

Intel[®] DPC++ Compatibility Tool

Migrating CUDA Codes to DPC++



intel[®]

Agenda

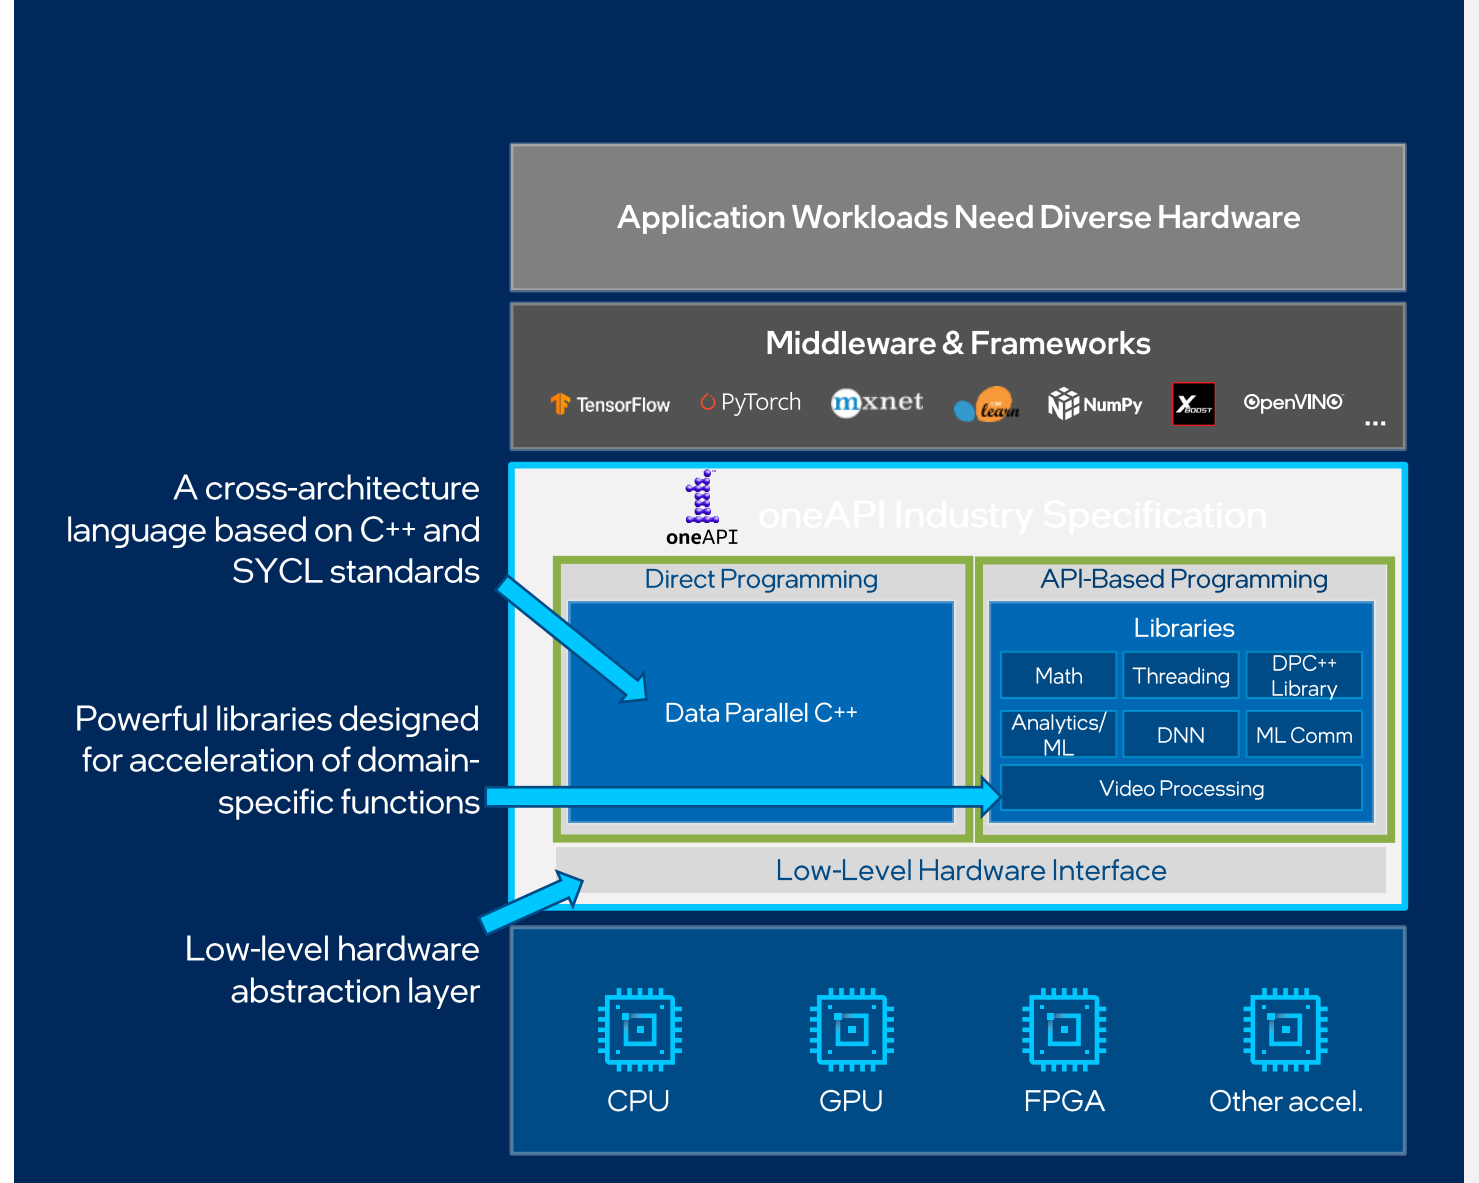
- oneAPI Brief Overview
- Intel® DPC++ Compatibility Tool Workflow
- Migration Flow and Vector-Add Example
- Demo Tutorial
 - Simple CUDA* File Project
 - Migrate Multi CUDA Files Project
- Best Known Methods for Migration
- Eclipse and Visual Studio Integration
- Key Takeaways

oneAPI Industry Initiative

Break the Chains of Proprietary Lock-in

Open to promote community and industry collaboration

Enables code reuse across architectures and vendors



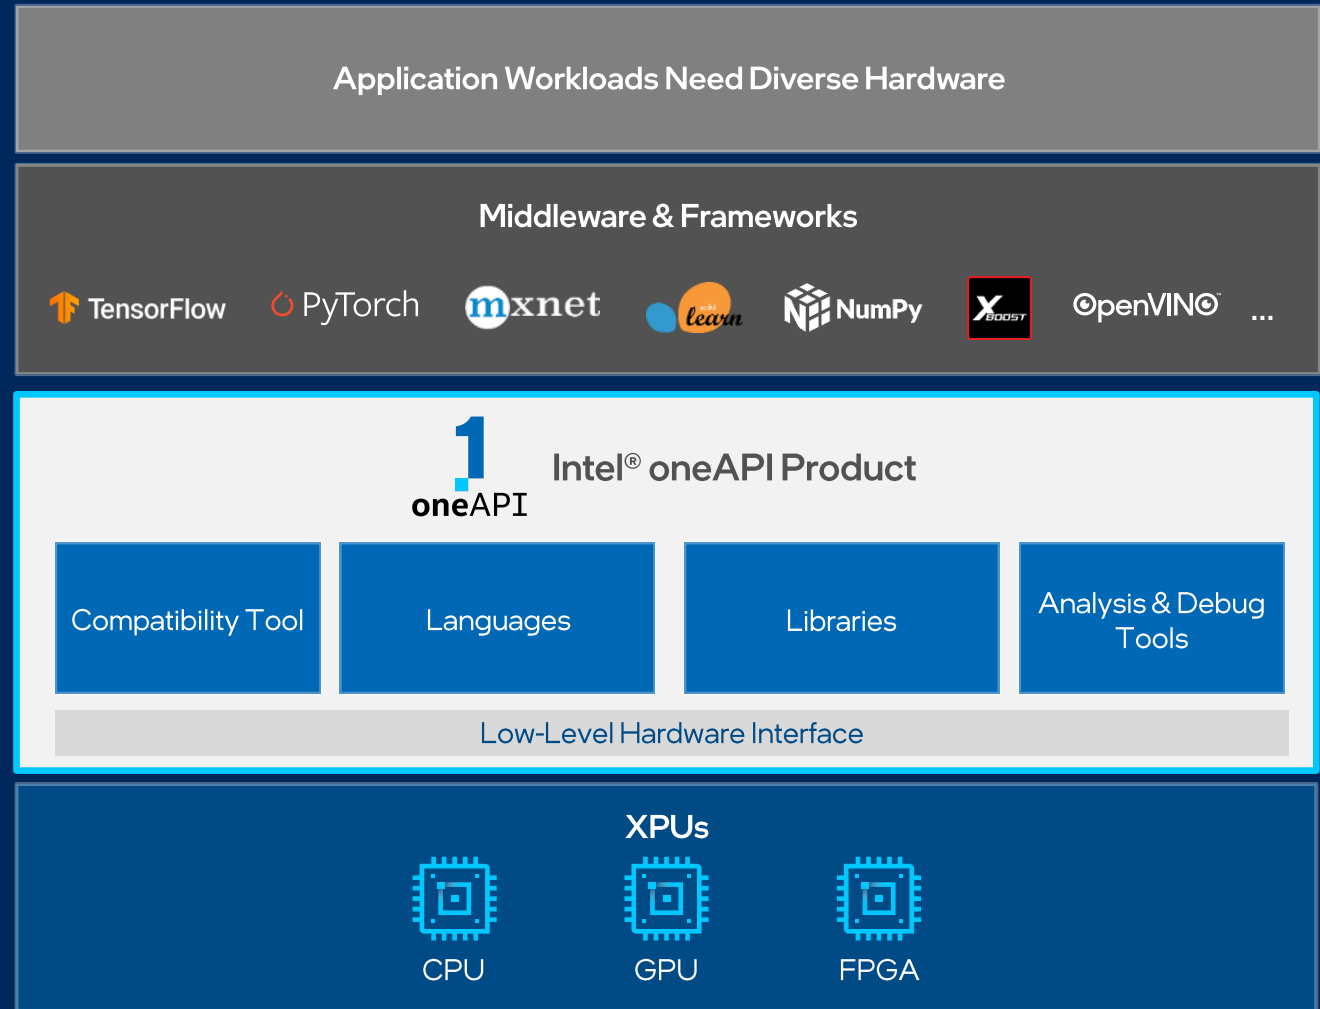
The productive, smart path to freedom for accelerated computing from the economic and technical burdens of proprietary programming models

Intel® oneAPI Product

Built on Intel's Rich Heritage of CPU Tools Expanded to XPU

A complete set of advanced compilers, libraries, and porting, analysis and debugger tools

- Accelerates compute by exploiting cutting-edge hardware features
- Interoperable with existing programming models and code bases (C++, Fortran, Python, OpenMP, etc.), developers can be confident that existing applications work seamlessly with oneAPI
- Eases transitions to new systems and accelerators—using a single code base frees developers to invest more time on innovation



[Available Now](#)

Intel® DPC++ Compatibility Tool

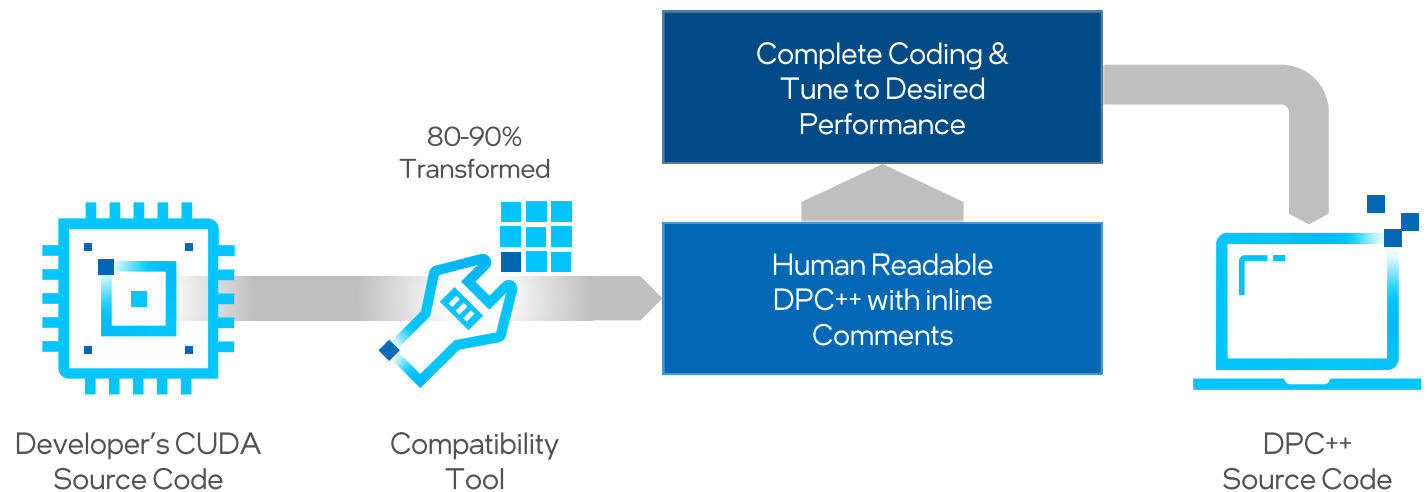
Minimizes Code Migration Time

Assists developers migrating code written in CUDA to DPC++ once, generating **human readable** code wherever possible

~80-90% of code typically migrates automatically

Inline comments are provided to help developers finish porting the application

Intel DPC ++ Compatibility Tool Usage Flow



Intel® oneAPI Base Toolkit

Accelerate Data-centric Workloads

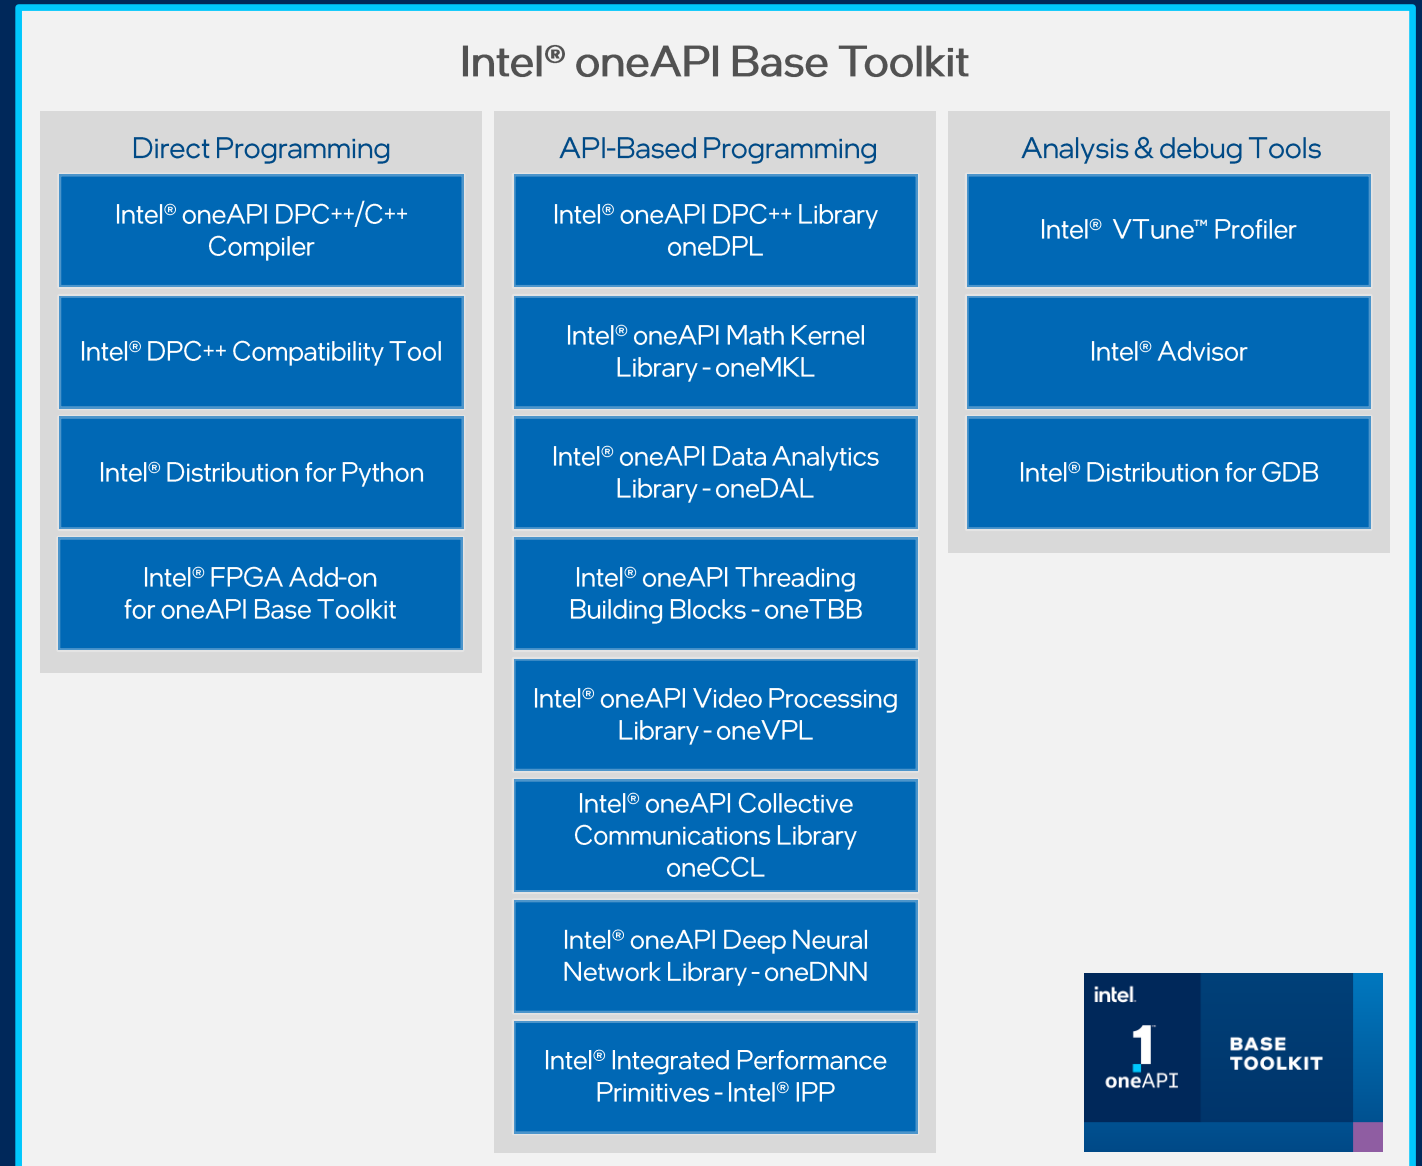
A core set of core tools and libraries for developing high-performance applications on Intel® CPUs, GPUs, and FPGAs.

Who Uses It?

- A broad range of developers across industries
- Add-on toolkit users since this is the base for all toolkits

Top Features/Benefits

- Data Parallel C++ compiler, library and analysis tools
- DPC++ Compatibility tool helps migrate existing code written in CUDA
- Python distribution includes accelerated scikit-learn, NumPy, SciPy libraries
- Optimized performance libraries for threading, math, data analytics, deep learning, and video/image/signal processing



Migrating Vector Add Example

Migrating Simple Example

- `dpct [options] [<source0>... <sourceN>]`
 - Ensure supported CUDA header files are available
 - May use `-cuda-include-path`
- Bult-in Usage Information
 - `dpct --help`

Vector-Add Example: Migration with Intel® DPC++ Compatibility Tool

CUDA

DPC++

```
#include <cuda.h>
#include <stdio.h>
#define VECTOR_SIZE 256
```

```
global void VectorAddKernel(float* A, float* B, float* C)
{
    A[threadIdx.x] = threadIdx.x + 1.0f;
    B[threadIdx.x] = threadIdx.x + 1.0f;
    C[threadIdx.x] = A[threadIdx.x] + B[threadIdx.x];
}
```

```
int main()
{
    float *d_A, *d_B, *d_C;

    cudaMalloc(&d_A, VECTOR_SIZE*sizeof(float));
    cudaMalloc(&d_B, VECTOR_SIZE*sizeof(float));
    cudaMalloc(&d_C, VECTOR_SIZE*sizeof(float));
}
```

```
#include <CL/sycl.hpp>
#include <dpct/dpct.hpp>
#define VECTOR_SIZE 256
```

```
void VectorAddKernel(float* A, float* B, float* C, sycl::nd_item<3>
item_ct1)
{
    A[item_ct1.get_local_id(2)] = item_ct1.get_local_id(2) + 1.0f;
    B[item_ct1.get_local_id(2)] = item_ct1.get_local_id(2) + 1.0f;
    C[item_ct1.get_local_id(2)] =
        A[item_ct1.get_local_id(2)] + B[item_ct1.get_local_id(2)];
}
```

```
int main()
{
    dpct::device_ext &dev_ct1 = dpct::get_current_device();
    sycl::queue &q_ct1 = dev_ct1.default_queue();
    float *d_A, *d_B, *d_C;

    d_A = sycl::malloc_device<float>(VECTOR_SIZE, q_ct1);
    d_B = sycl::malloc_device<float>(VECTOR_SIZE, q_ct1);
    d_C = sycl::malloc_device<float>(VECTOR_SIZE, q_ct1);
}
```

<https://github.com/oneapi-src/oneAPI-samples/tree/master/Tools/Migration/vector-add-dpct>

Vector-Add Migration Example (continued)

CUDA

DPC++

```
VectorAddKernel<<<1, VECTOR_SIZE>>>(d_A, d_B, d_C);
```

```
q_ct1.submit([&](sycl::handler &cgh) {  
    cgh.parallel_for(sycl::nd_range<3>(  
        sycl::range<3>(1, 1, VECTOR_SIZE),  
        sycl::range<3>(1, 1, VECTOR_SIZE)),  
        [=](sycl::nd_item<3> item_ct1) {  
            VectorAddKernel(d_A, d_B, d_C, item_ct1);  
        });  
});
```

4

```
float Result[VECTOR_SIZE] = { };  
cudaMemcpy(Result, d_C, VECTOR_SIZE*sizeof(float),  
           cudaMemcpyDeviceToHost);
```

```
float Result[VECTOR_SIZE] = { };  
q_ct1.memcpy(Result, d_C, VECTOR_SIZE * sizeof(float)).wait();
```

5

```
cudaFree(d_A);  
cudaFree(d_B);  
cudaFree(d_C);
```

```
sycl::free(d_A, q_ct1);  
sycl::free(d_B, q_ct1);  
sycl::free(d_C, q_ct1);
```

6

```
for (int i = 0; i < VECTOR_SIZE; i++) {  
    if (i % 16 == 0) {  
        printf("\n");  
    }  
    printf("%f ", Result[i]);  
}  
  
return 0;
```

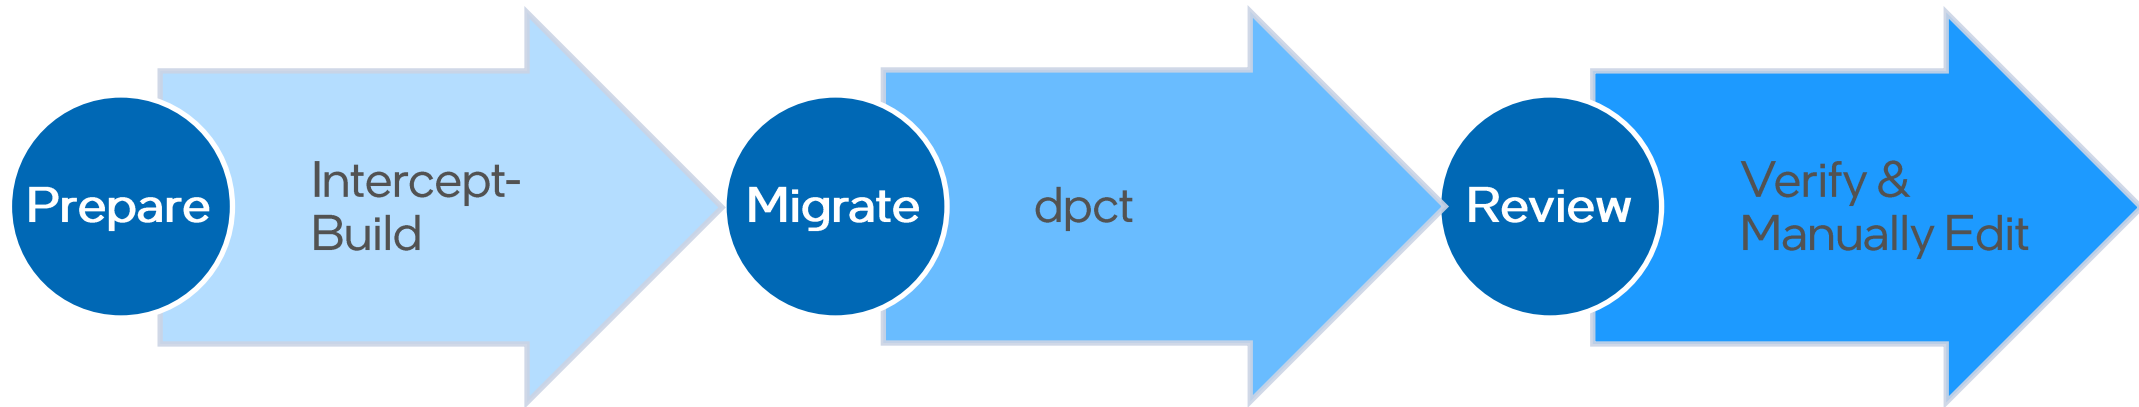
```
for (int i = 0; i < VECTOR_SIZE; i++) {  
    if (i % 16 == 0) {  
        printf("\n");  
    }  
    printf("%f ", Result[i]);  
}  
  
return 0;
```

7

Migrating Needleman Wunsch and HydroC Examples

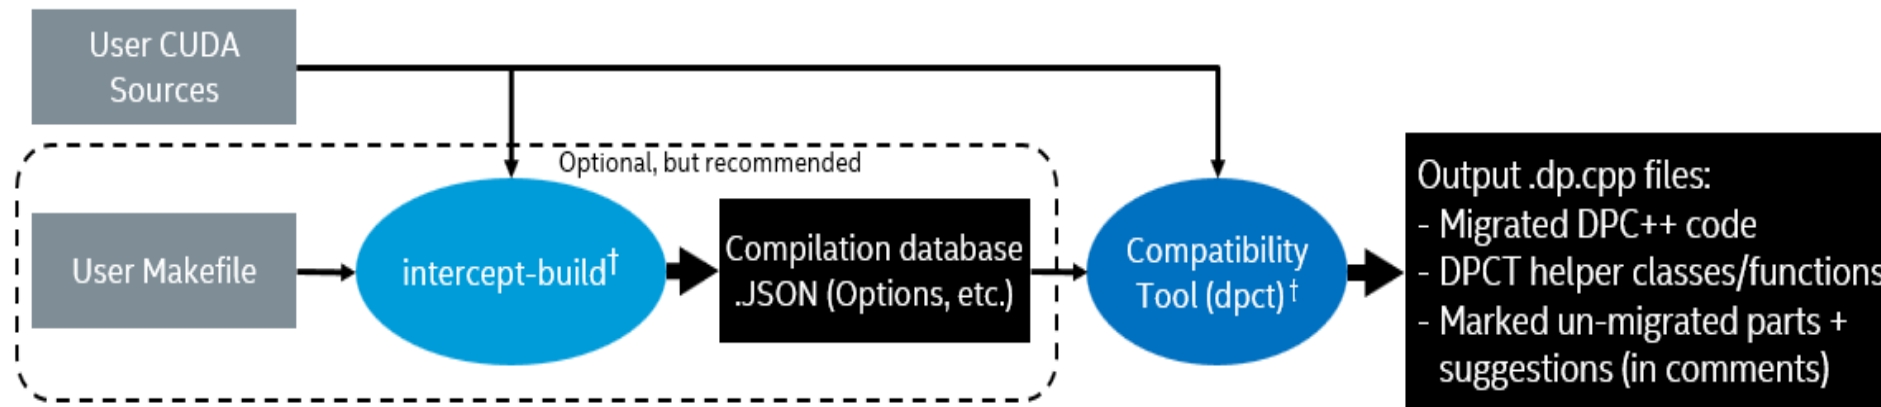
Migration Flow

Typical preparation steps for simple to complex projects



Intercept Build

- Use intercept-build to create a compilation database
 - For projects that use Make or Cmake
 - Keeps track of compilation options, settings, macro definitions, include paths, etc.
 - Creates a JSON file containing the build commands
- Run “make clean” before “intercept-build”



† Certain CUDA language header files may need to be accessible to the Intel® DPC++ Compatibility Tool

DPCT Basic Options

- `dpct [options] [<source0>... <sourceN>]`

DPCT Basic Options	
<code>--in-root</code>	Path to the root of the source tree to be migrated
<code>--out-root</code>	Path to root of generated files.
<code>-p</code>	Path to compile database JSON file
<code>--process-all</code>	Migrates/copies all files from <code>--in-root</code> directory to the <code>--out-root</code> directory, eliminating need to specify <code>.cu</code> files one by one
<code>--extra-arg</code>	Specify more Clang compiler options. e.g. <code>dpct --extra-arg="-std=c++14" --extra-arg="-l..."</code>
<code>--format-style</code>	Sets formatting style for output files. e.g. <code>=llvm, =google, =custom</code> (Uses <code>.clang-format</code> file)
<code>--format-range</code>	Code formatting applied to no code (<code>=none</code>), migrated code (<code>=migrated</code>), or all code (<code>=all</code>)

DPCT Recommended Options

- `dpct [options] [<source0>... <sourceN>]`

DPCT Options that Ease Migration/Debug	
<code>--keep-original-code</code>	Keep original CUDA code in the comments of generated DPC++ file. Allows easy comparison of original CUDA code to generated DPC++ code.
<code>--comments</code>	Insert comments explaining the generated code
<code>---always-use-async-handler</code>	Always create <code>cl::sycl::queue</code> with the async exception handler

- Many other options available use `dpct --help`

DPCT Namespace Usage

- DPCT namespace provides helper function and macros to assist the migration of input source code.
 - `dpct::`
- Implemented in header files (`include/dpct`)
- Intended to become part of your code.
- Examples: `dpct_malloc`, `dpct_memcpy`, `get_buffer`, `get_default_queue`, `get_default_context`
- Not recommended to use these when writing new DPC++ code

General Best Known Methods (BKMs)

- Migrate Incrementally
 - If you see *dpct* generate multiple errors when migrating a long list of CUDA source files in one run, do it one-by-one
- Start with a clean project - “make clean” before running “intercept-build make”
- Run *intercept-build make -k* to keep going when some targets can't be made when generating compilation database

Code Modifications Prior to Migration

- Ensure source files are syntactically correct
- Possibly needed due to differences between clang and nvcc
 1. Namespace qualification maybe needed in certain scenario with clang parser
 2. Additional forward class declarations may be needed by clang
 3. Space within the triple brackets of kernel invocation are tolerated by nvcc but not clang
 - e.g. `cuda_kernel<< <num_blocks, threads_per_block>> >(args...)`
- See [Compilation CUDA with clang](#) on llvm.org for more details.

Unified Shared Memory (USM) Usage

- DPC++ supports USM that allows pointer-based approach to manage host and device memory.
- USM produces less volume code compare to SYCL buffers
- The Compatibility Tool uses USM by default.
- May be trouble some for non-Intel compilers targeting non-Intel hardware

DPCT USM Option	
--usm-level	Sets Unified Shared Memory (USM) level. =Restricted: Use USM (default) =none: Uses helper functions and SYCL buffers

Diagnostics Reference

Compatibility Tool highlights issues with migration and code comments

*/path/to/file:20:1: warning:
DPCT10XX:0: text of the warning*

*//source code line for which warning
was generated*

ID	Message	Detailed Help	Suggestions to Fix
DPCT1000	An error handling if-stmt was detected but could not be rewritten. See the details in the resulting file comments.	<p>The CUDA* API return error codes that are consumed by the program logic. SYCL* uses exceptions to report errors and does not return the error code.</p> <p>When the error handling logic in the original code is simple (for example, a print error message and exit), the code is removed in the resulting Data Parallel C++ (DPC++) application. The expectation is that SYCL throws an exception, which is handled with the printing of an exception message and exiting (the exception handler is generated automatically by the Intel® DPC++ Compatibility Tool).</p> <p>This warning is generated when the Intel® DPC++ Compatibility Tool detects more complex error handling than it considers safe to remove.</p>	Review the error handling if-statement and try to rewrite it to use an exception handler instead.
DPCT1001	The statement could not be removed. See the details in the resulting file comments.	The Intel® DPC++ Compatibility Tool was not able to remove the code in the then clause of if-stmt. See DPCT1000.	See DPCT1000.
DPCT1002	A special case error handling if-stmt was detected. You may need to rewrite this code.	See DPCT1000	See DPCT1000.

[See Compatibility Tool – Diagnostics Reference](#)

Code Review or Rewrite Needed

[Diagnostic Reference](#)

- Error code logic replaced with (*,0) code or commented out
- Equivalent DPC++ API not available
- CUDA Compute Capability-dependent logic
- Hardware-dependent API (clock())
- Migration not supported for some API
- Execution time measurement logic
- Handling built-in vector type conflicts
- Migration of cuBLAS API (Review arguments list)

Demo: Simple CUDA Project Migration

- Rodinia Benchmark Suite v3.1 – Introduction
- Setting/Verifying the Environment for Intel® DPC++ Compatibility Tool
- Demo
 - Planning for Migration
 - Compatibility Tool Options
 - Migrating Needleman Wunsch Application

<http://rodinia.cs.virginia.edu/doku.php>

<https://software.intel.com/en-us/get-started-with-intel-dpcpp-compatibility-tool>

<https://software.intel.com/content/www/us/en/develop/documentation/intel-dpcpp-compatibility-tool-user-guide/top/migrate-a-project/migrate-a-project-on-linux.html>

<https://software.intel.com/en-us/intel-dpcpp-compatibility-tool-user-guide-usage-workflow-overview>

Demo: HydroC - Multi CUDA Files Project Migration

- Setting/Verifying the Environment for Intel[®] DPC++ Compatibility Tool
- Demo
 - Planning for Migration; Understanding the Application File ...
 - *intercept-build* Options
 - Compatibility Tool Options
 - Migrating HydroC Application

https://github.com/HydroBench/Hydro/tree/master/HydroC/cuHydroC_2DMpi/Src

<https://github.com/HydroBench/Hydro>

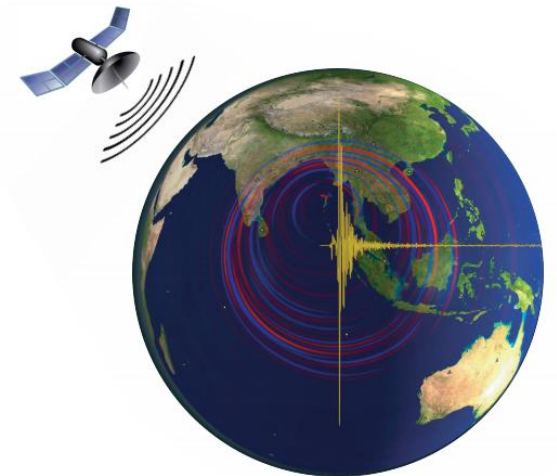
<https://github.com/HydroBench/Hydro/blob/master/License.txt>

SPECFEM3D_GLOBE

- SPECFEM3D_GLOBE simulates global and regional (continental-scale) seismic wave propagation
- Official repo: https://github.com/geodynamics/specfem3d_globe

```
github.com/AlDanial/cloc v 1.74 T=1.44 s (370.8 files/s, 156306.8 lines/s)
```

Language	files	blank	comment	code
Fortran 90	279	27677	41716	100021
C	81	3145	5405	20851
CUDA	88	1410	2286	10841
Ruby	61	554	192	4365
make	17	532	817	1887
C/C++ Header	5	284	370	995
C++	1	196	229	773
Markdown	1	31	0	102
SUM:	533	33829	51015	139835



SPECFEM3D_GLOBE – Migration to DPC++

Prepare

```
$ git clone --recursive --branch devel https://github.com/geodynamics/specfem3d\_globe.git
$ ./configure --with-cuda=cuda9 CUDA_LIB=${CUDA_ROOT}/lib64/ \
    CUDA_INC=${CUDA_ROOT}/include/ MPI_INC=${I_MPI_ROOT}/include/
$ intercept-build make -i
```

Migrate

```
$ dpct -p compile_commands.json
```

Review

Review diagnostics messages using [reference](#) and manually edit
Address other not-so-obvious issues

Diagnostics Messages Breakdown

DPCT{diagnostics#} (count)	Summary
DPCT1000 (6), DPCT1001(6), DPCT1003 (111), DPCT1009 (8), DPCT1010 (3), DPCT1024 (2)	Different scenarios for error handling
DPCT1005 (4), DPCT1012 (4), DPCT1017 (10), DPCT1019 (1), DPCT1022 (1), DPCT1026 (5), DPCT1027 (3), DPCT1051 (4)	Unavailable equivalent API's in SYCL* (e.g. device versions, certain device properties, timing logic)
DPCT1039 (9)	Handling atomics (global atomics by default, local will need intervention)
DPCT1049 (59)	Validating use of work-group sizes

Using plugins with IDE

Eclipse: Gaussian

The screenshot shows the Eclipse IDE interface. The main editor displays the file `gaussian.dp.cpp` with the following code snippet:

```
212     memset( &deviceProp, 0, sizeof(deviceProp));
213     /*
214     DPCT1003:0: Migrated API does not return error code.
215     may need to rewrite this code.
216     */
217     if (0 == (dpct::dev_mgr::instance()
218             .get_device(nDeviceIdx)
219             .get_device_info(deviceProp),
220             0))
221     {
222         printf("\nDevice Name \t\t - %s ", deviceProp.g
223              printf( "\n*****
224         printf("\nTotal Global Memory\t\t\t - %lu KB",
225              deviceProp.get_global_mem_size() / 1024)
226     /*
```

Below the code, the Intel(R) DPC++ Compatibility Tool shows a list of warnings:

Migrated Source File Location	Source File Location	Type	ID	Message	Actions
/home/intel...p, Line 216	/home/int... Line 210	warning	DPCT1003	Migrated API does not return error code. (*	Fix Help
/home/intel...p, Line 229	/home/int... Line 215	warning	DPCT1019	local_mem_size in SYCL is not a complete e	Fix Help
/home/intel...p, Line 242	/home/int... Line 221	warning	DPCT1022	There is no exact match between the maxG	Fix Help
/home/intel...p, Line 251	/home/int... Line 223	warning	DPCT1005	The device version is different. You need to	Fix Help
/home/intel...p, Line 266	/home/int... Line 230	warning	DPCT1009	SYCL uses exceptions to report errors and e	Fix Help
/home/intel...p, Line 271	/home/int... Line 230	warning	DPCT1010	SYCL uses exceptions to report errors and e	Fix Help
/home/intel...p, Line 569	/home/int... Line 462	warning	DPCT1010	SYCL uses exceptions to report errors and e	Fix Help

The Cheat Sheets window shows the "Intel(R) DPC++ Compatibility Tool Cheat Sheet" with an introduction and a specific warning: "DPCT1003: The migrated API does not return an error code, so (Typically, this happens because the CUDA API returns an error code an consumed by the program logic. SYCL uses exceptions to report errors and does not return the error cc The Intel(R) DPC++ Compatibility Tool inserts a (*, 0) operator, so that application could be compiled. This operator returns 0 and is inserted code is expected by the program logic and the new API does not return recommendation is to review all such places in the code.

Visual Studio 2019: Gaussian

```
219     int nDevCount = 0;
220
221     /* DPCT_ORIG   cudaGetDeviceCount( &nDevCount ); */
222     nDevCount = dpct::dev_mgr::instance().device_count();
223     printf( "Total Device found: %d", nDevCount );
224     for (int nDeviceIdx = 0; nDeviceIdx < nDevCount; ++nDeviceIdx )
225     {
226         memset( &deviceProp, 0, sizeof(deviceProp));
227         /* DPCT_ORIG   if( cudaSuccess ==
228            /*   cudaGetDeviceProperties(&deviceProp, nDeviceIdx) */
229            /*
230            DPCT1003:0: Migrated API does not return error code. (*, 0) is inserted.
231            You may need to rewrite this code.
232            */
233            if (0 == (dpct::dev_mgr::instance()
234                .get_device(nDeviceIdx)
235                .get_device_info(deviceProp,
236                    0))
237                {
238                /* DPCT_ORIG   printf( "\nDevice Name \t\t - %s ",
239                /*   deviceProp.name ); */
240                printf("\nDevice Name \t\t - %s ", deviceProp.get_name());
241                printf( "\n*****");
242                /* DPCT_ORIG   printf( "\nTotal Global Memory\t\t\t - %lu
243                /*   KB", deviceProp.totalGlobalMem/1024 ); */
244                printf("\nTotal Global Memory\t\t\t - %lu KB",
245                    deviceProp.get_global_mem_size() / 1024);
```

DPCT1003 : The migrated API does not return an error code, so (*, 0) is inserted

Typically, this happens because the CUDA API returns an error code and then it is consumed by the program logic.

SYCL uses exceptions to report errors and does not return the error code.

The Intel® DPC++ Compatibility Tool inserts a (*, 0) operator, so that the resulting application could be compiled. This operator returns 0 and is inserted if the return code is expected by the program logic and the new API does not return it. The recommendation is to review all such places in the code.

If in a DPC++ application you:

- Do not need the code that consumes the error code, remove the code and the (*, 0) operator.
- Need the code that consumes the error code, try to replace it with an exception handling code and use your logic in an exception handler.

Migrated Source File Location	Source File Location	Type	ID	Message	Actions
C:\temp\SampleCode\rodingia_3.1\cuda\gaiC:\temp\SampleCode\rodingia_3.1\cuda\gai	Warning	DPCT1003	Migrated API does not return error code. (*, 0) is inserted.	Help	
C:\temp\SampleCode\rodingia_3.1\cuda\gaiC:\temp\SampleCode\rodingia_3.1\cuda\gai	Warning	DPCT1019	local_mem_size in SYCL is not a complete equivalent of	Help	
C:\temp\SampleCode\rodingia_3.1\cuda\gaiC:\temp\SampleCode\rodingia_3.1\cuda\gai	Warning	DPCT1022	There is no exact match between the maxGridSize and the	Help	
C:\temp\SampleCode\rodingia_3.1\cuda\gaiC:\temp\SampleCode\rodingia_3.1\cuda\gai	Warning	DPCT1005	The device version is different. You need to rewrite	Help	
C:\temp\SampleCode\rodingia_3.1\cuda\gaiC:\temp\SampleCode\rodingia_3.1\cuda\gai	Warning	DPCT1009	SYCL uses exceptions to report errors and does not use	Help	
C:\temp\SampleCode\rodingia_3.1\cuda\gaiC:\temp\SampleCode\rodingia_3.1\cuda\gai	Warning	DPCT1010	SYCL uses exceptions to report errors and does not use	Help	
C:\temp\SampleCode\rodingia_3.1\cuda\gaiC:\temp\SampleCode\rodingia_3.1\cuda\gai	Warning	DPCT1010	SYCL uses exceptions to report errors and does not use the error	Help	

Summary

- OneAPI delivers a unified programming model to simplify development across diverse architectures
- Intel DPC++ Compatibility tool assists developers in migrating code written in CUDA to DPC++, increasing developer productivity
- DPC++ is an open specification for a portable, architecture-neutral language for expressing parallelism; it is based on industry standards

References

- [Intel® DPC++ Compatibility Tool Jupyter Tutorial](#)
- [Intel® DPC++ Compatibility Tool](#)
 - [User Guide](#)
 - [Get Started Guide](#)
 - [Release Notes](#)

Are You Ready to Try oneAPI?

1. Identify potential workloads/candidates for testing
 - a. Download DPCT and migrate code to DPC++ on-prem, if applicable
 - b. Test, tune and optimize your code or test samples in the [Intel® DevCloud](#)—a cloud-based development sandbox environment that gives you full access to the latest Intel® hardware and oneAPI software
<https://software.intel.com/devcloud/oneapi>
2. Learn more at <http://software.intel.com/oneapi> the channel to documentation, downloads, access to Intel® Devcloud, and access to support forum

Notices & Disclaimers

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at [intel.com](https://www.intel.com), or from the OEM or retailer.

The benchmark results reported herein may need to be revised as additional testing is conducted. The results depend on the specific platform configurations and workloads utilized in the testing, and may not be applicable to any particular user's components, computer system or workloads. The results are not necessarily representative of other benchmarks and other benchmark results may show greater or lesser impact from mitigations.

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 complete information visit www.intel.com/benchmarks.

INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS". NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO THIS INFORMATION INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

Copyright © 2020, Intel Corporation. All rights reserved. Intel, the Intel logo, Xeon, Core, VTune, and OpenVINO are trademarks of Intel Corporation or its subsidiaries in the U.S. and other countries. Khronos® is a registered trademark and SYCL is a trademark of the Khronos Group, Inc.

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

The Intel logo is centered on a solid blue background. It features the word "intel" in a white, lowercase, sans-serif font. A small blue square is positioned above the letter 'i'. To the right of the word "intel" is a registered trademark symbol (®).

intel®