

BOAST: a Metaprogramming Framework to Produce Portable and Efficient Computing Kernels for HPC Applications

Brice Videau^{*†}, Kevin Pouget^{*‡}, Luigi Genovese[§], Thierry Deutsch[§], Dimitri Komatitsch[†], Frédéric Desprez^{*¶},
Jean-François Méhaut^{*‡}

^{*}Laboratoire d'Informatique de Grenoble

[†]Centre National de la Recherche Scientifique

[‡]Université Joseph Fourier

[§]CEA/INAC/SP2M/L_Sim

[¶]Inria

Abstract—The portability of real HPC applications on new platforms is an open and very delicate problem. Especially the performance portability of the underlying computing kernels is problematic as they need to be tuned for each and every platform the application encounters. This paper presents BOAST, a metaprogramming framework dedicated to computing kernels. BOAST allows the description of a kernel and its possible optimizations using a DSL. BOAST runtime will then compare the different versions' performance as well as verify their exactness. BOAST is applied to three use cases: a Laplace kernel in OpenCL and two HPC applications BigDFT (electronic density computation) and SPECFEM3D (seismic and wave propagation).

I. INTRODUCTION

Porting and tuning HPC applications to new platforms is tedious and costly in term of human resources. Many developer teams are more concentrated on solving the physical problems faster and better with new algorithms and simulating new physics. This is why optimization, while important, is a task which is sometimes neglected or hidden for many developers. In the same way, Domain Specific Language (DSLs) have exactly this goal, the complete separation of the high-level code dedicated to simulate new physics and the low-level one which is more dependent on the HPC architecture. Nevertheless, decoupling the high and low levels does not avoid the need of optimizations which become more and more unavoidable with new architectures based on many cores and accelerators.

Unfortunately, portability efforts of the low level codes (i.e. computing kernels) are often lost when migrating to a new architecture. Worse, code may lose maintainability because several versions of some functionalities coexist, usually with a lot of duplication. Thus productivity of porting and tuning efforts is low as a huge fraction of those developments are never used after the platform they were intended for is decommissioned.

The primary goal of the BOAST metaprogramming framework is to simplify the optimization of computing kernels for HPC applications. The main idea is to use a high level

language to define an API of computing kernels i.e. sub-routines and their related arguments, but also describe the way to optimize it (unrolling, vectorization, parametrization, templating, tiling). BOAST will then test automatically the many possibilities of optimizations and choose the best one after validation from a reference version. Then BOAST can generate pieces of Fortran, C, OpenCL, or CUDA codes, test them for a specific target architecture.

With the definition of an API, BOAST can hide the implementation details at high level but also is a way to factorize different kernels depending only on parameters such as data types, length of filters, or boundary conditions. The maintainability and genericity of HPC codes is often limited. One of the reason is that producing generic code in Fortran 90/95 is difficult as the language does not really fit for it. Sometimes, adding genericity or functionalities can be quite costly and may degrade performance as optimization opportunities that come from over-specification are lost. BOAST can provide a solution to handle genericity and add new functionalities at low levels.

The paper is organized as follows. First in Section II we further motivate the need for metaprogramming frameworks. Secondly in Section III we illustrate the optimization problems related to scientific computing with the Laplace equation. Thirdly we present our solution, called BOAST, its architecture and its functionalities in Section IV. Fourthly we apply BOAST to three use cases: the Laplace equation motivating example, short multi-dimensional convolutions from the BigDFT quantum chemistry application, and the seismic wave propagation simulator SPECFEM3D where 80 GPU computing kernels have been ported improving considerably the maintainability and providing a single code base for OpenCL and CUDA kernel versions. Finally, before a conclusion and future works, we compare BOAST with other related works in Section VI.

II. BACKGROUND AND MOTIVATION

The motivation of the BOAST project comes from the need of the HPC communities which have full and large applications. These applications are difficult to port and maintain

for emerging architectures especially since new architectures are based on accelerators, longer vector units, or low power processors.

A. Evolution of HPC Architectures

Evolution of HPC Architectures is rapid and also diverse. In the last seven years no less than 7 architectures have been number one in the Top500 [1]:

- Sunway processor (TaihuLight, 2016)
- Intel Processor + Xeon Phi (Tianhe-2, 2013)
- AMD Processor + NVIDIA GPU (Titan, 2012)
- IBM BlueGene/Q (Sequoia, 2012)
- Fujitsu SPARC64 (K computer, 2011)
- Intel Processor + NVIDIA GPU (Tianhe-1, 2010)
- AMD Processor (Jaguar, 2009)

Being able to efficiently use those architectures on such a small period is challenging for scientists and application developers.

The race to exascale is not going to simplify the environment. All of the above machine architectures can be considered to build these supercomputers and other will appear soon. For instance European FP7 project DEEP considers using Accelerators (XEON Phi) while the European FP7 project Mont-Blanc considers using low-power embedded processor with integrated GPU.

Running existing applications on these new architectures is an open research subject as well as an ongoing porting effort. Thus, those projects have work packages dedicated to applications. Those work packages are dedicated to porting and optimizing Scientific applications on those new architectures. In the DEEP project six applications were selected for porting and optimizing, eleven were selected in the Mont-Blanc project.

B. Scientific Computing Applications

Scientific Computing Applications are usually developed by other science researchers such as physicists, chemists or meteorologists. Those codes are usually written in Fortran for historical and performance reasons. Codes can be quite huge (several thousands Lines of Code, LoC) with lots of functionalities. Nonetheless, they are usually based on computing kernels. Computing kernels are resource intensive and well-defined parts of a program that usually work on precisely defined data. Those kernels represent the most time-consuming part of an HPC application, and consequently, they are the prime target for optimizations.

Those applications are often developed by several individuals. Sometimes, some of those developers only work a few months on the application. Maintaining optimized code written by someone else is quite a challenge. Several languages and programming paradigms can also be used in a project and thus the maintainer must be knowledgeable in several areas of expertise. Portability problems can also be caused by the availability of an optimizing compiler for a specific language and architecture. C compilers are usually first available while Fortran may sometimes come later.

In Section V we will present two HPC applications that we used as use cases: SPECfem3D and BigDFT. They are both based on computing kernels and were selected as candidate applications in the Mont-Blanc project.

C. How Should Computing Kernels be Written?

The problematic here is to obtain computing kernels that present good performance on the architectures encountered by the application while staying portable after the optimization process took place. Indeed the application might encounter one of the many architectures that can be found in HPC. Investing manpower to optimize the application for a new architecture is reasonable, suffering hindrance from previous optimization work is not. Thus optimizations have to be as orthogonal as possible from one another so as to be easily activated and deactivated.

If this paradigm is followed by developers then they will rapidly be confronted with a huge optimization space to search. They will need to be able to test easily the performance impact of the chosen optimizations without running the full application. The same reasoning implies that kernels should be tested for non regression without running the full application.

What we want is computing kernels that are written in a *portable* manner, in a way that raises developer *productivity*, and they must present good *performance*.

III. BASIC EXAMPLE: OPENCL LAPLACE

During one of the Mont-Blanc face to face meeting a talk on OpenCL optimization on the Mali GPU was given. One of the case study was a Laplace filter [2]. This is a good example of an algorithm that is simple in its formulation but can be complex to optimize because many different optimizations are available and can be combined together. The correct combination of optimizations will depend on the target architecture.

A. The Laplace Filter

Listing 1 shows a simple implementation of the Laplace filter in C99. For sake of clarity, boundary conditions have been omitted.

```
void laplace(const int width, const int height,
             const unsigned char src[height][width][3],
             unsigned char dst[height][width][3]){
    for (int j = 1; j < height-1; j++) {
        for (int i = 1; i < width-1; i++) {
            for (int c = 0; c < 3; c++) {
                int tmp =
                -src[j-1][i-1][c] - src[j-1][i][c] - src[j-1][i+1][c]\
                - src[j][i-1][c] + 9*src[j][i][c] - src[j][i+1][c]\
                - src[j+1][i-1][c] - src[j+1][i][c] - src[j+1][i+1][c];
                dst[j][i][c] = (tmp<0 ? 0 : (tmp>255 ? 255 : tmp));
            }
        }
    }
}
```

Listing 1. Laplace C99 Filter.

A naive implementation in OpenCL is proposed in Listing 2. Each work item processes one pixel of the resulting image.

```
kernel laplace(const int width, const int height,
               global const uchar *src,
               global uchar *dst){
    int i = get_global_id(0);
```

```

int j = get_global_id(1);
for (int c = 0; c < 3; c++) {
    int tmp = -src[3*width*(j-1) + 3*(i-1) + c]\
        - src[3*width*(j-1) + 3*(i) + c]\
        - src[3*width*(j-1) + 3*(i+1) + c]\
        - src[3*width*(j) + 3*(i-1) + c]\
        + 9*src[3*width*(j) + 3*(i) + c]\
        - src[3*width*(j) + 3*(i+1) + c]\
        - src[3*width*(j+1) + 3*(i-1) + c]\
        - src[3*width*(j+1) + 3*(i) + c]\
        - src[3*width*(j+1) + 3*(i+1) + c];
    dst[3*width*j + 3*i + c] = clamp(tmp, 0, 255);
}
}

```

Listing 2. Laplace OpenCL Filter.

B. Possible Optimizations of the OpenCL Version

Several optimizations were proposed and successively applied to the previous implementation.

a) *Vectorization*: The first proposed optimization involves computing five pixels instead of one using the vectors available on the Mali architecture. The vectors have to be kept in memory to ensure an efficient execution. Using vectors of 16 elements, 15 useful components are simultaneously loaded and can be used in computation. This optimization yields speedups between 1.5 and 6 depending on the image size.

```

kernel laplace(const int width, const int height,
              global const uchar *src,
              global uchar *dst){
    int i = get_global_id(0);
    int j = get_global_id(1);

    uchar16 v11_ = vload16(0, src + 3*width*(j-1) + 3*5*i -3);
    uchar16 v12_ = vload16(0, src + 3*width*(j-1) + 3*5*i );
    uchar16 v13_ = vload16(0, src + 3*width*(j-1) + 3*5*i +3);
    uchar16 v21_ = vload16(0, src + 3*width*(j) + 3*5*i -3);
    uchar16 v22_ = vload16(0, src + 3*width*(j) + 3*5*i );
    uchar16 v23_ = vload16(0, src + 3*width*(j) + 3*5*i +3);
    uchar16 v31_ = vload16(0, src + 3*width*(j+1) + 3*5*i -3);
    uchar16 v32_ = vload16(0, src + 3*width*(j+1) + 3*5*i );
    uchar16 v33_ = vload16(0, src + 3*width*(j+1) + 3*5*i +3);

    int16 v11 = convert_int16(v11_);
    int16 v12 = convert_int16(v12_);
    int16 v13 = convert_int16(v13_);
    int16 v21 = convert_int16(v21_);
    int16 v22 = convert_int16(v22_);
    int16 v23 = convert_int16(v23_);
    int16 v31 = convert_int16(v31_);
    int16 v32 = convert_int16(v32_);
    int16 v33 = convert_int16(v33_);

    int16 res = v22 * (int16)9 - v11 - v12 - v13 - v21 - v23 -
        v31 - v32 - v33;
    res = clamp(res, (int16)0, (int16)255);
    uchar res_ = convert_uchar16(res);

    vstore8(res_.s01234567, 0, dst + 3*width*j + 3*5*i);
    vstore8(res_.s89ab, 0, dst + 3*width*j + 3*5*i + 8);
    vstore8(res_.scd, 0, dst + 3*width*j + 3*5*i + 12);
    dst[3*width*j + 3*5*i + 14] = res_.se;
}

```

Listing 3. Laplace OpenCL Filter Vectorized.

b) *Synthesizing Loads*: It is possible to reduce the number of loads since the vectors are overlapping. Listing 4 shows how the vectors are loaded in this case. This optimization yields marginal improvements.

```

uchar16 v11_ = vload16(0, src + 3*width*(j-1) + 3*5*i -3);
uchar16 v13_ = vload16(0, src + 3*width*(j-1) + 3*5*i +3);
uchar16 v12_ = uchar16(v11_.s3456789a, v13_.s56789abc);
uchar16 v21_ = vload16(0, src + 3*width*(j) + 3*5*i -3);
uchar16 v23_ = vload16(0, src + 3*width*(j) + 3*5*i +3);
uchar16 v22_ = uchar16(v21_.s3456789a, v23_.s56789abc);

```

```

uchar16 v31_ = vload16(0, src + 3*width*(j+1) + 3*5*i -3);
uchar16 v33_ = vload16(0, src + 3*width*(j+1) + 3*5*i +3);
uchar16 v32_ = uchar16(v31_.s3456789a, v33_.s56789abc);

```

Listing 4. Laplace OpenCL Filter Vectorized with Synthesized Loads.

c) *Temporary Variables Size*: Using *int* to store intermediary results is unnecessary. Listing 5 show how the code is modified to use smaller types. This yields a speedup of 1.3 for most image sizes.

```

short16 v11 = convert_short16(v11_);
short16 v12 = convert_short16(v12_);
short16 v13 = convert_short16(v13_);
short16 v21 = convert_short16(v21_);
short16 v22 = convert_short16(v22_);
short16 v23 = convert_short16(v23_);
short16 v31 = convert_short16(v31_);
short16 v32 = convert_short16(v32_);
short16 v33 = convert_short16(v33_);

short16 res = v22 * (short16)9 - v11 - v12 - v13 - v21 -
    v23 - v31 - v32 - v33;
res = clamp(res, (short16)0, (short16)255);

```

Listing 5. Laplace OpenCL Filter Vectorized Using *short*.

d) *More Optimizations*: Several other optimizations can be attempted at this point. For instance reducing or increasing the number of pixels each work item is processing. They can yield improved performance in some cases, especially when avoiding memory alignment problems. These optimizations will not be shown here but can be found in [2]. Others consist in tiling the problem space or improving the loads synthesis.

C. Improving the Methodology

The process described in the previous Subsection is quite tedious and requires some intimate knowledge of the target architecture. The results are impressive, as speedups of almost 10 are observed compared to the naive OpenCL implementation.

Nonetheless, the process is frustrating. The optimizations are never evaluated independently from one another. Some were arbitrarily configured (like the number of pixels chosen for instance). Testing the different combinations of those optimizations and the different parameters would be very costly in developer time and, moreover, like all repetitive and tedious tasks, error prone. Thus every created kernel would have to be thoroughly tested to ensure no error was done.

In the next Section we will present our answer to these issues: BOAST. BOAST is a framework dedicated to kernel description, optimization, regression testing, and autotuning.

IV. BOAST: USING CODE GENERATION IN APPLICATION AUTOTUNING

BOAST provides scientific application developers with a framework to develop and test application computing kernels [3].

Figure 1 illustrates the envisioned workflow and program structure. The user starts from an application kernel (either designed or implemented), and writes it in a dedicated language (Step 1). The language provides enough flexibility for the kernel to be meta-programmed with several orthogonal optimizations. Then the developer selects a kernel configuration and a target language. Those parameters define the output source code that will be generated by BOAST (Step

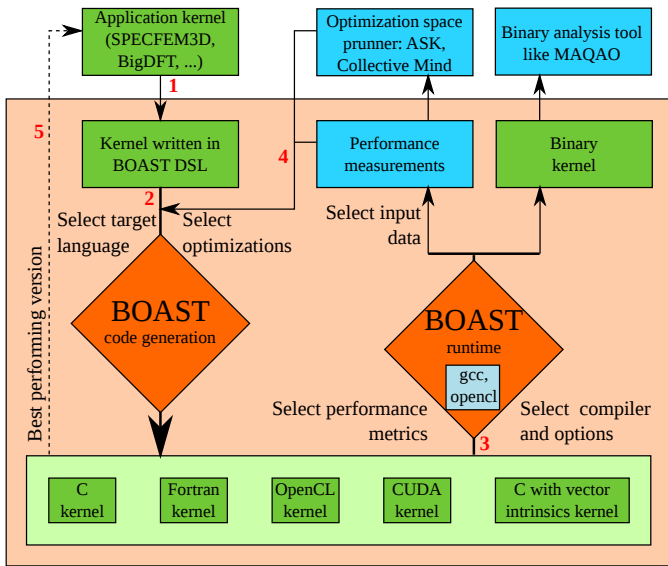


Fig. 1. Structure and Workflow of the BOAST framework.

2). The resulting code source is then built according to the user specified compiler and options (Step 3). If input data are available then the kernel can be benchmarked and tested for non regression. Based on the results, other optimizations can be selected (Step 4) or new optimizations can be added to the BOAST sources. The process can be repeated until a good candidate is found on the target platform. The resulting kernel is then added to the program (Step 5).

Several steps in the workflow can be automated: optimization and compiler flags exploration, non regression testing, benchmarking, ... This automation can be scripted by the user or by interfacing other dedicated tools. In order to achieve those goals three aspects should be considered: code description, code generation, and kernel execution run-time.

A. Kernel Description Language

Usually computing kernels are hotspots of an HPC application, and they take most of the time based on loop nests. A lot of efforts are dedicated to their tuning and the obtained result is often quite different from the original procedure. Several transformations can be applied to such kernels. And those optimizations are often applied manually as compiler may fail to recognise the opportunity.

There are many different loop optimization techniques [4]. We can cite loop skewing [5] (derives nested loops wavefronts) or loop interchange [6] (loop variables change places). The importance of correct loop imbrication on BLAS [7] operations is studied in [8], and shows performance increase of a factor up to 5 when using correct loop imbrication. The importance of code transformation is stressed in [9], where a selection of GPU kernels are ported to CPU and optimized.

BOAST kernel description language should be able to express all these optimizations. This gives us a set of constraints to implement in the language:

- Arbitrary number of variables have to be created and manipulated (types, attributes, ...).

- Procedures have to be abstracted (reunite Fortran and C like languages, attributes, ...).
- Functions must be available.
- Variables, constants and functions should be able to be composed in complex expressions.
- Basic control structures (for, while, if/else, ...) have to be abstracted.
- Powerful array management features (allowing several dimensions, transformations, indexing, ...).

In order to manipulate those abstractions we want to have a syntax similar to what programmers use. For instance, commonly used operators have to be available and behave as expected. It must also be possible to differentiate an action on an abstraction in the context of an expression and in the context of the management of this expression. $c = a + b$, an expression that affects the results of $a + b$ to c is not equal to $c \leftarrow a + b$, which saves the expression $a + b$ to a variable c . This is why an embedded Domain Specific Language approach was selected [10]. This feature allows for the coexistence of two languages: the host language and the Domain Specific Language (DSL). In our case, DSL allows the description of the kernel ($c = a + b$) while the host language provides the meta programming of the kernel ($c \leftarrow a + b$). Operator overloading of the host language will provide the familiar syntax programmers are accustomed to.

It was also important to have our constructs like *for* loops to have a syntax approaching those commonly found in programming languages. To this end, we needed a language which could seamlessly pass a block of code to a function. Ruby [11] is one such language. It has deep introspection capabilities as well. This is the main reason why it was selected for BOAST.

One of the added advantage of using a high-level scripting language like Ruby as the host language if its interfacing capabilities, providing an easy way to use the libraries needed during the development of the framework.

1) *BOAST Keywords*: In order to clearly differentiate what is going to be generated from what is related to manipulations in the host language four keywords were defined. They are *decl*, *pr*, *opn* and *close*. As the language is an Embedded Domain Specific Language (EDSL), these four keywords are methods in the BOAST namespace. Sample usage of these keywords will be found in the next Figures.

The *decl* method is used to declare variables or procedures and functions. The *pr* method calls the public *pr* method of objects it is called on. Each BOAST object is responsible for printing itself correctly depending on the BOAST configuration at the time the print public method is called. Calling directly the *pr* method of a BOAST object yields the same result. The *opn* method can be used to print the beginning of a control structure without an associated code block. The *close* method is the counterpart to the *opn* method. It is used to close a control structure without an associated code block.

2) *BOAST Abstractions*: BOAST defines several classes that are used to represent the structure of the code. These classes can be sorted in two groups, algebraic related and control flow related.

a) *Algebra*: The first and most fundamental abstraction is named Variable. Variables have a type, a name, and a set of

```

1 i = BOAST::Int("i") # or BOAST::Variable::new("i", Int)
2 k = BOAST::Int("k", :size => 8)
3 l = BOAST::Real("l", :dim => [ BOAST::Dim(-5, 21) ], :local
=> true )
4 BOAST::decl i, k, l
5 BOAST::pr i === 5
6 j = i + 5
7 BOAST::pr k === j * 2
8 BOAST::pr l[k] === 1.0
9 BOAST::register_funcall("sin")
10 BOAST::pr l[k+1] === BOAST::sin(j)

```

Listing 6. BOAST code.

```

1 integer(kind=4) :: i
2 integer(kind=8) :: k
3 real(kind=8), dimension
  (-5:21) :: l
4 i = 5
5 k = (i+5)*(2)
6 l(k) = 1.0_wp
7 l(k+1) = sin(i+5)

```

Listing 7. Fortran output.

```

1 int32_t i;
2 int64_t k;
3 double l[27];
4 i = 5;
5 k = (i+5)*(2);
6 l[k-(-5)] = 1.0;
7 l[k+1-(-5)] = sin(i+5);

```

Listing 8. C output.

Fig. 2. BOAST Code Snippet for Variables and Expressions.

named attributes. The existing attributes are mainly inspired from Fortran. Those attributes are not limited and can be arbitrarily enriched, allowing a lot of flexibility in Variable management.

The second abstraction is named Expression. It combines variables into algebraic or logic expressions. Most of the classical operators are overloaded for those two abstractions and thus the syntax of the expressions are rather straightforward. The exception is the assignment operator as it is important to differentiate between assigning an Expression or a Variable in the Ruby context and the assignment operator in the context of a BOAST Expression. Thus the assignment in a BOAST Expression is represented as the `===` operator, while the classical assignment is kept as the `=` operator. Function calls (FuncCall) are also abstracted and can be used in Expressions.

Figure 2 shows some basic usage of both Variables and Expressions as well as the `pr` and `decl` keywords. For clarity we stayed out of BOAST namespace so BOAST related class and methods are prefixed with `BOAST::`. Listing 6 shows the BOAST code that produces the Fortran (Listing 7) and C (Listing 8) output. First we define two variables `i` and `k` (note that `k` is 64 bit integer). The third variable named `l` is a one dimensional local array of length 27, it is indexed in the range `-5` to 21. All those Variables are affected to Ruby variables of the corresponding name.

On Line 4 we declare those three variables. Variable `i` is then affected the value 5. On Line 6 the `j` Ruby variable is used to store the BOAST expression `i + 5`. This variable will be used transparently through the rest of the program.

On Line 8 we use variable `k` to index into array `l` using the bracket operator. Would the array be multidimensional, the index would be comma separated, similar to Fortran notation.

On the last line, a call to the `sin` function is made through the creation of a FuncCall object. The possibility to use this method is declared using the `register_funcall` method.

b) Control Structures: The classical control structures are implemented. *If*, *For*, *While*, *Case* are abstractions in

```

BOAST::register_funcall("modulo")
i = BOAST::Int("i")
j = BOAST::Int("j")
BOAST::pr j === 0
BOAST::pr BOAST::For(i, 0, 100) {
  BOAST::pr BOAST::If( BOAST::modulo(i,7) == 0 ) {
    BOAST::pr j === j + 1
  }
}

```

Listing 9. BOAST code.

```

j=0
do i=0, 100, 1
  if (modulo(i, 7)==0) then
    j=j+1
  end if
end do

```

Listing 10. Fortran output.

```

1 j=0;
2 for(i=0; i<=100; i+=1){
3   if(modulo(i, 7)==0){
4     j=j+1;
5   }
6 }

```

Listing 11. C output.

Fig. 3. BOAST Code Snippet for control structures.

```

n = Int("n", :dir=>:in)
ndat = Int("ndat", :dir=>:in)
x = Real("x", :dir=>:in, :dim=>[Dim(0, n-1), Dim(ndat)])
y = Real("y", :dir=>:out, :dim=>[Dim(ndat), Dim(0, n-1)])
p = Procedure("magicfilter", [n, ndat, x, y])
open p
close p

```

Listing 12. BOAST code.

```

SUBROUTINE magicfilter(n, ndat, x, y)
  integer(kind=4), intent(in) :: n
  integer(kind=4), intent(in) :: ndat
  real(kind=8), intent(in), dimension(0:n-1), ndat) :: x
  real(kind=8), intent(out), dimension(ndat, 0:n-1) :: y
END SUBROUTINE magicfilter

```

Listing 13. Fortran output.

```

void magicfilter(const int32_t n, const int32_t ndat,
  const double * x, double * y) {
}

```

Listing 14. C output.

Fig. 4. BOAST Code Snippet for Procedure.

BOAST matching the behavior of corresponding control structures in other languages. An exception is the *For* in BOAST matches more closely the *For* in Fortran than the one in C.

Figure 3 shows some basic usages of the control structures. The example shows a C macro or function that behaves similarly to the Fortran modulo intrinsic. The sample script (inefficiently) computes and stores in `j` the number of multiples of 7 in the 0 to 100 range. It uses the *For* and *If* control structures. A Ruby block is passed to each of those constructs. This block is evaluated at the time the construct is printed. If several such constructs are needed (in an *if elsif else* case, for instance) they can be explicitly passed as parameters using the *lambda* Ruby keyword.

The last control structure is Procedure. It describes either procedures or functions. Code in Figure 4 presents the use of this abstraction. It illustrates the signature of a real kernel from BigDFT. This kernel uses an input array `x` and an output array `y`, both composed of double precision numbers. Those arrays have two dimensions which depend on input variables `n` and

ndat. We can see here the use of the *open* and *close* keywords that are used to print a control structure without an associated Ruby block. This time we placed ourselves inside the BOAST namespace.

The generated outputs in Fortran (Listing 13) and C (Listing 14) show the difference in meta-information that is kept between both versions.

B. BOAST Run-time

In the previous Section we presented the BOAST's language. This allows us to describe procedures and functions and to meta-program them using Ruby. Each version has to be compiled, linked, and executed to assess its performance in order to find the best version of a computing kernel. This can be very time consuming if this process cannot be automated. By enabling more versions for evaluation, the automation will bring improved portability, better performance, and in the end will improve the productivity of the developer.

In this section we will present the different aspects of BOAST run-time that allow this automation. Those aspects are multi-target language generation (*performance, portability*), kernel compilation (*productivity, performance*), kernel execution (*productivity, performance*) and last, kernel tracer, dumper and replay for non regression tests (*productivity*).

1) *Multi-target Language Generation*: Language availability and performance varies between platforms. It is thus important to express computing kernels in different languages, based on the availability and their respective merits on the target platform. Some languages have additional features, such as languages that target GPUs (OpenCL, CUDA), or languages that support multi-threading paradigms (OpenMP). The developer should be given tools to determine what kind of language is currently used in order to be able to use those additional features.

Similarly the target language must be changed with ease in order to compare different alternatives. Two methods are dedicated to this task: *set_lang* and *get_lang*. The target language can also be set through an environment variable before launching BOAST, allowing for easy command line scripting.

2) *Compilation*: Compilation of the generated kernels must also be very flexible because HPC application developers may encounter platforms with very diverse compilation environments. Proprietary and dedicated compilers are common on HPC infrastructures. Thus BOAST build system exhibits similar behavior to common build systems. Compilers and their compile/build options can be specified at several places. The list by increasing order of precedence includes: BOAST configuration file, environment variables, and at kernel build time.

This way the framework to test different compiler optimizations is completely available and performance study can include both kernel related parameters and compilation related parameters. This behavior is contained in the *CKernel* class of BOAST. When instantiating this class, a BOAST *Procedure* representing the entry point of the kernel is specified.

3) *Execution*: The next logical step is to benchmark the built kernel. BOAST offers a simple way to run a kernel that was successfully built without the need to fork another process. A built BOAST *CKernel* exposes a run method that accepts arguments corresponding to the BOAST *Procedure* used to instantiate the kernel. Arrays must be instances of *NArray* which are numerical arrays that use C arrays underneath.

Arrays which correspond to output parameters will be modified during the execution of the kernel so results can be checked. This allows for easy non regression testing. The run method also returns information (and result for kernels that are functions as well as output scalars) about the run. For instance, one such information includes the run time of the kernel and it is obtained using the system-wide real-time clock. Other probes can be inserted at compile-time, if needed. BOAST also supports PAPI [12] to capture hardware performance counters during each kernel execution.

4) *Kernel Replay*: The non regression methodology presented before is viable provided input data can be generated at run-time. For instance, in the case of the Laplace kernel, generating an input image and the corresponding reference output image, using a reference implementation, is easy. Unfortunately it is not always possible, because some applications have complex data patterns that are difficult to synthesize without running the full application. Thus BOAST offers a way to load binary data from the file system and use them as inputs of a kernel. Outputs can also be checked against those binary data thus enabling (almost) data oblivious non regression testing.

In order to use this methodology one has to be able to trace an application to get input data. Such a tracer, dedicated to CUDA and OpenCL, will be presented in the next subsection.

C. Non Regression Testing Using Trace Debugging

Debugging applications running on GPU environments is well-recognized as a hard and time-consuming activity. In complement with BOAST, we designed a trace-based debugging tool that simplifies this porting operation. The tool relies on BOAST support of multi-target code generation (Section IV-B1), used to validate an application from one GPU-programming framework to another or between different implementations using the same programming model. A case-study of the port of a CUDA application to OpenCL is presented in Section V-C.

The idea behind this tool is based on the assumption that the different GPU ports of the code should perform the very same operations, at least at the logical level (the APIs will have *implementation* differences, but they should offer nonetheless same functionalities). The usage of BOAST framework sustains this assumption, as both sets of kernels should be generated from the same source code.

Hence, the verification and validation of the new port can be narrowed down to asserting that both codes apply the same operations on the GPU. And the debugging part will consist in understanding what diverges. To that purpose, we developed GPUTrace, inspired by *strace* and *ltrace* tools: GPUTrace dynamically preloads a library between the

application and the GPU library, and collects the function name, execution range and argument values (input *and* output) of the relevant function calls. These information are traced in a unified format for all the APIs. This is a custom trace format. Listing 15 presents a sample output generated during SPECFEM3D tracing.

However, in contrast with `strace` and `ltrace`, `GPUTrace` has to be *state-full*. Indeed, most of API parameters are handles to opaque types. So, in order to generate meaningful traces, `GPUTrace` gathers information about these objects at creation time (handle value, buffer creation size and attributes, kernel name and prototype, *etc...*) and re-injects this information when the objects used. A state-less implementation of `GPUTrace` would highly rely on the GPU libraries introspection capabilities, which seem not possible at the moment.

```

1 New kernel: update_potential_kernel
2 ...
3 New buffer #94, 2Mb, READ_WRITE (0x74d0420)
4 New buffer #95, 2Mb, READ_WRITE (0x74d08d0)
5 New buffer #96, 2Mb, READ_WRITE (0x74d0d80)
6 ...
7 Buffer #94 written , 2314764b at +0b
8 Buffer #95 written , 2314764b at +0b
9 Buffer #96 written , 2314764b at +0b
10 ...
11 update_disp_veloc_kernel <128,1><4522,1><(
12     float *displ=<buffer #94 1.00e-24, 1.00e-24,
13         1.00e-24, ... >
14     float *veloc=<buffer #95 0.00e+00, 0.00e+00, 0.00e+00,
15         0.00e+00, ... >
16     float *accel=<buffer #96 0.00e+00, 0.00e+00, 0.00e+00,
17         0.00e+00, ... >
18     const int size=<578691>
19     const float deltat=<1.365914e-04>
20     const float deltatssqover2=<9.328606e-09>
21     const float deltatover2=<6.829570e-05>
22     <out> float *displ=<buffer #94 1.000e-24, 1.00e-24,
23         1.00e-24, 1.00e-24, ... >
24     <out> float *veloc=<buffer #95 0.000e+00, 0.00e+00,
25         0.00e+00, 0.00e+00, ... >
26     <out> float *accel=<buffer #96 0.000e+00, 0.00e+00,
27         0.00e+00, 0.00e+00, ... >
28 );

```

Listing 15. GPUTrace sample trace

Once two call traces have been generated by `GPUTrace`, the user can compare them with a graphical `diff` tool, and spot the different porting mistakes: two parameters reversed, an offset incorrectly applied, *etc...*

By default, `GPUTrace` only prints a unique identifier (the creation index) for the memory buffers. Additionally, `GPUTrace` supports several modifier flags. One flag can be activated to append the first bits of the buffer to the trace, printed in the right format, for visual inspection. Another flag can be set to dump the whole content of the buffer into a file, for a full inspection.

This last option is also useful to generate *replay buffers* for BOAST kernel execution. With a set of filters based on kernel names and execution counters, developers can precisely select which kernel execution parameters should be dumped, for further reuse as real-case benchmarks and non-regression testing.

V. BOAST USE CASES

In this section we will present the benefits of using BOAST on the Laplace motivating example as well as on two scientific applications. The first one, BigDFT [13], uses BOAST in order to develop new functionalities with performance portability in mind. The second one, SPECFEM3D [14], uses BOAST to factorize OpenCL and CUDA development, while having robust non regression tests.

A. Laplace Filter Kernel

Section III presented the Laplace motivating example. From this section we know that a number of optimizations can have an impact on the performance of the kernel on the Mali architecture. But, what is the impact in other architectures? And, are there any additional optimizations that can impact the performance?

1) *Optimization Space*: The list of already identified optimizations are vectorization, intermediary data type, number of pixels processed, and synthesizing loads. To create our generic implementation we decided to work at the component level rather than the pixel level. This approach leads to more flexibility and genericity when applying optimizations. We also decided to study the impact of another parameter which is the number of components to process on the column direction. This leads to being able to process tiles instead of only rows.

Here are the parameters we finally selected for our kernel optimization and their possible values:

- *x_component_number*: a positive integer
- *y_component_number*: a positive integer
- *vector_length*: 1, 2, 4, 8 or 16
- *temporary_size*: 2 or 4
- *synthesize_loads*: true or false
- *vector_recompute*: true or false

The last parameter is used when *x_component_number* is not divisible by *vector_length*. Two solutions are possible then, divide the remainder of the division in vectors of smaller sizes (*vector_recompute = false*) or load more data and compute useless values in vectors of the specified size (*vector_recompute = true*). This last option mimics the behavior of the ARM implementation, although when working at the component level it may not be a valuable thing to try.

2) *Performance Results*: Table I show the best results obtained by ARM on different images compared to the naive implementation and to the best version BOAST found. It shows that the generated version systematically outperforms the hand optimized version. As far as the optimization options are concerned, the results are disappointing: the same kernel configuration is the best for all image sizes. This kernel uses *x_component_number = 16*, *y_component_number = 1*, *vector_length = 16*, *temporary_size = 2* and *synthesize_loads = false*. *vector_recompute* because *x_component_number == vector_length*. Those results show that, when working on full vectors, synthesizing the loads is harmful to performance and the programmer is better of trusting the cache to load each vector in one cycle without compromising the bandwidth. The results shown here are the best of four runs for each configuration.

TABLE I
BEST PERFORMANCE OF ARM LAPLACE KERNEL.

| Image Size | Naive (s) | Best (s) | Accel. | BOAST (s) | Accel. |
|-------------|-----------|----------|--------|-----------|--------|
| 768 x 432 | 0.0107 | 0.00669 | x1.6 | 0.000639 | x16.7 |
| 2560 x 1600 | 0.0850 | 0.0137 | x6.2 | 0.00687 | x12.4 |
| 2048 x 2048 | 0.0865 | 0.0149 | x5.8 | 0.00715 | x12.1 |
| 5760 x 3240 | 0.382 | 0.0449 | x8.5 | 0.0325 | x11.8 |
| 7680 x 4320 | 0.680 | 0.0747 | x9.1 | 0.0581 | x11.7 |

TABLE II
BEST PERFORMANCE OF LAPLACE KERNEL ON SEVERAL ARCHITECTURES.

| Image Size | ARM | Intel | Ratio | NV | Ratio |
|-------------|----------|----------|-------|-----------|-------|
| 768 x 432 | 0.000639 | 0.000222 | x2.9 | 0.0000715 | x8.9 |
| 2560 x 1600 | 0.00687 | 0.00222 | x3.1 | 0.000782 | x8.8 |
| 2048 x 2048 | 0.00715 | 0.00226 | x3.2 | 0.000799 | x8.9 |
| 5760 x 3240 | 0.0325 | 0.0108 | x3.0 | 0.00351 | x9.3 |
| 7680 x 4320 | 0.0581 | 0.0192 | x3.0 | 0.00623 | x9.3 |

But what if we run our benchmark on other architectures? Table II shows the results obtained when running our BOAST implementation on other architectures. The chosen architectures include an Intel i7-2760QM CPU (Sandy Bridge architecture) that supports OpenCL 1.2 and a system with an NVIDIA gtx680 GPU that supports OpenCL 1.1. We see that the performance ratio between the different architectures is stable across image sizes.

The optimization parameters selected are not the same for those architectures. Indeed, the Intel CPU favors kernels that have the parameters, *x_component_number* = 16, *vector_length* = 8, *temporary_size* = 2 and *synthesize_loads* = false. Once again *vector_recompute* does not apply. *y_component_number* varies from 4 to 2 when image size increase, thus decreasing task granularity as the global work size increase. This result is unexpected and understanding why this happen could be interesting. The NVIDIA GPU favors processing square tiles: *x_component_number* = 4, *y_component_number* = 4, *vector_length* = 4, *temporary_size* = 2 and *synthesize_loads* = false. Once more, *vector_recompute* has no meaning.

3) *Laplace Conclusion*: In this Subsection we have shown the interest of BOAST in optimizing a well-known algorithm across different architectures. Chosen optimization combinations are highly dependent on the target architecture. But BOAST simplifies and speeds up the identification of the correct combination at the cost of metaprogramming. In the next Subsections we will show that our methodology also applies to real applications.

B. Creating an Auto-Tuned Convolution Library for BigDFT using BOAST

In 2005, the EU FP6-STREP-NEST BigDFT [13] project funded a consortium of four European laboratories (L_Sim - CEA Grenoble, Basel University - Switzerland, Louvain-la-Neuve University - Belgium and Kiel University - Germany), with the aim of developing a novel approach for DFT calculations based on Daubechies wavelets. Rather than simply building a DFT code from scratch, the objective of this

three-years project was to test the potential benefit of a new formalism in the context of electronic structure calculations.

As a matter of fact, Daubechies wavelets exhibit a set of properties which make them ideal for a precise and optimized DFT approach. In particular, their systematicity allows to provide a reliable basis set for high-precision results, whereas their locality (both in real and reciprocal space) is highly desired to improve the efficiency and the flexibility of the processing. Indeed, a localized basis set allows to optimize the number of degrees of freedom for a required accuracy [13], which is highly desirable given the complexity and inhomogeneity of the systems under investigation nowadays.

Despite that the application is mainly written in Fortran (360 kLOC of Fortran), it currently also includes 70 kLOC of C languages, accounting for more than 50% of the code base. It is a parallel application based on the standards MPI [15] and OpenMP [16]. It also supports CUDA [17] and OpenCL [18]. In the recent years this code has been used for many scientific applications, and its development and user consortium is continuously growing. Massively parallel computations are routinely executed with the BigDFT code, either in homogeneous or hybrid architectures. In 2009, the French Bull-Fourier award was attributed for the implementation of the hybrid version of BigDFT [19].

1) *BigDFT*: In the Kohn-Sham (KS) formulation of DFT, the electrons are associated to wavefunctions (orbitals), which are represented by arrays of floating point numbers. In wavelets formalism, the operators are written via convolutions with short, separable filters. The detailed description of how these operations are defined is beyond the scope of this paper and can be found in the BigDFT reference paper [13]. Convolutions are basic operations of lots of scientific application codes, like for example finite differences approaches, which are universally used in computational physics.

The CPU convolutions of BigDFT have thus been thoroughly optimized. In a recent paper [20], the optimization of the CPU convolutions of BigDFT has been extensively considered. One example of a specific convolution, called MagicFilter [21], can be seen in Listing 16. It applies a filter *filt* to the data set *in* and then stores the result in the data set *out* with a transposition [22].

```

1 double filt[16] = {F0, F1, ... , F15};
2 void magicfilter(int n, int ndat, double* in, double* out){
3     double temp;
4     int m;
5     for( j = 0; j < ndat; j++) {
6         for( i = 0; i < n; i++) {
7             temp = 0;
8             for( k = 0; k < 16; k++) {
9                 m = (i-7+k)%n;
10                temp += in[m + j*n] * filt[k];
11            }
12            out [j + i*ndat] = temp ;
13        }
14    }
15 }

```

Listing 16. MagicFilter.

As we can see, there are three nested loops working on arrays whose sizes vary. Various optimizations can be applied to this treatment and may focus on the loop structure, as well as on the size of the data.

2) *A Generic Convolution Library*: The number of convolution kernels needed in BigDFT has been continuously growing in the recent years. Various boundary conditions and functionalities have been added, making the BigDFT more and more powerful in terms of scientific applications. However, the cost of maintenance and development of the convolutions is always a delicate point to be considered while including a new functionality. The convolution patterns are usually rather similar, leading to code duplication and difficulties in code maintainability.

Therefore it appears very interesting to benefit from an automatic tool to drive the implementation and the generation of new convolutions. This would lead to an optimized code, adapted to different computing platforms, that is optimally factorized. In addition to this point, the help of such code generator is also important to build *new science*: the cost of implementing new convolutions would become so little that other functionalities (for example, the generalization of the BigDFT convolutions to Neumann boundary conditions or the usage of wavelet-on-the-interval basis) can be added with limited manpower.

For these reasons, a convolution library has been engineered with the help of BOAST. The detailed API of the library is beyond the scope of this article. The spirit is similar to BLAS-LAPACK API, where low-level operations are scheduled and called from high-level operations. The basic blocks will be comprised of unidimensional wavelet transforms and convolutions applied to multidimensional arrays. Combining those blocks will yield multidimensional transforms.

Each of the building blocks of these convolution libraries is optimally tuned by BOAST by choosing the sources providing the optimal kernel for the chosen computing platform. The sources of these kernels are then collected and compiled to meet the API of the library. A library written in this way might have an impact going largely beyond the community of BigDFT developers.

The interest in having a robust and optimally tuned library in this scientific field is therefore evident. Techniques are under investigation to also provide end users with fine-tuned binaries rather than the source codes, such that more aggressive inter-procedural optimization can be performed. Indeed, BOAST finds the optimal source code for a given kernel and compiler configuration but one could imagine using binary optimizer to the compiled binaries. Those optimizers could be coupled to BOAST and the output of the process would be the final binary rather than the source code. This experience therefore would be a first step toward the release of a tunable optimized convolution library oriented to computational physics communities.

3) *Performance Report*: Several kernels have already been implemented in BOAST for the convolution library. Figure 5 shows the performance of the wavelet transform operation as a factor of the unrolling length of the outer loop, the language used to implement it as well as the activation or not of the OpenMP parallelization. Results are given as a speedup compared to the sequential hand tuned implementation that can be found in BigDFT. These tests were run on the Intel Xeon X5550 that was used to hand optimize the code.

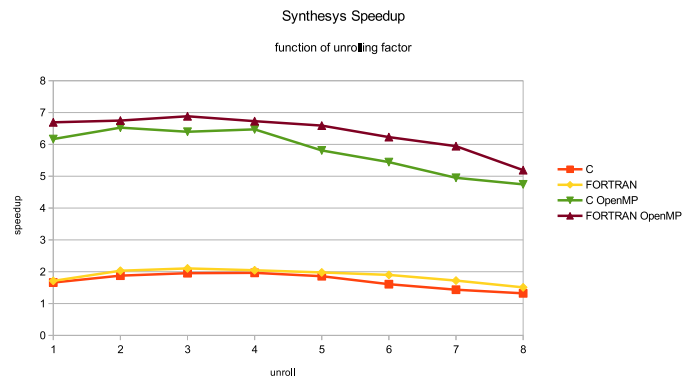


Fig. 5. Impact of Unrolling, Language Used and OpenMP on a Wavelet Transform Code.

We can see that this function is better optimized using Fortran and small unrolling factors. In the hand optimized version the unrolling factor was chosen much too high (a factor of 12 was used). This factor might have been optimal at the time the procedure was optimized (compiler version changed in the meantime) but since then the environment changed. Other optimizations have also been incorporated in the BOAST sources, like the systematic inner loop unrolling, and those could also help increase performance while limiting the interest of the outer loop unrolling.

Nonetheless what is interesting from the physicist point of view is that the generated source will give its better performance on a whole range of architectures/compiler combinations than that of the hand tuned code.

C. Porting SPECFEM3D to OpenCL using BOAST

In this last subsection, we study a free seismic wave propagation simulator, SPECFEM3D, in its GLOBE flavor¹. SPECFEM3D simulates seismic wave propagation at the local or regional scale based upon spectral-element method (SEM), with very good accuracy and convergence properties. It is a reference application for supercomputer benchmarking, thanks to its good scaling capabilities.

1) *SPECFEM3D*: When we started to work on the project (version v2.1 of July 2013) it supported graphics card GPU acceleration through NVidia CUDA. This GPU support came in addition to the MPI support implemented to enable multi-CPU parallel computing. Most of SPECFEM3D code-base is written in Fortran 2003 and only the GPU related parts are written in C. The split between CPU and GPU code was done at a rather fine grain, as the application counted more than 40 GPU kernels. Some of them were quite simple (*e.g.*, performing a few vector operations, but at massively parallel scale), while at the other side of the spectrum, some complex kernels took more than 80 parameters and perform very specific physical transformations.

Because of the complexity of wave propagation (and of the application architecture), it is hard to impossible to define unit tests. Hence, the application is validated by the accuracy of

¹Specfem3d Globe — CIG <http://www.geodynamics.org/cig/software/specfem3d-globe>

the results it produces, that is, the accuracy of seismograms of the simulated earthquakes, in comparison with the actual ones. Non-regression testing (after new developments) is also based on these seismograms, with measurements of the relative error between two identical simulations.

As part of the Mont-Blanc project, we had to port SPECFEM3D to OpenCL, so that it could be used to benchmark Mont-Blanc HPC platforms.

2) Porting to OpenCL:

a) *Porting kernels to BOAST:* Nvidia CUDA and OpenCL are based on the same programming model: a massively parallel accelerator running in disjoint, non addressable memory environment. Thanks to that proximity, we have been able to carry out most of the porting task with only a limited knowledge about SPECFEM3D internal physics.

This lack of SPECFEM3D internal knowledge led us to be particularly careful to the path we undertook for the port, as we would have been unable to understand how and why the application was not operating properly, if it was to fail.

Hence, our first milestone in the port was the translation of SPECFEM3D's CUDA kernels into BOAST EDSL. This way, we could ask BOAST framework to generate a CUDA version of the kernels, plug them back into SPECFEM3D and get (after fixing compile-time errors—prototypes and naming mistakes mainly) a first set of SPECFEM3D seismograms.

As we had expected, the seismograms were erroneous. But with the help of shell scripts and BOAST framework ability to store and provide the kernels' original source code, we built a set of SPECFEM3D binaries including only *one* BOAST-generated kernel, with all others reference-kernels. Running and validating all these binaries enabled use to pinpoint the misbehaving kernels. We finished the debugging with a side-by-side comparison that highlighted the coding mistakes.

b) *Porting run-time to OpenCL:* The second part of the port consisted in the translation of the CPU-side of the application, from CUDA API to OpenCL API. Most of the functions of the interfaces are very similar, with only naming-convention and data-structure distinctions. Hence, it was clear that automatic rewriting tools (namely `sed` regex and `emacs-lisp` functions) could be useful. To give an idea of the cost of a *manual* rewriting, we can count (in # of OpenCL API function calls): 70 kernel “function calls”, 790 arguments to set, 230 memory transfers, 160 buffer creations, and 270 releases.

Once the transformations were applied, compilation errors fixed, and OpenCL unsuccessful function calls solved, the application managed to complete its execution and generate a first set of seismograms. And again, as expected and feared, these seismograms were not valid as their shape was completely different from the reference ones.

As we had already validated BOAST-generated kernels (and trusted CUDA and OpenCL versions to be semantically identical), we knew that the bugs were now in the usage of the run-time, and we had to find a way to understand where SPECFEM3D's CUDA version of the code diverged from its OpenCL counterpart. To help us in that purpose, we had a strong assumption: both versions of the code were supposed to perform exactly the same

operations, with the same “logical” parameters (the APIs have *implementation* differences, for instance OpenCL has two memory transfer functions, `clEnqueueReadBuffer`, `clEnqueueWriteBuffer`, whereas CUDA has only one, with a direction parameter `cudaMemcpy(..., dir)`, but above that, it is the same functionality).

Hence, our idea for locating the execution problems was to make sure that both execution actually did the same thing. As the OpenCL results were invalid, we knew the executions would diverge at one or several points.

c) *Debugging OpenCL Execution: GPUTrace:* With the help of GPUTrace (described in Section IV-C), we could confront CUDA and OpenCL execution traces with a graphical diff tool, and spot the different porting mistakes: some parameters reversed, offsets incorrectly applied, *etc....*

One last problem remained, clearly highlighted by the seismograms not matching perfectly (they had a similar shape, but with a reduced intensity). We added more verbosity to GPUTrace output: first the initial bits of the GPU memory buffers, then their full content. The drift was visible in the trace, but it was nonetheless unclear where it started. We finally got it after hours of code review of BOAST kernels and OpenCL code. One kernel was *three*-dimensional, whereas the others were two-dimensional. But for all of them, only two dimensions were passed, and one was missing.

d) *Evaluation:* Our OpenCL/BOAST port of SPECFEM3D is now merged in SPECFEM3D's development tree and under test and extension by different research teams. On a platform with two K40x GPUs and N Intel Xeon processors, we measured similar results between the original CUDA version and our BOAST-CUDA version. With the same set of optimization flags, BOAST CUDA and OpenCL version reported similar execution time spans. The best execution speed was achieved with CUDA version though (25% higher than OpenCL), as one optimization parameter (`CUDA_LAUNCH_BOUND`) cannot be passed to OpenCL run-time, as of version 1.1. This parameter, in addition to specifying the work group size (which can be done in OpenCL), also constrains the number of work group that must run in parallel on a multiprocessor. This value is set to 7 in CUDA. This means that the compiler must be very conservative on register usage in order to allow this parallelism which allows better overlapping of communications end computations. This functionality is not supported in OpenCL.

By refactoring GPU kernel code in BOAST EDSL, the size of kernel code shrank by a factor of 1.8 (from 7500 to less than 4000 LOC, partly because of code duplication, also removing manually unrolled loops). This is beneficial for SPECFEM3D as it improves the readability and maintainability of its source code.

We have also been able to enhance SPECFEM3D's non-regression test-suite by adding per-kernel non-regression tests. This was done with the help of GPUTrace, that we used to capture all the input parameters of a particular *valid* kernel execution, as well as the output values. Then, during the non-regression testing, BOAST framework loads these buffer files, allocates GPU memory and initializes it through CUDA or

OpenCL run-time, and triggers the kernel execution. A comparison of the output values (for instance against a maximal error level) validates the non-regression.

In the same mindset, we provided SPECFEM3D test-suite with kernel performance evaluation mechanisms. These tests will help developers to try new optimizations in kernel code and measure their impact, without executing the whole application.

VI. RELATED WORK

Code generation and auto-tuning techniques are not new. Nonetheless, with recent developments in hardware, and the HPC landscape being as diverse as it is now, there is a renewed interest in the field. This related work section is split in three parts, focusing first on auto-tuning frameworks. Tools that provide a DSL to describe computing kernels will then be presented. Last, optimization space pruners and their ties to auto-tuning will be introduced.

A. Application Auto-Tuning

The most convenient way to obtain an application that can be auto-tuned on a given platform is to base this application on a widely used computing library. BLAS [23] and LAPACK [24] are such libraries. These libraries are either hand tuned for selected platforms or have auto-tuned implementations. Atlas [25] is an auto-tuned implementation of BLAS/LAPACK. ATLAS authors defined the *Automated Empirical Optimization of Software* methodology that we implemented with BOAST. Their kernel generation is performed using macro-functions in C.

Nonetheless many application formalisms cannot be reduced to standardized library or border on what could be considered edge cases for those libraries and not as optimized as more general cases. Orio [26] is an auto-tuning framework that has an approach close of that of BOAST but they are based on an annotated DSL describing loop transformations rather than a more generic EDSL. Halide [27] is an auto-tuning framework dedicated to image processing. It can also be used to describe other operations on memory buffers. Those two frameworks propose automated search space exploration to find the best version of a kernel. LGen [28] is a compiler that generates linear algebra programs for small fixed size problems. Knowing the problem size, it fully unrolls and vectorizes loops, yielding better performance than state of the art generic implementations.

B. Kernel Description DSL

The idea to describe computing kernels using a DSL has been already explored. SPIRAL [29] is a decade old generation framework for signal processing. It uses a proprietary DSL called SPL (Signal Processing Language), to describe a DSP algorithm. This DSL is then transformed into efficient programs in high level languages like C or Fortran. POET [30] also uses a DSL to describe custom code transformations, like loop unrolling, loop blocking and loop interchange. Those transformations can be parametrized in order to tune the

application. Orio [26] can be compared to POET as it aims at describing possible code transformations using a DSL. All those approaches are very different from ours as they put the emphasis on compilation techniques whereas BOAST relies on the user to express the different optimizations.

Halide [27] is closer in some ways to our approach as it uses an embedded C++ DSL to describe image processing algorithm. This DSL allows decoupling the algorithm description from its scheduling. Each pixel in the resulting image has a completely defined dependency tree with regard to pixels in the input image (and intermediary results). During generation, depending on memory and computing cost, some values are recomputed rather than fetched from memory. We used Halide to implement the *magicfilter* of BigDFT, but unfortunately results were 4 to 5 times slower than the one we obtained with BOAST. We speculate that the three-dimensional nature of our convolutions, as well as the filter length, can be considered edge cases in Halide and quite far from the intended target.

BEAST [31] is an autotuning framework that uses macro-processing in source files. Inside the macros, BEAST offers a DSL that can be used to specify values derived from a set of iterators. Those iterators explore the search space and can be constrained. Great care has been used to ensure that the iterators enumeration as well as the pruning are as efficient as possible.

Heterogeneous Programming Library [32] (HPL) is maybe the most similar to BOAST. It is an autotuning kernel library in C++ that targets OpenCL. It uses an EDSL in C++ to describe generic computing kernels. Like when using BOAST, optimization are left to the users and written directly using the EDSL. The syntax is very close to real C++ with the classical control structures having a trailing underscore. But whereas BOAST re-interprets the same code while changing the environment, in HPL each expression is stored in a tree and the tree is reevaluated in a different environment.

C. Optimization Space Pruners

Once auto-tuning techniques are used, the parameter space explodes and the systematic sampling rapidly becomes impossible to achieve. The cost of finding the optimal kernel parameters and environment parameters (compiler flags, used language, ...) is rapidly prohibitive. Dedicated frameworks have been developed to address this problem. Adaptive Sampling Kit (ASK) [33] is one such tool. It reduces the number of samples needed by creating a model of the performance and by minimizing the number of samples needed to find the parameters of this model. Several models and sampling techniques are implemented.

Collective Mind [34] proposes similar techniques to solve the problem but also stores the results and complete experimental setups in databases for future reference and reproducibility. This database approach also enables easy parallelization of the experimental process. Collective Mind also proposes collections of flags for many available compilers/versions easing the exploration process.

VII. CONCLUSION AND FUTURE WORKS

Application portability is an important issue that should be solved efficiently, especially given the large number of different processors now available for today's supercomputers. The work needed to get performance portability is a tedious task, even for experienced programmers. The availability of semi-automatic tools is therefore mandatory for the development of large simulation applications. Computing kernels' identification and optimization has to be carefully performed as they usually consume most of the computing resources.

In this paper we presented the BOAST infrastructure (DLS and runtime) that aims at describing kernels in a high level language and allows the comparison of the performance of different versions of the code in a simple and seamless way. We described its application to three use cases from the Mont-Blanc project. Results are encouraging as BOAST proved to be a powerful and flexible tool that allowed gains in performance compared to hand tuned codes. Performance portability of those codes is also improved.

Future development will focus on three main goals. First we find interesting to interface it with binary analysis tools like MAQAO [35] in order to build a feedback loop to guide optimization. On the autotuning side, interfacing with search space modellers/pruners in order to optimize the search of the optimal version of a kernel will allow us to gain some time in the optimization process. Finally, work should also continue on improving the support for vector codes. For instance producing a collection of small to medium useful vector patterns (transposition, for instance) in BOAST could really help users develop vectorized version of their algorithm. Other computing kernels and applications will also be ported on various architectures.

ACKNOWLEDGMENTS

The research leading to these results has received funding from the European Community's Seventh Framework Programme [FP7/2007-2013] under grant agreements N° 288777 and 610402.

REFERENCES

- [1] Top500.Org, "Top500," <http://www.top500.org>.
- [2] C. Adeniyi-Jones, "Optimal Compute on ARM Mali GPUs," http://www.cs.bris.ac.uk/home/simonm/montblanc/OpenCL_on_Mali.pdf.
- [3] J. Cronsioe, B. Videau, and V. Marangozova-Martin, "BOAST: Bringing Optimization through Automatic Source-to-Source Transformations," in *Embedded Multicore SoCs (MCSoc)*, 2013 IEEE 7th International Symposium on. IEEE, 2013, pp. 129–134.
- [4] M. E. Wolf and M. S. Lam, "A Loop Transformation Theory and an Algorithm to Maximize Parallelism," *IEEE Transactions on Parallel and Distributed Systems*, vol. 2, no. 4, pp. 452–471, 1991.
- [5] M. Wolfe, "Loops Skewing: The Wavefront Method Revisited," *International Journal of Parallel Programming*, vol. 15, no. 4, pp. 279–293, 1986.
- [6] J. R. Allen and K. Kennedy, "Automatic Loop Interchange," in *ACM SIGPLAN Notices*, vol. 19, no. 6. ACM, 1984, pp. 233–246.
- [7] C. L. Lawson, R. J. Hanson, D. R. Kincaid, and F. T. Krogh, "Basic Linear Algebra Subprograms for FORTRAN Usage," *ACM Transactions on Mathematical Software (TOMS)*, vol. 5, no. 3, pp. 308–323, 1979.
- [8] M. I. Soliman, "Performance Evaluation of Multi-Core Intel Xeon Processors on Basic Linear Algebra Subprograms," *Parallel Processing Letters*, vol. 19, no. 01, pp. 159–174, 2009.
- [9] D. Ye, A. Titov, V. Kindratenko, I. Ufimtsev, and T. Martinez, "Porting Optimized GPU Kernels to a Multi-Core CPU," in *Symposium on Application Accelerators in High-Performance Computing*, 2012.
- [10] P. Hudak, "Building Domain-specific Embedded Languages," *ACM Computing Surveys*, vol. 28, no. 4es, pp. 1–6, Dec 1996.
- [11] Y. Matsumoto and K. Ishituka, *Ruby Programming Language*. Addison Wesley Publishing Company, 2002.
- [12] P. Mucci, S. Browne, C. Deane, and G. Ho, "PAPI: A Portable Interface to Hardware Performance Counters," in *Proc. Dept. of Defense HPCMP Users Group Conference*. Citeseer, 1999, pp. 7–10.
- [13] L. Genovese, A. Neelov, S. Goedecker, T. Deutsch, S. A. Ghasemi, A. Willand, D. Caliste, O. Zilberberg, M. Rayson, A. Bergman *et al.*, "Daubechies Wavelets as a Basis Set for Density Functional Pseudopotential Calculations," *The Journal of chemical physics*, vol. 129, no. 1, 2008.
- [14] D. Komatitsch, "Fluid-Solid Coupling on a Cluster of GPU Graphics Cards for Seismic Wave Propagation," *Comptes Rendus Mécanique*, vol. 339, no. 2, pp. 125–135, 2011.
- [15] "The Message Passing Interface (MPI) Standard," 2012, <http://www.mcs.anl.gov/research/projects/mpl/>.
- [16] L. Dagum and R. Menon, "OpenMP: an Industry Standard API for Shared-Memory Programming," *Computational Science & Engineering, IEEE*, vol. 5, no. 1, pp. 46–55, Jan-Mar 1998.
- [17] NVIDIA, "NVIDIA Compute Unified Device Architecture," 2011, http://www.nvidia.com/object/cuda_home_new.html.
- [18] Khronos OpenCL consortium, "OpenCL: Open Computing Language," <http://www.khronos.org/opencl/>.
- [19] L. Genovese, M. Ospici, T. Deutsch, J.-F. Méhaut, A. Neelov, and S. Goedecker, "Density Functional Theory Calculation on Many-Cores Hybrid Central Processing Unit-Graphic Processing Unit Architectures," *The Journal of chemical physics*, vol. 131, no. 3, 2009.
- [20] B. Videau, V. Marangozova-Martin, L. Genovese, and T. Deutsch, "Optimizing 3D Convolutions for Wavelet Transforms on CPUs with SSE units and GPUs," in *Euro-Par 2013 Parallel Processing*. Springer, 2013, pp. 826–837.
- [21] L. Genovese, B. Videau, M. Ospici, T. Deutsch, S. Goedecker, and J.-F. Méhaut, "Daubechies Wavelets for High Performance Electronic Structure Calculations: the BigDFT Project." in *Compte-Rendu de l'Académie des Sciences, Calcul Intensif*. Académie des Sciences, 2010.
- [22] S. Goedecker, "Rotating a Three-Dimensional Array in an Optimal Position for Vector Processing: Case Study for a Three-Dimensional Fast Fourier Transform," *Computer Physics Communications*, vol. 76, no. 3, pp. 294–300, 1993.
- [23] J. J. Dongarra, J. Du Croz, S. Hammarling, and I. S. Duff, "A Set of Level 3 Basic Linear Algebra Subprograms," *ACM Transactions on Mathematical Software (TOMS)*, vol. 16, no. 1, pp. 1–17, 1990.
- [24] E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen, *LAPACK Users' Guide*, 3rd ed. Philadelphia, PA: Society for Industrial and Applied Mathematics, 1999.
- [25] R. C. Whaley and A. Petitet, "Minimizing Development and Maintenance Costs in Supporting Persistently Optimized BLAS," *Software: Practice and Experience*, vol. 35, no. 2, pp. 101–121, February 2005.
- [26] A. Hartono, B. Norris, and P. Sadayappan, "Annotation-Based Empirical Performance Tuning Using Orio," in *Proceedings of the 23rd IEEE International Parallel & Distributed Processing Symposium*, Rome, Italy, 2009, also available as Preprint ANL/MCS-P1556-1008. [Online]. Available: <http://www.mcs.anl.gov/uploads/xels/papers/P1556.pdf>
- [27] J. Ragan-Kelley, C. Barnes, A. Adams, S. Paris, F. Durand, and S. Amarasinghe, "Halide: A Language and Compiler for Optimizing Parallelism, Locality, and Recomputation in Image Processing Pipelines," *ACM SIGPLAN Notices*, vol. 48, no. 6, pp. 519–530, 2013.
- [28] D. G. Spampinato and M. Püschel, "A Basic Linear Algebra Compiler," in *Proceedings of Annual IEEE/ACM International Symposium on Code Generation and Optimization*. ACM, 2014, p. 23.
- [29] M. Püschel, J. M. Moura, B. Singer, J. Xiong, J. Johnson, D. Padua, M. Veloso, and R. W. Johnson, "SPIRAL: A Generator for Platform-Adapted Libraries of Signal Processing Algorithms," *International Journal of High Performance Computing Applications*, vol. 18, no. 1, pp. 21–45, 2004.
- [30] Q. Yi, K. Seymour, H. You, R. Vuduc, and D. Quinlan, "POET: Parameterized Optimizations for Empirical Tuning," in *Parallel and Distributed Processing Symposium, 2007. IPDPS 2007. IEEE International*. IEEE, 2007, pp. 1–8.
- [31] H. Anzt, B. Haugen, J. Kurzak, P. Luszczek, and J. Dongarra, "Experiences in autotuning matrix multiplication for energy minimization on GPUs," *Concurrency and Computation: Practice*

- and Experience*, vol. 27, no. 17, pp. 5096–5113, 2015, cpe.3516.
[Online]. Available: <http://dx.doi.org/10.1002/cpe.3516>
- [32] M. Viñas, Z. Bozkus, and B. B. Fraguera, “Exploiting Heterogeneous Parallelism With the Heterogeneous Programming Library,” *Journal of Parallel and Distributed Computing*, vol. 73, no. 12, pp. 1627–1638, 2013.
- [33] P. de Oliveira Castro, E. Petit, A. Farjallah, and W. Jalby, “Adaptive Sampling for Performance Characterization of Application Kernels,” *Concurrency and Computation: Practice and Experience*, vol. 25, no. 17, pp. 2345–2362, 2013.
- [34] G. Fursin, R. Miceli, A. Lokhmotov, M. Gerndt, M. Baboulin, D. Malony, Allen, Z. Chamski, D. Novillo, and D. D. Vento, “Collective Mind: Towards Practical and Collaborative Auto-Tuning,” *Scientific Programming*, vol. 22, no. 4, pp. 309–329, July 2014. [Online]. Available: <https://hal.inria.fr/hal-01054763>
- [35] L. Djoudi, D. Barthou, P. Carribault, C. Lemuét, J.-T. Acquaviva, and W. Jalby, “Exploring Application Performance: a New Tool for a Static/Dynamic Approach,” in *Proceedings of the 6th LACSI Symposium*, 2005.