

ProtoX: A First Look

Het Mankad*, Sanil Rao*, Phillip Colella†, Brian Van Straalen†, Franz Franchetti*

**Electrical and Computer Engineering Department Carnegie Mellon University, Pittsburgh, PA, USA*

†*Lawrence Berkeley National Laboratory, Berkeley, CA, USA*

*{hmankad, sanilr, franzf}@andrew.cmu.edu, †{bvstraalen, pcoella}@lbl.gov

Abstract—We present a first look at ProtoX, a code generation framework for stencil and pointwise operations that occur frequently in the numerical solution of partial differential equations. ProtoX has Proto as its library frontend and SPIRAL as the backend. Proto is a C++ based domain specific library which optimizes the algorithms used to compute the numerical solution of partial differential equations. Meanwhile, SPIRAL is a code generation system that focuses on generating highly optimized target code. Although the current design layout of Proto and its high level of abstractions provide a user friendly set up, there is still a room for improving its performance by applying various techniques either at a compiler level or at an algorithmic level. Hence, in this paper we propose adding SPIRAL as the library backend for Proto enabling abstraction fusion, which is usually difficult to perform by any compiler. We demonstrate the construction of ProtoX by considering the 2D Poisson equation as a model problem from Proto. We provide the final generated code for CPU, Multi-core CPU, and GPU as well as some performance numbers for CPU.

Index Terms—partial differential equations, stencil operations, code generation, Proto, SPIRAL

I. INTRODUCTION

There is a myriad of application areas in the field of scientific computing and engineering where numerical solutions to partial differential equations (PDEs) are required to be computed. Numerical methods like the finite difference method (FDM), finite element method (FEM), finite volume method (FVM) and multigrid method are used to approximate the solutions to these PDEs. The key components of these algorithms are the stencil and pointwise operations. Stencils are defined as a linear transformation,

$$S(x)_i = \sum_j \alpha_j x_{i+j}, \quad (1)$$

where $i, j \in \mathbb{Z}^D$ with D denoting the number of space dimensions, α denotes the weight and x is the multidimensional data array.

Typically these numerical algorithms are iterative in nature resulting in performing the stencil operations multiple times. Writing codes for these stencil based methods from scratch can be quite a cumbersome task for anyone. As such developers turn to libraries that provide stencil operations for them. One such library is Proto. It is a domain specific library written in C++ that provides a high level of abstraction for solving various PDEs using some of the aforementioned numerical methods.

Proto’s abstraction enables ease of programmability, but has drawbacks when it comes to performance. Many of Protos’ abstractions can be fused and optimized together,

resulting in better performance. However, abstraction fusion is something no compiler can easily perform. This results in additional burden on the library developers to manually introduce these optimizations. To enable abstraction fusion in Proto, we propose ProtoX, which is a C++ library based on Proto and runs a code generation system SPIRAL [1], [2] in the backend. The concept of using SPIRAL in the backend and a C/C++ based library in the front has shown positive results in the past [3], [4]. Some of the related works in the area of optimizing stencil computation with either automatic code generation or by optimizing data movement involved while performing stencil operation are discussed in the next section.

Contribution. Some of the main contributions of this work are:

- 1) A proof of concept of hooking SPIRAL code generation as backend into Proto is presented.
- 2) A Proto example is considered as a SPIRAL specification and a full program optimization with merged kernels across C++ functions are discussed for that example.
- 3) This work is a first look of ProtoX. Hence, initial CPU results are presented.

This work is organized as follows. Section II briefly discusses some of the related work in this area of research. Section III provides the necessary background material to understand this work. Here we discuss the design layout of both Proto and SPIRAL. Section IV introduces the structure of ProtoX with the help of the 2D Poisson problem and finally section V and VI provides with the concluding remarks, future work and acknowledgement.

II. RELATED WORK

Due to the frequent use of stencil operations in various scientific applications, there has been a significant amount of work done over the past couple of years in the area of generating optimized code for such computations. Some work focuses on developing stencil codes for a particular application, for example, Halide [5] is a C++ based DSL focusing on all image processing pipelines. Meanwhile, some focus on platform portability like LIFT [6] and SBLOCK [7]. There is also some work done towards improving performance of stencil codes by introducing new ways to optimize the data movement at a compiler level like [8]. PATUS [9] is an example of code generation framework for stencils that provide multi-core CPU and GPU code for such stencil operations. The work done by Holewinski et. al. in [10] showcases their efforts in the area of highly optimized code generation for

stencils but mainly targeting GPU accelerators. In [11] Li et. al. work on large scale stencil computation generating code for both spatial and temporal computation on heterogeneous many-core processors. Some of the more recent work in the area of code generation for stencil computation can be seen in ExaStencil [12], YASK [13] and Code Generation for In Place Stencils [14].

III. BACKGROUND

In order to understand the design of ProtoX we will first need to understand the design structure of Proto as well as SPIRAL. In this section we will briefly provide an overview of Proto and SPIRAL.

Proto. Proto is a C++ library designed to provide an intuitive interface that optimizes the designing and scheduling of an algorithm aimed at solving various PDEs numerically. In order to approximate the solution of a PDE, the domain is usually divided into either structured or non-structured grid and the equations are discretized using methods like FDM, FEM or FVM. Currently, Proto takes into account a multidimensional rectangular, structured grid with periodic boundary condition for all model problems.

There are four main C++ classes that implement a representation of the spatial rectangular grids are discussed below.

- **Point:** It represents points in \mathbb{Z}^D . It is used to indicate the location as well as offsets in a grid in Proto.
- **Box:** It denotes the rectangular subsets of \mathbb{Z}^D that has elements from the **Point** class as it nodes. Figure 1 indicates a typical 2D computational domain used in Proto which is divided into several Boxes $B_j, j = 0 \dots, M-1, M \in \mathbb{N}$. Here, each box in the **Box** class is represented as $B_j = [\ell_j, \mathbf{h}_j] \subset \mathbb{Z}^2$, where $(\ell_j \in \mathbb{Z}^2)$ and $(\mathbf{h}_j \in \mathbb{Z}^2)$ are it's lowest and highest corner points.
- **BoxData** $\langle \mathbb{T}, \mathbb{C}, \mathbb{D}, \mathbb{E} \rangle$: It represents the multidimensional data arrays on each **Point** in a **Box**. It is a templated class in C++. For a given point i in B_j , the **BoxData** class is defined as a function,

$$f(i) : B_j \rightarrow \mathcal{T}(\mathbb{T}, \mathbb{C}, \mathbb{D}, \mathbb{E}).$$

Here \mathbb{T} is the class template holder for the data which can be either \mathbb{R}, \mathbb{C} or \mathbb{Z} . Meanwhile, the dimensions of the range space is given by \mathbb{C}, \mathbb{D} and \mathbb{E} . Hence, based on the values of \mathbb{C}, \mathbb{D} and \mathbb{E} one can identify the given data to be a scalar, vector, matrix or 3^{rd} order tensor.

- **Stencil** $\langle \mathbb{T} \rangle$: Another important class in Proto is the **Stencil** class. This is where stencils are defined as self-contained objects to implement the operations described in (1). Stencils are applied to the **BoxData** to update the values in there. They are accompanied by the **Shift** class to carry out the grid shifts that are required in a stencil operation.

In order to apply the stencils or any other function in a pointwise manner, Proto uses the **forall** function. It takes in as an argument the function which is to be applied in a pointwise manner and the corresponding **BoxData** used for it.




Fig. 1: A pictorial representation of the **Point**, **Box** and **BoxData** classes in Proto is shown.

Operation	Description
$(A_{m \times n} \circ B_{p \times q})$	Operator composition: AB if $n = p$
$(A_{m \times n} \oplus B_{p \times q})$	Direct sum operation: $\begin{bmatrix} A & \\ & B \end{bmatrix}$
$(A_{m \times n} \otimes B_{p \times q})$	Tensor product: $\begin{bmatrix} A_{00}B & \cdots & A_{0,n-1}B \\ \vdots & \ddots & \vdots \\ A_{m-1,0}B & \cdots & A_{m-1,n-1}B \end{bmatrix}$
$\begin{bmatrix} - \\ - \end{bmatrix}$	Vertical stacking $\begin{bmatrix} A_{m \times n} \\ B_{p \times q} \end{bmatrix}$
$\begin{bmatrix} - \\ - \end{bmatrix}_{i=0}^k$	Iterative vertical stack
I_n	Identity matrix $I_{n \times n}$
$pw_{x \rightarrow f(x)}$	Pointwise operation
(a, b, c)	Row vector with three entries a, b and c

TABLE I: List of some operations used in SPIRAL for this work is shown here. The matrices $A_{m \times n}$ and $B_{p \times q}$ are considered as operators with $A : \mathbb{R}^n \rightarrow \mathbb{R}^m$ and $B : \mathbb{R}^q \rightarrow \mathbb{R}^p$.

SPIRAL. SPIRAL is a GAP based code generation system that was initially developed to automatically generate optimized C/C++ programs for linear transforms like the discrete Fourier Transform (DFT), discrete sine/cosine transform and many more [1], [2]. SPIRAL uses the Signal Processing Language (SPL) [15] to develop algorithms for these signal processing transforms. SPL is a declarative mathematical language that expresses linear transforms as matrix-vector product. Here, the matrix is considered to be an operator with one input and one output vector.

Operator Language (OL) is a generalization of SPL. It is the result of the ongoing efforts to expand the scope of SPIRAL beyond signal processing transforms [16]. As described in [16], it is a domain specific language (DSL) which is also a mathematical declarative language like the SPL. It is written in terms of different operators that describe the data flow and data layout of a given transform or a computational kernel of interest. This helps in manipulating the algorithm in such a way that enables introducing multiple optimization techniques at a high level of abstraction. Table I provides a list of some of the operations that will be useful in describing the data flow of the 2D Poisson problem discussed in this work.

After expressing the given problem specification in terms of OL, the next step is to rewrite it in terms of Σ -OL. It is an intermediate step in SPIRAL and an extension of Σ -SPL [2] which helps to incorporate loops into the problem specification

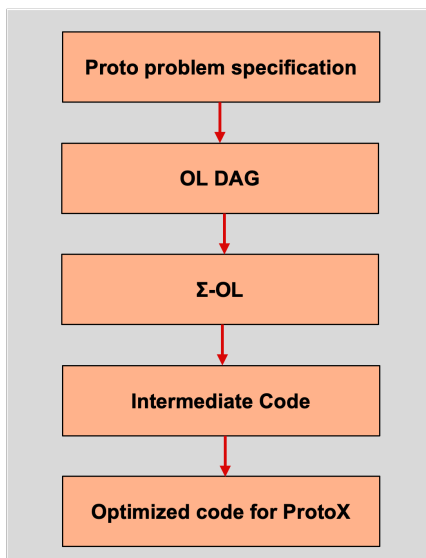


Fig. 2: ProtoX design layout starting with a problem specification from Proto to the final optimized code generated using SPIRAL

unlike OL. This enables loop merging and as a result allows abstraction fusion.

IV. PROTOX

In this section we will describe the structure of ProtoX. The idea is to interpret Proto as a Domain Specific Language (DSL) with the help of SPIRAL. This is done by first interpreting an example from Proto as a mathematical specification and then map the Proto program specification to an OL expression. It is then broken down into a Σ -OL expression which introduces loop fusion. This will help generate a highly optimized C++ code. Here we will explain these ideas with respect to the 2D Poisson equation. The design layout for ProtoX is shown in Fig 2

2D Poisson equation. The Poisson equation is given as,

$$\Delta\phi(x, y) = \rho(x, y), \quad x, y \in \Omega := [0, 1] \times [0, 1], \quad (2)$$

where ρ is a given function and ϕ is what we are solving for. Δ is the Laplace operator. We use the 5-pt stencil as a second order finite difference approximation of the Laplacian. The Jacobi iteration method is implemented in Proto to find the solution of (2). Let h denote the mesh spacing for the discretized domain Ω , then the Jacobi formula for a single iteration n is given as

$$\phi_i^n := \phi_i^{n-1} + \lambda(S(\phi)_i^{n-1} - \rho), \quad (3)$$

where $\lambda = h^2/4D$ and $S(\phi)_i = \sum_{s \in \mathbb{Z}^2} a_s \phi_{i+s}$ is the 5-point stencil applied to the input data ϕ_i^{n-1} .

Algorithm. In Proto, the domain space Ω is divided into several boxes and (2) is solved for each point in the box and then information is exchanged between the boxes to update




Fig. 3: DAG for the Poisson problem in Proto

the corresponding box data. The three main steps involved in this algorithm are

- 1) Applying the 5-pt Laplacian stencil to the given initial guess for ϕ .
- 2) Approximate the new value for ϕ using the Jacobi iteration method shown in (3).
- 3) Check the latest approximation of ϕ against the convergence criterion. Max norm is used in Proto to check for convergence.

We would like to note that in Proto each of the steps described in the algorithm above correspond to separate C++ function calls. Figure 4 provides the code sample that is used in Proto to solve (2). We can observe that the code reads in the same way as discussed in the algorithm above indicating an ease of use from a user perspective. However, this leads to intermediate data holders generating too much memory traffic. Hence, some optimization techniques need to be applied either at a compiler level or at an algorithmic level to overcome this issue.

```

1 // Defining the 5-pt Laplacian stencil
2 Stencil<double> laplace = Stencil<double>::Laplacian();
3 for (int iter = 0; iter < maxiter; iter++){
4   ...
5   // Solve for all boxes with each Box of size 64x64
6   for (auto dit=phi.begin(); dit != dit.end(); ++dit){
7     BoxData<double>& phiPatch = phi[*dit];
8     BoxData<double>& rhoPatch = rho[*dit];
9
10    // Compute the Laplacian
11    BoxData<double> temp = laplace(phiPatch, wgt);
12
13    // Jacobi iteration
14    forallInPlace(jacobiUpdate, phiPatch, temp,
15                 rhoPatch, lambda);
16  }
17  // Computing || ||_{inf}
18  double resmax=computeMaxResidualAcrossProcs(phi,
19  rho, dx);
20 }

```

Fig. 4: Sample Proto code for the 2D Poisson problem

Implementation in SPIRAL. Figure 3 provides the data flow for the algorithm discussed above. The first step in generating an optimized code using SPIRAL is to understand the data flow of the problem specification. Consequently, this data flow needs to be translated into OL with the help of the different operations shown in table I.

We will consider the case where the size of each `Box` in Proto for this problem specification is $n \times n$ with $m \times m$ total

elements in the box including ghost cells. The resulting OL expression corresponding to that data flow is shown below.

$$\text{Poisson}_{n,m,t}^{\ell,w,a} \rightarrow \left[\begin{array}{c} \text{Jacobi}_{n,m,w,l} \\ \|\cdot\|_{\infty}^{n,m,a} \end{array} \right] \circ \left(\left[\begin{array}{c} I_{n^2} \\ \text{Laplace} \end{array} \right] \oplus I_{n^2} \right) \circ X, \quad (4)$$

$$\text{Laplace}_{n,m,t} \rightarrow \text{Scatter}_{n^2 \times m^2} \circ [\text{Filt}(t)]_{i=0}^{m^2}, \quad (5)$$

$$\text{Jacobi}_{n,m,w,l} \rightarrow (1, w, -\lambda) \otimes I_{n^2}, \quad (6)$$

$$\|\cdot\|_{\infty}^{n,m,a} \rightarrow \text{Max} \circ pw_{x \rightarrow |x|}^{n \times n} \circ (0, 1/(a^2), -1) \otimes I_{n^2}. \quad (7)$$

Here X denotes the linearized input vector. For this problem specification it contains the initial data for ϕ and ρ with ϕ being of size $m \times m$ or m^2 and ρ is of size n^2 . Hence, X is of size $n^2 + m^2$. t denotes the filter taps corresponding to the 5-pt stencil. Its a 3×3 matrix with entries as $[0, 1, 0]$, $[1, -4, 1]$ and $[0, 1, 0]$ for the first, second and third row respectively. This matrix is flattened out and is iteratively applied to the input vector with the proper shifts. $\text{Scatter}_{n^2 \times m^2}$ denotes the *scatter matrix* [17] in SPIRAL. ℓ, w and a are scalar parameters used for the Jacobi iteration and the Laplacian.

The next step is to rewrite the OL breakdown rules shown in (4)-(7), in terms of Σ -OL. This is where loop merging is introduced which helps in fusing different Proto abstractions into one single loop resulting in a considerable amount of performance gain. Here the computation is done per point in the data set rather than the entire patch which enables performing several computations simultaneously. A sample of the Σ -OL expression for this problem specification is shown below.

$$\left(\sum_{j=0}^{N-1} S_{r_j} \text{MaxNorm} G_{s_j} \right) \circ \left(\sum_{j=0}^{N-1} S_{p_j} \text{Jacobi} G_{q_j} \right) \circ \left(\sum_{j=0}^{N-1} S_{u_j} \text{Laplace} G_{t_j} \right). \quad (8)$$

Here S and G are the scatter and gather functions that are used in SPIRAL to read and write data [2]. The subscripts t_j, u_j, q_j, p_j, s_j and r_j are the functions that indicate the exact number of points including their location that are being used to gather or read for computing a particular kernel (like the Laplace, Jacobi and MaxNorm in this case) and the number of points that are being scattered or written after the computation on the kernel is done.

Results. Once the Σ -OL representation is achieved, an intermediate code is generated which in the end generates the final optimized C/C++ code for the given problem specification. Figures 5, 6 and 7 provide the sample code for the 2D Poisson problem that is generated using SPIRAL. The code is generated for the `Box` size of 64×64 with the overall domain size being 256×256 . We can see that in comparison to the original code in Proto (see Fig. 4) which has multiple abstractions for the different computational kernels required to solve the Poisson problem, the SPIRAL generated code only

uses one abstraction and one loop. All the three kernels namely Laplace, Jacobi and MaxNorm are fused into one single loop. As a result this boosts the overall performance of the problem.

In Fig. 8, we provide a comparison of the run time between Proto and ProtoX. We compare three different box sizes - 64×64 , 128×128 and 256×256 . We keep the number of boxes fixed to be 4×4 . This makes the corresponding domain sizes of the problem as 256×256 , 512×512 and 1024×1024 respectively. We ran both the Proto as well as the ProtoX code for a fixed 100 Jacobi iterations. We can observe from the graph that ProtoX performs up to $2 \times$ faster than the base Proto code for the 2D Poisson problem. We would like to remark here that these results were obtained on a local CPU machine with 2.3GHz Quad-core Intel i7 processor.

In addition to working on code generation for different targets, we are also working on adding more examples to ProtoX from the current Proto library. We are currently working on the Euler equations that are used in gas dynamics. Here finite volume method is used to solve these equations. This problem is much more complex and intricate with many different types of stencil and pointwise operations required to solve it. This provides a lot of room for optimization for SPIRAL which in return gives a significant improvement in the overall performance of the code. Some of the initial CPU results are shown in Fig 9. We obtain up to $8 \times$ speedup over the base Proto code for the 2D Euler equation problem.

```

1 void Poisson_2D_fused(double *Y, double *X,
2 double weight1, double lambda1, double *rhs,
3 double a_h1, double *retvall){
4     for(int i1 = 0; i1 <= 4095; i1++) {
5         double s20, s21, s22;
6         int a48, b15;
7         b15 = ((66*(i1 / 64)) + (i1 % 64));
8         a48 = (b15 + 67);
9         s20 = X[a48];
10        s21 = ((X[(b15 + 1)] - (4.0*s20))
11        + X[(b15 + 66)] + X[(b15 + 68)]
12        + X[(b15 + 133)]);
13        s22 = rhs[i1];
14        Y[a48] = ((s20 + (weight1*s21))
15        - (lambda1*s22));
16        *(retvall) = (((*(retvall) >=
17        fabs(((1.0/(a_h1*a_h1))*s21)-s22))))
18        ? (*(retvall)) :
19        (fabs(((1.0/(a_h1*a_h1))*s21)
20        - s22)))));
21    }
22 }

```

Fig. 5: SPIRAL generated CPU code for the merged 2D Poisson equation for a `Box` of size 64×64 . All the abstractions from Proto have been fused into one single function call and one single for loop

V. CONCLUSION AND FUTURE WORK

This work demonstrates a proof of concept of having SPIRAL code generation as a backend to Proto using the 2D Poisson problem. It shows that by writing a Proto program as a SPIRAL specification, we can interpret Proto as a DSL and powerful code generation is possible. This work is an early prototype, so only CPU performance for some specific sizes has been shown. Considering that these are initial results,

```

1  #include <omp.h>
2  #include <math.h>
3  const int NUM_THREADS = 4;
4
5  //SPIRAL code for OpenMP
6  void possion_2d(double *Y, double *X, double weight1,
7  double lambda1, double a_h1, double *rhs,
8  double *retvall) {
9  #pragma omp parallel num_threads(4)
10 #reduction (max : retval)
11 {
12     { /* begin parallel loop */
13         int tid1 = omp_get_thread_num();
14         for(int il = tid1; il <= 4095; il += 4) {
15             int a51, b15;
16             double s28, s29, s30, s31, s32;
17             b15 = ((66*(il / 64)) + (il % 64));
18             a51 = (b15 + 67);
19             s28 = X[a51];
20             s29 = ((X[(b15 + 1)] - ((4.0)*(s28)))) +
21                 X[(b15 + 66)] + X[(b15 + 68)] +
22                 X[(b15 + 133)];
23             s30 = rhs[il];
24             s31 = ((s28 + ((weight1)*(s29))) -
25                 ((lambda1)*(s30)));
26             Y[a51] = s31;
27             s32 = max(*retvall,
28                 abs((((1.0) / ((a_h1)*(a_h1))))
29                 *(s29) - (s30)));
30             *(retvall) = s32;
31         }
32     } /* end parallel loop */
33 }
34 }

```

Fig. 6: SPIRAL generated code for OpenMP

```

1  #include "hip/hip_runtime.h"
2
3  __global__ void ker_code0(double *X, double *Y,
4  double weight1, double lambda1, double a_h1,
5  double *retvall) {
6  if (((256*blockIdx.x) + threadIdx.x) < 4096)) {
7  double s21, s22;
8  int a66, a67, b16;
9  a66 = (threadIdx.x + (256*blockIdx.x));
10 b16 = ((66*(a66 / 64)) + (threadIdx.x % 64));
11 a67 = (b16 + 67);
12 s21 = X[a67];
13 s22 = ((X[(b16 + 1)] - (4.0*s21)) + X[(b16 + 66)]
14 + X[(b16 + 68)] + X[(b16 + 133)]);
15 Y[a67] = ((s21 + (weight1*s22)) - (lambda1
16 *X[(a66 + 4356)]));
17 *(retvall) = (((*(retvall) >=
18 fabs((((1.0) / (a_h1*a_h1))*s21) - s22)))) ?
19 (*(retvall)):
20 (fabs((((1.0) / (a_h1*a_h1))*s21) - s22));
21 }
22 }
23 void possion_2d(double *Y, double *X, double weight1,
24 double lambda1, double a_h1, double *rhs,
25 double *retvall) {
26 dim3 b17(256, 1, 1), g1(17, 1, 1);
27 hipLaunchKernelGGL(ker_code0, dim3(g1), dim3(b17),
28 0, 0, X, Y, weight1, lambda1, a_h1, retvall);
29 }

```

Fig. 7: SPIRAL generated code for GPU

we can still observe that there is a reasonable amount of performance improvement happening for the given problem specification. The fused code for GPU and multi-core CPU is also shown. The run time performance results for these two architectures is part of the ongoing work.

Proto is a full featured packaged library that is being used for many scientific applications like astrophysics, particle physics, fluid and gas dynamics and many more. It has many

2D Poisson Performance on CPU

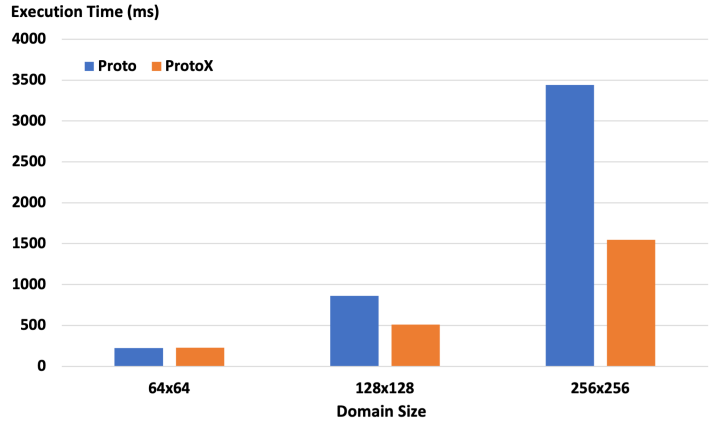


Fig. 8: Run time comparison between the reference Proto code and the SPIRAL backed ProtoX code generated code for CPU. Here we are comparing different BOX sizes ranging from 64×64 to 256×256 for a fixed 100 iterations. We can observe that the ProtoX is performing up to 2× faster than the base Proto code.

2D Euler Performance on Frontier




Fig. 9: Initial run time results for 2D Euler equations. This is part of the ongoing work of adding more examples to ProtoX from the current Proto library. We can observe here that SPIRAL generated fused code for the Euler equations present in ProtoX provides up to 8× better performance compared to the baseline Proto code for the same problem.

implementations of different types of numerical algorithms like the Adaptive Mesh Refinement (AMR), multigrid method and finite volume method. Hence, one of the future goals is to bring ProtoX at par with Proto in terms of the examples and different target platforms used in there. Other part of the future work is to make ProtoX interoperable with FFTX [3] to do cross library optimization.

VI. ACKNOWLEDGMENT

This project is funded by the DOE Office of Advanced Scientific Computing Research the Office of Science of the

REFERENCES

- [1] M. Puschel, J. Moura, J. Johnson, D. Padua, M. Veloso, B. Singer, J. Xiong, F. Franchetti, A. Gacic, Y. Voronenko, K. Chen, R. Johnson, and N. Rizzolo, "Spiral: Code generation for dsp transforms," *Proceedings of the IEEE*, vol. 93, no. 2, pp. 232–275, 2005.
- [2] F. Franchetti, Y. Voronenko, and M. Püschel, "Formal loop merging for signal transforms," *SIGPLAN Not.*, vol. 40, no. 6, p. 315–326, jun 2005.
- [3] F. Franchetti, D. G. Spampinato, A. Kulkarni, D. Thom Popovici, T. M. Low, M. Franusich, A. Canning, P. McCorquodale, B. V. Straalen, and P. Colella, "Fftx and spectralpack: A first look," in *2018 IEEE 25th International Conference on High Performance Computing Workshops (HiPCW)*, 2018, pp. 18–27.
- [4] S. Rao, A. Kutuluru, P. Brouwer, S. McMillan, and F. Franchetti, "Gbtlx: A first look," in *2020 IEEE High Performance Extreme Computing Conference (HPEC)*, 2020, pp. 1–7.
- [5] 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," vol. 48, no. 6, 2013.
- [6] B. Hagedorn, L. Stoltzfus, M. Steuer, S. Gorlach, and C. Dubach, "High performance stencil code generation with lift," ser. CGO 2018. New York, NY, USA: Association for Computing Machinery, 2018, p. 100–112.
- [7] T. Brandvik and G. Pullan, "Sblock: A framework for efficient stencil-based pde solvers on multi-core platforms," in *2010 10th IEEE International Conference on Computer and Information Technology*, 2010, pp. 1181–1188.
- [8] P. Basu, M. Hall, S. Williams, B. Van Straalen, L. Oliker, and P. Colella, "Compiler-directed transformation for higher-order stencils," in *2015 IEEE International Parallel and Distributed Processing Symposium*, 2015, pp. 313–323.
- [9] M. Christen, O. Schenk, and H. Burkhardt, "Patus: A code generation and autotuning framework for parallel iterative stencil computations on modern microarchitectures," in *2011 IEEE International Parallel and Distributed Processing Symposium*, 2011, pp. 676–687.
- [10] J. Holewinski, L.-N. Pouchet, and P. Sadayappan, "High-performance code generation for stencil computations on gpu architectures," in *Proceedings of the 26th ACM International Conference on Supercomputing*, ser. ICS '12. New York, NY, USA: Association for Computing Machinery, 2012, p. 311–320.
- [11] M. Li, Y. Liu, H. Yang, Y. Hu, Q. Sun, B. Chen, X. You, X. Liu, Z. Luan, and D. Qian, "Automatic code generation and optimization of large-scale stencil computation on many-core processors," in *Proceedings of the 50th International Conference on Parallel Processing*, ser. ICPP '21. New York, NY, USA: Association for Computing Machinery, 2021.
- [12] C. Lengauer, S. Apel, M. Bolten, A. Größlinger, F. Hannig, H. Köstler, U. Rüde, J. Teich, A. Grebhahn, S. Kronawitter, S. Kuckuk, H. Rittich, and C. Schmitt, "Exastencils: Advanced stencil-code engineering," in *Euro-Par 2014: Parallel Processing Workshops*, L. Lopes, J. Žilinskas, A. Costan, R. G. Cascella, G. Kecskemeti, E. Jeannot, M. Cannataro, L. Ricci, S. Benkner, S. Petit, V. Scarano, J. Gracia, S. Hunold, S. L. Scott, S. Lankes, C. Lengauer, J. Carretero, J. Breitbart, and M. Alexander, Eds. Cham: Springer International Publishing, 2014, pp. 553–564.
- [13] C. Yount, J. Tobin, A. Breuer, and A. Duran, "Yask—yet another stencil kernel: A framework for hpc stencil code-generation and tuning," in *2016 Sixth International Workshop on Domain-Specific Languages and High-Level Frameworks for High Performance Computing (WOLFHPC)*, 2016, pp. 30–39.
- [14] M. Essadki, B. Michel, B. Maugars, O. Zinenko, N. Vasilache, and A. Cohen, "Code generation for in-place stencils," in *Proceedings of the 21st ACM/IEEE International Symposium on Code Generation and Optimization*, ser. CGO 2023. New York, NY, USA: Association for Computing Machinery, 2023, p. 2–13.
- [15] J. Xiong, J. Johnson, R. Johnson, and D. Padua, "Spl: A language and compiler for dsp algorithms," *SIGPLAN Not.*, vol. 36, no. 5, p. 298–308, may 2001.
- [16] F. Franchetti, F. de Mesmay, D. McFarlin, and M. Püschel, "Operator language: A program generation framework for fast kernels," in *Domain-Specific Languages*, W. M. Taha, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 385–409.
- [17] M. Bolten, F. Franchetti, P. H. J. Kelly, C. Lengauer, and M. Mohr, "Algebraic description and automatic generation of multigrid methods in spiral," *Concurrency and Computation: Practice and Experience*, vol. 29, no. 17, p. e4105, 2017, e4105 cpe.4105.