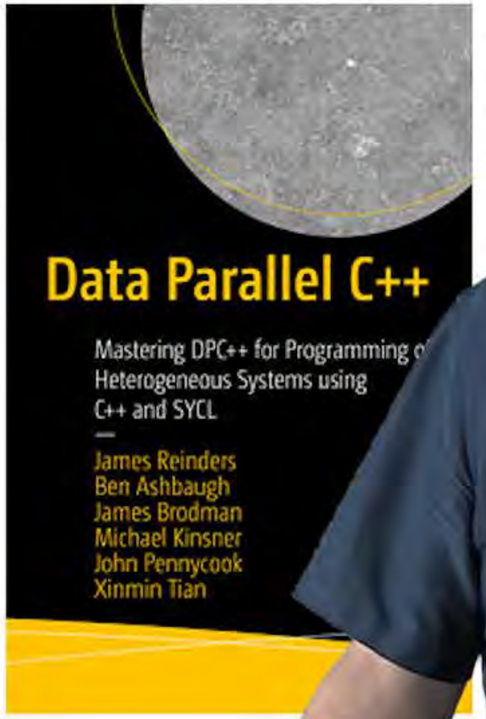

goal

bring

SYCL into LLVM

What is an XPU?

XPU ≈ *.* processing units

e.g., CPU, GPU, FPGA, DSP, ASIC

Our Quest

make

heterogeneous programming

ubiquitous.

Our Quest

make

heterogeneous programming

(XPU programming)

ubiquitous.

Can we *really* program XPUs?

1. Freedom:

   Use any XPU that I choose.

   (regardless of XPU type or vendor)

Can we *really* program XPUs?

1. Freedom
2. Value:

   Regardless of my XPU choice, I consistently can obtain a reasonable level of performance.

   (regardless of XPU type or vendor)

Can we *really* program XPUs?

1. Freedom
2. Value
3. Trustworthy:

   My coding choices can be made with confidence, and my code is maintainable.

SYCL embraces this vision

vision / goal:

- ✓ Freedom
- ✓ Value
- ✓ Trustworthy

SYCL explicitly supports XPU-specific codi
and tuning.

embracing also: oneAPI

vision / goal:

- ✓ Freedom
- ✓ Value
- ✓ Trustworthy

Like SYCL, oneAPI explicitly supports XPU-specific coding and tuning.

. processing units (XPUs)

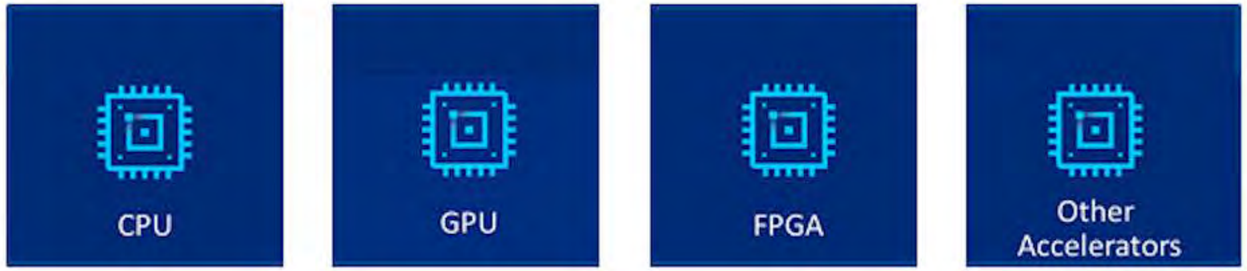

programming

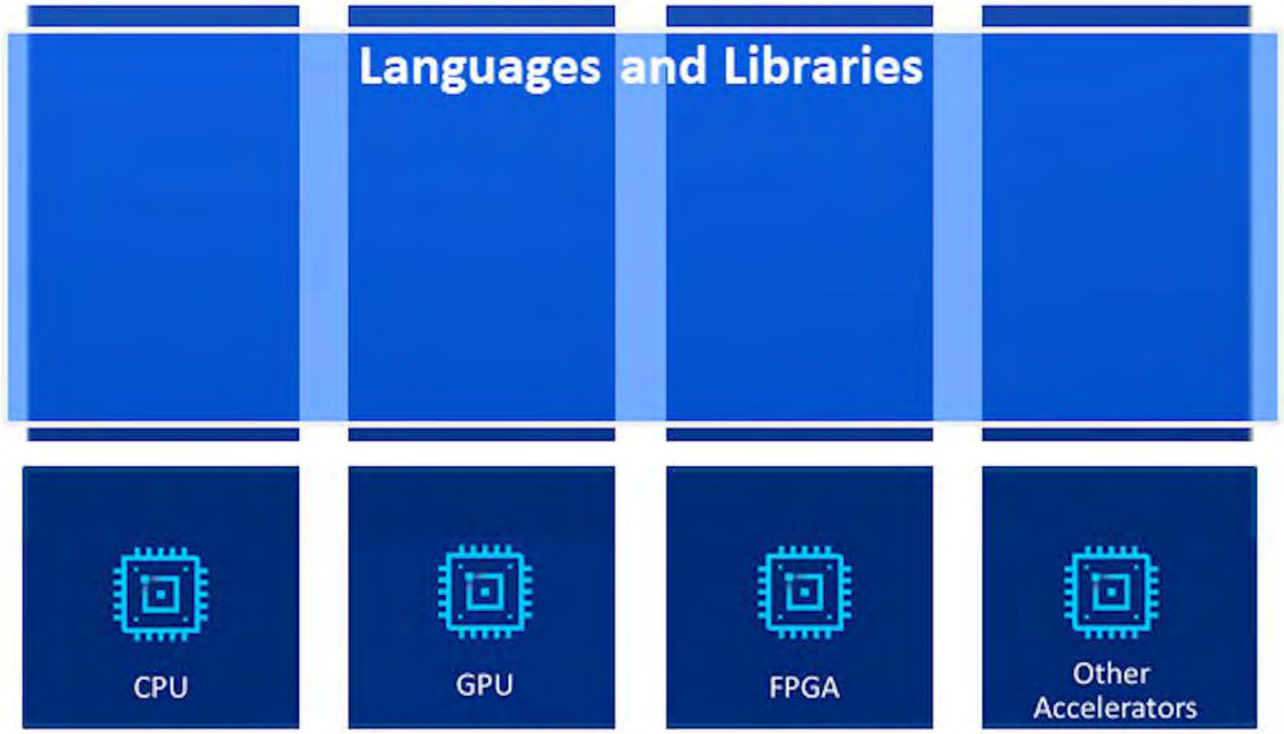

higher level programming

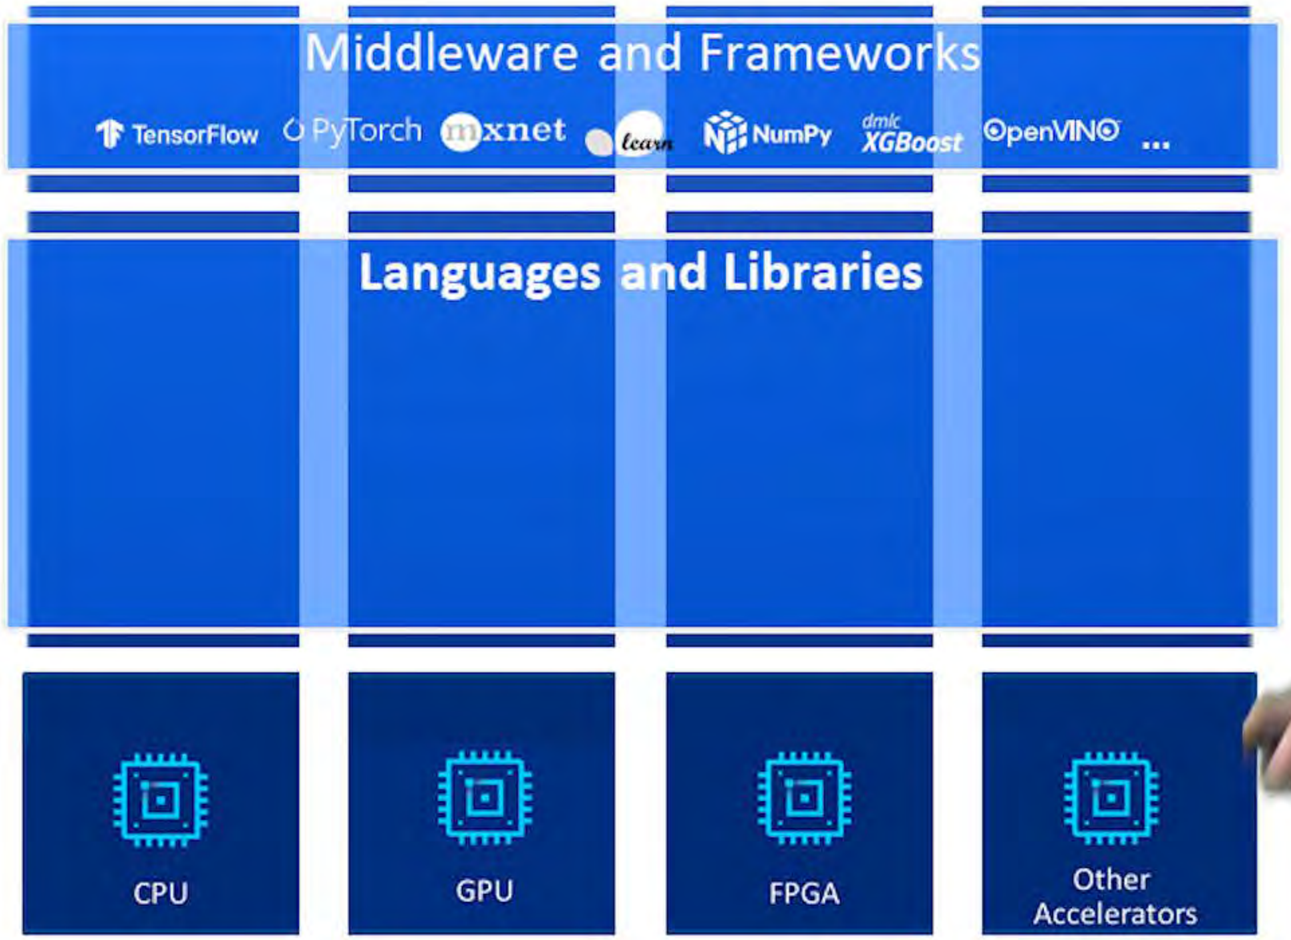

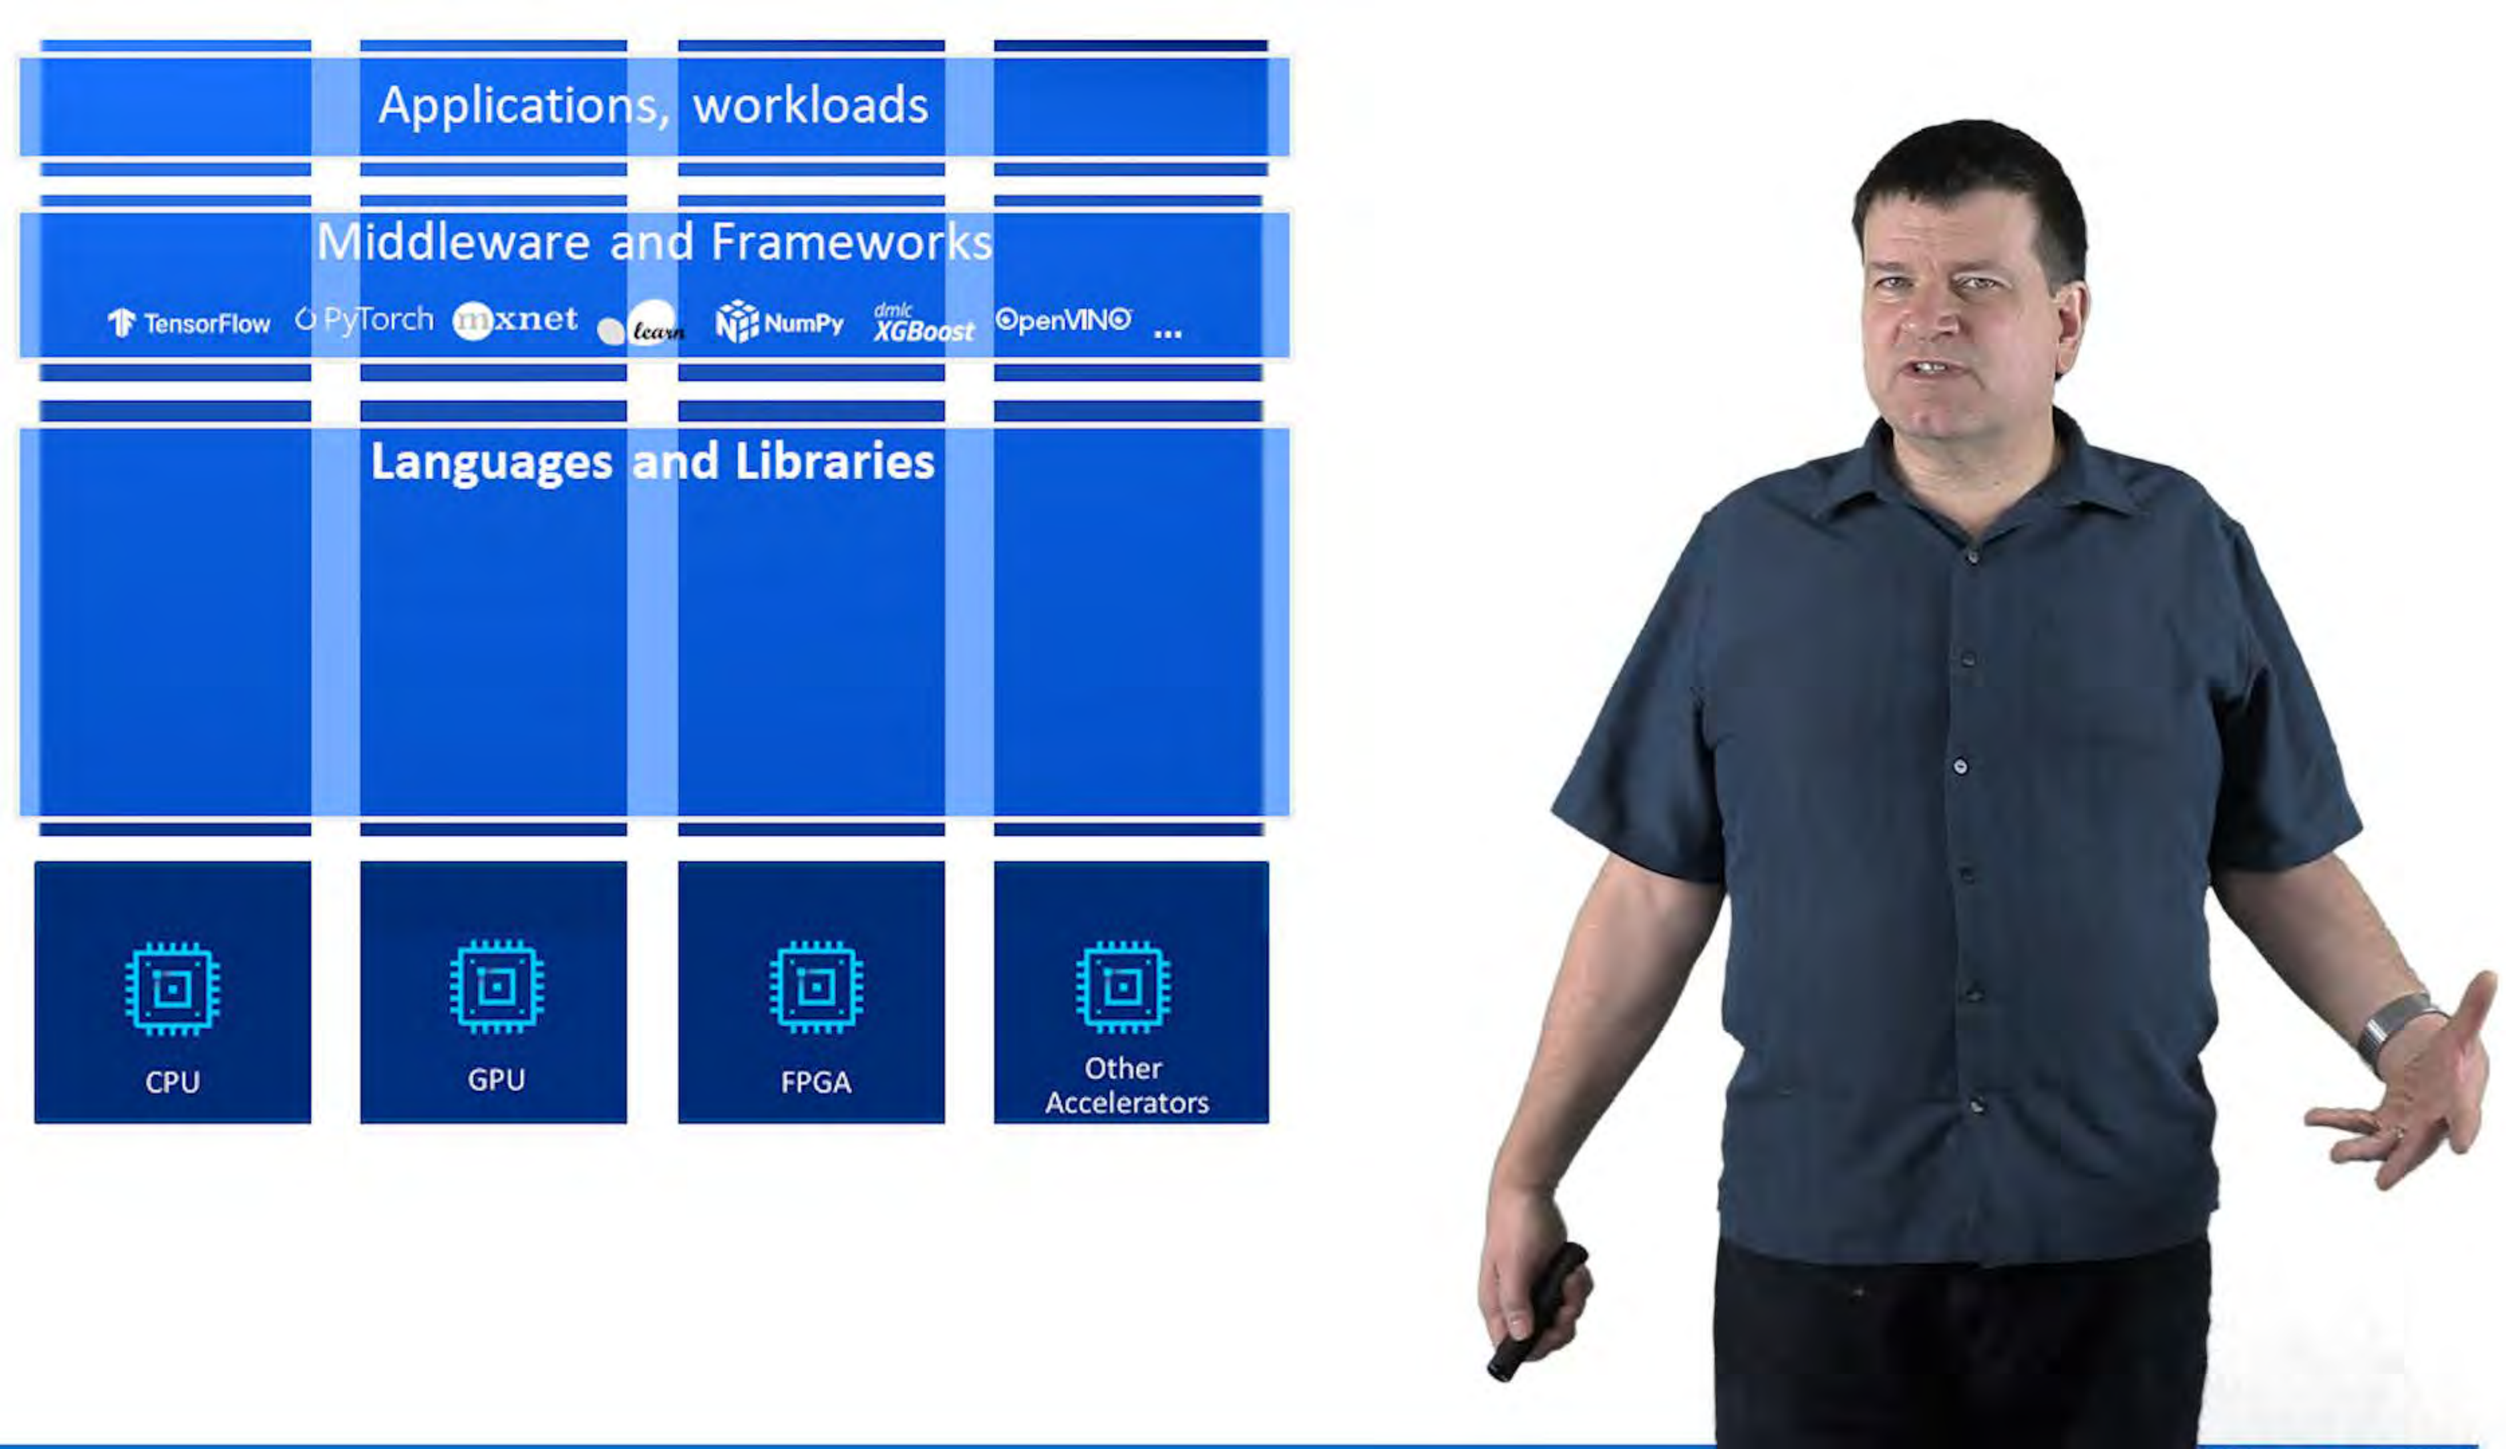
Applications, workloads
Middleware and Frameworks
TensorFlow
PyTorch
mxnet
learn
NumPy
dmlc XGBoost
OpenVINO
...
Languages and Libraries
CPU
GPU
FPGA
Other Accelerators

Applications, workloads

Middleware and Frameworks

TensorFlow PyTorch mxnet learn NumPy dmlc XGBoost OpenVINO ...

Languages and Libraries

oneAPI

An open specification and initiative to standardize programming of accelerated processing units (XPUs)

oneAPI

XPUs

CPU GPU FPGA Other Accelerators

Applications, workloads

Middleware and Frameworks

TensorFlow PyTorch mxnet scikit-learn NumPy dmlc XGBoost OpenVINO ...

Languages and Libraries

oneAPI

oneAPI

Intel's first product implementation of oneAPI released in 2020.
Free downloads.

XPUs

CPU GPU FPGA Other Accelerators

Applications, workloads

Middleware and Frameworks

TensorFlow PyTorch mxnet learn NumPy dmlc XGBoost OpenVINO ...

Intel® oneAPI

Languages

Hardware Abstraction Layer

XPUs

CPU GPU FPGA Other Accelerators

Applications, workloads
Middleware and Frameworks
TensorFlow
PyTorch
mxnet
Intel® oneAPI
Languages
Hardware Abstractio
XPUs
CPU
GPU
FPG

Data Parallel C++
Mastering DPC++ for Programming of Heterogeneous Systems using C++ and SYCL
James Reinders
Ben Ashbaugh
James Brodman
Michael Kinsner
John Pennycook
Xinmin Tian
Apress
open

Applications, workloads

Middleware and Frameworks

TensorFlow PyTorch mxnet learn NumPy dmlc XGBoost OpenVINO ...

Intel® oneAPI

Languages

Libraries

Hardware Abstraction Layer

XPUs

CPU GPU FPGA Other Accelerators

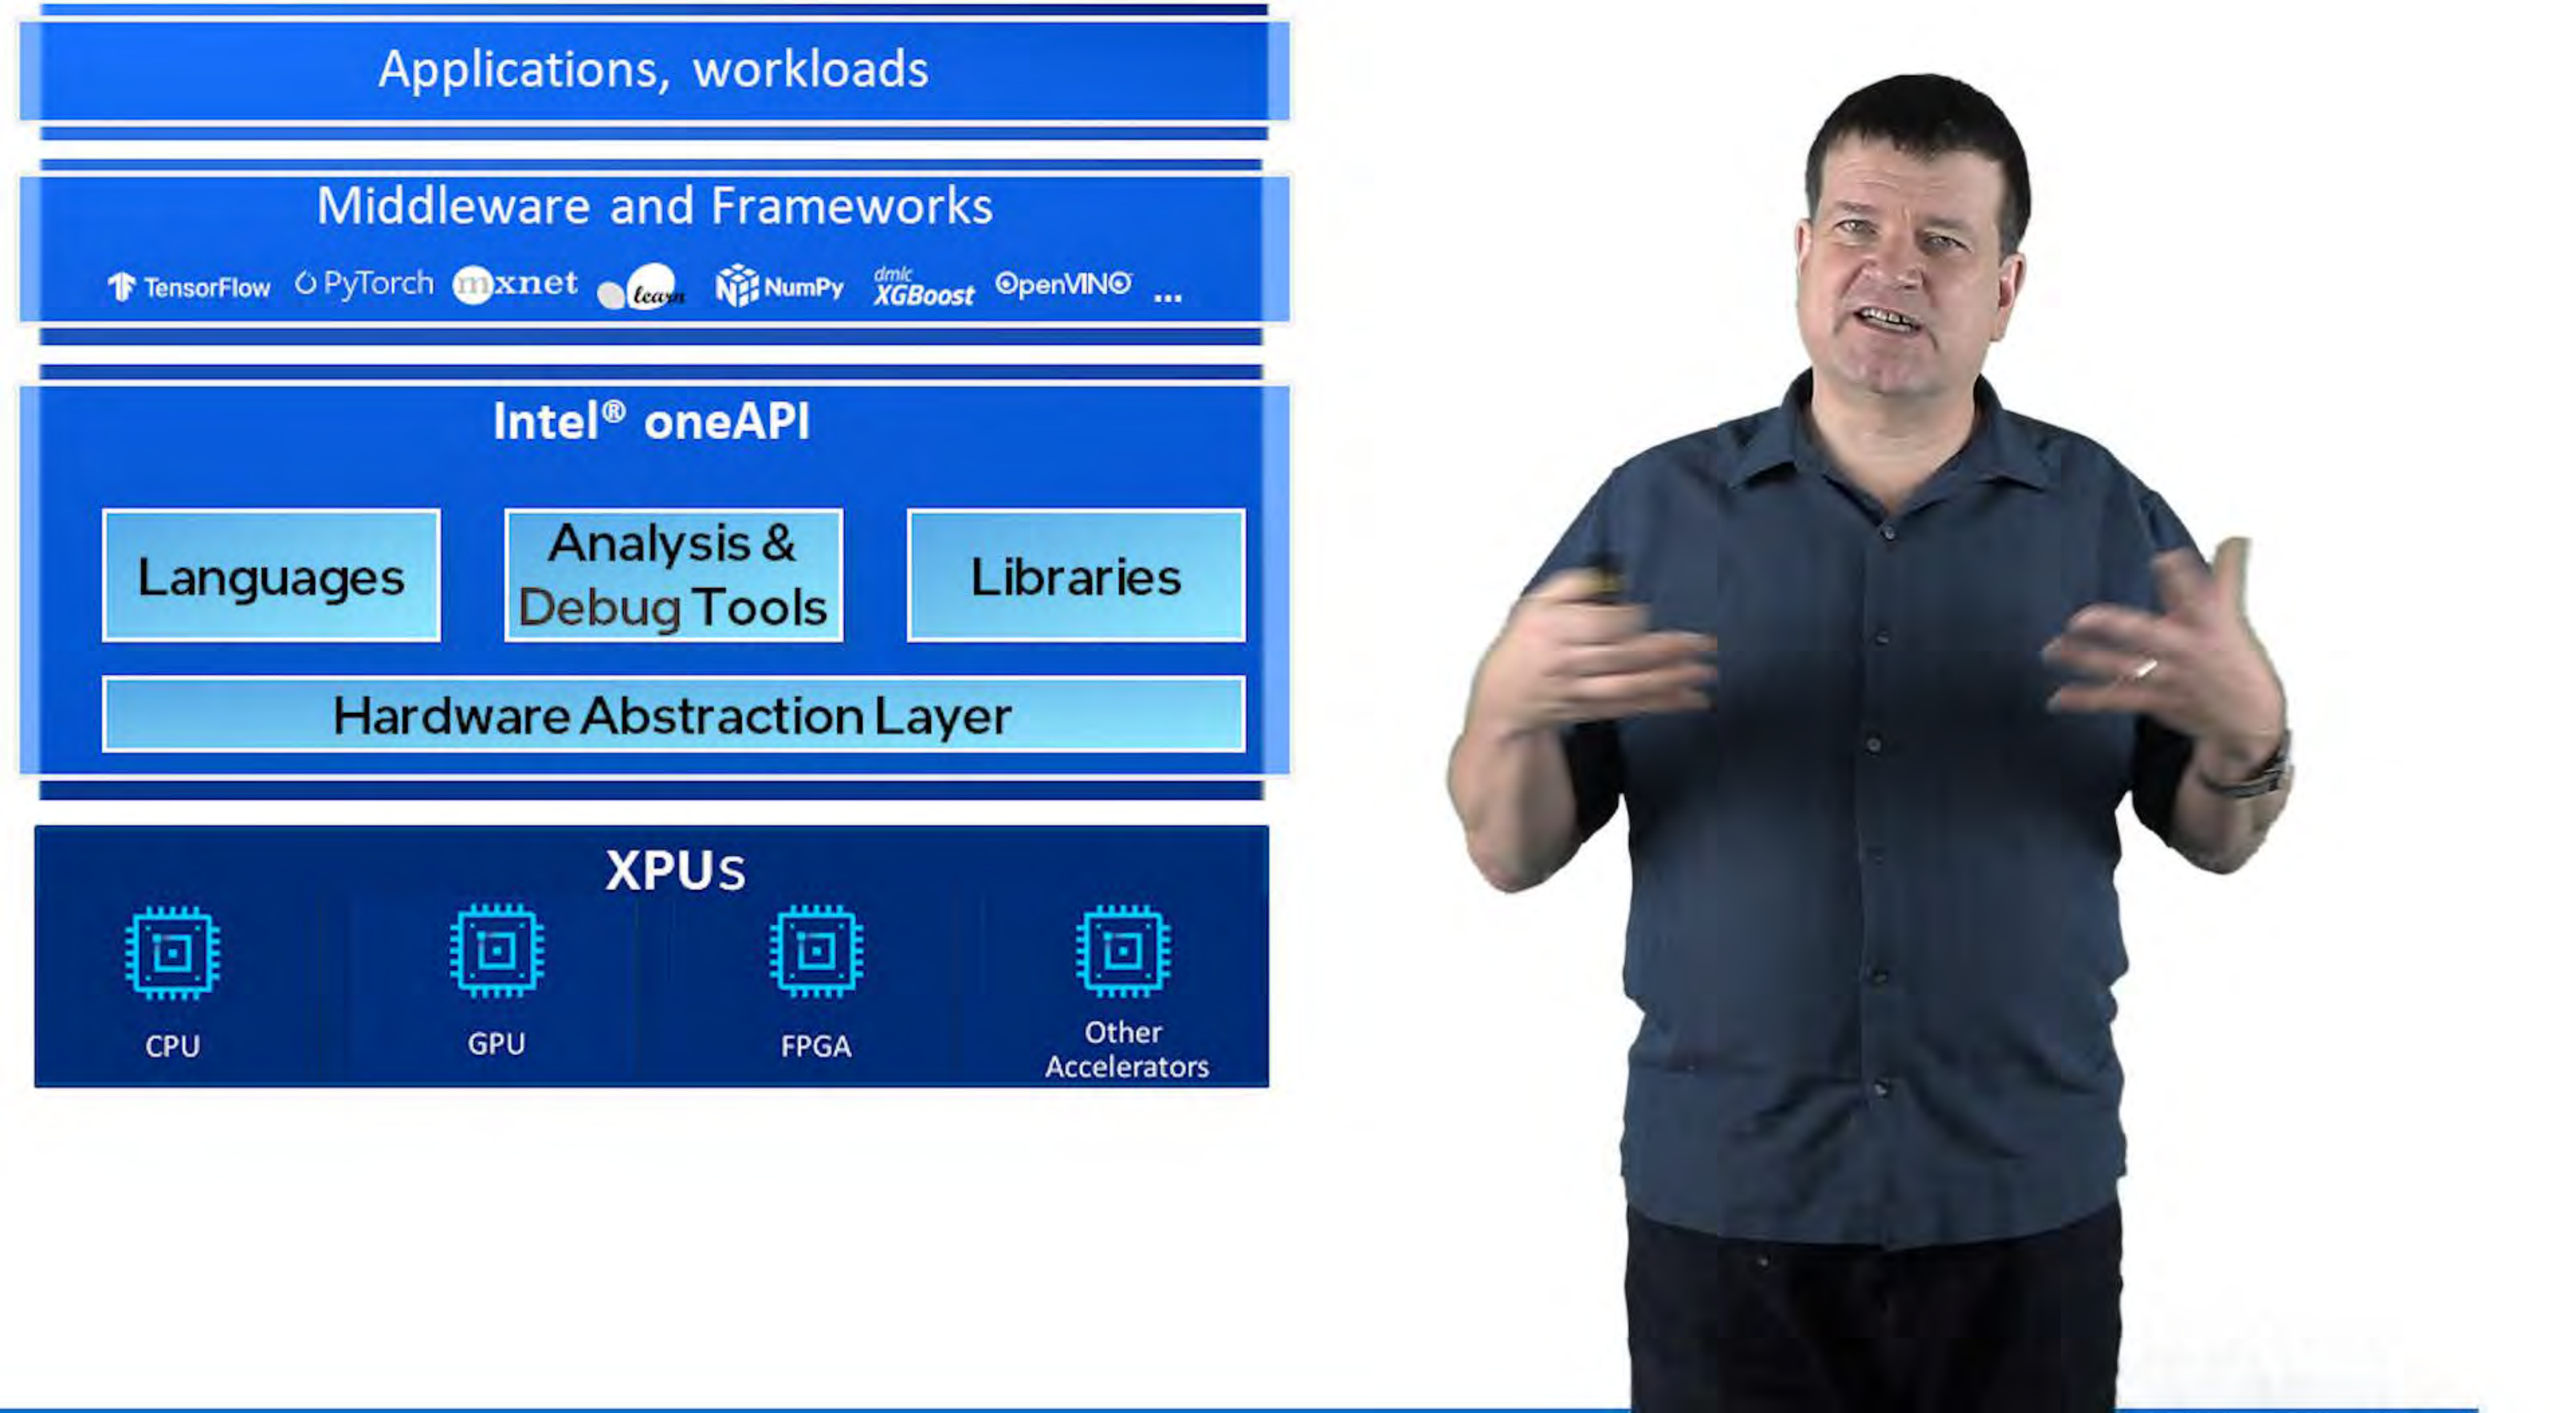
Applications, workloads
Middleware and Frameworks
TensorFlow
PyTorch
mxnet
learn
NumPy
dmlc XGBoost
OpenVINO
...
Intel® oneAPI
Languages
Analysis & Debug Tools
Libraries
Hardware Abstraction Layer
XPUs
CPU
GPU
FPGA
Other Accelerators

For instance… We see this coming:
a profound impact on FPGA programming
in addition to SYCL, we have familiar and powerful profilers, debuggers, library APIs…

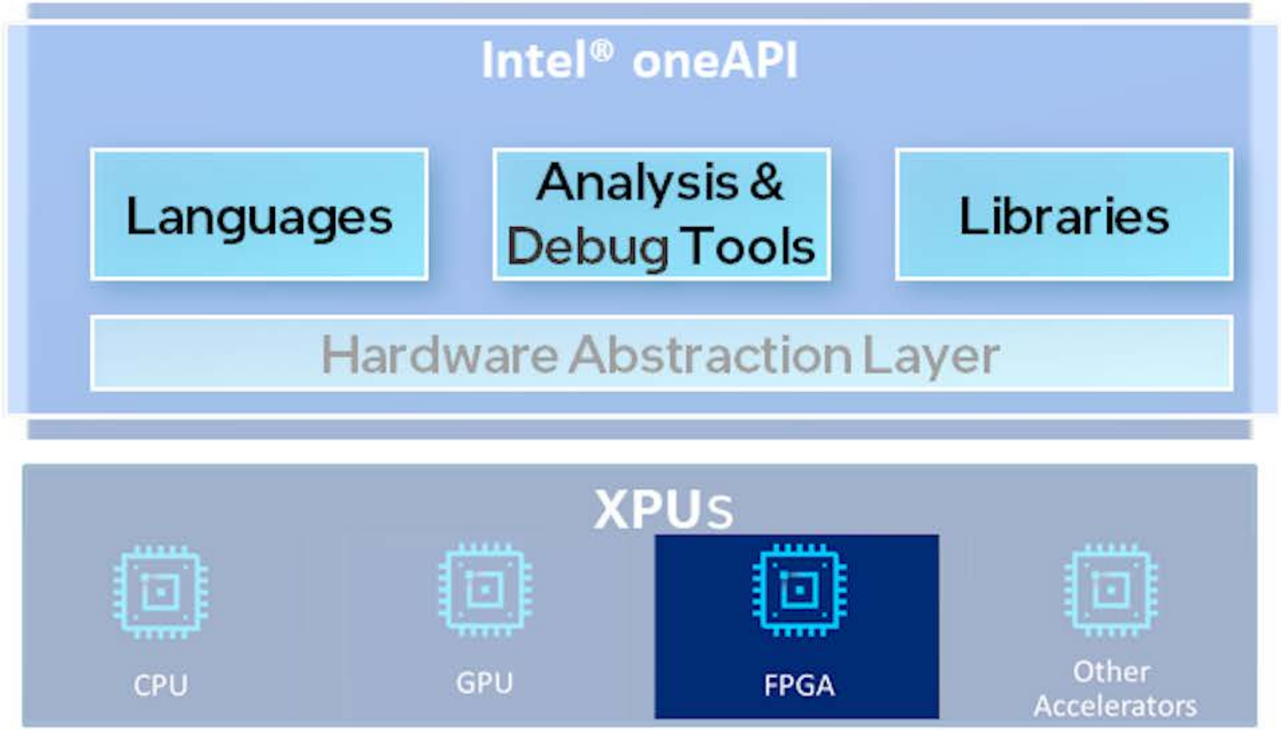

high enough performance

+

productivity

+

ability to adapt/adjust

=

better results in practice

Applications, workloads

Middleware and Frameworks

TensorFlow PyTorch mxnet learn NumPy dmlc XGBoost OpenVINO ...

Intel® oneAPI

Intel® oneAPI Toolkits

A complete set of proven developer tools expanded from CPU to XPU

oneAPI

Intel® oneAPI Base Toolkit

Native Code Developers

A core set of high-performance tools for building C++, Data Parallel C++ applications & oneAPI library-based applications

Add-on Domain-specific Toolkits

Specialized Workloads

Intel® oneAPI Tools for HPC

Deliver fast Fortran, OpenMP & MPI applications that scale

Intel® oneAPI Tools for IoT

Build efficient, reliable solutions that run at network's edge

Intel® oneAPI Rendering Toolkit

Create performant, high-fidelity visualization applications

Toolkits powered by oneAPI

Data Scientists & AI Developers

Intel® AI Analytics Toolkit

Accelerate machine learning & data science pipelines with optimized DL frameworks & high-performing Python libraries

Intel® Distribution of OpenVINO™ Toolkit

Deploy high performance inference & applications from edge to cloud

Latest version is 2021.2

Applications, workloads

Middleware and Frameworks

TensorFlow PyTorch mxnet learn NumPy dmlc XGBoost OpenVINO ...

Intel® oneAPI

Languages | Analysis & Debug Tools | Libraries

Hardware Abstraction Layer

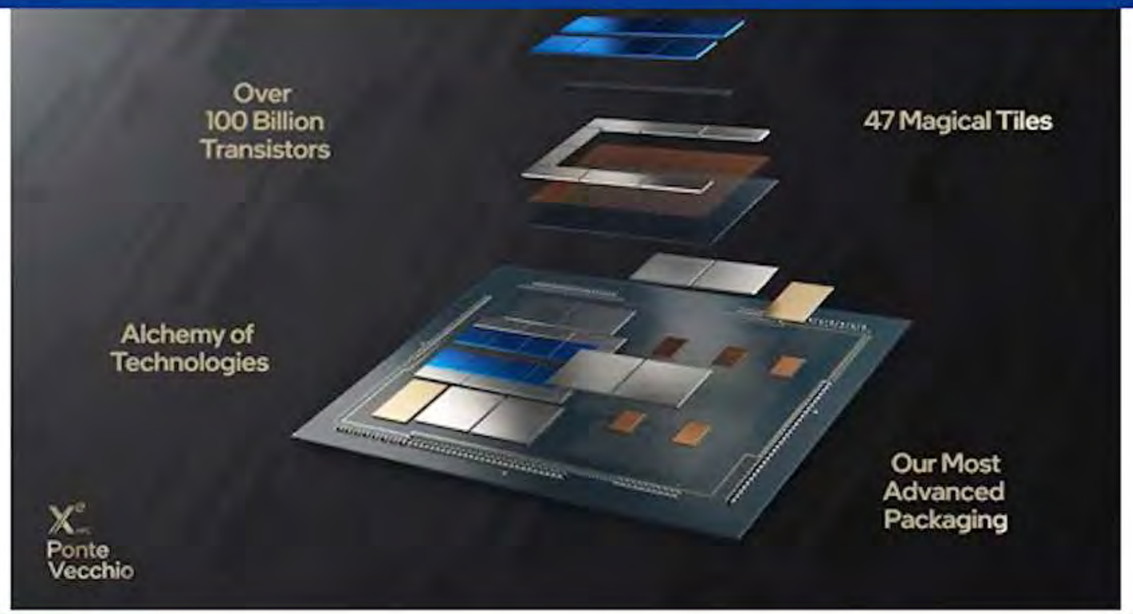

Applications, workloads

Middleware and Frameworks

TensorFlow PyTorch mxnet learn NumPy dmlc XGBoost OpenVINO ...

Intel® oneAPI

Languages | Analysis & Debug Tools | Libraries

Hardware Abstraction Layer

more info: https://www.alcf.anl.gov/aurora

Together – let's help the industry embrace heterogeneity

- to adopt the best XPUs for the job openly and universally

much of
oneAPI rests
solidly on SYCL

❤ SYCL 2020 ❤

We are committed to SYCL.

thank you
SYCL committee

(*and* thank you for

"my heros"

the Intel engineers,

who sold Intel on this)

intel®

Our support for an open XPU future is sincere, runs deep, and will help the entire industry.

Please work with me, to ensure our enthusiasm is always a positive force.

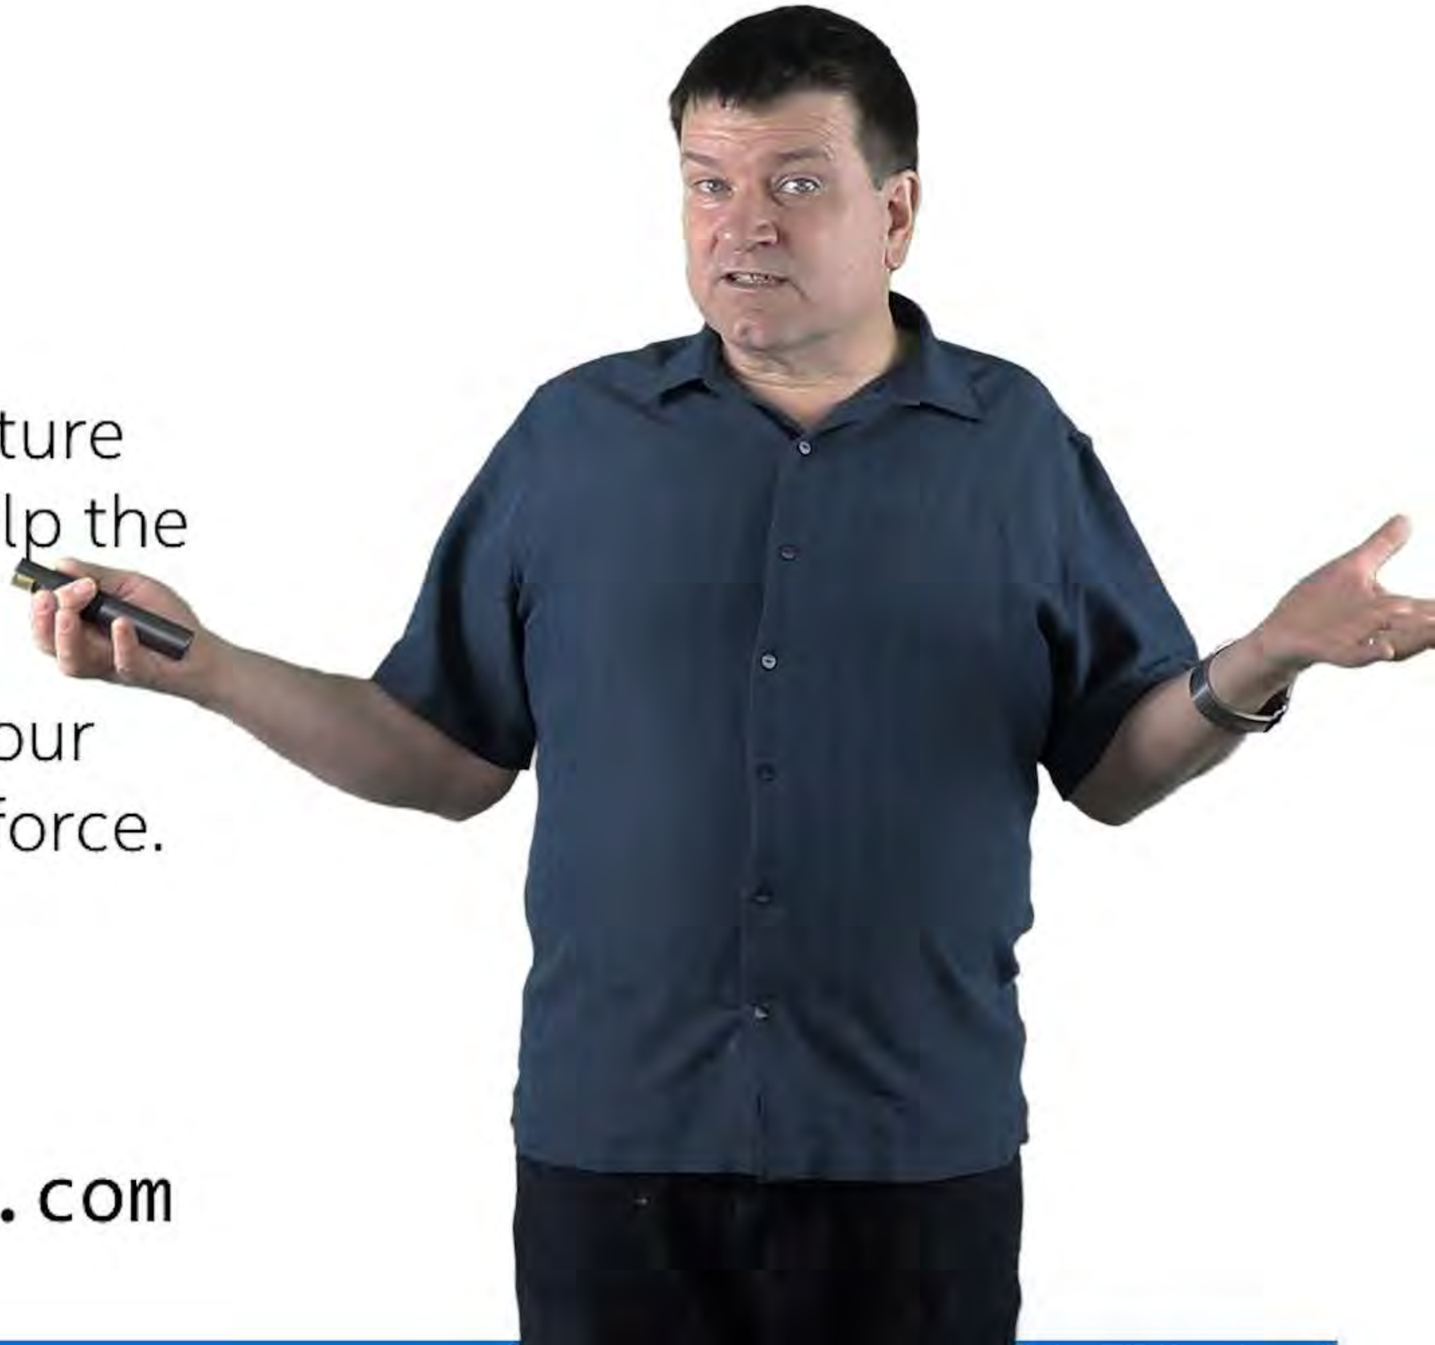

james.r.reinders@intel.com

Thank you

@jamesreinders #xpublog

james.r.reinders@intel.com

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 information go to www.intel.com/benchmarks.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. Check with your system manufacturer or retailer or learn more at intel.com.

Intel, the Intel logo, are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.
*Other names and brands may be claimed as the property of others.

© Intel Corporation.

intel

Thank you

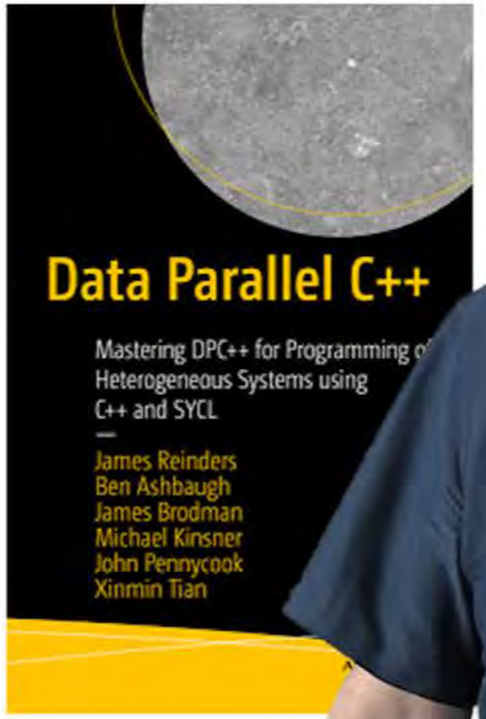

@jamesreinders #xpublog

james.r.reinders@intel.com