

oneAPI for IXPUG

Geoff Lowney
Intel Senior Fellow

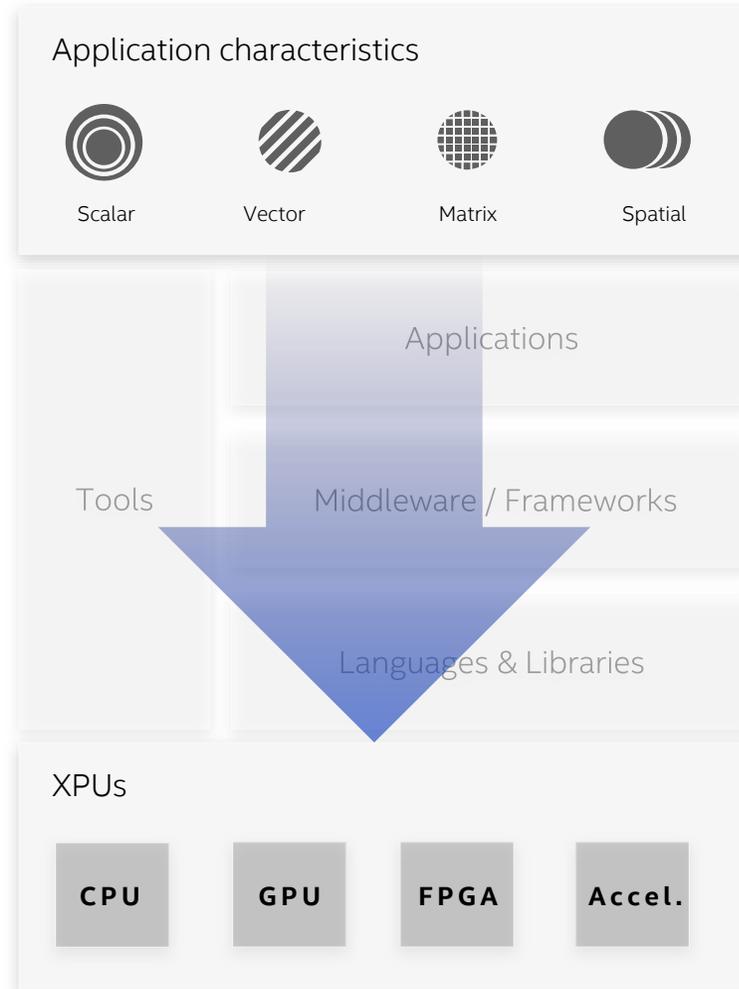
IXPUG Denver
November 20, 2019

We need your feedback

- A unique opportunity to steer the parallel programming ecosystem
- A problem worth solving
 - Multi-architecture, avoiding lock-in to 1 specific hardware architecture
 - Direct and library-based programming
 - Extending existing models
 - Performant
- Your input is important

PROGRAMMING CHALLENGES FOR MULTIPLE ARCHITECTURES

Variety of compute patterns



Diverse set of data-centric hardware required

No common programming language or APIs

Inconsistent tool support across platforms

Each platform requires unique software investment

What if

- For application developers
 - A single highly productive **language** and a rich set of **libraries** was available **everywhere**
- For HW vendors
 - Tap into **existing SW ecosystem**
 - Define their own “to the metal” language
- For tool developers
 - Leverage existing **compiler infrastructure** and tooling interfaces

oneAPI

A unified programming model to simplify development across diverse architectures

Common developer experience across architectures

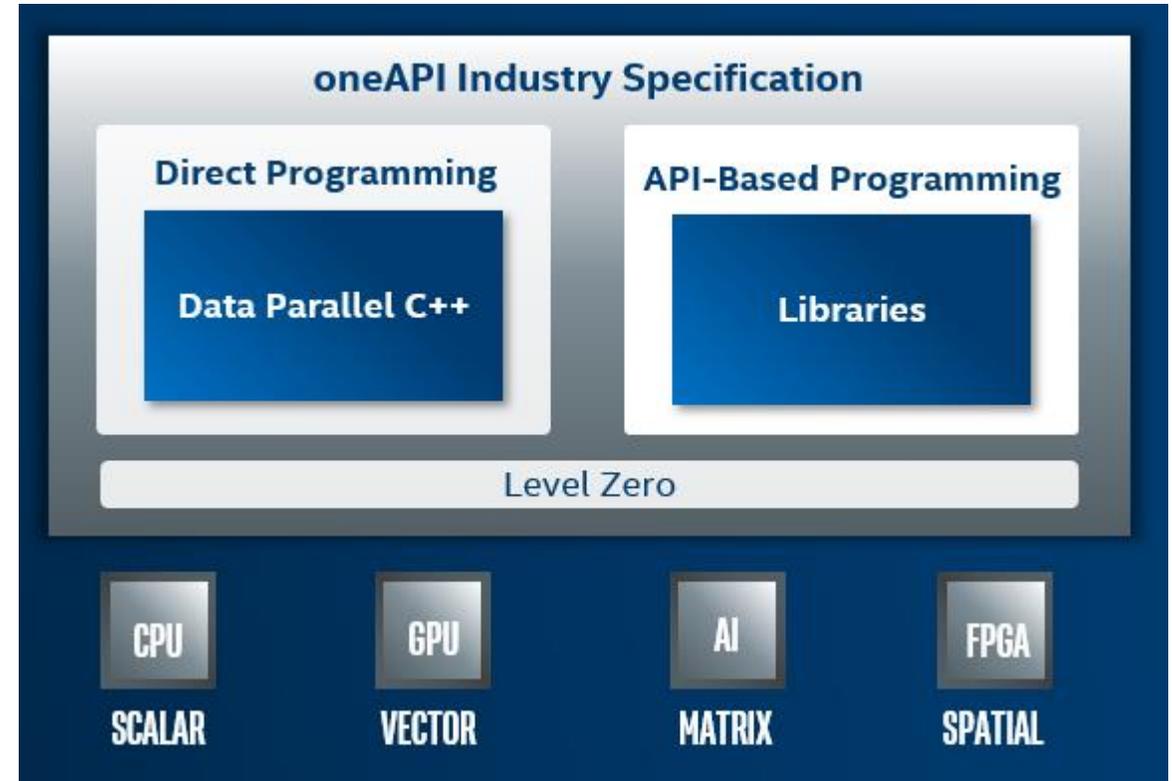
Unified and simplified language and libraries for expressing parallelism

Uncompromised native high-level language performance

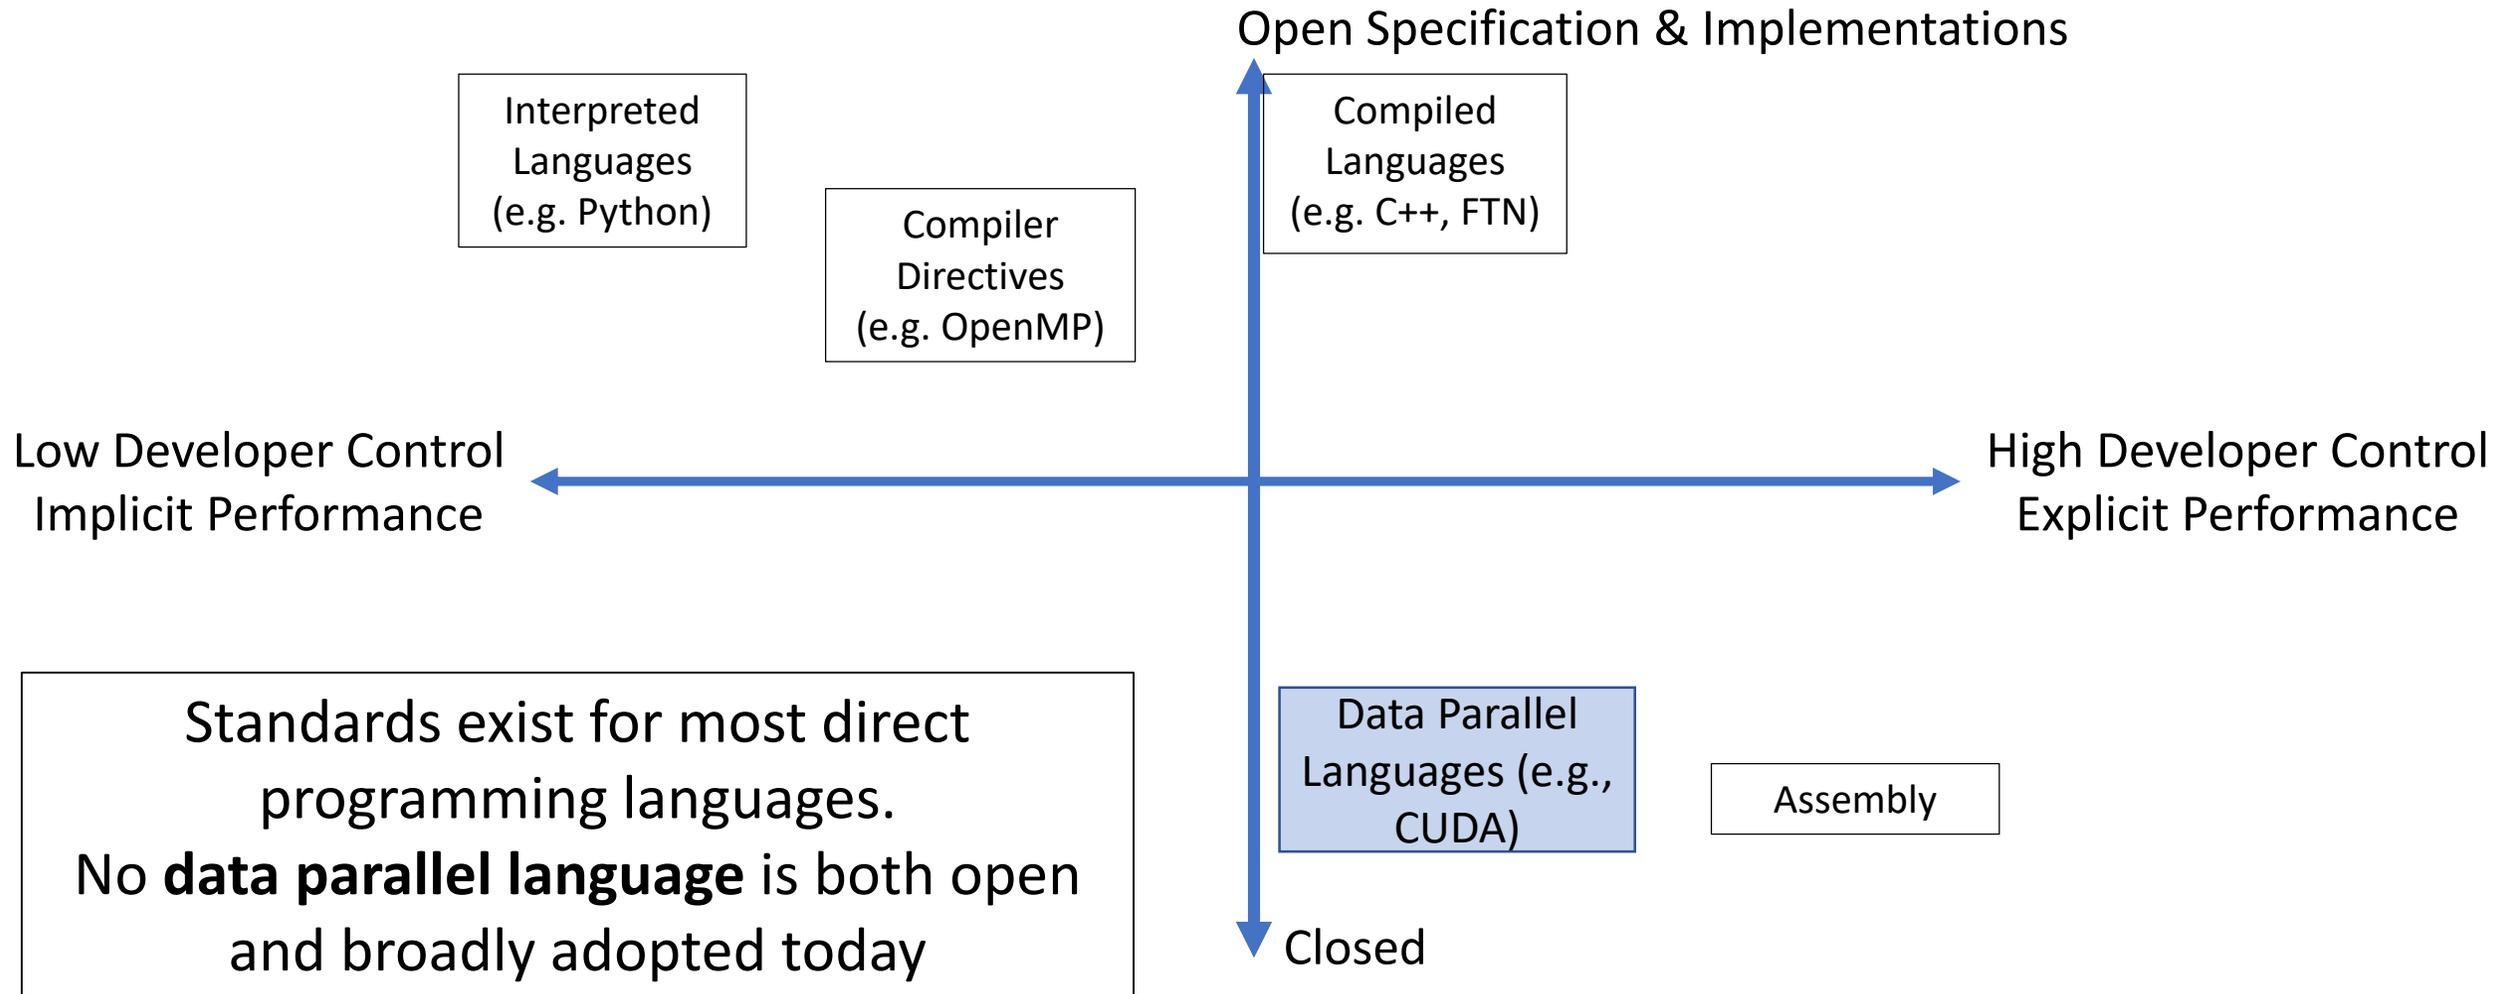
Interoperates with existing languages and libraries

Support for CPU, GPU, AI and FPGA

Based on industry standards and open specifications

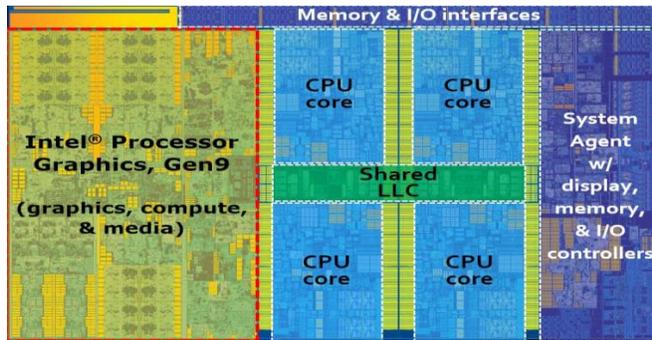


The direct programming problem

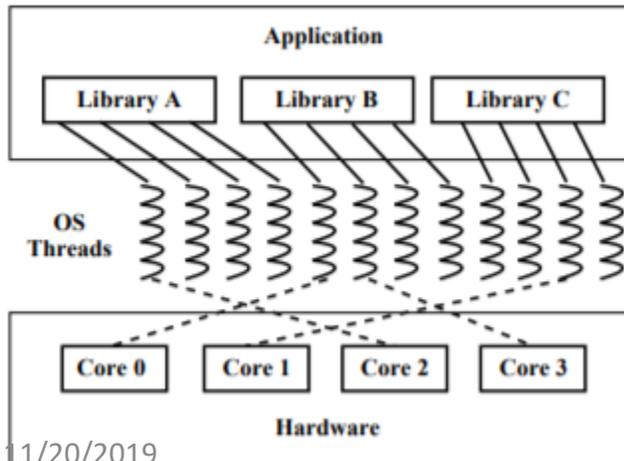


The API programming problem

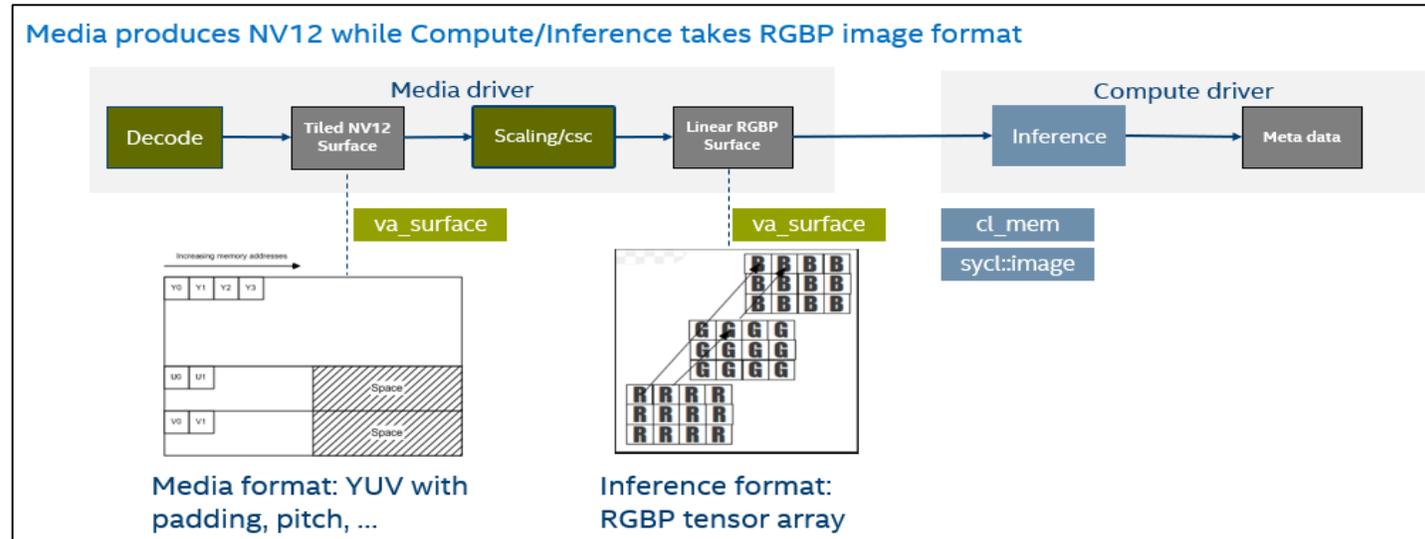
Where to run?



Sharing resources?



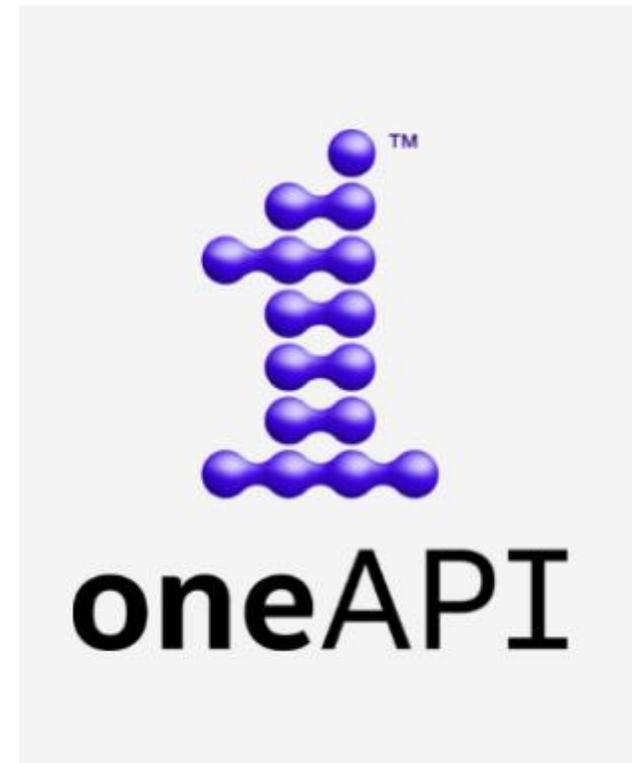
Sharing data?



oneAPI industry initiative

Open Industry Specification

- Specifies Language, APIs, Low level Hardware Interface
- Cooperative relationship with Khronos SYCL standard
- Promotes community and industry support
- Supports code reuse across architectures and vendors



<https://www.oneapi.com/>

ONEAPI INITIATIVE - ECOSYSTEM SUPPORT

allegro.ai

CINECA



GIGASPACES



Taboola



Hewlett Packard
Enterprise



Tencent 腾讯

RENIAI

Argonne
NATIONAL LABORATORY



codeplay®



sas



Indian Institute of
Technology Delhi



Tech
Mahindra

UNIVERSITY OF
CAMBRIDGE



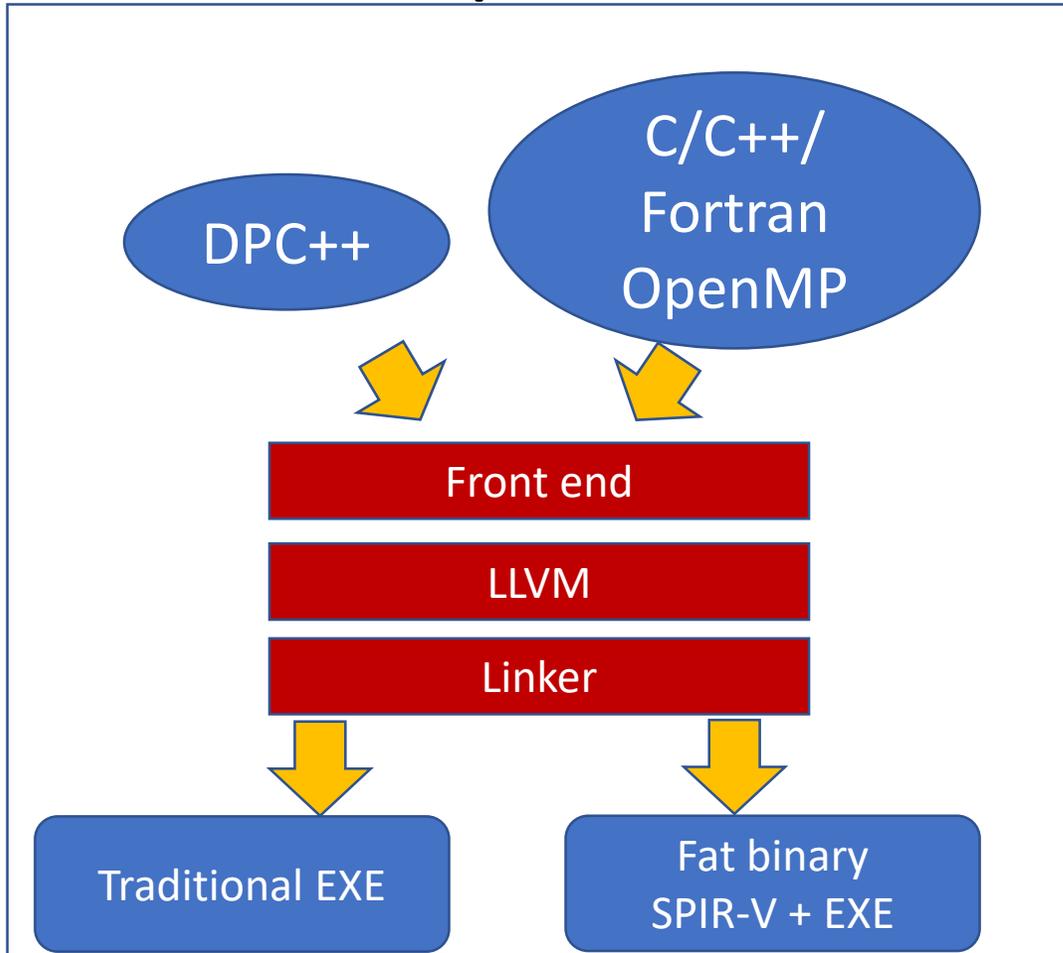
ZUSE
INSTITUTE
BERLIN

These organizations support the oneAPI initiative 'concept' for a single, unified programming model for cross-architecture development. It does not indicate any agreement to purchase or use of Intel's products.

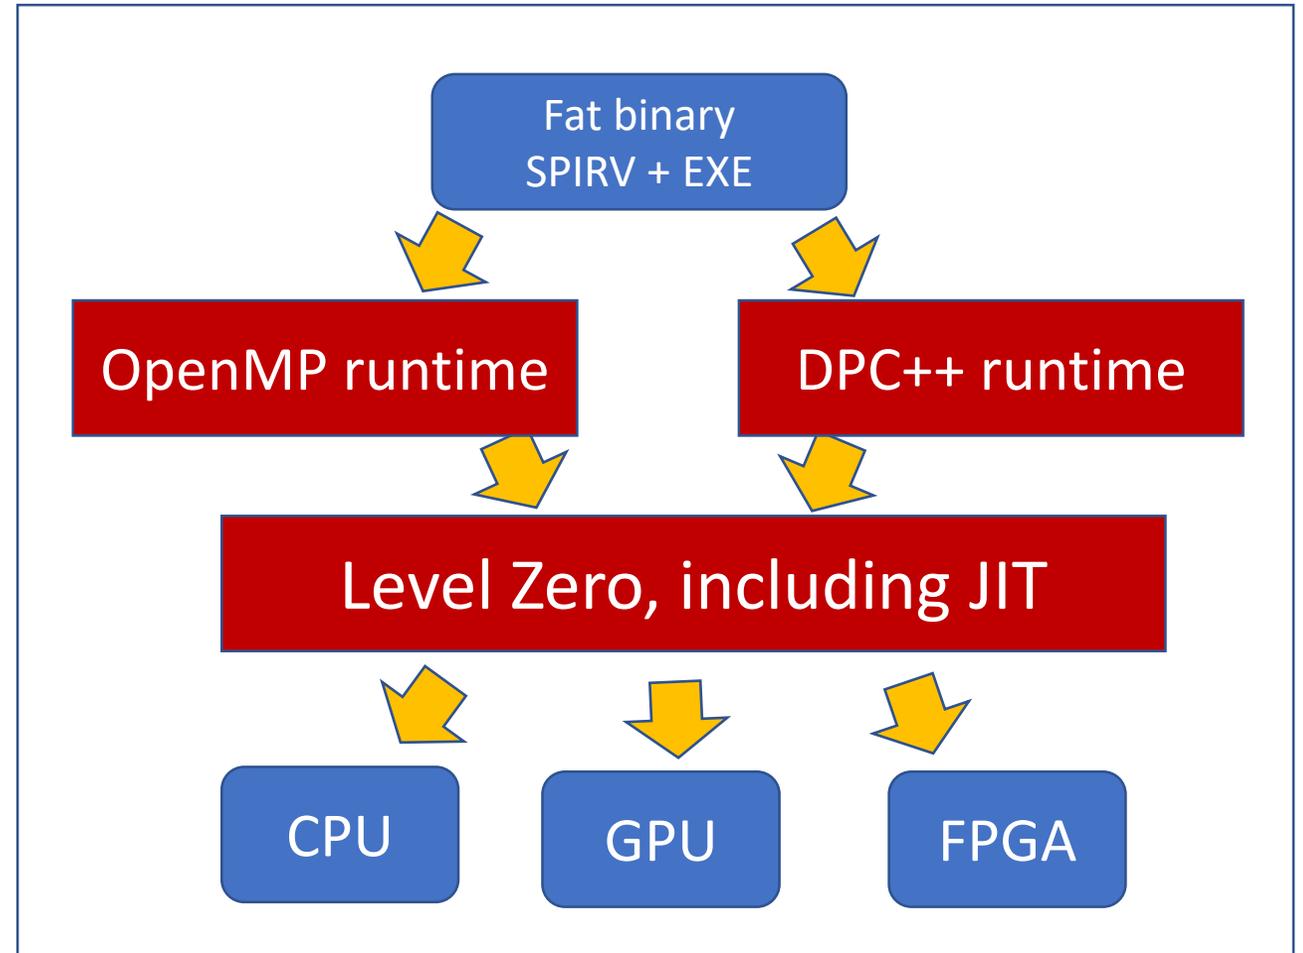
*Other names and brands may be claimed as the property of others.

Common stack for DPC++ and OpenMP

Compile-time



Run-time



Optimized MPI Operation with device RDMA

```
MPI_Send(
```

```
  dA,
```

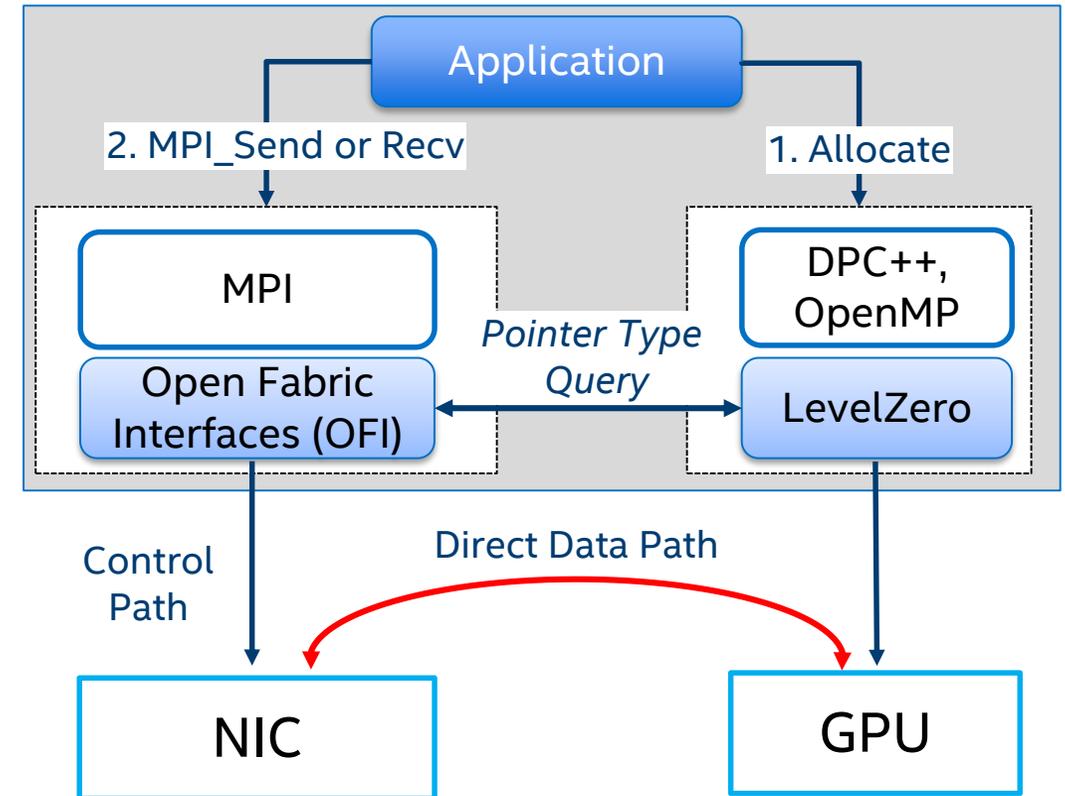
If address is in device memory,
direct path to NIC is used

```
  N,
```

```
  MPI_FLOAT,
```

```
  remote_rank, tag,  
  communicator);
```

```
}
```



ONEAPI AVAILABLE NOW ON INTEL DEVCLOUD

A development sandbox to develop, test and run your workloads across a range of Intel CPUs, GPUs, and FPGAs using Intel's oneAPI beta software

software.intel.com/en-us/devcloud/oneapi

Learn about oneAPI Toolkits

Learn Data Parallel C++

Evaluate Workloads

Build Heterogenous Applications

Prototype your project

NO DOWNLOADS | NO HARDWARE ACQUISITION | NO INSTALLATION | NO SET-UP AND CONFIGURATION

GET UP AND RUNNING FAST!

Notices and Disclaimers

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest forecast, schedule, specifications and roadmaps.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No product or component can be absolutely secure. Check with your system manufacturer or retailer or learn more at [intel.com].

Intel, the Intel logo is trademark 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.

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