oneAPI State of the Union

Tony Mongkolsmai
Software Architect

oneAPI DevSummit
June 13, 2023

Welcome

oneAPI Ecosystem

Google Microsoft accenture SIPEARL RISC-V Julia computing ArrayFire Durham University

MathWorks FUJITSU Adobe PHILIPS SIEMENS Healthineers SAMSUNG MEDISON mxnet 瀚博半导体 POWERED BY VASTSTREAM

cerebras stellar science RIKEN HUAWEI EURECOM Sandia National Laboratories UNIVERSITY of WASHINGTON Argonne NATIONAL LABORATORY Lawrence Livermore National Laboratory

KIT Karlsruher Institut für Technologie THE UNIVERSITY OF TENNESSEE KNOXVILLE Brookhaven National Laboratory ANACONDA. MANCHESTER 1824 UNIVERSITY of VIRGINIA nag

CoEs

Argonne NATIONAL LABORATORY University of BRISTOL सी डैक CDAC Computer Graphics Charles University Heidelberg 中国科学院 CHINESE ACADEMY OF SCIENCES LOBACHEVSKY UNIVERSITY Northern Illinois University OAK RIDGE National Laboratory

OLD DOMINION UNIVERSITY Stockholm University KTH TECHNISCHE UNIVERSITÄT DARMSTADT Berkeley UNIVERSITY OF CALIFORNIA UCDAVIS UNIVERSITY OF CAMBRIDGE Durham University ILLINOIS UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

University of Stuttgart THE UNIVERSITY OF TENNESSEE KNOXVILLE TACC SCI

oneAPI Going Mainstream

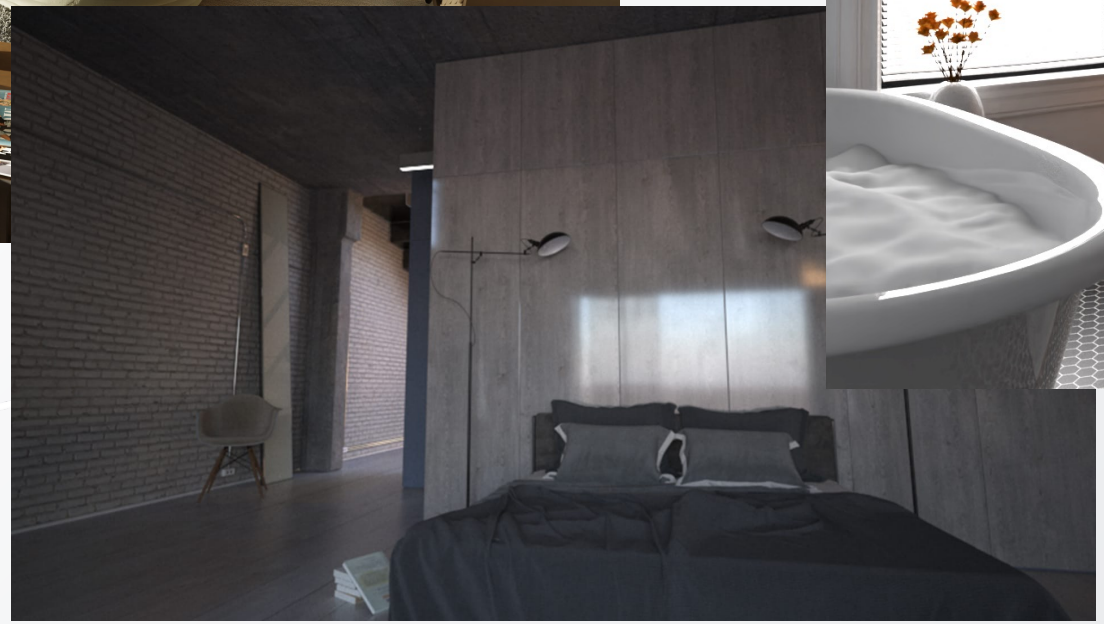

blender

chaos V-Ray

MOONRAY

oneAPI for AI

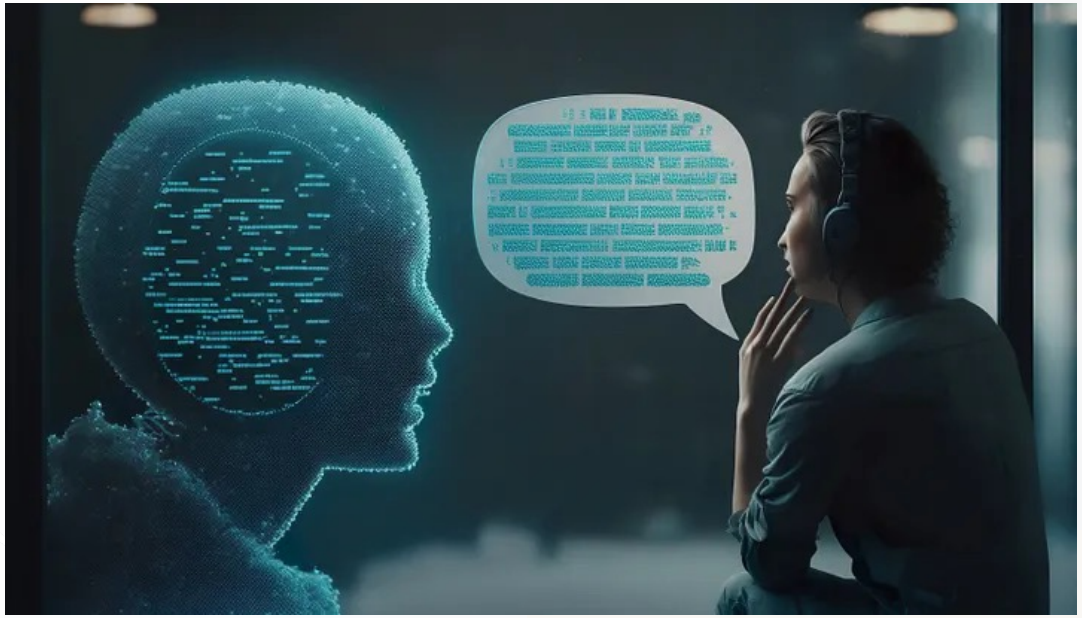

```
# load the Stable Diffusion model
pipe = StableDiffusionPipeline.from_pretrained("./stable-diffusion-v1-5",
                                               revision="fp16",
                                               torch_dtype=torch.float16)
# move the model to Intel Arc GPU
pipe = pipe.to("xpu")
```

```
# model is ready for submitting queries
pipe("an astronaut riding a horse on mars").images[0]
```

```
100% 50/50 [00:07<00:00, 6.99it/s]
```

Dolly

oneAPI for Everyone

oneAPI Innovator Program

István Z Reguly
oneAPI Innovator

Student Ambassadors for oneAPI

Harvey Johnson
oneAPI Student Ambassador

Joshua Shiells
oneAPI Student Ambassador

Melbin Martin
oneAPI Student Ambassador

Liftoff for Startups

DIGITAL CORTEX

sitemana

NayaOne

Success Stories

(don’t take our word for it!)

Ginkgo Project

- High-performance linear algebra library for many core systems focused on sparse linear systems
- Modern C++ library that supports GPU kernels in CUDA, HIP, and oneAPI SYCL
- Extensible and Open Source

Dr. Hartwig Anzt
University of Tennessee

DPEcho

- Leibniz-Rechenzentrum (LRZ) Project
- **D**ata **P**arallel **E**ulerian **C**onservative **H**igh **O**rder (DPEcho) for General-Relativity-Magneto-Hydrodynamic simulation (GR-MHD) to model turbulence, wave propagation, stellar winds and processes around black holes

Rafael Lago
Intel

Beesearch

- Collaboration between Beewant/Weaviate
- Unlocking the potential of large amounts of Unstructured Data
- Use AI to accurately identify unstructured information
- Use Vector Databases and Similarity search to scale up

Ahmed Joudad
CEO Beewant

Sebastian Witalec
Weaviate

Get Inspired!

Awesome

oneAPI

Industry and community showcase of projects

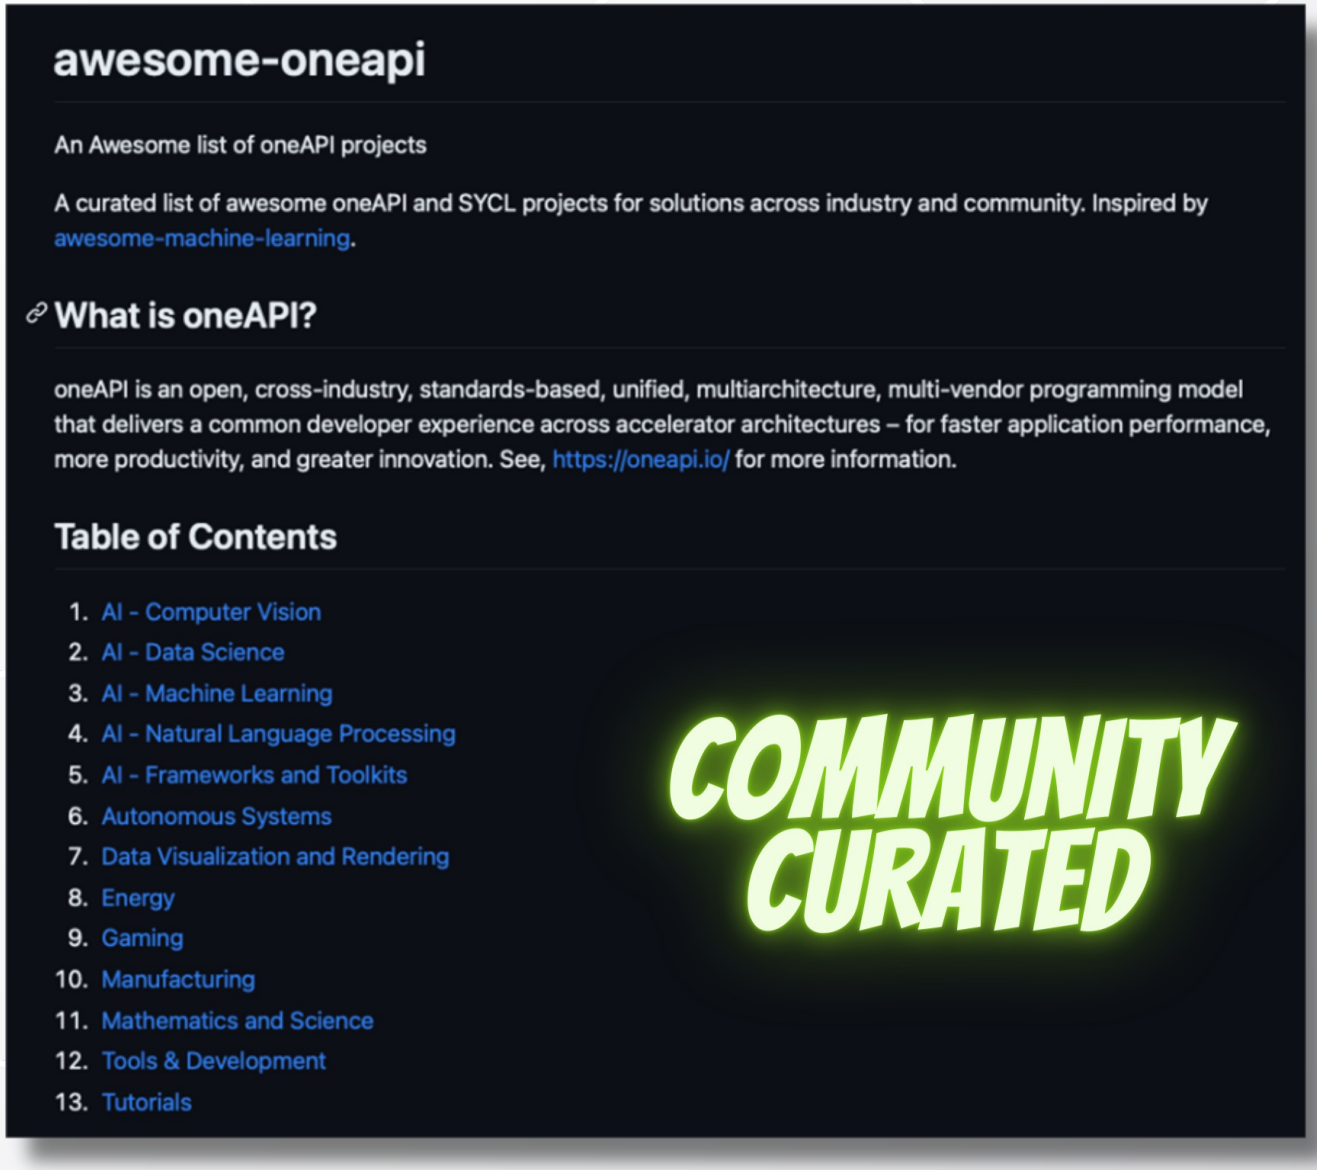

SYCLomatic

- Open source tool to help migrate CUDA code to oneAPI/SYCL
- 2832 contributors
- Just added support for 100+ CUDA APIs
- Working on SYCL 2020 support

https://github.com/oneapi-src/SYCLomatic

What’s New?

oneAPI Industry Specification

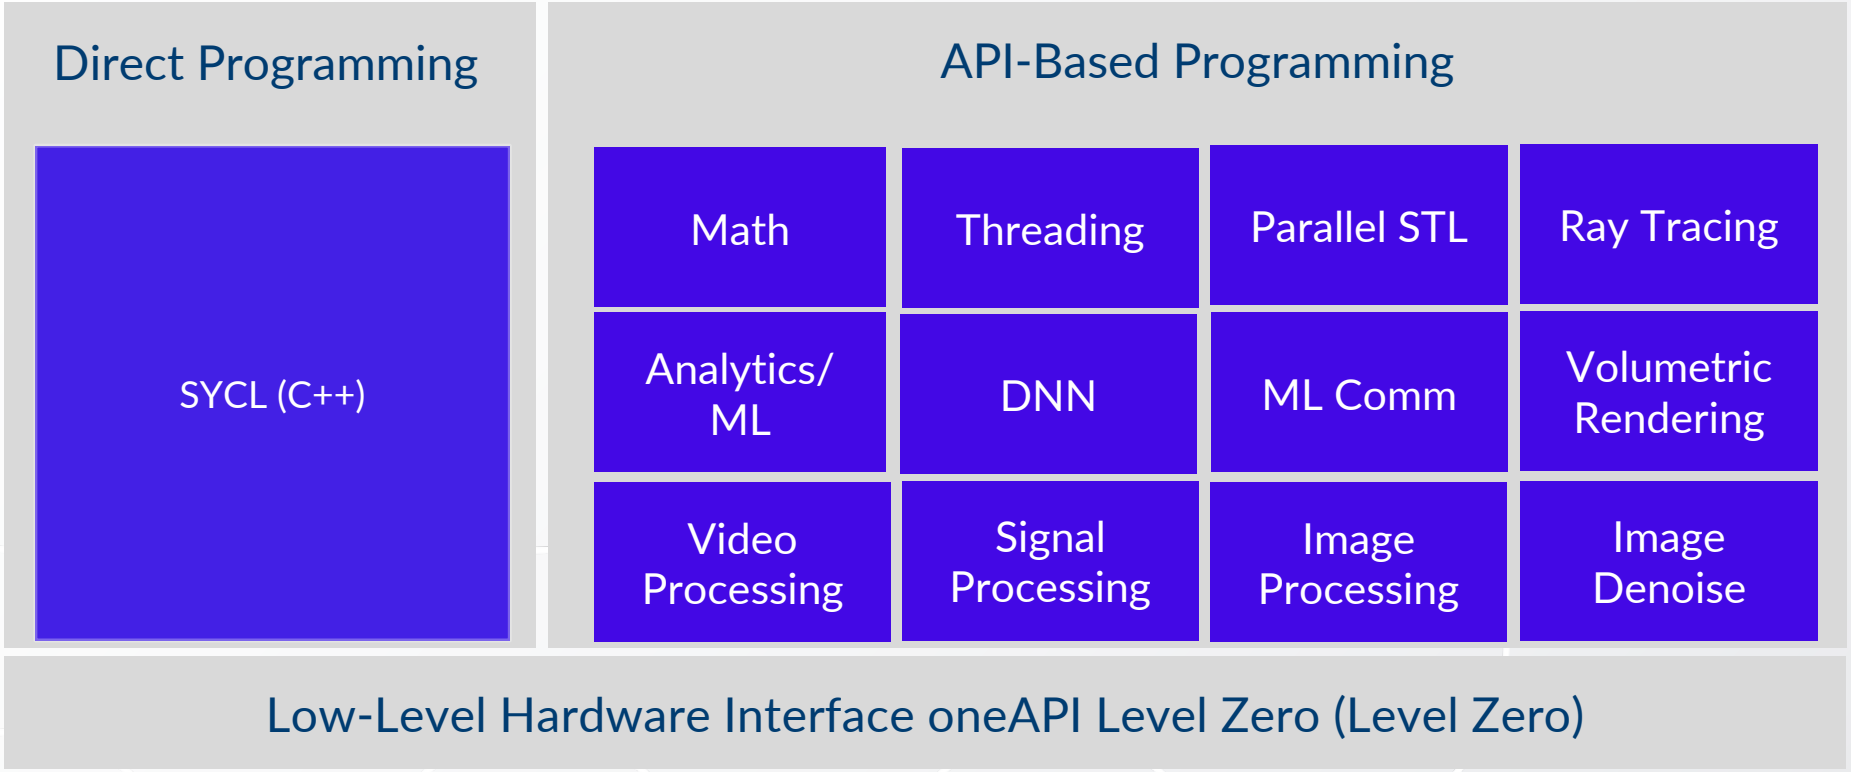

oneAPI Math Kernel Library (oneMKL) Specification

Features

- APIs for Dense/Sparse Linear Algebra, Fast Fourier Transforms, Vector Math, Vector Random Number Generation, Summary Statistics
- Open source oneMKL interfaces project - support for multiple hardware backends

What's New

- AMD Support
- NVIDIA DFT Support
- Intel Data Center Max CPU/GPU Support

Domain	Backend Support
Basic Linear Algebra Systems (BLAS)	• Intel oneMKL • NVIDIA cuBLAS • AMD rocBLAS • SYCL-BLAS
Linear Algebra Package (LAPACK)	• Intel oneMKL • NVIDIA cuSOLVER • AMD rocSOLVER
Random Number Generation	• Intel oneMKL • NVIDIA cuRAND • AMD rocRAND
Discrete Fast Fourier Transforms	• Intel oneMKL • NVIDIA cuFFT

oneAPI Deep Neural Network (oneDNN) Specification

Features

- Open source library implementation
- Supports key data type formats, including 16- and 32-bit floating points, bfloat16, and 8-bit integers

What's New

- Experimental Graph API and Sparsity Support
- Math Mode API to manage down converting to low precision data types
- Updated quantization scheme supporting int8
- Performance improvements on Intel/ARM CPUs and AMD/NVIDIA/Intel GPUs

Category	Functions
Compute intensive operations	• (De-)Convolution • Inner Product • RNN (Vanilla, LSTM, GRU) • GEMM
Memory bandwidth limited operations	• Pooling • Batch Normalization • Local Response Normalization • Layer Normalization • Elementwise • Binary elementwise • Softmax • Sum • Concat • Shuffle
Data manipulation	• Reorder

oneAPI Data Parallel C++ Library (oneDPL) Specification

Features

- Optimized C++ standard algorithms – parallel algorithms (C++17) and utilities
- Custom Utilities and Algorithms
- Execution Policies semantically aligned to C++ standard
- Built on underlying SYCL

Category	Functions
Buffer Wrappers	begin() end()
Iterators	counting_iterator discard_iterator permutation_iterator transform_iterator zip_iterator
Parallel Algorithms	exclusive_scan_by_segment inclusive_scan_by_segment reduce_by_segment binary_search lower_bound upper_bound

SYCL Compiler Advancements

- DPC++/C++ Compiler
  - Codeplay oneAPI for NVIDIA® GPUs and Codeplay oneAPI for AMD GPUs
  - Implementations for Proposed Extensions to SYCL
    - Joint Matrix
    - SYCL Graph
    - Bindless Texture
  - Updated to use SYCL 2020 Specification by default
- hipSYCL Compiler
  - Single pass compilation
  - Convenience of generating universal binaries -> --hipsycl-targets=generic
  - Intel GPU support no longer considered experimental
  - Unified device code representation for all backends

SYCL for Safety Critical Applications

- March 2023 - SYCL SC Working Group announced to develop C++-based heterogeneous parallel compute programming framework for safety-critical systems
- Will align with safety certification standards in avionics, automotive, industrial and medical fields

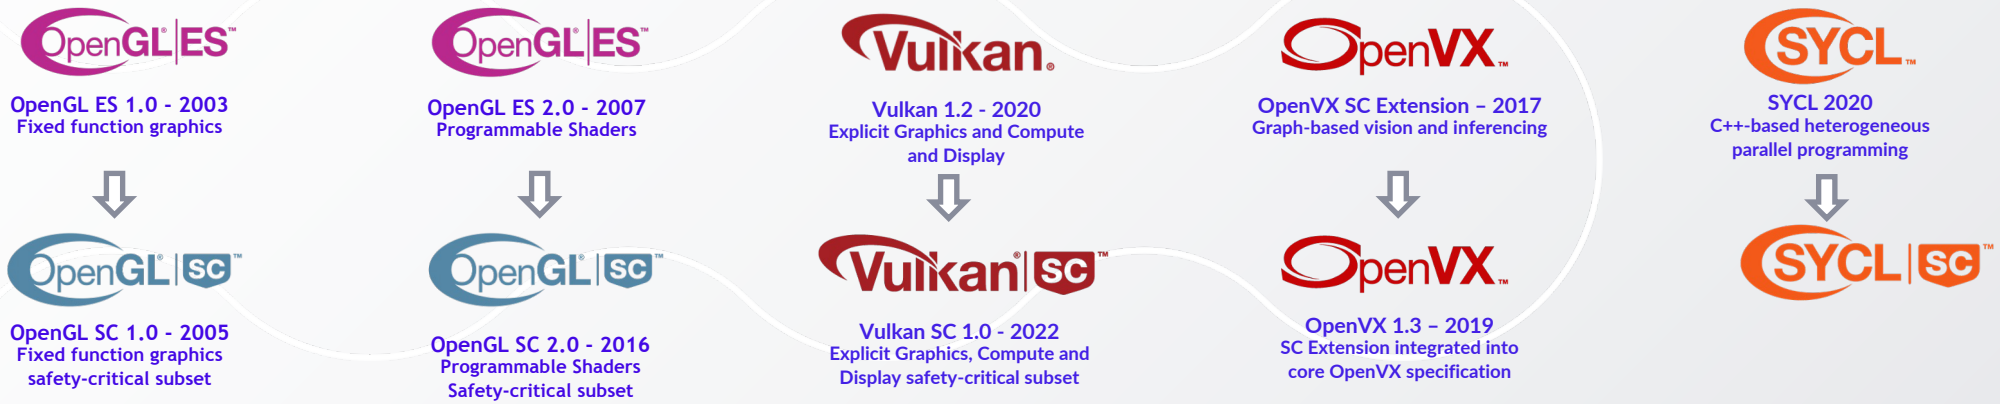

SYCL in Compiler Explorer

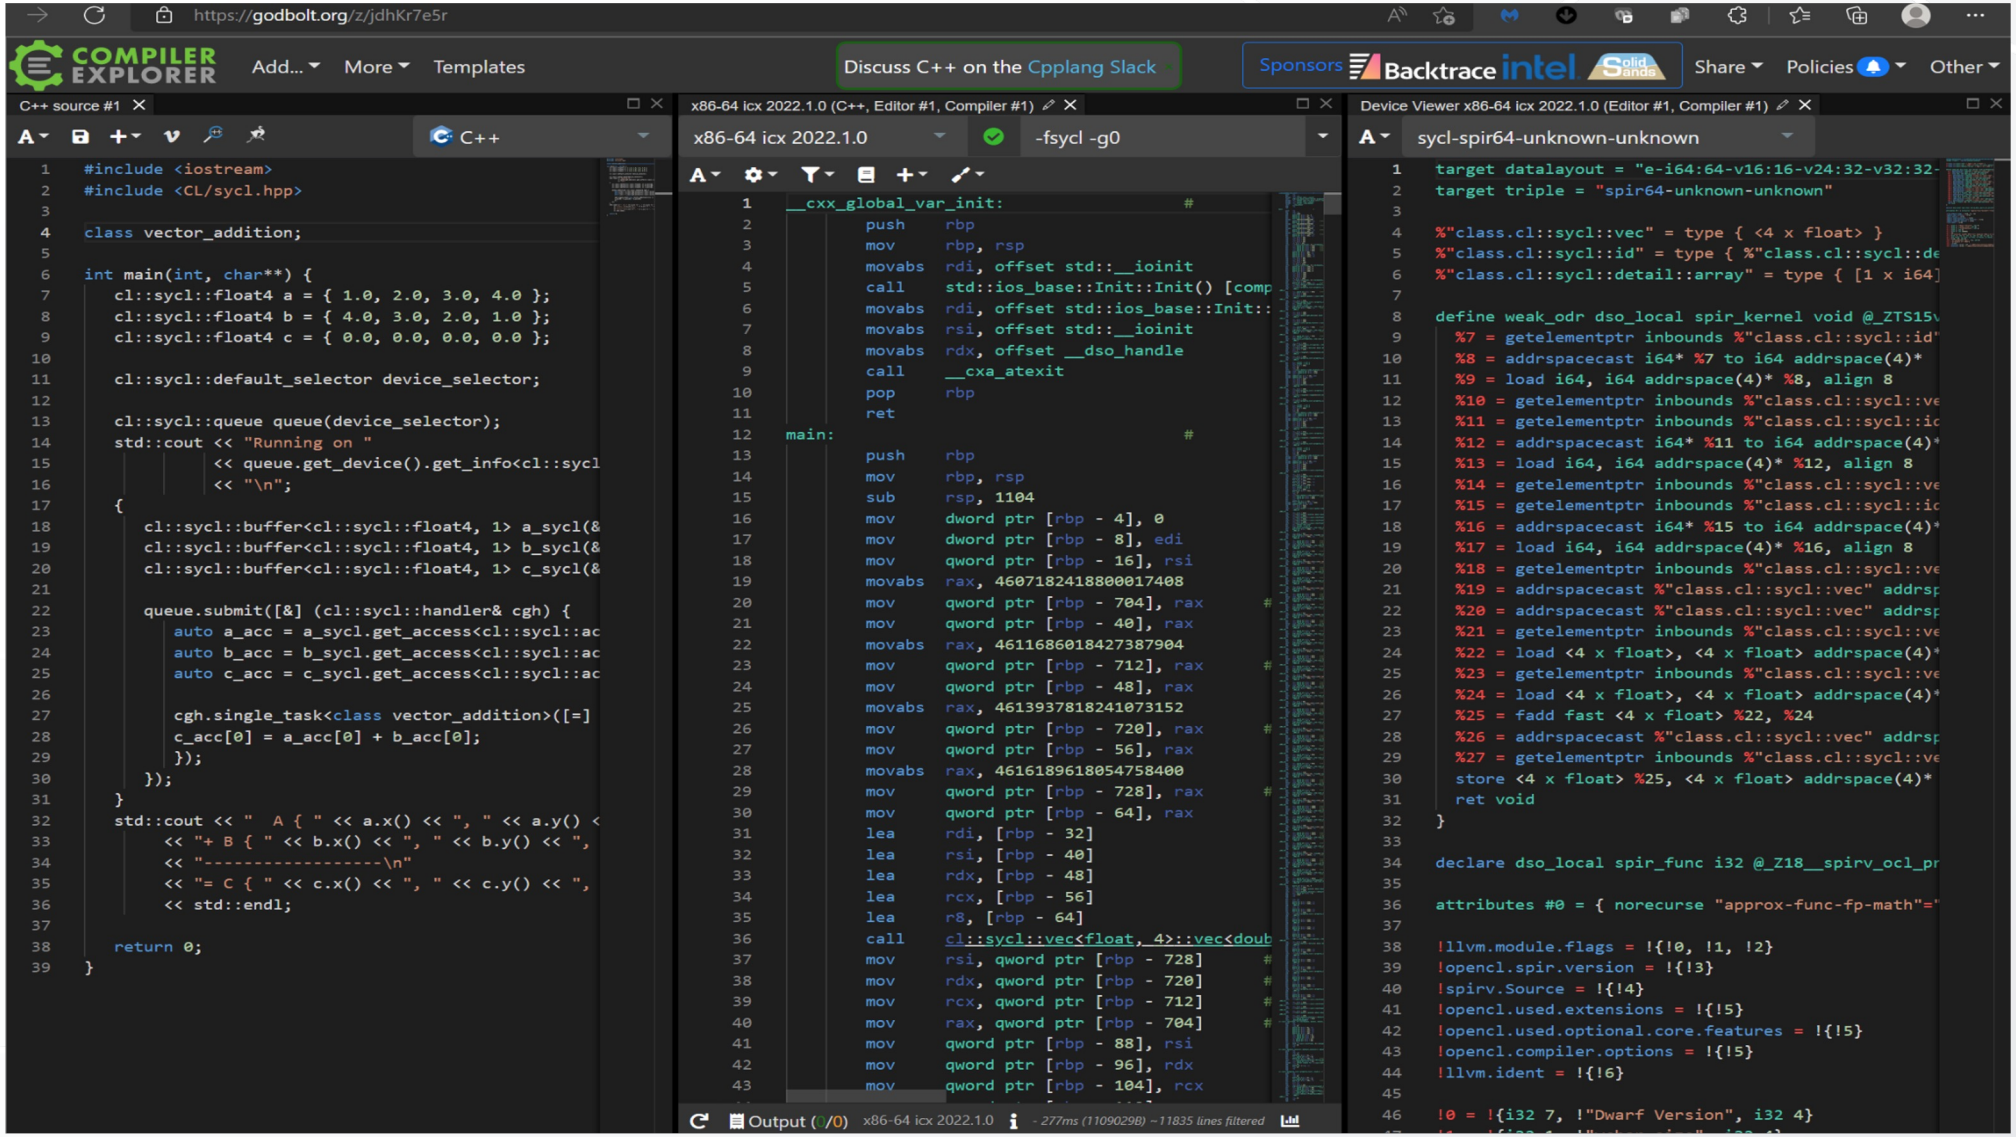

https://godbolt.org/z/jdhKr7e5r

oneAPI Construction Kit

- Simplifies process for hardware vendors to integrate with the oneAPI software stack
- For CPU + accelerator systems
- Ubuntu 20.04, Ubuntu 22.04, Windows
- RISC-V reference implementation

IWOCL 2020 Presentation

oneAPI Construction Kit Homepage

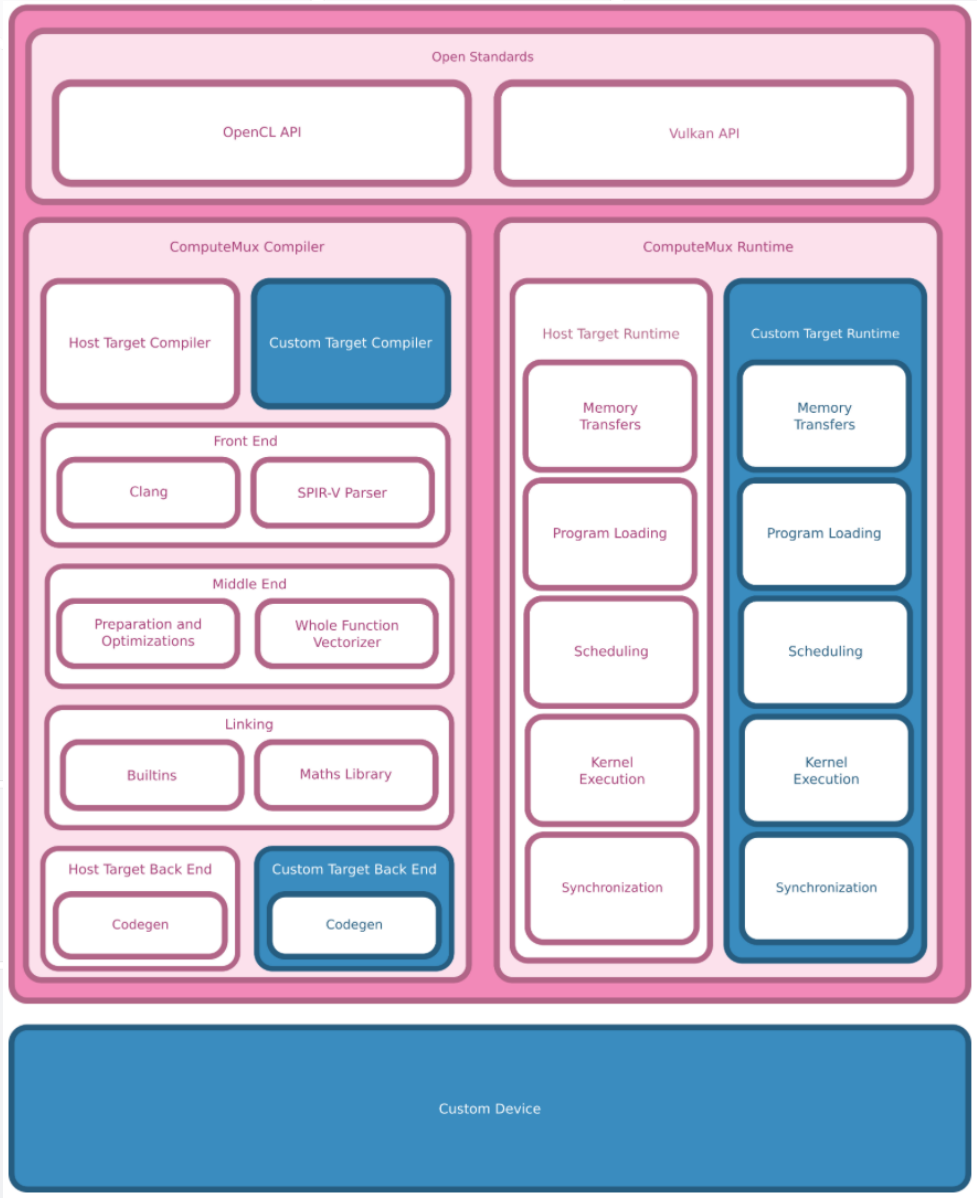

Community Matters

oneAPI Community Forum

For Software Developers

- Standards and industry defined libraries
- Future proof your software
- Enable an existing ecosystem of software and educational resources
- Develop with open standards for accelerator computing
- Enable software to run on multiple architectures

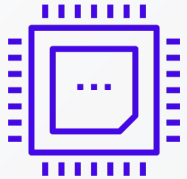

For Hardware Developers

- Fast track into market with an existing ecosystem
- Share development cost with open-source implementations
- Leverage an existing tested and optimized toolchain
- Enable an existing ecosystem of software and educational resources

Free and based on open standards

Special Interest Groups

Special Interest Groups influence the specifications and implementations

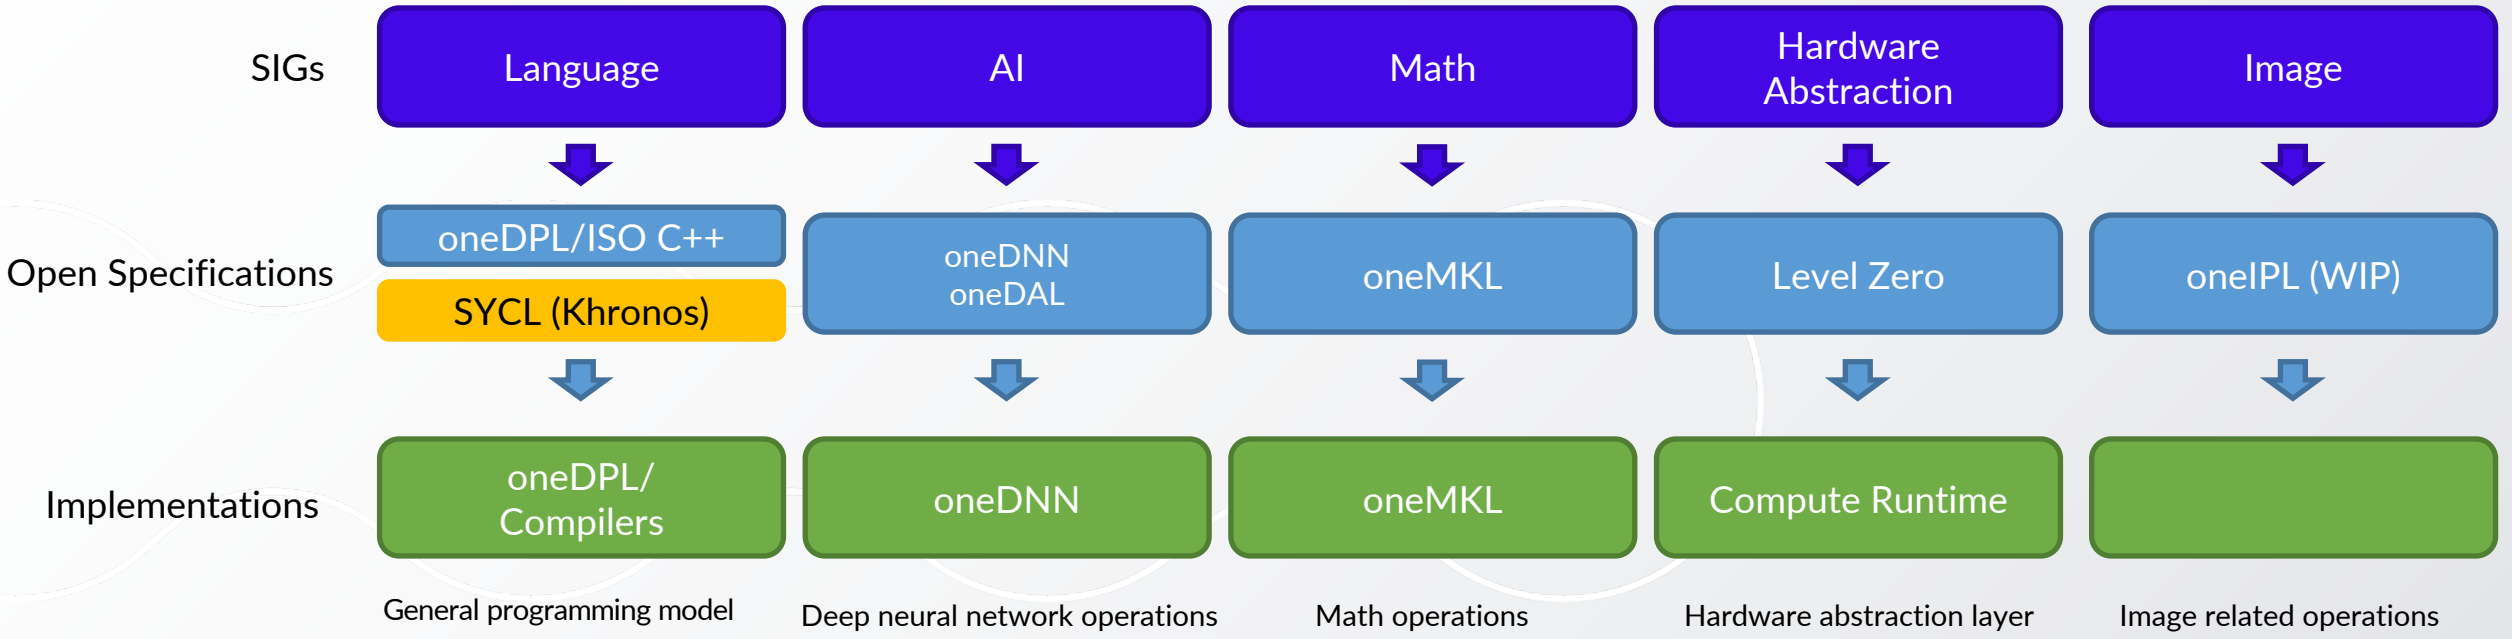

Get Involved!

Join	Learn	Try	Share
			
oneAPI Forum	Data Parallel C++ Book	Intel® Developer Cloud (Beta)	Awesome oneAPI GitHub

oneAPI Community Programs

Student Ambassadors

Innovator Program

Intel® Liftoff for Start-ups

Educator Program

Join our community!

Thank You