

DE LA RECHERCHE À L'INDUSTRIE



[www.cea.fr](http://www.cea.fr)

# Evaluating kernels on Xeon Phi to accelerate Gysela application


J. Bigot, M. Haefele, G. Latu

**CEA/DSM/IRFM & Maison de la Simulation**

and coworkers:


T. Cartier-Michaud, G. Dif-Pradalier, C. Ehrlacher, D. Estève,  
X. Garbet, P. Ghendrih, V. Grandgirard, C. Norscini, C. Passeron,  
F. Rozar, Y. Sarazin

**& INTEL Exascale lab & IPP+HLST Garching**



- Introduction to Gysela code
  
- Porting on Phi
  - Initial Setup
  - Micro-benchmarks
  - Interpolation kernels
  - 4D advections in Gys-ptoapp


- ▶ Main features in GYSELA (Gyrokinetic Semi-Lagrangian code):
  - ▶ Modelling of Tokamak plasma (targeting ITER)
  - ▶ Describing turbulence and transport (ITG instabilities)
    - turbulence *governs/limits plasma performance*
  - ▶ Main equations: Vlasov 5D, Poisson 3D (quasineutrality)
    - gyrokinetic setting (5D = 3D space + 2D velocity)
  - ▶ Heat & vorticity sources
    - (mimics heating system)
  - ▶ Collisional operator
  - ▶ Modelling fast particles
  - ▶ Adiabatic electron response



- ▶ Main unknown:  $\bar{f}^n(r, \theta, \varphi, v_{\parallel}, \mu)$

**Input** : *Physics parameters,  $\bar{f}^0$*

**Output** : *Diagnostics*



**for** time step  $n \geq 0$  **do**

Integrals:  $N_i^n(r, \theta, \varphi) = \int \int \bar{f}^n B(r, \theta) \mathcal{J}(k_{\perp} \rho_C) dv_{\parallel} d\mu;$

Push fields (**Poisson** Eq.):  $N_i^n(r, \theta, \varphi) \rightarrow \Phi^n(r, \theta, \varphi);$

Diagnostics for time step  $n;$

Push particles (**Vlasov** Eq. + **other** terms):  $\Phi^n(r, \theta, \varphi), \bar{f}^n \rightarrow \bar{f}^{n+1};$

**Algorithm 1:** Overall simplified Gysela algorithm

- ▶ **Fortran 90** code, hybrid **MPI+OpenMP**

- ▶ *Simplified* view of gyrokinetic Vlasov equation (dir. splitting):

$$\frac{\partial \bar{f}}{\partial t} + \frac{dr}{dt} \frac{\partial \bar{f}}{\partial r} + \frac{d\theta}{dt} \frac{\partial \bar{f}}{\partial \theta} + \frac{d\varphi}{dt} \frac{\partial \bar{f}}{\partial \varphi} + \frac{dv_{\parallel}}{dt} \frac{\partial \bar{f}}{\partial v_{\parallel}} = 0 \text{ (collisionless)}$$


- ▶ Solved through advections, **Semi-Lagrangian** scheme:

$$\partial_t \bar{f} + v_{\parallel} \partial_{\varphi} \bar{f} = 0 \text{ } (\hat{\varphi} \text{ operator})$$

$$\partial_t \bar{f} + \dot{v}_{\parallel} \partial_{v_{\parallel}} \bar{f} = 0 \text{ } (\hat{v}_{\parallel} \text{ operator})$$

$$\partial_t \bar{f} + \vec{v}_{GC} \cdot \vec{\nabla}_{\perp} \bar{f} = 0 \text{ } (\hat{r}\theta \text{ operator})$$

- ▶ Vlasov solver (**explicit** scheme) is composed of:
  - ▶ Successive directional splittings (advection steps)
  - ▶ Main cost of the application: interpolations (cubic splines)



- ▶  $\bar{f}$  conserved along characteristics
- ▶ Find the origin of the characteristics ending at the grid points (spatial grid)
- ▶ Interpolate value at origin  $X^*$  from known grid values: Cubic spline interpolation

3 steps for one advection:

- ▶ compute splines coefficients,
- ▶ compute feet (equations of motion),
- ▶ interpolate values.

**for** time step  $n \geq 0$  **do**

Integrals, Poisson, Diagnostics

Vlasov {

- 1D Advection in  $v_{\parallel}$  ( $\forall(\mu, r, \theta) = [local], \forall(\varphi, v_{\parallel}) = [*]$ );
- 1D Advection in  $\varphi$  ( $\forall(\mu, r, \theta) = [local], \forall(\varphi, v_{\parallel}) = [*]$ );
- Transposition of  $\bar{f}$ ;
- 2D Advection in  $(r, \theta)$  ( $\forall(\mu, \varphi, v_{\parallel}) = [local], \forall(r, \theta) = [*]$ );
- Transposition of  $\bar{f}$ ;
- 1D Advection in  $\varphi$  ( $\forall(\mu, r, \theta) = [local], \forall(\varphi, v_{\parallel}) = [*]$ );
- 1D Advection in  $v_{\parallel}$  ( $\forall(\mu, r, \theta) = [local], \forall(\varphi, v_{\parallel}) = [*]$ );


**Algorithm 2:** Parallel algo.: 2 domain decompositions

- ▶ no CFL for advections,  
comm. for transposition:


$$\Theta(N_r N_{\theta} N_{\varphi} N_{v_{\parallel}} N_{\mu})$$

$N_r = 512$ ,  $N_\theta = 512$ ,  $N_\varphi = 128$ ,  $N_{v_{||}} = 128$ ,  $N_\mu = 32$  (main unknown  $\bar{f}_n = 1$  TB)

Execution time for one Gysela run (Strong Scaling)



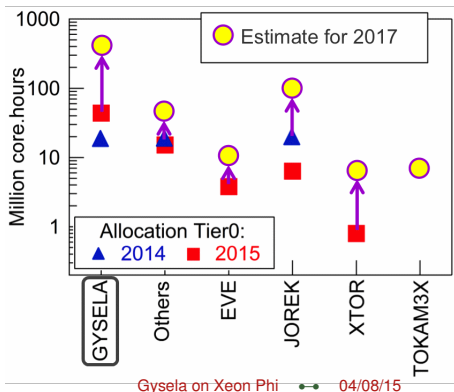
Relative efficiency for one Gysela run (Strong Scaling)



- ▶ Good result:  
78% relative efficiency on 64k cores  
(91% in Vlasov part)



- ▶ Fusion applications (and our institute CEA/DSM/IRFM) requires **computing power** for forthcoming years
- ▶ Supercomputers tends to provide more and more accelerators
  - candidates for next generation of parallel architectures
  - INTEL **Xeon Phi** and **GPGPUs** (AMD + Nvidia)



- Introduction to Gysela code
- Porting on Phi
  - Initial Setup
  - Micro-benchmarks
  - Interpolation kernels
  - 4D advections in Gys-protoapp

- Introduction to Gysela code
- **Porting on Phi**
  - Initial Setup
  - Micro-benchmarks
  - Interpolation kernels
  - 4D advections in Gys-protoapp

## Testbed: Helios machine (Fusion community, Japan)

Processor	Intel Xeon Sandy Bridge E5	Intel Xeon Phi 5110P
Clock frequency	2.1 - 2.8 GHz	1.05 - 1.238 GHz
Number of cores	8	60
Available memory	32 GB	8 GB
Peak performance (double precision)	173 GFlops/s	1011 GFlops/s
Sustainable memory bandwidth	40 GB/s	160 GB/s
Instruction execution model	out of order	in order
Simultaneous Multi Threading	2-way	4-way
Instruction set	x86-64 + 256bits-AVX	x86-64 + 512bits-SSE

### ▶ 2 programming models for Xeon Phi:

#### ▶ *offload* mode:

- ▶ Phi as an accelerator
- ▶ #pragma based

#### ▶ *native* mode:


- ▶ Phi as a linux node
- ▶ classical MPI + OpenMP

- ▶ Testbed (Helios machine - Fusion community, Japan)
  - ▶ Xeon Phi copro (5110P), 60 cores, 8GB mem., clock 1.05 Ghz
  - ▶ Sandy B. node (E5-2680), 2×8 cores, 64GB mem., clock 2.7 Ghz
- ▶ Initial assumptions on Xeon Phi
  - ▶ Easy to port code (x86 arch.)
  - ▶ Support OpenMP/MPI paradigm
  - ▶ How to get good performance ?
- ▶ Raw performance (×3 CPU peak, ×2 mem. BW)
  - ▶ SB: CPU peak 342 GFLOPS, mem bandwidth 70 GB/s (Stream triad)
  - ▶ Phi: CPU peak 1011 GFLOPS, mem bandwidth 130 GB/s (Stream triad)
- ▶ Approach in 4 steps:
  1. Direct port of a subset of Gysela: poor performance ☹️
  2. Memory and MPI benchmarks: inhomogeneous perf.
  3. Fallback: tune interpol. kernels (needed in Gysela), no MPI
  4. Try to put back a performant interpol. kernel into Gysela



- Introduction to Gysela code
- Porting on Phi
  - Initial Setup
  - **Micro-benchmarks**
  - Interpolation kernels
  - 4D advections in Gys-protoapp

# Memory Bandwidth



- ▶ Bandwidth on Xeon Phi
  - ▶ Up to 175 GB/s on the Xeon Phi with one thread/core
  - ▶ But 144 GB/s with 4 threads/core
  - ▶ With 2 or 3 threads/core, thread *affinity/pinning* does matter
  - ▶ x4 in mem. bandwidth compared to 1-socket S. Bridge
  - ▶ x2 in mem. bandwidth compared to 2-socket S. Bridge
  
- ▶ Latency Xeon Phi versus Sandy Bridge
  - ▶ **Similar** L1 latencies
  - ▶ **x4-x20 increase** otherwise on Xeon Phi (L2, L3, memory)
    - Cache reuse implementations have to target L1, L2
    - **Requires more efforts from the developer**
  
- ▶ Network performance (MPI communications)
  - ▶ Bandwidth **decreased** with Xeon Phi vs Sandy B.
  - ▶ Latency **increased** with Xeon Phi vs Sandy B.



- Introduction to Gysela code
  
- Porting on Phi
  - Initial Setup
  - Micro-benchmarks
  - **Interpolation kernels**
  - 4D advections in Gys-protoapp

- ▶ Parallelization strategy (no MPI) :
  - ▶ **native** mode choosen
    - ▶ because offload is slower (our tests on several configs)
    - ▶ avoid overhead due to Host-to-Phi data transfers (offload)
  - ▶ outer loops: **OpenMP**
  - ▶ inner block: loop vectorization through **SIMD directives**
- ▶ Code example: 1D advec - lagrange order 3 - on 4D data

```

1 #pragma omp parallel for collapse(3)
2   for (x1=0; x1<Nx1; x1++) {
3     for (x2=0; x2<Nx2; x2++) {
4       for (x3=0; x3<Nx3; x3++) {
5         #pragma vector nontemporal (f1)
6         #pragma vector always
7         for (x4=0; x4<Nx4; x4++) {
8           access_f(f1 ,x4,x3,x2,x1) = // OUTPUT data f1
9             coef1 * access_f(f0 ,x4-1,x3,x2,x1) + // INPUT data f0
10              coef2 * access_f(f0 ,x4 ,x3,x2,x1) +
11              coef3 * access_f(f0 ,x4+1,x3,x2,x1) +
12              coef4 * access_f(f0 ,x4+2,x3,x2,x1);
13         } } } }
    
```

- ▶ Vectorization through 512-bit MIC intrinsics  
only C language, Fortran is not accessible
- ▶ Code example: 1D advec - lagrange order 3 - on 4D data

```

1 #pragma omp parallel for collapse(3)
2 for (x1=0; x1<Nx1; x1++) {
3   for (x2=0; x2<Nx2; x2++) {
4     for (x3=0; x3<Nx3; x3++) {
5       for (x4=0; x4<Nx4; x4+=8) {
6         pthread = &(access_f(f0 ,x4 ,x3 ,x2 ,x1 ));
7         // read input data
8         tmp2 = _mm512_load_pd (pthread);
9         tmp1 = _mm512_loadunpacklo_pd (tmp1 , pthread -1);
10        tmp1 = _mm512_loadunpackhi_pd (tmp1 , pthread -1+8);
11        tmp3 = _mm512_loadunpacklo_pd (tmp3 , pthread +1);
12        tmp3 = _mm512_loadunpackhi_pd (tmp3 , pthread +1+8);
13        tmp4 = _mm512_loadunpacklo_pd (tmp4 , pthread +2);
14        tmp4 = _mm512_loadunpackhi_pd (tmp4 , pthread +2+8);
15        // 1+2+2+2=7 flop per loop iteration
16        tmpw = _mm512_mul_pd (tmp1 , coeff1 );
17        tmpw = _mm512_fmadd_pd (tmp2 , coeff2 , tmpw );
18        tmpw = _mm512_fmadd_pd (tmp3 , coeff3 , tmpw );
19        tmpw = _mm512_fmadd_pd (tmp4 , coeff4 , tmpw );
20        // write output data
21        _mm512_store_pd (&(access_f(f1 ,x4 ,x3 ,x2 ,x1 )), tmpw);
22      } } }

```

► Code example: 2D advec - lagrange order 3 - on 4D data

```

1 #pragma omp parallel for collapse(3)
2   for (x1=0; x1<Nx1; x1++) {
3     for (x2=0; x2<Nx2; x2++) {
4       for (x3=0; x3<Nx3; x3++) {
5         #pragma vector nontemporal (f1)
6         #pragma vector always
7         for (x4=0; x4<Nx4; x4++) {
8           access_f(f1 ,x4 ,x3 ,x2 ,x1) =
9             coefb1 * (coefa1 * access_f(f0 ,x4-1,x3-1,x2 ,x1) +
10                    coefa2 * access_f(f0 ,x4  ,x3-1,x2 ,x1) +
11                    coefa3 * access_f(f0 ,x4+1,x3-1,x2 ,x1) +
12                    coefa4 * access_f(f0 ,x4+2,x3-1,x2 ,x1) ) +
13             coefb2 * (coefa1 * access_f(f0 ,x4-1,x3  ,x2 ,x1) +
14                    coefa2 * access_f(f0 ,x4  ,x3  ,x2 ,x1) +
15                    coefa3 * access_f(f0 ,x4+1,x3  ,x2 ,x1) +
16                    coefa4 * access_f(f0 ,x4+2,x3  ,x2 ,x1) ) +
17             coefb3 * (coefa1 * access_f(f0 ,x4-1,x3+1,x2 ,x1) +
18                    coefa2 * access_f(f0 ,x4  ,x3+1,x2 ,x1) +
19                    coefa3 * access_f(f0 ,x4+1,x3+1,x2 ,x1) +
20                    coefa4 * access_f(f0 ,x4+2,x3+1,x2 ,x1) ) +
21             coefb4 * (coefa1 * access_f(f0 ,x4-1,x3+2,x2 ,x1) +
22                    coefa2 * access_f(f0 ,x4  ,x3+2,x2 ,x1) +
23                    coefa3 * access_f(f0 ,x4+1,x3+2,x2 ,x1) +
24                    coefa4 * access_f(f0 ,x4+2,x3+2,x2 ,x1) )
25           } } } }

```

- 1D advection (constant/small displacement) 1D interp lagrange 3
  - ▶ **Phi** perf: 46 GFLOPS (5% peek), BW: **106** GB/s (81% stream)
  - ▶ **SB** perf: 25 GFLOPS (7% peek), BW: **57** GB/s (81% stream)
- 2D advection (constant/small displacement) 2D interp lagrange 3
  - ▶ **Phi** perf: 250 GFLOPS (25% peek), BW: **111** GB/s (85% stream)
  - ▶ **SB** perf: 134 GFLOPS (39% peek), BW: **59** GB/s (84% stream)

A factor  $\times 2$  is obtained on Phi compared to one full SB node  
match expected behaviour 😊

Performance on Phi is varying much (10% is common) 😞  
with domain size, and from one run to the other

1. 3D advection (constant displacement) 3D interp lagrange 3, 4D data
    - ▶ **Phi** perf: **228** GFLOPS (23% peek), BW: 25 GB/s (19% stream)
    - ▶ **SB** perf: **156** GFLOPS (46% peek), BW: 17 GB/s (25% stream)
  
  2. 4D advection (constant displacement) 4D interp lagrange 3, 4D data
    - ▶ **Phi** perf: **160** GFLOPS (16% peek), BW: 4.3 GB/s (3.3% stream)
    - ▶ **SB** perf: **145** GFLOPS (42% peek), BW: 3.9 GB/s (5.6% stream)
- **Hard/long** to get good perf. on complex kernels on Phi ☹️
  - **3D** stencil easier to optimize than **4D** stencil (complex memory pattern)
  - Speedup up to **x2** in best cases (Phi versus one SB node) 😊
  - Small modifications OR changing compiler version
  - bad vectorization by the compiler on Phi → slowdown by **x4** ☹️

- ▶ Prefetch (load data in advance) accelerates computation
  - especially on memory-bound kernels
  - save 20% exec time on 1d/2d kernels
- ▶ Cache blocking (loop tiling) is crucial
  - especially on compute-bound kernels
  - save exec time on 3d kernels (50% reduction on exec. time)
- ▶ Tune aligned data, avoid cache trashing
  - save 20% exec time on 1d/2d kernels
- ▶ Comparing similar **C** and **Fortran** kernels
  - not clear tendency
  - give better or worse exec. time depending on the kernel

- ▶ Internal compiler optim. impact perf.  
(much more on Phi than on SB) ☹️
  - compiler does not give comprehensive feedbacks
  - looking at **generated assembly** code is painful but helpful
  - splitting the body of loop into multiple loops lead to effective speedups
  
- ▶ Writing “assembly” version may speedup computation (C code only)
  - 512-bit intrinsics help especially on compute-bound kernels 😊
  
- ▶ Phi works well with 170 up to 240 *well-pinned* threads versus 16 threads for SB ☹️



- Introduction to Gysela code
- Porting on Phi
  - Initial Setup
  - Micro-benchmarks
  - Interpolation kernels
  - 4D advections in Gys-protoapp

- ▶ Goal: feasibility of porting Gysela on Phi & rough estimate of the performance
- ▶ Approach: design a simplified version of Gysela, named **Gys-protoapp**

## Gys-protoapp code

- ▶ Remove non-essential parts of the Gysela code
  - ▶ diagnostics, alternative implementations, collisions, sources (keep the smallest set of numerical kernels Vlasov+Poisson)
  - ▶ 50k loc (Gysela) → 14k loc (proto-app)
- ▶ Remove a lot of MPI communication schemes
  - ▶ Restrict a single  $\mu$  value (4D problem instead of 5D)
  - ▶ Single node execution (works with `mpirun -np 1`)
  - ▶ A simulation  $N_r = 128, N_\theta = 256, N_\phi = 32, N_{v||} = 64$ : 10 hours on one SB node
- ▶ Add a new Vlasov solver (4D advection algorithm)
  - ▶ Computation intensive kernel, well-suited for Xeon Phi

- ▶ Usual Vlasov solver uses directional splitting (i.e. 1D and 2D advection operators - mem. bound):  
 $(\hat{v}_{\parallel}/2, \hat{\varphi}/2, \hat{r}\hat{\theta}, \hat{\varphi}/2, \hat{v}_{\parallel}/2)$

- ▶ Design a new 4D advection approach (compute bound):

$\eta(r = *, \theta = *, \varphi = *, v_{\parallel} = *) \leftarrow$  compute **spline coeff.** from the  
 4D function  $f^n(r = *, \theta = *, \varphi = *, v_{\parallel} = *)$ ;

**for** All grid points  $(r_i, \theta_j, \varphi_k, v_{\parallel l})$  **do**

$(r_i, \theta_j, \varphi_k, v_{\parallel l})^* \leftarrow$  **foot** of characteristic that ends at  $(r_i, \theta_j, \varphi_k, v_{\parallel l})$  ;  
 $f^{n+1}(r_i, \theta_j, \varphi_k, v_{\parallel l}) \leftarrow$  **interpolate**  $f^n$  at location  $(r_i, \theta_j, \varphi_k, v_{\parallel l})^*$  using  $\eta$  ;

**Algorithm 3:** 4D semi-Lagrangian scheme

- ▶ Three main kernels (rough profiling given):
  - 4D spline interpolator (51 % of exec. time)
  - feet of characteristics (36 %)
  - spline coeff computation (10 %)
  
- ▶ 4D interpolator - vectorization opportunities:
  - 4D tensor product with stencil of size 4,  
per grid point: 595 FLOP, read 1 float, write 1 float
  - very high computational intensity

- ▶ Parallelization/Optimization strategy (no MPI) :
  - ▶ Phi native mode, OpenMP
  - ▶ Cache friendly: loop blocking (2 levels in  $v_{||}$  and  $\varphi$ )
  - ▶ Reuse feet stored into L2 cache (temporal locality)
  
- ▶ Code structure of *spline 4D advection* :

```

1 do ith_blk=0, nb_blk_th ! loop blocking in theta
2   do ivpar=0, Nvpar
3     call feet_computations_with_openmp(...)
4     !$OMP PARALLEL DO COLLAPSE(2)
5       do iphi=0, iNphi
6         do ith=ith_blk*th_bsize, (ith_blk+1)*th_bsize-1
7           call interpolations_vectorized_kernel(...);
8         end do
9       end do
10    end do
11  end do
    
```

```

1 #define R.BSIZE 8
2 subroutine interpolations_vectorized_kernel(..., spline coeff.)
3 do ir_outer=0, Nr, R.BSIZE
4   ! retrieve grid cell containing the foot, compute spline basis
5   !dir$ simd
6     do ir_inner=0,R.BSIZE-1
7       ir=ir_outer+ir_inner
8       r_foot=...; th_foot=...; vpar_foot=...; phi_foot=...;
9       ir_star      = map_on_grid(r_foot)
10      ith_star      = map_on_grid(th_foot)
11      ivpar_star    = map_on_grid(vpar_foot)
12      iphi_star     = map_on_grid(phi_foot)
13      sbasis(1:16)  = compute_spline_basis(*_star, *_foot)
14    end do
15    ! interpolate in combining spline basis and spline coeff.
16    psum(0:R.BSIZE-1) = 0.
17    do <nest_of_four_loops>
18    !dir$ simd
19      do ir_inner=0,R.BSIZE-1
20        coeff = load spline coeff. located at *_star (with unit stride)
21        psum(ir_inner) = psum(ir_inner) + coeff(...) * sbasis(...)
22      end do
23    end do
24    f1(ir_outer:ir_outer+R.BSIZE-1,ith,iphi,ivpar)=psum(0:R.BSIZE-1)
25  end do
26 end subroutine interpolations_vectorized_kernel

```

- ▶ 4D advection (**variable** displacement) 4D **cubic spline**, 4D data  
 $N_r = 128, N_\theta = 128, N_\varphi = 128, N_{v_{||}} = 64$ 
  - ▶ **Phi** perf: **80** GFLOPS (7% peek), BW: 2.7 GB/s (2.% stream)
  - ▶ **SB** perf: **33** GFLOPS (9% peek), BW: 1.1 GB/s (1.6% stream)
- Variable displacements → unpredictable mem. access (prefetch pb)
- Reduced performance compare to previous kernels  
 → variable displacements: costs induced by **integer computations**,  
 memory indirections  
 → memory accesses cannot always be well aligned
- Sensivity to intel compiler version  
 → SIMD instructions employed and optimizations performed are varying
- Quite a long way to get this optimized version ...

Put back the 4D kernel in Gys-protoapp on Xeon Phi:

- ▶ From first port on Phi, to optim. version, factor  $\times 14$  on exec. time ☺
  - $N_r = 128, N_\theta = 128, N_\varphi = 32, N_{v_{||}} = 64$
  - First port (one call to Vlasov solver): 45 s
  - Optim. version (one call to Vlasov solver): 3.2 s
- ▶ Execution time:  $\times 2$  larger on Phi than on SB (16 cores) ☹
  - Amdhal's law: others computations should be optimized also ...
    - 1) computation of the feet characteristics
    - 2) spline coeff. computations
- ▶ Optimizations was useful for running on SB ☺
  - Overall execution time: reduced by 30% up to 45% on typical cases
  - vectorization directives have some interesting collateral effect
  - Tuned 4D advection is competitive compared to classical Strang splitting



- ▶ Achieving good performance on Phi:
  - not impossible 😊, but harder than on Sandy Bridge
  - successful on simple interpolation kernel 😊
  - needs: vectorization, fine grain parallelism, cache, prefetch
  - interact with the compiler (look at the generated assembly code)
  - easier if *only* one small kernel needs to be optimized
  
- ▶ Gys-protoapp (reduced Gysela application):
  - Xeon Phi still 2× slower than Sandy Bridge (16 cores) 😞
  - Sandy Bridge perf. of Gys-protoapp improved (30%-45%) 😊

Paper in CEMRACS'14 proceedings, accessible at

<http://arxiv.org/abs/1503.04645>

Acknowledgments: T. Guillet (Intel Exascale lab) for fruitful discussions,  
IFERC-CSC, Rokkasho, Japan (Helios supercomputer),  
RZG, Garching, Germany (Mick machine)

**Input** :  $\bar{f}^*(r, \theta, \varphi, v_{||}, \mu)$

**Output** :  $\bar{f}^\circ(r, \theta, \varphi, v_{||}, \mu)$

```

for  $\underline{\mu}$  do in parallel MPI
  for  $\underline{r}$  do in parallel MPI
    for  $\underline{\theta}$  do in parallel MPI
      for  $\underline{\theta}$  do in parallel OpenMP
        for  $\underline{v}_{||}$  do
           $\Delta\varphi \leftarrow v_{||} \Delta t$ 
          Compute spline representation of  $\bar{f}^*(r, \theta, \varphi = *, v_{||}, \mu)$ 
          for  $\underline{\varphi}$  do
             $\bar{f}^\circ(r, \theta, \varphi, v_{||}, \mu) =$ 
              spline_interpolate( $\bar{f}^*(r, \theta, \varphi - \Delta\varphi, v_{||}, \mu)$ )

```

**Algorithm 4:** Advection in variable  $\varphi$  on  $\bar{f}^*$

# MPI communication performance

Ping pong benchmark from the Intel MPI Benchmark (IMB)

Host0	CPU1	0.69		
	MIC0	4.90	2.73	
	MIC1	4.31	7.56	3.12
Host1	CPU1	2.20		
	MIC0	4.71	9.04	
	MIC1	4.66	7.93	6.92
	CPU1		MIC0	MIC1
		Host0		

Latencies ( $\mu$ s)

Host0	CPU1	5029		
	MIC0	456	2016	
	MIC1	1609	416	2004
Host1	CPU1	5729		
	MIC0	418	273	
	MIC1	1608	418	969
	CPU1		MIC0	MIC1
		Host0		

Bandwidth (MB/s)

- Similar results on supermic (LRZ, Garching, Germany), Eurora (Cineca, Italy), Robin (Bull R&D, Grenoble, France)
- In green: Typical performance for Infiniband
- In red: Low and non homogeneous network performance