

# Gyrokinetic Simulations on Many- and Multi-core Architectures with the Global Electromagnetic Particle-In-Cell Code ORB5

Noé Ohana<sup>a</sup>, Claudio Gheller<sup>a</sup>, Emmanuel Lanti<sup>a</sup>, Andreas Jocksch<sup>b</sup>, Stephan Brunner<sup>a</sup>, Laurent Villard<sup>a</sup>

<sup>a</sup>Swiss Plasma Center (SPC), École Polytechnique Fédérale de Lausanne (EPFL), Station 13, 1015 Lausanne, Switzerland

<sup>b</sup>Swiss National Supercomputing Centre (CSCS), Via Trevano 131, 6900 Lugano, Switzerland

---

## Abstract

Gyrokinetic codes in plasma physics need outstanding computational resources to solve increasingly complex problems, requiring the effective exploitation of cutting-edge HPC architectures. This paper focuses on the enabling of ORB5, a state-of-the-art, first-principles-based gyrokinetic code, on modern parallel hybrid multi-core, multi-GPU systems. ORB5 is a Lagrangian, Particle-In-Cell (PIC), finite element, global, electromagnetic code, originally implementing distributed MPI-based parallelism through domain decomposition and domain cloning.

In order to support multi/many cores devices, the code has been completely refactored. Data structures have been re-designed to ensure efficient memory access, enhancing data locality. Multi-threading has been introduced through OpenMP on the CPU and adopting OpenACC to support GPU acceleration. MPI is further used in combination with the two approaches. The performance results obtained using the full production ORB5 code on the Summit system at ORNL, on Piz Daint at CSCS and on the Marconi system at CINECA are presented, showing the effectiveness and performance portability of the adopted solutions: the same source code version was used to produce all results on all architectures.

*Keywords:* Plasma physics, High performance computing, Gyrokinetics, Multithreading, GPU

---

## 1. Introduction

In magnetized plasmas, the quality of confinement is well known to be degraded by the presence of turbulence arising from various small-scale instabilities, causing cross-field transport to exceed by large factors the neoclassical estimates. This is the main cause of the large size required for a fusion reactor based on the magnetic confinement principle, and consequently of the high cost of such devices. Understanding the fundamental properties of this turbulence is therefore of paramount importance. On the theoretical side, the gyrokinetic theory has been established as the most appropriate and complete description of turbulence in hot, low-collisionality core plasmas. While this dynamical reduction of the full Vlasov-Maxwell system eliminates the fastest scales, solving the resulting gyrokinetic equations remains a problem of formidable complexity requiring sophisticated numerical approaches [1]: gyrokinetic equations consist of nonlinear equations for time-dependent distribution functions in 5D phase space coupled to an integro-differential system of time-dependent field equations in 3D. In order to solve increasingly complex problems, outstanding computational resources are required. The effective exploitation of state-of-the-art HPC architectures represents therefore a mandatory task to tackle the most challenging questions in plasma physics.

The pinnacle of modern HPC performance is reached by many-core and heterogeneous computing (based, for instance, on GPU or FPGA devices), and current trends suggest that some form of heterogeneous computing will continue to be prevalent in emerging HPC systems. Therefore, the ability to fully exploit new heterogeneous and many-core architectures is of paramount importance towards achieving optimal performance on modern HPC systems. On the other hand, with the increasing size and complexity of numerical simulations, it is of primary importance for scientists to be able to exploit all available hardware in emerging HPC environments to achieve maximum computational throughput and efficiency. Exploiting these various forms of novel hybrid architectures is non-trivial however, due to the challenges presented by mixed hardware computing and the increasing levels of architectural parallelism. New algorithms and numerical and computational solutions are required.

In the last years a massive effort has been devoted to enable scientific applications to hybrid systems in all scientific areas. For plasma physics, we can mention the GPU enabling work made for the GENE code [2], the CGYRO code [3], the GTC code [4, 5] and the XGC code [6, 7]. These works have shown that the effort made at porting gyrokinetic codes to GPU was beneficial: for example, the GTC code was ported on Titan and an acceleration factor up to 3 could be achieved by using the GPUs [8]. We note, however, that a thorough benchmark comparison of these

---

*Email address:* noe.ohana@alumni.epfl.ch (Noé Ohana)

gyrokinetic codes is beyond the scope of this paper.

Our target code is ORB5 [9, 10, 11, 12], a state-of-the-art, first-principles-based gyrokinetic code. It is a Lagrangian, delta- $f$ , Particle-In-Cell (PIC), finite element, global, electromagnetic code developed by the Swiss Plasma Center (SPC) in collaboration with the Max Planck IPP in Garching and Greifswald and the University of Warwick. Some of its unique features include the use of high order (up to cubic) B-spline basis functions for the field representation, based on the variational Galerkin formulation that allows for the straightforward use of curvilinear toroidal magnetic coordinates including the magnetic axis, flow-conserving source and noise control operators, an enhanced control variate and pullback schemes that solve the cancellation problem for electromagnetic simulations, an adaptive scheme that adjusts the number of Larmor points for the gyroaveraging operator and a field-aligned Discrete Fourier filter eliminating unphysical modes. Last but not least, even though it uses a  $\delta f$  representation as control variate, the ORB5 code is equivalent to a truly full- $f$  code, with the exception that the polarization density is linearized.

ORB5 has recently been completely refactored with new data structures and its parallelism enhanced. Particle data structures are now structure of arrays instead of array of structures as it was found to increase the performance by privileging contiguous memory access, both on the CPU and on the GPU. Also, a new data structure was introduced for the representation of points on the Larmor ring that resulted in enhanced modularity and performance. Originally a pure MPI code based on domain decomposition and domain cloning [13, 14], hybrid MPI/OpenMP and MPI/OpenACC parallel programming models have been introduced. Various multithreading algorithmic solutions have been implemented for the different kernels, in particular for the gyro-averaged charge and current deposition and field assignment.

A remarkable feature of this development is that it has resulted in a single source code version that can be run either on CPU-only or on GPU-equipped HPC systems. The refactoring has been accomplished adopting a modular approach, making the code easily extensible and maintainable. Continuous integration is supported, through the usage of Git for code versioning [15] and Jenkins for automated testing and verification [16].

There are other approaches for programming HPC systems in a portable way than the purely directive-based one adopted here. The most popular ones are library based. Kokkos/Cabana is one option which has been used for porting a plasma physics application [17]. It provides a multidimensional array data structure. A library with a concept similar to Kokkos is Raja [18]. SYCL is a combination of a library with a specialized compiler [19]. We preferred directive-based approach to avoid large parts of our applications becoming library specific.

Being a community-driven code, ORB5 needs to be effectively usable on a large variety of computing systems,

so that the scientists can exploit all the available supercomputing resources. Since the development of applications codes such as ORB5 spans a period of time longer than the typical timescale of HPC architecture evolution, our prime objective is to design a portable code which provides good performance on different HPC platforms, rather than fine-tuning the performance on a specific architecture. In this work, we present the results obtained on three cutting-edge supercomputing platforms: Summit at ORNL (Oak Ridge National Laboratory), Marconi at Cineca, and Piz Daint at CSCS (Swiss National Supercomputing Center). The performed tests are representative of typical production runs, and only the compilation procedure accounts for specific, architecture-related optimization, adaptable by any standard user. Too specific architecture-related customization has been avoided, privileging portability and usability instead of extreme performance optimization. All results in this paper have been obtained with a single source code version which is actually the full production code version. The numerical and physical parameters of the data sets correspond to real production cases. Most importantly, all timings reported include all host-device and all across-node data transfers.

The details of the ORB5 code and its parallel implementation will be given in Section 2, followed, in Section 3, by the presentation of the results obtained for different tests and benchmarks. Results will be discussed in Section 4 together with the main conclusions.

## 2. The ORB5 code

ORB5 is a global gyrokinetic PIC code solving for the Vlasov-Maxwell set of equations in Tokamak geometry. Its model is derived from a variational formulation with consistent ordering [20, 21]. It accounts for electromagnetic perturbations around a realistic MHD equilibrium [22] or an approximated “ad-hoc” one. Different options are available for the field solver, namely the long wavelength approximation, Padé approximation, or arbitrary wavelength [23]. One can simulate multiple gyrokinetic species, and drift-kinetic, adiabatic or hybrid electrons. Among others, the code also features inter- and intra- species collisions [24], heat sources [25], and strong flows [26]. A detailed description of ORB5 capabilities can be found in [12].

### 2.1. Algorithms and MPI implementation

In ORB5 a control variate scheme is used, meaning that it represents the distribution functions of plasma species as the sum of a time-independent function  $f_0$  and a time-dependent perturbation  $\delta f$ . The latter is discretized with numerical markers (also called particles in this paper). Those markers evolve in fields discretized with 3D finite element B-splines up to third order. ORB5 uses the magnetic coordinates  $(s, \theta^*, \varphi)$ , where  $s$  is the radial coordinate,  $\varphi$  the toroidal coordinate, and  $\theta^*$  the poloidal coordinate modified from the geometric poloidal angle so that

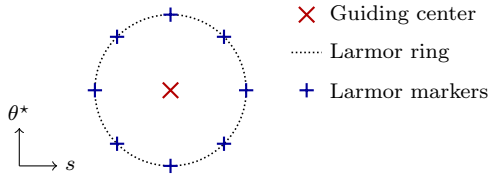


Figure 1: Larmor markers around a guiding center in the poloidal plane. The number of Larmor markers per guiding center can be fixed or proportional to the ring size.

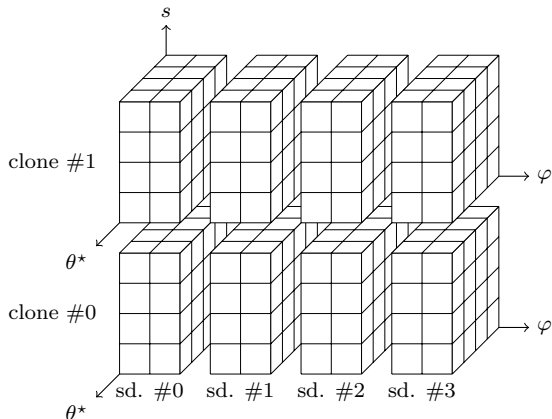


Figure 2: MPI 2-level parallelization of ORB5 in subdomains (sd.) and clones. Toroidal dimension can be decomposed into subdomains, and each of those subdomains can be cloned to split the particle workload over many MPI tasks.

field lines are straight in this coordinate system. The inner radial boundary condition can be Dirichlet, ‘free’ or unicity if the magnetic axis is part of the domain. The outer boundary condition is Dirichlet. The matrix problems resulting from the 3D finite element discretization of field equations are solved in discrete Fourier space in the poloidal and toroidal dimensions, allowing to take advantage of the anisotropy of the fluctuations to filter out unphysical perturbation modes which are not aligned with the field lines. This Fourier filtering reduces the required number of markers of one order of magnitude for the same noise level [27]. Moreover, thanks to the background and profiles axisymmetry, toroidal mode numbers are decoupled from each other in the field solvers, offering a trivial level of parallelism. In order to control the noise inherent to PIC simulations, ORB5 includes a Krook operator, a coarse graining technique, and quadtree smoothing. The electromagnetic cancellation problem in Ampère’s law can be cured using an enhanced control variate scheme [28] or the pullback scheme [29]. The gyro-averaging operations make use of numerical markers along the Larmor rings of each guiding center, as shown in Figure 1. The number of Larmor points per guiding center can be fixed or adaptive (scaling with the ring size).

ORB5 has two levels of MPI parallelism (Figure 2). The first one is domain decomposition along the toroidal dimension. Each subdomain works with field and particle data of a toroidal slice. This slicing is uniform in order

to guarantee a balanced workload. Hence, the number of subdomains is constrained to be a divider of the number of toroidal grid cells, which is frequently a power of 2 for best Fast Fourier Transform performance. There are three kinds of communications between subdomains. First, they send and receive particle data leaving or entering the subdomain. Particles can jump across several subdomains in a time step so it is an any-to-any communication. Due to the axisymmetry of the underlying physical problem, the number of markers per subdomain is guaranteed to be statistically uniform at all times. Second, subdomains communicate guard cell field data with their neighbors due to the finite element overlap over several grid cells. Third, parallel data transposition is required before performing local (to a CPU) Fourier transform in the toroidal direction. This operation is an all-to-all communication which swaps the partitioned dimension between toroidal and poloidal ones.

The second level of MPI parallelization in ORB5 is domain cloning. It consists in replicating subdomain field data on different tasks, thus splitting the particle workload without changing the grid resolution. MPI communications among clones are reductions of the grid data before carrying out the field solvers, and broadcasts afterwards. Domain cloning was introduced since the very first version of the ORB5 code. It allowed for running with more MPI tasks than the number of toroidal grid points. In the newly refactored, multithreaded version of ORB5 (see below), domain cloning is shown to be useful on multi-socket nodes, i.e. the optimum configuration is to use one clone per socket.

Instead of (or in addition to) domain cloning, we may consider 2D or 3D domain decomposition: the potential scalability gain of these options was investigated in Ref.[30]. While 1D decomposition in the toroidal direction poses no problem for load balancing, additional decomposition in the radial and poloidal directions is more challenging: ensuring particle load balance requires equal-volume domains, but due to the polar-like grid, this leads to off-balance grids. Moreover, the finite Larmor radius (FLR), lying essentially in the poloidal plane, introduces technical difficulties when a Larmor ring straddles two or more domains: this requires additional guard cells but unlike those required by the overlap of finite elements, the number of necessary guard cells depends on the largest Larmor radius of all markers in the domain and is not necessarily the same throughout the plasma, nor is it the same in all simulations. Although subject of investigation, the 2D or 3D domain decomposition has not yet been implemented in the ORB5 code.

The time integrator is a fourth-order Runge-Kutta.

## 2.2. Rewriting and refactoring strategy

Developers of legacy codes are confronted to the recurrent dilemma: should the code be entirely rewritten, or is refactoring a better option? We have adopted a two-pronged approach that combines both aspects of rewriting

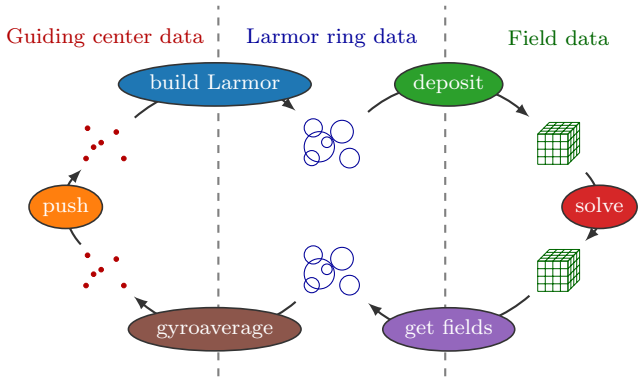


Figure 3: Main stages of an ORB5 time step. The 6-fold scheme differs from the usual PIC 4-fold scheme because of the introduction of the Larmor ring data.

and refactoring. First, we have extracted the fundamental kernels of the original ORB5 code and rewrote them from scratch. These then served as ‘test-bed’ codes, which we called ‘PIC-engine’ [31] and ‘GK-engine’ [32] for the pure PIC and for the specific gyrokinetics aspects, respectively. With respect to ORB5, the GK-engine misses the toroidal geometry, electromagnetic fluctuations, noise controls, heat sources, and multiple kinetic species, among other features. Meanwhile, the PIC-engine does not even include any field solver, nor gyrokinetic effects. In these test-beds, several options for data structures and multi-threading were implemented and tested. These include in particular the challenging particle-to-grid and grid-to-particle kernels. Then, the most performing options of the PIC-engine and GK-engine were implemented into the full production ORB5 code.

A standard PIC time step consists of 4 stages: (1) the deposition of particle charges and currents on the grid; (2) the field solve on the grid; (3) the interpolation of the field at particle positions; and (4) the particle push. As compared to standard PIC schemes, the finite Larmor radius (FLR) of the gyrokinetic description introduces an additional non-locality. In the original ORB5, data structures were attached to guiding centers and stages (1) and (3) were treated as follows: a loop on guiding centers, inside of which a loop on Larmor markers was performed. As the number of Larmor markers can be different for every guiding center, this nested loop combination performed weakly and proved difficult to optimize. Moreover, trying to improve data locality by sorting according to guiding center positions does not ensure locality of Larmor points.

Therefore, another strategy was adopted for the refactoring of the ORB5 code: a new data structure was introduced, named ‘Larmor ring data’, containing the positions in real space of the Larmor markers and guiding center weights. The main stages of an integration substep now consist of 6 stages instead of 4 and are depicted in Figure 3. With this, the operations related to FLR (‘build Larmor’ and ‘gyroaveraging’) can be completely separated from those of particle-to/from-grid (‘deposit’ and ‘get fields’)

for which now all optimizations pertinent to ‘standard’ (non-gyrokinetic) PIC algorithms can then be straightforwardly applied, including improving data locality by sorting the Larmor point according to their position in the grid. This strategy was first tested in the GK-engine and then, since it proved beneficial, was introduced in the refactored ORB5 code. The introduction of an intermediate Larmor data structure makes the code more efficient as well as more modular, at the expense of a higher memory footprint. Indeed, loops over Larmor points are taken out of the loops over guiding centers, leading to better vectorization. It also allows for the sorting of the Larmor markers by cells before depositing them, improving data locality and avoiding random memory accesses. Indeed, without sorting, the deposition routine loops over the markers and then over the finite elements they contribute to, while with sorting the outermost loop can be on the grid cells and the innermost one on the particles in that cell.

### 2.3. Multi/Many threads acceleration

Multi-threading for multi-core CPUs has been implemented by using the OpenMP application programming interface [33]. The current implementation follows version 3.1 standard. OpenMP represents a consolidated and stable solution to exploit shared memory devices, ensuring good scalability up to many cores and high portability, being supported by all the major hardware and software providers and implemented in the most common compilers of, in particular, Fortran, C and C++ programming languages. It is based on a set of directives which instruct the code on how to split the work among the available threads, with minimum impact on the source code, although specific customization has to be done in order to ensure good performance and scalability.

In ORB5, OpenMP is used exclusively in routines looping over the particles. Such loop iterations are in general independent from one another, which make them trivially parallelizable over OpenMP threads. When it comes to accumulation of some marker quantity on a grid, such as charge and current depositions or binning diagnostics, one has to take care about race conditions (several threads updating the same memory location simultaneously). For scalar quantities, we use the OpenMP built-in `reduction` clause. For arrays, we make each OpenMP thread accumulate to a temporary replica, and we reduce them afterwards. This manual operation was found to be more efficient than the `reduction` clause, especially for large arrays. Those techniques have been found to be much more efficient than `atomic` directives. The deposition algorithms using particle sorting avoid this grid replication by multithreading the outer loop over grid cells so that the inner loop over markers within a cell can be executed serially.

The GPU implementation has followed a similar approach, adopting the OpenACC programming model [34] to enable the GPU to carry out the most time-consuming

parts of the code. The current implementation reads revision 2.5. Just like OpenMP, OpenACC is a directive-based approach, which allows keeping the source code almost unchanged, leaving to the compiler the task of creating the GPU version. Contrary to OpenMP however, OpenACC is a new standard under continuous development and only few compilers, namely PGI and GNU, fully support its latest releases. Hence portability is still limited in this respect. On the HPC systems typically targeted by ORB5 (and, in particular, used in this work), the PGI compiler is expected to be available. The adoption of OpenACC is therefore not limiting its usability.

OpenACC provides two main classes of directives: those to manage work sharing and those to control data transfer. Both are designed with performance as the main objective. The OpenACC approach was tested on the simplified PIC-engine kernels [31] and several options regarding the data layout and multithreading were tried there, the most performant being implemented in ORB5 afterwards.

In ORB5, all the computations involving particles have been offloaded to the GPU: more precisely, these are the stages ‘push’, ‘build Larmor’, ‘deposit’, ‘get fields’ and ‘gyroaverage’, with only the ‘solve’ operations performed on the host. Loops over particles are particularly well suited for GPUs since millions of independent iterations can be executed by thousands of threads. Communication in between GPU and CPU is minimized by keeping particle data on the GPU only. More precisely, all Larmor ring data are kept strictly on the GPU, while only a fraction of the guiding center data, related to the markers that change domain, has to be transferred via MPI at every time step from GPU to GPU. Having a single kernel using particles on the CPU would annihilate the speed-up brought by the GPU due to the time required to transfer particle data between device and host.

The ‘solve’ stage consists of a computational part, namely the backsolve operations of the field solver and the MPI communication of grid data required by the toroidal domain decomposition and toroidal discrete Fourier transform, in addition to reductions of grid data over the clones. Field backsolve operations have not yet been ported to GPUs as they were initially not found to be the critical parts for our application. Since particle-to/from-grid operations are performed on the GPU, field data has to sit on both device and host. Therefore field data is transferred from the GPU to the CPU after the deposition, and back to the GPU after the field solve.

In the current ORB5 implementation, most of the loops multithreaded with OpenMP and OpenACC are the same. Therefore, when enabling OpenACC to benefit from GPUs, OpenMP is disabled. The situation could evolve in the future, for instance if the ‘solver’ stage gets multithreaded on the host with OpenMP. One can also think of getting rid of either OpenMP or OpenACC directives through the adoption of more recent standards. Indeed, OpenMP 4.5 provides GPU instrumentation directives,

and OpenACC 2.7 enables one to treat CPUs as devices. Those standards were not properly implemented by the compilers at the time of this work.

Concerning race conditions in deposition routines, the grid replication method used with OpenMP threads is not suitable anymore with GPU threads due to their large number. We therefore use OpenACC `atomic` directives around accumulation instructions in order to make them thread safe. The introduction of OpenACC `atomics` for our application on NVIDIA GPUs, where the atomic update has been effectively implemented at hardware level on Pascal and Volta architectures, does not lead to any significant time overhead compared to racy depositions. It is to be noted that these race conditions can be totally avoided by sorting particles as mentioned in the previous section (2.2) and multithreading the loop over grid cells.

### 3. Results

As stated above, all the results presented in this paper have been obtained with a single source code version, which is the full production ORB5 code version. Also, the simulation parameters used in this paper are representative of production runs. Most importantly, all timings reported include all host-device and node-to-node data transfers.

In order to analyze the performance and the scalability of our parallel, multi-threaded implementation of ORB5, we have performed a number of tests and benchmarks exploiting three different state-of-the-art architectures available at Cineca, CSCS and ORNL supercomputing centers. Cineca is the Italian national HPC facility. We have used the A3 partition of its *Marconi* flagship system, made of 3224  $2 \times 24$ -cores Skylake CPU nodes. We have exploited the GPU partition of *Piz Daint* at CSCS, the Swiss National Supercomputing Center. This partition is a Cray XC50 machine, with 5704 nodes equipped each with a 12-cores Intel Haswell CPU and one NVIDIA Tesla P100 GPU. Finally, we have run on the *Summit* supercomputer at Oak Ridge National Laboratory in the USA, deploying 4608 nodes, each equipped with two 21-cores IBM POWER9 CPUs and 6 NVIDIA Tesla V100 GPUs. Table 1 summarizes the characteristics of the systems and indicates the version of the adopted compilers as well as invoked compiler flags.

As physical and numerical simulation parameters we use gyrokinetic deuterium ions with an adaptive number of Larmor points as well as “heavy” kinetic electrons only 200 times lighter than ions. The pullback scheme is used to solve for the electromagnetic fluctuations. The Poisson equation makes use of the long wavelength approximation. A density-, zonal flow- and parallel flow- conserving Krook operator is used as noise control and heat source. The number of time steps has been set to be a small fraction of an actual simulation.

The single node tests of section 3.1, as well as those of sections 3.2 and 3.3 adopt a mesh of  $N_s \times N_{\theta^*} \times N_{\varphi} =$

System	Computing Node	$N_{\text{nodes}}$	Interconnect	Compiler	Compiler flags
Marconi	2 Intel Xeon 8160 CPU (Skylake)	3224	Fat-Tree Intel Omnipath	Intel 17.0.4.196	-O2 -xCORE-AVX2 -fopenmp
Piz Daint	Intel Xeon E5-2690 v3 (Haswell) CPU + NVIDIA Tesla P100 GPU	5704	Dragonfly Aries	PGI 18.10 + Cuda 9.1	-O3 -fast -Minline -Munroll -Mvect=levels:7,nosizelimit -Mcuda -acc -ta=tesla,cc60
Summit	2 IBM POWER9 CPUs + 6 NVIDIA Tesla V100 GPUs	4608	Fat-Tree EDR 100G InfiniBand	PGI 18.10 + Cuda 9.2	-O3 -fast -Minline -Munroll -Mvect=simd [-mp=nonuma   -Mcuda -acc -ta=tesla,cc70]

Table 1: Technical characteristics of the three adopted HPC platforms, the adopted compilers and the corresponding compilation options. Two sets of options have been used on Summit, enabling either OpenMP or OpenACC exclusively.

$512 \times 1024 \times 4$  cells and  $20 \cdot 10^6$  particles for each of the two simulated species (deuterium and electrons). This setup is representative of one subdomain, typically fitting on a single node, of a medium-size tokamak production run. Meshes and particle numbers relevant for production runs on multiple nodes will be used in section 3.4.

For all the tests, we have adopted as performance metric the time-to-solution and the fractions of time spent in the different parts of the code. Only time spent in the time loop has been considered, thus excluding initialization. In a typical production run there are at least tens of thousands of time steps, so that the initialization time is a negligible fraction of the total.

### 3.1. Single node MPI and OpenMP performance

ORB5 allows specifying different configurations of parallelism exploiting the domain decomposition, cloning and multithreading (on the CPU or on the GPU). A first set of tests has focused on the CPU implementation, analyzing how different configurations influence the performance and selecting Skylake as the reference architecture (the comparison among different architectures will be given in the next section).

Figure 4 shows the breakdown of the time to solution in the different components of the code, as introduced in section 2.1. Each bar represents a different combination of number of MPI subdomains (domains in which the grid is partitioned along the toroidal direction), number of MPI clones (replicas of a domain), and number of OpenMP threads per MPI task. The product of numbers of subdomains, clones and threads equals the number of CPU cores (48 in this case).

The first three bars show the timings for one MPI subdomain, the second two bars for 2 subdomains, and the last two bars for 4 subdomains. The first bar of each group uses pure MPI parallelization, without multithreading. The number of OpenMP threads is then increased, decreasing the number of MPI clones accordingly.

With 1 subdomain, going from 48 clones to 2 clones times 24 threads leads to an overall speed-up of 43% (intermediate steps are not shown, but the performance is monotonically improving with the number of threads). Shared memory, besides optimizing the memory usage by avoiding data replica, improves the performance compared to

clones, which have to account also for the overhead of MPI reduce and broadcast operations. When running with OpenMP, the threads of a task are bound to contiguous cores. Going to full OpenMP parallelization with 48 threads means that they will expand over the two sockets of the Skylake CPU. Such configuration performs worse than keeping one MPI task per socket, because threads access memory resources on the remote socket.

Although not yet itself OpenMP parallel, the field solver is also faster when threads are adopted for the other parts of the code. In that case, fewer clones are instanced, resulting in a more efficient memory usage. Memory contention between several clones, concurrently solving for their field replica, has in fact proved to lead to a significant slow-down of the calculation.

Charge and current deposition is much faster with OpenMP threads than MPI clones, especially as a result of improved reduction. The ‘Get field’ routine is faster as well because field is interpolated from a single instance per socket in memory, not from many clones. The push loop slow-down with OpenMP threads is due to the embedded calculation of diagnostics within the particle loop which leads to false sharing, with thread replicas allocated in the same cache line. This will be fixed in the coming release of the code, leading to a further improvement of the overall performance with respect to pure MPI setups.

Skylake’s capability to run two threads on a single core (so-called hyperthreading) was not found to bring any significant speed-up to the overall timing of the code. It was not the case on other architectures either. The results are not shown here, but some parts of the code were accelerated by a small fraction of the order of 10%, while others were slowed down by a similar fraction. Getting a significant speed-up would require fine-tuning the hyperthreading on the different kernels of the code, which we decided not to do in order to stay consistent with our portability policy.

With 2 subdomains, the pure MPI parallelization with 24 clones is 20% faster than with 1 subdomain. In particular, the field solver timing is reduced almost by a factor 2, the mesh size per domain being halved. The deposition is a bit faster as well because data locality is improved by having sorted the particles in toroidal domains and the MPI reduction over clones is alleviated. Trading MPI clones for



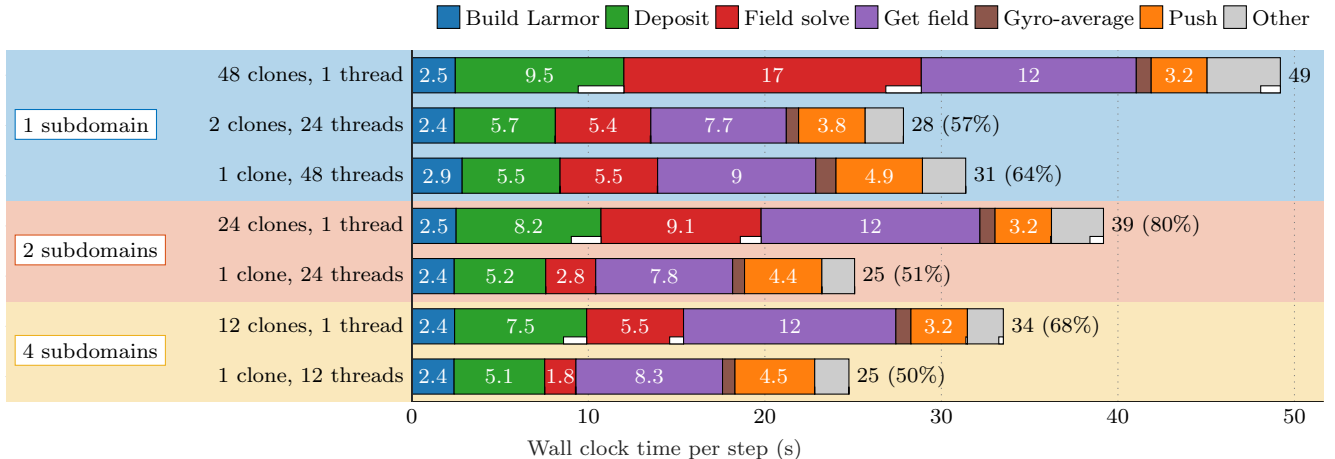


Figure 4: Comparison of timings for different MPI and OpenMP parallelization combinations on a Skylake node (2x24-cores). Timing ‘Other’ includes in particular sources and diagnostics. The ‘Field solve’ includes discrete Fourier transforms and backsolve operations. White bars represent time spent in MPI routines. The percentages of the total time are relative to the top bar.

OpenMP threads leads to an extra 36% speed-up for the same reasons as the 1-subdomain case.

The best performance is obtained with 4 subdomains, with a 32% speed-up on the pure MPI parallelization compared to the 48 clones case, and a 50% speed-up with OpenMP. The number of subdomains cannot be increased further for this test case with  $N_\varphi = 4$  because it has reached the limit of a single toroidal grid cell per subdomain. Alleviating this limit would require 2D or 3D domain decomposition [30] but this option is not implemented in ORB5 yet.

### 3.2. Running on different architectures

The different available computing architectures have been exploited, selecting for each of these the best possible combination of number of subdomains, clones and threads, fully exploiting the flexibility given by the various levels of parallelism implemented in ORB5, taking into account a number of constraints of the code (e.g. the number of grid points in the toroidal direction must be divisible by the number of subdomains) and specific node features in terms of number of cores and GPUs. These tests have been performed with the aim of comparing how different architectures perform in practice for our application code. On the other hand, exactly the same source code was used, i.e. there was no fine tuning, architecture-dependent optimization. For all GPU-equipped architectures, one GPU is used per MPI rank. We did not investigate other options of e.g. OpenMP threads linked to different GPUs, multiple MPI ranks oversubscribing the same GPU, or one MPI task directing multiple GPUs, which might bring further optimization. The resulting configurations are summarized in Figure 5, where the breakdown of the timings is presented. The first two bars show the results on the Skylake and the POWER9 CPUs, using MPI and OpenMP. The fraction of time spent in the various kernels is approximately the same for both CPUs. The difference in performance measured for all the code components between

the two architectures can be partially justified with the larger number of cores available on Skylake and with the usage of the PGI compiler on the POWER9, which may provide sub-optimal optimization compared to the native IBM XLF compiler. Further investigation is required in order to fully understand this difference.

The third and fourth bar of Figure 5 show the performance when GPUs are used. Due to the memory constraints, the test requires at least 4 GPUs to run. We chose to fully exploit the Summit node with its 6 GPUs and, consistently, 6 Piz Daint nodes. Two subdomains times three clones were set in both cases. OpenMP multithreading was switched off, since our current implementation of ORB5 does not support the simultaneous usage of OpenACC and OpenMP. On both systems, the field solver is the most time-consuming kernel. This is in fact the only part of the code which has no GPU implementation, running on the CPU only. For this field solver kernel, as seen in Section 3.1, the usage of the clones penalizes the POWER9+GPU setup compared to the POWER9 one, due to the presence in the former of collective MPI operations, necessary to coordinate the work on the different clones. All the other kernels run on the GPU. Comparing the P100 and V100 GPUs, we find an overall performance difference of a factor 1.5 (subtracting the time spent on the field solver using the CPU only), which results from the combination of the faster computing capabilities and the higher memory bandwidth of the V100 architecture.

The analysis of the runtime of ORB5 on a Summit node running with or without GPUs, show that the usage of the accelerators for this application reduces the computing time by a factor of  $\sim 5$  overall, and by a factor of  $\sim 8$  when only GPU-enabled kernels are considered. The kernel benefiting most from the GPU acceleration is ‘get field’ with a factor  $\sim 30$  speed-up. The push routine is accelerated by a factor  $\sim 12$ , the gyro-averaging by  $\sim 10$ , the Larmor construction by  $\sim 7$  and the deposition by  $\sim 5$ . Let us recall that these timings include the CPU-GPU data

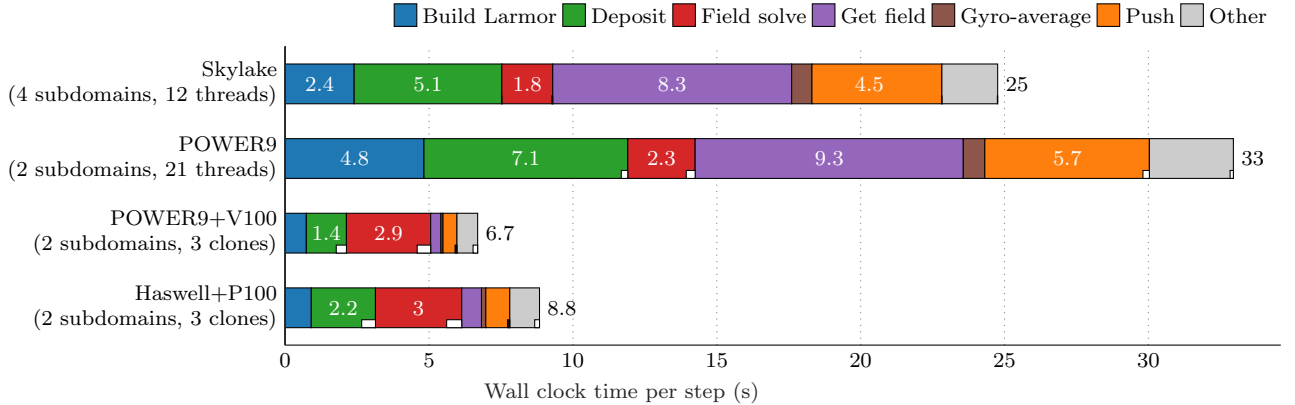


Figure 5: ORB5 performance on different architectures: Skylake 2×24-cores CPU, POWER9 2×21-cores CPU with and without 6 V100 GPUs, and 6 Haswell 12-cores CPUs with one P100 GPU each. Timing ‘Other’ includes in particular sources and diagnostics. The ‘Field solve’ includes discrete Fourier transforms and backsolve operations. White bars represent time spent in MPI routines.

transfers and that, for reference, the peak performance ratio of GPU acceleration on a Summit node, in the absence of any data transfer, is  $\sim 40$ .

### 3.3. Particle sorting

Data locality is expected to be highly beneficial for multi-core, multi-thread applications. For this reason, as mentioned in Section 2.1, sorting of markers with respect to the field cells has been introduced with the ORB5 refactoring. A bucket sort algorithm is adopted, this choice resulting from the analysis performed in [35], where several algorithmic variations and GPU implementations were investigated. One should however note that the results of the cited publication were performed on GPUs using Kepler architecture, which do not implement hardware atomic operations in double precision.

Figure 6 shows the effect of sorting on Marconi (Skylake) and Summit (POWER9+V100). On Marconi, the deposition routine is 59% faster with sorting than without and ‘get field’ is 31% faster. Unfortunately, the sorting itself becomes the most time-consuming part of the code (28% of the overall time) and counterbalances the speed-up it brought. In the end, the code is 7% slower with sorting for the present simulation parameters.

Even though the deposition and ‘get field’ timings are not the dominant ones on Summit with the GPUs, particle sorting has been tried and found to improve deposition by 66% (ignoring MPI reduction time) and ‘get field’ by 36%. The overall performance is however again decreased by 3% due to the cost of sorting itself.

Those results demonstrate the potential of the adoption of particle sorting to enforce memory locality, which however can be fully exploited only if specific optimization of the sorting kernel is accomplished. Indeed, only the most promising sorting algorithm found in the PIC-engine studies [35] has been implemented in ORB5 but it may not be the best for ORB5.

### 3.4. Parallel scalability

Figure 7 shows the parallel scalability of ORB5 up to 256 nodes on Marconi (7a), up to 2048 nodes of Piz Daint with or without using the GPUs (7b) and up to 4096 compute nodes on Summit machine (7c) with or without using the GPUs. It is worth noticing that the code could run up to more than 24000 GPUs in the largest set-up on Summit.

To obtain the results of Figure 7, four data sets are considered, which are scaled by a factor of 8 from one another for both particle and grid numbers. For each data set a strong scaling is performed, represented with a different color. On Marconi, Figure 7a, strong and weak scalings are close to ideal up to the maximum tested (we could not run beyond 256 nodes, i.e. 12288 MPI tasks, due to limited availability of resources). The scaling appears even super-ideal, but that is an artifact due to poor performance of the single-node case, which suffers from poor cache usage due to the large memory footprint. This effect is also seen, although much smaller, on the single-node Piz Daint results, but is absent from the single-node Summit results. Note that the single-node data set considered in section 3.1 has a much smaller grid (by a factor of 32), which was chosen to correspond to a toroidal slice of the full data set at 32 nodes, and does not suffer from this poor performance problem on any of the architectures considered. Increasing the number of nodes for CPU-only cases on Piz Daint, Figure 7b and Summit, Figure 7c, the excellent strong scaling properties are confirmed for all data sets up to 24576 CPU cores on Piz Daint, respectively 172032 CPU cores on Summit.

For all data sets and number of nodes considered on both Piz Daint and Summit platforms, the use of GPU is significantly accelerating the code. The smallest data set (20M markers per species,  $32 \times 256 \times 128$  cells, blue symbols and lines on Figure 7c) on a single Summit node runs 4.3 times faster using its 6 GPUs than with its CPUs only. This factor is slightly different than the one found in section 3.2 since the grid resolution is not the same. The strong scalability of this case without GPUs is nearly ideal



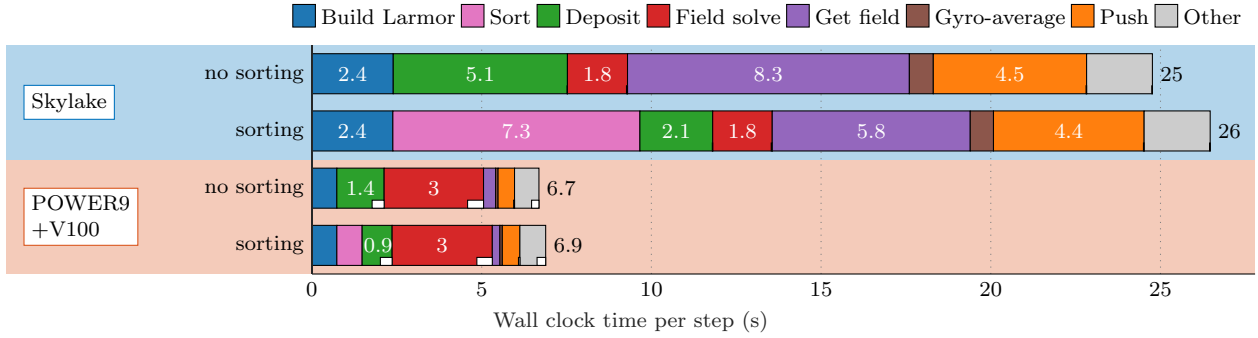
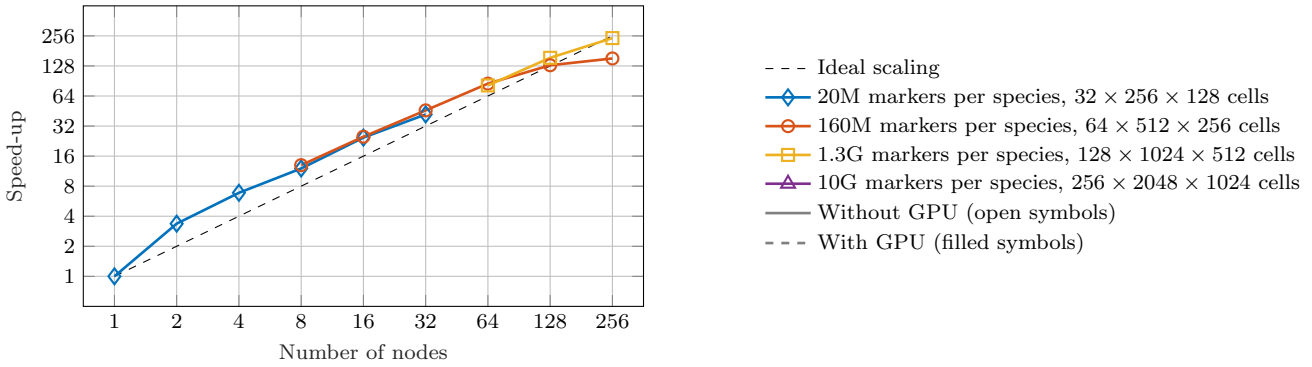
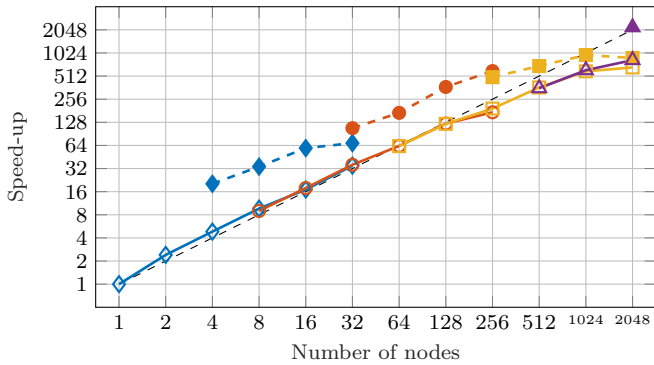


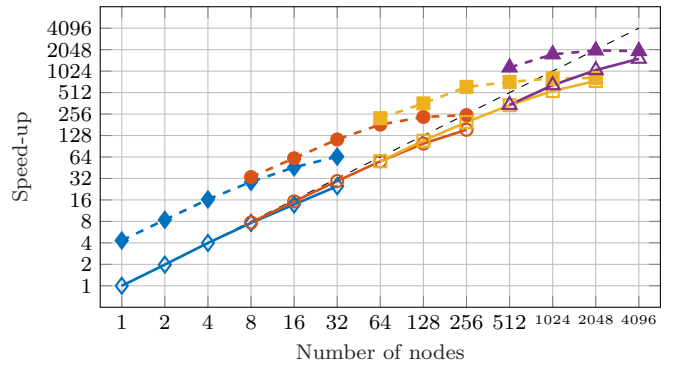
Figure 6: Comparison with and without sorting on Skylake (4 subdomains times 12 threads) and Power+V100 architectures (2 subdomains times 3 clones). Timing ‘Other’ includes in particular sources and diagnostics. The ‘Field solve’ includes discrete Fourier transforms and backsolve operations. White bars represent time spent in MPI routines.



(a) Marconi, 2 clones times 24 threads per node. Speed-up is normalized to the smallest case on a single node, which takes 20s per time step.



(b) Piz Daint, 12 threads per node for CPU-only runs, and a single task per node for CPU+GPU runs. Speed-up is normalized to the smallest case on a single node without GPU, which takes 115s per time step.



(c) Summit, 2 clones times 21 threads per node for CPU-only runs, and 6 clones per node for CPU+GPU runs. Speed-up is normalized to the smallest case on a single node without GPU, which takes 28s per time step.

Figure 7: Strong and weak scalings on Marconi (a), Piz Daint (b) and Summit (c), supercomputers. Each series represents a strong scaling. The set of first points from each strong scaling provides a weak scaling, multiplying the number of nodes and the problem size by a factor 8 for both particles and grids between consecutive colors. Continuous lines with open symbols indicate results using CPUs only, while dashed lines with filled symbols make use of the GPUs. We use as many subdomains as number of nodes up to the limit of the number of toroidal grid cells, beyond which clones are added.

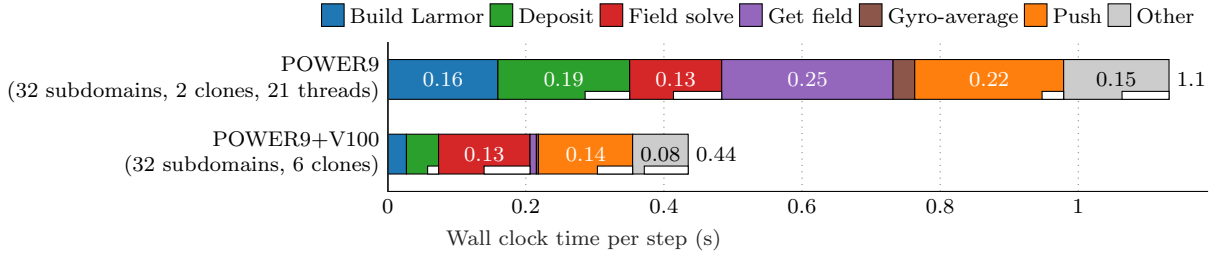


Figure 8: Breakdown of ORB5 kernel timings on the small data set (20M guiding centers per species,  $32 \times 256 \times 128$  cells), using 32 nodes on Summit, with or without GPUs. Timing ‘Other’ includes in particular sources and diagnostics. The ‘Field solve’ includes discrete Fourier transforms and backsolve operations. White bars represent time spent in MPI routines.

up to 8 nodes, and then slightly decreases to reach a parallel efficiency of 78% at 32 nodes with respect to 1 node. The parallel efficiency when using the GPUs is only 47%. This drop in efficiency can be explained with the higher MPI communication to computing ratio when GPUs are used, combined to the impact of the non-accelerated part of the code, essentially the field solve kernel, taking a higher fraction of the wall clock time, as shown in Figure 8. The presented breakdown of the timings refers to the small data set, but it is representative also of larger configurations, for which MPI communication gets increasingly heavier, resulting in a poorer parallel scalability compared to the CPU only runs. For instance, the second largest case (1.3G markers per species,  $128 \times 1024 \times 512$  cells, yellow symbols and lines on Figure 7c) running on 2048 nodes does not benefit much from the use of GPUs because it is dominated by communications. On the largest case (10G markers per species,  $256 \times 2048 \times 1024$  cells, violet symbols and lines on Figure 7c), the parallel efficiency of the CPU-only version at 4096 nodes with respect to 512 nodes is still 55%.

Between the first points of each strong scalability test, the problem size and the used computational resources have been scaled by a factor of 8 (twice more cells in each direction and 8 times more markers), hence they can be used to evaluate weak scalability. On Summit, weak scalability is linear for the first three tests, while it drops on the largest test, both for CPU-only and using the GPUs, due to the big communication overhead when large system configurations are used. The largest case is 512 times heavier than the smallest one and it is 33% slower on 512 nodes than the small one on 1 node with CPU only, or 48% slower using the GPUs.

The single node performance of Piz Daint on the smallest case (Figure 7b) without GPU, is about 4 times lower than on Summit, which is to be expected because it embeds a 12 cores Haswell CPU versus the  $2 \times 21$  cores of the POWER9 on Summit. Due to memory constraints, at least 4 nodes are required for this case to fit the GPU memory. Using the GPUs on Piz Daint brings around a factor 4 speed-up on the linear phases of the strong scalings, leading to a performance similar to Summit without GPUs at the same number of nodes. Considering the number of nodes, the overall weak and strong scalability on Piz

Daint appears to be even slightly better than that of Summit. However, one should consider that Piz Daint nodes are much lighter (less CPU cores and GPUs per node) than Summit ones, so the number of MPI tasks (hence the associated overhead) is smaller.

This is further shown in Figure 9, by comparing Summit and Piz Daint scalings with respect to the number of MPI tasks instead of the number of nodes, for the GPU cases only (note that there is one MPI task per GPU). The fact that 4 Daint P100 GPUs perform better than 6 Summit V100 GPUs in the smallest case is due to the usage of 4 subdomains instead of 6 clones. However, when the computing resources grow, Piz Daint strong scalability tends to degrade earlier than that of Summit. This is because Summit benefits from more intra-node MPI communication while Piz Daint uses inter-node MPI communication. This is particularly clear comparing the scalability of the two systems for the second largest test (yellow markers).

#### 4. Discussion and Conclusions

The tests presented in Section 3 allowed us to investigate the performance of the ORB5 code on different state-of-the-art HPC architectures, in typical production simulation setups, without any architecture-specific custom optimization (except those provided by the compiler), hence in a typical user-based scenario, using the same source code.

The parallel performance can be optimized at the application level, thanks to the different options supported by ORB5 for parallelism (domain decomposition and cloning and multithreading), whose tuning allows one to identify the most effective combination of computing resources. Often this is not driven only by the performance, but also by the memory usage and data locality. For instance, the six GPUs on a Summit fat node can be effectively exploited by dividing the domain into two subdomains (one per socket), replicated into three clones. Since the domain decomposition requires a number of MPI tasks power of two, without clones it would have been impossible to exploit all the six available GPUs. At the same time, cloning allows one to reduce the data size to fit the GPU memory.

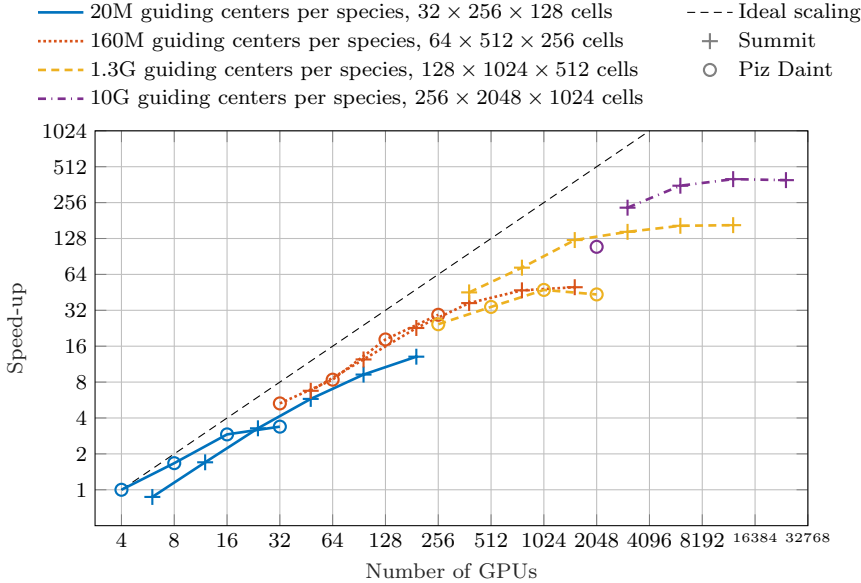


Figure 9: Comparison of strong scalings on Summit (cross symbols) and Piz Daint (circles symbols) supercomputers using GPUs. Data is the same as in Figure 7 with the total number of GPUs along the x-axis. Speed-up is normalized to the smallest case on 4 GPUs of Piz Daint, which takes 5.7s per time step.

As a general outcome, we have found that single-node parallelization optimization shows that subdomain decomposition should be prioritized, followed by OpenMP multithreading, and finally domain cloning. MPI parallelization is also useful to guarantee memory locality on multi-socket machines, avoiding to have OpenMP threads expanding over more than one socket. This priority hierarchy may be different for other simulation parameters and higher number of nodes, but it anyway provides insight for a first guess before fine-tuning optimization. The usage of OpenMP multithreading guarantees acceptable performance up to a few tens of threads both on the Skylake and on the POWER9 architectures, provided affinity is enforced and the socket boundaries are not crossed.

As expected, on a single node, the usage of the GPU leads to an important speed-up compared to non-accelerated nodes. However, such speed-up is far from being the nominal one. This is due to two main concurrent factors: the presence of the non-accelerated field solver and the overhead due to CPU-GPU data movement. The former is planned to be solved in the coming releases of the code. Since originally the solver represented a negligible part of the computation, it was left aside while focusing on the multithreaded implementation of the most time-consuming components. With all other code components accelerated, the field solver becomes an actual bottleneck and now needs to be ported on the GPU. Once this will be accomplished, also the data-movement problem will be strongly alleviated, since charge densities and forces will be permanently resident on the GPU, with no more need of transferring them at each time iteration. Further data-movement overheads due, for example, to the calculation of diagnostics or to MPI communication are expected to be effectively hidden through asynchronous processes, or

through the adoption of GPUDirect.

Numerical kernels could be further optimized on the GPU by rewriting them in Cuda. To ensure the code remains future-proof, only the parts exempt from physics should be considered, such as particle deposition and field evaluation, particle sorting and field solver (all the physics is in the matrix assembly). We know from the PIC-engine studies [31] that such work could bring a few tens percents speed-up, but it is not our priority as it would affect kernels that are either taking a small fraction of the wall clock time, or planned to be upgraded in Fortran.

The challenge of accelerating PIC’s critical operations of particle-to-field and field-to-particle has also been tackled introducing new algorithms using particle sorting. Even though the overall gain is not yet positive, the ‘deposition’ and ‘get field’ operations are significantly faster on CPU and GPU architectures. Previous studies done with the PIC-engine [31, 35] had shown better performance with the sorting because they did not benefit from the hardware atomic implementation brought by the newest generations of GPUs in their reference simulations without sorting. In the future, if the cost of sorting could be reduced, a net performance gain would be achieved on ORB5. One could for instance track markers leaving cells in order to avoid full sort at each iteration.

Scalability has been investigated focusing on the two GPU-equipped systems, for which it represents a major challenge. When CPU-only tests are run, both Summit and Piz Daint show remarkable linear strong scalability up to very large configurations, losing efficiency only at the largest set-up, as shown also looking at weak scaling. When GPUs are used, the ratio between computing and communication time drops, and good strong scalability is much harder to achieve. Our tests show that ORB5

can take advantage of the Summit fat nodes architecture, strong scaling efficiently to large numbers of GPUs, with improved scalability compared to Piz Daint, that has a larger MPI communication overhead, each node having only one GPU. A second advantage of Summit is that bigger problem sizes can be handled by a single node, with an overall improvement of weak scaling. On the other hand, the Summit node resulted to be less flexible than the “fine grained” Piz Daint node. In order to fully exploit the resources of the Summit node we had to make use of clones, which penalized the performance compared to Piz Daint, where pure domain decomposition could be used. This resulted in better performance of Piz Daint when few GPUs (nodes) are used.

Overall, ORB5 has proved to be a prime example of a code designed for the coming generations of heterogeneous HPC architectures. Based on MPI, OpenMP and OpenACC the same codebase can be used efficiently on distributed systems as well as on multi/many threads devices. It can exploit state-of-the-art supercomputing systems up to the largest scales, and could be tested up to more than 24000 GPUs. It is easily portable, supporting the Intel, PGI, Cray and GNU (not shown in the paper) compilers, strongly facilitating the work of the scientific user, and making accessible a broad spectrum of HPC solutions to effectively perform large scale simulations tackling the most challenging problems in plasma physics. The directive-based tools used in this work can be considered as mature enough and widely available for other legacy codes needing similar refactoring. Among all features needed for this project, only one was problematic, namely OpenACC’s GPUDirect. This way of exchanging data directly from a GPU to another was not yet fully supported by the network constructor, but it is expected to be within the next couple of years. Moreover, the only missing feature we had wished to have when porting the code to GPUs, namely OpenACC support for reduction over arrays, is now provided by standard 2.7.

## Acknowledgments

Simulations presented in this paper were performed at the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725, Project ID is PHY137, at CSCS (Swiss National Supercomputing Center), in the User Lab program, project ID s909, and at CINECA (Italian National HPC facility), in the EUROfusion program, project ID FUA23. We acknowledge all those facilities.

We are also grateful to Stephane Ethier from PPPL (Princeton Plasma Physics Laboratory) for very fruitful discussions about porting production codes to GPUs.

This work has been carried out within the framework of the EUROfusion Consortium and has received funding from the Euratom research and training program 2014-2018 and 2019-2020 under grant agreement No 633053.

The views and opinions expressed herein do not necessarily reflect those of the European Commission. This work has been partly supported by the Swiss National Science Foundation.

## Bibliography

- [1] X. Garbet, Y. Idomura, L. Villard, and T.H. Watanabe. Gyrokinetic simulations of turbulent transport. *Nucl. Fusion*, 50(4):043002, apr 2010.
- [2] T. Dannert, A. Marek, and M. Rampp. Porting Large HPC Applications to GPU Clusters: The Codes GENE and VERTEX. *Adv. Parallel Comput.*, 25:305–314, 2014.
- [3] I Sfiligoi, J Candy, and M Kostuk. CGYRO Performance on Power9 CPUs and Volta GPUs. In *Lect. Notes Comput. Sci. (including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics)*, volume 11203 LNCS, pages 365–372. Springer, Cham, jun 2018.
- [4] X. Meng, X. Zhu, P. Wang, Y. Zhao, X. Liu, B. Zhang, Y. Xiao, W. Zhang, and Z. Lin. Heterogeneous programming and optimization of gyrokinetic toroidal code and large-scale performance test on th-1a. In Julian Martin Kunkel, Thomas Ludwig, and Hans Werner Meuer, editors, *Supercomputing*, pages 81–96, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.
- [5] K. Madduri, K. Ibrahim, S. Williams, E.-J. Im, S. Ethier, J. Shalf, and L. Oliker. Gyrokinetic toroidal simulations on leading multi- and manycore HPC systems. In *Proc. 2011 Int. Conf. High Perform. Comput. Networking, Storage Anal. - SC ’11*, page 1, New York, New York, USA, 2011. ACM Press.
- [6] E. D’Azevedo, S. Abbott, T. Koskela, P. Worley, S.-H. Ku, S. Ethier, E. Yoon, M. S. Shephard, R. Hager, J. Lang, et al. The fusion code xgc: Enabling kinetic study of multiscale edge turbulent transport in iter. In *Exascale Scientific Applications*, pages 529–552. Chapman and Hall/CRC, 2017.
- [7] S. Abbott and E. D’Azevedo. Physics based optimization of particle-in-cell simulations on gpus. In *APS Meeting Abstracts*, 2016.
- [8] W. Tang and Z. Lin. Global gyrokinetic particle-in-cell simulation. In *Exascale Scientific Applications*, pages 507–528. Chapman and Hall/CRC, 2017.
- [9] T.-M. Tran, K. Appert, M. Fivaz, G. Jost, J. Vaclavik, and L. Villard. Global gyrokinetic simulation of ion-temperature-gradient-driven instabilities using particles. In J. W. Connor, E. Sindoni, and J. Vaclavik, editors, *Th. of Fusion Plasmas*, volume 18, page 45. Società Italiana di Fisica, Bologna, 1999.
- [10] S. Jolliet, A. Bottino, P. Angelino, R. Hatzky, T.-M. Tran, B.F. Mcmillan, O. Sauter, K. Appert, Y. Idomura, and L. Villard. A global collisionless PIC code in magnetic coordinates. *Comput. Phys. Commun.*, 177(5):409–425, sep 2007.
- [11] A. Bottino, T. Vernay, B. Scott, S. Brunner, R. Hatzky, S. Jolliet, B.F. McMillan, T.-M. Tran, and L. Villard. Global simulations of tokamak microturbulence: Finite- $\beta$  effects and collisions. *Plasma Phys. Control. Fusion*, 53(12):124027, dec 2011.
- [12] E. Lanti, N. Ohana, N. Tronko, T. Hayward-Schneider, A. Bottino, B.F. McMillan, A. Mishchenko, A. Scheinberg, A. Biancalani, P. Angelino, S. Brunner, J. Dominski, P. Donnel, C. Gheller, R. Hatzky, A. Jocksch, S. Jolliet, Z.X. Lu, J.P. Martin Collar, I. Novikau, E. Sonnendrücker, T. Vernay, and L. Villard. Orb5: A global electromagnetic gyrokinetic code using the PIC approach in toroidal geometry. *Computer Physics Communications*, page 107072, nov 2019.
- [13] C.C. Kim and S.E. Parker. Massively Parallel Three-Dimensional Toroidal Gyrokinetic Flux-Tube Turbulence Simulation. *J. Comput. Phys.*, 161(2):589–604, jul 2000.
- [14] R. Hatzky. Domain cloning for a particle-in-cell (PIC) code on a cluster of symmetric-multiprocessor (SMP) computers. *Parallel Comput.*, 32(4):325–330, apr 2006.
- [15] Git web homepage. <https://git-scm.com>.
- [16] Jenkins web homepage. <https://jenkins.io>.

- [17] A. Scheinberg, S. Ethier, C.-S. Chang, G. Chen, R. Bird, S. Slatery, and P. Worley. Kokkos and Fortran in the Exascale Computing Project plasma physics code XGC. In *Proc. 2019 Int. Conf. High Perform. Comput. Networking, Storage Anal. - SC '19*, page 3. ACM Press, 2019.
- [18] D. A. Beckingsale, T. R. W. Scogland, J. Burmark, R. Hornung, H. Jones, W. Killian, A. J. Kunen, O. Pearce, P. Robinson, and B. S. Ryu. RAAJ: Portable Performance for Large-Scale Scientific Applications. In *2019 IEEE/ACM International Workshop on Performance, Portability and Productivity in HPC (P3HPC)*, pages 71–81. IEEE, nov 2019.
- [19] R. Keryell. SYCL: A Single-Source C++ Standard for Heterogeneous Computing. In *2019 IEEE/ACM International Workshop on Performance, Portability and Productivity in HPC (P3HPC)*, 2019.
- [20] H. Sugama. Gyrokinetic field theory. *Phys. Plasmas*, 7(2):466–480, feb 2000.
- [21] N. Tronko, A. Bottino, C. Chandre, and E. Sonnendrücker. Hierarchy of second order gyrokinetic Hamiltonian models for particle-in-cell codes. *Plasma Phys. Control. Fusion*, 59(6):064008, jun 2017.
- [22] H. Lütjens, A. Bondeson, and O. Sauter. The CHEASE code for toroidal MHD equilibria. *Comput. Phys. Commun.*, 97(3):219–260, sep 1996.
- [23] J. Dominski, B.F. McMillan, S. Brunner, G. Merlo, T.-M. Tran, and L. Villard. An arbitrary wavelength solver for global gyrokinetic simulations. Application to the study of fine radial structures on microturbulence due to non-adiabatic passing electron dynamics. *Phys. Plasmas*, 24(2):022308, feb 2017.
- [24] T. Vernay, S. Brunner, L. Villard, B.F. McMillan, S. Jolliet, T.-M. Tran, A. Bottino, and J.P. Graves. Neoclassical equilibria as starting point for global gyrokinetic microturbulence simulations. *Phys. Plasmas*, 17(12):122301, dec 2010.
- [25] B.F. McMillan, S. Jolliet, T.-M. Tran, L. Villard, A. Bottino, and P. Angelino. Long global gyrokinetic simulations: Source terms and particle noise control. *Phys. Plasmas*, 15(5):052308, may 2008.
- [26] B.F. McMillan, P. Hill, A. Bottino, S. Jolliet, T. Vernay, and L. Villard. Interaction of large scale flow structures with gyrokinetic turbulence. *Phys. Plasmas*, 18(11):112503, nov 2011.
- [27] S. Jolliet. *Gyrokinetic Particle-In-Cell Global Simulations of Ion-Temperature-Gradient and Collisionless-Trapped-Electron-Mode Turbulence in Tokamaks*. PhD thesis, EPFL, Lausanne, 2009.
- [28] A. Mishchenko, A. Bottino, R. Hatzky, E. Sonnendrücker, R. Kleiber, and A. Könies. Mitigation of the cancellation problem in the gyrokinetic particle-in-cell simulations of global electromagnetic modes. *Phys. Plasmas*, 24(8):081206, aug 2017.
- [29] A. Mishchenko, A. Bottino, A. Biancalani, R. Hatzky, T. Hayward-Schneider, N. Ohana, E. Lanti, S. Brunner, L. Villard, M. Borchardt, R. Kleiber, and A. Könies. Pullback scheme implementation in ORB5. *Comput. Phys. Commun.*, 238:194–202, may 2019.
- [30] A. Jocksch, N. Ohana, E. Lanti, A. Scheinberg, S. Brunner, C. Gheller, and L. Villard. Prediction of the inter-node communication costs of a new gyrokinetic code with toroidal domain. In R. Wyrzykowski, J. Dongarra, E. Deelman, and K. Karczewski, editors, *Parallel Processing and Applied Mathematics*, pages 370–380, Cham, 2018. Springer International Publishing.
- [31] F. Hariri, T.-M. Tran, A. Jocksch, E. Lanti, J. Progsch, P. Messmer, S. Brunner, C. Gheller, and L. Villard. A portable platform for accelerated pic codes and its application to gpus using openacc. *Computer Physics Communications*, 207:69 – 82, 2016.
- [32] N. Ohana, A. Jocksch, E. Lanti, T.-M. Tran, S. Brunner, C. Gheller, F. Hariri, and L. Villard. Towards the optimization of a gyrokinetic Particle-In-Cell (PIC) code on large-scale hybrid architectures. In *J. Phys. Conf. Ser.*, volume 775, page 012010. IOP Publishing, nov 2016.
- [33] OpenMP web homepage. <https://www.openmp.org/>.
- [34] OpenACC web homepage. <https://www.openacc.org/>.
- [35] A. Jocksch, F. Hariri, T.-M. Tran, S. Brunner, C. Gheller, and L. Villard. A bucket sort algorithm for the particle-in-cell method on manycore architectures. In R. Wyrzykowski, E. Deelman, J. Dongarra, K. Karczewski, J. Kitowski, and K. Wiatr, editors, *Parallel Processing and Applied Mathematics*, pages 43–52, Cham, 2016. Springer International Publishing.