InKS, a Programming Model to Decouple Algorithm from Optimization in HPC Codes

Ksander Ejjaouani, Olivier Aumage, Julien Bigot, Michel Mehrenberger, Hitoshi Murai, Masahiro Nakao, Mitsuhisa Sato

► To cite this version:

Ksander Ejjaouani, Olivier Aumage, Julien Bigot, Michel Mehrenberger, Hitoshi Murai, et al.. InKS, a Programming Model to Decouple Algorithm from Optimization in HPC Codes. *Journal of Supercomputing*, 2019, 10.1007/s11227-019-02950-2 . hal-02281963

HAL Id: hal-02281963

<https://hal.inria.fr/hal-02281963>

Submitted on 9 Sep 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

InKS, a Programming Model to Decouple Algorithm from Optimization in HPC Codes

Ksander Ejjaouani · Olivier Aumage ·
Julien Bigot · Michel Méhrenberger ·
Hitoshi Murai · Masahiro Nakao ·
Mitsuhisa Sato

Abstract Existing programming models tend to tightly interleave algorithm and optimization in HPC simulation codes. This requires scientists to become experts in both the simulated domain and the optimization process and makes the code difficult to maintain or port to new architectures. In this paper, we propose the INKS programming model that decouples these concerns with two distinct languages: INKS_{pia} to express the simulation algorithm and INKS_{pso} for optimizations. We define INKS_{pia} and evaluate the feasibility of defining INKS_{pso} with three test-languages: $\text{INKS}_{\text{o/C++}}$, $\text{INKS}_{\text{o/loop}}$, $\text{INKS}_{\text{o/XMP}}$. We evaluate the approach on synthetic benchmarks (NAS and heat equation) as well as on a more complex example (6D Vlasov-Poisson solver). Our evaluation demonstrates the soundness of the approach as it improves the separation of algorithmic and optimization concerns at no performance cost. We also identify a set of guidelines for the later full definition of the INKS_{pso} language.

Keywords programming model · separation of concerns · HPC · DSL

Ksander Ejjaouani
Maison de la Simulation, CEA, CNRS, Univ. Paris-Sud, UVSQ, Université Paris-Saclay
Inria
Gif-sur-Yvette, France
E-mail: ksander.ejjaouani@inria.fr

Olivier Aumage
Inria, LaBri, Bordeaux, France
E-mail: olivier.aumage@inria.fr

Julien Bigot
Maison de la Simulation, CEA, CNRS, Univ. Paris-Sud, UVSQ, Université Paris-Saclay
Gif-sur-Yvette, France
E-mail: julien.bigot@cea.fr

Michel Mehrenberger
Université de Marseille, Marseille, France
E-mail: mehrenbe@math.unistra.fr

Hitoshi Murai · Masahiro Nakao · Mitsuhisa Sato
Riken CCS, Kobe, Japan
E-mail: {h-murai,masahiro.nakao,msato}@riken.jp

1 Introduction

It is more and more common to identify simulation as the *third pillar of science* together with theory and experimentation. Parallel computers provide the computing power required by the more demanding of these simulations. The complexity and heterogeneity of these architectures do however force scientists to write complex code (using vectorization, parallelization, accelerator specific languages, etc.). These optimizations heavily depend on the target machine, the code has to be adapted whenever it is ported to a new architecture.

As a result, scientists have to become experts in the art of computer optimizations in addition to their own domain of expertise. It is very difficult in practice to maintain a code targeting multiple distinct architectures. One fundamental cause for this situation is the tight interleaving of two distinct concerns imposed by most programming models. On the one hand, the algorithm comes from the expertise of the domain scientists and does not depend on the target architecture. On the other hand, optimizations form another domain of expertise and has to be adapted for a given architecture. Thereby, both algorithm and optimizations concerns are expressed within a single code. This mix impedes simulation codes maintainability and readability while hindering developers productivity.

Many approaches have been proposed to improve this situation in the form of libraries or languages [17, 4, 20, 19]. Approaches based on automated optimization processes typically isolate the algorithmic aspects well, but restrict their domain of applicability and/or the range of supported optimizations. Approaches based on optimization tools and libraries enable optimization specialists to express common optimizations efficiently but leave others mixed with the algorithm.

In this paper, we propose the Independent Kernel Scheduling (INKS) programming model to separate algorithm from optimization choices in HPC simulation codes. We define the INKS_{pia} language used to express the algorithm of an application independently of its optimization. This separation aims to improve the readability and maintainability of codes while easing portability and new optimization expression. This approach is used for common optimizations while $\text{INKS}_{\text{o/C++}}$ for less common optimizations. Such a program can then be optimized using $\text{INKS}_{\text{o/XMP}}$ and $\text{INKS}_{\text{o/loop}}$, two domain-specific languages (DSLs) which ask for optimization information only. While these DSLs target some common optimizations, $\text{INKS}_{\text{o/C++}}$ can be used for less common ones.

This paper makes the following contributions: **1)** it defines the INKS programming model and its *platform-independent algorithmic* language INKS_{pia} ; **2)** it proposes an implementation of INKS and tests the INKS_{psa} approach with three optimization DSLs, $\text{INKS}_{\text{o/C++}}$, $\text{INKS}_{\text{o/loop}}$ and $\text{INKS}_{\text{o/XMP}}$; and **3)** it evaluates the approach on the synthetic NAS parallel benchmarks [3] and on the 6D Vlasov-Poisson solving with a semi-Lagrangian method.

The remaining of the paper is organized as follows. Section 2 analyzes related works. Section 3 describes INKS and its implementation. Section 4 shows

the use of INKS on a 6D Vlasov-Poisson solver. Section 5 evaluates our approach. Section 6 concludes the paper.

2 Related Work

We now present approaches currently available to scientific programmers to help implementing simulation applications. A first widely used approach is based on imperative languages such as Fortran or C. Libraries like MPI extend this to distributed memory with message passing. Abstractions very close to the execution machine make fine-tuning possible to achieve good performance on any specific architecture. It does however require encoding complex optimizations directly in the code. As there is no language support to separate the algorithm and architecture-specific optimizations, tedious efforts have to be applied [13] to support performance portability. Algorithm and optimizations are instead often tightly bound together in codes.

A second approach is offered by tools (libraries, frameworks or language extensions) that encode classical optimizations. *OpenMP* [5], *REPARA* [8] or *Kokkos* [4] support common shared-memory parallelization patterns. For example, Kokkos offers multidimensional arrays and iterators for which efficient memory mappings and iteration orders are selected independently. *UPC* [9] or *XMP* [17] support the partitioned global address space paradigm. For example, in XMP, directives describe array distribution and communications between nodes. These tools offer gains of productivity when the optimization patterns they offer fit the requirements. The separation of optimizations from the main code base also eases porting between architectures. Even if expressed more compactly optimizations do however remain mixed with the algorithm. For instance, in OpenMP or REPARA, parallel concerns are specified on top of an existing code which already carries optimization choices, such as loops order.

A third approach pushes this further with tools that automate the optimization process. For example, *PaRSEC* [12] or *StarPU* [1] support the many-tasks paradigm. In StarPU, the user expresses its code as a directed acyclic graph (DAG) of tasks with data dependencies that is automatically scheduled at runtime depending on the available resources. Other examples are *SkeTo* [21] or *Lift* [19] that offer algorithmic skeletons. Lift offers a limited set of parallel patterns whose combinations are automatically transformed by an optimizing compiler. Automating optimization improves productivity and clearly separate these optimizations which improves portability. The tools do however not cover the whole range of potential optimizations such as the choice of work granularity inside tasks in StarPU for example. The algorithm remains largely interleaved with optimization choices even with this approach.

A last approach is based on DSLs that restrict the developer to the expression of the algorithm only while optimizations are handled independently, such as *Pochoir* [20] or *PATUS* [6], DSLs for stencil problems. In Pochoir, the user specifies a stencil (computation kernel and access pattern), boundary

conditions and a space-time domain while all optimizations are handled by a compiler. These approaches ensure a very good separation of concerns. The narrower the target domain is, the more efficient domain and architecture-specific optimizations are possible. However, it makes it less likely for the tool to cover the needs of a whole application. On contrary, the wider the target domain is, the less efficient optimizations are possible.

To summarize, one can consider a continuum of approaches from very general approaches where the optimization process is manual to more and more domain specific where the optimization process can be automated. The more general approaches support a large range of optimizations and application domains but yield high implementation costs and low separation of concerns and portability. The more automated approaches reduce implementation costs and offer good separation of concerns and portability but restrain the range of supported domains and optimizations. Ideally, one would like to combine all these advantages: **1)** the domain generality of imperative languages, **2)** the ease of optimization offered by dedicated tools and **3)** the separation of concerns and performance portability offered by DSLs with **4)** the possibilities of fine and manual optimizations offered by both imperative languages and dedicated tools. The following section describes the INKS programming model that aims to combine these approaches to offer such a solution.

3 The INKS programming model

This section first introduces the design of our INKS programming model, based on the use of distinct languages to express the algorithm and optimization choices separately. It then presents a prototype implementation of the model composed of INKS_{pia} , the algorithm language, and two INKS_{pso} optimization languages. The simulation *algorithm* consists in the set of values computed, the formula used to produce them as well as the simulation inputs and outputs. *Optimization choices* include all that is not the algorithm: e.g. the computing unit selected for each computation, their ordering, the memory location for each value, etc. Multiple optimization choices can differ in performance but simulation results depend on the algorithm only. The INKS approach is summarized in Figure 1. The INKS_{pia} language is used to express the algorithm with no concern for optimization choices. A compiler can automatically generate non-optimized choices from an INKS_{pia} specification, mostly for testing purposes. The INKS_{pso} language is used to define optimizations only, while other information is gathered from the INKS_{pia} code.

Fig. 1: The INKS model

Many versions of the optimization choices can be devised for a single algorithm; for example to optimize for multiple targets.

We now describe the INKS_{pia} language, and proposes three preliminary DSLs to validate strategies for the design of the INKS_{pso} language: $\text{INKS}_{\text{o/XMP}}$ handles domain decomposition, $\text{INKS}_{\text{o/loop}}$ focuses on efficient loop nest and $\text{INKS}_{\text{o/C++}}$ enables to use C++ and write arbitrary complex optimization.

3.1 The INKS_{pia} language.

In INKS_{pia} [2] values are stored in infinite multidimensional arrays based on dynamic single assignment (DSA, each coordinate can only be written once). Memory placement of each coordinate is left unspecified. Computations are specified by *kernel* procedures that **1)** take as parameters data arrays and integer coordinates; **2)** specify the coordinate they might read and will write in each array; and **3)** define either a C++ or INKS implementation. An INKS implementation defines kernels *validity domains*: coordinates where C++ kernels can generate values in arrays. Kernel execution order is left unspecified. The simulation entry point is a kernel marked *public*. Listing 1 presents a simple INKS_{pia} code: a 1D 3-point stencil computation. The simulation consists of one 2D logical array, **Array** (line 12), two kernels, **stencil3** (line 1) and **boundary** (line 6), and is parameterized by two integers, **X** and **T** (line 11). Line 12 specifies that the simulation starts with a subset of **Array** (every values in the space dimension, at the first time-step) and expects, as output, another subset: every values in the space dimension at the last time-step. To do so, it can call the **stencil3** and **boundary** kernels with a specific set of values for their integer parameters **x** and **t** (validity domain) and with **Array** as their logical array parameter, as expressed on line 15 and 16.

A INKS_{pia} code specifies a parameterized task graph (PTG) [7]. This graph is encoded using the Polyhedron Model [10]. It provides a compact representation of static control programs. This representation covers a large range of problems but imposes a few limitations. Mostly, all the problem parameters must be known at launch time, it does for example not support adaptive mesh or time-steps. It is still possible to express these concerns outside INKS_{pia} and call the INKS implementation multiple times with different parameters.

Once the algorithm is specified in INKS_{pia} , the goal is to write the optimization choices meaning choosing a memory layout and a scheduling of the kernels. Two approaches exist to produce these choices: the automatic compiler or INKS_{pso} . All approaches rely on the *Kernel* file. This file is generated by the INKS_{pia} compiler which translates the INKS_{pia} kernels to C++ functions. The INKS_{pia} compiler can produce generic optimization choices with a valid but non-optimized computations scheduling and memory allocations to execute them. Scheduling and memory layouts are computed using the *Integer Set Library* [22] and recent works on modular mapping in the polyhedron model [14]. Arbitrarily complex versions of optimization choices can also be written manually in plain C++. These functions can be called from any ex-

```

1 kernel stencil3(x, t) : (
2   double A(2) {in: (x-1:x+1, t-1) | out: (x, t)}
3 )
4 #CODE(C) A(x, t+1) = 0.5*A(x, t-1) + 0.25*(A(x-1, t-1)+A(x+1, t-1)); #END
5
6 kernel boundary(x, t) : (
7   double A(2) {in: (x, t-1) | out: (x, t)}
8 )
9 #CODE(C) A(x,t) = A(x, t-1); #END
10
11 public kernel inks_stencil(X, T) : (
12   double Array(2) {in: (0:X, 0) | out: (0:X, T-1)}
13 )
14 #CODE(INKS)
15   stencil3 (1:X-1, 1:T) : (Array),
16   boundary {(0, 1:T), (X-1, 1:T)} : (Array)
17 #END

```

Listing 1: 1D stencil computation on a 2D domain in INKS_{pia}

isting code whose language supports the C calling convention. However, that approach requires information present in INKS_{pia} to be repeated. The INKS_{pso} DSL thus interface the optimization process with INKS_{pia} , offering the optimization specialists to specify optimization only.

3.2 INKS_{pso} DSLs implementation

Implementing a complete optimization DSL for the INKS model is a long-term objective. In this study, we present the two DSLs $\text{INKS}_{\text{o/XMP}}$ and $\text{INKS}_{\text{o/loop}}$. Both DSLs solely describe optimizations, and get missing information from INKS_{pia} code.

$\text{INKS}_{\text{o/XMP}}$ (illustrated in Listing 2) handles distributed memory domain decomposition by combining C and directives based on XMP and adapted for INKS. The compiler replaces these directives by C and XMP code. The **inks decompose** directive supports static or dynamic allocation of logical arrays described in the algorithm. The domain size is extracted from INKS_{pia} source and the user only has to specify its mapping onto memory. As in XMP, $\text{INKS}_{\text{o/XMP}}$ supports domain decomposition mapped onto an XMP topology. In INKS_{pia} code, there are no concerns for memory optimization such as dimension ordering or memory reuse. Therefore, $\text{INKS}_{\text{o/XMP}}$ supports dimension reordering and folding which consists in reusing the same memory address for subsequent indices in a given dimension. The **exchange** directive supports halo exchanges. The user specifies which dimension should be exchanged and which computational kernel will be executed after the exchange. From these informations and the INKS_{pia} kernel specification, the $\text{INKS}_{\text{o/XMP}}$ compiler then computes the halo size. While XMP requires halo values to be stored contiguously with the domain, $\text{INKS}_{\text{o/XMP}}$ supports a dynamic halo extension where halo values are stored in dedicated, dynamically allocated buffers to reduce memory footprint.

$\text{INKS}_{\text{o/loop}}$ (illustrated in Listing 3) offers to specify manually loop nests for which the compiler generates plain C++ loops. Plain C++ is usable in combination with $\text{INKS}_{\text{o/loop}}$. The **loop** keyword introduces a nest optimiza-

```

1  /**InKSo/XMP specification***/
2  #pragma xmp nodes p[nNodes]
3  void inks_stencil(T& Input, size_t X, size_t T){
4      #pragma inks decompose % Array // allocation of "Array" algorithmic array
5      (2*2, 1) // dimension reordering, time dimension is fold by 2 and not distributed
6      with t onto p // block decomposition mapped on the XMP topology
7
8      for(int t=1; t<T; t++){
9          boundary(Array, 0, t);
10         for(int x=1; x<X-1; x++){
11             stencil3(Array, x, t);
12         }
13     }
14 }
15
16 /**XMP + C result***/
17 #pragma xmp nodes p[nNodes]
18 #pragma xmp template t[:][:] // xmp 1d logical array
19 #pragma xmp distribute t[*][block] onto p
20 void inks_stencil(T& Input, size_t X, size_t T){
21     #pragma xmp align Array[t][x] with t[t][x]
22     #pragma xmp template_fix t[2][X]
23     Array = (double(*)[X]) xmp_malloc(xmp_desc_of(f6d), 2, X);
24     for(int t=1; t<T; t++){
25         boundary(Array, 0, t);
26         for(int x=1; x<X-1; x++){
27             stencil3(Array, x, t);
28         }
29     }
30 }

```

Listing 2: InKS_{o/XMP} optimization choices of the InKS_{pia} code presented in Figure 1 and the C + XMP result

```

1  /** stencil.iloop InKSo/Loop optimization choices file ***/
2  loop stencil3_loop(t, X) : stencil3 { // set "t" value
3      // "Set" not specified -> loop bounds are computed, with a fixed "t"
4      Order: x; // order of the loop
5  }
6  /** C++ result ***/
7  void stencil3_loop(T& Array, size_t t, size_t X){
8      for (int x=0; x<X; x++){
9          stencil3(Array, x, t);
10     }
11 /** InKSo/Loop result used ***/
12 void inks_stencil(T& Input, size_t X, size_t T){
13     for(int t=1; t<T; t++){
14         boundary(Array, 0, t);
15         stencil3_loop(Array, t, X);
16     }
17 }
18 }

```

Listing 3: InKS_{o/loop} optimization choices of a loop from the InKS_{pia} code presented in Figure 1 with its C++ result and use

tion with a name, the list of parameters from the algorithm on which the loop bounds depend and a reference to the optimized kernel. Loop bounds can be automatically extracted from InKS_{pia}, but the **Set** keyword makes it possible to restrict these bounds. The **Order** keyword specifies the iteration order on the dimensions named according to the InKS_{pia} code. The **Block** keyword enables the user to implement blocking. It takes as parameters the size of block for the loops starting from the innermost one. If there are less block

Fig. 2: The 6D Vlasov-Poisson algorithm

sizes than loops, the remaining loops are not blocked. The **Buffer** keyword supports copying data in a local buffer before computation and back after to ensure data continuity and improve vectorization. The compiler uses data dependencies from the `INKSpia` code to check the validity of the loops order and generate vectorization directives whenever possible.

4 The 6D Vlasov/Poisson problem

The 6D Vlasov-Poisson equation, presented in (1), describes the movement of particles in a plasma and the resulting electric field. We study its resolution for a single species on a 6D Cartesian mesh with periodic boundary conditions. We solve the Poisson part using a fast Fourier transform (FFT) and rely on a Strang splitting (order 2 in time) for the Vlasov part. This leads to 6 1D advectons: 3 in space dimensions (x_1, x_2, x_3) and 3 in velocity dimensions (v_1, v_2, v_3) . Each 1D advection relies on a Lagrange interpolation of degree 4. In the space dimensions, we use a semi-Lagrangian approach where the stencil is not applied around the destination point but at the foot of characteristics, only known at runtime, as described in more details in [18].

The main unknown is f (`f6D` in the code), the distribution function of particles in 6D phase space. Due to the Strang splitting, a first half time-step of advectons is required after `f6D` initialization but before the main time-loop. These advectons need the electric field E as input. E is obtained through the FFT-based Poisson solver that in turn needs the charge density ρ as input. ρ is computed by a reduction of `f6D`. The main time-loop is composed of 4 steps: advectons in space dimensions, computation of the charge density (reduction) and electric field (Poisson solver) and advectons in velocity dimensions. The algorithm of the simulation is presented in Figure 2.

The remaining of the section presents two optimizations implemented in the *Selalib* [16] version of the 6D Vlasov-Poisson problem.

The 6D nature of `f6D` requires a lot of memory, but the regularity of the problem means it can be distributed in blocks with good load-balancing.

Fig. 3: Dynamic halo exchange representation on a 2D domain

Halos are required to hold values of neighbors for the advections. Connected halo zones would increase the number of points in all dimensions and consume too much memory. Split advections mean that halos are required in a single dimension at a time though. We therefore use dynamic halos composed of two buffers, one for each boundary of the advected dimension (denoted “*right*” and “*left*”). Figure 3 shows this optimization on a 2D domain. Listing 4 shows the INKS_o/XMP implementation of this strategy on 6D Vlasov-Poisson.

```

1  /** INKSo/XMP code for the f6d domain decomposition and dynamic halo exchange */
2  #pragma xmp nodes p6d[pV3][pV2][pV1][pZ][pY][pX]
3  #pragma inks decompose dynamic
4  % f6d // dynamic allocation of f6d algorithmic array
5  (8:, 7:, 6, 5, 4, 3, 2, 1) // dimension reordering, dim 7 and 8 are folded
6  with t6d onto p6d // block decomposition mapped on the XMP topology
7  // Dynamic halo exchange on the 4th dimension, halo sizes are computed automatically
8  // R and L are now allocated buffers and contain the halo values
9  #pragma inks exchange periodic f6d(4, advection4) to R and L
10 foo(R, L);
11 /** Generated C+XMP code for the f6d domain decomposition */
12 double (f6d*) [][][][][][]; // need to declare f6d global, valid in xmp
13 #pragma xmp nodes p6d[pV3][pV2][pV1][pZ][pY][pX] // xmp 6d cartesian node topology
14 #pragma xmp template t6d[:][:][:][:][:][:] // xmp 6d logical array
15 #pragma xmp distribute t6d [block][block][block][block][block][block] onto p6d
16 // map element of f6d to element of t6d
17 #pragma xmp align f6d[n][m][l][k][j][i] with t6d[n][m][l][k][j][i]
18 #pragma xmp template_fix t6d[N][M][L][K][J][I]
19 f6d = (double(*)[M][L][K][J][I]) xmp_malloc(xmp_desc_of(f6d), N, M, L, K, J, I);

```

Listing 4: 6D decomp. & halo exchange in INKS_o/XMP and resulting code

Advections account for the main computational cost of the problem, up to for 95% of the sequential execution time. Six loops surround the stencil computation of each advection and in a naive version, the use of a modulo to handle periodicity and application along non-contiguous dimensions slow down the computation. To enable vectorization and improve cache use, we copy **f6d** elements into contiguous buffers along with the *left* and *right* halos. Advections are applied on these buffers before copying them back into **f6d**. Blocking further improves performance by copying multiple elements at a time. Listing 5 corresponds to the INKS_o/loop implementation of these optimizations in a sequential version and presents the generated code. The Poisson solver

```

1  /** advect.iks InKSpia algorithmic file **/
2  kernel advection3 (i, j, k, l, m, n, t, K, step) : (
3      double f6d {in: (i, j, 0:K, l, m, n, t, step-1) |out: (i, j, k, l, m, n, t, step)},
4      int disp {in: (n)},
5      double coef{in: (n, 0:4)}
6  )
7  #CODE (C)
8  f6d(i, j, k, l, m, n, t, step) =
9      coef(n, 0) * f6d(i, j, k-1+disp(n), l, m, n, t, step-1)
10     + coef(n, 1) * f6d(i, j, k+0+disp(n), l, m, n, t, step-1)
11     + coef(n, 2) * f6d(i, j, k+1+disp(n), l, m, n, t, step-1)
12     + coef(n, 3) * f6d(i, j, k+2+disp(n), l, m, n, t, step-1);
13 #END
14 public kernel main_code(t, I, J, K, L, M, N, Niter) : (
15     double coef(2),
16     int disp(1),
17     double f6d(6) {in: (i, j, 0:K, l, m, n, t, step-1) |out: (i, j, k, l, m, n, t, step)}
18 )
19 #CODE(INKS)
20 advection3 (0:I, 0:J, 0:K, 0:L, 0:M, 0:N, 1:Niter, K, 2) : (f6d, disp, coef)
21 /* ... */
22 #END
23 /** advect.iloop InKSo/Loop optimization choices file **/
24 loop advection3_loops(t, I, J, K, L, M, N, Niter) : advection3 { // set "t" value
25     // "Set" not specified -> loop bounds are computed, with a fixed "t"
26     Order: n, m, l, k, i, j; // order of the loop
27     Block: 16; // blocking on the inner dimension j
28     Buffer: f6d(3); // copy the third dimension of f6d to a 1d buffer
29 }

```

Listing 5: A loop nest in INKSpia optimized in INKSo/loop

relies on a remapping scheme, where the domain decomposition is modified between each FFT execution.

5 Evaluation

This section evaluates the INKS model on the NAS benchmark, a simple stencil code and the 6D Vlasov-Poisson problem. We have implemented the algorithm of the following programs using INKSpia.

- 4 sequential NAS kernels (IS, FT, EP and MG), C++ version [11] as reference;
- finite difference 3D heat resolution (7-point stencil) ([15] as reference);
- 6D Vlasov-Poisson, using Fortran/MPI Selalib [16] as reference.

For the 3d heat equation solver, two strategies were implemented: one uses double buffering (Heat/Buf) and the other implements a cache oblivious strategy (Heat/Obl). To evaluate the INKS programming model on Vlasov-Poisson 6D, we have conducted three experiments. The first one is based on optimization choices written in plain C++ and cover the whole simulation. Then, we evaluate INKSo/XMP and INKSo/loop on Vlasov-Poisson separately as they target different optimizations and are not usable together currently. A first experiment focuses on the sequential aspects with the intra-node optimization of the v_1 advection using either INKSo/loop or INKSo/C++. A second experiment focuses on the parallel aspects with the charge density computation, the Poisson solver and a halo exchange optimized either with C/XMP or with INKSo/XMP.

Algorithm		InKS _{o/C++}	InKS _{o/loop}	InKS _{o/XMP}
Heat	Buf		3 versions: Same as reference, 2D and 3D cache blocking	
	Obl			
NAS	EP			
	FT			
	MG			
	IS			
VP6D	Initialization			
	Advections		adv_v1	
	Reduction			
	Poisson solver			

Table 1: Summary of the InKS_{o/C++}, InKS_{o/loop} and InKS_{o/XMP} optimization choices implemented based on six InKS_{pia} algorithm: the 3D heat solver, 4 NAS kernels (EP, FT, MG, IS) and the Vlasov-Poisson 6D solver.

Benchmark	Execution time (second)			Complexity	
	Reference	InKS _{o/C++}	Rel. dev.	Ref.	InKS
NAS/FT	49.16 (±0.16)	44.00 (±0.05)	11.66%	6	5
NAS/IS	1.91 (±0.00)	1.90 (±0.01)	0.53%	55	52
NAS/MG	4.61 (±0.02)	4.29 (±0.01)	7.46%	20	12
NAS/EP	53.52 (±0.02)	55.35 (±0.37)	-3.30%	19	19
Heat/Buf	2.76 (±2.78)	2.93 (±3.53)	-6.22%	5	3
Heat/Obl	1.20 (±1.16)	1.18 (±0.96)	1.81%	22	13
VP6D	25.92 (±0.21)	21.92 (±0.16)	18.28%	N/A	N/A
Benchmark	Reference (C++)	InKS _{o/loop}	Rel. dev.	Reference (Fortran)	
NAS/IS	1.91 (±0.00)	1.91 (±0.01)	0.00%	N/A	
NAS/MG	4.61 (±0.02)	4.32 (±0.02)	6.84%	N/A	
VP6D	4.18 (±0.05)	4.23 (±0.06)	-1.10%	3.63 (±0.09)	
Benchmark	Reference	InKS _{o/XMP}	Efficiency (# cores)		
NAS/EP	53.52 (±0.02)	14.00 (±0.13)	70.5% (4 cores)		

Table 2: Execution time of the InKS_{o/C++}, InKS_{o/loop} and InKS_{o/XMP} implementations of the sequential NAS benchmark, class B - Time/iteration of the 3D heat equation (7point stencil), size (1024^3) - Time/iteration of the InKS_{o/C++} and Fortran implementation of the sequential 6D Vlasov-Poisson, size (32^6). Median and standard deviation of 10 executions - GNU Complexity score of the implementation.

The table 1 summarizes the optimization choices we have implemented. All codes are compiled with Intel 18 compiler (`-O3 -xHost`), Intel MPI 2018 and executed on the *Irene* cluster (TGCC, France) equipped with 192 GB RAM, two Skylake 8168 CPUs per node and a EDR InfiniBand interconnect. The table 2 summarizes the optimization choices we have implemented.

The NAS CG kernel relies on indirections not expressible in the polyhedron model of InKS_{pia}. Its implementation would thus have to rely on a large C++ kernel whose optimization would be mixed with the algorithm. For the same reason, the NAS FT kernels was only partially implemented. InKS_{pia} can however be used to express all other NAS kernels, the 3D heat equation solver as well as the 6D Vlasov-Poisson algorithm. Even if not as expressive as C or Fortran, InKS_{pia}, through the polyhedron model, handles static controls programs. This covers the needs of a wide range of simulation domains and offers abstractions close to the execution machine rather than from a specific simulation domain. More specifically, it supports programs expressible as Parameterized Task Graphs [7]: a directed acyclic graph of tasks (here kernels)

in a compact representation whose dependence and scheduling are parameterized by integers fixed at the INKS entry point execution. Among others, it can express computations such as FFTs or stencils with input coordinates unknown at compile-time, as in 6D Vlasov-Poisson.

INKS separates the specification of algorithm and optimization in distinct files. Multiple optimization strategies can be implemented for a single algorithm, as shown for the 3D heat equation where each relies on a specific memory layout and scheduling. Similarly, based on the INKS_{pia} algorithm implemented for the IS, EP and MG NAS kernels, we have developed multiple optimization choices based on $\text{INKS}_{\text{o/C++}}$ and either $\text{INKS}_{\text{o/loop}}$ or $\text{INKS}_{\text{o/XMP}}$. For more complex cases such as the 6D Vlasov-Poisson problem, we were able to develop four optimization choices based on one INKS_{pia} algorithm: (1) $\text{INKS}_{\text{o/XMP}}$, (2) C with XMP, (3) $\text{INKS}_{\text{o/loop}}$ and (4) $\text{INKS}_{\text{o/C++}}$. This proves, to some extent, that the separation of concerns is respected.

Finding the right metric to evaluate the easiness of writing a code is a difficult question. As illustrated in Listings 1 and 5 however, algorithm expression in INKS_{pia} is close to the most naive C implementation where loops are replaced by INKS validity domains with no worry for optimization. The polyhedron model used to represent the INKS_{pia} code is compatible with a subset of the C language. Therefore, we could have choose the C as the algorithm language. However, the optimization process would have suffer from such a choice. Indeed, while INKS_{pia} is designed to remove optimizations through the use of DSA and fine granularity kernels, C enables its users to reuse memory and specify a total scheduling. Although it is possible to analyze the code and to retrieve the DSA form and the partial order of the code using the polyhedron model, it will be complicated for optimization specialists to find which loops can be broke or which arrays can be expended from a C code.

Concerning the specification of optimization choices, it is close to their expression in C++. Table 2 compares the GNU complexity score of INKS optimizations to the reference code. INKS scores are slightly better, which indicates that our language is not more complex than C++. The difference comes from the extraction of the computational kernels, placed in the algorithm, which hides parts of the complexity. In addition, the use of $\text{INKS}_{\text{o/C++}}$ to write optimizations let optimization specialists reuse their preexisting knowledge of this language. Similarly writing the $\text{INKS}_{\text{o/C++}}$ version of optimization choices for the INKS_{pia} version of Vlasov-Poisson 6D is close to the expression to the same optimization in Fortran, in the reference version. These considerations should not hide the fact that some information has to be specified both in the INKS_{pia} and $\text{INKS}_{\text{o/C++}}$ files with this approach leading to more code overall.

On contrary, $\text{INKS}_{\text{o/XMP}}$ and $\text{INKS}_{\text{o/loop}}$ enables the developer to specify optimization choices only while algorithmic information is extracted from INKS_{pia} code. This is illustrated by Listing 2 presenting the $\text{INKS}_{\text{o/XMP}}$ 6D domain decomposition and the XMP result. Both are equivalent, but the $\text{INKS}_{\text{o/XMP}}$ expects only optimization choices parameters. Hence, one can test another memory layout, such as a different dimension ordering, by changing

only a few parameters, while multiple directives must be modified in XMP. Moreover, using the domain decomposition directive offered by $\text{INKS}_{\text{o/XMP}}$ and XMP code, one can derive a simple parallel code from a INKS_{pia} algorithm, as we did with the 3D Heat solver. Note that since $\text{INKS}_{\text{o/XMP}}$ is a wrapper for XMP using INKS_{pia} to retrieve some information, it is usable in any codes that can be expressed using INKS_{pia} and optimized with XMP. Similarly, with $\text{INKS}_{\text{o/loop}}$ (Listing 5), developers can easily test different optimization choices that would be tedious in plain C++. This is what we did with the 3D Heat solver based on the double buffering technique. As shown on Figure 1, using $\text{INKS}_{\text{o/loop}}$, we have implemented 3 versions of the loops: the same as reference, one based on a 2D cache-blocking and a last one using a 3D cache-blocking by modifying almost nothing in the $\text{INKS}_{\text{o/loop}}$ code. Since $\text{INKS}_{\text{o/XMP}}$ and $\text{INKS}_{\text{o/loop}}$ are respectively usable with C and C++, INKS does not restrict the expressible optimization choices: one can still implement optimizations not handled by our DSLs in C/C++. Moreover, operations such as halo size computation or vectorization capabilities detection are automatized using the INKS_{pia} code. In summary, the approach enables optimization specialists to focus on their specialty which make the development easier.

Regarding performance, the INKS approach makes it possible to express optimizations that do not change the algorithm. Optimizations of the four NAS parallel benchmarks and 3D heat equation solver in INKS were trivial to implement and their performance match or improve upon the reference as presented in Table 2. Investigation have shown that Intel ICC 18 does not vectorize properly the reference versions of **MG** and **FT**. The use of the Intel **ivdep** directive as done on the INKS versions leads to slightly better performance.

For the full Vlasov-Poisson 6D problem, INKS_{pia} and C++ enable us to implement the same complex optimization strategies written in the reference version. Therefore, our implementation match the reference in terms of performance, as shown on Table 2. For the v_1 advection, both the $\text{INKS}_{\text{o/C++}}$ and $\text{INKS}_{\text{o/loop}}$ optimizations of the INKS code achieve performance similar to the reference as shown in Table 2. For the parallel aspects, the $\text{INKS}_{\text{o/XMP}}$ optimization offers performance similar to XMP as shown on Figure 4. MPI is faster on all cases compared to both XMP and $\text{INKS}_{\text{o/XMP}}$. At the moment, it seems that XMP does not optimize local copies which slows down the Poisson solver. Besides, some XMP directives are based on MPI RMA which make the comparison with MPI Send/Receive complex. Still, MPI is much harder to program: more than 350 lines of MPI and Fortran are required to handle domain decomposition, remapping for FFT and halo exchange in Selalib while 50 lines in XMP and 15 in $\text{INKS}_{\text{o/XMP}}$.

However, INKS_{pia} , $\text{INKS}_{\text{o/loop}}$ and $\text{INKS}_{\text{o/XMP}}$ have limitations. As mentioned earlier, INKS_{pia} can not be used to express non static control program, such as the **CG** or the full **FT** NAS kernel. It also makes more complex the optimization of arrays passed in parameters (as input or output). $\text{INKS}_{\text{o/loop}}$ offers very limited option, making it unusable with complex loops such as the one in the 3D heat solver optimized using the cache-oblivious strategy. Similarly, loop nests calling multiple INKS_{pia} computational kernels can not be expressed

Fig. 4: Weak and strong scaling for 3 parts of the Vlasov-Poisson solver up to 64 nodes (1 process/node) on a 32^6 grid divided among processes (strong scaling) or 16^6 grid per process (weak scaling). Median of 10 executions.

using $\text{INKS}_{\text{o/loop}}$. Moreover, it can not respects the constraints imposed by a specific memory mapping and offers only a few optimization. Adding new optimization strategies would require to add new keywords, such as unrolling, loop fusion, etc. $\text{INKS}_{\text{o/XMP}}$ offers memory allocations but no controls on the access to this memory, letting optimization specialists in charge of the good use of this memory. Besides, it is usable only with XMP.

Although we want to address some issues in the INKS_{pia} language, $\text{INKS}_{\text{o/loop}}$ and $\text{INKS}_{\text{o/XMP}}$ were preliminary tests before a real INKS_{ps} implementation. The goal was to propose a way to express optimizations and retrieving the algorithmic information in INKS_{pia} code. These two DSLs enables us to highlight a set of guidelines for the full definition of the INKS_{ps} language. First, this definition must be based on the concepts that INKS_{ps} must express, i.e. the optimizations related to memory (allocations and layouts) and to computations scheduling. With $\text{INKS}_{\text{o/XMP}}$ we tried some tests on the memory aspects, providing a directive to allocate a logical array and to reorder its dimensions. $\text{INKS}_{\text{o/loop}}$ focus on computations scheduling by adding strong constraint on the order of the computations. Secondly, in order to express these optimization concepts only, it must be bound to its algorithmic counterpart, express in INKS_{pia} . In $\text{INKS}_{\text{o/XMP}}$, the domain decomposition directive makes reference directly to logical arrays described in INKS_{pia} , making possible the computation of the size of each dimension and the halo size depending of the data being accessed. In $\text{INKS}_{\text{o/loop}}$, the **Order** keyword refers to the structuring variable of a computational kernel defined in INKS_{pia} . Finally, we must work on the mean to express these optimization concepts with the most generality. For this part, we think that working with existing tools is probably mandatory for most complex strategy, such as XMP for the domain decomposition. Using another existing tool for computations scheduling, making it more general. Similarly, a declarative language such as $\text{INKS}_{\text{o/loop}}$ may not be the best strategy: to handle the interaction between memory and computation, an imperative language could be easier.

6 Conclusion and future works

In this paper, we have presented the INKS programming model to separate the algorithm and the optimization choices and its implementation supporting two DSLs: $\text{INKS}_{\text{o/loop}}$ for loop optimizations and $\text{INKS}_{\text{o/XMP}}$ for domain decomposition. We have evaluated INKS on synthetic benchmarks and on the Vlasov-Poisson solving. We have demonstrated its generality and advantages in terms of separation of concerns to improve maintainability and portability while offering performance on par with existing approaches.

While this paper demonstrates the interest of the INKS model, it still requires some work to further develop it. We will improve the optimization DSLs; base $\text{INKS}_{\text{o/loop}}$ on existing tools and ensure interactions with $\text{INKS}_{\text{o/XMP}}$. This will be done within the scope of the development of a complete INKS_{pso} DSL enabling its users to manage memory placement and computations scheduling. This planned DSL will truly separates algorithm from optimization choices in static programs while its compiler could offer static analysis of the program to ensure correctness. We also want to target different architectures to demonstrate the portability gains of the INKS model.

References

1. Augonnet C, Thibault S, Namyst R, Wacrenier PA (2011) StarPU: A unified platform for task scheduling on heterogeneous multicore architectures. *Concurr Comput : Pract Exper* 23(2), DOI 10.1002/cpe.1631
2. Aumage O, Bigot J, Ejjaouani K, Mehrenberger M (2017) InKS, a programming model to decouple performance from semantics in simulation codes. Tech. rep., Inria
3. Bailey DH, Barszcz E, Barton JT, Browning DS, Carter RL, Dagum L, Fatoohi RA, Frederickson PO, Lasinski TA, Schreiber RS, et al. (1991) The NAS parallel benchmarks. *Journal of Supercomputing Applications*
4. Carter Edwards H, Trott CR, Sunderland D (2014) Kokkos. *J Parallel Distrib Comput* 74(12), DOI 10.1016/j.jpdc.2014.07.003
5. Chandra R, Dagum L, Kohr D, Maydan D, McDonald J, Menon R (2001) *Parallel Programming in OpenMP*. Morgan Kaufmann
6. Christen M, Schenk O, Burkhart H (2011) PATUS: A Code Generation and Autotuning Framework for Parallel Iterative Stencil Computations on Modern Microarchitectures. In: *Parallel & Distributed Processing Symposium (IPDPS) 2011*, IEEE, DOI 10.1109/ipdps.2011.70
7. Cosnard M, Jeannot E (1999) Compact dag representation and its dynamic scheduling. *Journal of Parallel and Distributed Computing* 58(3), DOI 10.1006/jpdc.1999.1566
8. Danelutto M, García J, Miguel Sanchez L, Sotomayor R, Torquati M (2016) Introducing parallelism by using repara c++11 attributes. pp 354–358, DOI 10.1109/PDP.2016.115

9. El-Ghazawi T, Carlson W, Sterling T, Yelick K (2005) *UPC: Distributed Shared Memory Programming*. Wiley
10. Feautrier P, Lengauer C (2011) *Polyhedron Model*, Springer. DOI 10.1007/978-0-387-09766-4_502
11. Griebler D, Loff J, Mencagli G, Danelutto M, Fernandes LG (2018) Efficient NAS benchmark kernels with c++ parallel programming. In: 2018 26th Euromicro International Conference on Parallel, Distributed and Network-based Processing (PDP), DOI 10.1109/PDP2018.2018.00120
12. Hoque R, Herault T, Bosilca G, Dongarra J (2017) Dynamic task discovery in PaRSEC: A data-flow task-based runtime. In: 8th Workshop on Latest Advances in Scalable Algorithms for Large-Scale Systems, ACM
13. Höhnerbach M, Ismail AE, Bientinesi P (2016) The vectorization of the Tersoff multi-body potential: An exercise in performance portability. In: International Conference for High Performance Computing, Networking, Storage and Analysis, IEEE
14. Isoard A (2016) Extending polyhedral techniques towards parallel specifications and approximations. PhD thesis, École doctorale en Informatique et Mathématiques de Lyon
15. Kamil S (2012) StencilProbe: A microbenchmark for stencil applications. Accessed: 2017-08-25
16. Kormann K, Reuter K, Rampp M, Sonnendrücker E (2016) Massively parallel semi-lagrangian solution of the 6d vlasov-poisson problem.
17. Lee J, Sato M (2010) Implementation and performance evaluation of XscalableMP: A parallel programming language for distributed memory systems. In: International Conference on Parallel Processing Workshops
18. Mehrenberger M, Steiner C, Marradi L, Crouseilles N, Sonnendrücker E, Afeyan B (2013) Vlasov on GPU (VOG project). ESAIM: Proc 43, DOI 10.1051/proc/201343003
19. Steuwer M, Remmelg T, Dubach C (2017) LIFT: A functional data-parallel IR for high-performance gpu code generation. In: 2017 IEEE/ACM International Symposium on Code Generation and Optimization (CGO)
20. Tang Y, Chowdhury RA, Kuszmaul BC, Luk CK, Leiserson CE (2011) The Pochoir stencil compiler. In: 23rd Symposium on Parallelism in Algorithms and Architectures, ACM, SPAA '11, DOI 10.1145/1989493.1989508
21. Tanno H, Iwasaki H (2009) Parallel skeletons for variable-length lists in SkeTo skeleton library. In: Proceedings of the 15th International Euro-Par Conference on Parallel Processing, Springer, Euro-Par '09, DOI 10.1007/978-3-642-03869-3_63
22. Verdoolaege S (2010) isl: An integer set library for the polyhedral model. In: Fukuda K, Hoeven Jvd, Joswig M, Takayama N (eds) *Mathematical Software – ICMS 2010*, Springer