

Benchmarking and software sustainability at the exascale era

Preparing BigDFT for Aurora

Christoph Bauinger, Luigi Genovese



Notices & Disclaimers

Performance varies by use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

See backup for configuration details. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document, with the sole exception that code included in this document is licensed subject to the Zero-Clause BSD open source license (OBSD), <http://opensource.org/licenses/OBSD>.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

Motivation: Aurora

- Computer system at Argonne National Laboratory (ANL)
- >2 Exaflops FP64 theoretical peak performance
- 21,248 Intel Xeon CPU
- 63,744 Intel Data Center GPU

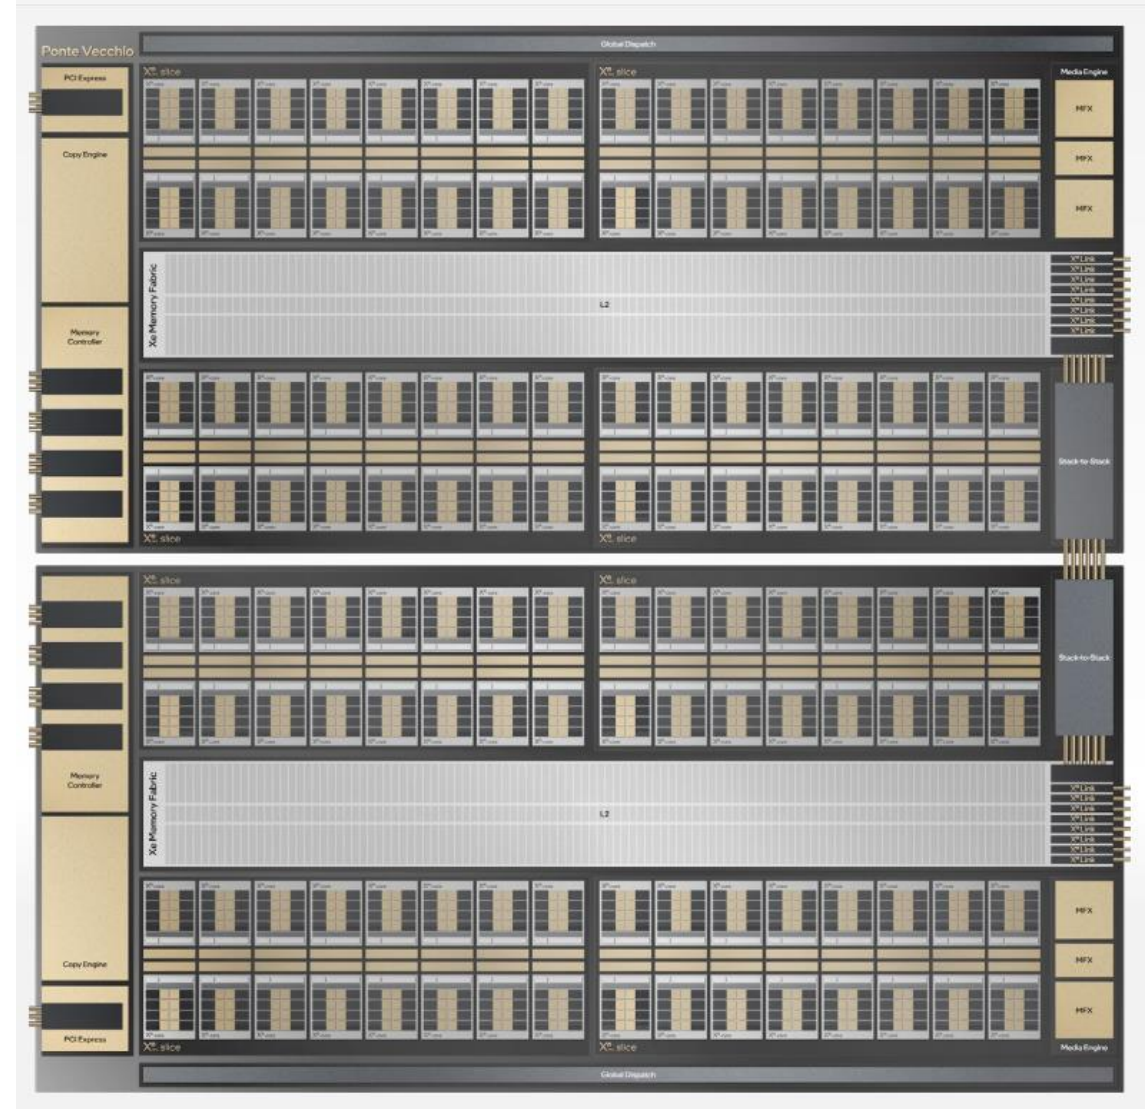


<https://www.anl.gov/article/us-department-of-energys-incite-program-seeks-proposals-for-2024-to-advance-science-and-engineering>, public domain

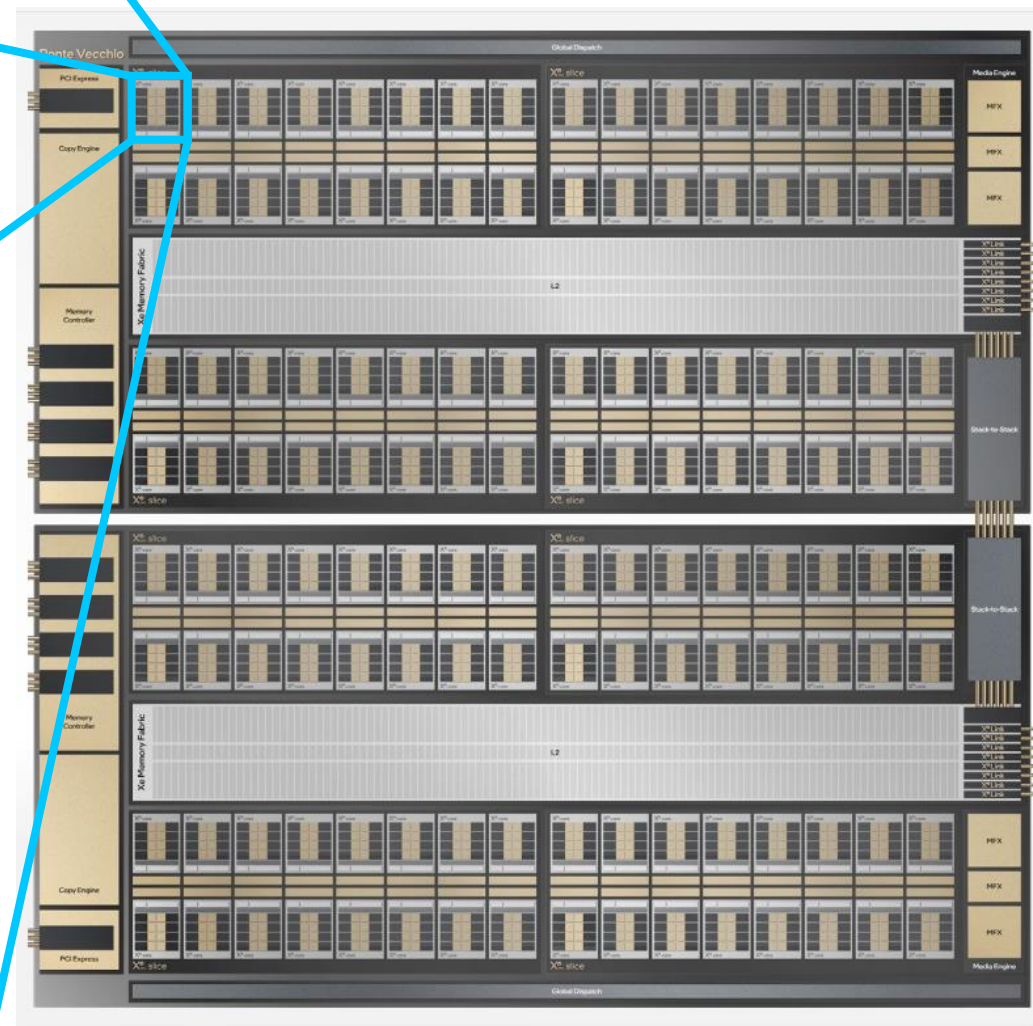
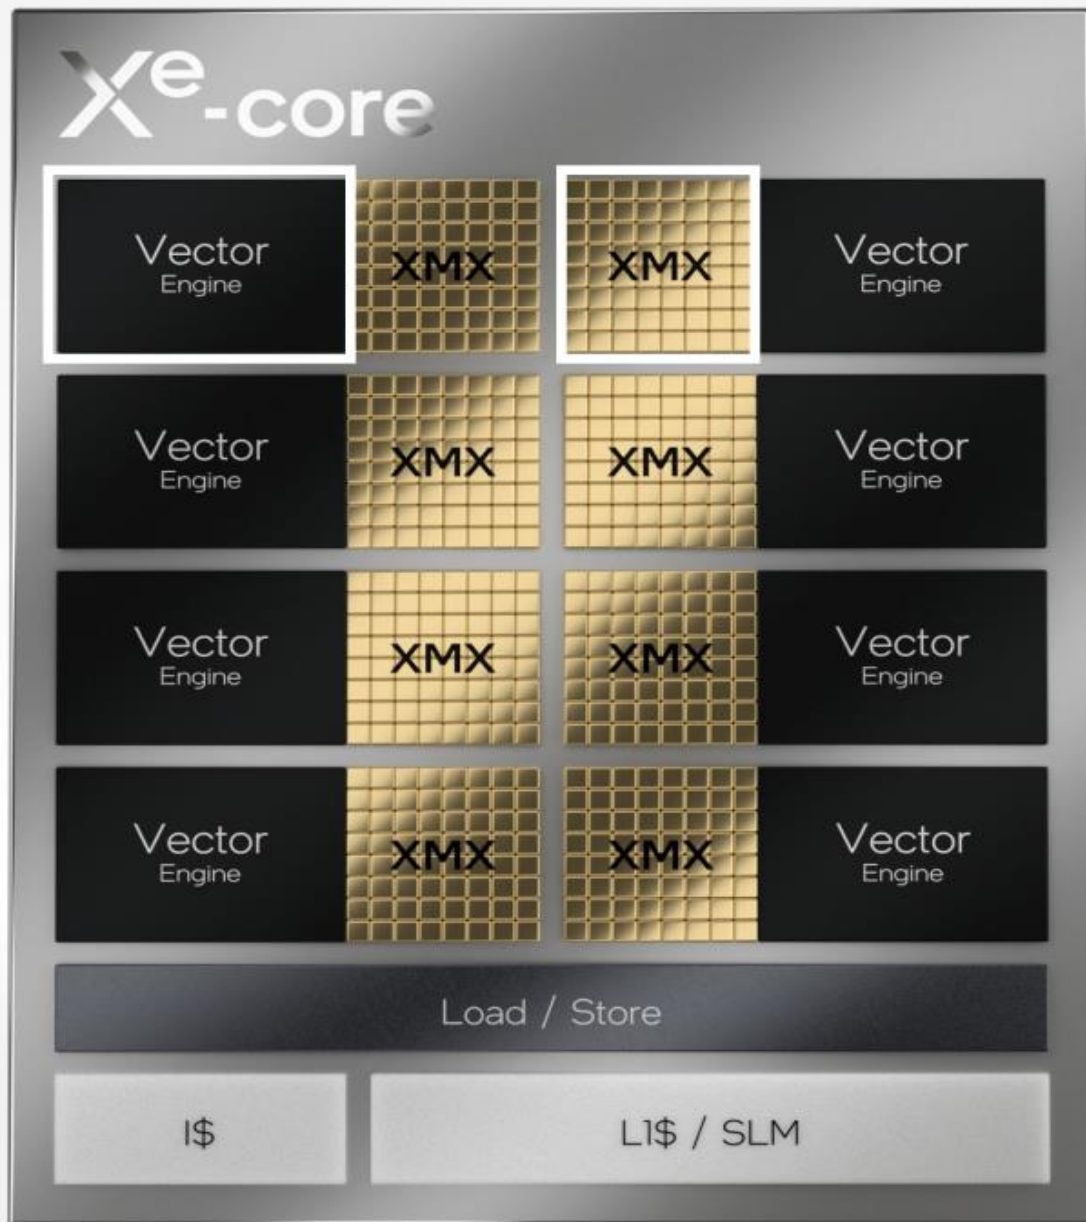
<https://www.alcf.anl.gov/aurora>

Intel Data Center GPU (PVC)

- 128GB HBM memory
- ~52 TFLOPS processing power (FP32 and FP64)
- TDP: 600W
- 2 Stacks



<https://www.intel.com/content/www/us/en/docs/oneapi/optimization-guide-gpu/2023-2/intel-xe-gpu-architecture.html#XE-HPC-2-STACK-DATA-CENTER-GPU-MAX>



<https://www.intel.com/content/www/us/en/docs/oneapi/optimization-guide-gpu/2023-2/intel-xe-gpu-architecture.html#XE-HPC-2-STACK-DATA-CENTER-GPU-MAX>

SYCL

- Standard by Khronos Group, initially released 2014
- Enables cross-platform single-source C++ code for heterogeneous and offload processors
- Implemented in Intel oneAPI DPC++
- Supported by several libraries (oneMKL, oneDNN, Kokkos, AMReX, etc.)

```
1  #include <sycl/sycl.hpp>
2
3  int main(int argc, char** argv)
4  {
5      sycl::queue Q(sycl::gpu_selector_v);
6      double *val_in = sycl::malloc_shared<double>(32, Q);
7      double *val_out = sycl::malloc_shared<double>(32, Q);
8
9      Q.fill(val_in, 1.1, 32);
10
11     Q.parallel_for(sycl::nd_range<1>(32, 32), [=](auto it) {
12         ...
13         const int i = it.get_global_linear_id();
14
15         val_out[i] = std::sqrt(val_in[i]);
16         ...
17     }).wait();
18
19     sycl::free(val_in, Q);
20     sycl::free(val_out, Q);
21
22     return 0;
23 }
```

A simple SYCL example

BigDFT

- Code which implements Kohn-Sham Density Functional Theory (KS-DFT)
- Fortran code with C++, OpenCL, and CUDA
- “PBE0” functional GPU enabled (CUDA, OpenCL) since 2009
- CUDA implementation shown to scale to several thousand NVIDIA nodes

Laura E Ratcliff, A Degomme, José A Flores-Livas, Stefan Goedecker, and Luigi Genovese. 2018. Affordable and accurate large-scale hybrid-functional calculations on GPU-accelerated supercomputers. *Journal of Physics: Condensed Matter* 30, 9 (feb 2018), 095901. <https://doi.org/10.1088/1361-648X/aaa8c9>

Laura E Ratcliff, Luigi Genovese, Hyowon Park, Peter B Littlewood, and Alejandro Lopez-Bezanilla. 2021. “Exploring metastable states in UO₂ using hybrid functionals and dynamical mean field theory”. *Journal of Physics: Condensed Matter* 34, 9 (dec 2021), 094003. <https://doi.org/10.1088/1361-648X/ac3cf1>

BigDFT – Fock operator evaluation

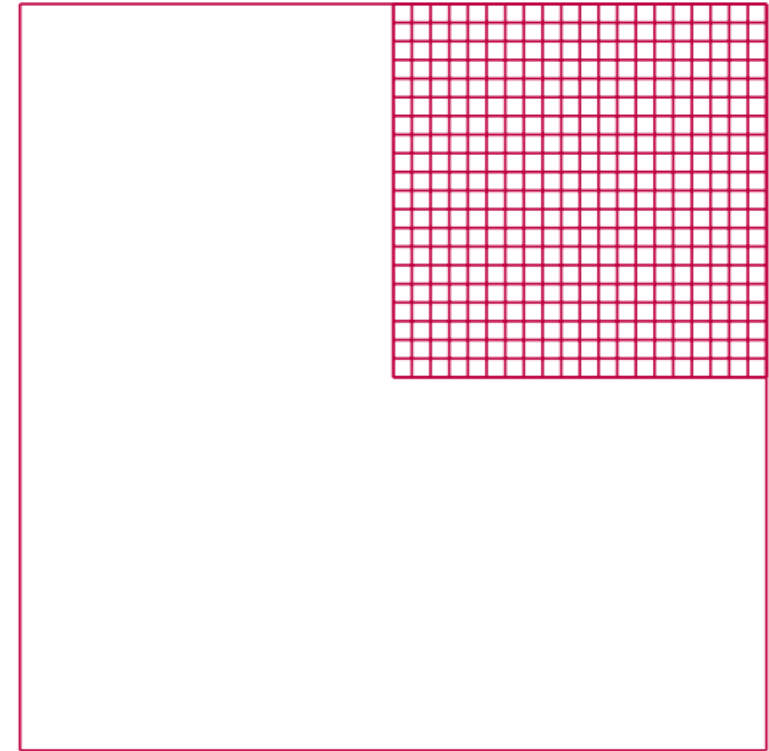
▪ Need to evaluate the Fock operator

$$[\hat{D}_X \psi_i](\mathbf{r}) = \sum_j \int d\mathbf{r}' \underbrace{\frac{\psi_j^*(\mathbf{r}') \psi_i(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|}}_{V_{i,j}(\mathbf{r})} \psi_j(\mathbf{r})$$

1. Compute $\rho_{i,j}(\mathbf{r}) = \psi_j^*(\mathbf{r}) \psi_i(\mathbf{r})$
2. Solve Poisson equation $\nabla^2 V_{i,j}(\mathbf{r}) = \rho_{i,j}(\mathbf{r})$
3. Multiply $V_{i,j}(\mathbf{r})$ with $\psi_j(\mathbf{r})$ to get Fock operator

BigDFT – Poisson Solver

- Uses convolution with Green function in Frequency Domain
- Performs convolution **as batched 1D FFTs**
- Can handle different boundary conditions
- Requires 0-padding in case of free boundary conditions



2D example of zero-padding

Nazim Dugan, Luigi Genovese, and Stefan Goedecker. 2013. A customized 3D GPU Poisson solver for free boundary conditions. Computer Physics Communications 184, 8 (2013), 1815–1820. <https://doi.org/10.1016/j.cpc.2013.02.024>

BigDFT – Poisson Solver

- Uses convolution with Green function in Frequency Domain
- Performs convolution **as batched 1D FFTs**
- Can handle different boundary conditions
- Requires 0-padding in case of free boundary conditions

Nazim Dugan, Luigi Genovese, and Stefan Goedecker. 2013. A customized 3D GPU Poisson solver for free boundary conditions. Computer Physics Communications 184, 8 (2013), 1815–1820. <https://doi.org/10.1016/j.cpc.2013.02.024>

```
1  if (<Boundary condition in x is free>)  
2      <zero-pad in x direction>  
3  1DFFT_X(Ny*Nz, Sx) //real-to-complex FFT  
4  if (<Boundary condition in y is free>)  
5      <zero-pad in y direction>  
6  <transpose data>  
7  1DFFT_Y((Sx/2+1)*Nz, Sy)  
8  if (<Boundary condition in z is free>)  
9      <zero-pad in z direction>  
10 <transpose data>  
11 1DFFT_Z((Sx/2+1)*Sy, Sz)  
12  
13 <Convolution kernel multiplication>  
14  
15 inverse_1DFFT_Z((Sx/2+1)*Sy, Sz)  
16 if (<Boundary condition in z is free>)  
17     <remove padding in z direction>  
18 <inverse transpose>  
19 inverse_1DFFT_Y((Sx/2+1)*Nz, Sy)  
20 if (<Boundary condition in y is free>)  
21     <remove padding in y direction>  
22 <inverse transpose>  
23 inverse_1DFFT_X(Ny*Nz, Sx) //complex-to-real  
24 if (<Boundary condition in z is free>)  
25     <remove padding in x direction>
```

Pseudo code of Poisson Solver in BigDFT

Porting BigDFT to SYCL – compatibility tool

- Existing accelerated code is CUDA and uses CuFFT
- We used Intel's DPC++ compatibility tool to automatically translate CUDA to SYCL
- We introduced an interface layer which chooses at runtime the CUDA or SYCL code
- oneMKL instead of CuFFT

```
1 //CUDA
2 __global__ void post_computation_kernel(int nx, int ny, int nz,
3     double *rho, double *data1, int shift1, double *data2,
4     int shift2, double hfac) {
5     int tj = threadIdx.x;
6     int td = blockDim.x;
7     int blockData = (nx*ny*nz) / (gridDim.x*gridDim.y);
8     int jj = (blockIdx.y*gridDim.x + blockIdx.x)*blockData;
9
10    for (int k=0; k<blockData/td; k++) {
11        int idx = jj + tj + k*td;
12        data1[idx+shift1] = data1[idx+shift1] +
13            hfac*rho[idx]*data2[idx+shift2];
14    }
15 }
16
17 //SYCL
18 void post_computation_kernel(int nx, int ny, int nz,
19     double *rho, double *data1, int shift1, double *data2,
20     int shift2, double hfac, const sycl::nd_item<3> &item) {
21     int tj = item.get_local_id(2);
22     int td = item.get_local_range(2);
23     int blockData = (nx*ny*nz) /
24         (item.get_group_range(2)*item.get_group_range(1));
25     int jj = (item.get_group(1)*item.get_group_range(2) +
26         item.get_group(2))*blockData;
27
28     for (int k=0; k<blockData/td; k++) {
29         int idx = jj + tj + k*td;
30         data1[idx+shift1] = data1[idx+shift1] +
31             hfac*rho[idx]*data2[idx+shift2];
32     }
33 }
```

CUDA kernel and automatically translated SYCL kernel.

Porting BigDFT to SYCL – BBFFT

- Goal: Further increasing SYCL performance on CPU
- Transposing and padding data is memory-bandwidth intensive and problematic on CPU
- Introduced double-batched FFT library

```
1  if (<Boundary condition in x is free>)
2      <zero-pad in x direction>
3  1DFFT_X(Ny*Nz, Sx) //real-to-complex FFT
4  if (<Boundary condition in y is free>)
5      <zero-pad in y direction>
6  <transpose data>
7  1DFFT_Y((Sx/2+1)*Nz, Sy)
8  if (<Boundary condition in z is free>)
9      <zero-pad in z direction>
10 <transpose data>
11 1DFFT_Z((Sx/2+1)*Sy, Sz)
12
13 <Convolution kernel multiplication>
14
15 inverse_1DFFT_Z((Sx/2+1)*Sy, Sz)
16 if (<Boundary condition in z is free>)
17     <remove padding in z direction>
18 <inverse transpose>
19 inverse_1DFFT_Y((Sx/2+1)*Nz, Sy)
20 if (<Boundary condition in y is free>)
21     <remove padding in y direction>
22 <inverse transpose>
23 inverse_1DFFT_X(Ny*Nz, Sx) //complex-to-real
24 if (<Boundary condition in z is free>)
25     <remove padding in x direction>
```

Remove additional memory accesses, where applicable

Porting BigDFT to SYCL – BBFFT

- Bbfft allows zero-padding on-the-fly with callbacks
- Avoids transposing data due to double-batching
- In addition, utilize symmetries to half memory requirements for kernel
- Halved memory requirements
- More than halved computing time on 4th generation Xeon CPU

```
1  if (<Boundary condition in x is free>)
2      <set callback to add zeros>
3  1DFFT_X(Ny*Nz, Sx) //real-to-complex FFT
4  if (<Boundary condition in y is free>)
5      <set callback to add zeros>
6  1DFFT_Y((Sx/2+1)*Nz, Sy)
7  if (<Boundary condition in z is free>)
8      <set callback to add zeros>
9  1DFFT_Z((Sx/2+1)*Sy, Sz)
10
11  <Convolution kernel multiplication>
12
13  inverse_1DFFT_Z((Sx/2+1)*Sy, Sz)
14  inverse_1DFFT_Y((Sx/2+1)*Nz, Sy)
15  inverse_1DFFT_X(Ny*Nz, Sx) //complex-to-real
```

Pseudo code of Poisson Solver with bbfft in BigDFT

Benchmarking results – Test systems

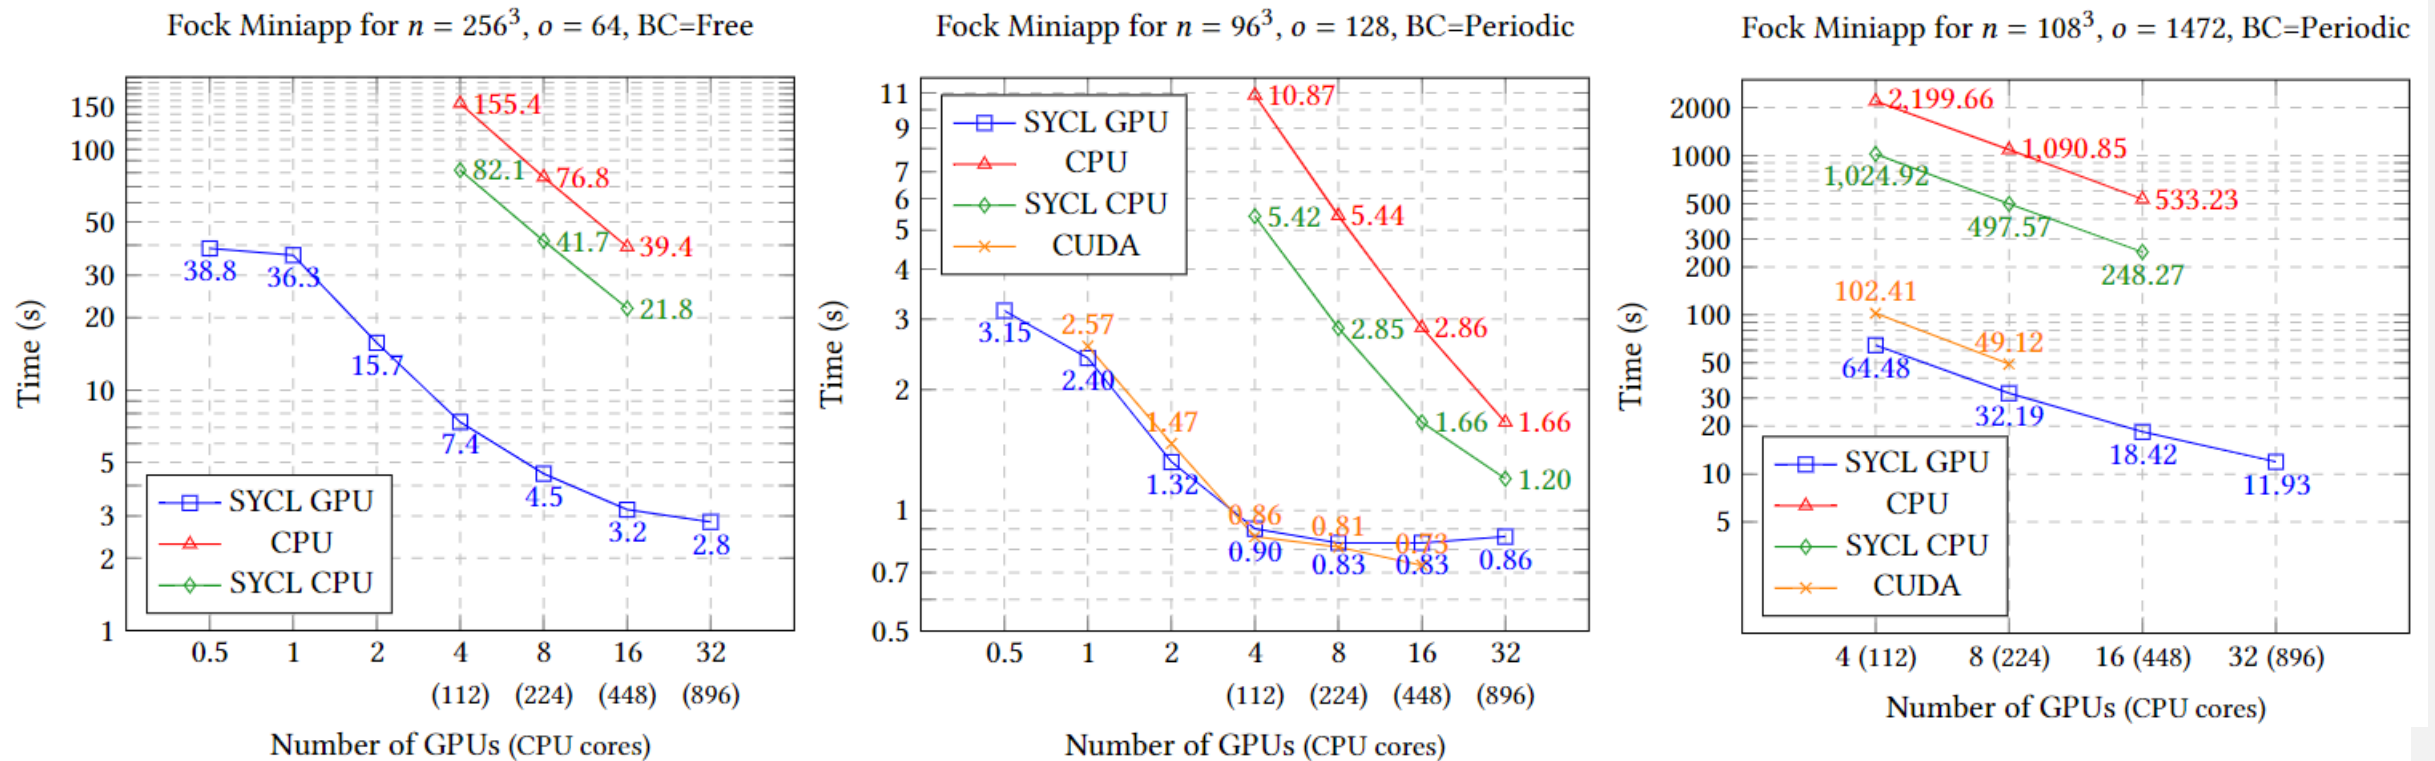
- PVC Test system: Florence
 - Lenovo SD650-I V3 servers with 2 Intel Xeon Platinum 8480+ Processors and 4 Intel Data Center GPU Max Series, 512 GB RAM
 - Each PVC is power-capped to 450 W
 - GPUs connected with XeLink
- NVIDIA Test system:
 - 2 Intel Xeon Platinum 8360Y Processors and 2 NVIDIA A100 40GB GPUs, 256 GB RAM
 - No NVLink
- Host code compiled with ifx, icx, icpx, intel oneAPI 2023.1
- Intel MPI 2021.9 in all cases

<https://ark.intel.com/content/www/us/en/ark/products/231746/intel-xeon-platinum-8480-processor-105m-cache-2-00-ghz.html>

<https://www.intel.com/content/www/us/en/developer/articles/technical/intel-data-center-gpu-max-series-overview.html?wapkw=tuscany>

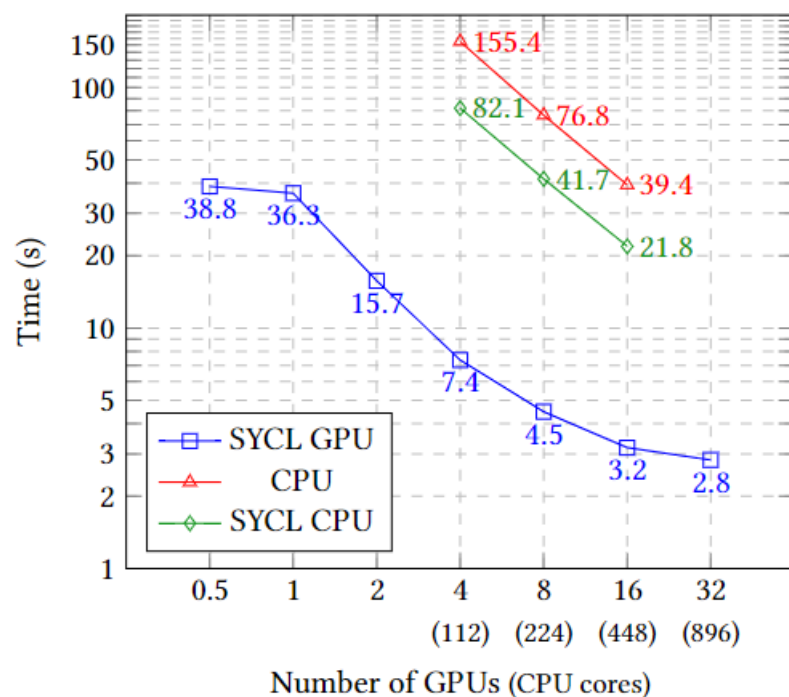
Benchmarking results – Fock miniapp

- Fock miniapp is a small test code to evaluate performance and correctness of the Fock operator evaluation

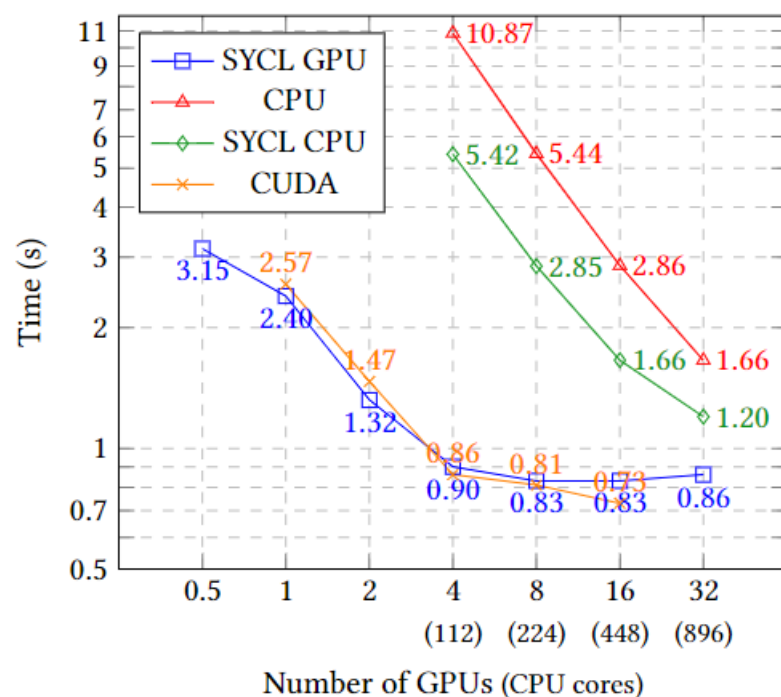


Benchmarking results – Fock miniapp

Fock Miniapp for $n = 256^3$, $o = 64$, BC=Free

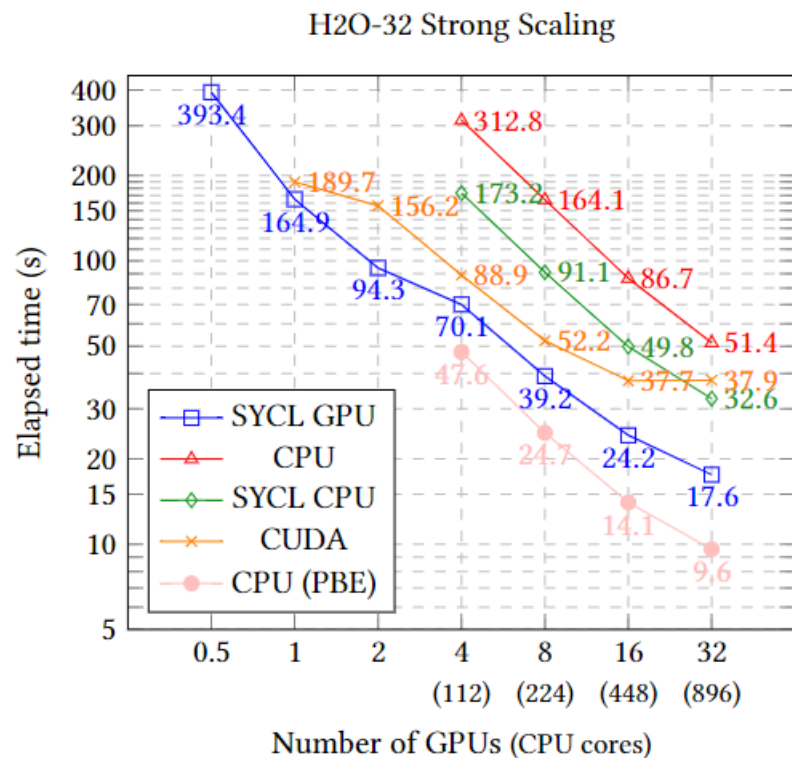


Fock Miniapp for $n = 96^3$, $o = 128$, BC=Periodic

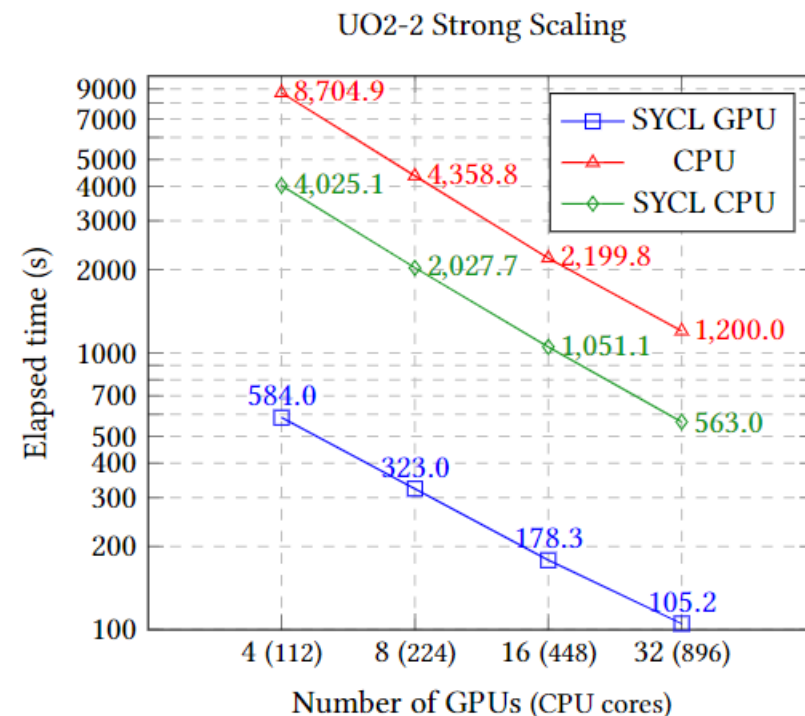


Case	Read (GB/s)		Write (GB/s)		Read+Write	
	L3	HBM	L3	HBM	L3	HBM
$n = 256^3, o = 64$	0.02	468	336	315	336	783
$n = 96^3, o = 128$	1172	66	1744	48	3516	114

Benchmarking results – Full App



32 H2O molecules. Performance delta between NVIDIA and Intel GPUs purely on the CPU side due to different processors



UO2 molecules with 1432 orbitals and a grid size of $n=108 \times 108 \times 108$. Elapsed time is for 3 iterations.

Performance Portability

- Goal: Run SYCL code on all backends with great performance
 - Intel CPU
 - Intel GPU
 - AMD CPU
 - AMD GPU
 - NVIDIA GPU
- Problem: Lacking support from bbfft and oneMKL

Summary and Future Work

- Shown that CUDA to SYCL port is easy even in big projects
- Shown that SYCL+PVC performs and scales well, indicating readiness for Aurora
- Shown that SYCL+CPU is a viable alternative to existing Fortran+OMP implementation
- Continue working on AMD and NVIDIA support with our SYCL implementation
- Run large simulations on Aurora

The Intel logo is centered on a solid blue background. It features the word "intel" in a white, lowercase, sans-serif font. A small, light blue square is positioned above the first vertical stroke of the letter 'i'. To the right of the word "intel" is a small white registered trademark symbol (®).

intel®