



SYCL State of the Union Keynote SYCLcon 2022 Specification Release



Michael Wong

SYCL V

Codeplay Distinguished

ISOCPP Foundation

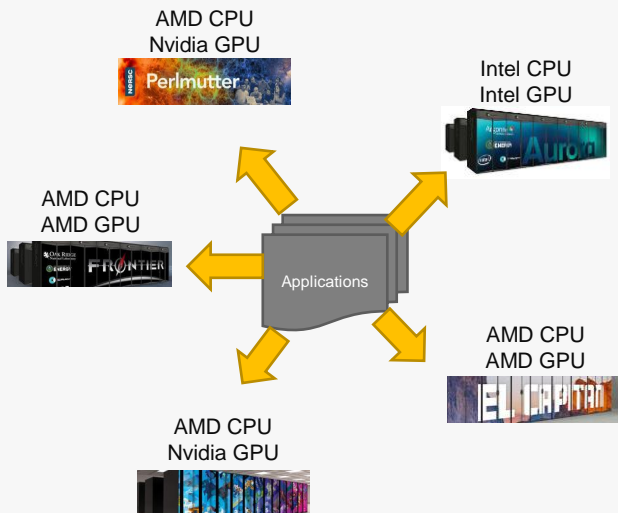
ISO C++ Direction

michael@codeplay.com | wongmichael.co



Programming Models Must Persist

US National Laboratory Supercomputers 2021-2023



- HPC and now exascale computing requires programming models that endure for future workloads, > 20 years
- But Hardware changes frequently, constant improvement
- Programming models, have to be stable but also support latest HW,

Requires an open interface, across architectures with multiple implementations



SYCL 2020 Launched February 2021

Expressiveness and simplicity for heterogeneous programming in modern C++

Closer alignment and integration with ISO C++ to simplify porting of standard C++ applications

Improved programmability, smaller code size, faster performance

Based on C++17, backwards compatible with SYCL 1.2.1

Backend acceleration API independent

New Features


Unified Shared Memory | Parallel Reductions | Subgroup Operations | Class template Argument Deduction



Significant SYCL adoption in Embedded, Desktop and HPC Markets

SYCL Projects cumulative growth

[Projects - SYCL.tech](https://www.sycl.tech)



Benchmarks



Frameworks



Libraries




Scientific




SYCL user and developer Phenomenal Growth

Star history



stackoverflow questions annually on SYCL+oneapi



Some open-source SYCL implementations/prototypes on GitHub



Github SYCL source files is over 30000

"#include <CL/sycl.hpp>" / Pull requests Issues Marketplace Explore

Repositories	0
Code	11K
Commits	13
Issues	117
Discussions	3
Packages	0
Marketplace	0
Topics	0

21,446 code results

jfuentes/heterogeneous-computing-website
english/content/unit2/oneapi.md

```
18 {{< embed-pdf url="pdf/11-dpc++.pdf" >}}
19
20 **E
21
22 *Exc
23
24 ...
25 #inc
26 cons
27 usir
```

17.5k mentions of #include <CL/sycl.hpp> in Jan, 2022

21 K in April, 2022

"#include <SYCL/sycl.hpp>" / Pull requests Issues Marketplace Explore

Repositories	0
Code	1K
Commits	277
Issues	152
Discussions	3
Packages	0
Marketplace	0

8,804 code results

xfong/sycl
library/include/sycl_engine.hpp

```
1 #ifndef RUNTIME_ENGINE_SYCL_SYCL_HPP
2 #define RUNTIME_ENGINE_SYCL_SYCL_HPP
3 #include <CL/sycl.hpp>
4 namespace sycl {
5
6 #endif // RUNTIME_ENGINE_SYCL_SYCL_HPP
7
8 #include "runtime_engine_sycl.hpp"
```

8.2k mentions of #include <SYCL/sycl.hpp> in Jan, 2022

8.8 K in Apr, 2022



Market needs SYCL: easy to build SYCL on any device

2016



- 21 projects
- 2 implementations in work
- 3-4 platforms
- SYCL was a TSG of OpenCL
- 10 members attending TSG
- <10 GitHub stars
- <10 StackOverflow questions
- No Safety Critical
- No book, a few articles
- < 100 GitHub include code
- No Supercomputing presence
- No automotive
- No Clang/LLVM
- No common one-stop website
- No common teaching material
- No HPC

2022



- 147 projects
- >12 implementations in work
- 10+ platforms
- SYCL is independent WG of Khronos
- 20+ members attending WG
- >600 GitHub stars
- >600 StackOverflow questions
- Safety Critical Exploratory Group
- 1 book, many articles
- ~30000 GitHub include code
- SC BoF 5 years in a role
- Renesas R-Car
- Clang/LLVM DPC++ active
- [SYCL.tech](https://www.sycl.tech)
- SYCL Academy
- ~7 HPC systems



SYCL is mainstream

- Open Standards and Open Source implementations
- Open cross-company collaboration
- Co-design for all forms of extreme heterogeneity
- SYCL Survey coming

Q* What language functionality would you like to see more broadly supported?

- ☐ C++20 as the core language
- ☐ C++23 as the core language
- ☐ Unified Shared Memory (USM)
- ☐ Unnamed kernel lambda functions
- ☐ Hierarchical Parallelism improvements
- ☐ None
- ☐ Other (Comment)




Market needs SYCL to Evolve (call to action)

- More workloads, NAMD, GROMACS, ROOTS
- Compile more ISO C++ as C++ advances
- **CTS, SDK, Compiler explorer**
- SYCL Graphs
- SPEC Accel/HPC Benchmarks
- Wikipedia, FAQ
- Safety Critical
- Educational videos
- Public CTS nodes
- More ecosystem, more libraries
- SYCL MLIR
- Public CTS does
- SYCL interpreter, Jupyter notebook
- More Vendor adoption
- Better tooling, profilers, debuggers, analyzers
- Data movement is still King, Will get worse with sparsity
- CUDA to SYCL conversion
- Parallelism Survey 2022 at NASA

SYCL Projects
2022

Programming Languages
(244 entries)



Others:

- Ruby (3 entries)
- Julia (2)
- CUDA/OpenMP (1)
- IDL (1)
- Tcl/tk (1)
- Shell scripting (1)
- Don't know

- Fortran/C
- Python is
- SYCL/DPC++ (by FUN3D)



Market needs SYCL to succeed in democratizing


Q* Which of these provisional and vendor extensions would you like to see become Khronos extensions?

- ☐ Scoped Parallelism: [View via hipSYCL](#)
- ☐ Multi-device queue: [View via hipSYCL](#)
- ☐ Command group properties: [View via hipSYCL](#)
- ☐ Buffer-USM interop: [View via hipSYCL](#)
- ☐ Pipes: [View via Intel](#)
- ☐ Accessor restrict property: [View via Intel](#)
- ☐ Accessor properties: [View via Intel](#)
- ☐ Enqueue barrier: [View via Intel](#)
- ☐ Bfloat16: [View via Intel](#)
- ☐ Properties: [View via Intel](#)
- ☐ None



From the recent Programming Systems Research Forum at DOE Feb 2022

Please rate your interest in participating in each of these potential breakout group topics. Pick at most six (6) topics. We are saving two breakout groups free to suggest a topic in the future.
46 responses



Compilers with AI/ML

Performance debugging

Language Parallelism

Distributed Memory

Verification Correctness

Sparsity Irregularity

Performance portability

Domains ne

Task Based



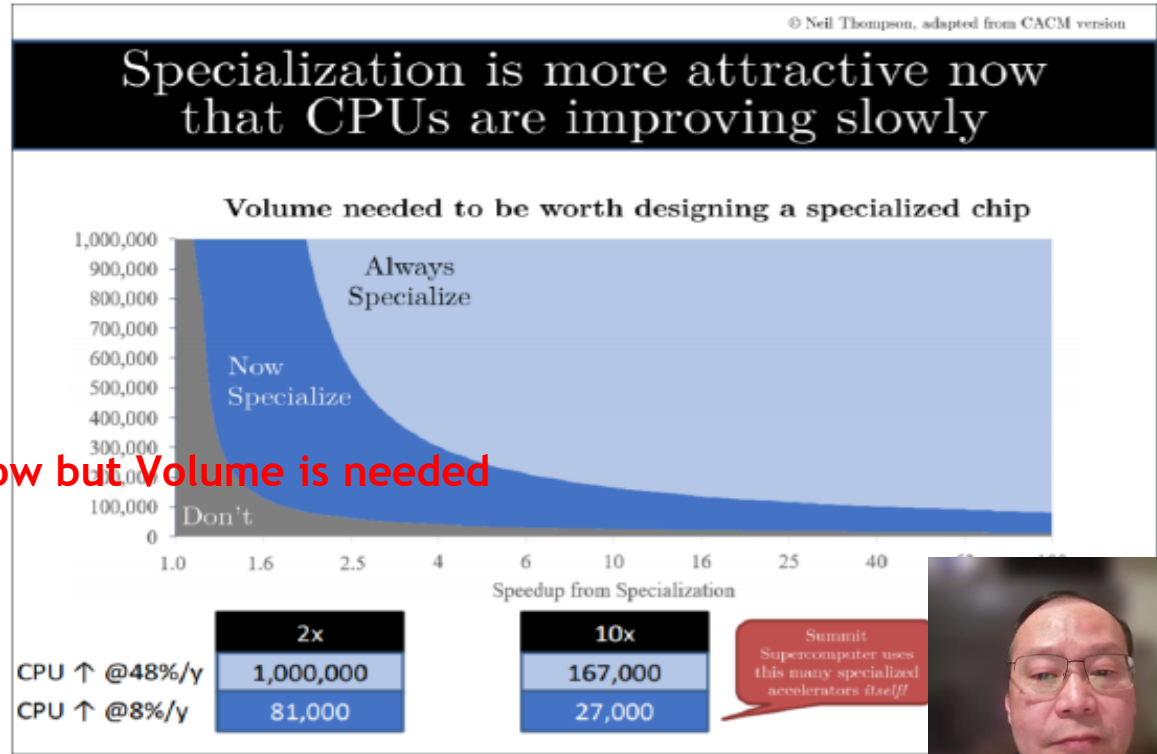
From ASCR Workshop Mar 2021

ASCR Workshop on Reimagining Codesign


<https://www.orau.gov/ASCR-CoDesign/>

From Neil Thompson's keynote:

Specialization is more attractive now but Volume is needed



From Flops to TOPS in ML



<https://www.kdnuggets.com/2020/05/tops-just-hype-dark-ai-silicon-disguise.html>



工
▼

OpenAI booth at NeurIPS 2019 in Vancouver, Canada
Image Credit: Khan Johnson / VentureBeat

Last week, OpenAI published a paper detailing GPT-3, a machine learning model that achieves strong results on a number of natural language benchmarks. A 175 billion parameters model where a parameter affects data's prominence in an overall prediction, it's the largest of its kind. And with a memory size exceeding 350GB, it's one of the priciest, costing an estimated 12 million to train.



- 30% over PyTorch
- 20% over Tensorflow + XLA
- 8% over DeepSpeed

02.00554v1 [cs.LG] 31 Jan 2021

Operator class	% flop	% Runtime
Tensor contraction	99.80	61.0
Statistical normalization	0.17	25.5
Element-wise	0.03	13.5

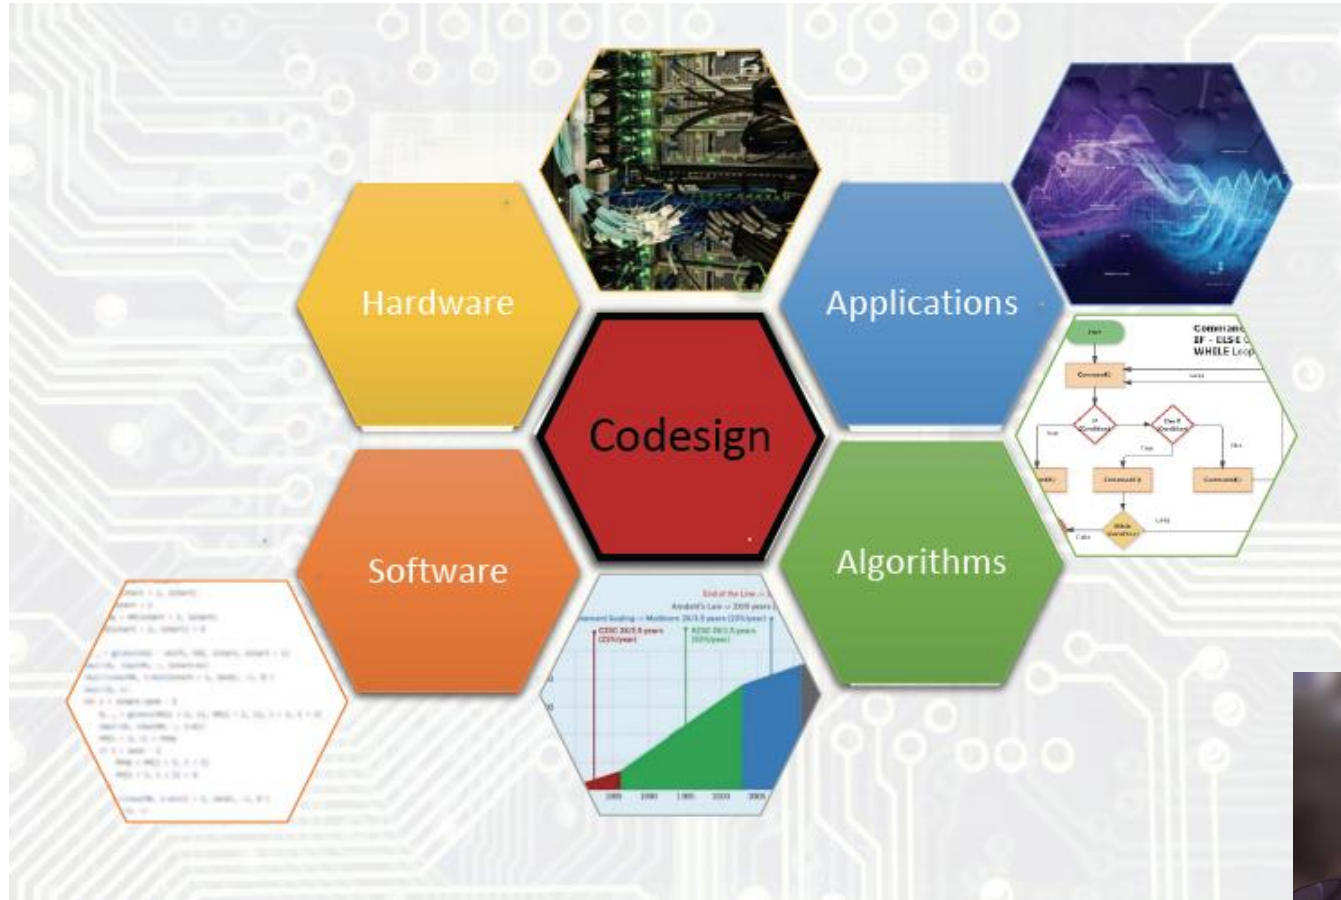
Microsoft, 2020b), number of training samples (Raffel et al., 2019; Liu et al., 2019), and total iterations (Liu et al., 2019; Kaplan et al., 2020). These all significantly increase compute requirements. Indeed, transformers are becoming the dominant task for machine learning compute where training a model can cost tens of thousands to millions of dollars, even cause environmental concerns (Strubell et al., 2019). These trends will only accelerate with pushes toward models with tens of billions to trillions of parameters (X

The growing energy and performance costs of deep learning have driven the need for neural networks by selectively pruning components. Similarly to the way that neurons in the brain generalize just as well, if not better than, the original dense neural networks, the memory footprint of regular networks is fit for mobile devices, as well as the speed of inference. In this paper, we survey progress on sparsity in deep learning, with a specific focus on the design of training algorithms for pruning neural networks, different training strategies to achieve model sparsity, and practice. Our work distills ideas from more than 300 research papers and who wish to utilize sparsity today, as well as to researchers whose goal is to include the necessary background on mathematical methods in sparsity at early structure adaptation, the intricate relations between sparsity and model performance, and the challenges of training sparse networks that could serve as a baseline for comparison of different sparse network sparsity can improve future workloads and outline major open problems.


The supreme goal of all theory is to make the irreducible basic elements possible without having to surrender the adequate representation of

This work is licensed under a Creative Commons Attribution 4.0 International License

We need Codesign with extreme Heterogeneity



Future SYCL: Emerging transformative technologies



- Chiplets and Superchips
- Licensable IP for Server-class processors (ARM)
- Open Source Hardware and Open Silicon Compilers (RISC-V)
- Photonic Resource Disaggregation (Ayar Labs TeraPhy, ARPA-E ENLITENED, and DARPA PIPES)
- Standardized Accelerator Interfaces (CCIX, Coherent PCIe, CXL, UCIe, HSA, Intel Level-0 & DSA)
- Advanced Hardware description Languages and Hardware generator
- New Accelerators (Ex: AMD/Xilinx Versal, SambaNova, Graphcore, Cerebras...)
- Advanced Packaging Technologies for Heterogeneous Integration (HIR) <https://eps.ieee.org/technology/heterogeneous-integration-roadmap/2019-edition.html>
- Open Source and Extensible Compiler Framework (MLIR)
- AI integrated applications
- Programming Abstractions, Frameworks, and Languages



Parallel Industry Initiatives



C++11



C++14



C++17



C++20



C++23



SYCL 1.2
C++11 Single source
programming




SYCL 1.2.1
C++11 Single source
programming




SYCL 2020
C++17 Single source
programming
Many backend options




SYCL 202X
C++20 Single source
programming
Many backend options



OpenCL 1.2
OpenCL C Kernel
Language



OpenCL 2.1
SPIR-V in Core



OpenCL 2.2



OpenCL 3.0



2011

2015

2017

2020

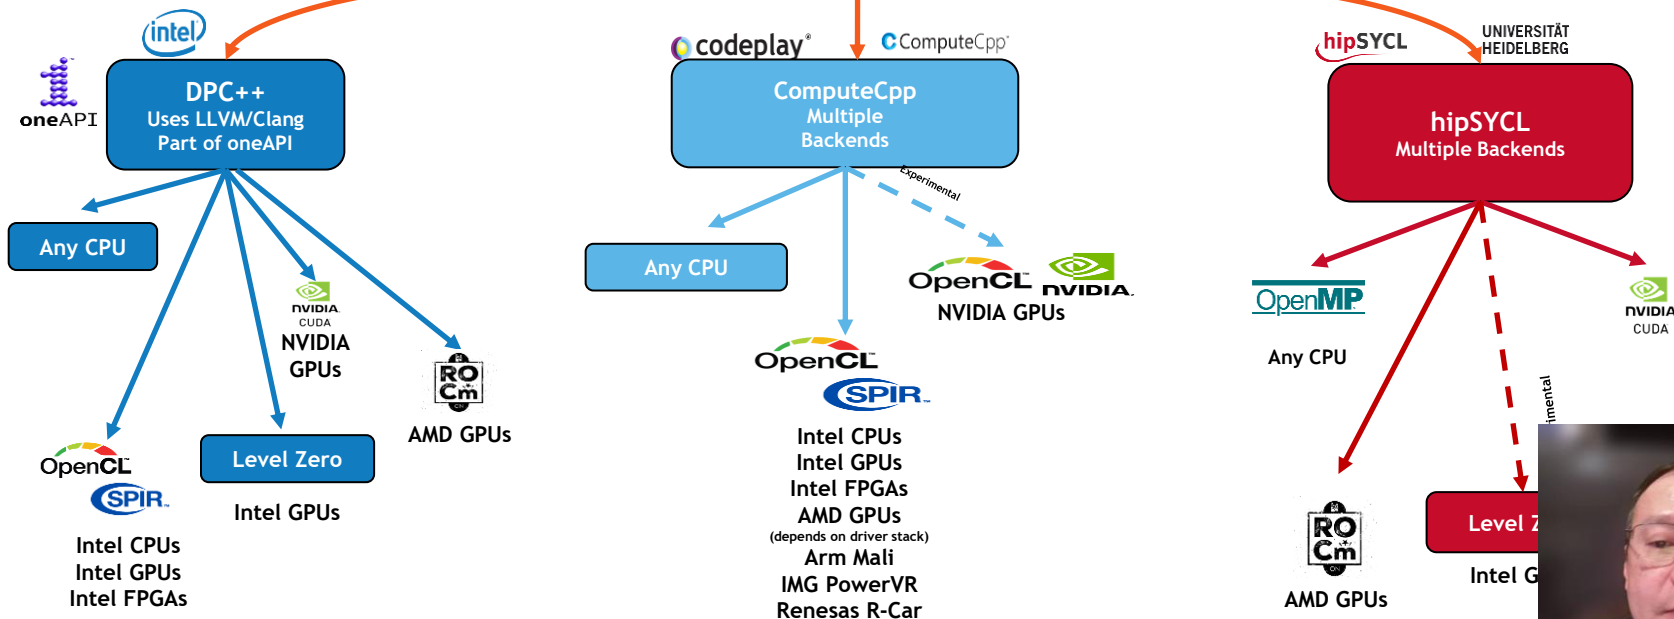


SYCL Implementations in Development (2022/05/01)

SYCL, OpenCL and SPIR-V, as open industry standards, enable flexible integration and deployment of multiple acceleration technologies

SYCL
Source Code

SYCL enables Khronos to influence ISO C++ to (eventually) support heterogeneous compute

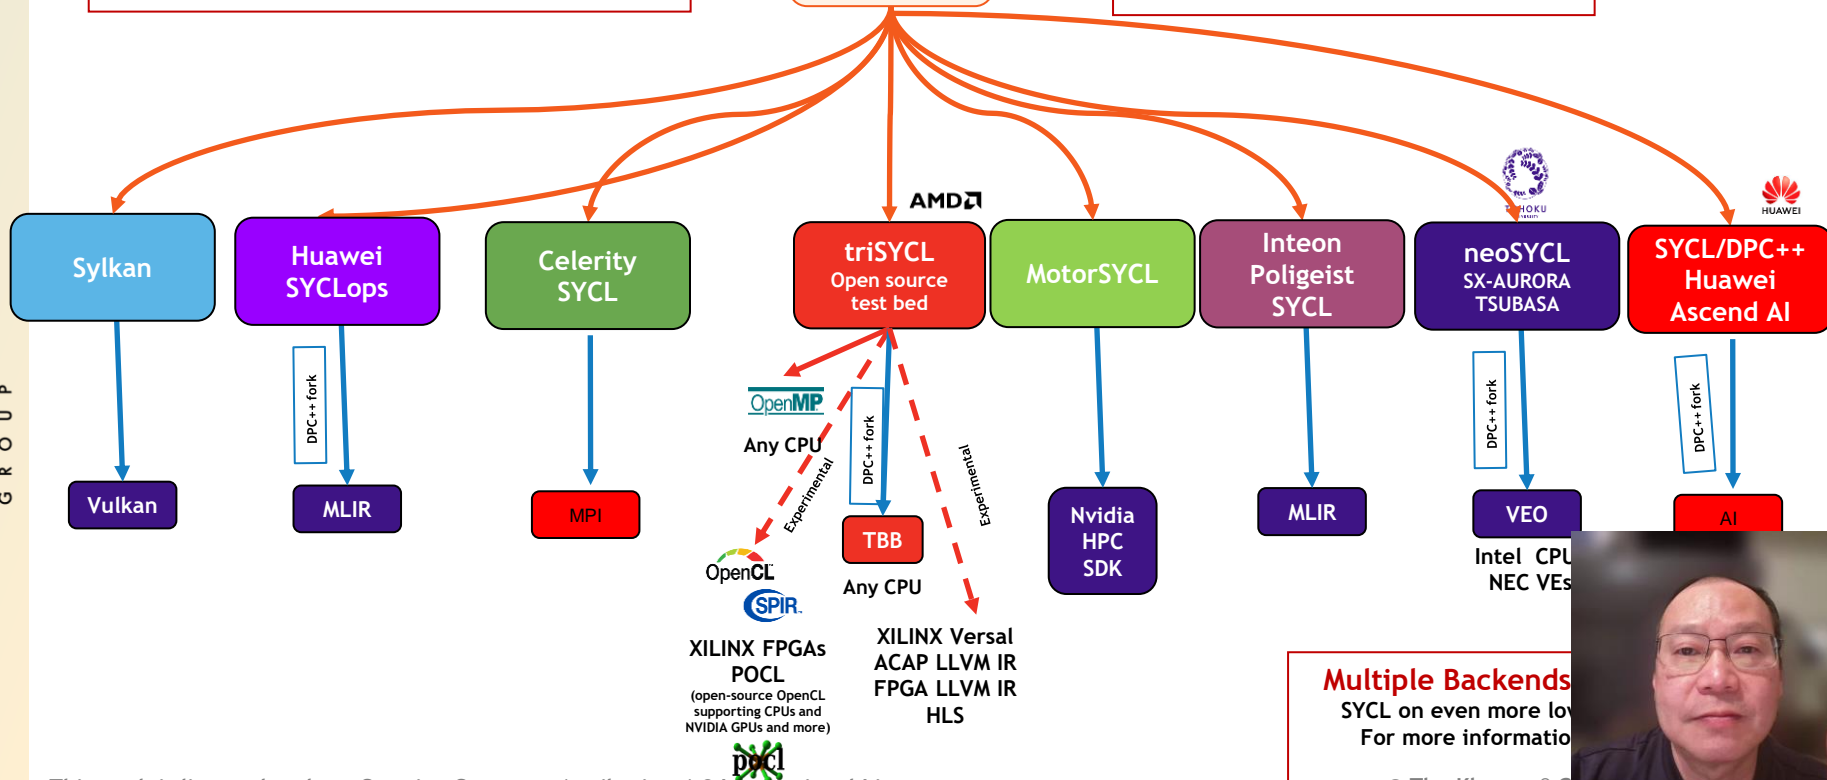


SYCL Experimental Development (2022/05/01)


SYCL, OpenCL and SPIR-V, as open industry standards, enable flexible integration and deployment of multiple acceleration technologies

SYCL
Source Code

SYCL enables Khronos to influence ISO C++ to (eventually) support heterogeneous compute



Migration advice by Platform (2022/05/05): For a device kind, which SYCL compiler?



ALTERA



AMD

habana

ambria

Google

groq

GRAPHCORE



Tenstorrent

SambaNova



XILINX

- NVIDIA GPUs
 - DPC++ supported today by Codeplay, open-source, optimized
- AMD GPUs
 - hipSYCL open-source today is better and is supporting European Processor initiative in LUMI & Karolina
 - DPC++ experimental support today by Codeplay, open-source
- Intel GPUs
 - DPC++ supported by Intel
- New hardware:
 - Codeplay supports a range of new hardware with ComputeAorta + ComputeCpp, including RISC-V custom hardware and GPU IP from Imagination Technologies and ARM, as well as Renesas R-Car for automotive
 - Also can do it on their own (DIY)
- FPGA
 - DPC++ compiler for Intel FPGAs
 - AMD/Xilinx has triSYCL prototype for FPGA & CGRA
- AI/ML/NN
 - All the major implementations/companies can customize support or you can build it DIY
- RISC-V
 - Codeplay has developed ACORAN stack for RISC-V as an accelerator
- CPU
 - All major SYCL implementations



SYCL Ecosystem, Research and Benchmarks 2022/05/01

Implementations

neoSYCL
SX-AURORA TSUBASA



oneAPI
oneAPI



HUAWEI

ComputeCpp™

hipSYCL

Research

Celerity
High-level C++ for Accelerator Clusters




PYTORCH



ATLAS
EXPERIMENT

Flashlight

Tutorials/Benchmarks/ Books



Future Workshops



Distributed and Heterogeneous Programming in C++ (DHPCC++22)



oneAPI Developer Summit @ SYCLcon 2022

Recent Workshops



Portable Heterogeneous Programming with SYCL (PHPS22)



SYCL Workshop with ENCCS for the Karolina Supercomputer

KHRONOS GROUP

STREAM
HPC

Argonne
NATIONAL LABORATORY

QUALCOMM

arm

AMD

SYCL



intel

codeplay




University of
BRISTOL

Working Group Members

Creative Commons Attribution 4.0 International License

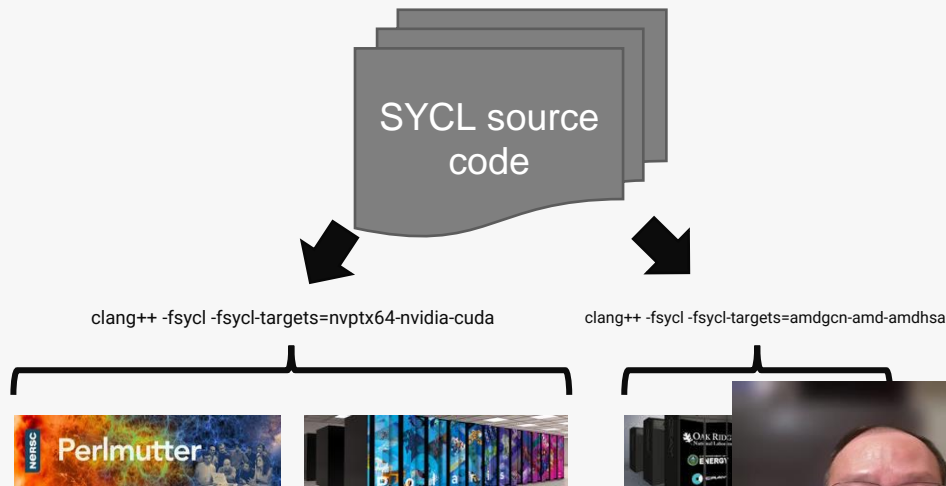
© The Khronos Group



Nvidia and AMD Support in DPC++

- Extending DPC++ to target Nvidia and AMD GPUs
- Supporting Perlmutter, Polaris and Frontier supercomputers
- Open source and available to everyone
- Codeplay commercially supports these implementations

Different targets using a simple compiler flag


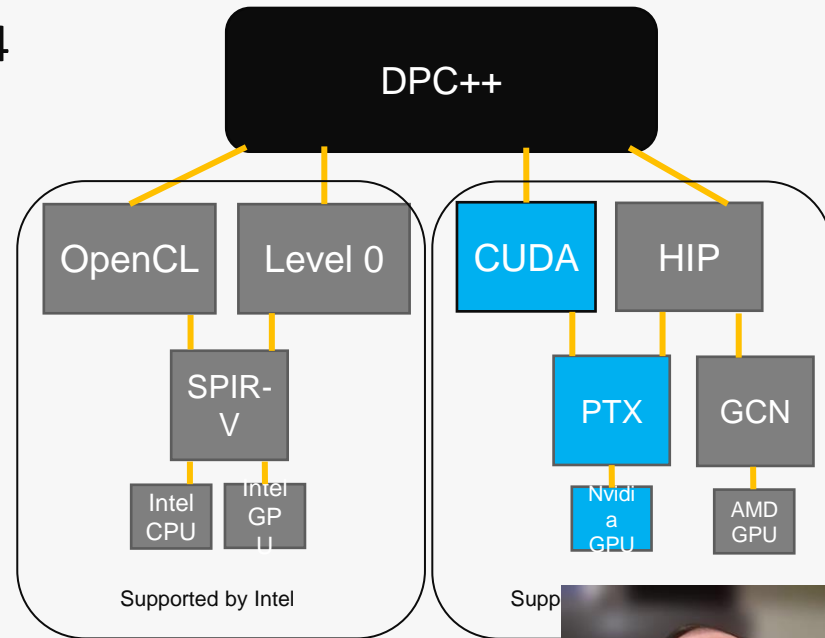


<https://www.codeplay.com/oneapiforcuda>
Resources for AMD coming soon




CUDA as a backend

- The DPC++ runtime currently implements 4 SYCL backends:
 - OpenCL
 - Level Zero
 - CUDA
 - HIP
- Moving up a level of abstraction, you untether your codebase from specific hardware
- Developers have flexibility to prioritize performance



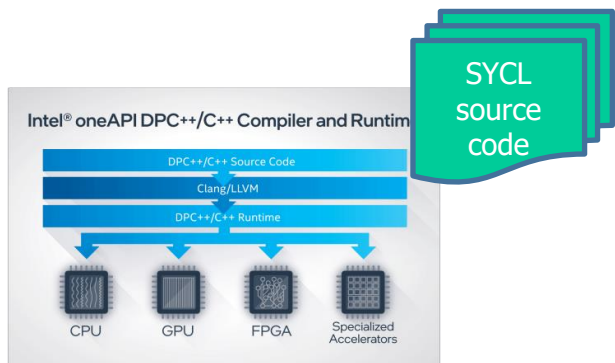
Performance comparison



- SYCL uses the same C++ performance model as CUDA, so it achieves very similar performance for the same code
- SYCL makes it easy to write *parameterizable* code that adapts the algorithms to underlying hardware: this enables automatic optimizations that increase performance



oneAPI and SYCL

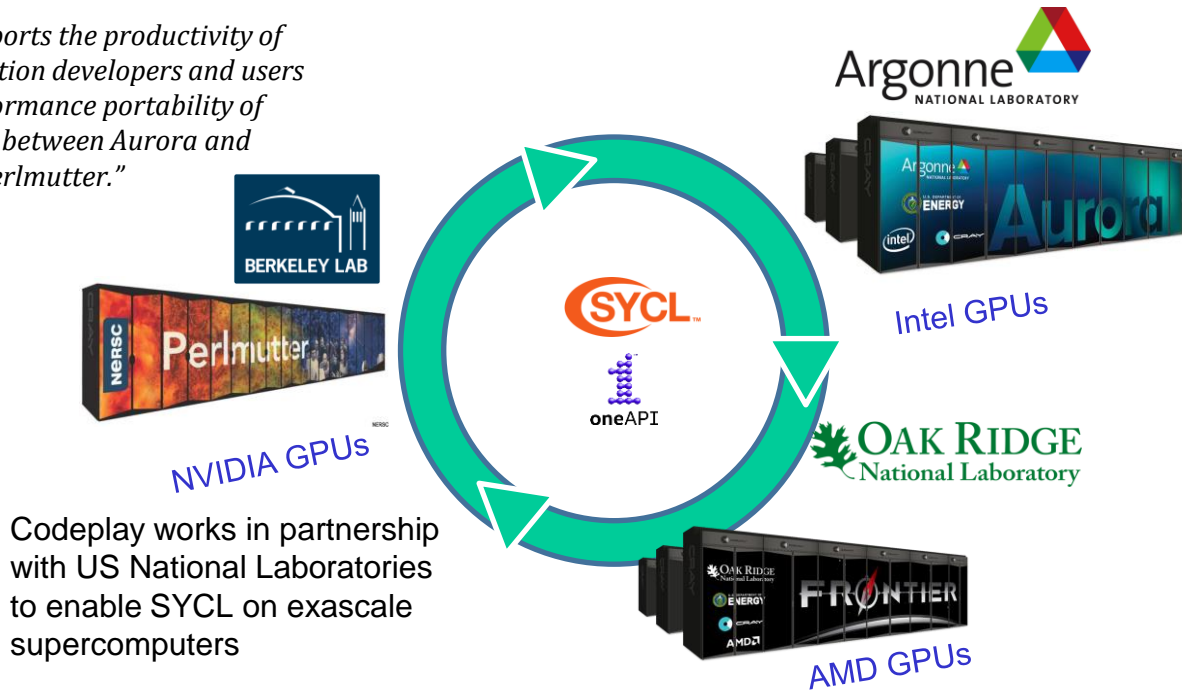


- SYCL sits at the heart of oneAPI
- Provides an open standard interface for developers
- Defined by the industry



SYCL Enables Supercomputers

"this work supports the productivity of scientific application developers and users through performance portability of applications between Aurora and Perlmutter."



Enables a broad range of software frameworks and applications



Towards a Portable Pipeline in Drug Discovery

- The LIGATE project aims at building a portable drug discovery pipeline
 - Funded from the European High-Performance Computing Joint Undertaking Joint Undertaking (JU)
- HPC is heading toward specialization and extreme heterogeneity
- SYCL enables to write platform independent code, while keeping native-comparable performance



<https://www.ligateproject.eu/>

This work is licensed under a Creative Commons Attribution 4.0 International License

© The Khronos® G




SYCL as a universal programming model for HPC

Starting with US National Labs

Across Europe, Asia are many Petascale and pre-exascale systems

- With many variety of CPUs GPUs FPGAs, custom devices
- Often with interconnected usage agreements
- Europe EPI: **hipSYCL in Leonardo, Lumi and Karolina**



Easier, Short Path to Heterogeneous Programming

Coming Soon: An Open Source CUDA* to SYCL* Code Migration Tool - Project SYCLomatic, DPCT

SYCL* with oneAPI open, cross-architecture, standards-based programming

Allows developers to expand the value of their investments across architectures

Provides choice in hardware & freedom from proprietary, single-vendor lock-in

Intel is providing a CUDA* to SYCL* migration tool:
Project SYCLomatic

Enables developers a productive path to create single-source, portable code for hardware targets regardless of vendor

Simplifies development while delivering performance, reduces time & costs for code maintenance


A community to share, collaborate & contribute software technologies

Available on GitHub by end of May

github.com/oneapi-src/SYCLomatic github.com/oneapi-src/SYCLomatic-test

Use the tool, please provide feedback!


Simplify Heterogeneous Development
From Diverse Programming Approaches



To Productive, Performant
Cross-architecture, Cross-vendor Programming



Project SYCLomatic: CUDA to SYCL Code



SYCLomatic Tool Usage Workflow

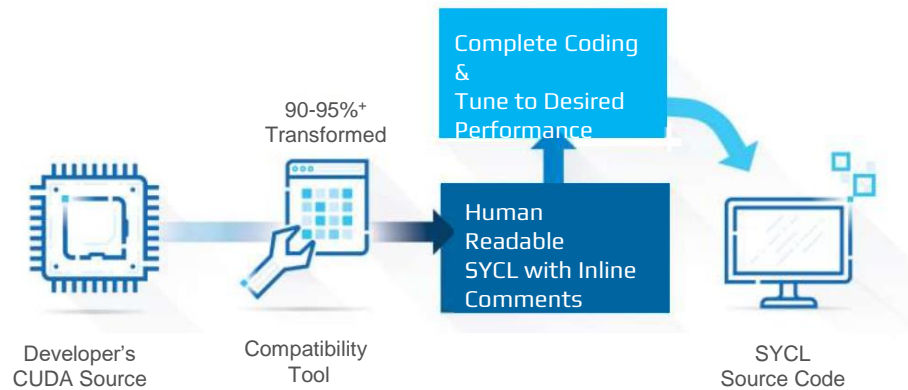
(newly open sourced by the end of May)

Collect compilation options of the Developer's CUDA* source from project build scripts, eg. Makefile, vcxproj file

Assist developers migrating code written in CUDA to SYCL by generating SYCL code wherever possible

Typically, 90%-95%+ of CUDA code automatically migrates to SYCL code

Inline comments are provided to help developer complete and tune the code



* Intel estimates as of September 2021. Based on measurements on a set of 70 HPC benchmarks and samples, with examples like Babel, SHOC, PENNANT. Results may vary.

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




Case Study: Zuse Institute Berlin





From CUDA to DPC++ and back to Nvidia GPUs... and FPGAs

A oneAPI case study with the tsunami simulation easyWave



- Originally written for NVIDIA GPUs with CUDA
- Auto-converted to SYCL (DPC++) with Intel's Compatibility Tool
- Resulting SYCL code runs on: NVIDIA GPU, Intel GPU, Intel CPU, Intel FPGA [+ any new processor with SYCL support]

Even on NVIDIA GPU, SYCL code is only 4% slower (with Codeplay developed open-source compiler)




N-Body

- Simulates gravitational interaction in a fictional galaxy

$$\vec{F}_i = - \sum_{i \neq j} G \frac{(\vec{r}_i - \vec{r}_j)}{|\vec{r}_i - \vec{r}_j|^3}$$

<https://github.com/topics/n-body-simulator>



Familiar Standard C++ Code

CUDA

```
void DiskGalaxySimulator::stepSim() {  
    // Compute updated positions  
    constexpr int wg_size = 256;  
    int nblocks = ((getNumParticles() - 1) / wg_size) + 1;  
  
    // Profiling info - rather than using the CUDA event recording  
    // approach, we are instead measuring the time from before kernel  
    // submission until host synchronization. This is more portable via  
    // dpct.  
    auto start = std::chrono::steady_clock::now();  
    for (size_t i = 0; i < params.maxIterationsPerFrame; i++) {  
        particle_interaction<<<nblocks, wg_size>>>(pos_d, pos_next_d, vel_d,  
                                                    params);  
        std::swap(pos_d, pos_next_d);  
    }  
    gpuErrchk(cudaDeviceSynchronize());  
    auto stop = std::chrono::steady_clock::now();  
    lastStepTime =  
        std::chrono::duration<float, std::milli>(stop - start)  
            .count();  
  
    // Sync data  
    recvFromDevice();  
}
```

Code changes
are minimal

SYCL

```
void DiskGalaxySimulator::stepSim() {  
    // Compute updated positions  
    constexpr int wg_size = 256;  
    int nblocks = ((getNumParticles() - 1) / wg_size) + 1;  
  
    auto start = std::chrono::steady_clock::now();  
    for (size_t i = 0; i < params.maxIterationsPerFrame; i++) {  
        dpct::get_default_queue().submit([i](sycl::handler &cgh) {  
            auto pos_d_ct0 = pos_d;  
            auto pos_next_d_ct1 = pos_next_d;  
            auto vel_d_ct2 = vel_d;  
            auto params_ct3 = params;  
  
            cgh.parallel_for(sycl::nd_range<1>(sycl::range<1>(nblocks) *  
                                                sycl::range<1>(wg_size),  
                                                sycl::range<1>(wg_size)),  
                            [=](sycl::nd_item<1> item_ct1) {  
                                particle_interaction(pos_d_ct0, pos_next_d_ct1,  
                                                    vel_d_ct2, params_ct3,  
                                                    item_ct1);  
                            });  
        });  
        std::swap(pos_d, pos_next_d);  
    }  
    gpuErrchk((dpct::get_current_device().queues_wait_and_throw(), 0));  
    auto stop = std::chrono::steady_clock::now();  
    lastStepTime =  
        std::chrono::duration<float, std::milli>(stop - start)  
            .count();  
  
    // Sync data  
    recvFromDevice();  
}
```



Performance Achieved

N-body demo running with CUDA on device: NVIDIA GeForce RTX 3060
Kernel time: 10.20 ms

N-body demo running with DPC++ on device: NVIDIA GeForce RTX 3060
Kernel time: 8.79 ms

www.codeplay.com/oneapiforcuda/



JOIN SYCL Safety-Critical Exploratory Forum Now

Exploring real-world industry requirements for
open and *royalty-free* high-level compute APIs
suitable for safety-critical markets

Proven Khronos
Exploratory Process to
ensure industry
requirements are fully
understood before starting
standardization initiatives

More information and
signup instructions
<https://www.khronos.org/syclsc>

Khronos SYCL Safety-Critical Exploratory Forum



Online discussion forum and
weekly Zoom calls

**No detailed design activity
to protect participants IP**

Explore if consensus can be built around an
agreed **Scope of Work** document

Discuss what standardization activities can
best execute actions in the Scope of Work

Scope of
Work
Document

Agreed SOW
document released
from NDA and made
public

Initiation of
Khronos Working
Group to execute
the SOW

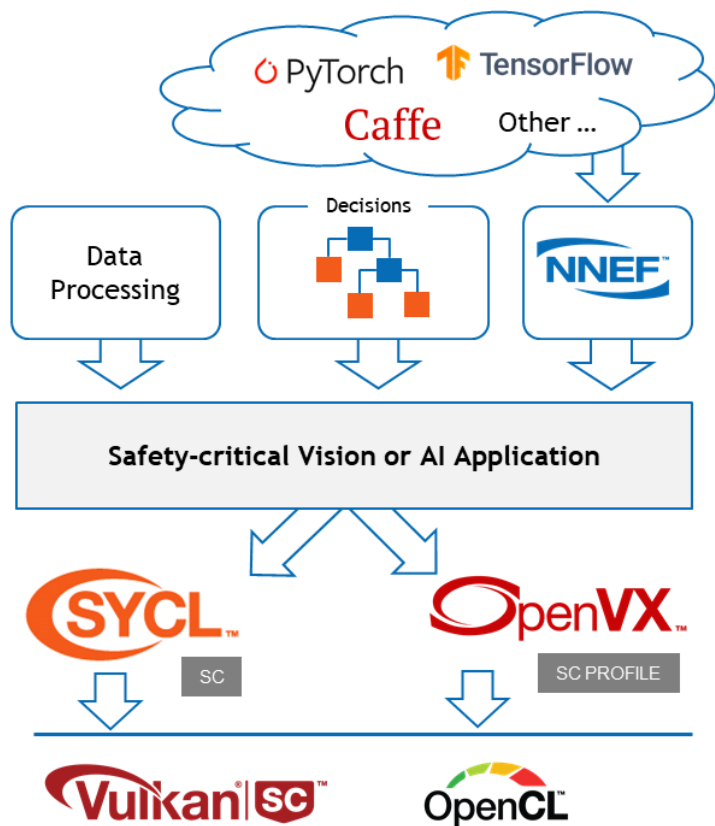
Any company is
welcome to join

No cost or IP
Licensing obligations

Project NDA to cover
Exploratory Forum
Discussions



SYCL SC in the Khronos SC Ecosystem




Neural network models are trained in the cloud using a variety of platforms.

Once the model is trained it is exported and converted to NNEF before being passed to a safety-critical API for inferencing.

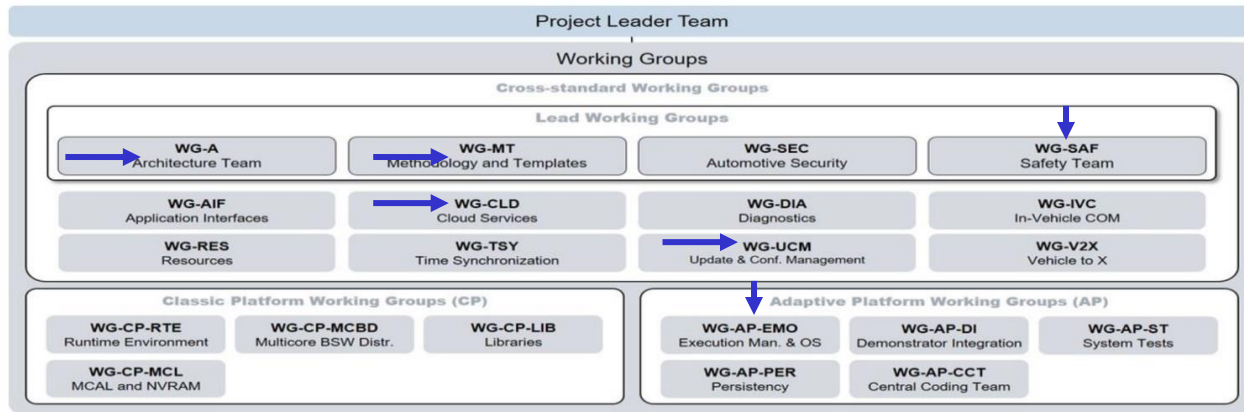
OpenVX provides high level APIs for Vision and AI with a safety-critical profile, enabling applications to quickly deploy trained NN models.

SYCL SC provides a general parallel programming API for accelerated compute at the C++ level. A typical AI application pipeline will combine the discreet functionality exposed by OpenVX with proprietary algorithms written using SYCL SC involving data pre-processing and post-processing, as well as complex decision making.

Vulkan SC or **OpenCL** are lower, execution-level APIs that can be used to accelerate higher-level APIs like SYCL SC & OpenVX



Khronos AUTOSAR Liaison What next?



CROSS-STANDARD LEAD WORKING GROUPS (FO, CP, AP)



Design guidelines for using parallel processing technologies on Adaptive Platform
AUTOSAR AP R19-11

Document Title	Design guidelines for using parallel processing technologies on Adaptive Platform
Document Owner	AUTOSAR
Document Responsibility	AUTOSAR
Document Identification No	884
Document Status	published
Part of AUTOSAR Standard	Adaptive Platform
Part of Standard Release	R19-11

3D Graphics

Desktop, Mobile and Web

3D Assets

Authoring and Delivery

Portable XR

Augmented and Virtual Reality

Parallel Computation


Vision, Inferencing, Machine Learning




Case Study: Automotive AI



- Codeplay provides a full open-standards software stack (SYCL/OpenCL/SPIR-V) for the Renesas R-Car AI system-on-chip
- This enables easy transition from NVIDIA GPU platforms to R-Car platform



 **Acoran**



"For Renesas, SYCL is a key enabler for automotive ADAS/AD software developers that allows them to easily use the highly-efficient, heterogeneous accelerators of the R-Car SoC Series through the open Khronos standard,"



Final words

- Programming Models Must Persist
- SYCL 2020 Launched February 2021
- SYCL user and developer Phenomenal Growth
- Market needs SYCL: easy to build SYCL on any device
- SYCL is mainstream
- Market needs SYCL to Evolve (call to action)
- **SYCL can serve even more Extreme Heterogeneity**
- **Data Movement is still King**
- **We need Codesign with extreme Heterogeneity**
- Future SYCL: Emerging transformative technologies
- SYCL can be a part of a standard programming model for all HPC ,Embedded AI/ML, and Automotive
- SYCL is an open standard with multiple company contributions, lots of European/Asia projects



Enabling Industry Engagement

- SYCL working group values industry feedback
 - <https://community.khronos.org/c/sycl>
 - <https://sycl.tech>
- SYCL Academy
 - <https://github.com/codeplaysoftware/syclacademy>
- SYCL FAQ
 - <https://www.khronos.org/blog/sycl-2020-what-do-you-need-to-know>
- SYCL Survey coming

- **Advisory Panel**
Chaired by Tom
Deakin of U of Bristol
- Regular meetings to give feedback on roadmap and draft specifications

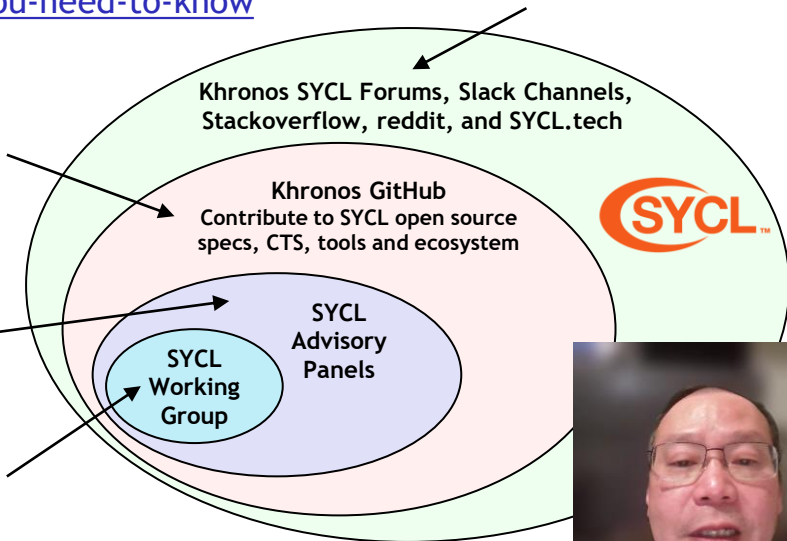
Public contributions to Specification, Conformance Tests and software
<https://github.com/KhronosGroup/SYCL-CTS>
<https://github.com/KhronosGroup/SYCL-Docs>
<https://github.com/KhronosGroup/SYCL-Shared>
<https://github.com/KhronosGroup/SYCL-Registry>
<https://github.com/KhronosGroup/SyclParallelSTL>

Invited Experts

<https://www.khronos.org/advisors/>

Khronos members

<https://www.khronos.org/members/>
<https://www.khronos.org/registry/SYCL/>



Open to all!

<https://community.khronos.org/www.khr.io/slack>
<https://app.slack.com/client/TDMDFS87M/CE9UX4CHG>
<https://community.khronos.org/c/sycl/>
<https://stackoverflow.com/questions/tagged/sycl>
<https://www.reddit.com/r/sycl>
<https://github.com/codeplaysoftware/syclacademy>
<https://sycl.tech/>

