

oneAPI DevSummit at ISC 21 | June 22, 2021

DESIGN, DEVELOPMENT, AND VALIDATION OF A DPC++ BACKEND FOR OCCA

ANOOP MADHUSOODHANAN PRABHA
CEDRIC ANDREOLLI
PHILLIPE THIERRY
Intel Corporation

SAUMIL SUDHIR PATEL
KRIS ROWE

Argonne National Laboratory

SHOUT-OUTS

- David Medina (OCCA lead developer/maintainer)
- Tim Warburton (Virginia Tech, CEED)
- Michael D'mello (Intel Center of Excellence at ALCF)

This research was supported by the Exascale Computing Project (17-SC-20-SC), a joint project of the U.S. Department of Energy's Office of Science and National Nuclear Security Administration, responsible for delivering a capable exascale ecosystem, including software, applications, and hardware technology, to support the nation's exascale computing imperative.



U.S. DEPARTMENT OF
ENERGY

Argonne National Laboratory is a
U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC.

intel[®]

Argonne
NATIONAL LABORATORY

75
1946-2021

THE BIG PICTURE

What is
OCCA?

Why
DPC++?

OCCA
API

OKL
Translation

Validation

What's
Next?



U.S. DEPARTMENT OF
ENERGY

Argonne National Laboratory is a
U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC.

intel[®]

Argonne
NATIONAL LABORATORY

75
1946-2021

THE BIG PICTURE

What is
OCCA?

Why
DPC++?

OCCA
API

OKL
Translation

Validation

What's
Next?



U.S. DEPARTMENT OF
ENERGY

Argonne National Laboratory is a
U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC.

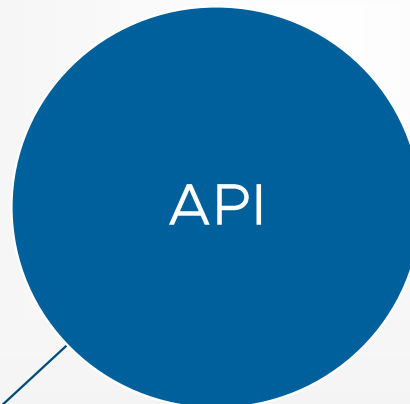
intel[®]

Argonne
NATIONAL LABORATORY

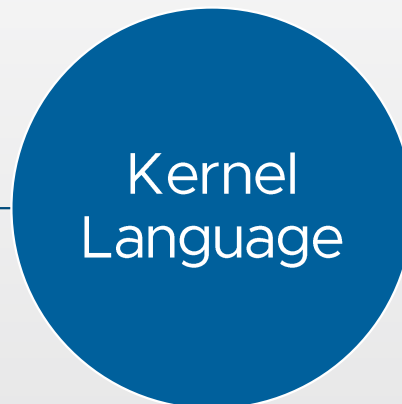
75
1946-2021

WHAT IS OCCA?

Open
Concurrent
Compute
Architecture



- Unified models for device, memory, etc.
- Backend selection at runtime:
 - Serial, OpenMP
 - CUDA, HIP, OpenCL, Metal
- Lightweight wrappers around backend APIs
- C/C++, Fortran



- Directive based extension to C/C++
- Transparent translation to backend code
- JIT compilation + caching
- Alternatively, write kernels directly with backend specific code

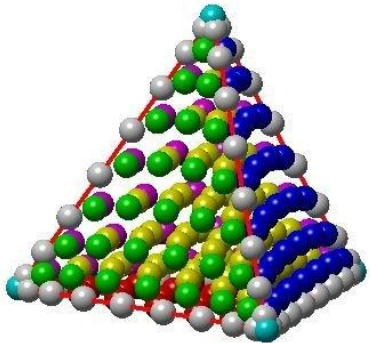

<https://github.com/libocca/occa>



- Hardware info
- Available backends
- Environment variables
- Translate/compile kernels

WHO USES OCCA?

Major applications in the public and private sectors

			
<p>Parallel Numerical Algorithms Group</p> <ul style="list-style-type: none">• libParanumal	<p>Center for Exascale Discretizations</p> <ul style="list-style-type: none">• NekRS• MFEM• Laghos	<p>Shell</p> <ul style="list-style-type: none">• Full wave propagation library	<p>Naval Postgraduate School</p> <ul style="list-style-type: none">• NUMA



U.S. DEPARTMENT OF
ENERGY

Argonne National Laboratory is a
U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC.

intel®

Argonne
NATIONAL LABORATORY



75

1946-2021

WHERE IS OCCA USED?

Laptops to TOP500 supercomputers

First full-core, pin-resolved CFD simulation of an SMR using **NekRS** on **Summit at OLCF** (#2 on TOP500)

Nuclear Engineering and Design 378 (2021) 111143

Contents lists available at [ScienceDirect](#)




Nuclear Engineering and Design


journal homepage: www.elsevier.com/locate/nucengdes

Feasibility of full-core pin resolved CFD simulations of small modular reactor with momentum sources

Jun Fang^{a,*}, Dillon R. Shaver^a, Ananias Tomboulides^{a,b}, Misun Min^a, Paul Fischer^a, Yu-Hsiang Lan^a, Ronald Rahaman^a, Paul Romano^a, Sofiane Benhamadouche^c, Yassin A. Hassan^d, Adam Kraus^e, Elia Merzari^e



Future Exascale Systems



USING OCCA

Example application

Application Code
(Host Only)

```
1  #include <occa.hpp>
2  int main()
3  {
4      int N = 256;
5      std::vector<float> a_host(N, 1.0);
6      std::vector<float> b_host(N, 1.0);
7      std::vector<float> c_host(N);
8
9      occa::device device("{mode: 'dpcpp', device_id: 0, platform_id: 0}");
10
11     occa::memory a_device = device.malloc<float>(N);
12     occa::memory b_device = device.malloc<float>(N);
13     occa::memory c_device = device.malloc(N * sizeof(float));
14
15     a_device.copyFrom(a_host.data());
16     b_device.copyFrom(b_host.data());
17
18     occa::kernel addVectors = device.buildKernel("addVectors.okl", "addVectors");
19
20     addVectors(N, a_device, b_device, c_device);
21
22     c_device.copyTo(c_host.data());
23     device.finish();
24
25     return 0;
26 }
```

Device selection

Allocate device memory

Host-to-device transfer

Create kernel

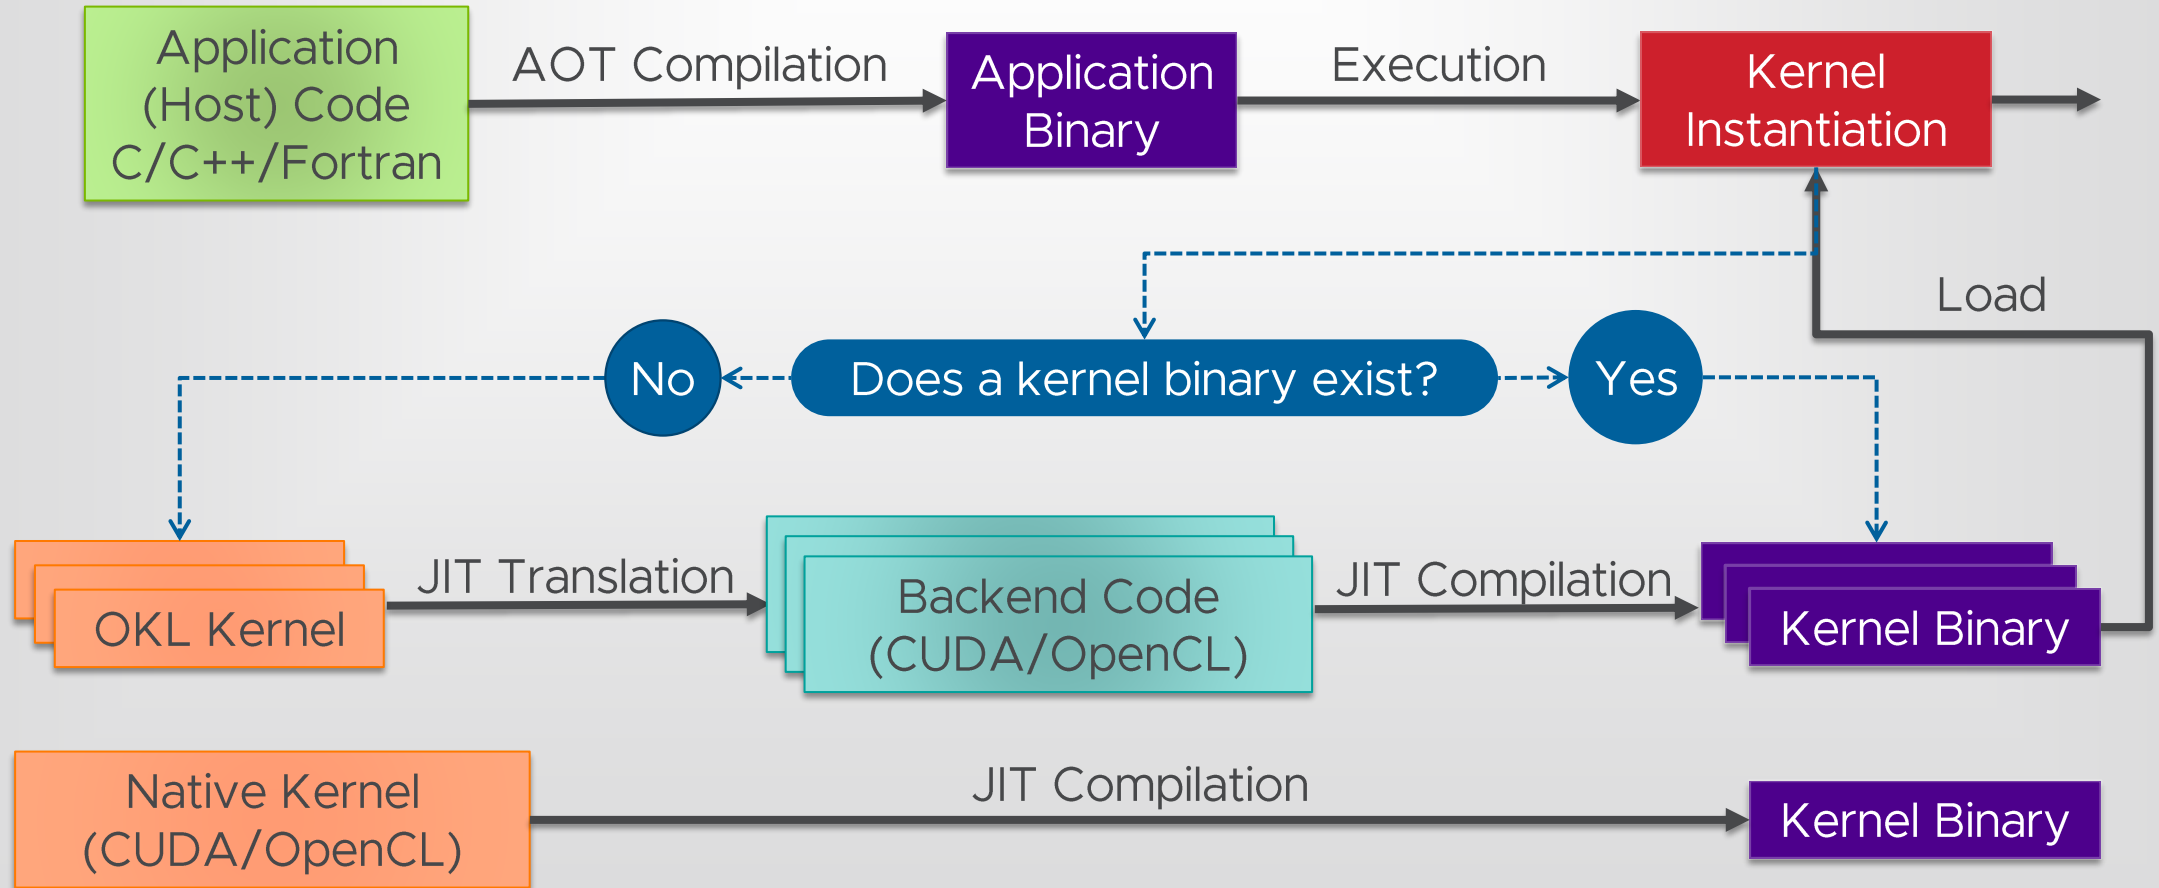
Launch kernel

Device-to-host transfer



USING OCCA

Running an application



THE BIG PICTURE

What is
OCCA?

Why
DPC++?

OCCA
API

OKL
Translation

Validation

What's
Next?



U.S. DEPARTMENT OF
ENERGY

Argonne National Laboratory is a
U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC.

intel[®]

Argonne
NATIONAL LABORATORY



75
1946-2021

BETTER TOGETHER

A serendipitous collaboration

ALCF



intel®



Shell



OCCA
DPC++
Backend

- NekRS already uses OCCA
- Needed a backend for Aurora
- After exploring options decided on DPC++

- Shell has used OCCA from its beginning
- Needed C/C++ and Fortran support
- Expressed interested in a DPC++ backend



U.S. DEPARTMENT OF
ENERGY

Argonne National Laboratory is a
U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC.

intel®

Argonne
NATIONAL LABORATORY



75
1946-2021

Why not use DPC++ directly?

Applications already use
OCCA extensively

We support but do not
“own” these applications



Implementing a DPC++ backend for OCCA was the
optimal path given the constraints



U.S. DEPARTMENT OF
ENERGY

Argonne National Laboratory is a
U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC.

intel[®]

Argonne
NATIONAL LABORATORY

75
1946–2021

DEVELOPMENT APPROACH

OCCA API

- Map OCCA platform, execution, and memory models into DPC++
- Implement core OCCA runtime classes
- Verify using native DPC++ kernels

OKL


- Map OCCA programming model into DPC++
- Implement OKL to DPC++ source translator
- Verify translation using OCCA command line tool

Validate

- Microbenchmark kernels
- Mini-apps
- Full CSE applications



THE BIG PICTURE



U.S. DEPARTMENT OF
ENERGY

Argonne National Laboratory is a
U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC.

intel[®]

Argonne
NATIONAL LABORATORY



75
1946-2021

OCCA API

Architecture overview

- Frontend provides unified models for
 - Device
 - Memory
 - Kernel
 - Stream
 - StreamTag
- OCCA uses a PIMPL design pattern
- Backends (called modes) extend base classes, implement virtual functions



U.S. DEPARTMENT OF
ENERGY

Argonne National Laboratory is a
U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC.

intel®

Argonne
NATIONAL LABORATORY

75
1946–2021

OCCA API

Backend model

*Class names used here differ slightly from actual names in the OCCA source code

device

- malloc
- buildKernel
- createStream
- tagStream
- waitFor
- finish

memory

- copyTo
- copyFrom
- addOffset

kernel

- setRunDims
- pushArg
- run
- operator()

stream

streamTag



U.S. DEPARTMENT OF
ENERGY

Argonne National Laboratory is a
U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC.

intel[®]

Argonne
NATIONAL LABORATORY



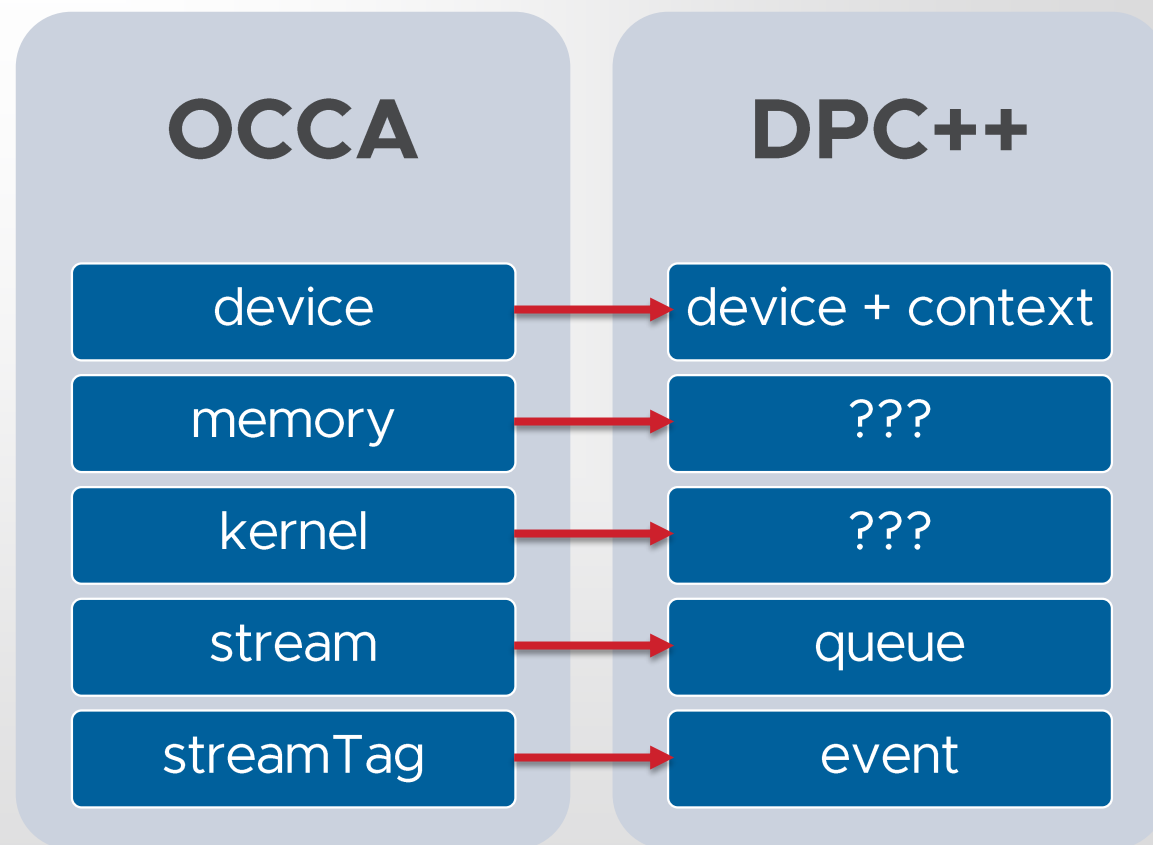
75

1946-2021

OCCA API

Mapping OCCA models to DPC++

- There is no explicit abstraction for a
 - **Platform**
 - Equivalent to an occa mode
 - Defined via device selection
 - **Context**
 - Assume one device per context



U.S. DEPARTMENT OF
ENERGY

Argonne National Laboratory is a
U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC.

intel[®]

Argonne
NATIONAL LABORATORY

75
1946-2021

OCCA API

DPC++ backend: memory

OBSERVATIONS

1. Device memory management is **explicit** via `malloc/free` functions
2. Host/device memory transfers are **explicit** via `copyTo/copyFrom`
3. Other backends (CUDA, OpenCL) pass memory to kernels via pointers
4. Extensive changes to OCCA's internal API would be required to use DPC++ buffers/accessors

DESIGN CHOICES

- Use DPC++ USM functions for device memory management
- Use **explicit USM** (`malloc_device`) by default for performance reasons
- More flavours of memory via extra argument to `occa::malloc`
 - *Host*: use `malloc_host`
 - *Shared*: use `malloc_shared`



OCCA API

DPC++ backend: kernels

OBSERVATIONS

1. Virtual functions are not supported in DPC++ kernels
2. We would like code in a **single translation unit**
3. OCCA supports OKL and native (backend specific code) kernels
4. Kernel launch bounds
 - Are extracted from OKL kernels
 - Must be set for native kernels

DESIGN CHOICES

- Require native kernels to be defined as **extern “C”** functions with
 - A **queue** as the first argument
 - An **nd_range** as the second
- Translate OKL kernels in the same way, adding a **queue** and **nd_range** to the argument list
- Run kernels via **sys::runFunction**

OCCA API

DPC++ backend

device

- malloc
- buildKernel
- createStream
- tagStream
- waitFor
- finish

- sycl::device
- sycl::context

memory

- copyTo
- copyFrom
- addOffset
- void***

kernel

- setRunDims
- pushArg
- run
- operator()
- void (*f)(...)**

stream


- sycl::queue**

streamTag

- sycl::event**



THE BIG PICTURE



U.S. DEPARTMENT OF
ENERGY

Argonne National Laboratory is a
U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC.

intel[®]

Argonne
NATIONAL LABORATORY

75
1946-2021

OKL

Programming Model

- **OCCA Kernel Language (OKL)** can be used to write portable kernels
- OKL **attributes** are keywords prefixed by the “@” symbol
- The OKL spec defines attributes for
 - Loop parallelism
 - Memory spaces
 - Synchronization
 - Atomics
- Defined constants can be passed at runtime before JIT compilation

```
1  @kernel void matrixBlockMultiply(  
2      @ restrict const double *A @dim(M, K),  
3      @ restrict const double *B @dim(K, N),  
4      @ restrict double *C @dim(M, N),  
5      const int M,  
6      const int N,  
7      const int K)  
8  {  
9      for (int nb = 0; nb < N; nb += BLOCK_N; @outer)  
10     {  
11         for (int mb = 0; mb < M; mb += BLOCK_M; @outer)  
12         {  
13             for (int nt = 0; nt < BLOCK_N; ++nt; @inner)  
14             {  
15                 for (int mt = 0; mt < BLOCK_M; ++mt; @inner)  
16                 {  
17                     double C_ij = 0.0;  
18                     for (int k = 0; k < K; ++k)  
19                     {  
20                         C_mn += A(mb + mt, k) * B(k, nb + nt);  
21                     }  
22                     C(mb + nt, nb + nt) = C_mn;  
23                 }  
24             }  
25         }  
26     }  
27 }
```



OCCA Kernel (OKL)

```
1 @kernel void matrixBlockMultiply(  
2     @ restrict const double *A @dim(M, K),  
3     @ restrict const double *B @dim(K, N),  
4     @ restrict double *C @dim(M, N),  
5     const int M,  
6     const int N,  
7     const int K)  
8 {  
9     for (int nb = 0; nb < N; nb += BLOCK_N; @outer)  
10    {  
11        for (int mb = 0; mb < M; mb += BLOCK_M; @outer)  
12        {  
13            for (int nt = 0; nt < BLOCK_N; ++nt; @inner)  
14            {  
15                for (int mt = 0; mt < BLOCK_M; ++mt; @inner)  
16                {  
17                    double C_ij = 0.0;  
18                    for (int k = 0; k < K; ++k)  
19                    {  
20                        C_mn += A(mb + mt, k) * B(k, nb + nt);  
21                    }  
22                    C(mb + nt, nb + nt) = C_mn;  
23                }  
24            }  
25        }  
26    }  
27 }
```

JIT Translation

CUDA

```
1 extern "C" __global__ void _occa_matrixBlockMultiply_0(const double *__restrict__ A,  
2                                                         const double *__restrict__ B,  
3                                                         double *__restrict__ C,  
4                                                         const int M,  
5                                                         const int N,  
6                                                         const int K)  
7 {  
8     {  
9         int nb = 0 + (8 * blockIdx.y);  
10    {  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27 }
```

OpenCL

```
1 #pragma OPENCL EXTENSION cl_khr_fp64 : enable  
2  
3 __kernel void _occa_matrixBlockMultiply_0(__global const double * restrict A,  
4 __global const double * restrict B,  
5 __global double * restrict C,  
6 const int M,  
7 const int N,  
8 const int K) {  
9     {  
10        int nb = 0 + (8 * get_group_id(1));  
11        {  
12            int mb = 0 + (32 * get_group_id(0));  
13            {  
14                int nt = 0 + get_local_id(1);  
15                {  
16                    int mt = 0 + get_local_id(0);  
17                    double C_ij = 0.0;  
18                    for (int k = 0; k < K; ++k) {  
19                        C_mn += A[mb + mt + (M * k)] * B[k + (K * (nb + nt))];  
20                    }  
21                    C[mb + nt + (M * (nb + nt))] = C_mn;  
22                }  
23            }  
24        }  
25    }  
26 }
```



OKL

DPC++ backend

- OCCA has extensive infrastructure for parsing/transforming source code
- An abstract class interface is provided for translating OKL to backend code
- Developer prescribes how to
 - Translate kernel argument and return types
 - Map attributes
- Recall: we required OCCA kernel functions to take a `queue` and `nd_range` as arguments
- Using lambdas for the DPC++ command group and kernel functions is straightforward
- Define the `nd_range`:
 - Local range = `@inner` dimensions
 - Global range = `@inner` x `@outer`



```

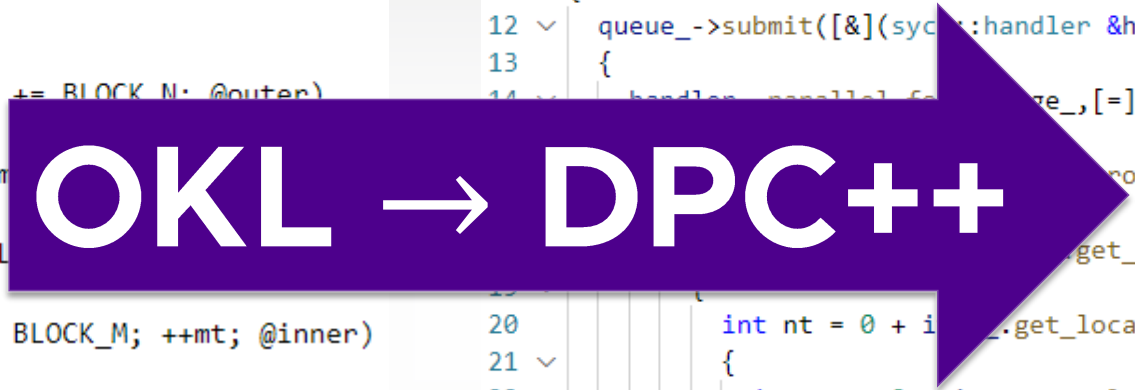
1 @kernel void matrixBlockMultiply(
2     @ restrict const double *A @dim(M, K),
3     @ restrict const double *B @dim(K, N),
4     @ restrict double *C @dim(M, N),
5     const int M,
6     const int N,
7     const int K)
8 {
9     for (int nb = 0; nb < N; nb += BLOCK_N; @outer)
10    {
11        for (int mb = 0; mb < M; mb += BLOCK_M; @middle)
12        {
13            for (int nt = 0; nt < BLOCK_N; nt += BLOCK_N; @inner)
14            {
15                for (int mt = 0; mt < BLOCK_M; ++mt; @inner)
16                {
17                    double C_ij = 0.0;
18                    for (int k = 0; k < K; ++k)
19                    {
20                        C_mn += A(mb + mt, k) * B(k, nb + nt);
21                    }
22                    C(mb + nt, nb + nt) = C_mn;
23                }
24            }
25        }
26    }
27 }

```

```

1 #include <CL/sycl.hpp>
2
3 extern "C" void _occa_matrixBlockMultiply_0(sycl::queue *queue_,
4                                             sycl::nd_range<3> *range_,
5                                             const double *__restrict__ A,
6                                             const double *__restrict__ B,
7                                             double *__restrict__ C,
8                                             const int &M,
9                                             const int &N,
10                                            const int &K)
11 {
12     queue_>submit([&](sycl::handler &handler_)
13     {
14         handler.parallel_for(range_, [=](sycl::nd_item<3> item_)
15         {
16             int nt = 0 + item_.get_local_id(1);
17             {
18                 int mt = 0 + item_.get_local_id(2);
19                 double C_ij = 0.0;
20                 for (int k = 0; k < K; ++k)
21                 {
22                     C_mn += A[mb + mt + (M * k)] * B[k + (K * (nb + nt))];
23                 }
24                 C[mb + mt + (M * (nb + nt))] = C_mn;
25             }
26         });
27     });
28 }

```




```

1 @kernel void matrixBlockMultiply(
2     @ restrict const double *A @dim(M, K),
3     @ restrict const double *B @dim(K, N),
4     @ restrict double *C @dim(M, N),
5     const int M,
6     const int N,
7     const int K)
8 {
9     for (int nb = 0; nb < N; nb += BLOCK_N; @outer)
10    {
11        for (int mb = 0; mb < M; mb += BLOCK_M; @outer)
12        {
13            for (int nt = 0; nt < BLOCK_N; ++nt; @inner)
14            {
15                for (int mt = 0; mt < BLOCK_M; ++mt; @inner)
16                {
17                    double C_ij = 0.0;
18                    for (int k = 0; k < K; ++k)
19                    {
20                        C_mn += A(mb + mt, k) * B(k, nb + nt);
21                    }
22                    C(mb + nt, nb + nt) = C_mn;
23                }
24            }
25        }
26    }
27 }

```

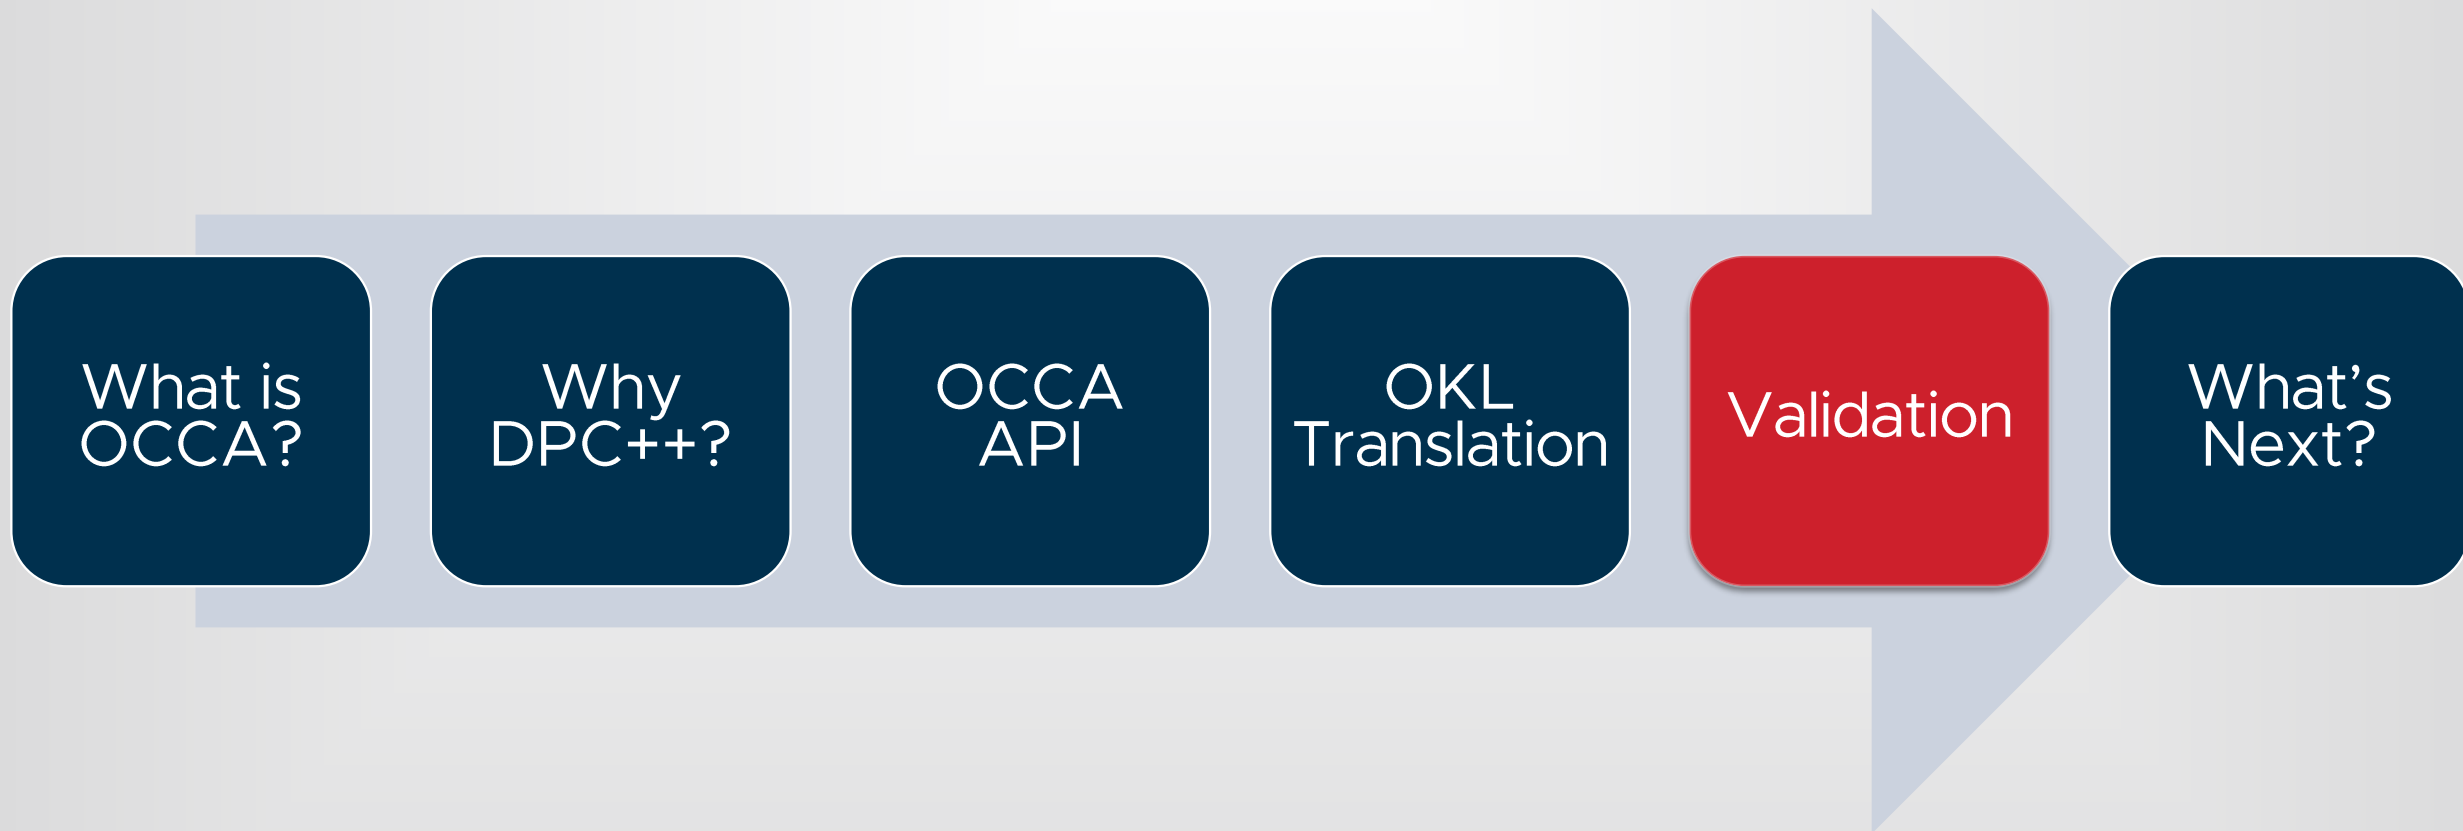
```

1 #include <CL/sycl.hpp>
2
3 extern "C" void _occa_matrixBlockMultiply_0(sycl::queue *queue_,
4                                             sycl::nd_range<3> *range_,
5                                             const double *__restrict__ A,
6                                             const double *__restrict__ B,
7                                             double *__restrict__ C,
8                                             const int &M,
9                                             const int &N,
10                                            const int &K)
11 {
12     queue_>submit([&](sycl::handler &handler_)
13     {
14         handler_.parallel_for(*range_, [=](sycl::nd_item<3> item_)
15         {
16             int nb = 0 + (8 * item_.get_group(1));
17             {
18                 int mb = 0 + (32 * item_.get_group(2));
19                 {
20                     int nt = 0 + item_.get_local_id(1);
21                     {
22                         int mt = 0 + item_.get_local_id(2);
23                         double C_ij = 0.0;
24                         for (int k = 0; k < K; ++k)
25                         {
26                             C_mn += A[mb + mt + (M * k)] * B[k + (K * (nb + nt))];
27                         }
28                         C[mb + mt + (M * (nb + nt))] = C_mn;
29                     }
30                 }
31             }
32         });
33     });
34 }

```



THE BIG PICTURE



U.S. DEPARTMENT OF
ENERGY

Argonne National Laboratory is a
U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC.

intel[®]

Argonne
NATIONAL LABORATORY



75
1946-2021

VALIDATION

OCCA API + OKL

- OCCA's test harness
- Benchmark kernels
 - stream
 - reduction
 - matrix transpose
 - matrix multiply
- Mini-apps: NekBench, ISO3DFD
- NekRS: example cases (e.g., “turbPipe”)
- Initial performance is comparable to
 - Using DPC++ directly
 - The OCCA OpenCL backend

All tests were performed on Intel GPUs using the oneAPI Base Toolkit



U.S. DEPARTMENT OF
ENERGY

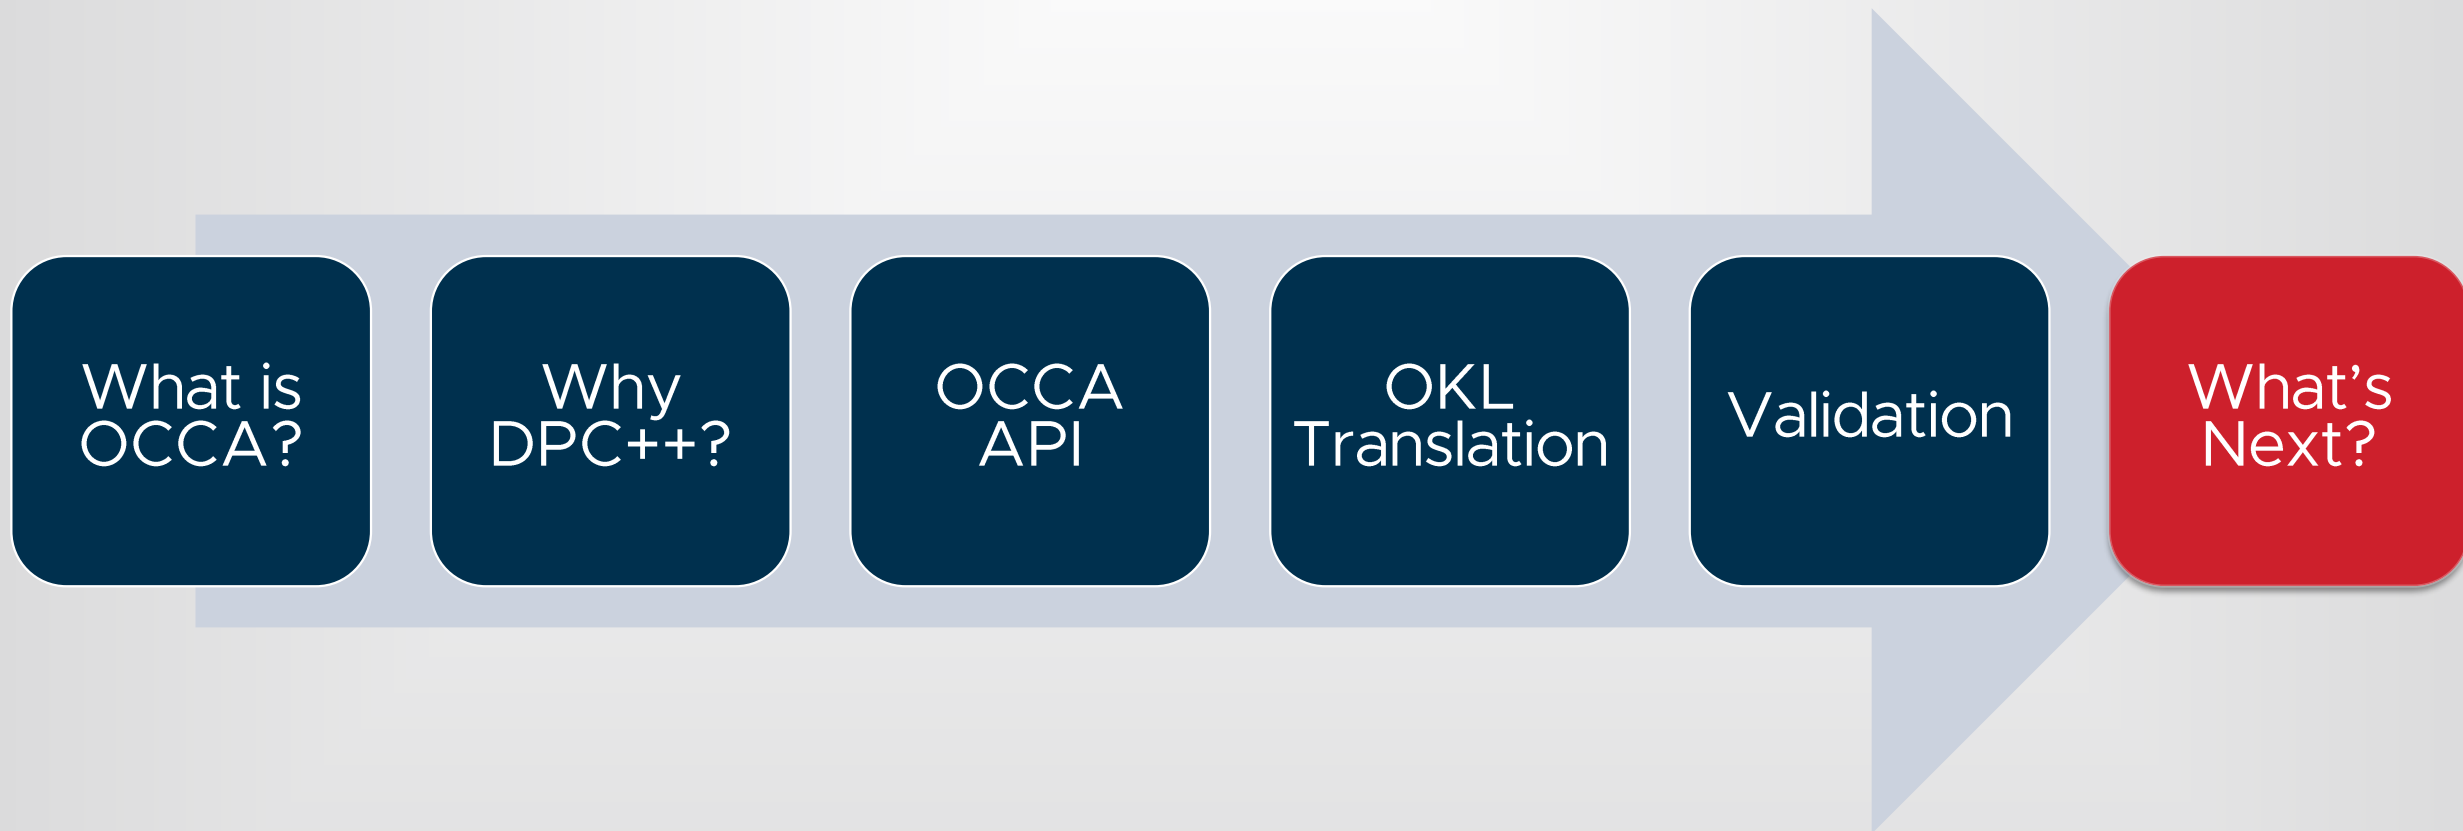
Argonne National Laboratory is a
U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC.

intel[®]

Argonne
NATIONAL LABORATORY

75
1946–2021

THE BIG PICTURE



U.S. DEPARTMENT OF
ENERGY

Argonne National Laboratory is a
U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC.

intel[®]

Argonne
NATIONAL LABORATORY



75
1946-2021

NEXT STEPS

Current Efforts

- OCCA DPC++ backend available on GitHub in near future
- Technical report on development of DPC++ backend
- OCCA performance benchmarks
- **Performance study**—across vendor hardware, all backends
 - *What is the price paid for using a portability framework?*

Future Work

- Propose API for performance features
 - subdevices
 - thread-block/group collectives
 - asynch mem copies (i.e., global ↔ shared/local)

To learn more about OCCA visit <https://libocca.org>

For source code and ways to get involved checkout <https://github.com/libocca/occa>



U.S. DEPARTMENT OF
ENERGY

Argonne National Laboratory is a
U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC.

intel®

Argonne 
NATIONAL LABORATORY

75
1946-2021

TAKEAWAYS

OCCA is used by mission critical applications in the public and private sector

A new OCCA DPC++ backend was the shortest path to running on Intel GPUs

The DPC++ USM model was central to in designing the new OCCA backend

Local memory use in kernels required workarounds: no free function is available currently

OCCA's test suite, benchmark kernels, and NekRS were used for validation

Understanding and optimizing performance is the next focus



U.S. DEPARTMENT OF
ENERGY

Argonne National Laboratory is a
U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC.

intel®

Argonne
NATIONAL LABORATORY

75
1946-2021



Argonne National Laboratory is a U.S. Department of Energy laboratory managed by UChicago Argonne, LLC.



EXASCALE
COMPUTING
PROJECT



CEED
EXASCALE DISCRETIZATIONS



This research was supported by the Exascale Computing Project (17-SC-20-SC), a joint project of the U.S. Department of Energy's Office of Science and National Nuclear Security Administration, responsible for delivering a capable exascale ecosystem, including software, applications, and hardware technology, to support the nation's exascale computing imperative.