

27th POP Webinar

Performance Analysis of OpenMP Target Offloading in Score-P

2024-05-28 | Jan André Reuter



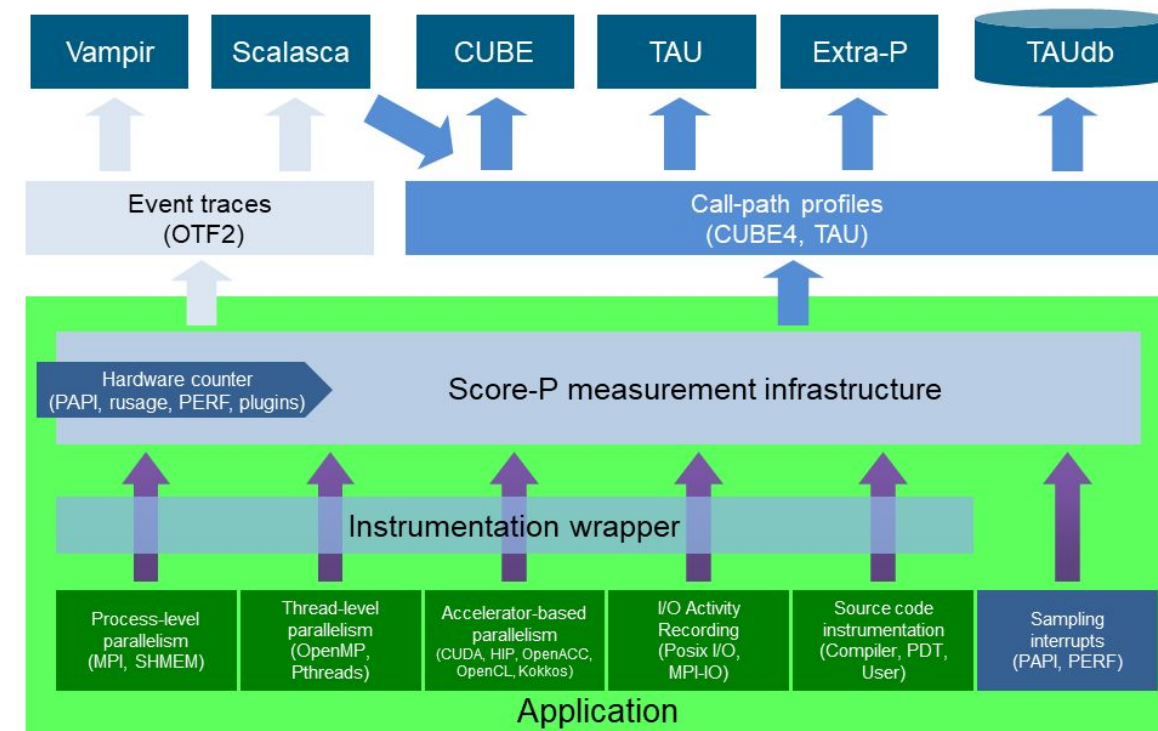
AGENDA

- What is the Score-P infrastructure?
- The OpenMP Tools Interface and our support of it
- OpenMP and offloading
- Handling offloading events on the host
- Handling offloading events from accelerators
- Results
- What if my runtime has only limited support?
- Final words

WHAT IS SCORE-P?

The Score-P instrumentation and measurement infrastructure

- Score-P is a highly scalable **instrumentation** tool
- Support for multi-process, thread-parallel and accelerator-based paradigms
- Support for additional metrics (I/O, HW counters, ...)
- Flexible **measurement** without re-compilation:
 - **Profile** generation (CUBE4 .cubex format)
 - Event **trace** recording (OTF2 format)
- Support for C, C++, Fortran and Python



INSTRUMENTING APPLICATIONS

A very high level overview on how to use Score-P



```
~/Sources/OpenMP/hello-world scorep-nvc -mp=multicore hello-world.c
~/Sources/OpenMP/hello-world OMP_NUM_THREADS=4 ./a.out
Hello World from thread 1
Hello World from thread 0
Hello World from thread 2
Hello World from thread 3
~/Sources/OpenMP/hello-world scorep-score -r scorep-*/profile.cubex
```

```
Estimated aggregate size of event trace:          790 bytes
Estimated requirements for largest trace buffer (max_buf): 790 bytes
Estimated memory requirements (SCOREP_TOTAL_MEMORY): 11MB
(hint: When tracing set SCOREP_TOTAL_MEMORY=11MB to avoid intermediate flushes
or reduce requirements using USR regions filters.)
```

flt	type	max_buf[B]	visits	time[s]	time[%]	time/visit[us]	region
	ALL	789	22	0.00	100.0	47.58	ALL
	OMP	628	16	0.00	45.4	29.68	OMP
	COM	120	5	0.00	28.8	60.20	COM
	SCOREP	41	1	0.00	25.9	270.87	SCOREP
	OMP	340	4	0.00	33.8	88.49	!\$omp parallel @hello-world.c:14
	OMP	96	4	0.00	2.2	5.80	!\$omp implicit barrier @hello-world.c:14
	COM	96	4	0.00	0.7	1.81	hello_world
	OMP	96	4	0.00	6.6	17.31	!\$omp critical @hello-world.c:6
	OMP	96	4	0.00	2.7	7.14	!\$omp critical sblock @hello-world.c:6
	SCOREP	41	1	0.00	25.9	270.87	a.out
	COM	24	1	0.00	28.1	293.74	main

```
~/Sources/OpenMP/hello-world # Create filter if needed
~/Sources/OpenMP/hello-world OMP_NUM_THREADS=4 \
SCOREP_ENABLE_TRACING=true \
SCOREP_FILTERING_FILE=initial_scorep.filter \
SCOREP_TOTAL_MEMORY=11MB \
./a.out
```

```
Hello World from thread 0
Hello World from thread 1
Hello World from thread 3
Hello World from thread 2
```

Build systems like CMake may need additional steps. See [scorep-wrapper](#) `--help` for more info.

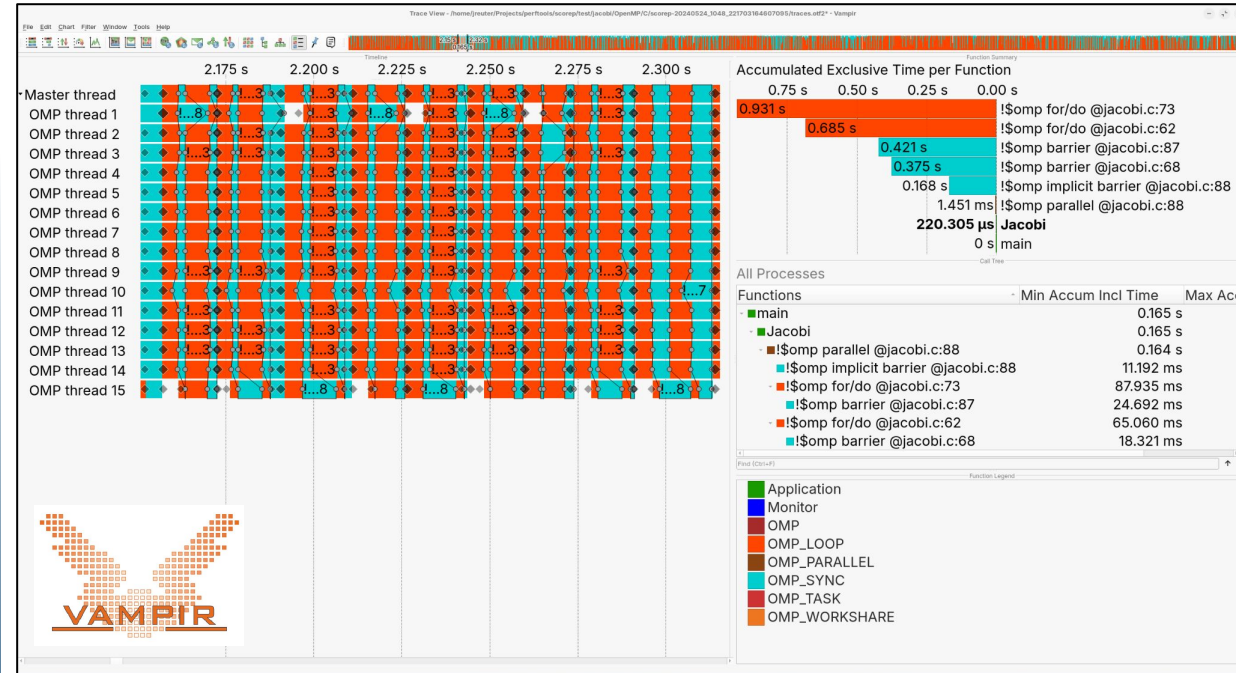
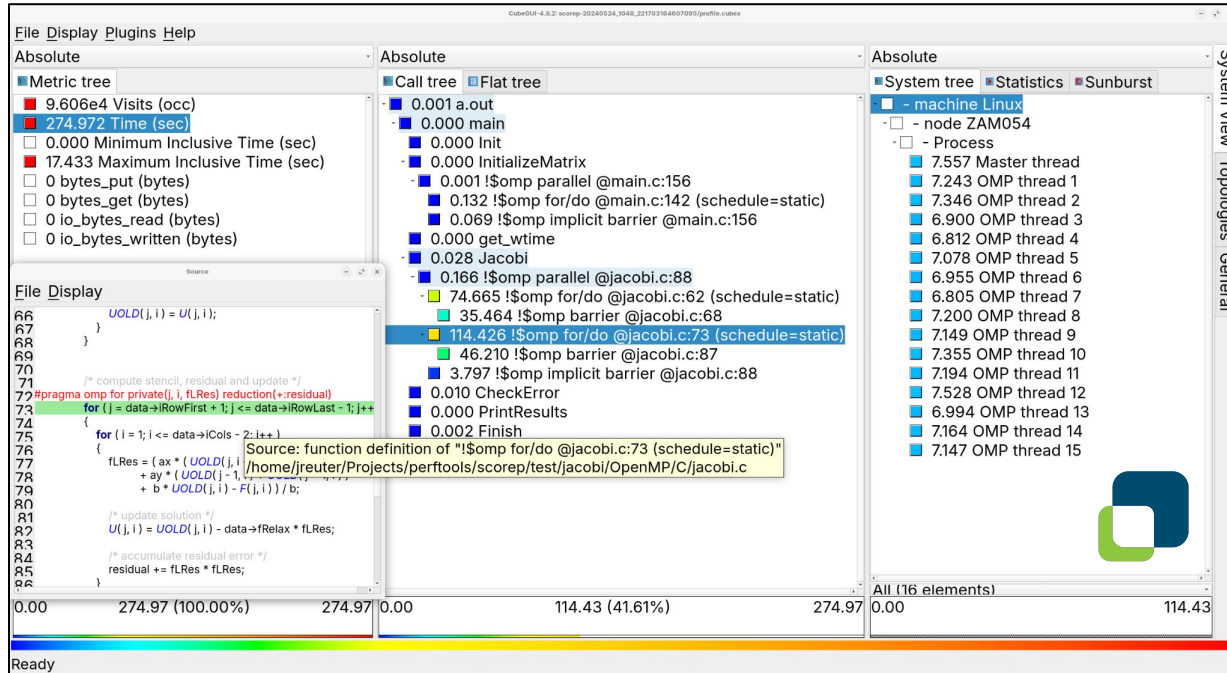
More information



Using POP Tools: Score-P, Scalasca
Bernd Mohr

VIEW YOUR RESULTS

A short look at Cube and Vampir



- Open source viewer for .cubex profiles
- More information:
<https://www.scalasca.org/scalasca/software/cube-4.x/>

- Commercial viewer for .otf2 traces
- More information:
<https://vampir.eu>

OPENMP SUPPORT IN SCORE-P

Two ways to collect information about OpenMP



- Source-to-source instrumentation tool
- Independent from compiler used
- Instrumentation up to OpenMP 3.x
- Various limitations
 - Code sometimes has to be prepared for OPARI2

OPENMP SUPPORT IN SCORE-P

Two ways to collect information about OpenMP



- Source-to-source instrumentation tool
- Independent from compiler used
- Instrumentation up to OpenMP 3.x
- Various limitations
 - Code sometimes has to be prepared for OPARI2



- Standardised tool interface since OpenMP 5.0
- Enables development of tools using any implementation of the OpenMP API
- Support for the latest and greatest OpenMP features
- Continuously expanded with new versions

WHAT IS THE OPENMP TOOLS INTERFACE?

- Interface for *first-party* tools, linked or loaded into the OpenMP program

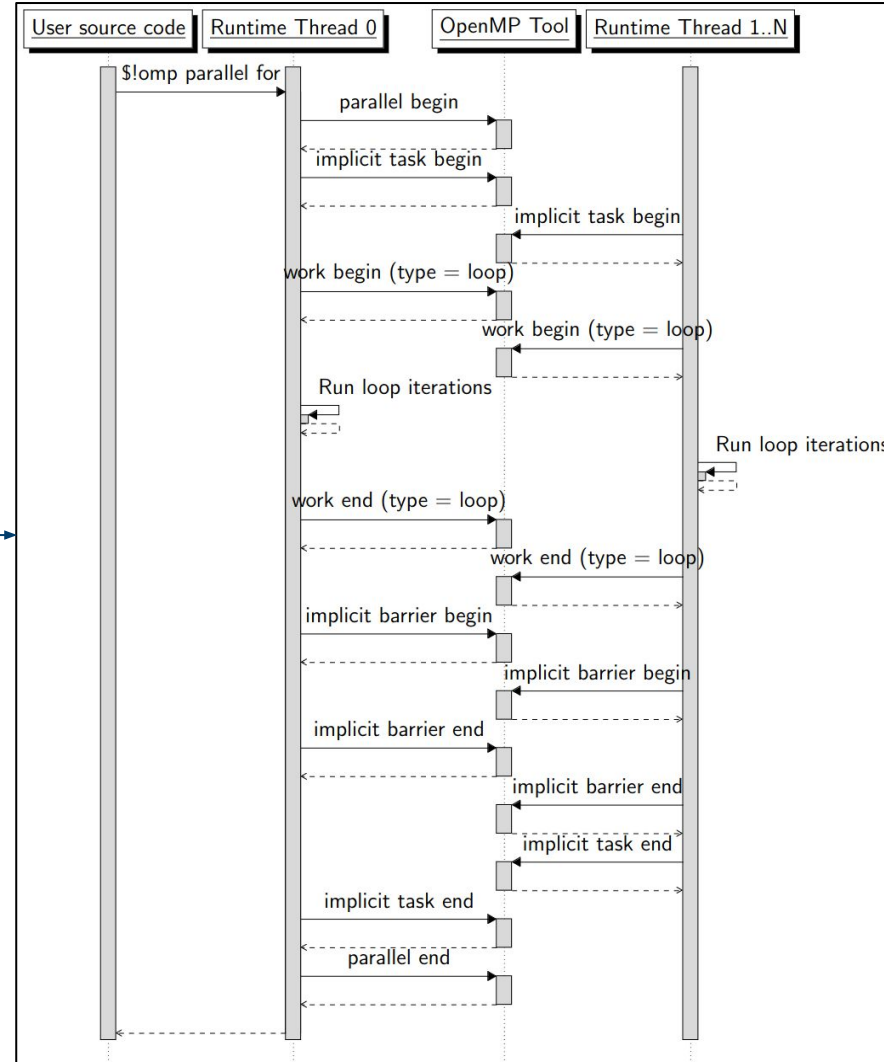
```
$ clang -fopenmp my-program.c -lmy-tool  
$ clang -fopenmp my-program.c my-tool.c  
$ OMP_TOOL_LIBRARIES=my-tool.so ./my-program
```

- Defined by the OpenMP standard, implemented by OpenMP runtimes
- Tools have to implement functions to interact with OMPT (e.g. `ompt_start_tool`)
- One interaction method: *callbacks*, invoked for runtime events

OPENMP TOOLS INTERFACE CALLBACKS

Just a part of what tools see for user code

```
int main( void )
{
    const int N = 100;
    #pragma omp parallel for
    for( int i = 0; i < N; ++i)
    {
        // Some code...
    }
    return 0;
}
```



- Tools receive a lot of events for user code
- With this, we are able to record events for OpenMP directives

SCORE-P AND OMPT: THE PRESENT ...



- First support in version 8.0 (released in December 2022)
- Selectable via `scorep --thread=omp:ompt`
- Support tried to match the available features in OPARI2 (focused on OpenMP 3.x)
- Small feature additions and several bug fixes in 8.x:
 - Recording loop schedules
 - Recording `omp_test_lock` events

OpenMP directive	Support
OpenMP 3.x	yes
cancel	no
task depend	no
task detach	no
taskgroup	yes
taskloop	no
teams	no
scope	no

... AND THE FUTURE (V9.0)

- Current OpenMP features planned for Score-P v9.0
 - Improvements to support of task directives

```

#include <stdio.h>

void foo() {
    int x = 0, y = 2;
    #pragma omp task depend( inout: x ) shared( x )
    x++;
    // 1st child task
    #pragma omp task shared( y )
    y--;
    // 2nd child task
    #pragma omp taskwait depend( in: x )
    // 1st taskwait
    printf( "x=%d\n", x );
    #pragma omp taskwait
    // 2nd taskwait
    printf( "y=%d\n", y );
}

int main() {
    #pragma omp parallel
    #pragma omp single
    foo();
    return 0;
}
    
```

```

- 0.001 a.out
- 0.000 main
- 0.026 !$omp parallel @tasking.c:18
- 0.000 !$omp single @tasking.c:17
- 0.000 !$omp single sblock @tasking.c:17
- 0.000 foo
  0.000 !$omp create task @tasking.c:5
  0.000 !$omp create task @tasking.c:7
- 0.000 !$omp taskwait @tasking.c:10
  0.000 !$omp task @tasking.c:7
  0.000 !$omp taskwait @tasking.c:12
- 0.217 !$omp barrier @tasking.c:17
  0.000 !$omp task @tasking.c:5
  0.223 !$omp implicit barrier @tasking.c:18
    
```

OpenMP directive	Support
OpenMP 3.x	yes
cancel	partial
task depend	partial
task detach	yes
taskgroup	yes
taskloop	yes
teams	yes
scope	no

... AND THE FUTURE (V9.0)

- Current OpenMP features planned for Score-P v9.0
 - Improvements to support of task directives
 - Improved support for `cancel` directive

OpenMP directive	Support
OpenMP 3.x	yes
<code>cancel</code>	partial
<code>task depend</code>	partial
<code>task detach</code>	yes
<code>taskgroup</code>	yes
<code>taskloop</code>	yes
<code>teams</code>	yes
<code>scope</code>	no

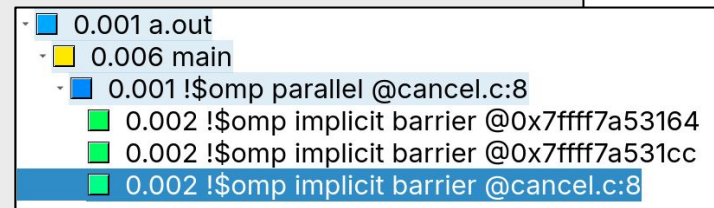
```
#include <omp.h>

int main( void )
{
    if( !omp_get_cancellation() ) return 1;

    int cancelled = 0;
    #pragma omp parallel
    {
        if( omp_get_thread_num() == 0 )
        {
            cancelled = 1;
            #pragma omp cancel parallel
        }

        #pragma omp cancellation point parallel
        // Do some very long calculation, if region is not cancelled
    }

    return cancelled ? 0 : 1;
}
```



```
- 0.001 a.out
- 0.006 main
- 0.001 !$omp parallel @cancel.c:8
  0.002 !$omp implicit barrier @0x7fff7a53164
  0.002 !$omp implicit barrier @0x7fff7a531cc
  0.002 !$omp implicit barrier @cancel.c:8
```

... AND THE FUTURE (V9.0)

- Current OpenMP features planned for Score-P v9.0
 - Improvements to support of task directives
 - Improved support for `cancel` directive
 - Support for `teams` directive

OpenMP directive	Support
OpenMP 3.x	yes
<code>cancel</code>	partial
<code>task depend</code>	partial
<code>task detach</code>	yes
<code>taskgroup</code>	yes
<code>taskloop</code>	yes
<code>teams</code>	yes
<code>scope</code>	no

```
int main( void )
{
    int sum = 0;

    #pragma omp teams num_teams( 4 )
    #pragma omp distribute parallel for num_threads( 2 ) reduction( +: sum )
    for( int i = 0; i < 100; ++i )
    {
        sum++;
    }

    return sum == 100 ? 0 : 1;
}
```

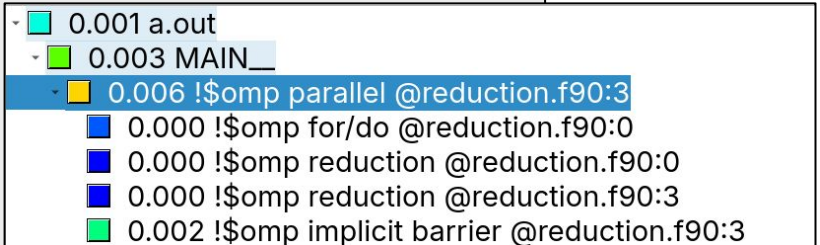
```
0.002 a.out
- 0.003 main
  - 0.003 !$omp teams @teams.c:5
    - 0.000 !$omp distribute @teams.c:6
      - 0.000 !$omp parallel @teams.c:6
        - 0.000 !$omp for/do @teams.c:6
          - 0.001 !$omp implicit barrier @teams.c:6
            - 0.001 !$omp implicit barrier @teams.c:5
```

... AND THE FUTURE (V9.0)

- Current OpenMP features planned for Score-P v9.0
 - Improvements to support of task directives
 - Improved support for `cancel` directive
 - Support for `teams` directive
 - Recording of reduction clause

```
PROGRAM REDUCTION3
  N = 0
  !$OMP PARALLEL DO REDUCTION(+: N)
  DO I = 1, 100
    N = N + I
  END DO

  WRITE (*,*) N
END PROGRAM REDUCTION3
```



OpenMP directive	Support
OpenMP 3.x	yes
cancel	partial
task depend	partial
task detach	yes
taskgroup	yes
taskloop	yes
teams	yes
scope	no

COMPILER SUPPORT FOR OMPT

We do have strict requirements

- OpenMP runtimes have to implement the tools interface
- However, runtimes may have bugs or features not fully implemented
- We test the OpenMP runtime during configuration to prepare for known issues

Compiler	Host Host	Events Target	Accelerator Device	Events Tracing
AOMP 18.0-1	Full	Full		Full
CCE 17.0.0	None	Full		None
Clang 18.1.6	Full	Partial		None
GCC 14.1	None	None		None
NVHPC 24.5	Full	Full		None
oneAPI 2024.1	Full	Partial		None
ROCm 6.1	Full	Full		Partial

OMPT support: yes	OMPT support: yes	OMPT support: no
OMPT header: yes	OMPT header: yes	OMPT header: yes
OMPT tool: yes	OMPT tool: yes	OMPT tool: yes
OMPT C support: yes	OMPT C support: yes	OMPT C support: no, overdue events not dispatched
OMPT C++ support: yes	OMPT C++ support: yes	OMPT C++ support: no, overdue events not dispatched
OMPT Fortran support: yes	OMPT Fortran support: yes	OMPT Fortran support: no, overdue events not dispatched
OMPT critical checks passed: \ yes	OMPT critical checks passed: \ yes	
OMPT remediable checks passed: \ yes	OMPT remediable checks passed: \ no, wrong_test_lock_mutex, missing_work_loop_schedule detected	
OMPT is default: yes	OMPT is default: yes	

CCE 16.0.1

NVHPC 24.5

oneAPI 2024.1

RIGOROUS TESTING

Ensuring that Score-P is able to work with compiler runtimes



- Large internal test suite, based on OpenMP examples and additional smoke tests
- Score-P is regularly tested with GCC, LLVM/Clang, NVHPC, ROCm & oneAPI
- Allows easy testing of new compilers as soon as they release
 - NVHPC 24.5, released on May 22nd, did show differences in runtime for example
- Contributed to several bug reports to compiler vendors (~65 bugs since Dec. 2022) and workarounds in Score-P

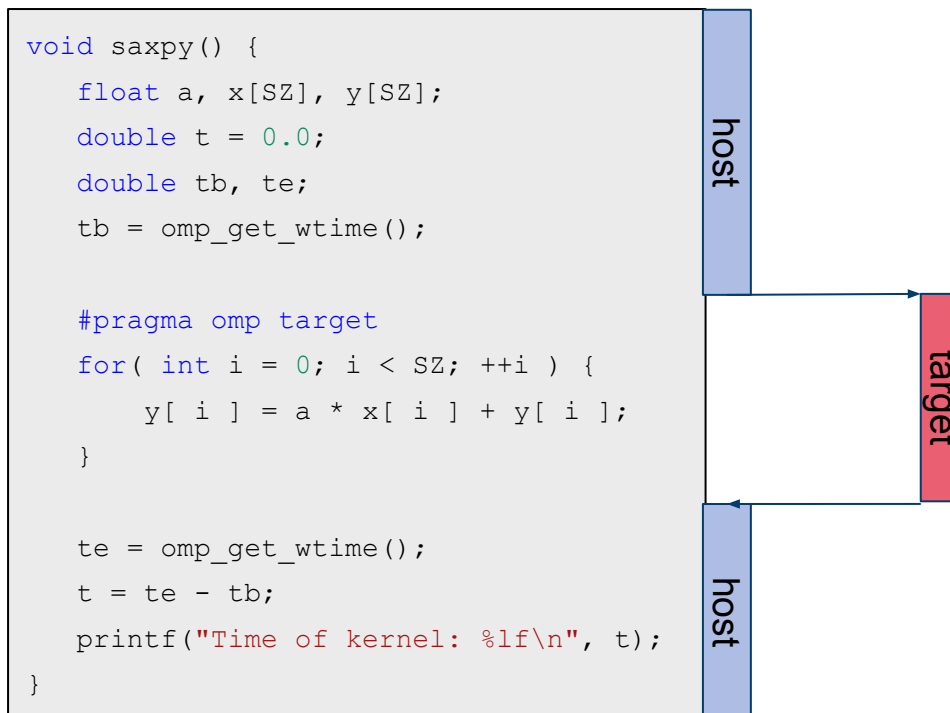
	C	C++	Fortran
Host Examples	138	7	165
Offload Examples	63	6	63
Teams Smoke Tests	30	0	0
General Smoke Tests	42	4	0

HOW TO HANDLE OFFLOADING TO ACCELERATORS?

OFFLOADING WORK TO ACCELERATORS

What is offloading in OpenMP?

- OpenMP introduced `target` directives in OpenMP 4.0
- Expanded in later spec. versions, including functions like `omp_target_alloc`



More information:



26th POP Webinar – Asynchronous
GPU Programming in OpenMP

Tuesday, 30 April 2024, 15:00 CEST

Christian Terboven
RWTH Aachen

Michael Klemm
AMD

OFFLOADING WORK TO ACCELERATORS

Splitting data transfers and kernels

```
void saxpy( float a, float* x, float* y, int n ) {  
    #pragma omp target teams distribute parallel for  
    for( int i = 0; i < n; ++i ) {  
        y[ i ] = a * x[ i ] + y[ i ];  
    }  
}  
  
void vecadd( float* x, float* y, float* data_out, int n )  
{  
    #pragma omp target teams distribute parallel for  
    for( int i = 0; i < n; ++i ) {  
        data_out[ i ] = x[ i ] + y[ i ];  
    }  
}  
  
void example() {  
    float a, x[N], y[N], b[N], data_out[N];  
  
    #pragma omp target data map ( to: x[:N], y[:N], b[:N], a ) \  
        map( from: data_out[:N] )  
    {  
        saxpy( a, x, y, N );  
        vecadd( y, b, data_out, N );  
    }  
}
```

Simple optimizations:

- Reduce numbers of data transfers
- Reuse data, if possible
- Transfer data beforehand, maybe asynchronously

More information:



26th POP Webinar – Asynchronous GPU Programming in OpenMP

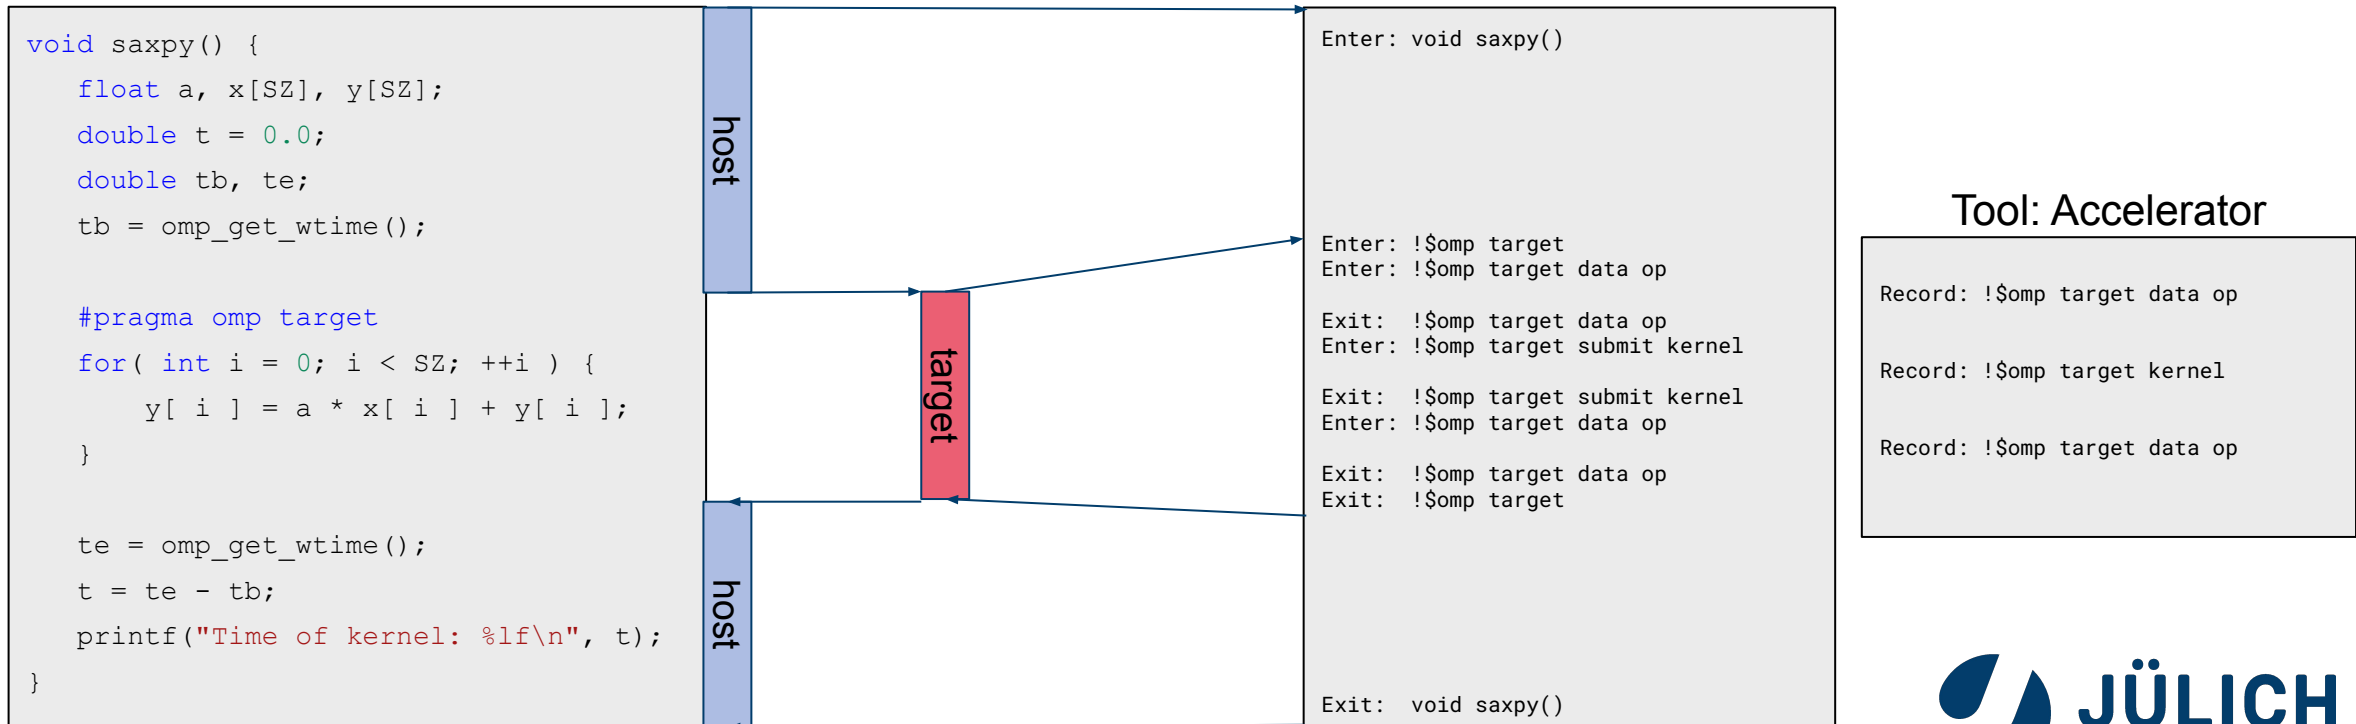
Tuesday, 30 April 2024, 15:00 CEST

Christian Terboven
RWTH Aachen

Michael Klemm
AMD

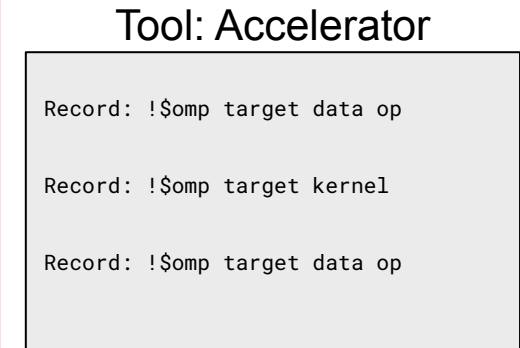
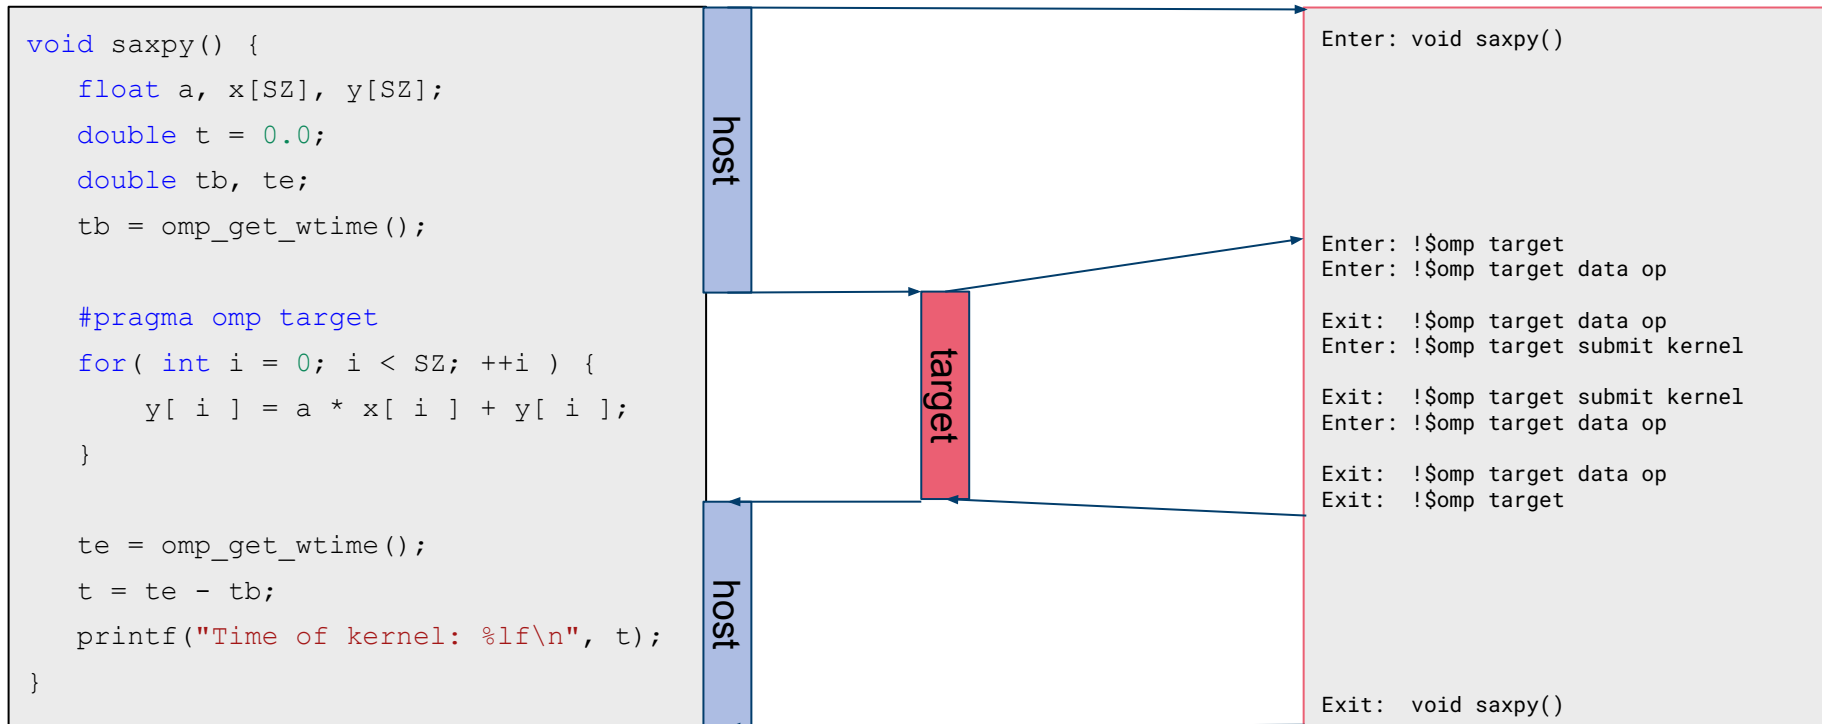
ACCELERATOR EVENTS AND TOOLS

- Unlike host directives, offloading needs to be handled twice to get all information
 - i. Host side dispatching events and waiting for completion
 - ii. Accelerator actually handling the event



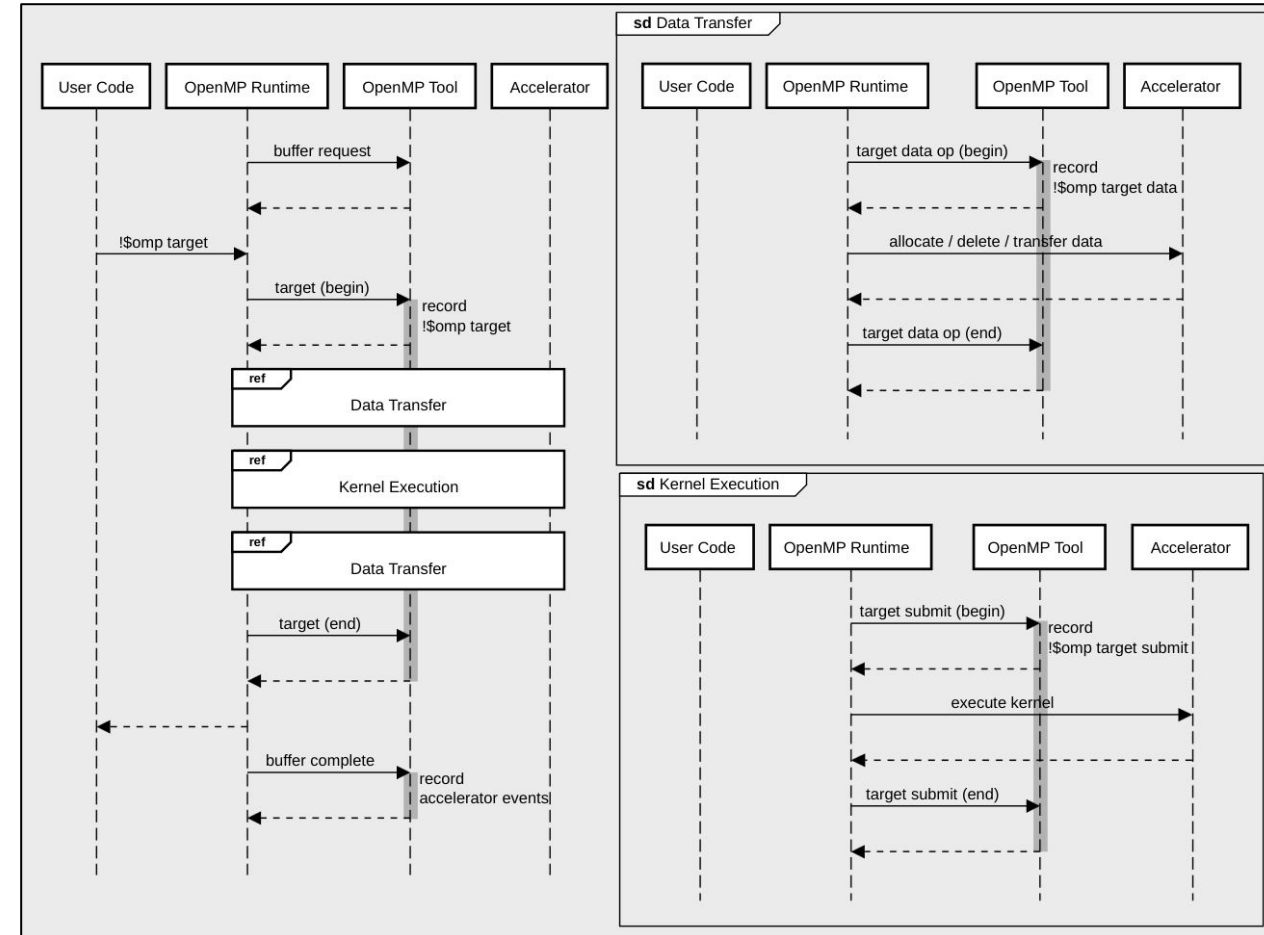
ACCELERATOR EVENTS AND TOOLS

- Unlike host directives, offloading needs to be handled twice to get all information
 - Host side dispatching events and waiting for completion
 - Accelerator actually handling the event



TARGET DIRECTIVES AND CALLBACKS

- Host side events are similar to existing ones (e.g. `parallel_begin`)
- We receive one begin and end event for the following scenarios:
 - `target` directives
 - data transfers
 - submit of a kernel
- This includes directives and function calls



TARGET DIRECTIVES AND CALLBACKS

Getting more into the details

```
typedef void (*ompt_callback_target_emi_t) (  
    ompt_target_t      kind,  
    ompt_scope_endpoint_t endpoint,  
    int                device_num,  
    ompt_data_t*       task_data,  
    ompt_data_t*       target_task_data,  
    ompt_data_t*       target_data,  
    const void*        codeptr_ra );
```

```
typedef void (*ompt_callback_submit_emi) (  
    ompt_scope_endpoint_t endpoint,  
    ompt_data_t*          target_data,  
    ompt_id_t*            host_op_id,  
    unsigned int          requested_num_teams );
```

```
typedef void (*ompt_callback_target_data_op_emi) (  
    ompt_scope_endpoint_t endpoint,  
    ompt_data_t*          target_task_data,  
    ompt_data_t*          target_data,  
    ompt_id_t*            host_op_id,  
    ompt_target_data_op_t otype,  
    void*                 src_addr,  
    int                   src_device_num,  
    void*                 dest_addr,  
    int                   dest_device_num,  
    size_t                 bytes,  
    const void*           codeptr_ra );
```

- Three callbacks, giving us the important information

TARGET DIRECTIVES AND CALLBACKS

Getting more into the details

```
typedef void (*ompt_callback_target_emi_t) (  
    ompt_target_t      kind,  
    ompt_scope_endpoint_t endpoint,  
    int                device_num,  
    ompt_data_t*      task_data,  
    ompt_data_t*      target_task_data,  
    ompt_data_t*      target_data,  
    const void*       codeptr_ra );
```

```
typedef void (*ompt_callback_submit_emi) (  
    ompt_scope_endpoint_t endpoint,  
    ompt_data_t*         target_data,  
    ompt_id_t*          host_op_id,  
    unsigned int         requested_num_teams );
```

```
typedef void (*ompt_callback_target_data_op_emi) (  
    ompt_scope_endpoint_t endpoint,  
    ompt_data_t*         target_task_data,  
    ompt_data_t*         target_data,  
    ompt_id_t*          host_op_id,  
    ompt_target_data_op_t optype,  
    void*               src_addr,  
    int                 src_device_num,  
    void*               dest_addr,  
    int                 dest_device_num,  
    size_t              bytes,  
    const void*         codeptr_ra );
```

- Three callbacks, giving us the important information
- Start / end of the operation

TARGET DIRECTIVES AND CALLBACKS

Getting more into the details

```
typedef void (*ompt_callback_target_emi_t) (  
    ompt_target_t      kind,  
    ompt_scope_endpoint_t endpoint,  
    int                device_num,  
    ompt_data_t*       task_data,  
    ompt_data_t*       target_task_data,  
    ompt_data_t*       target_data,  
    const void*        codeptr_ra );
```

```
typedef void (*ompt_callback_submit_emi) (  
    ompt_scope_endpoint_t endpoint,  
    ompt_data_t*         target_data,  
    ompt_id_t*           host_op_id,  
    unsigned int         requested_num_teams );
```

```
typedef void (*ompt_callback_target_data_op_emi) (  
    ompt_scope_endpoint_t endpoint,  
    ompt_data_t*         target_task_data,  
    ompt_data_t*         target_data,  
    ompt_id_t*           host_op_id,  
    ompt_target_data_op_t otype,  
    void*                src_addr,  
    int                  src_device_num,  
    void*                dest_addr,  
    int                  dest_device_num,  
    size_t               bytes,  
    const void*          codeptr_ra );
```

- Three callbacks, giving us the important information
- Start / end of the operation
- Information about what exactly is done

TARGET DIRECTIVES AND CALLBACKS

Getting more into the details

```
typedef void (*ompt_callback_target_emi_t) (  
    ompt_target_t      kind,  
    ompt_scope_endpoint_t endpoint,  
    int                device_num,  
    ompt_data_t*       task_data,  
    ompt_data_t*       target_task_data,  
    ompt_data_t*       target_data,  
    const void*        codeptr_ra );
```

```
typedef void (*ompt_callback_submit_emi) (  
    ompt_scope_endpoint_t endpoint,  
    ompt_data_t*         target_data,  
    ompt_id_t*           host_op_id,  
    unsigned int         requested_num_teams );
```

```
typedef void (*ompt_callback_target_data_op_emi) (  
    ompt_scope_endpoint_t endpoint,  
    ompt_data_t*         target_task_data,  
    ompt_data_t*         target_data,  
    ompt_id_t*           host_op_id,  
    ompt_target_data_op_t otype,  
    void*                src_addr,  
    int                  src_device_num,  
    void*                dest_addr,  
    int                  dest_device_num,  
    size_t                bytes,  
    const void*          codeptr_ra );
```

- Three callbacks, giving us the important information
- Start / end of the operation
- Information about what exactly is done
- Source code position

TARGET DIRECTIVES AND CALLBACKS

Getting more into the details

```
typedef void (*ompt_callback_target_emi_t) (  
    ompt_target_t      kind,  
    ompt_scope_endpoint_t endpoint,  
    int                device_num,  
    ompt_data_t*       task_data,  
    ompt_data_t*       target_task_data,  
    ompt_data_t*       target_data,  
    const void*        codeptr_ra );
```

```
typedef void (*ompt_callback_submit_emi) (  
    ompt_scope_endpoint_t endpoint,  
    ompt_data_t*         target_data,  
    ompt_id_t*           host_op_id,  
    unsigned int         requested_num_teams );
```

```
typedef void (*ompt_callback_target_data_op_emi) (  
    ompt_scope_endpoint_t endpoint,  
    ompt_data_t*         target_task_data,  
    ompt_data_t*         target_data,  
    ompt_id_t*           host_op_id,  
    ompt_target_data_op_t otype,  
    void*                src_addr,  
    int                  src_device_num,  
    void*                dest_addr,  
    int                  dest_device_num,  
    size_t               bytes,  
    const void*          codeptr_ra );
```

- Three callbacks, giving us the important information
- Start / end of the operation
- Information about what exactly is done
- Source code position
- Unique information per target region and operation

CORRELATION OF EVENTS

- We do need to correlate two things:
 - Host callbacks between each other
 - Host callbacks to accelerator events
- Done via `target_data` and `host_op_id`
- What needs to be transferred?

```
typedef void (*ompt_callback_submit_emi) (  
    ompt_scope_endpoint_t endpoint,  
    ompt_data_t*          target_data,  
    ompt_id_t*           host_op_id,  
    unsigned int         requested_num_teams );
```

CORRELATION OF EVENTS

- We do need to correlate two things:
 - Host callbacks between each other
 - Host callbacks to accelerator events
- Done via `target_data` and `host_op_id`
- What needs to be transferred?

```
typedef struct scorep_ompt_target_data_t {  
    const void* codeptr_ra;  
    ompt_id_t   target_id;  
    bool        supports_device_tracing;  
} scorep_ompt_target_data;
```

```
typedef void (*ompt_callback_submit_emi) (  
    ompt_scope_endpoint_t endpoint,  
    ompt_data_t*          target_data,  
    ompt_id_t*            host_op_id,  
    unsigned int          requested_num_teams );
```

CORRELATION OF EVENTS

- We do need to correlate two things:
 - Host callbacks between each other
 - Host callbacks to accelerator events
- Done via `target_data` and `host_op_id`
- What needs to be transferred?

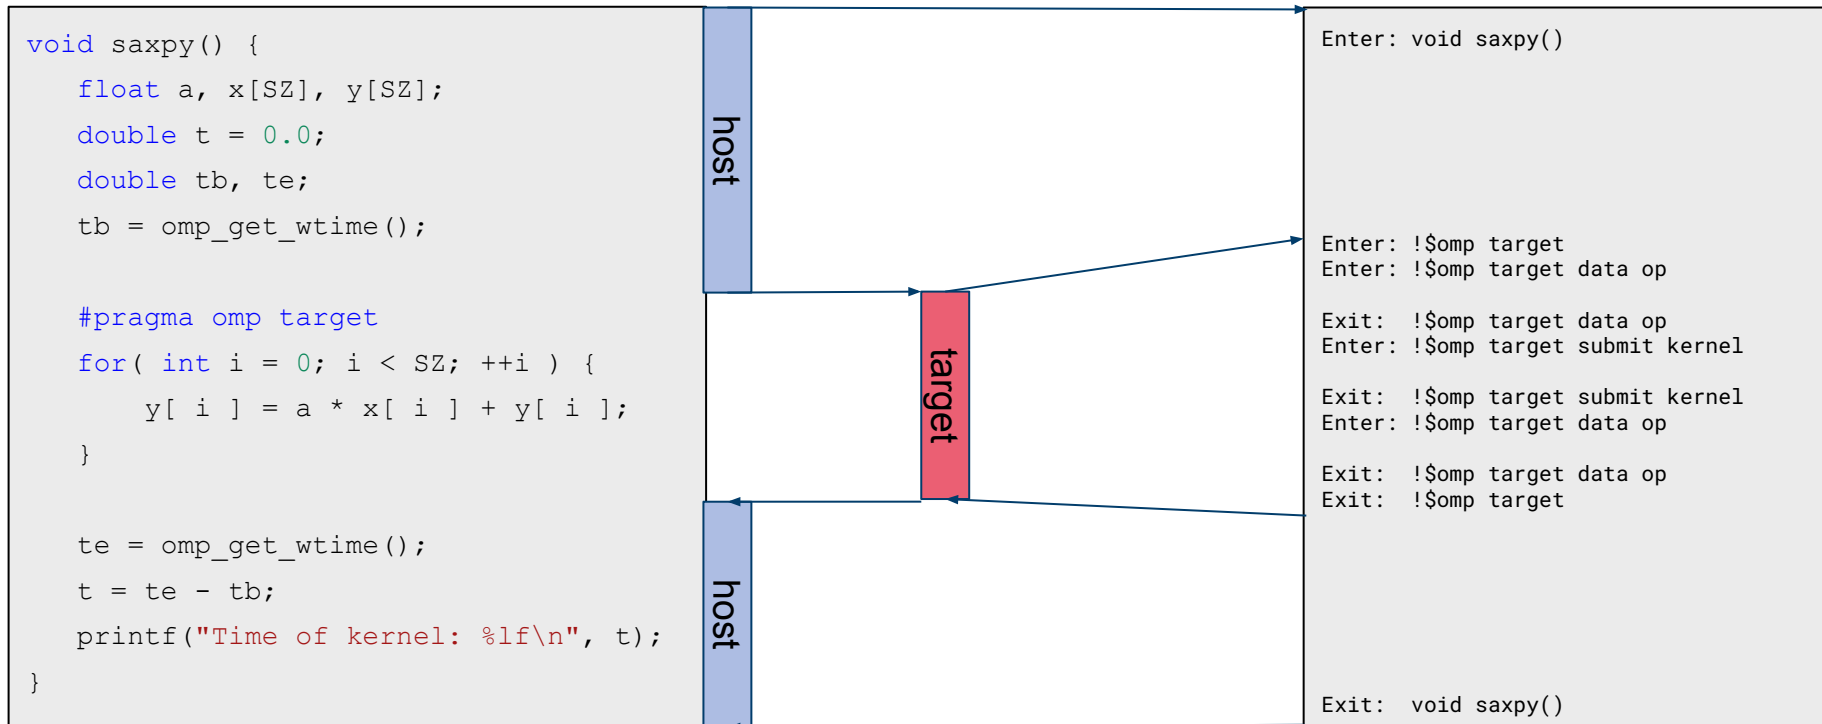
```
typedef struct scorep_ompt_target_data_t {  
    const void* codeptr_ra;  
    ompt_id_t   target_id;  
    bool       supports_device_tracing;  
} scorep_ompt_target_data;
```

```
/* Use the following 64-bit layout for mapping host_op_id:  
 * -----  
 * 00000000 00000000 00000000 | 00000000 00000000 00000000 00000000  
 * hostLocationId           | hostOpId  
 * -----  
 */
```

```
typedef void (*ompt_callback_submit_emi) (  
    ompt_scope_endpoint_t endpoint,  
    ompt_data_t*          target_data,  
    ompt_id_t*           host_op_id,  
    unsigned int          requested_num_teams );
```

ACCELERATOR EVENTS AND TOOLS

- Unlike host directives, offloading needs to be handled twice to get all information
 - i. Host side dispatching events and waiting for completion
 - ii. Accelerator actually handling the event



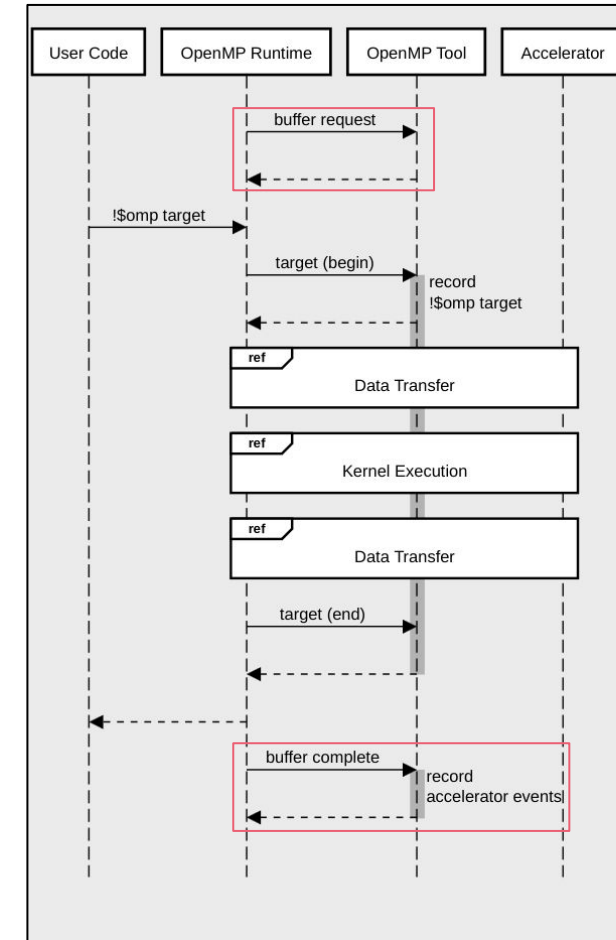
THE DEVICE TRACING INTERFACE

What is it?

- Buffer-based handling of accelerator events (similar to CUPTI, rocTracer)
- When a device is initialized, we can enable this interface
- Runtime will ask for buffers, record events, flush full buffers

What tools need to do:

- Sort buffers, as runtimes are not required to sort buffer
- Convert timestamps, either manually or via `ompt_translate_time`
- Iterate through buffer and write events



THE DEVICE TRACING INTERFACE

What does a record contain?

- Each record contains:
 - What type of event is recorded
 - When the record was recorded
 - Which thread recorded the record
 - The mapped target_data of the tool
 - The actual record of the callback
- Actual records may contain more information, like end timestamp and our set host_op_id

```
typedef struct ompt_record_ompt_t {  
    ompt_callbacks_t type;  
    ompt_device_time_t time;  
    ompt_id_t thread_id;  
    ompt_id_t target_id;  
    union {  
        [...]  
        ompt_record_target_t target;  
        ompt_record_target_data_op_t target_data_op;  
        ompt_record_target_kernel_t target_kernel;  
    } record;  
} ompt_record_ompt_t;
```

```
typedef struct ompt_record_target_kernel_t {  
    ompt_id_t host_op_id;  
    unsigned int requested_num_teams;  
    unsigned int granted_num_teams;  
    ompt_device_time_t end_time;  
} ompt_record_target_kernel_t;
```

BRINGING BOTH TOGETHER

What can we record with callbacks and the device tracing interface?

```
typedef void (*ompt_callback_target_emi_t) (  
    ompt_target_t      kind,  
    ompt_scope_endpoint_t endpoint,  
    int                device_num,  
    ompt_data_t*      task_data,  
    ompt_data_t*      target_task_data,  
    ompt_data_t*      target_data,  
    const void*       codeptr_ra );
```

- 1 !\$omp target @0x00204a69 (callee id=6, target type=target)
- 1 !\$omp target submit @0x00204a69 (callee id=2)
- 0 PER PROCESS METRICS
- 0 KERNELS
- 1 !\$omp target (kernel execution) @0x00204a69 (callee id=2, teams requested=1, teams granted=1)

```
typedef void (*ompt_callback_submit_emi) (  
    ompt_scope_endpoint_t endpoint,  
    ompt_data_t*         target_data,  
    ompt_id_t*          host_op_id,  
    unsigned int         requested_num_teams );
```

```
typedef struct ompt_record_ompt_t {  
    ompt_callbacks_t  type;  
    ompt_device_time_t time;  
    ompt_id_t         thread_id;  
    ompt_id_t         target_id;  
    union {  
        [...]  
        ompt_record_target_t      target;  
        ompt_record_target_data_op_t target_data_op;  
        ompt_record_target_kernel_t target_kernel;  
    } record;  
} ompt_record_ompt_t;
```

```
typedef struct ompt_record_target_kernel_t {  
    ompt_id_t host_op_id;  
    unsigned int requested_num_teams;  
    unsigned int granted_num_teams;  
    ompt_device_time_t end_time;  
} ompt_record_target_kernel_t;
```

HOW DO I USE THE NEW FEATURES?

HOW TO USE THE NEW FEATURES?

There are no additional steps needed!

- The OpenMP Tools Interface will be the default with Score-P v9.0
- Example:

```
~/Sources/OpenMP/saxpy > scorep --version                               jreuter@zam226
Score-P 9.0-dev

~/Sources/OpenMP/saxpy > scorep-amdclang -fopenmp --offload-arch=gfx1101 saxpy.c

~/Sources/OpenMP/saxpy > SCOREP_ENABLE_TRACING=true ./a.out          jreuter@zam226
Time of kernel: 0.009119
[Score-P] src/measurement/scorep_definition_cube4.c:771: Warning: Given metric name "Open
MP Memory" was changed to "OpenMP_Memory" for CubePL processing, i.e., .spec file and Cub
eGUI derived metrics processing. Given name still used for display. Note, profiling only.
```

- If device tracing cannot be activated, a message is shown

```
~/Sources/OpenMP/saxpy > ./a.out
[Score-P] src/adapters/ompt/scorep_ompt_events_device_tracing.inc.c:85: Warning: Device Tracing interface could not be initial
ized and will be disabled for device NVIDIA GeForce MX550 (0). This will lead to accelerator events not showing up in the resu
lts. Lookup function is NULL.
Time of kernel: 0.074238
```

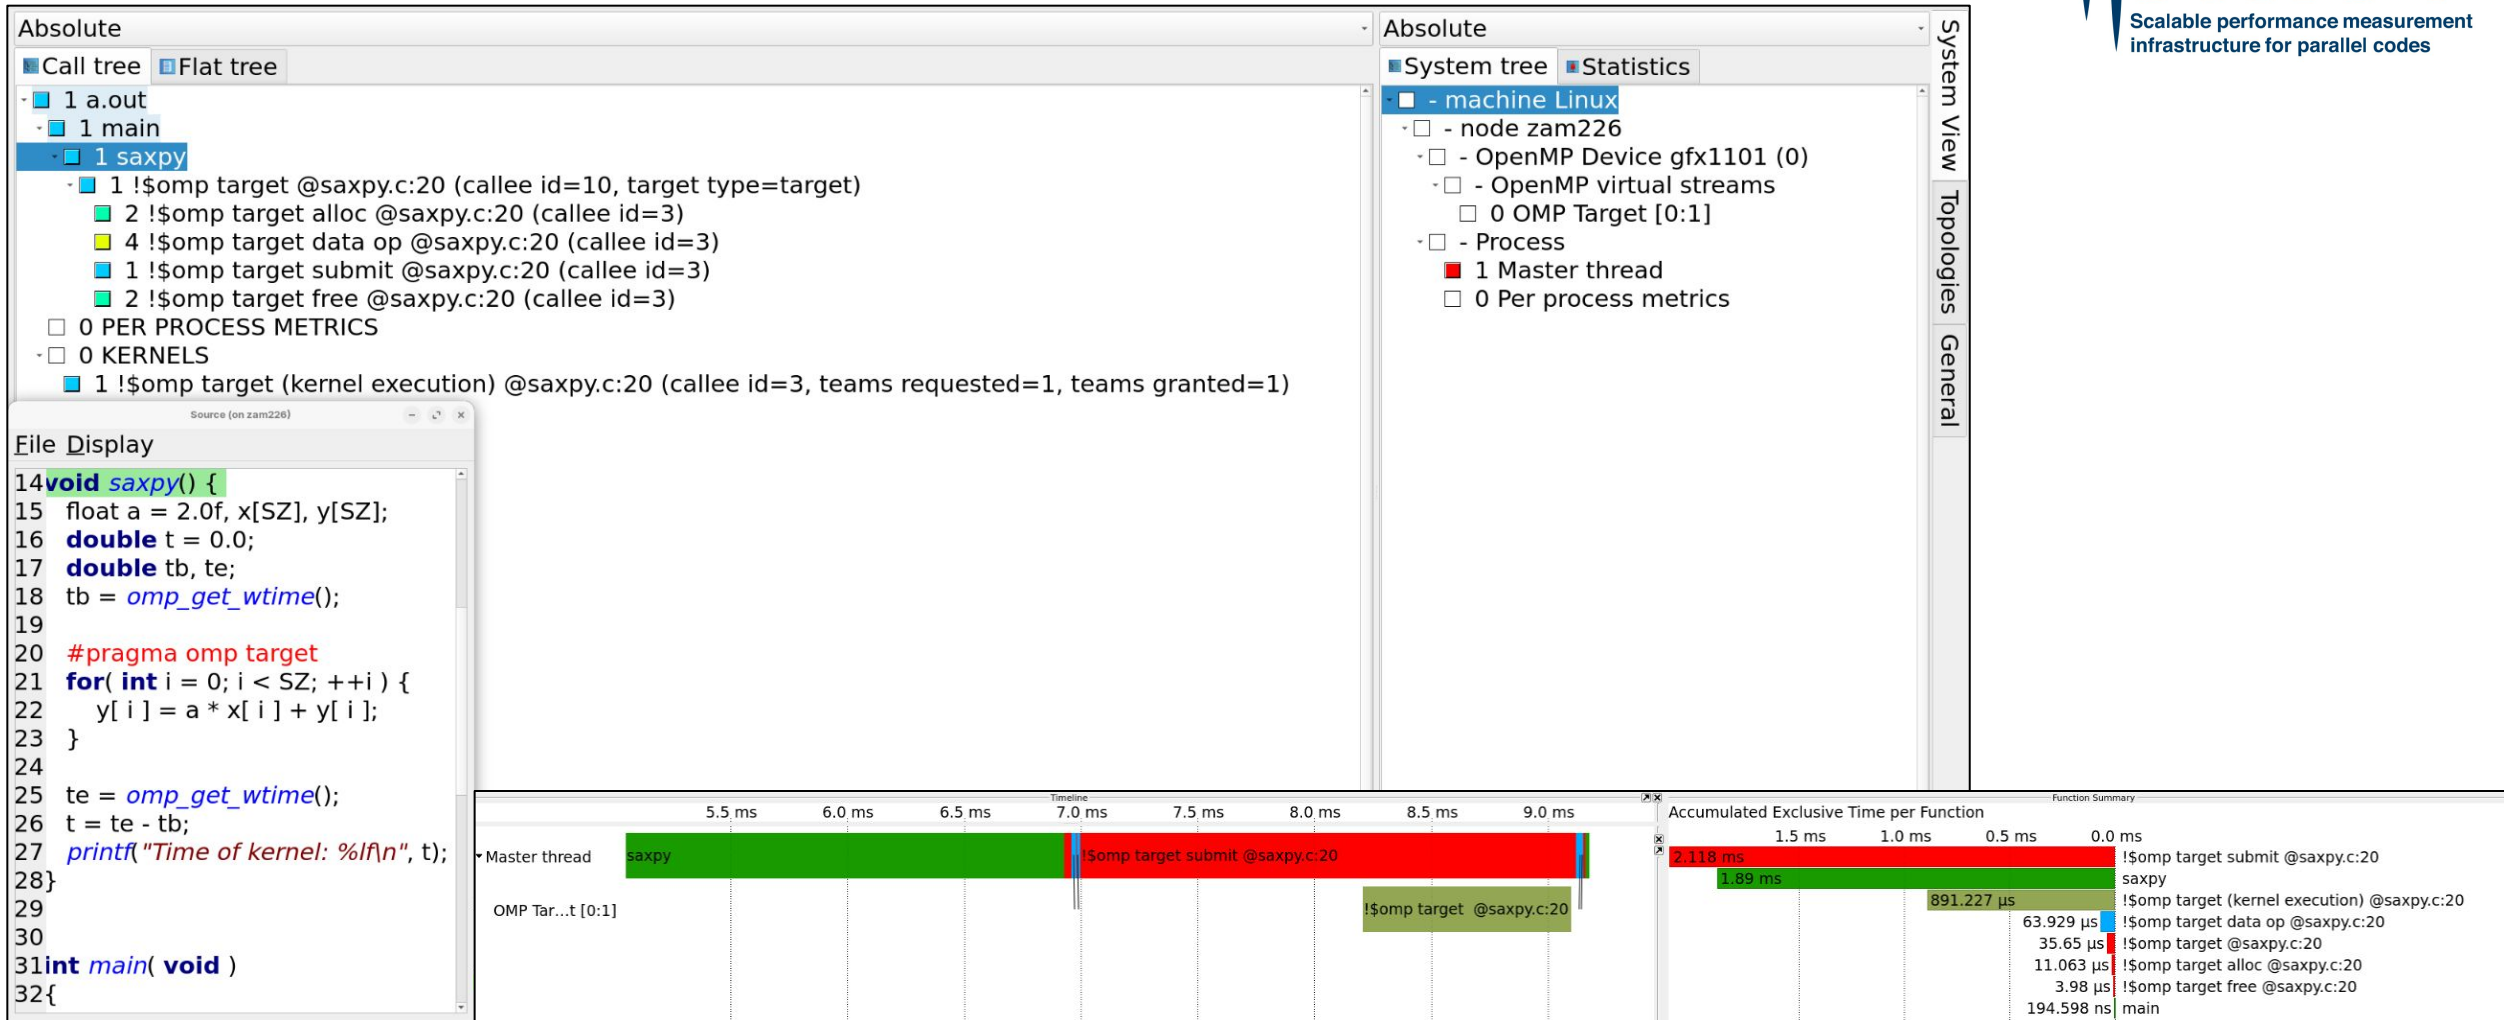
```
#include <stdio.h>
#include <omp.h>
#define SZ 1000

void init(float* arr, size_t size) {
    for(int i = 0; i < size; ++i) {
        arr[ i ] = i;
    }
}

void saxpy() {
    float a = 2.0f, x[SZ], y[SZ];
    double t = 0.0;
    double tb, te;
    tb = omp_get_wtime();
    #pragma omp target
    for( int i = 0; i < SZ; ++i ) {
        y[ i ] = a * x[ i ] + y[ i ];
    }
    te = omp_get_wtime();
    t = te - tb;
    printf("Time of kernel: %lf\n", t);
}

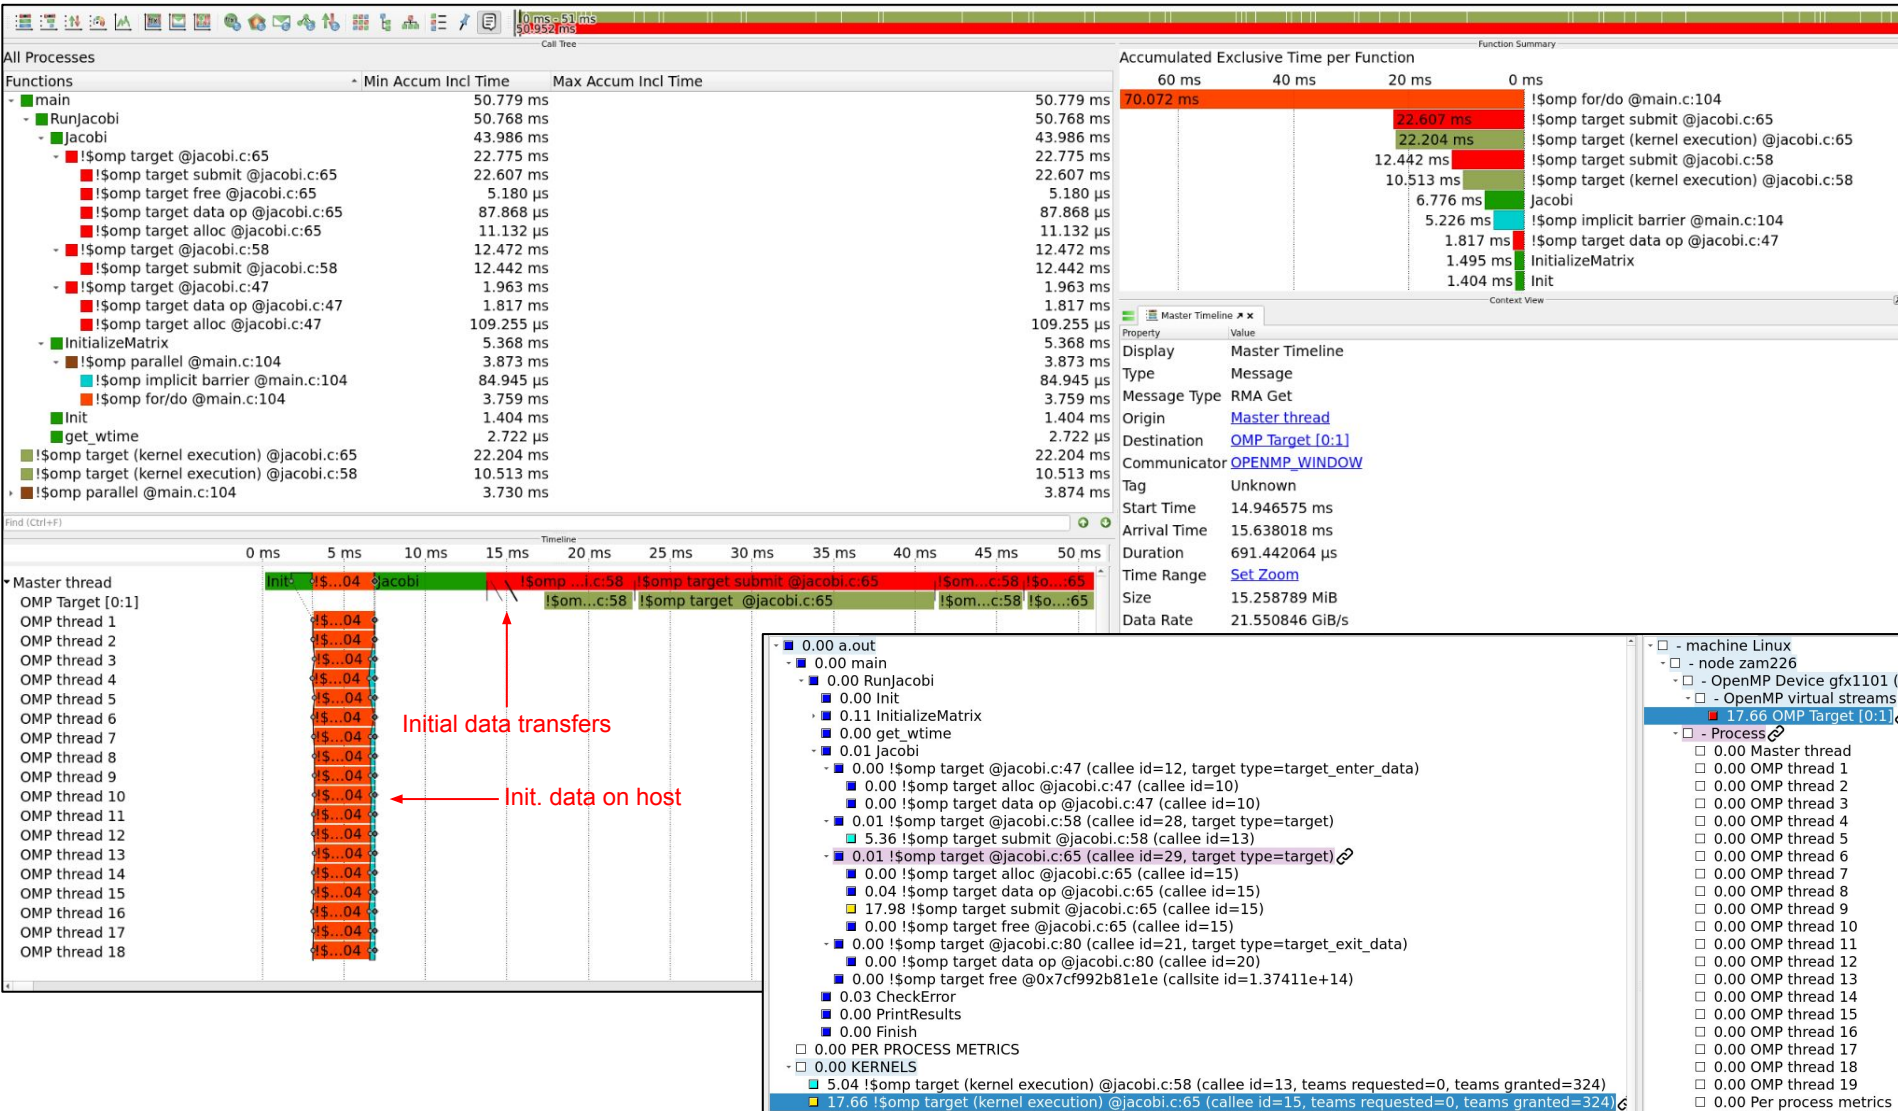
int main( void ) {
    saxpy();
}
```

RESULTS: LETS START SMALL ...



... INCREASE THE COMPLEXITY ...

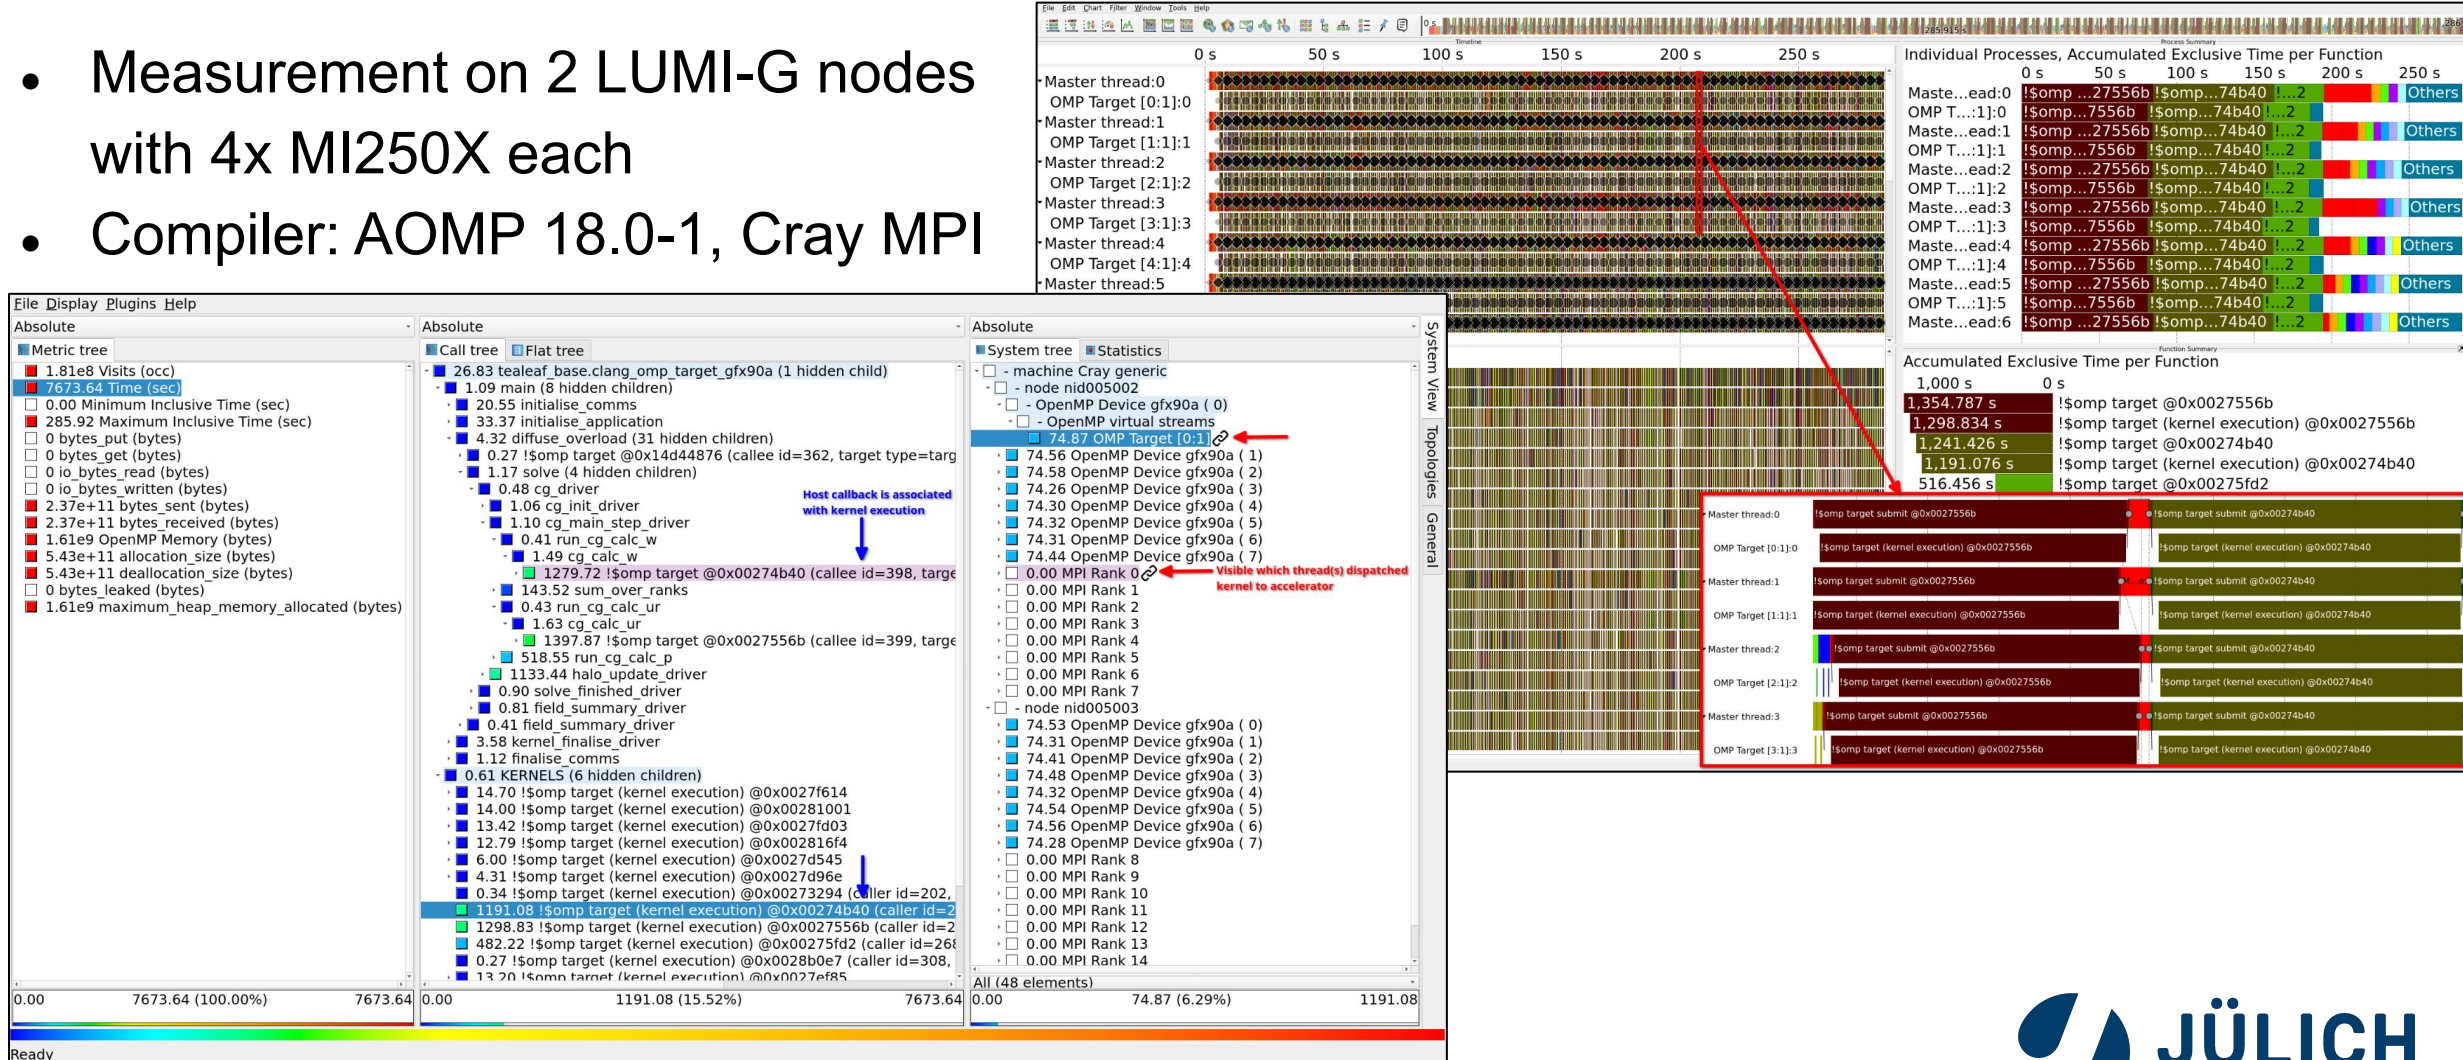
- Jacobi example used in Score-P testing
- System info:
 - Ubuntu 22.04
 - ROCm 6.1.0
 - RX 7700 XT



... AND SCALE IT UP!

Running the SPEC HPC tealeaf benchmark on multiple LUMI-G nodes

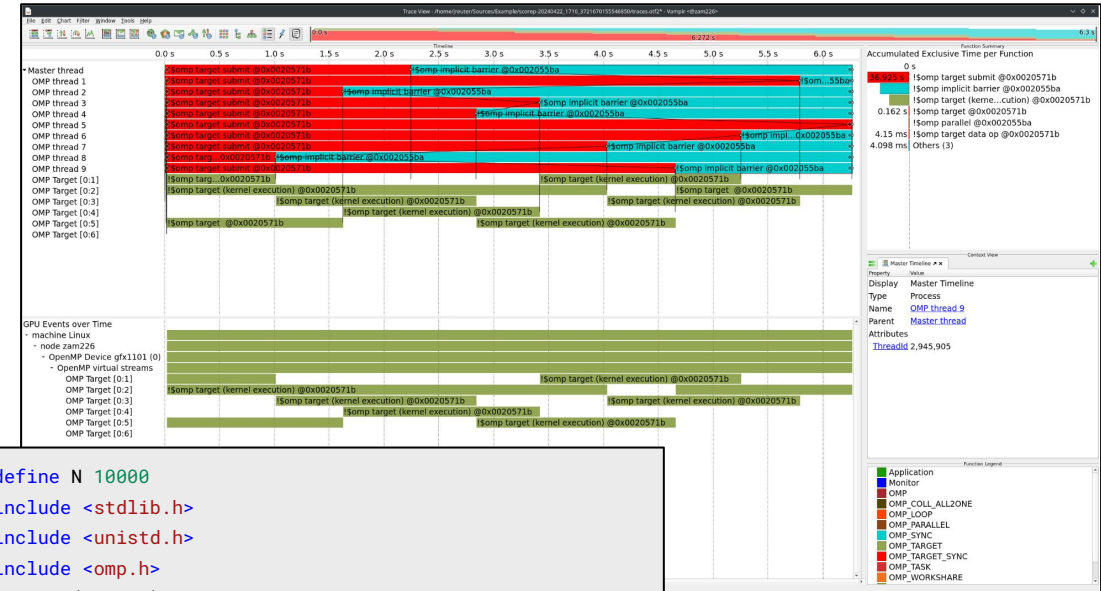
- Measurement on 2 LUMI-G nodes with 4x MI250X each
- Compiler: AOMP 18.0-1, Cray MPI



LIMITATIONS

The OpenMP Tools Interface only offers so much... (as of spec. 5.2)

- The device tracing interface doesn't give us any information about the low-level stream for any event
 - To handle overlapping events, we create OpenMP virtual streams
 - May lead to more OpenMP virtual streams than actually created by runtime



```
#define N 10000
#include <stdlib.h>
#include <unistd.h>
#include <omp.h>
int main( void )
{
    #pragma omp parallel
    {
        size_t* sum = malloc( sizeof( size_t ) * N );
        #pragma omp target teams \
            distribute parallel for simd \
            map(tofrom: sum[:N])
        for( size_t i = 0; i < N; ++i )
        {
            sum[i] = i;
        }
        free( sum );
    }
}
```


LIMITATIONS

Compiler and runtime limitations

- The device tracing interface doesn't give us any information about the low-level stream for any event
 - To handle overlapping events, we create OpenMP virtual streams
 - May lead to more OpenMP virtual streams than actually created by runtime
- OpenMP runtimes still have runtime issues
 - We generally recommend the latest releases, as they are the most stable

```
- For best support, it is advised to use the latest compiler versions
to ensure best support of the device tracing interface. In our testing,
ROCM 6.1.0 and AOMP 19.0-0 offer the best support, with the following
limitations:
- AOMP 19.0-0 will report incorrect times for data transfers
  between devices.
- AOMP 18.0-1 may dead lock for short programs when multiple
  accelerators are initialized.
- ROCm 6.1.0 and earlier and AOMP 18.0-0 and earlier do not support
  multiple devices per rank. If kernels are executed on more than one
  device per process, execution may abort. Otherwise events may
  be associated with the wrong accelerator.
- AOMP 18.0-0 incorrectly maps identifiers between callbacks and the
  device tracing interface. This leads to data transfers being shown
  incorrectly between the host threads and devices.
- ROCm 5.7.1 and earlier and AOMP 17.0-3 and earlier do not support
  accessing a device from multiple threads. This may lead to issues
  where events are associated with the incorrect host thread.
- ROCm 5.6 to 5.7.1 do not dispatch all callbacks for `#pragma omp
  target enter/exit data`. Score-P will abort due to timestamp issues.
- When utilizing multiple accelerators with ROCm 5.5, execution will
  dead lock at the end of the program execution when Score-P calls
  `stop_trace` for the device tracing interface.
- ROCm 5.4 and earlier are not supported due to not having a way to
  translate the device time to host time.
```

Snippet from our OPEN_ISSUES for AMD (others to follow)

WHAT ABOUT OTHER RUNTIMES?

Not all runtimes do support device tracing

- Score-P will output a warning, reminding that no accelerator data will be collected by OMPT
- However, host callbacks are still recorded!

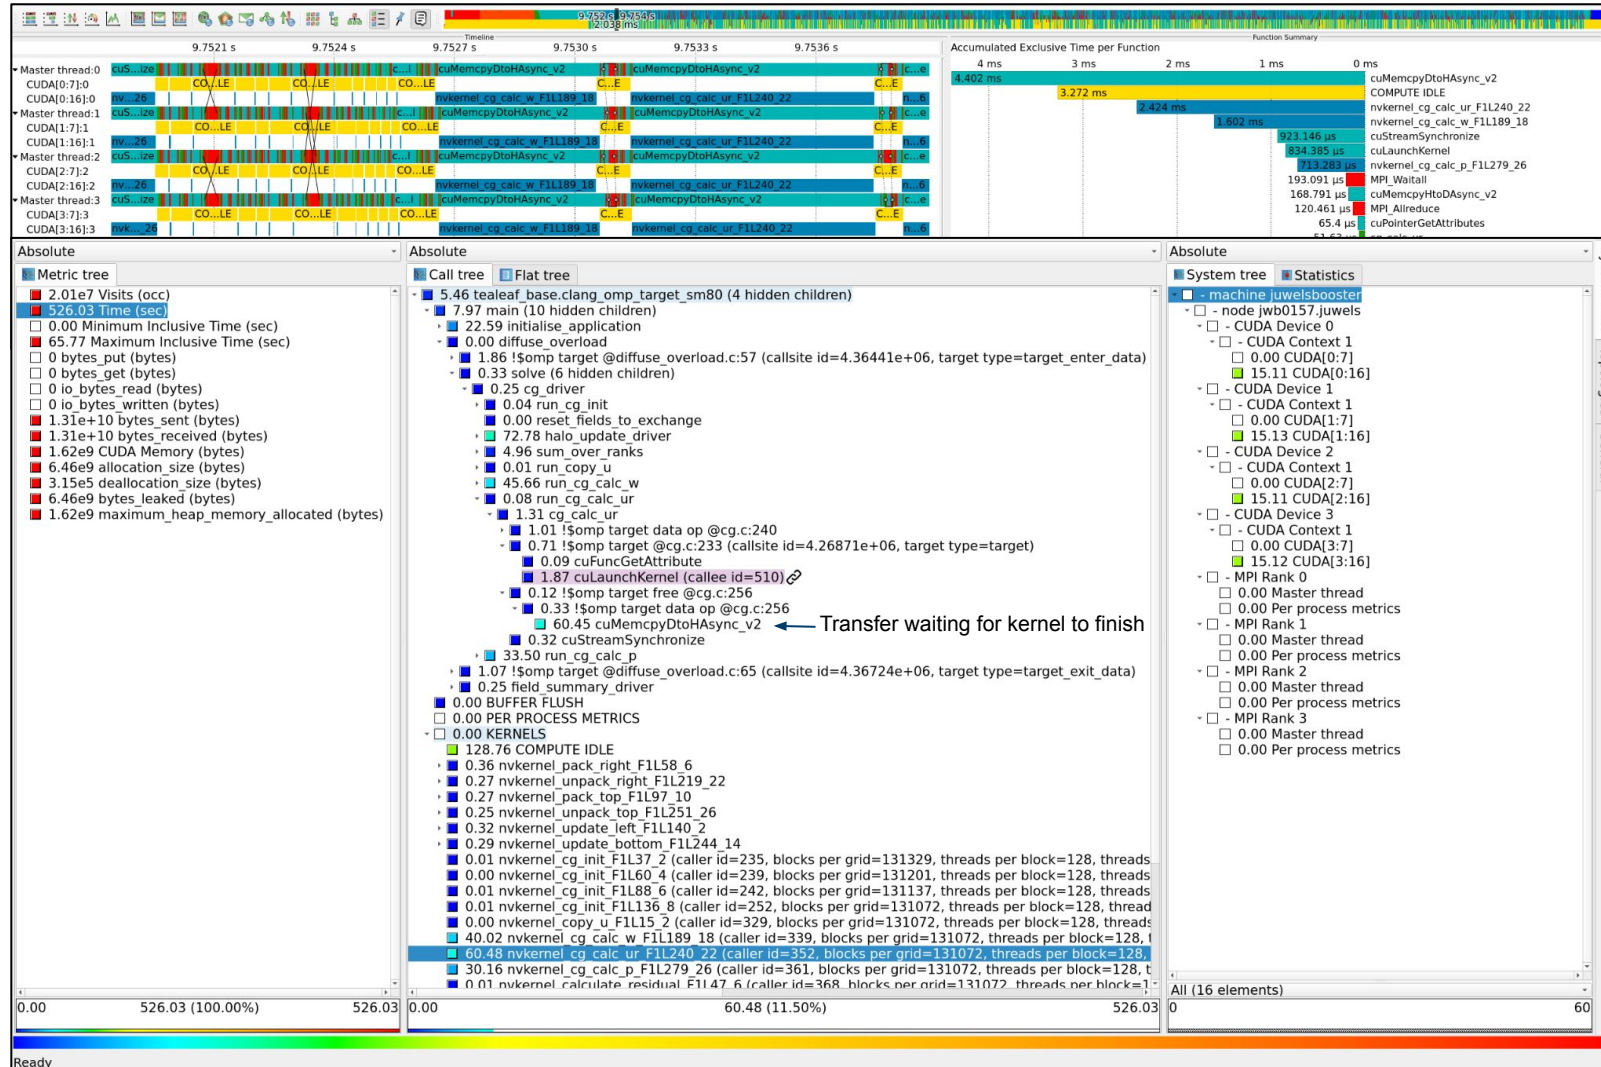
Solution: Use native GPU adapters, if possible!

```
$ scorep --thread=omp:ompt -cuda nvc -mp=gpu \  
my-code.c -o my-code.out  
$ SCOREP_ENABLE_TRACING=true ./my-code.out
```

Compiler	Host	Events	Accelerator Events
	Host	Target	Device Tracing
AOMP 18.0-1	Full	Full	Full
CCE 17.0.0	None	Full	None
Clang 18.1.6	Full	Partial	None
GCC 14.1	None	None	None
NVHPC 24.5	Full	Full	None
oneAPI 2024.1	Full	Partial	None
ROCm 6.1	Full	Full	Partial

COMBINING OMPT AND CUPTI

Showing off results on JUWELS Booster with NVHPC 23.7 on one node with 4x A100



- OMPT adapter is still able to record host events
- Some host events might show longer times than expected
 - Synchronization points of low-level runtime
- Native accelerator adapter records kernels and data transfers

FINAL WORDS

A short overview of what was shown in this talk



- With Score-P v9.0, we will expand our OpenMP support in several ways
- Most important: Users will be able to record OpenMP target events
 - AMD compilers sufficiently support the OpenMP Tools Interface
 - For other compilers: Native accelerator adapters required to get events
- Some compromises had to be made, partially because of the 5.2 specification
- Available implementation already works on several different systems and on small and large scale

OBTAIN SCORE-P AND GET IN CONTACT



- Visit our web page:
<https://score-p.org>
- Check out our public GitLab mirror:
<https://gitlab.com/score-p/scorep>
- Available on several different platforms:



<https://go.fzj.de/scorep-ompt-device-tracing>
Get the Score-P development version

**THANKS FOR YOUR ATTENTION!
QUESTIONS?**