


The 11th International workshop on OpenCL and SYCL

IWOCL & SYCLcon 2023



Transforming Fortran weather and climate applications to OpenCL using PSyclone

Sergi Siso, Hartree Centre STFC UKRI


with Andrew Porter and Rupert Ford, Hartree Centre STFC UKRI

April 18–20, 2023 | University of Cambridge, UK

iwocl.org

The future (and present) of HPC is heterogeneous

50 Years of Microprocessor Trend Data



Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
 New plot and data collected for 2010-2021 by K. Rupp

Source: <https://github.com/karlrupp/microprocessor-trend-data>

Top 500 list November 2022


Position	Name	Processor	Linpack (PFlop/s)
#1	Frontier	AMD EPYC 64 cores AMD Instinct MI250X	1,102
#2	Fugaku	Fujitsu A64FX 48C	442
#3	Lumi	AMD EPYC 64 cores AMD Instinct MI250X	309
#4	Leonardo	Xeon Platinum 8358 32C Nvidia A100 SMX4	174
#5	Summit	IBM POWER9 22C Nvidia V100	148
#7	Taihulight	Sunway SW26010 260C	93
#10	Tianhe-2A	Intel Xeon E5-2692v2 12C MATRIX-2000	61

And upcoming Intel GPUs, Nvidia CPUs, RISC-V, FPGAs, ...

Large number of HPC applications use Fortran

Archer2 Usage by Language

March-August 2022



Code	Language	Percentage use
VASP	Fortran	27.29%
Unidentified	Unknown	22.35%
CP2K	Fortran	6.28%
GROMACS	C	4.53%
CASTEP	Fortran	4.03%
Met Office UM	Fortran	3.10%
SENGA	Fortran	2.92%
Nektar++	C++	2.81%
NEMO	Fortran	2.46%
LAMMPS	C++	2.40%
PDNS3D	Fortran	1.89%
iIMB	Fortran	1.71%

Source: <https://cpufun.substack.com/p/is-fortran-a-dead-language> - Jim Cownie

<https://www.archer2.ac.uk/support-access/status.html#:~:text=0.0-,Historical%20usage%20data,-Period>

Software Sustainability

- HPC scientific applications are large and complex software projects.
 - Coupling of many different areas of expertise.
 - Large number of contributors from multiple institutions.
 - Some have millions of LOC.
- Productivity, readability, maintainability are essential for the sustainability of large software projects.
- Community effort: hard to maintain multiple implementations.

Ideally single source, with performance and parallelisation details abstracted.

Performance Portability Strategies

- **Maintain multiple implementations:** e.g., CUDA, HIP, OpenCL. Requires re-implementing the application in a new programming model and maintaining it over time.
- **Compiler hints/keywords:** e.g., OpenMP, OpenACC. Provide descriptive constructs. The compiler has flexibility to decide how to implement them for the target architecture.
- **Compile-time abstractions:** e.g., SYCL, Kokkos, Raja. Use C++ template metaprogramming to abstract the parallelisation API, the parallel execution order, how data structures are laid out in memory and on which space data resides.
- **Task-based parallelism:** e.g., Legion, Cabana, OmpSs, DaCe. Data-centric programming. The developer describes the dependencies between tasks and a runtime system decide how to execute them.

and what about Fortran?

Fortran has limited heterogeneous programming capabilities and lacks the powerful compile-time mechanisms that C++ performance portability frameworks use.


- **OpenMP 5 and OpenACC:** Still used differently on CPU and GPUs. Irregular vendor and compiler support.
- **CUDA Fortran:** Proprietary, single vendor and compiler support.
- **HPF/do concurrent:** Not widely adopted. Irregular compiler support.
- **Pre-processor macros:** Sometimes used in HPC codes but impacts software sustainability.

Can performance portability be achieved by source-to-source transformations?

PSyclone: a code generation and transformation system for weather and climate Fortran applications



This work: a new PScyclone backend for OpenCL




Portability != Performance Portability

- A direct mapping to a portable language backend is not enough!
- CPU, GPUs and especially FPGAs require different implementations.
- Performance portability can be improved by providing a list of code transformations (a PSyclone recipe) specific to each target platform.

PSyIR: PSyclone Intermediate Representation

- It is a **mutable representation** intended to be **programmatically manipulated** through transformations or PSyclone scripts.
- It **provides utilities** like DAG visualisations and automatic insertion of performance/debugging calipers to aid HPC experts.
- It gracefully **supports incomplete code information** like unsupported Fortran features and unresolved datatypes.
- It is itself **domain-agnostic**, but it is **extensible** to create the domain-specific DSLs that will be used by the applications.

The PSyclone workflow



Parallelisation and optimisation encoded as transformation 'recipes'.

Separation of concerns

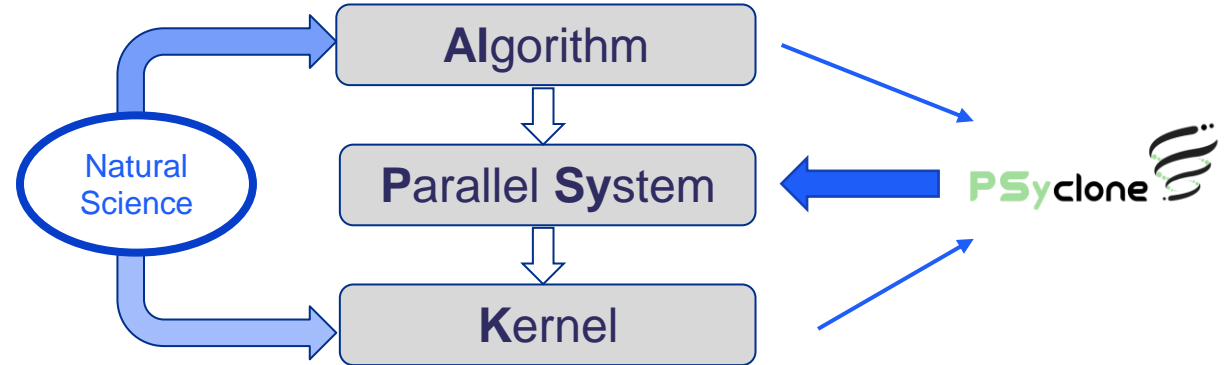
Perf portability: diff. recipes for each architecture

Visible "readable" output
Standard debugging/profiling

Standard output composable w. other s2s, vendor compilers

PSyKAI: a kernel-based model for Fortran

- Used in LFRic and NemoLite2D



```

1 call invoke(continuity(ssha_t, sshn_t, sshn_u, sshn_v, &
2             hu, hv, un, vn) &
3             momentum_u(ua, un, vn, hu, hv, ht, &
4                       ssha_u, sshn_t, sshn_u, sshn_v), &
5             momentum_v(va, un, vn, hu, hv, ht, &
6                       ssha_v, sshn_t, sshn_u, sshn_v), &
7             bc_ssh(istp, ssha_t), &
8             bc_solid_u(ua), &
9             bc_solid_v(va), &
10            bc_flather_u(ua, hu, sshn_u), &
11            bc_flather_v(va, hv, sshn_v), &
12            copy(un, ua), &
13            copy(vn, va), &
14            copy(sshn_t, ssha_t), &
15            next_sshu(sshn_u, sshn_t), &
16            next_sshv(sshn_v, sshn_t) &
17 )
  
```

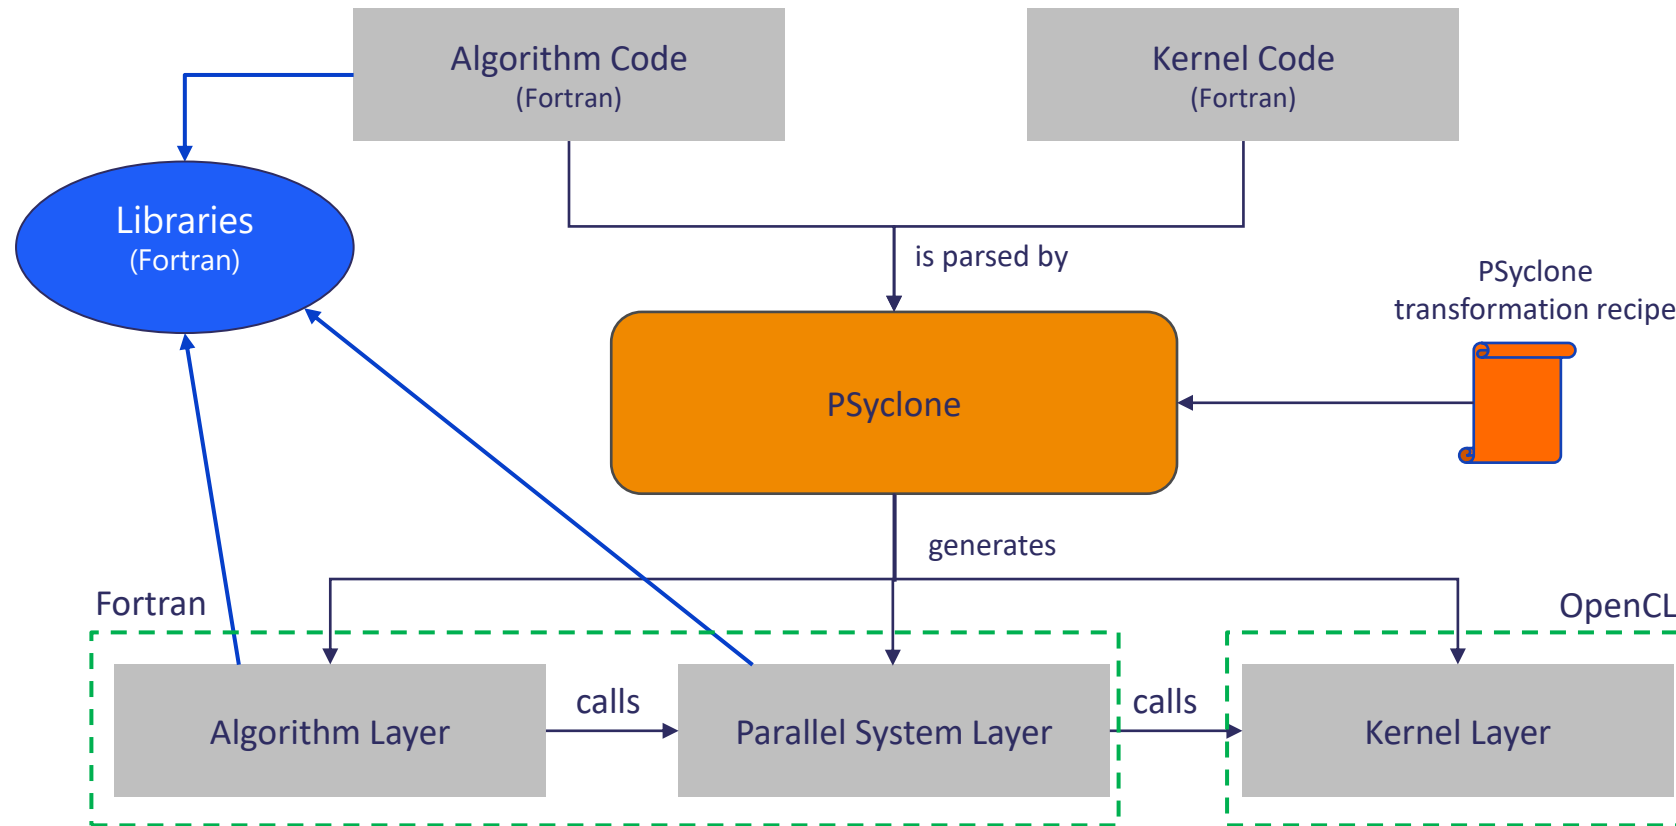
Listing 1: PSyKAI Algorithm layer

```

1 type, extends(kernel_type) :: bc_solid_u
2   type(go_arg), dimension(2) :: meta_args = &
3     (/ go_arg(GO_READWRITE, GO_CU, GO_POINTWISE), &
4       go_arg(GO_READ, GO_GRID_MASK_T) &
5     /)
6   integer :: ITERATES_OVER = GO_ALL_PTS
7   integer :: index_offset = GO_OFFSET_NE
8 contains
9   procedure, nopass :: code => bc_solid_u_code
10 end type bc_solid_u
11
12 subroutine bc_solid_u_code(ji, jj, ua, tmask)
13   integer, intent(in) :: ji, jj
14   integer, dimension(:, :), intent(in) :: tmask
15   real(go_wp), dimension(:, :), intent(inout) :: ua
16   if(tmask(ji, jj) * tmask(ji+1, jj) == 0) then
17     ua(ji, jj) = 0._go_wp
18   end if
19 end subroutine bc_solid_u_code
  
```

Listing 2: PSyKAI Kernel layer

Mapping PSyKAI to OpenCL



Simple example without optimisations

```
1 subroutine example_kernel_code(ji, jj, array1, array2)
2   use parameter_mode, only: scalar_parameter
3   implicit none
4   real(8), dimension(:, :), intent(inout) :: array1
5   real(8), dimension(:, :), intent(in) :: array2
6   array1(ji, jj) = array1(ji, jj) + array2(ji+1, jj) * &
7     scalar_parameter
8 end subroutine example_kernel_code
```

Listing 3: Example of a Fortran PSyKAI kernel

```
1 # For each kernel in the Parallel System layer ...
2 for kern in schedule.kernels():
3   # Convert the Globals to Arguments, since OpenCL
4   # kernels do not have access to Fortran global
5   # variables.
6   globals_to_arguments.apply(kern)
7
8 # Transform the whole Parallel System to use OpenCL
9 opencil_trans.apply(schedule)
```

Listing 4: PSyclone script to generate unoptimized OpenCL



```
1 ! Initialize OpenCL device, compile kernels and set up
2 ! buffers if it's the first time executing this PSy-layer
3 if (.first_time.) initialize_device_kernel_and_buffers( &
4     example_kernel, array1, array2)
5
6 ! Set up arguments in case they changes from previous
7 ! execution of this PSy-layer
8 array1_cl_mem = TRANSFER(array1, array1_cl_mem)
9 array2_cl_mem = TRANSFER(array2, array2_cl_mem)
10 ! Call clSetKernelArg for each OpenCL kernel argument
11 CALL kernel_set_args(example_kernel, array1_cl_mem, &
12     array2_cl_mem, scalar_parameter)
13
14 ! Launch the kernel
15 ierr = clEnqueueNDRangeKernel(cmd_queues(1), &
16     example_kernel, 2, &
17     (/xstart - 1, ystart - 1/), &
18     (/xlen - xstart - 1, ylen - ystart - 1/), &
19     C_NULL_PTR, 0, C_NULL_PTR, C_NULL_PTR)
```

Listing 5: Generated Fortran OpenCL PSy layer

```
1 __kernel void example_kernel_code(
2   __global double * restrict array1,
3   __global double * restrict array2,
4   double scalar_parameter
5 ){
6   int LEN1 = get_global_size(0);
7   int ji = get_global_id(0);
8   int jj = get_global_id(1);
9   array1[ji+jj*LEN1] = array1[ji+jj*1LEN1] + \
10     array2[(ji+1)+jj*1LEN1] * scalar_parameter;
11 }
```

Listing 6: Generated OpenCL code

*simplified representation of the generated OpenCL code

OpenCL Optimisations

- Kernel Blocking
- Boundary Masking

```
1 # For each kernel in the Parallel System layer ...
2 for kern in schedule.kernels():
3     # Convert the Globals to Arguments, since OpenCL
4     # kernels do not have access to Fortran global
5     # variables.
6     globals_to_arguments.apply(kern)
7
8     # Make kernels traverse the whole domain and mask
9     # out the computations in the boundary values
10    move_boundaries_trans.apply(kern)
11
12    # Provide a block size
13    koptions['local_size'] = 64
14
15    kern.set_openccl_options(koptions)
16
17 # Transform the whole Parallel System to use OpenCL
18 openccl_trans.apply(schedule)
```

Listing 7: PSyclone script to generate OpenCL



```
1 ! Set up arguments in case they changes from previous
2 ! execution of this PSy-layer
3 globalsize = (/grid%nx, grid%ny/)
4 localsize = (/64, 1/)
5 array1_cl_mem = TRANSFER(array1, array1_cl_mem)
6 array2_cl_mem = TRANSFER(array2, array2_cl_mem)
7 ! Call clSetKernelArg for each OpenCL kernel argument
8 CALL kernel_set_args(example_kernel, array1_cl_mem, &
9 & array2_cl_mem, scalar_parameter, &
10 & xstart - 1, xstop - 1, ystart - 1, ystop - 1)
11
12 ! Launch the kernel
13 ierr = clEnqueueNDRangeKernel(cmd_queues(1), &
14 & example_kernel, 2, C_NULL_PTR, &
15 & C_LOC(globalsize), C_LOC(localsize), 0, &
16 & C_NULL_PTR, C_NULL_PTR)
```

Listing 8: Generated Fortran OpenCL Parallel-System layer

```
1 __attribute__((reqd_work_group_size(64, 1, 1)))
2 __kernel void example_kernel_code(
3     __global double * restrict array1,
4     __global double * restrict array2,
5     double scalar_parameter,
6     int xstart, int xstop, int ystart, int ystop
7 ) {
8     int LEN1 = get_global_size(0);
9     int ji = get_global_id(0);
10    int jj = get_global_id(1);
11    if (((ji < xstart) || (ji > xstop)) || \
12        ((jj < ystart) || (jj > ystop))) {
13        return;
14    }
15    array1[ji+jj*LEN1] = array1[ji+jj*1LEN1] + \
16        array2[(ji+1)+jj*1LEN1] * scalar_parameter;
17 }
```

Listing 9: Generated OpenCL code

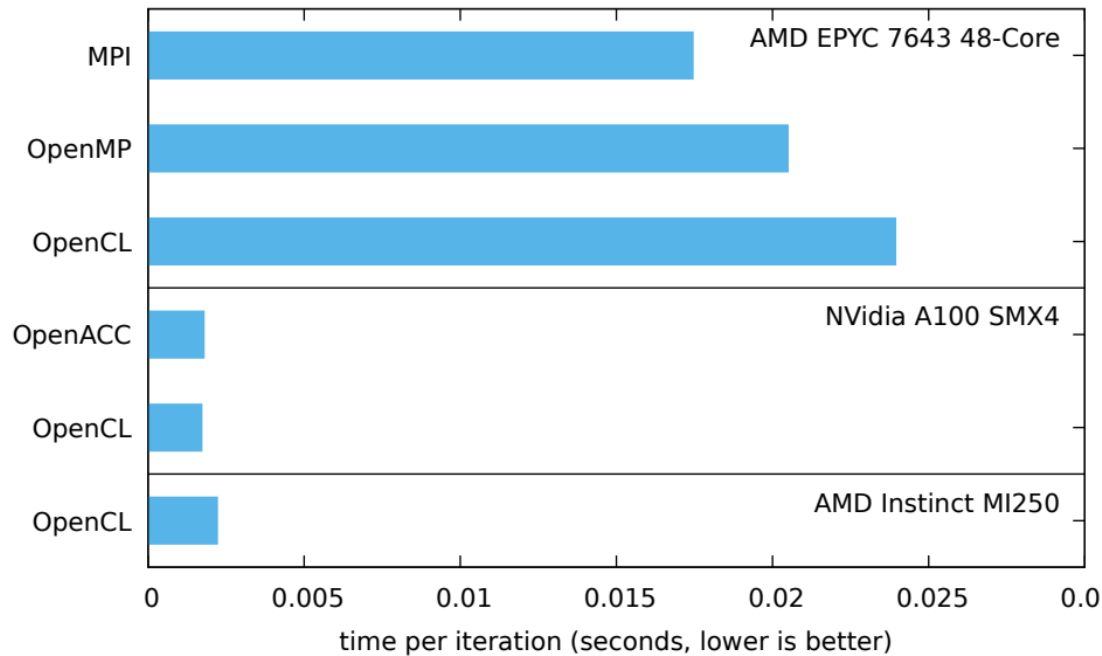
NemoLite2D (<https://github.com/stfc/PSycloneBench/tree/master/benchmarks/nemo/nemolite2d>)

- A vertically-averaged version of the free-surface component of the NEMO model.
- Implements a continuity equation for the update of the sea-surface height and two vertically-integrated momentum equations for the two velocity components.

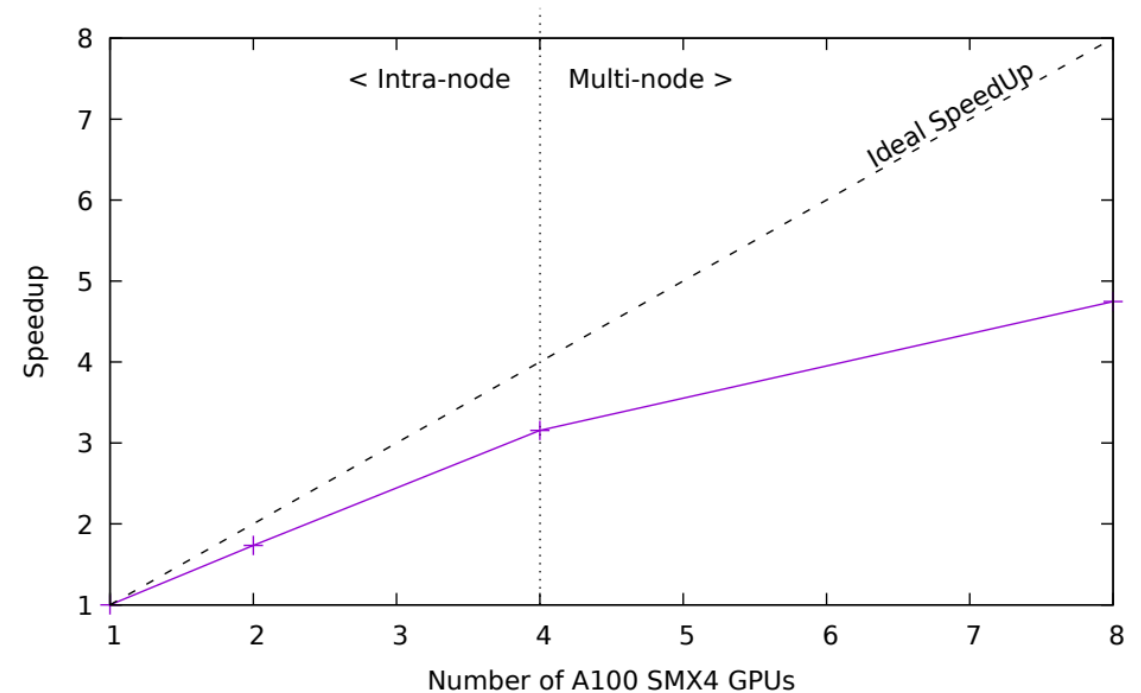
Kernel Name	Description	Fortran Lines
continuity	Eulerian forward time stepping method.	21
momentum_u	Compute advection, viscosity, coriolis force, ... on the u-field.	112
momentum_v	Compute advection, viscosity, coriolis force, ... on the v-field.	117
bc_ssh	Clamped boundary conditions.	26
bc_solid_u	Solid boundary conditions for u-velocity.	11
bc_solid_v	Solid boundary conditions for v-velocity.	11
bc_flather_u	Flather open boundary condition for u-velocity.	29
bc_flather_v	Flather open boundary condition for v-velocity.	29
copy	Copy the value from one array to another.	10
next_sshu	Time update of the u-field.	21
next_sshv	Time update of the v-field.	21

Performance Results

Performance comparison of NEMOLite2D (size 2048²) with multiple parallel programming models on multiple devices



Strong scalability of NEMOLite2D (size 6000²) with hybrid MPI and OpenCL



Best results for each device corresponds to 63, 60, 32 % of peak bandwidth respectively

- 48-core AMD EPYC 7643 CPU using the Intel OpenCL Runtime for CPUs and compiled with gfortran 9.4
- NVIDIA A100 SMX4 GPU using the NVIDIA OpenCL drivers and compiled with nvfortran 22.5
- AMD Instinct MI250 GPU using the ROCM 5.4 OpenCL drivers and compiled with gfortran 9.4

Dynamic Evaluation of Runtime Invariants

```
1  subroutine example_kernel(ji, jj, array, ct1, ct2)
2  implicit none
3  integer, intent(in) :: ji, jj, ct1, ct2
4  real(go_wp), dimension(:,,:), intent(inout) :: array
5  integer :: i
6
7  do i = 1, ct1
8      va(ji,jj) = va(ji,jj) / ct2
9  enddo
10 end subroutine kernel
```

Listing 10: Test kernel for dynamic optimizations

```
1 # Transform the Parallel System to use OpenCL
2 # with captured runtime invariants provided as
3 # pre-processor constants
4 opencl_trans.apply(
5     schedule, {'define_runtime_invariants': True})
```

Listing 11: PSyclone recipe with OpenCL dynamic optimizations

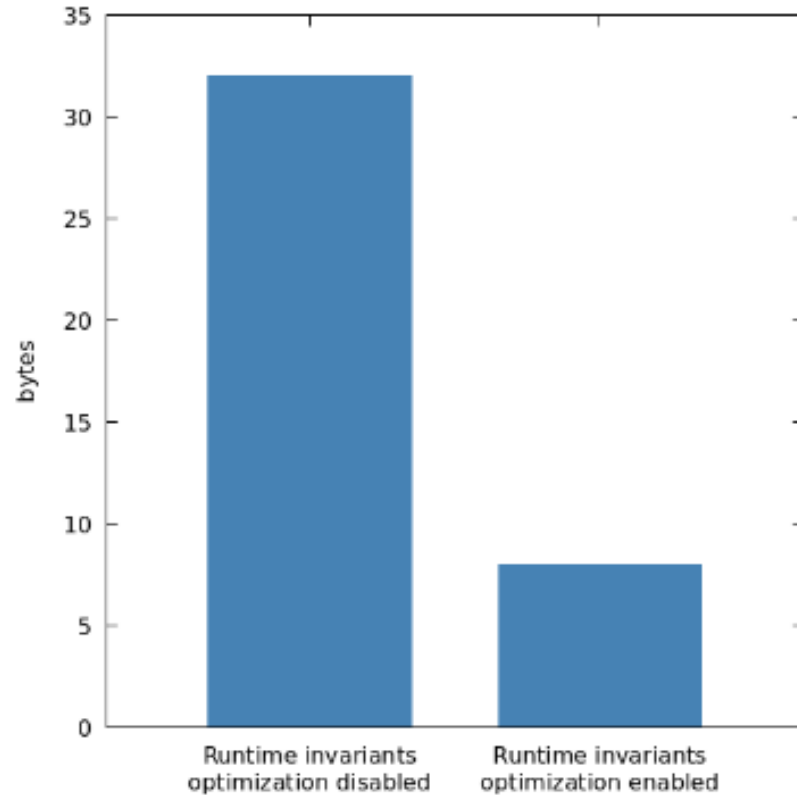


The generated OpenCL code replaces ct1 and ct2 with undeclared file scope symbols

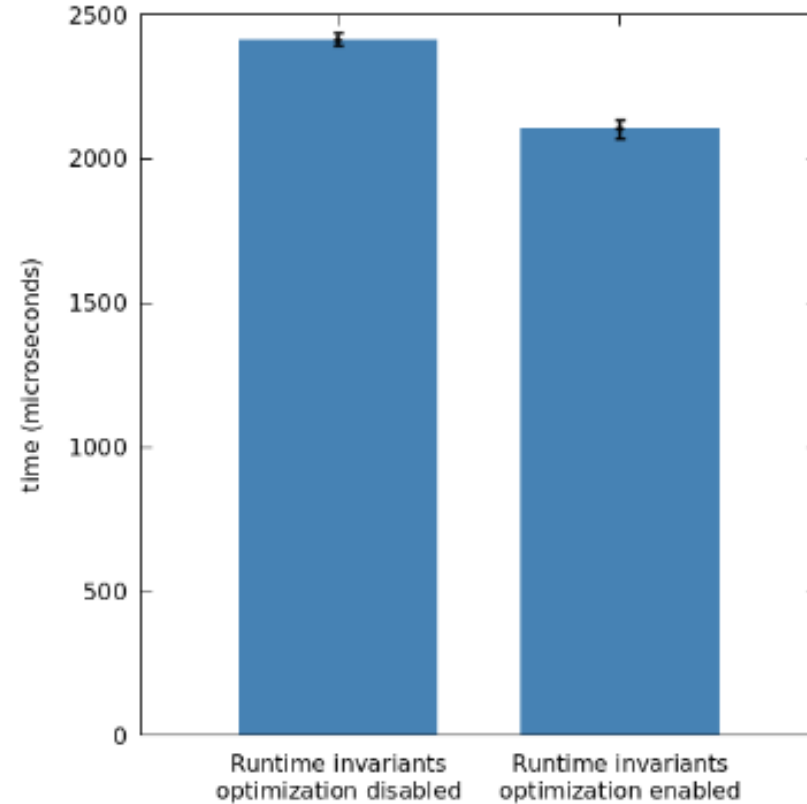
```
1 CHARACTER(LEN=4096) compiler_flags
2 WRITE (compiler_flags, *) ""
3 WRITE (compiler_flags, '(A,A,I0)') TRIM(compiler_flags),
4     " -Dxstart_example_kernel=", xstart
5 WRITE (compiler_flags, '(A,A,I0)') TRIM(compiler_flags),
6     " -Dxstop_example_kernel=", xstop
7 WRITE (compiler_flags, '(A,A,I0)') TRIM(compiler_flags),
8     " -Dystart_example_kernel=", ystart
9 WRITE (compiler_flags, '(A,A,I0)') TRIM(compiler_flags),
10    " -Dystop_example_kernel=", ystop
11 WRITE (compiler_flags, '(A,A,F0)') TRIM(compiler_flags),
12    " -Dcaptured_ct1=", ct1
13 WRITE (compiler_flags, '(A,A,F0)') TRIM(compiler_flags),
14    " -Dcaptured_ct2=", ct2
15 kernel_names(1) = 'example_kernel'
16
17 ! OpenCL Runtime Compilation
18 CALL add_kernels(1, kernel_names, &
19     compiler_flags=compiler_flags)
```

Listing 12: OpenCL driver code generated to capture runtime invariant values

Dynamic Evaluation of Runtime Invariants



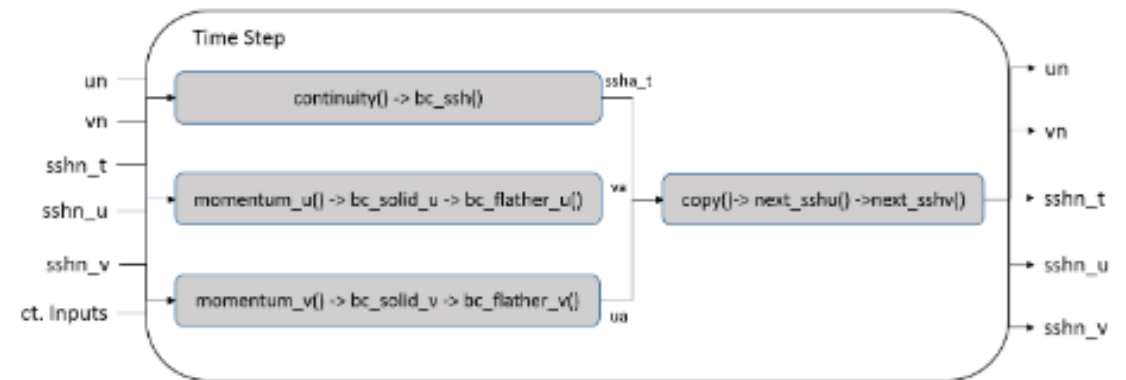
Number of bytes sent by `clSetKernelArg`



Execution time of the test kernel on an 8-core Intel Xeon Silver 4215

Targeting FPGAs: the EuroEXA project

- Required significant transformations from CPU/GPU code:
 - Functional parallelism (OCL queues)
 - Duplicate kernels
 - Inline loops into kernel (taskify)
 - Buffer burst memory operations to local memory.



Performance on Xilinx U200 FPGA

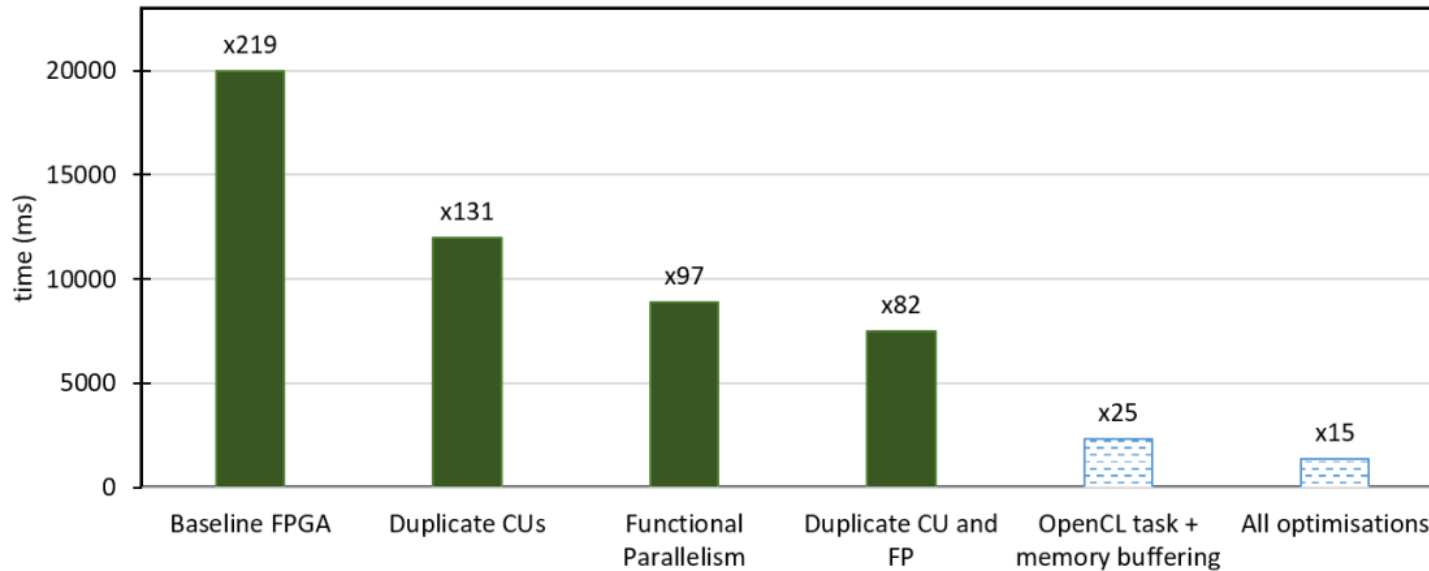



Figure: Execution time (Y axis) and slowdown compared to 1 Xeon Silver 4215 core (inverse of speedup - top of the bars) of multiple OpenCL optimizations on a Xilinx U200 FPGA. Solid bars are as generated by PSyclone, dashed bars required manual tweaks.

Current limitations:

- Only using 1 DDR memory bank and 1 SLR in the Xilinx U200 (out of 4 DDR memory banks and 3 SLR)
- Not using OpenCL pipes for faster communication between kernels.
- Not using OpenCL vendor extensions such as `xcl_dataflow`, `xcl_pipeline_loop` or `xcl_pipeline_workitems`.

General applicability of Fortran-to-OpenCL

Application specific



Standard Fortran

- NemoLite2D
- PSyKAI
- Any Numerical Operations

← Today

NEMO example:

- Infer kernels from Fortran array notation and dependency analysis
- Deduces domain specific knowledge from loop patterns and naming conventions of the code style-guide

```
5: If[]
  BinaryOperation[operator:'AND']
    Reference[name:'ln_wave']
    Reference[name:'ln_sdw']
  Schedule[]
    0: Loop[type='levels', field_space='None', it_space='None']
      Literal[value:'1', Scalar<INTEGER, UNDEFINED>]
      Reference[name:'jpkm1']
      Literal[value:'1', Scalar<INTEGER, UNDEFINED>]
      Schedule[]
        0: Loop[type='lat', field_space='None', it_space='None']
          Literal[value:'1', Scalar<INTEGER, UNDEFINED>]
          Reference[name:'jpi']
          Literal[value:'1', Scalar<INTEGER, UNDEFINED>]
          Schedule[]
            0: Loop[type='lon', field_space='None', it_space='None']
              Literal[value:'1', Scalar<INTEGER, UNDEFINED>]
              Reference[name:'jpi']
              Literal[value:'1', Scalar<INTEGER, UNDEFINED>]
              Schedule[]
                0: InlinedKern[]
                  Schedule[]
                    0: Assignment[]
                      ArrayReference[name:'zun']
                      Reference[name:'ji']
                      Reference[name:'jj']
                      Reference[name:'jk']
                      BinaryOperation[operator:'MUL']
```

Conclusion

PSyclone enables automatic **Fortran to OpenCL** transformation for codes adhering to the **PSyKAI kernel-based** parallelism model. **Separation of concerns** and **performance-portability** are achieved by providing a recipe of code transformations.

Future work

- Generalise solution to support any numerical operations
- More performance portability IR transformations
- SYCL backend



Science and
Technology
Facilities Council

Hartree Centre

Thank you

sergi.siso@stfc.ac.uk