





Science and
Technology
Facilities Council



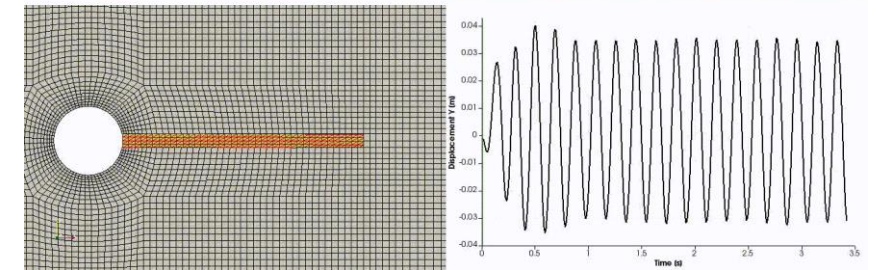

Using SYCL to accelerate the Multiscale Universal Interface coupling library

Mayank Kumar, Wendi Liu, Stephen M. Longshaw, Omar Mahfoze
UKRI Science & Technology Facilities Council, Daresbury Laboratory, UK



Motivation



- MUI is a portable, easy to use library that can couple softwares that simulate processes involving multiple phases and scales
 - Fluid-structure interaction (ParaSiF)
 - Molecular dynamics and DSMC (MNF)
- Exa-scale implementations of most scientific softwares are either ongoing or already complete
- Without a parallel implementation MUI potentially can be the limiting step in the complete simulation workflow



Science and
Technology
Facilities
Council

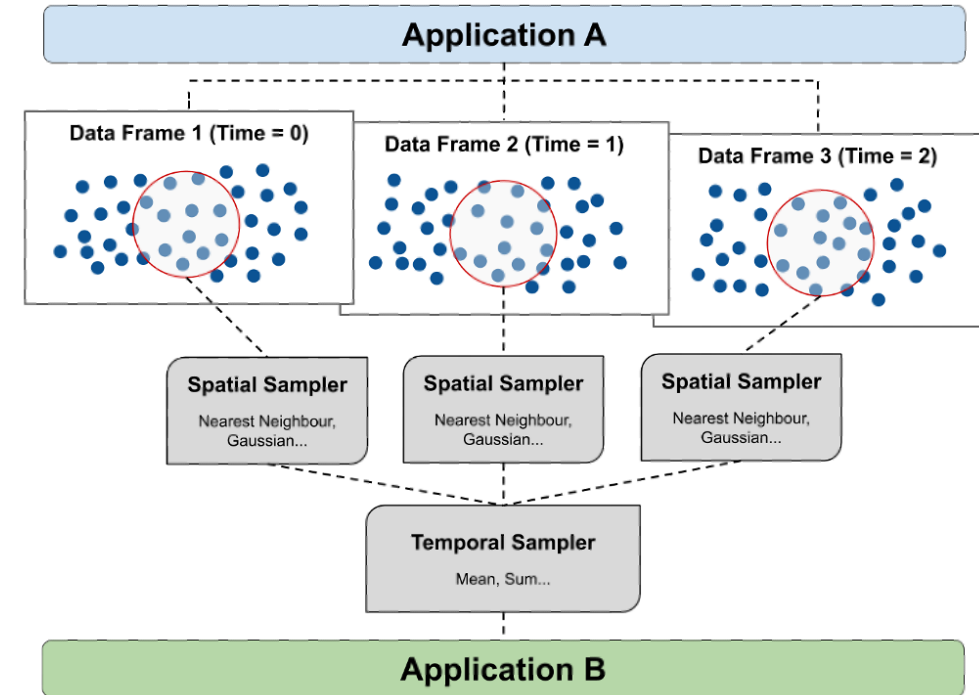
Time: 0.007 [s]

VoF: water
1.0
0.8
0.5
0.2
0.0
y
x



MUI

- A concurrent interface for coupling heterogeneous solvers^{1,2}
- A header-only code-coupling library written in C++
- MUI transfers data through MPI MPMD as a cloud of points, submitted as frames of time and provides spatial and temporal data sampling
- Scales to over 100,000 MPI ranks with parallel efficiency in excess of 90% depending on the codes



- Spatial Samplers
 - Exact
 - Linear
 - Quintic
 - RBF
- Temporal Samplers
 - Exact
 - Mean
 - Sum
- Algorithms
 - Fixed Relaxation
 - Aitken

RBF filter for FSI applications

Radial Basis Function

- General form of RBF³ on maintaining high-order consistent/conservative black-box coupling

$$s(\mathbf{r}) = \sum_{i=1}^N \alpha_i \varphi(\|\mathbf{r} - \mathbf{r}_i\|) + p(\mathbf{r})$$

- In MUI context:

$$s_i = \sum_{j=1}^M H_{i,j} \alpha_j \quad \text{OR} \quad \mathbf{s} = \mathbf{H} \boldsymbol{\alpha}$$

$$\mathbf{H} = \mathbf{A}_{as} \mathbf{C}_{ss}^{-1}$$

Since \mathbf{C}_{ss} is symmetric matrix:

$$\mathbf{H}^T = (\mathbf{A}_{as} \mathbf{C}_{ss}^{-1})^T$$

$$= (\mathbf{C}_{ss}^{-1})^T \mathbf{A}_{as}^T$$

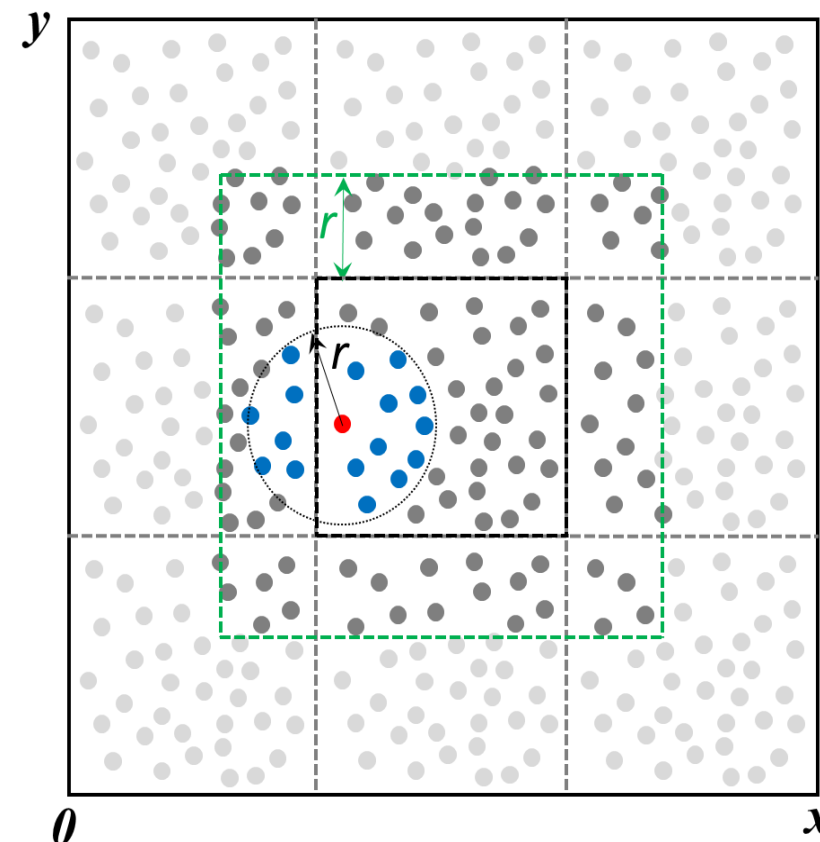
$$= \mathbf{C}_{ss}^{-1} \mathbf{A}_{as}^T$$

In which: $\mathbf{H}^T = \mathbf{C}_{ss}^{-1} \mathbf{A}_{as}^T$ can

be written as: $\mathbf{C}_{ss} \mathbf{H}^T = \mathbf{A}_{as}^T$

Conjugate Gradient solver

can be used to solve \mathbf{H}^T

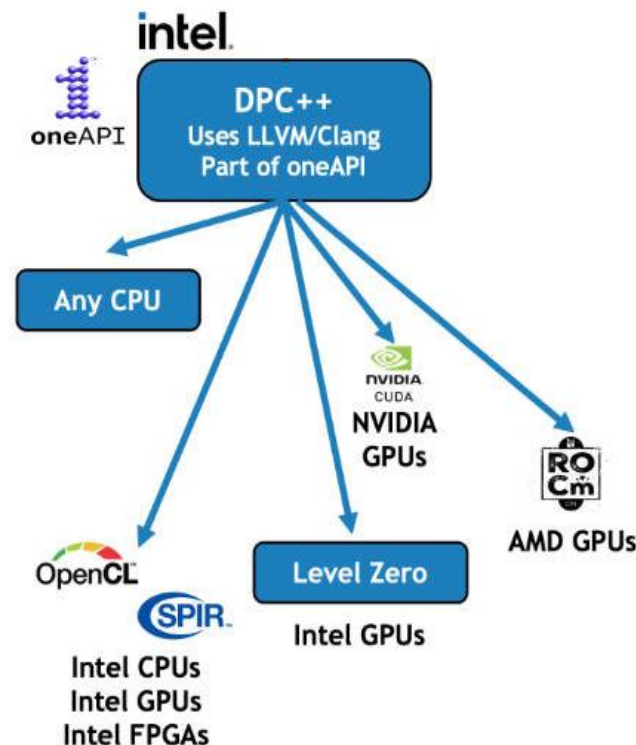


SYCL Programming model

SYCL (pronounced 'sickle') is a royalty-free, cross-platform abstraction layer.

Enables code for heterogeneous and offload processors to be written using modern ISO C++

Provides APIs and abstractions to find devices (e.g. CPUs, GPUs, FPGAs) on which code can be executed, and to manage data resources and code execution on those devices



SYCL USM

Allows reading and writing data using conventional pointers

Memory allocations can be explicit (device, host) or managed by SYCL runtime (shared)

SYCL BUFFER

Container for data accessed by both host and device

Memory managed by SYCL runtime using accessors

SYCL Implementation

CG Algorithm⁴

$$r_0 = b - Ax_0$$

$$p_0 = r_0$$

$$i=0$$

while ($r_i > \text{tol}$)

{

$$y_i = Ap_i$$

$$\alpha = r_i \cdot r_i / p_i \cdot y_i$$

$$x_{i+1} = x_i + \alpha p_i$$

$$r_{i+1} = r_i - \alpha y_i$$

$$\beta = r_{i+1} \cdot r_{i+1} / r_i \cdot r_i$$

$$p_{i+1} = \beta p_i + r_i$$

$$i = i+1$$

Serial Implementation

```
dotp_vec_vec()
{
    size_t rows = size_row;
    product = 0.;
    for (idx = 0; idx < rows; idx++)
    {
        product +=
            vec1_value[idx] * vec2_value[idx];
    });
    return (product);
}
```

Hotspots in MUI

- Generation of C_{ss} and A_{as} matrix
- Solution of H^T

SYCL Implementation

```
sycl_dotp_vec_vec()
{
    size_t rows = size_row;
    VTYPE *prod;
    prod = (VTYPE *)malloc(1);
    prod[0] = 0.;
    VTYPE *dotp;
    dotp = sycl::malloc_device<VTYPE>(1, defaultQueue);
    defaultQueue.memcpy(dotp, prod(sizeof(VTYPE))).wait();
    auto chg = [&](sycl::handler &hc)
    {
        hc.parallel_for(sycl::range(rows), [=](sycl::id<1> idx)
        {
            auto product = 0.;
            product = vec1_value[idx] * vec2_value[idx];
            auto v = sycl::atomic_ref<
                VTYPE, sycl::memory_order::relaxed,
                sycl::memory_scope::device,
                sycl::access::address_space::global_space>(*dotp);
            v.fetch_add(product);
        });
    });
    defaultQueue.submit(chg).wait();

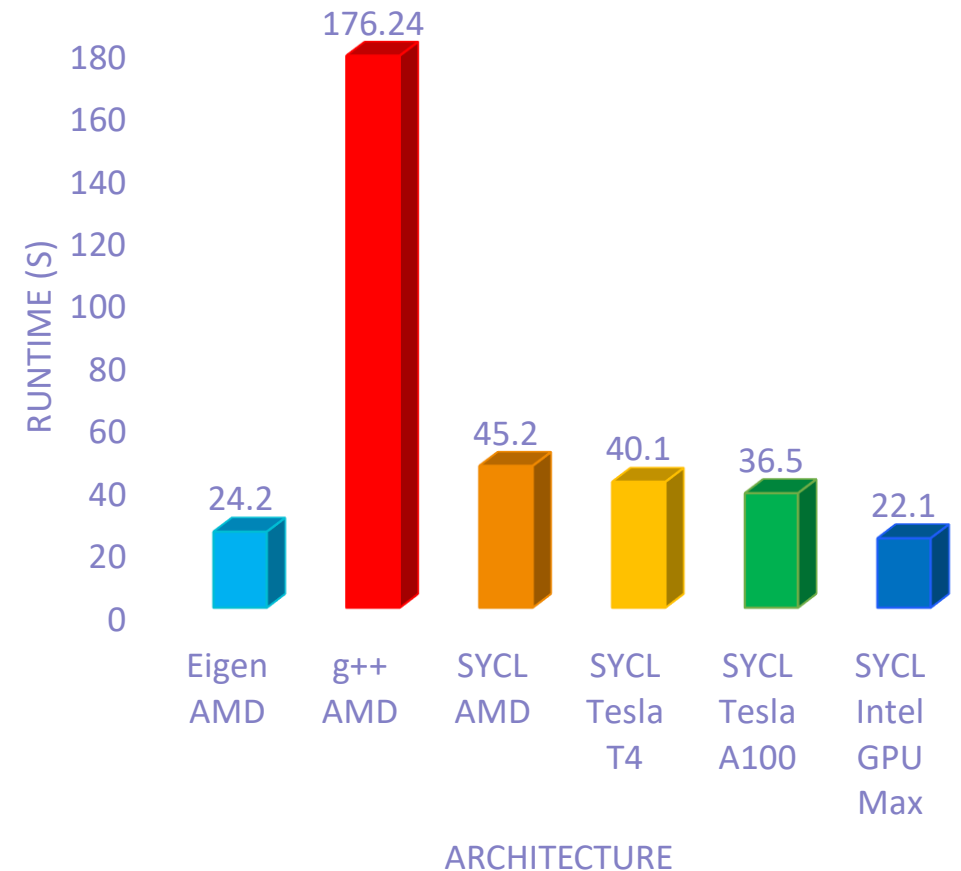
    defaultQueue.memcpy(prod, dotp, (sizeof(VTYPE))).wait();
    return (prod[0]);
}
```



Results

- Using a test matrix of 576*576 the SYCL code was tested for multiple architectures.
- The results are compared with benchmark linear solver Eigen running on 32 core AMD processor
- Profiling results show GPU occupancy less than 10%.

Matrix Vector product	52.7 %
Dot product	42.5%



Conclusion and future work

- SYCL provides a promising approach for the architecture independent parallelisation of MUI
- Further optimization of the linear algebra library
- Parallelization of C_{ss} and A_{as} matrix generation

References

- [1] Tang, Y.H., Kudo, S., Bian, X., Li, Z. and Karniadakis, G.E., 2015. Multiscale universal interface: a concurrent framework for coupling heterogeneous solvers. *Journal of Computational Physics*, 297, pp.13-31.
- [2] Tang, Y.H., Kudo, S., Longshaw, S.M., Liu, W., Skillen, A., Mahfoze, O., Bian, X., Li, Z., Karniadakis, G.E., Nash, R.W., Richardson, C., Manson-Sawko, R., 2023. The Multiscale Universal Interface 2.0, <https://doi.org/10.5281/zenodo.10054600>
- [3] Rendall, T.C.S. and Allen, C.B. 2009. Improved radial basis function fluid–structure coupling via efficient localized implementation. *Int. J. Numer. Meth. Engng.*, 78: 1188-1208. <https://doi.org/10.1002/nme.2526>
- [4] Baratta, I., Richardson, C., & Wells, G. 2022. Performance analysis of matrix-free conjugate gradient kernels using SYCL. Apollo - University of Cambridge Repository. <https://doi.org/10.17863/CAM.83190>