Introduction to OpenMP - GPUs

International HPC Summer School

Ludovic Capelli

EPCC

July 8, 2024

Individuals

The material presented in this lecture is inspired from content developed by:

- John Urbanic
- Michael Klemm

Contributors

In this material:

- The slides are created using LaTeX *Beamer* available at `https://ctan.org/pkg/beamer`.
- The sequence diagrams are created using an extended version of the *pgf-umlsd* package available at `https://ctan.org/pkg/pgf-umlsd`.

License - Creative Commons BY-NC-SA-4.0[1]

Non-Commercial you may not use the material for commercial purposes.

Shared-Alike if you remix, transform, or build upon the material, you must distribute your contributions under the same license as the original.

Attribution you must give appropriate credit, provide a link to the license, and indicate if changes were made. You may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your use.

[1] You can find the full documentation about this license at:
`https://creativecommons.org/licenses/by-nc-sa/4.0/`

Nota bene

Format

This set of slides is designed to be covered as a **live presentation**, and thus contains little text and explanations. It is **less suitable** to be studied as a stand-alone document.

Moment of truth

Before we start:

- Connect to the Bridges node

```
ssh your_username@bridges2.psc.edu
```

- This repository is on Github.
- Clone the repository used in this session

```
git clone https://github.com/capellil/
    IHPCSS_Introduction_to_OpenMP_GPU_examples
```

Moment of truth

- You will see it contains multiple folders, numbered.
- Each of which will be an illustration to a concept we will see.

Moment of truth

- Load the modules needed for the GPU ecosystem.

```
1 module load nvhpc/22.9;
2 module load openmpi/4.0.5-nvhpc22.9;
```

- Put these two lines in your `.bashrc` file.

Moment of truth

- Go to the repository, inside the first folder.

```
cd repository/<language>/1.Preamble
```

- Compile the source code in it

```
make
```

- Run the executable on the compute node

```
sbatch submission_bridges.slurm
```

Table of content

Table of Contents

What are GPUs?

What are GPUs?

- Stands for **G**raphics **P**rocessing **U**nits.

What are GPUs?

- Stands for **G**raphics **P**rocessing **U**nits.
- Devices specialised in... graphics processing.

What are GPUs?

- Stands for **G**raphics **P**rocessing **U**nits.
- Devices specialised in... graphics processing.
- Typically has a given operation to apply to a lot of inputs
  - points in a mesh
  - pixels in a picture

What are GPUs?

- Stands for **G**raphics **P**rocessing **U**nits.
- Devices specialised in... graphics processing.
- Typically has a given operation to apply to a lot of inputs
  - points in a mesh
  - pixels in a picture
- The GPU architecture allows them to complete this task very effectively.

Use case

- Imagine you are developing a graphic editing software.
- You are currently implementing a filter allowing to darken a picture by making every pixel 20% darker.
- The same calculation is applied to every pixel, regardless of its current colour.

How CPUs would do it

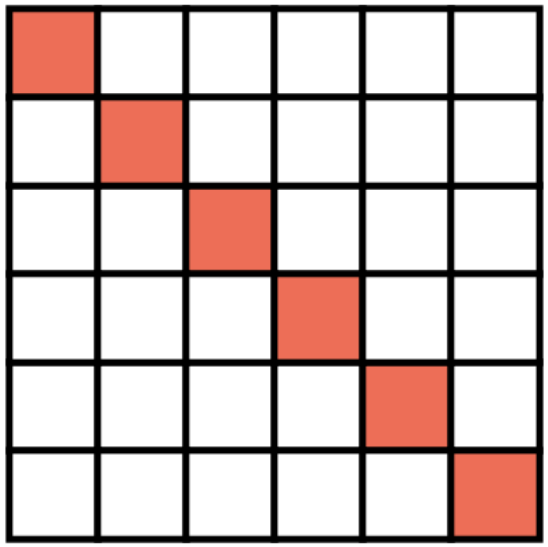

How CPUs would do it.

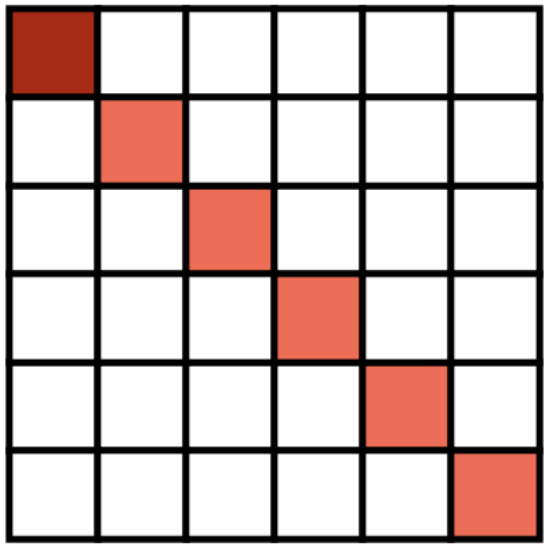

How CPUs would do it

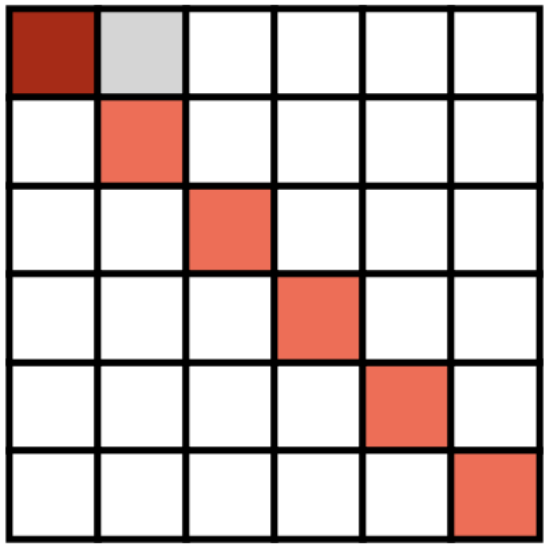

How CPUs would do it

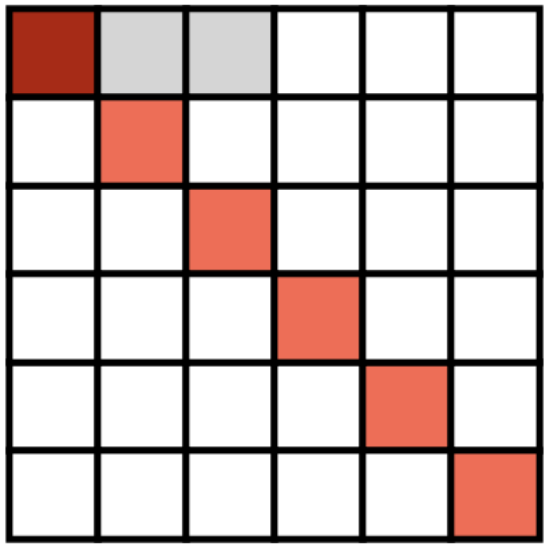

How GPUs work.

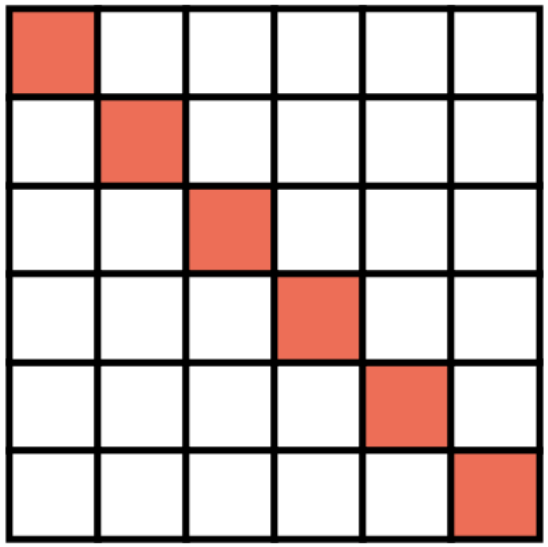

How GPUs work.

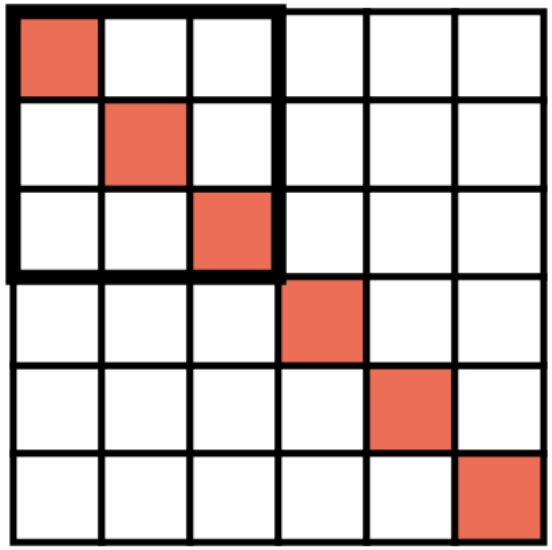

How GPUs work.

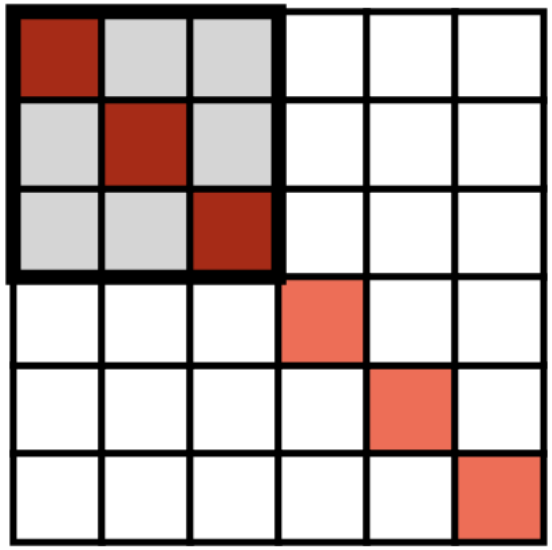

How GPUs work.

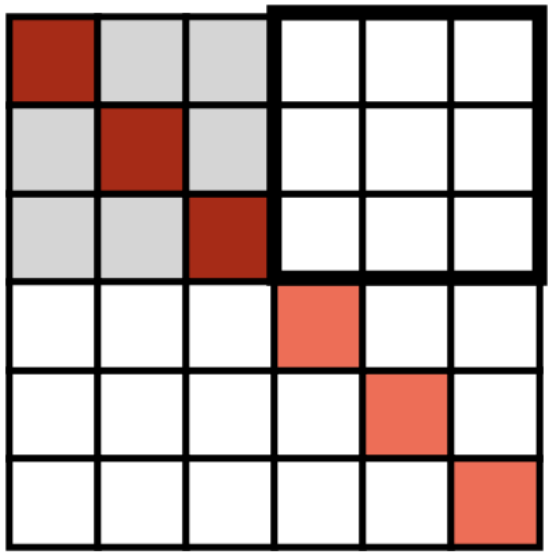

How GPUs work.

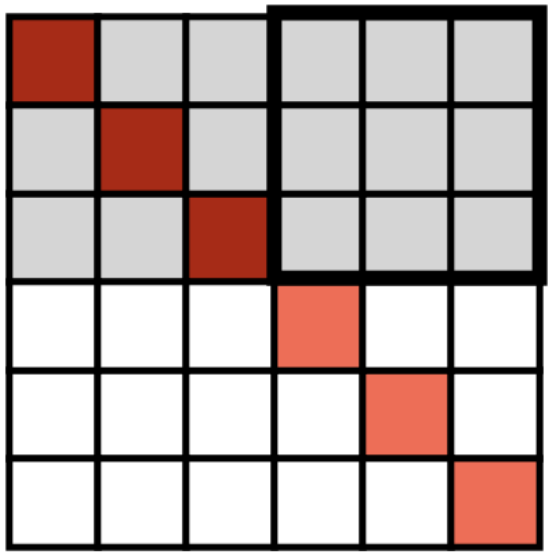

How GPUs work.

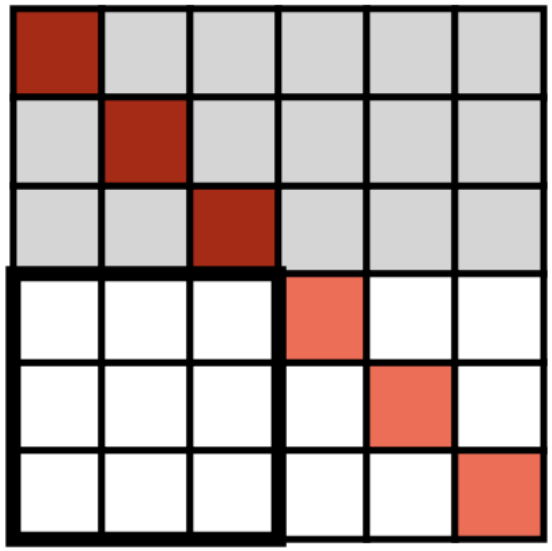

How GPUs work.

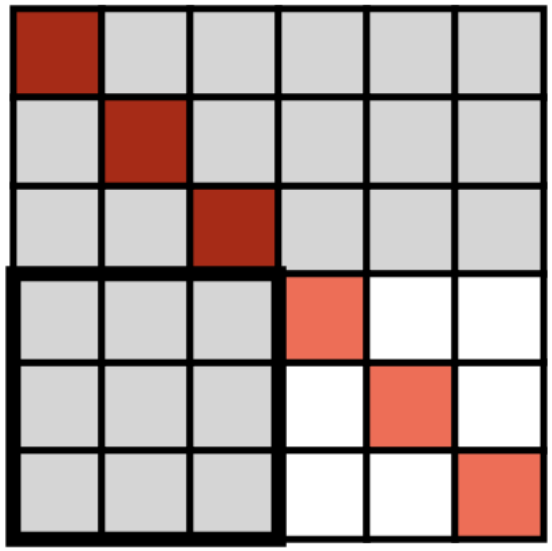

How GPUs work.

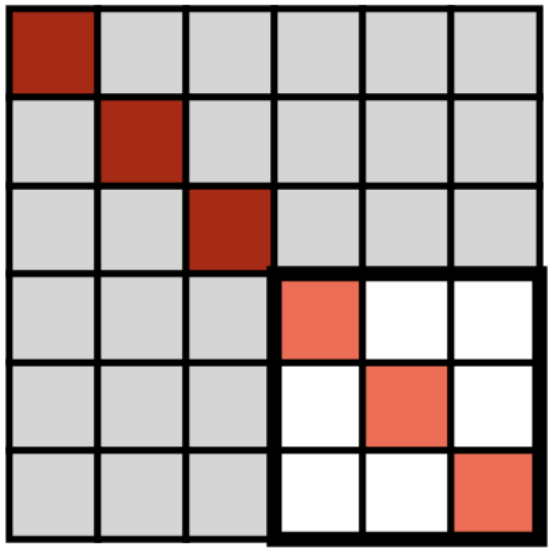

How GPUs work.

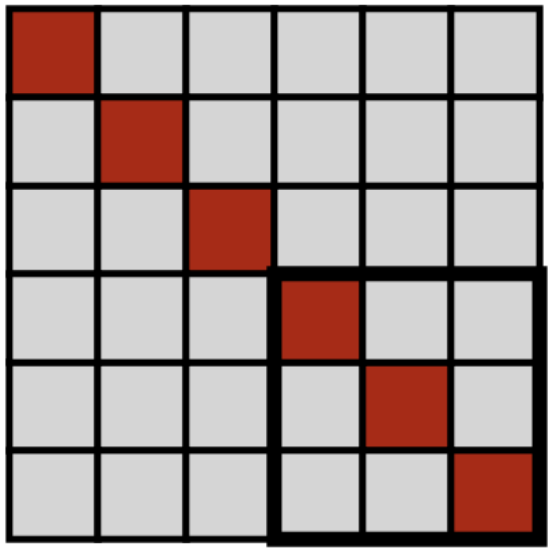

How GPUs work.

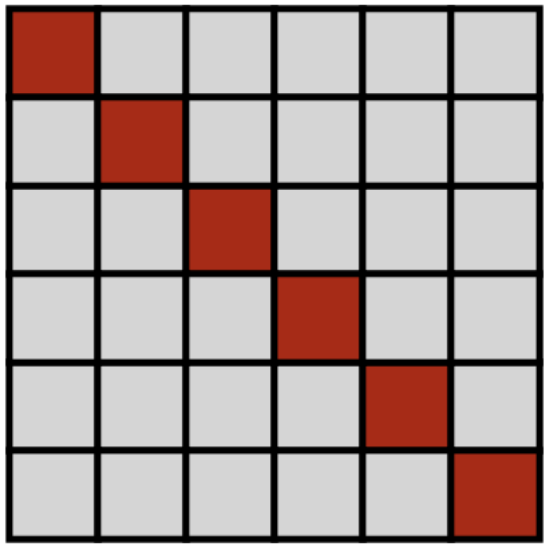

How GPUs work.

Demonstration by Mythbusters and NVIDIA at https://www.youtube.com/watch?v=-P28LKWTzrI.

GPU specs

- On Bridges2, the GPU nodes have NVIDIA V100-32GB SXM2.

GPU specs

- On Bridges2, the GPU nodes have NVIDIA V100-32GB SXM2.
  - Core clock: 1.3GHz, boost at 1.5GHz.

GPU specs

- On Bridges2, the GPU nodes have NVIDIA V100-32GB SXM2.
  - Core clock: 1.3GHz, boost at 1.5GHz.
  - 5,120 cores.

GPU specs

- On Bridges2, the GPU nodes have NVIDIA V100-32GB SXM2.
  - Core clock: 1.3GHz, boost at 1.5GHz.
  - 5,120 cores.
  - Bandwidth: 900GB/s

GPU specs

- On Bridges2, the GPU nodes have NVIDIA V100-32GB SXM2.
  - Core clock: 1.3GHz, boost at 1.5GHz.
  - 5,120 cores.
  - Bandwidth: 900GB/s
  - Max Thermal Design Power: 250W.

GPU specs

- On Bridges2, the GPU nodes have NVIDIA V100-32GB SXM2.
  - Core clock: 1.3GHz, boost at 1.5GHz.
  - 5,120 cores.
  - Bandwidth: 900GB/s
  - Max Thermal Design Power: 250W.
- Each GPU node on Bridges2 has 8 such GPUs...
  - Total of over 40K+ GPU cores per node
  - On top of the two Intel Xeon Gold 6248 "Cascade Lake" CPUs per node

Advantages of GPUs

Advantages of GPUs

- Computation tanks: 15.7 TFLOPS per V100GPU, over 120 TFLOPS per GPU node.

Advantages of GPUs

- Computation tanks: 15.7 TFLOPS per V100GPU, over 120 TFLOPS per GPU node.
- Higher power efficiency: 250W per GPU.

Advantages of GPUs

- Computation tanks: 15.7 TFLOPS per V100GPU, over 120 TFLOPS per GPU node.
- Higher power efficiency: 250W per GPU.
- Low price per core:
  - £5,400 for 1 CPU that has 20 cores, so ∼£270 per core.

Advantages of GPUs

- Computation tanks: 15.7 TFLOPS per V100GPU, over 120 TFLOPS per GPU node.
- Higher power efficiency: 250W per GPU.
- Low price per core:
  - £5,400 for 1 CPU that has 20 cores, so ∼£270 per core.
  - £4,000 for 1 GPU that has 5,000 cores, so ∼£1 per core.

Advantages of GPUs

- Computation tanks: 15.7 TFLOPS per V100GPU, over 120 TFLOPS per GPU node.
- Higher power efficiency: 250W per GPU.
- Low price per core:
  - £5,400 for 1 CPU that has 20 cores, so ∼£270 per core.
  - £4,000 for 1 GPU that has 5,000 cores, so ∼£1 per core.

Why even bothering with CPUs then?

Limitations of GPUs

[2] though alleviated now with unified memory, which in turns raises new challenges as now the CPUs and the GPUs can have race conditions between them

Limitations of GPUs

- Memory transfers[2]

[2] though alleviated now with unified memory, which in turns raises new challenges as now the CPUs and the GPUs can have race conditions between them

Limitations of GPUs

- Memory transfers[2]
- Do not handle divergence well (needs to get cores idle)

[2] though alleviated now with unified memory, which in turns raises new challenges as now the CPUs and the GPUs can have race conditions between them

Limitations of GPUs

- Memory transfers[2]
- Do not handle divergence well (needs to get cores idle)
- Simpler **I**nstruction **S**et **A**rchitecture (ISA): operations issued do not support as advanced calculations

[2] though alleviated now with unified memory, which in turns raises new challenges as now the CPUs and the GPUs can have race conditions between them

Limitations of GPUs

- Memory transfers[2]
- Do not handle divergence well (needs to get cores idle)
- Simpler **I**nstruction **S**et **A**rchitecture (ISA): operations issued do not support as advanced calculations
- Non-coalesced memory accesses

[2] though alleviated now with unified memory, which in turns raises new challenges as now the CPUs and the GPUs can have race conditions between them

Opportunity to seize

- Despite having limitations, GPUs still offer a highly attractive solution to certain types of workloads, especially in scientific simulations.

Opportunity to seize

- Despite having limitations, GPUs still offer a highly attractive solution to certain types of workloads, especially in scientific simulations.
- Decided to harness that graphical pipeline for more general computing tasks.

Opportunity to seize

- Despite having limitations, GPUs still offer a highly attractive solution to certain types of workloads, especially in scientific simulations.
- Decided to harness that graphical pipeline for more general computing tasks.
- Birth of GPGPU: General Purpose Graphical Processing Unit.

Solutions

CUDA First on it, free, amazing documentation, but only available for NVIDIA GPUs. Imposed their terms: CUDA cores, CUDA threads, warps etc...

Solutions

CUDA First on it, free, amazing documentation, but only available for NVIDIA GPUs. Imposed their terms: CUDA cores, CUDA threads, warps etc...

OpenACC **Open Acc**elerators. Directive-based solution for GPGPU. Limited support in compilers.

Solutions

CUDA First on it, free, amazing documentation, but only available for NVIDIA GPUs. Imposed their terms: CUDA cores, CUDA threads, warps etc...

OpenACC **Open Acc**elerators. Directive-based solution for GPGPU. Limited support in compilers.

OpenMP NVIDIA support on their GPUs noticeably improved recently with OpenMP 5.0, this is why we will use it, now that you already know how to use it for CPUs.

Table of Contents

Host vs device[3]

Host The device on which the OpenMP program begins execution.

[3] See Section 1.2.1 in OpenMP standard version 5.2

Host vs device[3]

Host The device on which the OpenMP program begins execution.

Target A device onto which code and data may be offloaded from the host device.

[3] See Section 1.2.1 in OpenMP standard version 5.2

`target` construct

- **Transfers** the control flow to the target device

`target` construct

- **Transfers** the control flow to the target device
- This transfer is:

`target` construct

- **Transfers** the control flow to the target device
- This transfer is:
  - sequential

`target` construct

- **Transfers** the control flow to the target device
- This transfer is:
  - sequential
  - synchronous

`target` construct

- **Transfers** the control flow to the target device
- This transfer is:
  - sequential
  - synchronous
- **In theory**, this can be combined with any OpenMP construct.

`target` construct

- **Transfers** the control flow to the target device
- This transfer is:
  - sequential
  - synchronous
- **In theory**, this can be combined with any OpenMP construct.
- **In practice**, there is only a useful subset of OpenMP features for a target device such as a GPU, e.g., no I/O, limited use of base language features.

Offloading to a device - Example

```
1 // Block A: executed on host
2 // Block B: executed on host
3 // Block C: executed on host
```

```
1 ! Block A: executed on host
2 ! Block B: executed on host
3 ! Block C: executed on host
```

Offloading to a device - Example

```
1 // Block A: executed on host
2 #pragma omp target
3 {
4     // Block B: executed on target
5 }
6 // Block C: executed on host
```

```
1 ! Block A: executed on host
2 !$OMP TARGET
3     ! Block B: executed on target
4 !$OMP END TARGET
5 ! Block C: executed on host
```

Currently on the host or device?

- When you will start to have functions, and nested functions, you may not know at some point whether you are on the host or on a device.
- You can check by calling the function `omp_is_initial_device()`.

Currently on the host or device?

```c
if(omp_is_initial_device())
{
    printf("I am on the host\n");
}
else
{
    printf("I am on the device.\n");
}
```

```fortran
IF(omp_is_initial_device())
    WRITE(*, '(A)') 'I am on the host.'
ELSE
    WRITE(*, '(A)') 'I am on the device.'
END IF
```

Time to practise: `2.HelloWorld`

Update the source code provided such that the for loop marked is executed on the GPU.

Tips

You will need:

- the `target` construct

Table of Contents

Implicit mapping - CPUs

```c
int a = 123;
int b[2] = {456, 789};
#pragma omp parallel
{
    a *= 2;
    b[0]++;
    b[1]++;
}
```

```fortran
INTEGER :: a = 123
INTEGER, DIMENSION(0:1) :: b := (/456, 789/)
!$OMP PARALLEL
    a *= 2
    b(0) = b(0) + 1
    b(1) = b(1) + 1
!$OMP END PARALLEL
```

Implicit mapping - CPUs

```c
int a = 123;
int b[2] = {456, 789};
#pragma omp parallel default(none) shared(a, b)
{
    a *= 2;
    b[0]++;
    b[1]++;
}
```

```fortran
INTEGER :: a = 123
INTEGER, DIMENSION(0:1) :: b = (/456, 789/)
!$OMP PARALLEL DEFAULT(NONE) SHARED(a, b)
    a *= 2;
    b(0) = b(0) + 1
    b(1) = b(1) + 1
!$OMP END PARALLEL
```

Implicit mapping - GPUs

```
1  int a = 123;
2  int b[2] = {456, 789};
3  #pragma omp target
4  {
5      a *= 2;
6      b[0]++;
7      b[1]++;
8  }
```

```
1  INTEGER :: a = 123
2  INTEGER, DIMENSION(0:1) :: b := (/456, 789/)
3  !$OMP TARGET
4      a *= 2
5      b(0) = b(0) + 1
6      b(1) = b(1) + 1
7  !$OMP END TARGET
```

Implicit mapping - GPUs

```
1 int a = 123;
2 int b[2] = {456, 789};
3 #pragma omp target "firstprivate(a)" shared(b)
4 {
5     a *= 2;
6     b[0]++;
7     b[1]++;
8 }
```

```
1 INTEGER :: a = 123
2 INTEGER, DIMENSION(0:1) :: b := (/456, 789/)
3 !$OMP TARGET "FIRSTPRIVATE(a)" SHARED(b)
4     a *= 2
5     b(0) = b(0) + 1
6     b(1) = b(1) + 1
7 !$OMP END TARGET
```

Implicit mapping

By default:

- scalar variables are implicitly mapped as `firstprivate`[4][5].

[4] See Section 2.19.7 in OpenMP standard version 5.0

[5] In OpenMP 4.0, it was automatically mapped with `tofrom`.

Implicit mapping

By default:

- scalar variables are implicitly mapped as `firstprivate`[4][5].
- statically allocated array variables are implicitly mapped as `map(tofrom)`.

[4] See Section 2.19.7 in OpenMP standard version 5.0

[5] In OpenMP 4.0, it was automatically mapped with `tofrom`.

Implicit mapping

By default:

- scalar variables are implicitly mapped as `firstprivate`[45].
- statically allocated array variables are implicitly mapped as `map(tofrom)`.

Warning

When implicitly mapping a scalar, the target device gets an initialised copy of that scalar. However, each thread on the target device creates its own local copy of it, and do not interact with the one that will may be transferred back to the host device.

[4] See Section 2.19.7 in OpenMP standard version 5.0

[5] In OpenMP 4.0, it was automatically mapped with `tofrom`.

Implicit mapping - Arrays

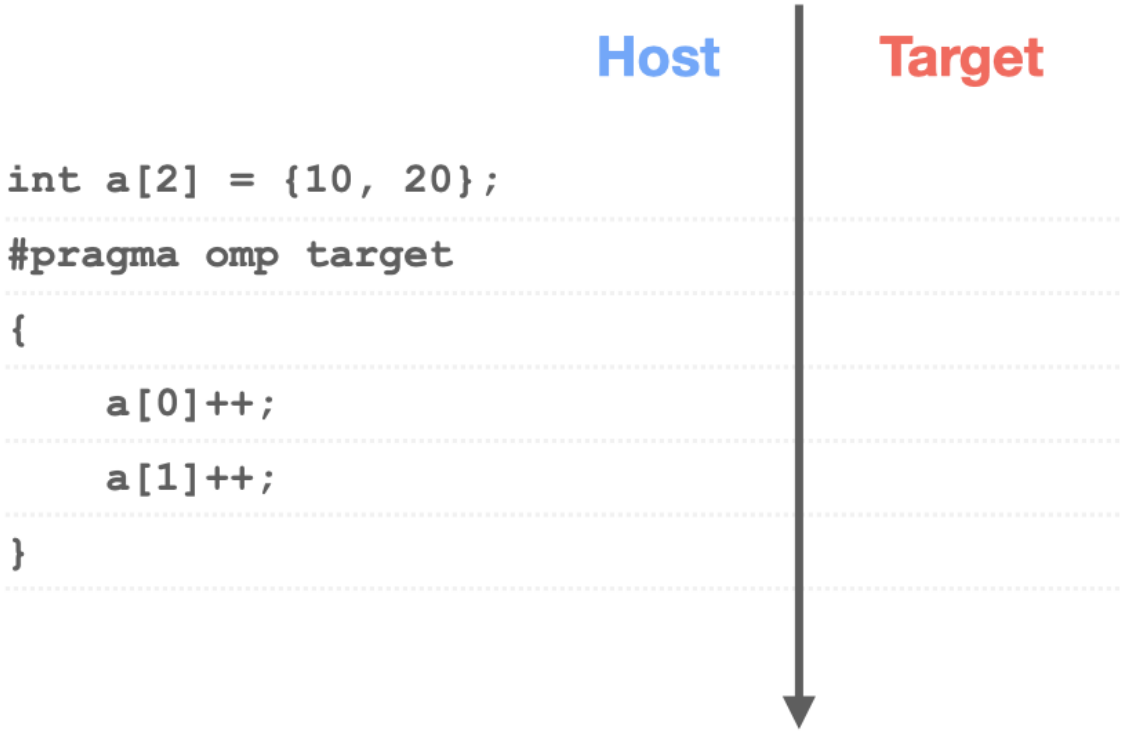

Implicit mapping - Arrays

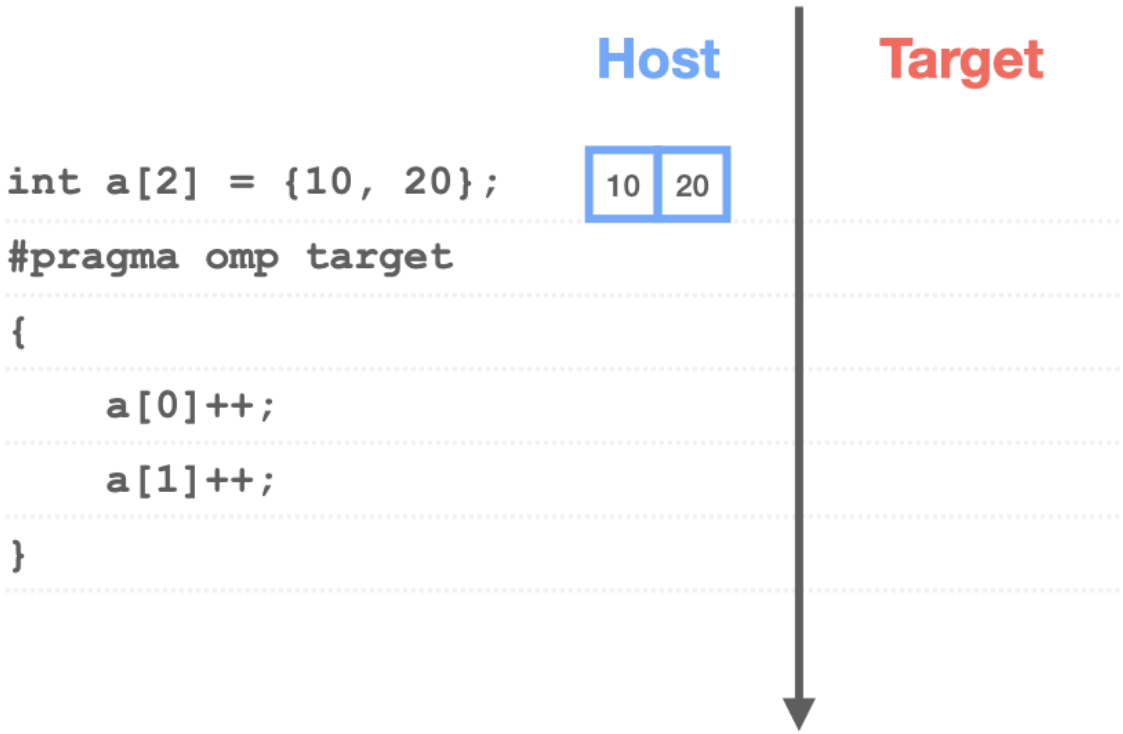

Implicit mapping - Arrays

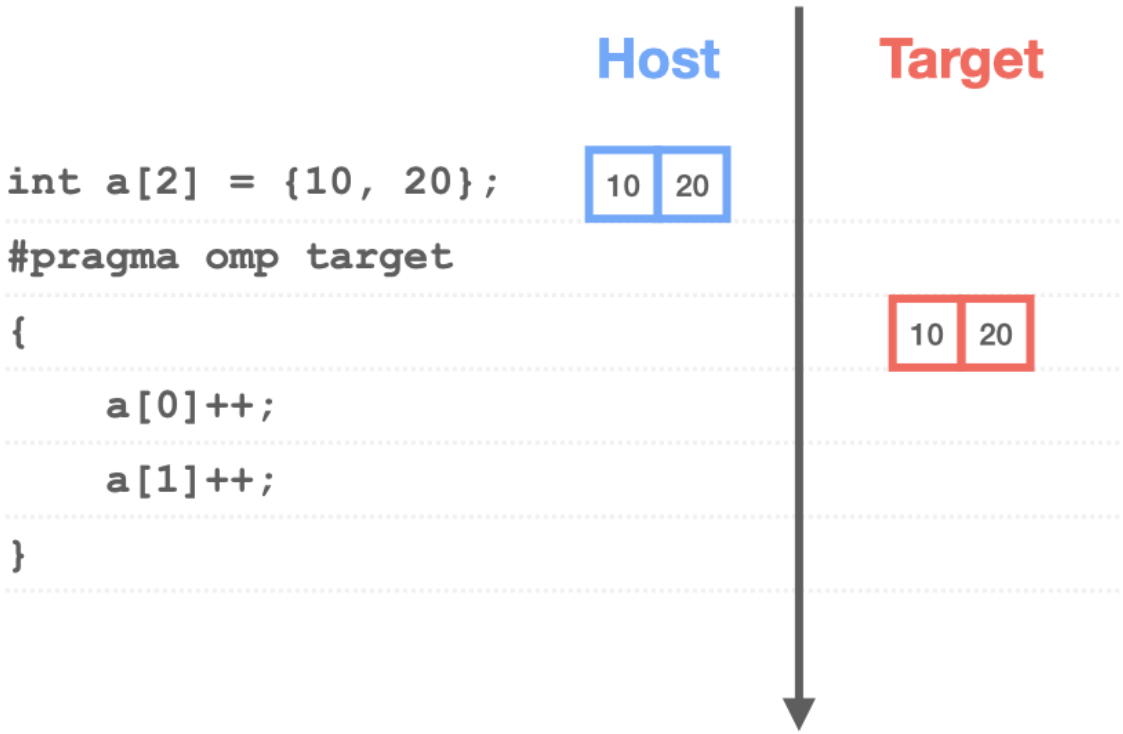

Implicit mapping - Arrays

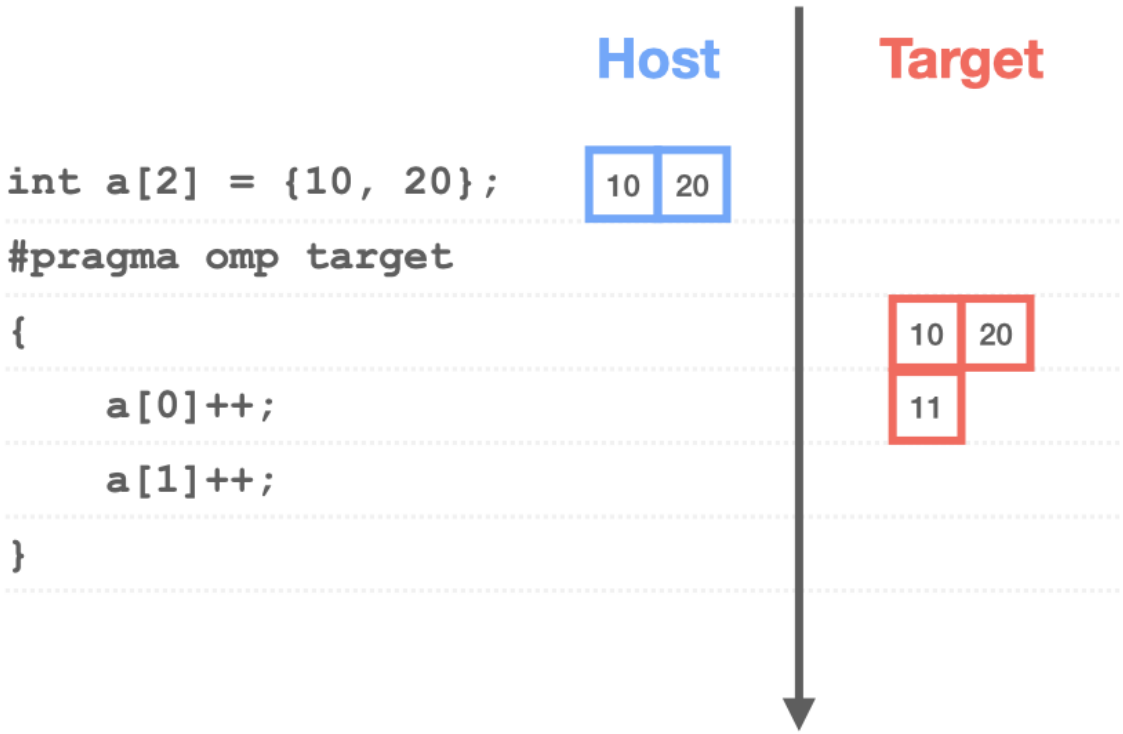

Implicit mapping - Arrays

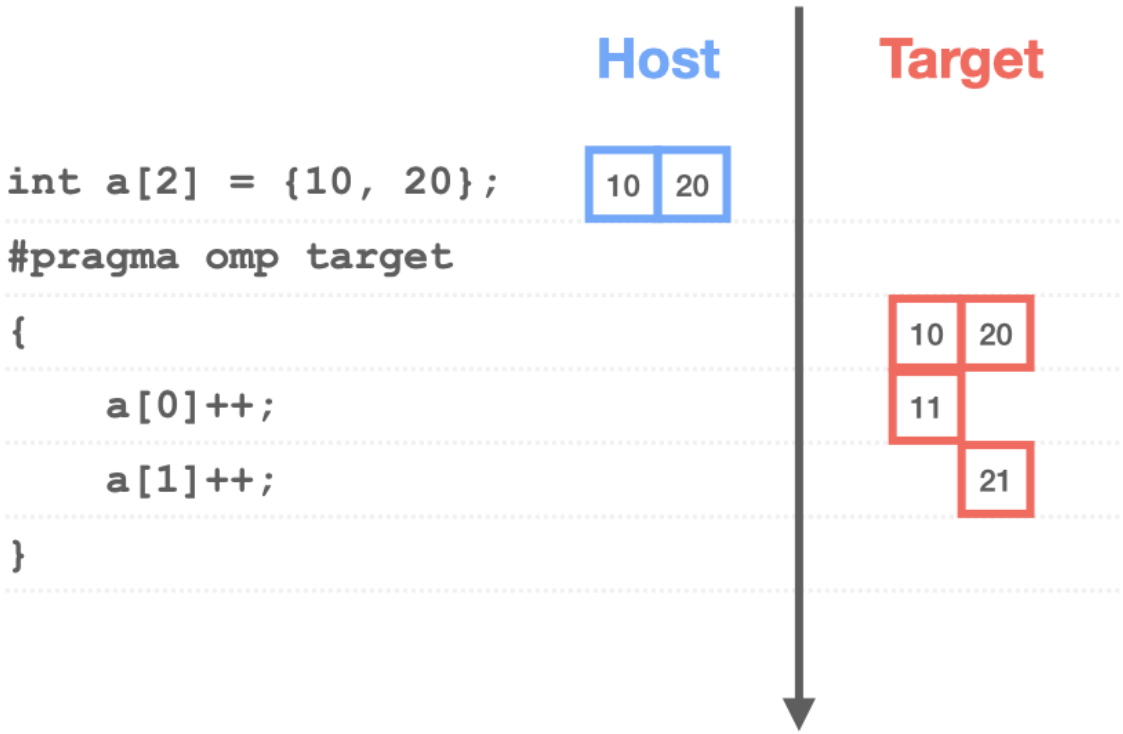

Implicit mapping - Arrays

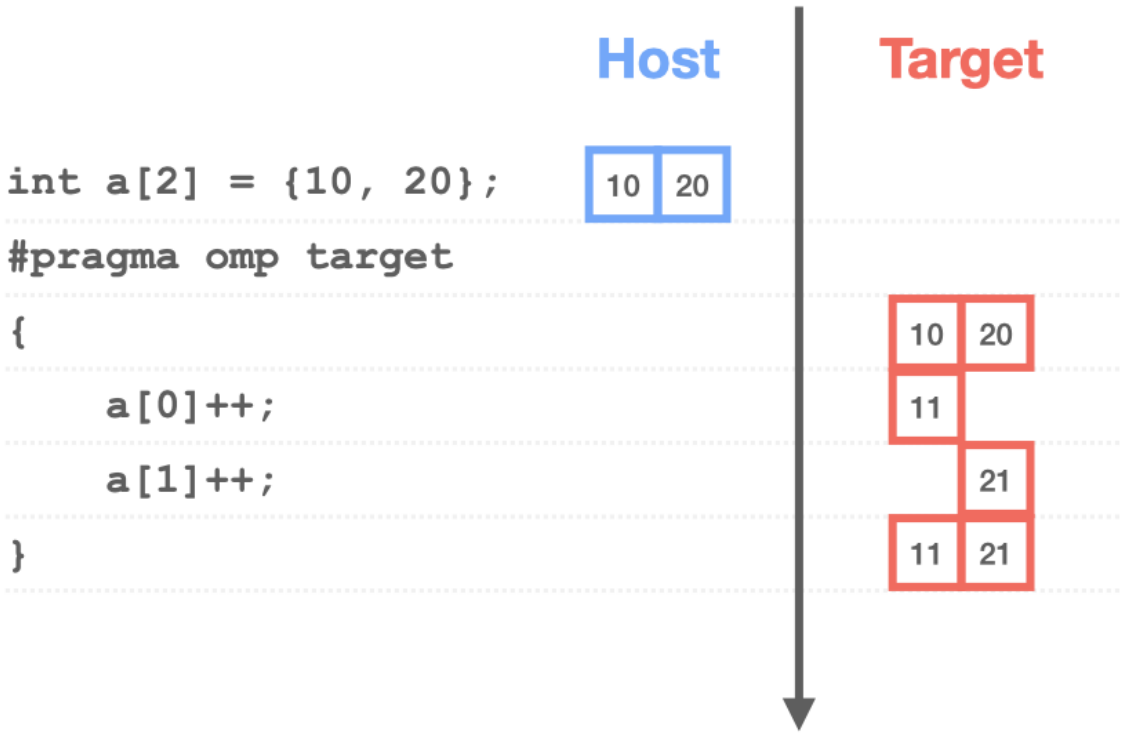

Implicit mapping - Arrays

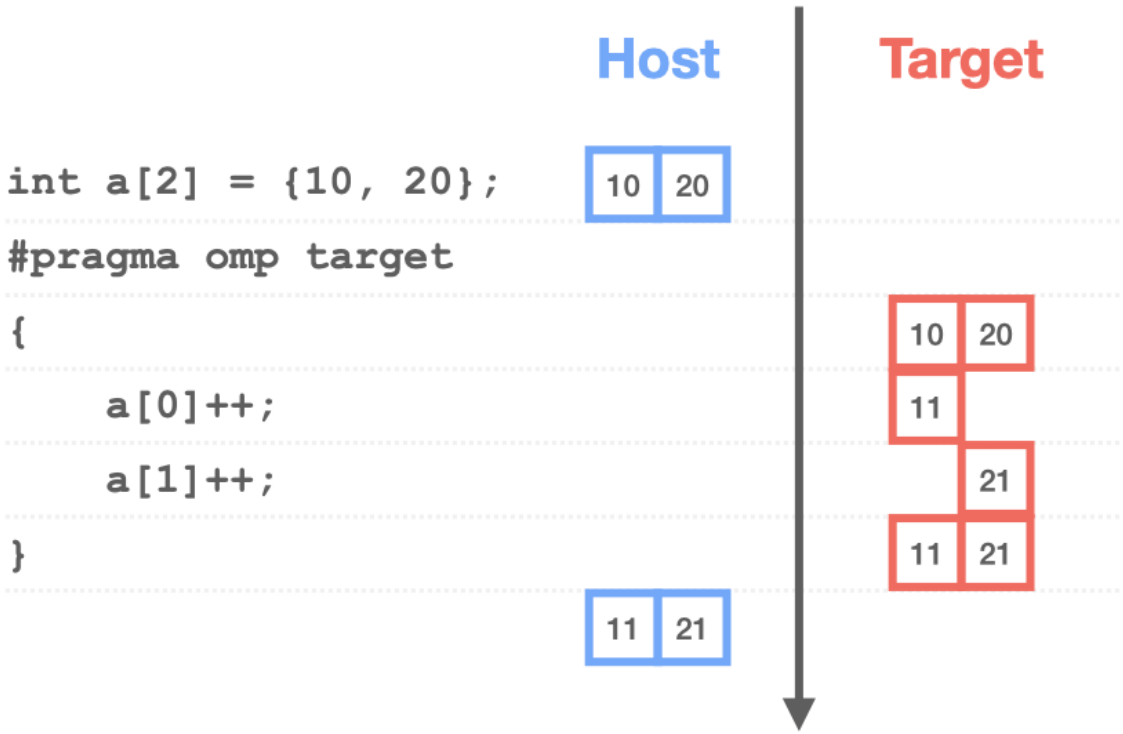

Implicit mapping - Scalars

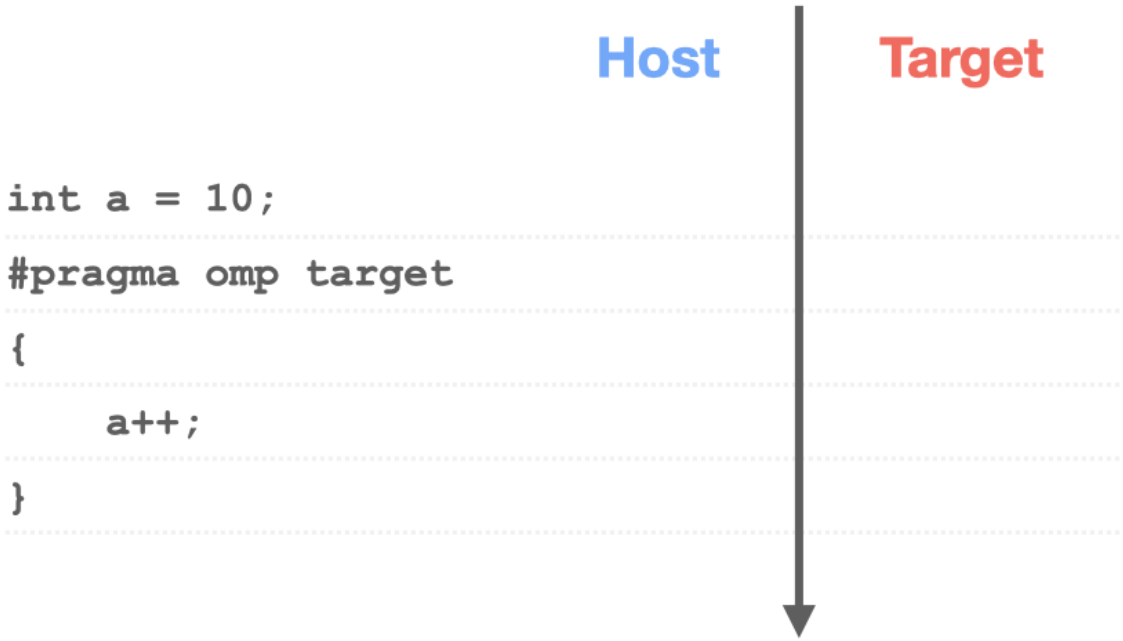

Implicit mapping - Scalars

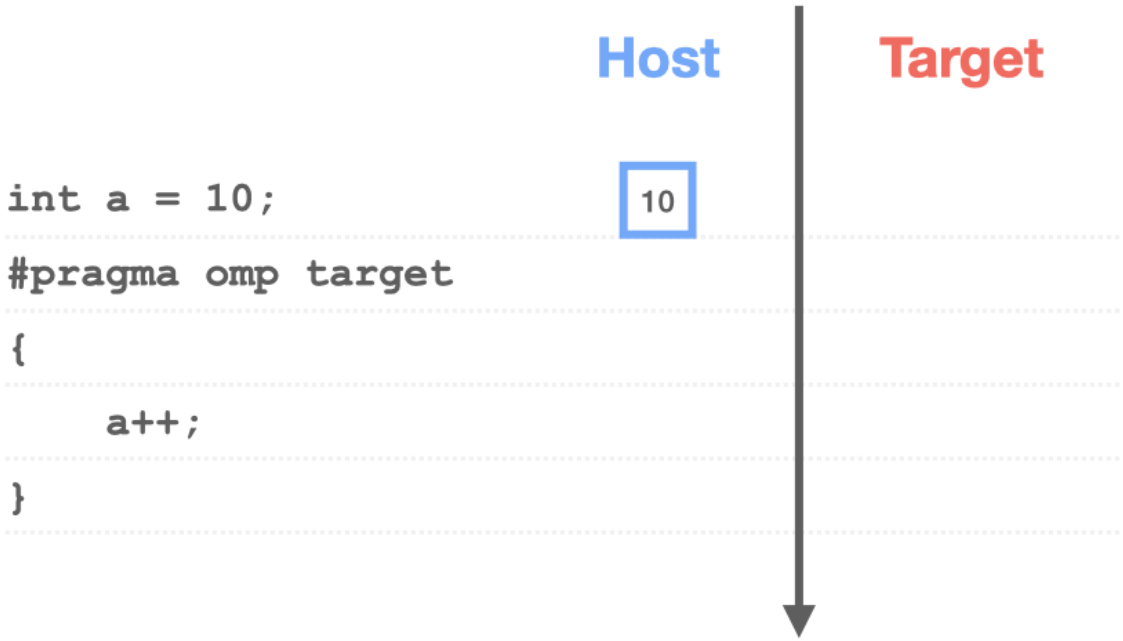

Implicit mapping - Scalars

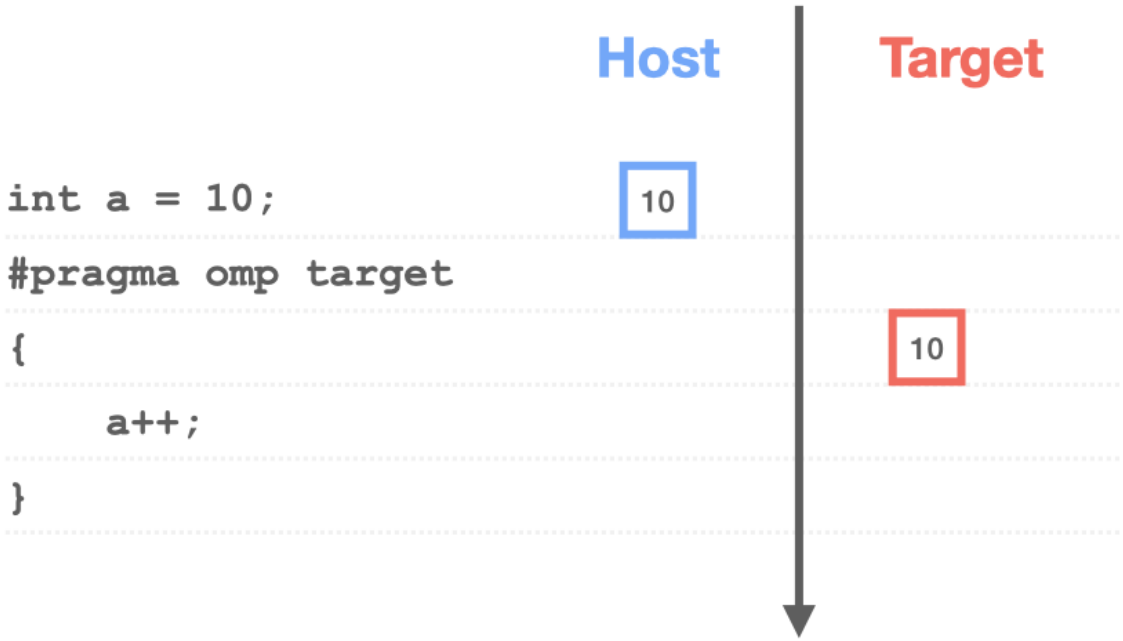

Implicit mapping - Scalars

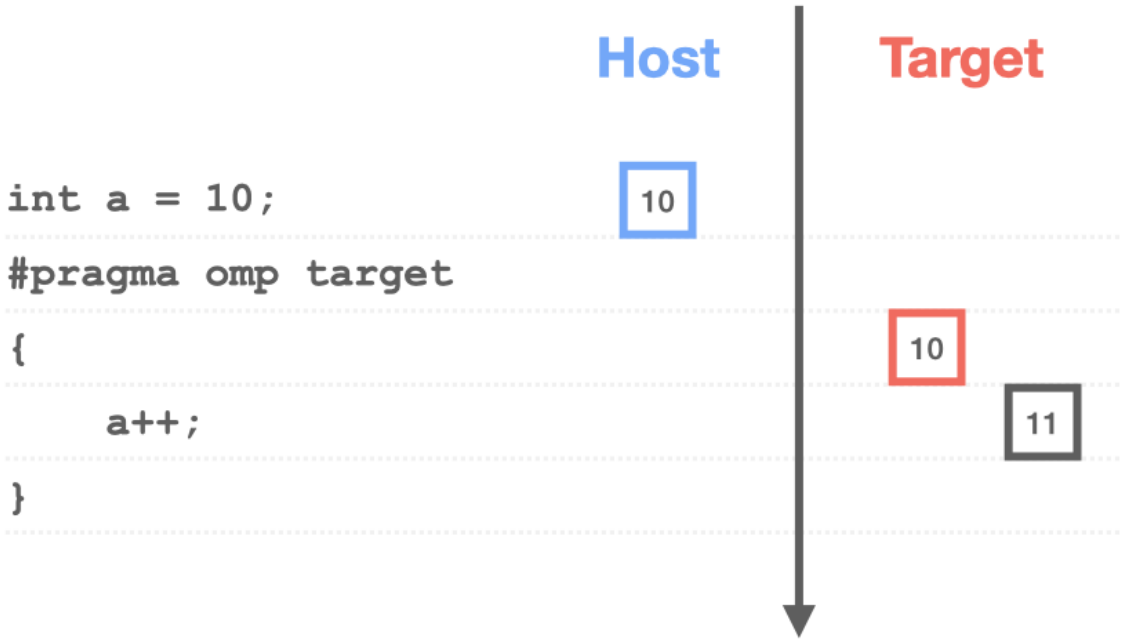

Implicit mapping - Scalars

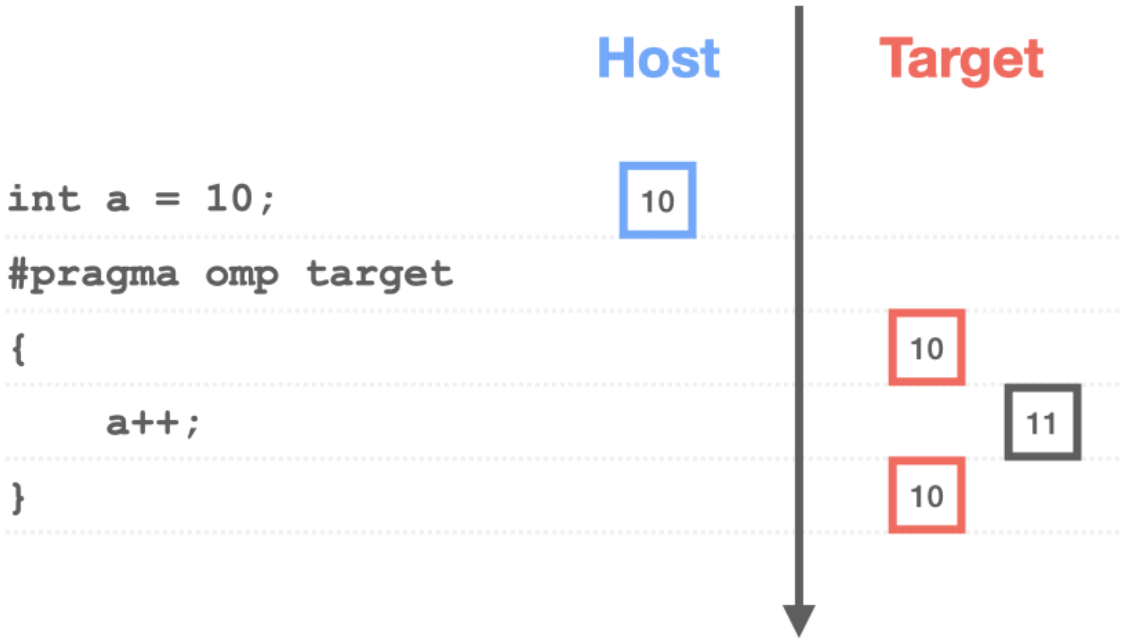

Implicit mapping - Scalars

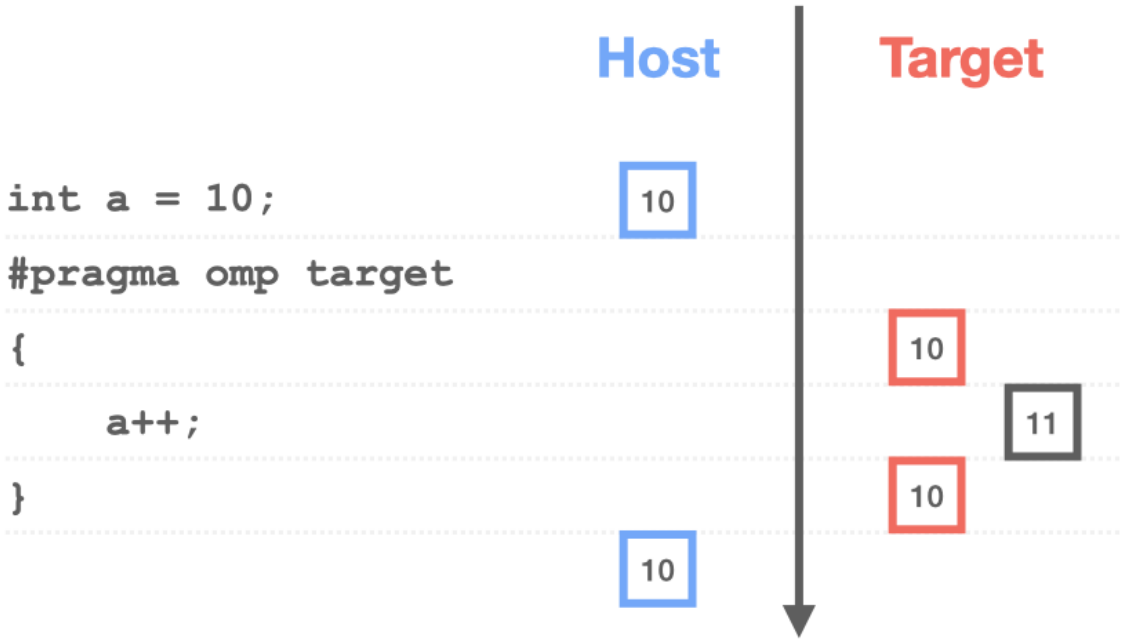

Data movements

- Relying on implicit mapping is risky:

Data movements

- Relying on implicit mapping is risky:
  - scalars / arrays: different default mapping by default

Data movements

- Relying on implicit mapping is risky:
  - scalars / arrays: different default mapping by default
  - defaults can change, it happened in OpenMP 4.5.

Data movements

- Relying on implicit mapping is risky:
  - scalars / arrays: different default mapping by default
  - defaults can change, it happened in OpenMP 4.5.
- Conversely, one can use explicit mapping to determine exactly how data is mapped to the target device.

Data movements

- Relying on implicit mapping is risky:
  - scalars / arrays: different default mapping by default
  - defaults can change, it happened in OpenMP 4.5.
- Conversely, one can use explicit mapping to determine exactly how data is mapped to the target device.
- This can be achieved using the `map` clause.

Data movements

- Relying on implicit mapping is risky:
  - scalars / arrays: different default mapping by default
  - defaults can change, it happened in OpenMP 4.5.
- Conversely, one can use explicit mapping to determine exactly how data is mapped to the target device.
- This can be achieved using the `map` clause.
- Unlike its CPU counterpart, there is no `default(none)` clause for target offloading, which would force every variable to be explicitly mapped.

Explicit mapping: `map` clause

```
1 #pragma omp target map(map-type: variable-name)
```

```
1 !$OMP TARGET MAP(map-type: variable-name)
2 !$OMP END TARGET
```

With the map-type being of the following:

- `to`
- `from`
- `tofrom`

Explicit mapping: `map` clause

```
1 #pragma omp target map(map-type: variable-name)
```

```
1 !$OMP TARGET MAP(map-type: variable-name)
2 !$OMP END TARGET
```

With the map-type being of the following:

- `to`
- `from`
- `tofrom`
- `alloc`
- `delete`

`map(tofrom:my_var)` clause value

`map(tofrom:my_var)` clause value

- Value of data before target is sent to the device for the target region.

`map(tofrom:my_var)` clause value

- Value of data before target is sent to the device for the target region.
- Final value of the variable in target region brought back to the device.

`map(to:my_var)` clause value

`map(to:my_var)` clause value

- This clause value is used when passing variables to target constructs that need a access to a host-initialised variable, and/or whose updates done on the device are not needed back on the host.

`map(to:my_var)` clause value

- This clause value is used when passing variables to target constructs that need a access to a host-initialised variable, and/or whose updates done on the device are not needed back on the host.
- The variable `my_var` begins the target region with the value it had before entering the region.

`map(to:my_var)` clause value

- This clause value is used when passing variables to target constructs that need a access to a host-initialised variable, and/or whose updates done on the device are not needed back on the host.
- The variable `my_var` begins the target region with the value it had before entering the region.
- The value of variable `my_var` will be left unchanged on the host after the target region.

map(to:my_var) clause value - Example

```
1 int a = 123;
2 int b;
3 #pragma omp target map(to:a,b)
4 {
5     a *= 2;
6     b = 100;
7 }
```

```
1 INTEGER :: a := 123
2 INTEGER :: b
3 !$OMP TARGET MAP(to:a,b)
4     a *= 2;
5     b = 100;
6 !$OMP END TARGET
```

`map(from:my_var)` clause value

- This clause value is used when passing variables to target constructs that will be initialised or update on the device, and the final value is needed back on the host.
- The variable `my_var` begins the target region with an uninitialised value.
- The value of variable `my_var` leaves the target region with the final value it had at the end of it.

map(from:my_var) clause value

```
1 int a;
2 int b = 123;
3 #pragma omp target map(from:a, b)
4 {
5     a = 123;
6 }
```

```
1 INTEGER :: a
2 INTEGER :: b := 123
3 !$OMP TARGET MAP(from:a, b)
4     a = 123
5 !$OMP END TARGET
```

Mapping of arrays

```
1 int* a = (int*)malloc(sizeof(int) * 10);
2 #pragma omp target map(tofrom:a[0:10])
3 {
4     a[5] = 123;
5 }
```

```
1 INTEGER, ALLOCATABLE :: a
2 ALLOCATE(a(0:9))
3 !$OMP TARGET MAP(tofrom:a(0:10))
4     a(5) = 123
5 !$OMP END TARGET
```

Mapping of arrays

- You can also map array variables.
- The notation requires to pass the interval of the array you want to map.[6]
- First, you pass the index of the first element to map.
- Second, you pass the **length** of the section to map.

Caution

You do not pass the start index and the end index, but the start index and the length of the array section to map.

[6] Because you can map a part of the array if you so wish.

Different map types

If you would like to use multiple types of mapping, you need to declare them in distinct `map` clauses.

```
1 int a = 123;
2 int b;
3 #pragma omp target map(to:a) map(from:b)
4 {
5     a *= 2;
6     b = 100;
7 }
```

```
1 INTEGER :: a := 123
2 INTEGER :: b
3 !$OMP TARGET MAP(to:a,b)
4     a *= 2;
5     b = 100;
6 !$OMP END TARGET
```

Time to practise: `3.DataMapping`

Update the source code provided to minimise the data movements of the different variables involved in calculations.

Tips

You will need:

- the `target` construct
- the `to` clause
- the `tofrom` clause
- the `from` clause

Table of Contents

Why bother about data environments?

Quick look at the bandwidths:

- Inside the GPU: 900GB/s
- Between the CPU and the GPU: 2GB/s (i.e.: 1 PCI-E)

Moral of the story: you want to move data as little as possible.

target data

`target data`

- The `target data` construct creates a data environment, with variables that will persist throughout the `target data` region.

`target data`

- The `target data` construct creates a data environment, with variables that will persist throughout the `target data` region.
- Data environments allow to avoid redundant data movements between `target` constructs.

`target data`

- The `target data` construct creates a data environment, with variables that will persist throughout the `target data` region.
- Data environments allow to avoid redundant data movements between `target` constructs.
- Said otherwise, `target` constructs enclosed in a `target data` construct inherit their data environment.

target data construct

```
1  int a = 0;
2  // Data moves to the device, stays until the end.
3  #pragma omp target data map(tofrom:a)
4  {
5      // Some host code
6      #pragma omp target map(tofrom:a)
7      {
8          a = 45;
9      }
10     printf("a = %d\n", a);
11     #pragma omp target map(tofrom:a)
12     {
13         a++;
14     }
15     // Some host code
16 }
17 printf("a = %d\n", a);
```

target data construct

```
1  INTEGER :: a := 0
2  ! Data moves to the device, stays until the end.
3  !$OMP TARGET DATA map(tofrom:a)
4      ! Some host code
5      !$OMP TARGET map(tofrom:a)
6          a = 45
7      !$OMP END TARGET
8      WRITE(*, '(A,I0)') "a = ", a
9      !$OMP TARGET map(tofrom:a)
10         a = a + 1
11     !$OMP END TARGET
12     ! Some host code
13 !$OMP END TARGET DATA
14 WRITE(*, '(A,I0)') "a = ", a
```

Arbitrary data environment management points

- Sometimes, we may want to use data environment in a less structured way.
- Typically happens when ones starts to have function calls, and potential orphan constructs.
- We may have at some point one variable that we know we need to pass to the device.
- We may want to issue the data transfer asynchronously.

Arbitrary data environment management points

```
1 #pragma omp target enter data map(alloc:a)
2 // Any code
3 #pragma omp target exit data map(delete:a)
```

```
1 !$OMP TARGET ENTER DATA MAP(alloc:a)
2 // Any code
3 !$OMP TARGET EXIT DATA MAP(delete:a)
```

Note:

- for `target enter data`, it maps data to the device, clause can be `to` or `alloc` only.
- for `target exit data`, it unmaps data from the device, clause can be `from` or `delete` only.
- `alloc` and `delete` allow to manipulate data on the device without requiring data movements.

`map(alloc:my_var)` clause value

- Allocates a variable on the device, does not issue any data movement between the host and the device.

map(alloc:my_var) clause value

```c
int a = 123;
#pragma omp target enter data map(alloc:a)
#pragma omp target
{
    a = 246;
    // Something with A.
}
```

```fortran
INTEGER :: a := 123
!$OMP TARGET ENTER DATA MAP(alloc:a)
!$OMP TARGET
    a = 246
    ! Something with A.
!$OMP END TARGET
```

`map(delete:my_var)` clause value

- Deletes a variable on the device, does not issue any data movement between the host and the device.

map(delete:my_var) clause value

```
1 // Some code...
2 #pragma omp target exit data map(delete:a)
```

```
1 ! Some code...
2 !$OMP TARGET EXIT DATA MAP(delete:a)
```

`target update` construct

- Data environment perform data movements at the start and at the end.
- However, sometimes, we may want to update the host copy from the device copy, or vice-versa, in the middle.
- This is where the update construct comes in.

`target update` construct

The `printf` at line 8 relies on an outdated host copy of variable `a`.

```
1  int a;
2  #pragma omp target data map(from:a)
3  {
4      #pragma omp target map(from:a)
5      {
6          a = 123;
7      }
8      printf("a = %d\n", a);
9      #pragma omp target map(from:a)
10     {
11         a++;
12     }
13 }
```

target update construct

The targetupdate construct will fetch its value from the device.

```
1  int a;
2  #pragma omp target data map(from:a)
3  {
4      #pragma omp target map(from:a)
5      {
6          a = 123;
7      }
8      #pragma omp target update from(a)
9      printf("a = %d\n", a);
10     #pragma omp target map(from:a)
11     {
12         a++;
13     }
14 }
```

target update construct

Conversely, we may want, punctually, to update the value of a variable mapped to a device.

```
1  int a = 123;
2  int result = 0;
3  #pragma omp target data map(to:a) map(from:result)
4  {
5      #pragma omp target map(to:a) map(from:result)
6      {
7          result +=  a * 100;
8      }
9      a++;
10     #pragma omp target map(to:a) map(from:result)
11     {
12         result +=  a * 100;
13     }
14 }
```

target update construct

```
1  int a = 123;
2  int result = 0;
3  #pragma omp target data map(to:a) map(from:result)
4  {
5      #pragma omp target map(to:a) map(from:result)
6      {
7          result +=  a * 100;
8      }
9      a++;
10     #pragma omp target update to(a)
11     #pragma omp target map(to:a) map(from:result)
12     {
13         result +=  a * 100;
14     }
15 }
16
```

Time to practise: `4.DataEnvironment`

Tips

You will need the following:

- `target enter data`
- `target exit data`
- `target map`
- `allocate`

Table of Contents

Hardware

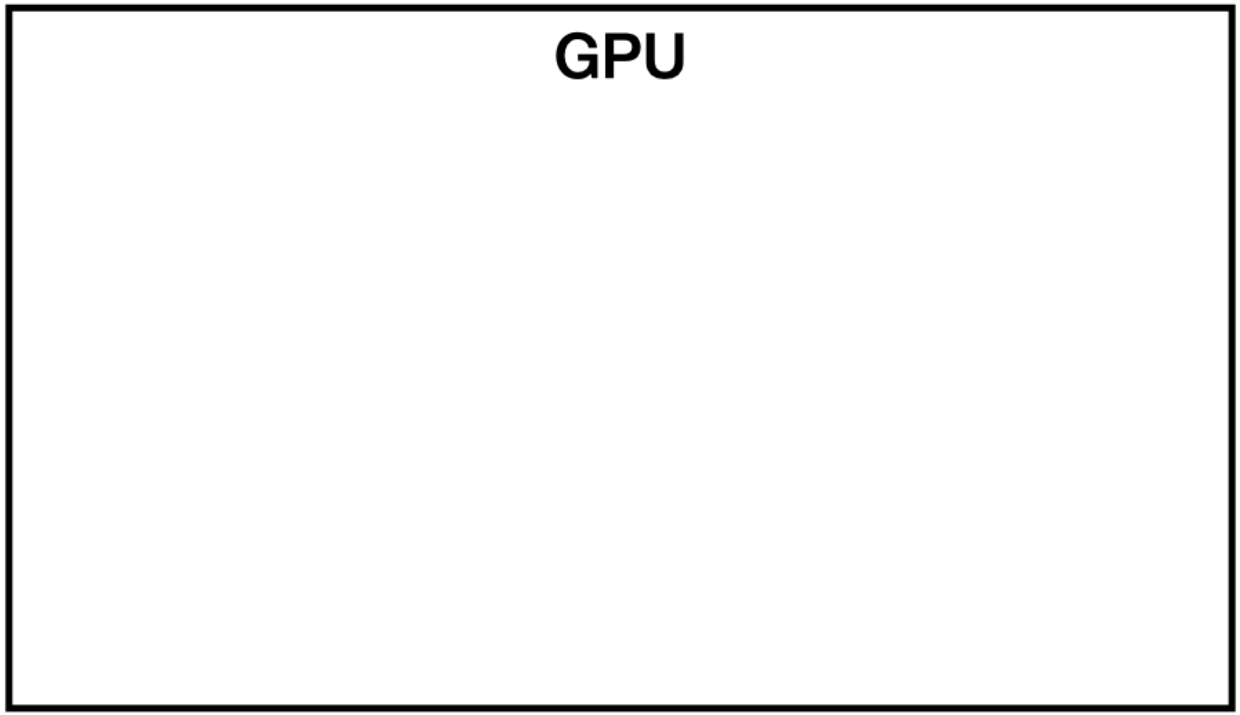

Hardware

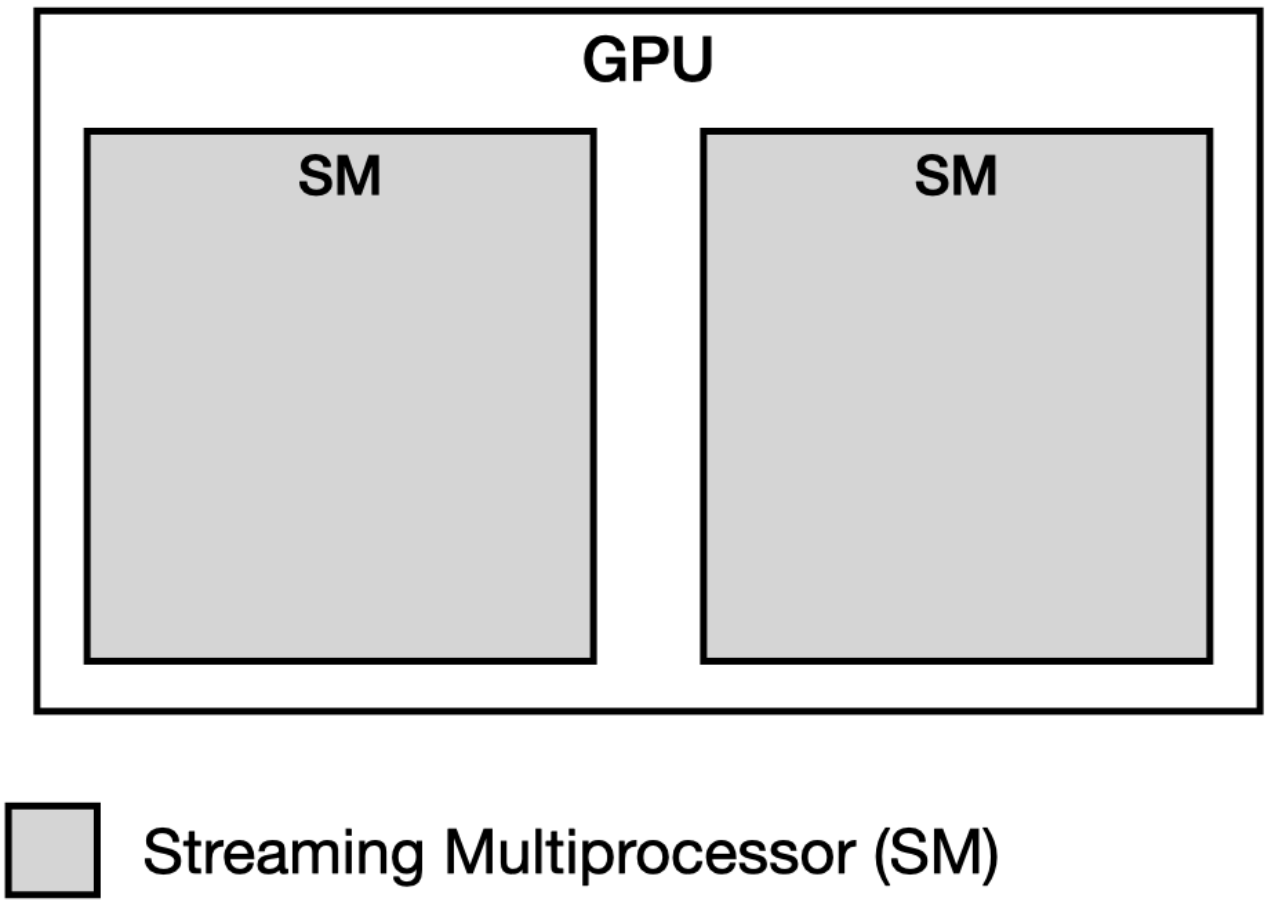

Hardware

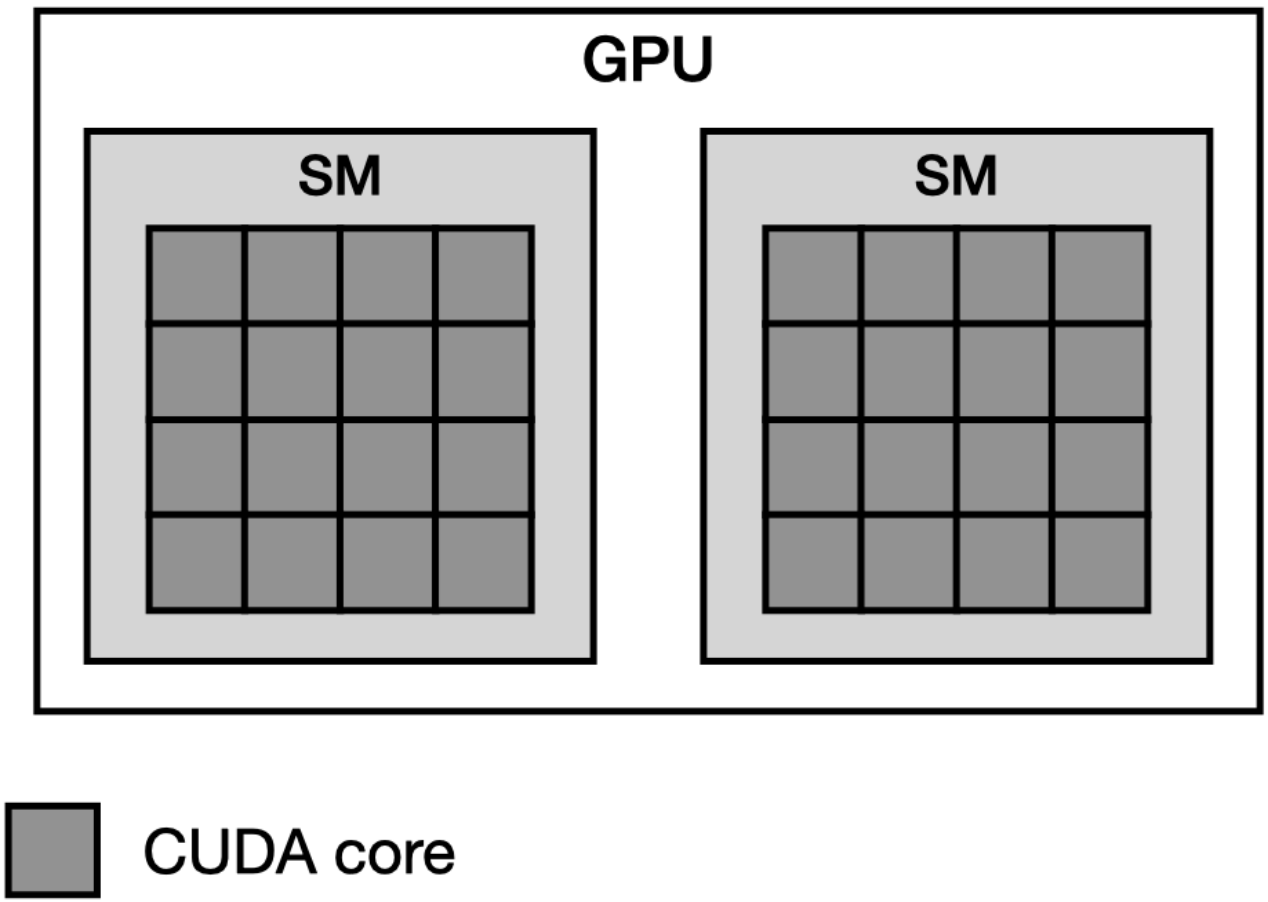

teams

- The `teams` construct forks the execution into multiple teams of threads.

teams

- The `teams` construct forks the execution into multiple teams of threads.
- The code in the associated structured block will be executed by the master thread of each team.

`teams`

- The `teams` construct forks the execution into multiple teams of threads.
- The code in the associated structured block will be executed by the master thread of each team.
- Each team will therefore has its own "thread 0".

teams

Similarly to CPU OpenMP, you can:

- get your team number, using the `omp_get_team_num()` function.
- get the total number of teams, using the `omp_get_num_teams()` function.

teams

```
1  #pragma omp target
2  {
3      #pragma omp teams
4      {
5          int my_team_number = omp_get_team_num();
6          printf("I am the master thread for team %d.\n",
       my_team_number);
7      }
8  }
```

```
1  !$OMP TARGET
2      !$OMP TEAMS
3          int my_team_number = omp_get_team_num();
4          printf("I am the master thread for team %d.\n",
       my_team_number);
5      !$OMP END TEAMS
6  !$OMP END TARGET
```

Composite `target teams` construct

```
1 #pragma omp target teams
2 {
3     int my_team_number = omp_get_team_num();
4     printf("I am the master thread for team %d.\n",
      my_team_number);
5 }
```

```
1 !$OMP TARGET TEAMS
2     int my_team_number = omp_get_team_num();
3     printf("I am the master thread for team %d.\n",
      my_team_number);
4 !$OMP END TARGET TEAMS
```

`teams`

- Similarly to the `parallel` construct, you can control the number of teams spawned.

teams

- Similarly to the `parallel` construct, you can control the number of teams spawned.
- Use the `num_teams` clause.

teams

- Similarly to the `parallel` construct, you can control the number of teams spawned.
- Use the `num_teams` clause.
- When used, this clause indicates that **at most** `<X>` teams will be spawned.

teams

```
1 #pragma omp target teams num_teams(X)
2 {
3     int my_team_number = omp_get_team_num();
4     printf("I am the master thread for team %d.\n",
      my_team_number);
5 }
```

```
1 !$OMP TARGET NUM_TEAMS(X)
2     int my_team_number = omp_get_team_num();
3     printf("I am the master thread for team %d.\n",
      my_team_number);
4 !$OMP END TARGET
```

teams

- Inside your `teams` construct, by default you just have the master thread running, just like in classic serial program execution.

`teams`

- Inside your `teams` construct, by default you just have the master thread running, just like in classic serial program execution.
- When wanting to fork threads in, you can use the `parallel` construct.

teams

```
1  #pragma omp target
2  {
3      // One team of one thread
4      #pragma omp teams
5      {
6          // Multiple teams of one thread
7          #pragma omp parallel
8          {
9              // Multiple teams of multiple threads
10         }
11     }
12 }
```

Loops!

```
1 for(int i = 0; i < N; i++)
2 {
3     a[i] = i + 123;
4 }
```

```
1 DO i = 0, N - 1
2     a(i) = i + 123
3 END DO
```

Loops!

```
1 #pragma omp target teams
2 {
3     for(int i = 0; i < N; i++)
4     {
5         a[i] = i + 123;
6     }
7 }
```

```
1 !$OMP TARGET TEAMS
2     DO i = 0, N - 1
3         a(i) = i + 123
4     END DO
5 !$OMP END TARGET TEAMS
```

Multi-level parallelism

- Just like in a `parallel` construct, threads do not automatically split iterations unless they are asked to.

Multi-level parallelism

- Just like in a `parallel` construct, threads do not automatically split iterations unless they are asked to.
- The `distribute` clause is equivalent to the `for` / `do` clause for CPUs.

Multi-level parallelism

- Just like in a `parallel` construct, threads do not automatically split iterations unless they are asked to.
- The `distribute` clause is equivalent to the `for` / `do` clause for CPUs.
- It allows the iteration set to be split across the master threads of all teams.

Loops!

```
1 #pragma omp target teams
2 {
3     #pragma omp distribute
4     for(int i = 0; i < N; i++)
5     {
6         a[i] = i + 123;
7     }
8 }
```

```
1 !$OMP TARGET TEAMS
2     !$OMP DISTRIBUTE
3     DO i = 0, N - 1
4         a(i) = i + 123
5     END DO
6 !$OMP END TARGET TEAMS
```

Composite `target teams distribute` construct

```
1 #pragma omp target teams distribute
2 for(int i = 0; i < N; i++)
3 {
4     a[i] = i + 123;
5 }
```

```
1 !$OMP TARGET TEAMS DISTRIBUTE
2 DO i = 0, N - 1
3     a(i) = i + 123
4 END DO
```

Scheduling

- On CPUs, the `schedule` clause is available.

Scheduling

- On CPUs, the `schedule` clause is available.
- On GPUs, the `dist_schedule` clause is available.

Scheduling

- On CPUs, the `schedule` clause is available.
- On GPUs, the `dist_schedule` clause is available.
- Only `static` schedule available however.

Scheduling

- On CPUs, the `schedule` clause is available.
- On GPUs, the `dist_schedule` clause is available.
- Only `static` schedule available however.
- Can still specify a particular `chunksize`.

Scheduling

```
1 #pragma omp target teams distribute
2 for(int i = 0; i < N; i++)
3 {
4     a[i] = i + 123;
5 }
```

```
1 !$OMP TARGET TEAMS DISTRIBUTE
2 DO i = 0, N - 1
3     a(i) = i + 123
4 END DO
```

Scheduling

```
1 #pragma omp target teams distribute dist_schedule(static
      , 24)
2 for(int i = 0; i < N; i++)
3 {
4     a[i] = i + 123;
5 }
```

```
1 !$OMP TARGET TEAMS DISTRIBUTE DIST_SCHEDULE(STATIC, 24)
2 DO i = 0, N - 1
3     a(i) = i + 123
4 END DO
```

Nested loops

```
1 for(int i = 0; i < N; i++)
2 {
3     for(int j = 0; j < N; j++)
4     {
5         a[i][j] = i * j + 123;
6     }
7 }
```

```
1 DO i = 0, N - 1
2     DO j = 0, j < N - 1
3         a(i,j) = i * j + 123
4     END DO
5 END DO
```

Nested loops

```
1 #pragma omp target teams distribute
2 for(int i = 0; i < N; i++)
3 {
4     for(int j = 0; j < N; j++)
5     {
6         a[i][j] = i * j + 123;
7     }
8 }
```

```
1 !$OMP TARGET TEAMS DISTRIBUTE
2 DO i = 0, N - 1
3     DO j = 0, j < N - 1
4         a(i,j) = i * j + 123
5     END DO
6 END DO
```

Nested loops

```
1 for(int i = 0; i < N; i++)
2 {
3     #pragma omp target teams distribute
4     for(int j = 0; j < N; j++)
5     {
6         a[i][j] = i * j + 123;
7     }
8 }
```

```
1 DO i = 0, N - 1
2     !$OMP TARGET TEAMS DISTRIBUTE
3     DO j = 0, j < N - 1
4         a(i,j) = i * j + 123
5     END DO
6 END DO
```

Nested loops

```
1 #pragma omp target teams distribute
2 for(int i = 0; i < N; i++)
3 {
4     #pragma omp parallel for
5     for(int j = 0; j < N; j++)
6     {
7         a[i][j] = i * j + 123;
8     }
9 }
```

```
1 !$OMP TARGET TEAMS DISTRIBUTE
2 DO i = 0, N - 1
3     !$OMP PARALLEL DO
4     DO j = 0, j < N - 1
5         a(i,j) = i * j + 123
6     END DO
7 END DO
```

Massive loop

```
1 for(int i = 0; i < SUPER_BIG_N; i++)
2 {
3     a[i] = i + 123;
4 }
```

```
1 DO i = 0, SUPER_BIG_N - 1
2     a(i) = i + 123
3 END DO
```

Massive loop

```
1 #pragma omp target teams distribute parallel for
2 for(int i = 0; i < SUPER_BIG_N; i++)
3 {
4     a[i] = i + 123;
5 }
```

```
1 !$OMP TARGET TEAMS DISTRIBUTE PARALLEL DO
2 DO i = 0, SUPER_BIG_N - 1
3     a(i) = i + 123
4 END DO
```

Time to practise: 5.MultiLevelParallelism

Tips

You will need:

- the `target` construct
- the `teams` construct
- the `distribute` construct
- the `parallel` construct
- the `for` construct

Table of Contents

Multi-GPU

- Nowadays, GPU nodes contain more than 1 GPU.

Multi-GPU

- Nowadays, GPU nodes contain more than 1 GPU.
- On the AWS node, there are 4 GPUs.

Multi-GPU

- Nowadays, GPU nodes contain more than 1 GPU.
- On the AWS node, there are 4 GPUs.
- Fully using a GPU is great.

Multi-GPU

- Nowadays, GPU nodes contain more than 1 GPU.
- On the AWS node, there are 4 GPUs.
- Fully using a GPU is great.
- Fully using all GPUs is more fun.

Find our way in a multi-GPU environment.

Find our way in a multi-GPU environment.

- Implicitly, one of the devices is selected by default.

Find our way in a multi-GPU environment.

- Implicitly, one of the devices is selected by default.
- You can know which one, by calling `omp_get_default_device()`.

Find our way in a multi-GPU environment.

- Implicitly, one of the devices is selected by default.
- You can know which one, by calling `omp_get_default_device()`.
- You can change which one, by calling `omp_set_default_device(new_device)`.

Devices number and identifiers

[7] Does not seem to be supported yet, as of `nvhpc/22.2`.

Devices number and identifiers

- The host is a device too, you can get its number by calling `omp_get_initial_device()`.

[7] Does not seem to be supported yet, as of `nvhpc/22.2`.

Devices number and identifiers

- The host is a device too, you can get its number by calling `omp_get_initial_device()`.
- You can get the number of the device you are currently running on by calling `omp_get_device_num()`.[7]

[7] Does not seem to be supported yet, as of `nvhpc/22.2`.

Devices number and identifiers

- The host is a device too, you can get its number by calling `omp_get_initial_device()`.
- You can get the number of the device you are currently running on by calling `omp_get_device_num()`.[7]
- You can get the total number of target devices by calling `omp_get_num_devices()`.

[7] Does not seem to be supported yet, as of `nvhpc/22.2`.

Devices number and identifiers

- The host is a device too, you can get its number by calling `omp_get_initial_device()`.
- You can get the number of the device you are currently running on by calling `omp_get_device_num()`.[7]
- You can get the total number of target devices by calling `omp_get_num_devices()`.

Note

`omp_get_num_devices()` returns the number of **target devices**, it does **not** include the host device.

[7] Does not seem to be supported yet, as of `nvhpc/22.2`.

Multi-GPU work split

- Send different parts of an array to different GPUs.

Multi-GPU work split

- Send different parts of an array to different GPUs.
- Have multiple GPUs work in parallel on a given arrray.

Addressing GPUs - The `device` clause

Addressing GPUs - The `device` clause

- Can ask for a target construct to be executed on a specific device.

Addressing GPUs - The `device` clause

- Can ask for a target construct to be executed on a specific device.
- Just use the `device` clause on a `target` construct.

Addressing GPUs - The `device` clause

- Can ask for a target construct to be executed on a specific device.
- Just use the `device` clause on a `target` construct.

```
1 // This will execute on the first GPU, device #0.
2 #pragma omp target device(0)
3 {
4     // Code to execute on device 0
5 }
```

```
1 ! This will execute on the first GPU, device #0.
2 !$OMP TARGET DEVICE(0)
3     ! Code to execute on device 0
4 !$OMP END TARGET
```

Time to practise: 6.MultiGPU

Guidance

- Can ask for a target construct to be executed on a specific device.
- Try to get 2 GPUs, and get a target construct executed on each.
- Also try to get a target construct executed on the host device.
- Also try to get a target construct executed on a device that does not exist.

Table of Contents

Asynchronous

- By default, target constructs are synchronous.

Asynchronous

- By default, target constructs are synchronous.
- You cannot process beyond it until it is finished.

Asynchronous

- By default, target constructs are synchronous.
- You cannot process beyond it until it is finished.
- The `nowait` clause allows to execute a target construct in an asynchronous fashion.

Asynchronous

```
1 #pragma omp target nowait
2 {
3       // ...
4 }
```

```
1 !$OMP TARGET NOWAIT
2       // ...
3 !$OMP END TARGET NOWAIT
```

Asynchronous

```
1 #pragma omp target nowait
2 {
3     // ...
4 }
```

```
1 !$OMP TARGET NOWAIT
2     // ...
3 !$OMP END TARGET NOWAIT
```

- When using the `nowait` clause, execution does not wait for the end of the `target` construct to continue its progress.

Asynchronous

```
1 #pragma omp target nowait
2 {
3     // ...
4 }
```

```
1 !$OMP TARGET NOWAIT
2     // ...
3 !$OMP END TARGET NOWAIT
```

- When using the `nowait` clause, execution does not wait for the end of the `target` construct to continue its progress.
- Dependencies between `target` constructs are expressed using the `depend` clause.

Asynchronous

```
1 #pragma omp target nowait
2 {
3     // ...
4 }
```

```
1 !$OMP TARGET NOWAIT
2     // ...
3 !$OMP END TARGET NOWAIT
```

- When using the `nowait` clause, execution does not wait for the end of the `target` construct to continue its progress.
- Dependencies between `target` constructs are expressed using the `depend` clause.
- If a `target` construct with a `nowait` clause does not have `depend` clauses, it is assumed to have no dependencies and can execute freely at any time.

Asynchronous

`depend(in:a)` I will read from variable `a`.

Asynchronous

`depend(in:a)` I will read from variable `a`.

`depend(out:a)` I will write to variable `a`.

Asynchronous

`depend(in:a)`	I will read from variable `a`.
`depend(out:a)`	I will write to variable `a`.
`depend(inout:a)`	I will read from, and write to, variable `a`.

Structured synchronisation

```
1  #pragma omp target nowait map(from:a) depend(out:a)
2  {
3      a = 456;
4  }
5  #pragma omp target nowait map(to:a) map(from:b) \
6                           depend(in:a) depend(out:b)
7  {
8      b = a + 3;
9  }
10 #pragma omp target map(to:a) map(from:c) \
11                          depend(in:a) depend(out:c)
12 {
13     c = a + 10
14 }
```

Structured synchronisation

```
1  !$OMP TARGET NOWAIT MAP(FROM:a) DEPEND(OUT:a)
2      a = 456;
3  !$OMP END TARGET
4  !$OMP TARGET NOWAIT MAP(TO:a) MAP(FROM:b) &
5                      DEPEND(IN:a) DEPEND(OUT:b)
6      b = a + 3;
7  !$OMP END TARGET
8  !$OMP TARGET MAP(FROM:a) MAP(TO:c) &
9                DEPEND(IN:a) DEPEND(OUT:c)
10     c = a + 10
11 !$OMP END TARGET
```

Construct synchronisation[8]

```
1  #pragma omp single
2  {
3      #pragma omp target nowait map(from:a) depend(out:a)
4      {
5          a = 456;
6      }
7      #pragma omp target nowait map(to:a) map(from:b) \
8                              depend(in:a) depend(out:b)
9      {
10         b = a + 3;
11     }
12     #pragma omp target nowait map(to:a) map(from:c) \
13                             depend(in:a) depend(out:c)
14     {
15         c = a + 10
16     }
17 }
```

Construct synchronisation[9]

```
1  !$OMP SINGLE
2  !$OMP TARGET NOWAIT MAP(FROM:a) DEPEND(OUT:a)
3      a = 456;
4  !$OMP END TARGET
5  !$OMP TARGET NOWAIT MAP(TO:a) MAP(FROM:b) &
6                      DEPEND(IN:a) DEPEND(OUT:b)
7      b = a + 3;
8  !$OMP END TARGET
9  !$OMP TARGET NOWAIT MAP(FROM:a) MAP(TO:c) &
10                 DEPEND(IN:a) DEPEND(OUT:c)
11     c = a + 10
12 !$OMP END TARGET
13 !$OMP END SINGLE
```

[9] Example assumes everything in `parallel` construct

Explicit synchronisation

```
1  #pragma omp target nowait map(from:a) depend(out:a)
2  {
3      a = 456;
4  }
5  #pragma omp target nowait map(to:a) map(from:b) \
6                      depend(in:a) depend(out:b)
7  {
8      b = a + 3;
9  }
10 #pragma omp target nowait map(to:a) map(from:c) \
11                     depend(in:a) depend(out:c)
12 {
13     c = a + 10
14 }
15 #pragma omp taskwait
```

Explicit synchronisation

```
1  !$OMP TARGET NOWAIT MAP(FROM:a) DEPEND(OUT:a)
2      a = 456;
3  !$OMP END TARGET
4  !$OMP TARGET NOWAIT MAP(TO:a) MAP(FROM:b) &
5                      DEPEND(IN:a) DEPEND(OUT:b)
6      b = a + 3;
7  !$OMP END TARGET
8  !$OMP TARGET NOWAIT MAP(FROM:a) MAP(TO:c) &
9                DEPEND(IN:a) DEPEND(OUT:c)
10     c = a + 10
11 !$OMP END TARGET
12 !$OMP TASKWAIT
```

Time to practise: `7.Asynchronous`

Execute the different statements shown in different `target` constructs, using asynchronous execution.

Tips

You will need:

- `target` construct
- `nowait` clause
- `depend` clause
- `map` clause

Table of Contents

Summary

You now know how to:

- offload your workload to GPUs.
- control the mapping of data between host and targets.
- handle multiple GPUs.
- issue kernels in asynchronous fashion.

To be seen

- Reverse offloading
- Coalescing memory accesses.
- Metadirectives
- Directive cancellation
- Conditional execution
- And a lot, lot more...

To be seen

- Reverse offloading
- Coalescing memory accesses.
- Metadirectives
- Directive cancellation
- Conditional execution
- And a lot, lot more...

The best place to learn more about OpenMP and how it works, to get the specifications and so on is the OpenMP forum website.

Your opinion matters

Figure: Link to feedback survey

Feel free to connect

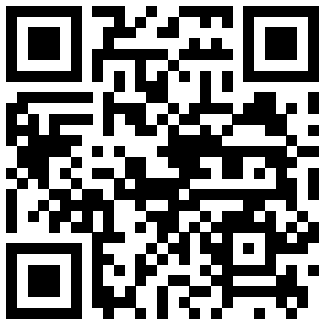

Figure: Link to LinkedIn profile