

Hybrid Programming

Hybrid Programming

- Hybrid programming here stands for the interaction of OpenMP with a lower-level programming model, e.g.
 - OpenCL
 - CUDA
 - HIP

Hybrid Programming

- Hybrid programming here stands for the interaction of OpenMP with a lower-level programming model, e.g.
 - OpenCL
 - CUDA
 - HIP
- OpenMP supports these interactions
 - Calling low-level kernels from OpenMP application code
 - Calling OpenMP kernels from low-level application code

Example: Calling saxpy

```
void example(){
    float a = 2.0;
    float * x;
    float * y;

    // allocate the device memory
    #pragma omp target data map(to:x[0:count]) map(tofrom:y[0:count])
    {
        compute__1(n, x);
        compute__2(n, y);
            saxpy(n, a, x, y)
        compute__3(n, y);
    }
}
```

Example: Calling saxpy

```
void example(){
    float a = 2.0;
    float * x;
    float * y;

    // allocate the device memory
    #pragma omp target data map(to:x[0:count]) map(from:y[0:count])
    {
        compute__1(n, x);
        compute__2(n, y);
        saxpy(n, a, x, y)
        compute__3(n, y);
    }
}
```

```
void saxpy(size_t n, float a,
           float * x, float * y){
    #pragma omp target teams distribute \
        parallel for simd
    for (size_t i = 0; i < n; ++i){
        y[i] = a * x[i] + y[i];
    }
}
```

Example: Calling saxpy

```
void example(){
    float a = 2.0;
    float * x;
    float * y;

    // allocate the device memory
    #pragma omp target data map(to:x[0:count]) map(from:y[0:count])
    {
        compute__1(n, x);
        compute__2(n, y);
        saxpy(n, a, x, y)
        compute__3(n, y);
    }
}
```

Let's assume that we want to implement the saxpy() function in a low-level language.

```
void saxpy(size_t n, float a,
           float * x, float * y){
    #pragma omp target teams distribute \
        parallel for simd
    for (size_t i = 0; i < n; ++i){
        y[i] = a * x[i] + y[i];
    }
}
```

HIP Kernel for saxpy()

```
___global___ void saxpy__kernel(size__t n, float a, float * x, float * y){
    size__t i = threadIdx.x + blockIdx.x * blockDim.x;
    y[i] = a * x[i] + y[i];
}

void saxpy__hip(size__t n, float a, float * x, float * y){
    assert(n % 256 == 0);
    saxpy__kernel<<<n/256,256,0,NULL>>>(n, a, x, y);
}
```

HIP Kernel for saxpy()

```
___global___ void saxpy__kernel(size__t n, float a, float * x, float * y){  
    size__t i = threadIdx.x + blockIdx.x * blockDim.x;  
    y[i] = a * x[i] + y[i];  
}
```

```
void saxpy__hip(size__t n, float a, float * x, float * y){  
    assert(n % 256 == 0);  
    saxpy__kernel<<<n/256,256,0,NULL>>>(n, a, x, y);  
}
```



These are device pointers!

HIP Kernel for saxpy()

- Assume a HIP version of the SAXPY kernel:

```
___global___ void saxpy__kernel(size__t n, float a, float * x, float * y){
    size__t i = threadIdx.x + blockIdx.x * blockDim.x;
    y[i] = a * x[i] + y[i];
}


void saxpy__hip(size__t n, float a, float * x, float * y){
    assert(n % 256 == 0);
    saxpy__kernel<<<n/256,256,0,NULL>>>(n, a, x, y);
}
```

These are device pointers!

- We need a way to translate the host pointer that was mapped by OpenMP directives and retrieve the associated device pointer.

Pointer Translation /1

- When creating the device data environment, OpenMP creates a mapping between
 - the (virtual) memory pointer on the host and
 - the (virtual) memory pointer on the target device.



Pointer Translation /1

- When creating the device data environment, OpenMP creates a mapping between
 - the (virtual) memory pointer on the host and
 - the (virtual) memory pointer on the target device.
- This mapping is established through the data-mapping directives and their clauses.

Host memory


x: 0xabcd

```
01010101011010
01111010110101
00010101010101
01010101010201
01011010000100
10101010101010
00110011100110
```

```
#pragma omp target data \
    map(to:x[0:n])
...
!$omp end target data
```


Pointer Translation /1

- When creating the device data environment, OpenMP creates a mapping between
 - the (virtual) memory pointer on the host and
 - the (virtual) memory pointer on the target device.
- This mapping is established through the data-mapping directives and their clauses.




```

#pragma omp target data \
    map(to:x[0:n])
...
!$omp end target data
  
```



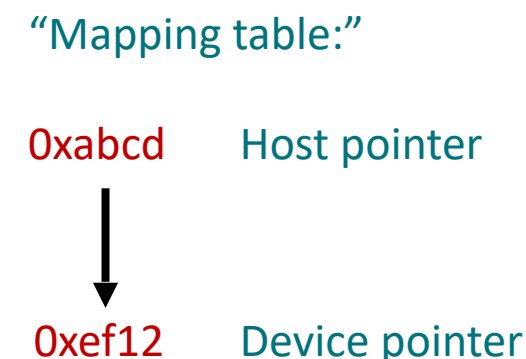
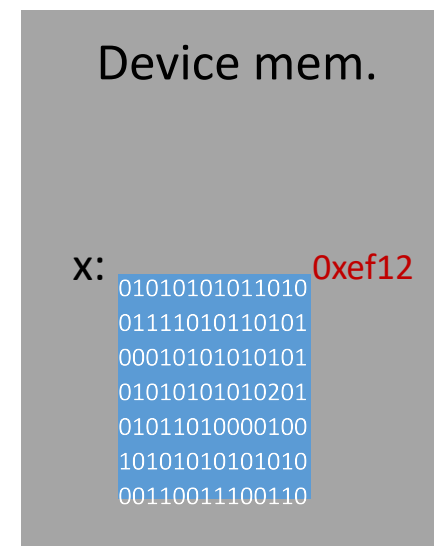
Pointer Translation /1

- When creating the device data environment, OpenMP creates a mapping between
 - the (virtual) memory pointer on the host and
 - the (virtual) memory pointer on the target device.
- This mapping is established through the data-mapping directives and their clauses.



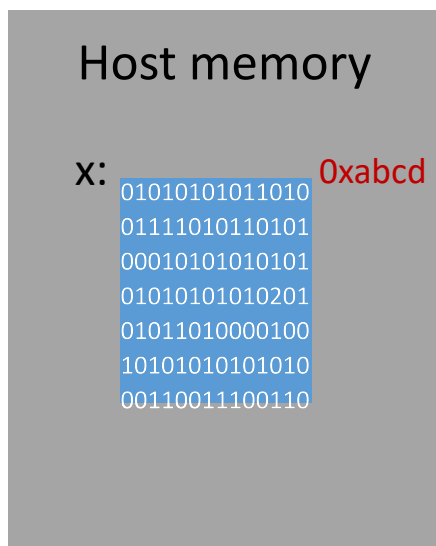
```

#pragma omp target data \
    map(to:x[0:n])
...
!$omp end target data
    
```



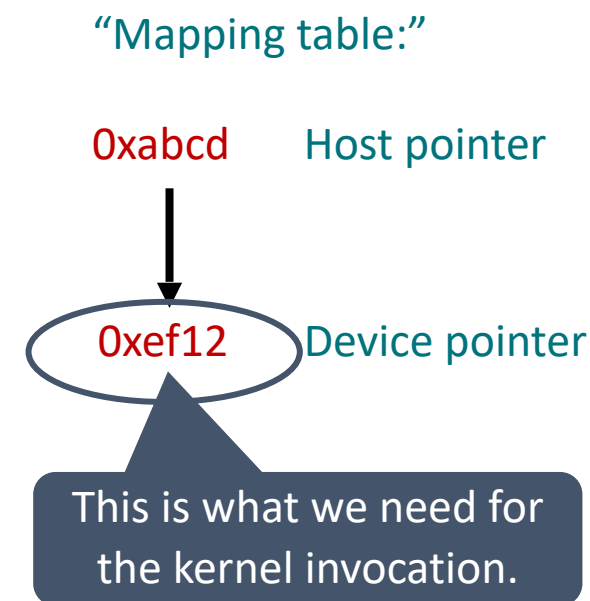
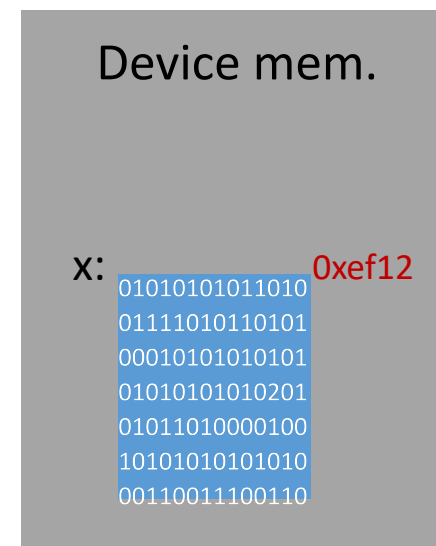
Pointer Translation /1

- When creating the device data environment, OpenMP creates a mapping between
 - the (virtual) memory pointer on the host and
 - the (virtual) memory pointer on the target device.
- This mapping is established through the data-mapping directives and their clauses.



```

#pragma omp target data \
    map(to:x[0:n])
...
!$omp end target data
    
```



Pointer Translation /2

- The target data construct defines the `use__device__ptr` clause to perform pointer translation.
 - The OpenMP implementation searches for the host pointer in its internal mapping tables.
 - The associated device pointer is then returned.

```
type * x = 0xabcd;  
#pragma omp target data use__device__ptr(x)  
{  
    example__func(x); // x == 0xef12  
}
```

- Note: the pointer variable shadowed within the target data construct for the translation.

Putting it Together...

```
void example(){
    float a = 2.0;
    float * x = ...; // assume: x = 0xabcd
    float * y = ...;

    // allocate the device memory
    #pragma omp target data map(to:x[0:count]) map(tofrom:y[0:count])
    {
        compute__1(n, x); // mapping table: x:[0xabcd,0xef12], x = 0xabcd
        compute__2(n, y);
        #pragma omp target data use__device__ptr(x,y)
        {
            saxpy__hip(n, a, x, y) // mapping table: x:[0xabcd,0xef12], x = 0xef12
        }
        compute__3(n, y);
    }
}
```


Advanced Task Synchronization

Asynchronous API Interaction

- Some APIs are based on asynchronous operations
 - MPI asynchronous send and receive
 - Asynchronous I/O
 - HIP, CUDA and OpenCL stream-based offloading
 - In general: any other API/model that executes asynchronously with OpenMP (tasks)

Asynchronous API Interaction

- Some APIs are based on asynchronous operations
 - MPI asynchronous send and receive
 - Asynchronous I/O
 - HIP, CUDA and OpenCL stream-based offloading
 - In general: any other API/model that executes asynchronously with OpenMP (tasks)
- Example: HIP memory transfers

```
do_something();  
hipMemcpyAsync(dst, src, nbytes, hipMemcpyDeviceToHost, stream);  
do_something_else();  
hipStreamSynchronize(stream);  
do_other_important_stuff(dst);
```

Asynchronous API Interaction

- Some APIs are based on asynchronous operations
 - MPI asynchronous send and receive
 - Asynchronous I/O
 - HIP, CUDA and OpenCL stream-based offloading
 - In general: any other API/model that executes asynchronously with OpenMP (tasks)
- Example: HIP memory transfers

```
do_something();  
hipMemcpyAsync(dst, src, nbytes, hipMemcpyDeviceToHost, stream);  
do_something__else();  
hipStreamSynchronize(stream);  
do_other__important__stuff(dst);
```


- Programmers need a mechanism to marry asynchronous APIs with the parallel task model of OpenMP
 - How to synchronize completions events with task execution?

Try 1: Use just OpenMP Tasks

```
void hip__example(){
#pragma omp task //task A
{
    do__something();
    hipMemcpyAsync(dst, src, nbytes, hipMemcpyDeviceToHost, stream);
}
#pragma omp task //task B
{
    do__something__else();
}
#pragma omp task //task C
{
    hipStreamSynchronize(stream);
    do__other__important__stuff(dst);
}
}
```

Try 1: Use just OpenMP Tasks


```
void hip__example(){  
#pragma omp task //task A  
{  
do__something();  
hipMemcpyAsync(dst, src, nbytes, hipMemcpyDeviceToHost, stream);  
}  
#pragma omp task //task B  
{  
do__something__else();  
}  
#pragma omp task //task C  
{  
hipStreamSynchronize(stream);  
do__other__important__stuff(dst);  
}  
}
```



Race condition between the tasks A & C,
task C may start execution before
task A enqueues memory transfer.

Try 1: Use just OpenMP Tasks

```
void hip__example(){  
#pragma omp task //task A  
{  
    do__something();  
    hipMemcpyAsync(dst, src, nbytes, hipMemcpyDeviceToHost, stream);  
}  
#pragma omp task //task B  
{  
    do__something__else();  
}  
#pragma omp task //task C  
{  
    hipStreamSynchronize(stream);  
    do__other__important__stuff(dst);  
}  
}
```



Race condition between the tasks A & C,
task C may start execution before
task A enqueues memory transfer.


- This solution does not work!

Try 2: Use just OpenMP Tasks Dependences

```
void hip__example(){
#pragma omp task depend(out:stream) //task A
{
    do__something();
    hipMemcpyAsync(dst, src, nbytes, hipMemcpyDeviceToHost, stream);
}
#pragma omp task //task B
{
    do__something__else();
}
#pragma omp task depend(in:stream) //task C
{
    hipStreamSynchronize(stream);
    do__other__important__stuff(dst);
}
}
```


Try 2: Use just OpenMP Tasks Dependences

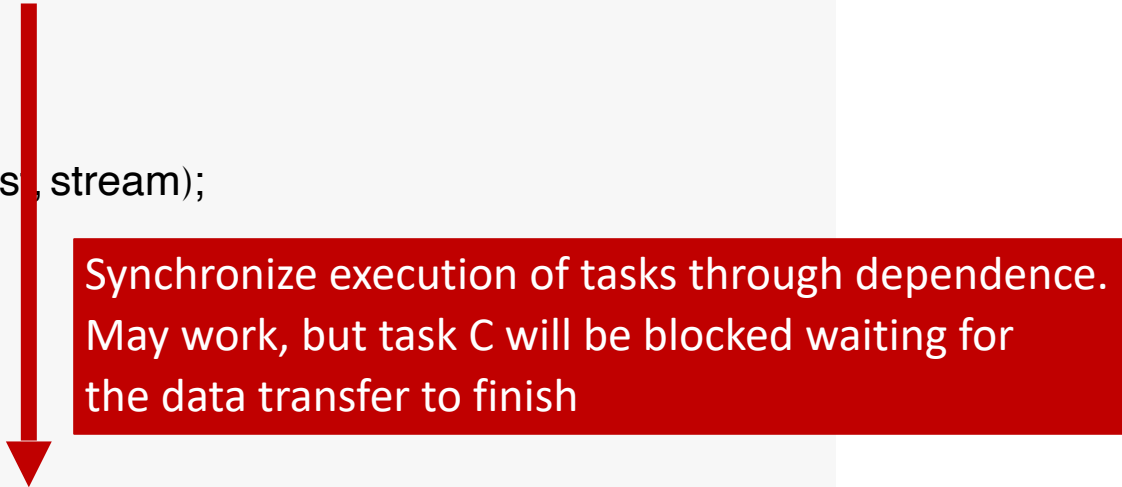
```
void hip__example(){  
#pragma omp task depend(out:stream) //task A  
{  
do__something();  
hipMemcpyAsync(dst, src, nbytes, hipMemcpyDeviceToHost, stream);  
}  
#pragma omp task //task B  
{  
do__something__else();  
}  
#pragma omp task depend(in:stream) //task C  
{  
hipStreamSynchronize(stream);  
do__other__important__stuff(dst);  
}  
}
```



Synchronize execution of tasks through dependence.
May work, but task C will be blocked waiting for
the data transfer to finish

Try 2: Use just OpenMP Tasks Dependences

```
void hip__example(){
#pragma omp task depend(out:stream) //task A
{
    do__something();
    hipMemcpyAsync(dst, src, nbytes, hipMemcpyDeviceToHost, stream);
}
#pragma omp task //task B
{
    do__something__else();
}
#pragma omp task depend(in:stream) //task C
{
    hipStreamSynchronize(stream);
    do__other__important__stuff(dst);
}
}
```



Synchronize execution of tasks through dependence.
May work, but task C will be blocked waiting for
the data transfer to finish

- This solution may work, but
 - takes a thread away from execution while the system is handling the data transfer.
 - may be problematic if called interface is not thread-safe

OpenMP Detachable Tasks

- OpenMP 5.0 introduces the concept of a detachable task
 - Task can detach from executing thread without being “completed”
 - Regular task synchronization mechanisms can be applied to await completion of a detached task
 - Runtime API to complete a task

OpenMP Detachable Tasks

- OpenMP 5.0 introduces the concept of a detachable task
 - Task can detach from executing thread without being “completed”
 - Regular task synchronization mechanisms can be applied to await completion of a detached task
 - Runtime API to complete a task
- Detached task events: `omp__event__t` datatype

OpenMP Detachable Tasks

- OpenMP 5.0 introduces the concept of a detachable task
 - Task can detach from executing thread without being “completed”
 - Regular task synchronization mechanisms can be applied to await completion of a detached task
 - Runtime API to complete a task
- Detached task events: `omp__event__t` datatype
- Detached task clause: `detach(event)`

OpenMP Detachable Tasks

- OpenMP 5.0 introduces the concept of a detachable task
 - Task can detach from executing thread without being “completed”
 - Regular task synchronization mechanisms can be applied to await completion of a detached task
 - Runtime API to complete a task
- Detached task events: `omp__event__t` datatype
- Detached task clause: `detach(event)`
- Runtime API: `void omp__fulfill__event(omp__event__t *event)`

Detaching Tasks

```
omp__event__t *event;
void detach__example(){
#pragma omp task detach(event)
{
    important__code();
}

#pragma omp taskwait
}
```

Detaching Tasks

```
omp__event__t *event;  
void detach__example(){  
#pragma omp task detach(event)  
  {  
    important__code();  
  }  
#pragma omp taskwait  
}
```

①

1. Task detaches

Detaching Tasks

```
omp__event__t *event;  
void detach__example(){  
#pragma omp task detach(event)  
  {  
    important__code();  
  }  
#pragma omp taskwait  
}
```

①

②

1. Task detaches
2. taskwait construct cannot complete

Detaching Tasks

```
omp__event__t *event;  
void detach__example(){  
#pragma omp task detach(event)  
{  
    important__code();  
}  
#pragma omp taskwait  
}
```

①

②④

Some other thread/task:

```
omp__fulfill__event(event);
```

③

1. Task detaches
2. taskwait construct cannot complete

3. Signal event for completion

Detaching Tasks

```
omp__event__t *event;  
void detach__example(){  
#pragma omp task detach(event)  
{  
    important__code();  
}  
#pragma omp taskwait  
}
```

①

②④

Some other thread/task:

```
omp__fulfill__event(event);
```

③

1. Task detaches
2. taskwait construct cannot complete
3. Signal event for completion
4. Task completes and taskwait can continue

Putting It All Together

```
void callback(hipStream__t stream, hipError__t status, void *cb__dat){
    omp__fulfill__event((omp__event__t *) cb__data);
}

void hip__example(){
    omp__event__t *hip__event;
#pragma omp task detach(hip__event) // task A
    {
        do__something();
        hipMemcpyAsync(dst, src, nbytes, hipMemcpyDeviceToHost, stream);
        hipStreamAddCallback(stream, callback, hip__event, 0);
    }
#pragma omp task // task B
    do__something__else();

#pragma omp taskwait
#pragma omp task // task C
    {
        do__other__important__stuff(dst);
    }
}
```

Putting It All Together

```
void callback(hipStream__t stream, hipError__t status, void *cb__dat){
    omp__fulfill__event((omp__event__t *) cb__data);
}

void hip__example(){
    omp__event__t *hip__event;
#pragma omp task detach(hip__event) // task A
    {
        do__something();
        hipMemcpyAsync(dst, src, nbytes, hipMemcpyDeviceToHost, stream);
        hipStreamAddCallback(stream, callback, hip__event, 0);
    }
#pragma omp task // task B
    do__something__else();

#pragma omp taskwait
#pragma omp task // task C
    {
        do__other__important__stuff(dst);
    }
}
```

1. Task A detaches

Putting It All Together

```

void callback(hipStream__t stream, hipError__t status, void *cb__dat){
    omp__fulfill__event((omp__event__t *) cb__data);
}

void hip__example(){
    omp__event__t *hip__event;
#pragma omp task detach(hip__event) // task A
    {
        do__something();
        hipMemcpyAsync(dst, src, nbytes, hipMemcpyDeviceToHost, stream);
        hipStreamAddCallback(stream, callback, hip__event, 0);
    } ①
#pragma omp task // task B
    do__something__else();

#pragma omp taskwait ②
#pragma omp task // task C
    {
        do__other__important__stuff(dst);
    } }

```

1. Task A detaches
2. taskwait does not continue

Putting It All Together

```
void callback(hipStream__t stream, hipError__t status, void *cb__dat){
  ③ omp__fulfill__event((omp__event__t*) cb__data);
}

void hip__example(){
  omp__event__t *hip__event;
  #pragma omp task detach(hip__event) // task A
  {
    do__something();
    hipMemcpyAsync(dst, src, nbytes, hipMemcpyDeviceToHost, stream);
    hipStreamAddCallback(stream, callback, hip__event, 0);
  } ①
  #pragma omp task // task B
  do__something__else();

  #pragma omp taskwait ②
  #pragma omp task // task C
  {
    do__other__important__stuff(dst);
  } }
}
```



1. Task A detaches
2. taskwait does not continue
3. When memory transfer completes, callback is invoked to signal the event for task completion

Putting It All Together

```

void callback(hipStream__t stream, hipError__t status, void *cb__dat){
  ③ omp__fulfill__event((omp__event__t*) cb__data);
}

void hip__example(){
  omp__event__t *hip__event;
#pragma omp task detach(hip__event) // task A
  {
    do__something();
    hipMemcpyAsync(dst, src, nbytes, hipMemcpyDeviceToHost, stream);
    hipStreamAddCallback(stream, callback, hip__event, 0);
  } ①
#pragma omp task // task B
  do__something__else();

#pragma omp taskwait ②④
#pragma omp task // task C
  {
    do__other__important__stuff(dst);
  } }

```



1. Task A detaches
2. taskwait does not continue
3. When memory transfer completes, callback is invoked to signal the event for task completion
4. taskwait continues, task C executes

Removing the taskwait Construct

```
void callback(hipStream__t stream, hipError__t status, void *cb__dat){
    omp__fulfill__event((omp__event__t *) cb__data);
}

void hip__example(){
    omp__event__t *hip__event;
#pragma omp task depend(out:dst) detach(hip__event) // task A
    {
        do__something();
        hipMemcpyAsync(dst, src, nbytes, hipMemcpyDeviceToHost, stream);
        hipStreamAddCallback(stream, callback, hip__event, 0);
    }
#pragma omp task // task B
    do__something__else();

#pragma omp task depend(in:dst) // task C
    {
        do__other__important__stuff(dst);
    }
}
```

Removing the taskwait Construct

```
void callback(hipStream__t stream, hipError__t status, void *cb__dat){
    omp__fulfill__event((omp__event__t *) cb__data);
}

void hip__example(){
    omp__event__t *hip__event;
#pragma omp task depend(out:dst) detach(hip__event) // task A
    {
        do__something();
        hipMemcpyAsync(dst, src, nbytes, hipMemcpyDeviceToHost, stream);
        hipStreamAddCallback(stream, callback, hip__event, 0);
    }
#pragma omp task // task B
    do__something__else();

#pragma omp task depend(in:dst) // task C
    {
        do__other__important__stuff(dst);
    }
}
```

1. Task A detaches and task C will not execute because of its unfulfilled dependency on A

Removing the taskwait Construct

```

void callback(hipStream__t stream, hipError__t status, void *cb__dat){
  ② omp__fulfill__event((omp__event__t *) cb__data);
}
void hip__example(){
  omp__event__t *hip__event;
#pragma omp task depend(out:dst) detach(hip__event) // task A
  {
    do__something();
    hipMemcpyAsync(dst, src, nbytes, hipMemcpyDeviceToHost, stream);
    ① hipStreamAddCallback(stream, callback, hip__event, 0);
  }
#pragma omp task // task B
  do__something__else();

#pragma omp task depend(in:dst) // task C
  {
    do__other__important__stuff(dst);
  }
}

```



1. Task A detaches and task C will not execute because of its unfulfilled dependency on A
2. When memory transfer completes, callback is invoked to signal the event for task completion

Removing the taskwait Construct

```

void callback(hipStream__t stream, hipError__t status, void *cb__dat){
  ② omp__fulfill__event((omp__event__t *) cb__data);
}
void hip__example(){
  omp__event__t *hip__event;
#pragma omp task depend(out:dst) detach(hip__event) // task A
  {
    do__something();
    hipMemcpyAsync(dst, src, nbytes, hipMemcpyDeviceToHost, stream);
    ① hipStreamAddCallback(stream, callback, hip__event, 0);
  }
#pragma omp task // task B
  do__something__else();

#pragma omp task depend(in:dst) // task C
  {
    do__other__important__stuff(dst);
  }
}

```



1. Task A detaches and task C will not execute because of its unfulfilled dependency on A
2. When memory transfer completes, callback is invoked to signal the event for task completion
3. Task A completes and C's dependency is fulfilled

Case Study: NWChem TCE CCSD(T)

TCE: Tensor Contraction Engine

CCSD(T): Coupled-Cluster with Single, Double,
and perturbative Triple replacements

NWChem

- Computational chemistry software package
 - Quantum chemistry
 - Molecular dynamics
- Designed for large-scale supercomputers
- Developed at the EMSL at PNNL
 - EMSL: Environmental Molecular Sciences Laboratory
 - PNNL: Pacific Northwest National Lab
- URL: <http://www.nwchem-sw.org>

Finding Offload Candidates

- Requirements for offload candidates
 - Compute-intensive code regions (kernels)
 - Highly parallel
 - Compute scaling stronger than data transfer, e.g., compute $O(n^3)$ vs. data size $O(n^2)$

Example Kernel (1 of 27 in total)

```

subroutine sd__t__d1__l(h3d,h2d,h1d,p6d,p5d,p4d,
|   h7d,triplex,t2sub,v2sub)
c Declarations omitted.
double precision triplex(h3d*h2d,h1d,p6d,p5d,p4d)
double precision t2sub(h7d,p4d,p5d,h1d)
double precision v2sub(h3d*h2d,p6d,h7d)
!$omp target „presence?(triplex,t2sub,v2sub)”
!$omp teams distribute parallel do private(p4,p5,p6,h2,h3,h1,h7)
do p4=1,p4d
do p5=1,p5d
do p6=1,p6d
do h1=1,h1d
do h7=1,h7d
do h2h3=1,h3d*h2d
triplex(h2h3,h1,p6,p5,p4)=triplex(h2h3,h1,p6,p5,p4)
| - t2sub(h7,p4,p5,h1)*v2sub(h2h3,p6,h7)
end do
end do
end do
end do
end do
end do
!$omp end teams distribute parallel do
!$omp end target
end subroutine

```

- All kernels have the same structure
- 7 perfectly nested loops
- Some kernels contain inner product loop (then, 6 perfectly nested loops)
- Trip count per loop is equal to “tile size” (20-30 in production)

Example Kernel (1 of 27 in total)

```

subroutine sd__t__d1__l(h3d,h2d,h1d,p6d,p5d,p4d,
  l      h7d,triplex,t2sub,v2sub)
c Declarations omitted.
double precision triplex(h3d*h2d,h1d,p6d,p5d,p4d)
double precision t2sub(h7d,p4d,p5d,h1d)
double precision v2sub(h3d*h2d,p6d,h7d)
!$omp target „presence?(triplex,t2sub,v2sub)“
!$omp teams distribute parallel do private(p4,p5,p6,h2,h3,h1,h7)
  do p4=1,p4d
    do p5=1,p5d
      do p6=1,p6d
        do h1=1,h1d
          do h7=1,h7d
            do h2h3=1,h3d*h2d
              triplex(h2h3,h1,p6,p5,p4)=triplex(h2h3,h1,p6,p5,p4)
              l - t2sub(h7,p4,p5,h1)*v2sub(h2h3,p6,h7)
            end do
          end do
        end do
      end do
    end do
  end do
!$omp end teams distribute parallel do
!$omp end target
end subroutine

```

- All kernels have the same structure
- 7 perfectly nested loops
- Some kernels contain inner product loop (then, 6 perfectly nested loops)
- Trip count per loop is equal to “tile size” (20-30 in production)
- Naïve data allocation (tile size 24)
 - Per-array transfer for each **target** construct
 - triplex: 1458 MB
 - t2sub, v2sub: 2.5 MB each

Example Kernel (1 of 27 in total)

```

subroutine sd__t__d__l(h3d,h2d,h1d,p6d,p5d,p4d,
  | h7d,triplex,t2sub,v2sub)
c Declarations omitted.
double precision triplex(h3d*h2d,h1d,p6d,p5d,p4d)
double precision t2sub(h7d,p4d,p5d,h1d)
double precision v2sub(h3d*h2d,p6d,h7d)
!$omp target „presence?(triplex,t2sub,v2sub)“
!$omp teams distribute parallel do private(p4,p5,p6,h2,h3,h1,h7)
do p4=1,p4d
do p5=1,p5d
do p6=1,p6d
do h1=1,h1d
do h7=1,h7d
do h2h3=1,h3d*h2d
triplex(h2h3,h1,p6,p5,p4)=triplex(h2h3,h1,p6,p5,p4)
| - t2sub(h7,p4,p5,h1)*v2sub(h2h3,p6,h7)
end do
end do
end do
end do
end do
!$omp end teams distribute parallel do
!$omp end target
end subroutine

```

1.5GB data transferred
(host to device)

1.5GB data transferred
(device to host)

- All kernels have the same structure
- 7 perfectly nested loops
- Some kernels contain inner product loop (then, 6 perfectly nested loops)
- Trip count per loop is equal to “tile size” (20-30 in production)
- Naïve data allocation (tile size 24)
 - Per-array transfer for each **target** construct
 - triplex: 1458 MB
 - t2sub, v2sub: 2.5 MB each

Invoking the Kernels / Data Management

- Simplified pseudo-code

```

1$omp target enter data map(alloc:triplexx(i:tr__size))
c  for all tiles
  do ...
    call zero__triplexx(triplexx)
    do ...
      call comm__and__sort(t2sub, v2sub)
1$omp target data map(to:t2sub(t2__size)) map(to:v2sub(v2__size))
  if (...)
    call sd__t__d1__1(h3d,h2d,h1d,p6d,p5d,p4d,h7,triplexx,t2sub,v2sub)
  end if
c  same for sd__t__d1__2 until sd__t__d1__9
1$omp target end data
  end do
  do ...
c  Similar structure for sd__t__d2__1 until sd__t__d2__9, incl. target data
  end do
  call sum__energy(energy, triplexx)
  end do
1$omp target exit data map(release:triplexx(i:size))

```

- Reduced data transfers:

Invoking the Kernels / Data Management

■ Simplified pseudo-code

```

!$omp target enter data map(alloc:triplexx(i:tr__size))
c  for all tiles
do ...
  call zero__triplexx(triplexx)
do ...
  call comm__and__sort(t2sub, v2sub)
!$omp target data map(to:t2sub(t2__size)) map(to:v2sub(v2__size))
  if (...)
    call sd__t__d1__1(h3d,h2d,h1d,p6d,p5d,p4d,h7,triplexx,t2sub,v2sub)
  end if
c  same for sd__t__d1__2 until sd__t__d1__9
!$omp target end data
end do
do ...
c  Similar structure for sd__t__d2__1 until sd__t__d2__9, incl. target data
end do
  call sum__energy(energy, triplexx)
end do
!$omp target exit data map(release:triplexx(i:size))

```

Allocate 1.5GB data once,
stays on device.

■ Reduced data transfers:

■ triplexx:

- allocated once
- always kept on the target

Invoking the Kernels / Data Management

■ Simplified pseudo-code

```

!$omp target enter data map(alloc:triplexx(i:tr__size))
c  for all tiles
do ...
  call zero__triplexx(triplexx)
do ...
  call comm__and__sort(t2sub, v2sub)
!$omp target data map(to:t2sub(t2__size)) map(to:v2sub(v2__size))
  if (...)
    call sd__t__d1__1(h3d,h2d,h1d,p6d,p5d,p4d,h7,triplexx,t2sub,v2sub)
  end if
c  same for sd__t__d1__2 until sd__t__d1__9
!$omp target end data
end do
do ...
c  Similar structure for sd__t__d2__1 until sd__t__d2__9, incl. target data
end do
  call sum__energy(energy, triplexx)
end do
!$omp target exit data map(release:triplexx(i:size))

```

Allocate 1.5GB data once,
stays on device.

Update 2x2.5MB of data for
(potentially) multiple kernels.

■ Reduced data transfers:

- triplexx:
 - allocated once
 - always kept on the target
- t2sub, v2sub:
 - allocated after comm.
 - kept for (multiple) kernel invocations

Invoking the Kernels / Data Management

■ Simplified pseudo-code

```

!$omp target enter data map(alloc:triplexx(1:tr__size))
c for all tiles
do ...
  call zero__triplexx(triplexx)
do ...
  call comm__and__sort(t2sub, v2sub)
!$omp target data map(to:t2sub(t2__size)) map(to:v2sub(v2__size))
  if (...)
    call sd__t__d1__(h3d,h2d,h1d,p6d,p5d,p4d,h7,triplexx,t2sub,v2sub)
  end if
c same for sd__t__d1__2 until sd__t__d1__9
!$omp target end data
end do
do ...
c Similar structure for sd__t__d2__1 until sd__t__d2__9, incl. target data
end do
  call sum__energy(energy, triplexx)
end do
!$omp target exit data map(release:triplexx(1:size))

```

Allocate 1.5G
stays on

Update 2x2.5
(potentially) r

```

subroutine sd__t__d1__(h3d,h2d,h1d,p6d,p5d,p4d,
  h7d,triplexx,t2sub,v2sub)
c Declarations omitted.
double precision triplexx(h3d*h2d,h1d,p6d,p5d,p4d)
double precision t2sub(h7d,p4d,p5d,h1d)
double precision v2sub(h3d*h2d,p6d,h7d)
!$omp target „presence?(triplexx,t2sub,v2sub)”
!$omp teams distribute parallel do private(p4,p5,p6,h2,h3,h1,h7)
  do p4=1,p4d
  do p5=1,p5d
  do p6=1,p6d
  do h1=1,h1d
  do h7=1,h7d
  do h2h3=1,h3d*h2d
    triplexx(h2h3,h1,p6,p5,p4)=triplexx(h2h3,h1,p6,p5,p4)
    t2sub(h7,p4,p5,h1)*v2sub(h2h3,p6,h7)
  end do
  end do
  end do
  end do
  end do
  end do
!$omp end teams distribute parallel do
!$omp end target
end subroutine

```

Invoking the Kernels / Data Management

■ Simplified pseudo-code

```

!$omp target enter data map(alloc:triplexx(1:tr__size))
c for all tiles
do ...
  call zero__triplexx(triplexx)
do ...
  call comm__and__sort(t2sub, v2sub)
!$omp target data map(to:t2sub(t2__size)) map(to:v2sub(v2__size))
  if (...)
    call sd__t__d1__(h3d,h2d,h1d,p6d,p5d,p4d,h7,triplexx,t2sub,v2sub)
  end if
c same for sd__t__d1__2 until sd__t__d1__9
!$omp target end data
end do
do ...
c Similar structure for sd__t__d2__1 until sd__t__d2__9, incl. target data
end do
  call sum__energy(energy, triplexx)
end do
!$omp target exit data map(release:triplexx(1:size))

```

Allocate 1.5G
stays on

Update 2x2.5
(potentially) r

```

subroutine sd__t__d1__(h3d,h2d,h1d,p6d,p5d,p4d,
  h7d,triplexx,t2sub,v2sub)
c Declarations omitted.
double precision triplexx(h3d*h2d,h1d,p6d,p5d,p4d)
double precision t2sub(h7d,p4d,p5d,h1d)
double precision v2sub(h3d*h2d,p6d,h7d)
!$omp target „presence?(triplexx,t2sub,v2sub)”
!$omp teams distribute parallel do private(p4,p5,p6,h2,h3,h1,h7)
  do p4=1,p4d
  do p5=1,p5d
  do p6=1,p6d
  do h1=1,h1d
  do h7=1,h7d
  do h2h3=1,h3d*h2d
    triplexx(h2h3,h1,p6,p5,p4)
    v2sub(h7,p4,p5,h1)*v2sub(h2,h3)
  end do
  end do
  end do
  end do
  end do
  end do
!$omp end teams distribute parallel do
!$omp end target
end subroutine

```

Presence check determines that arrays
have been allocated in the device data
environment already.