



# FFTc: An MLIR Dialect for Developing HPC Fast Fourier Transform Libraries

Yifei He<sup>(✉)</sup>, Artur Podobas, Måns I. Andersson, and Stefano Markidis

KTH Royal Institute of Technology, Stockholm, Sweden  
{yifeihe, podobas, mansande, markidis}@kth.se

**Abstract.** Discrete Fourier Transform (DFT) libraries are one of the most critical software components for scientific computing. Inspired by FFTW, a widely used library for DFT HPC calculations, we apply compiler technologies for the development of HPC Fourier transform libraries. In this work, we introduce FFTc, a domain-specific language, based on Multi-Level Intermediate Representation (MLIR), for expressing Fourier Transform algorithms. We present the initial design, implementation, and preliminary results of FFTc.

**Keywords:** MLIR · Fast Fourier Transform Compiler · DSL

## 1 Introduction

HPC libraries for computing Discrete Fourier Transforms (DFT) are critical computational building blocks for enabling signal processing, data analysis, and the solution of Partial Differential Equations (PDE). In particular, Fast Fourier Transform (FFT) algorithms solve DFT via  $\mathcal{O}(n \log n)$  calculations, where  $n$  is the input size against the naive DFT implementation corresponding to a matrix-vector multiply with complex numbers requiring  $\mathcal{O}(n^2)$  calculations.

Several algorithms for FFT have been designed, including the notorious Cooley-Tukey recursive scheme to the Stockham and Pease algorithms [11]. FFT algorithms can be expressed using a factorized formulation, e.g., the entire FFT operation is expressed as the multiplication of matrices, and different algorithms will correspond to various factorization forms. These matrices are largely sparse, and their final computation will still rely only on  $\mathcal{O}(n \log n)$  operations. Therefore, from an abstraction point of view, we can express any FFT algorithms in terms of matrix multiplications. Most importantly for this work, different factorizations are better suited than others for achieving high-performance on a given system. For instance, Stockham FFT factorization is an excellent fit for accelerators while other factorizations containing block matrices are a good fit for hierarchical memory systems. For this reason, to be capable of expressing and generating automatically and optimizing different FFT algorithms for different architectures is critical for producing high-performance FFT libraries.

FFTW [8] is among the most successful implementations of FFT libraries. Inspired by the FFTW design and development, in this work, we propose a new

framework, called FFTc (FFT compiler), for the automatic generation of FFT algorithms using the MLIR and LLVM compiler infrastructure. To achieve this, we design a new language to express FFT algorithms using different formulations. The major contributions of this paper are the following:

- We design and provide a first initial development of a domain-specific language for the automatic code generation of FFT algorithms, leveraging MLIR and LLVM infrastructure.
- We collect and analyze the preliminary performance results from small-size one-dimensional FFT and compare the performance with the FFTW performance.

## 2 Background

The goal of this work is to develop a DSL for FFT calculation. A direct computation of the Fourier transform is the multiplication of a DFT matrix by the input vector  $x$ . We can define the  $DFT_N$  matrix as:

$$DFT_{N,m,n} = (\omega_N)^{mn}, \quad \text{where } \omega_N = \exp(-2\pi i/N) \quad \text{for } 0 \leq m, n < N. \quad (1)$$

The most famous FFT algorithm was introduced in 1965 by Cooley and Tukey. This algorithm relies on the recursive nature of DFT i.e. several small DFTs can describe a large DFT. In this paper, we use a matrix-formalism to represent FFT algorithms where a matrix-factorization of the DFT matrix into sparse and structured matrices describes each FFT algorithm. For example the Cooley-Tukey factorization of  $DFT_4$ :

$$DFT_4 = \underbrace{\begin{bmatrix} 1 & & & \\ & 1 & & \\ & & 1 & \\ & & & 1 \end{bmatrix}}_{DFT_2 \otimes I_2} \begin{bmatrix} 1 & & & \\ & 1 & & \\ & & 1 & \\ & & & -i \end{bmatrix} \underbrace{\begin{bmatrix} 1 & 1 & & \\ & 1 & -1 & \\ & & 1 & 1 \\ & & & 1 & -1 \end{bmatrix}}_{I_2 \otimes DFT_2} \begin{bmatrix} 1 & & & \\ & 1 & & \\ & & 1 & \\ & & & 1 \end{bmatrix}, \quad (2)$$

where the  $I$  is the identity matrix. Here, we see the use of  $DFT_2$  in the formulation of  $DFT_4$ . In the example, we see the sparse (zeros in the matrices are omitted for clarity) and structured nature of the algorithm. The Cooley-Tukey general-radix decimation-in-time algorithm for  $N$  inputs can be written as:

$$DFT_N = (DFT_K \otimes I_M) D_M^N (I_K \otimes DFT_M) \Pi_K^N \quad \text{with } N = MK, \quad (3)$$

where  $\Pi_K^N$  is a stride permute and  $D_M^N$  is a diagonal matrix of *twiddle*-factors. Different FFT algorithms, such as Stockham and Pease FFT can be expressed using different factorization schemes.

In this work, we use the LLVM (originally for Low-Level Virtual Machine) compiler infrastructure for the development of the FFT domain-specific language. LLVM is a collection of compiler and toolchain technologies: it consists of

a set of modular compiler components, including the Clang front-ends, optimizer, code generator, debugger, linker, and OpenMP runtime. Particularly important for developing portable HPC code, the LLVM compiler technologies support many targets, including x86, Arm, and GPU systems [9].

The LLVM project also includes Multi-Level Intermediate Representation (MLIR), a project aiming at supporting the building of domain-specific compilers, and combining existing compiler infrastructure together. While MLIR (and the XLA compiler) was initially developed by Google for machine learning workloads, MLIR is widely used today for the development of domain-specific languages beyond machine and deep learning. To solve domain-specific problems, MLIR offers the infrastructure to define and introduce high-level abstractions and transforms [10]. The main mechanism to extend MLIR is the development of *dialects* that allow defining new operations, attributes, and types. In addition, MLIR allows using multiple dialects that can be used together within one module. Examples of existing MLIR dialects are the affine, LLVM, GPU, vector, SPIR-V dialects. In this work, we design and develop an MLIR dialect to express FFT libraries.

### 3 Related Work

Several efforts exist for the development of high-performance FFT libraries. The inspiration for developing an FFT DSL is FFTW [7], which is the most widely used open-source FFT library. At its heart, FFTW is an FFT compiler, based on Objective Caml, to generate Directed Acyclic Graphs (DAG) of FFT algorithms and performs algebraic optimization on them. FFTW uses a planner at runtime to recursively decompose the DFT problem into sub-problems. These sub-problems are solved directly by optimized, straight-line code that is automatically generated by a special-purpose compiler, called *genfft* [8]. An additional DSL for numerical kernels including FFT is SPIRAL. SPIRAL [6] is a program generation system for linear transforms and other mathematical functions that produces HPC code in C. SPIRAL also supports FFTs [5]: it applies pattern match and rewriting to generate optimal FFT formulation for different hardware, such as multicore systems. Then, SPIRAL maps the matrix formula to high-performance C code.

### 4 Methodology: A Domain-Specific Language for FFT

This section describes FFTc— a custom Domain-Specific Language (DSL) for describing Fast Fourier Transforms (FFT). Our aspiration with FFTc is to increase the productivity of algorithm developers without any loss in performance while at the same time being able to target multiple different backends (CPUs/GPUs/etc.) with the same input source code. In short, FFTc aims to increase *productivity*, *portability*, and (hopefully) *performance*.

The execution model and compilation pipeline are shown in Fig. 1. The current implementation supports the parts in dark color; the remaining parts will be the focus in the near future.

The FFTc compilation pipeline has five core parts: (a) is the translation from the DSL to the Abstract Syntax Tree (AST), (b) is generating the MLIR out of AST, (c) stands for progressive lowering from FFT dialect to LLVM dialect, going through different levels of abstraction represented by dialects, (d) emits LLVM IR out of the MLIR’s LLVM dialect, (e) is the LLVM middle-end compilation and code generation.

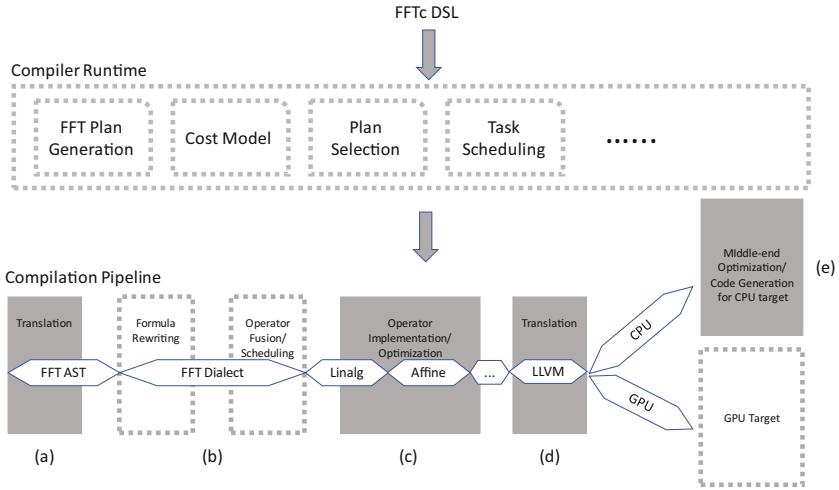


Fig. 1. Compilation Pipeline

#### 4.1 The FFTc Language and Grammar

The goal of FFTc is to create an input language that resembles (as close as possible) that of mathematics, which we believe will help end-users in being more productive without losing familiarity with the code they are writing. An example source code of ours is seen in Listing 1.1, where we have aimed to keep them as similar to abstract mathematical expressions as possible, such as Eq. (3). We support the *Kronecker product* through the binary operation ‘ $\otimes$ ’, the matrix-matrix multiplication using ‘ $\cdot$ ’, and the matrix multiplication with the twiddle matrix through the `twiddle`. Furthermore, we have a set of unary operations, such as creating the `identity` matrix, and calculating the `dft`. Finally, we have support for `permute`. In short, we currently support all necessary language constructs to describe FFTs in a factorized form. Additionally, our grammar supports the correct right-associative binding of (e.g.,) matrix multiplication, which is different from the traditional left-associative binding of binary operators. A subset of the grammatical construct (in Backus-Naur form) is shown in

Listing 1.2. The grammatical construct is based on (and extended) from LLVM’s Kaleidoscope language tutorial [2].

```

1 var InputReal <4, 1> = [[1], [2], [3], [4]];
2 var InputImg <4, 1> = [[1], [2], [3], [4]];
3 var InputComplex = createComplex(InputReal, InputImg);
4 var result = (DFT(2) ⊗ I(2)) · twiddle(4,2) ·
5             (I(2) ⊗ DFT(2)) · Permute(4,2) · InputComplex
;

```

Listing 1.1. DSL FFT language

```

1 expression -> additive-expr ('+' | '-') additive-expr
2 additive-expr -> (multiplicative-expr ( '*' | '/' )
   multiplicative-expr)
3 multiplicative-expr -> (FFT-expr ( '*' | '/' ) FFT-expr) *
4 FFT-expr -> (primary ( '⊗' | '·' ) FFT-expr) *
5 primary -> identifierexpr | numberexpr | parenexpr |
   tensorliteral

```

Listing 1.2. FFTc Language Grammar Extension in Backus-Naur form

## 4.2 FFTc Compilation Pipeline

The FFTc compilation pipeline shown in Fig. 1 is based on the MLIR’s tutorial project [4]. The compilation starts at the frontend (Fig. 1:a), where the lexical analysis, parsing, and building of an Abstract Syntax Tree (AST) based on our custom DSL language take place. The FFT dialect is the first state of MLIR generated from the AST (Fig. 1:b). Then a series of lowering passes are applied (Fig. 1:c) on the FFT dialect in order to expand many of the custom operators (e.g., the Kronecker product) into a lowered state. For example, a matrix multiplication, written in our language using “·”, will be expanded to a three-level nested loop implementing said matrix multiplication. Furthermore, we can apply several existing MLIR optimization passes (such as Affine) in order to further optimize the transformed kernels. Finally, near the end of the pipeline (Fig. 1:c), we lower our representation to the LLVM Intermediate Representation (IR), after which we inject the code into the LLVM backend for compilation towards machine code (Fig. 1:d). We explain this pipeline in more detail next.

**4.2.1 Phase 1: Translation** The FFT dialect is the first dialect in the compilation pipeline. The FFT dialect provides the basic building blocks for different kinds of FFT algorithms and defines the complex tensor data type and operations.

- **FFT dialect data type:** The FFT dialect operates on the double tensor and complex tensor as well as scalar integer as attributes. There is `createComplex` to generate the complex tensor from the double tensor of real and imaginary parts.

- **FFT dialect operations:** We define the operations needed to implement various kinds of popular FFTs. Examples of such operators are the Kronecker product and matrix-matrix multiplication. We also define the `DFT`, `Identity`, and `permute` matrix generator. These make it a lot easier to construct the FFT algorithm with the similar notation and syntax in mathematics. The map of operations from FFTc DSL to MLIR FFT dialect is shown in Table 1.

With the FFT dialect implementation described above, we can generate MLIR out of AST, as shown in (a) to (b) in Fig. 1. Figure 2 shows an example of the FFT dialect IR that is translated from size 4 recursive FFT in Listing 1.1.

**Table 1.** From FFTc DSL to FFT Dialect MLIR

FFTC DSL	FFT Dialect
<code>createComplex(A, B)</code>	<code>fft.createCT(a,b)</code>
<code>A · B</code>	<code>fft.matmul a, b :</code>
<code>A ⊗ B</code>	<code>fft.kroneckerproduct a, b</code>
<code>twiddle (a,b)</code>	<code>fft.twiddle (a , b)</code>
<code>I(size)</code>	<code>fft.identity (a)</code>
<code>DFT(size)</code>	<code>fft.dft(a)</code>
<code>Permute (a ,b)</code>	<code>fft.Permute(a, b)</code>

**4.2.2 Phase 2: Operator Implementation/Optimization** MLIR supports different levels of abstraction through dialects. We lower the FFT dialect to a mix of dialects. Then, we can reuse the analysis/transform passes embedded in those dialects. We run shape inference to prepare for later transforms and perform progressive lowering to a mix of dialects to implement and optimize FFT operations.

- **Shape Inference:** In the FFTc DSL, all the operations operate on generic tensors. We do not need to explicitly specify the shape of tensor data. This reduces the efforts of the programmers. However, carrying shape information in the IR can simplify the workload of analysis and transform passes, as well as code generation. We can obtain the shape of input tensors during the initialization of constants. Later, we propagate the shapes through the computation to every operation involved. We implement a specific shape inference function for each operation based on the input augments, such as for the Kronecker product. All dimensions of the output tensor would be the multiplication of the corresponding dimensions of two input tensors.
- **Progressive Lowering:** The compilation pipeline generates the actual implementations of the operations, which we defined through progressive lowering. To reuse existing optimizations in MLIR’s dialects, we lower the FFT dialect to a mix of dialects, comprising of Affine, Arithmetic, Complex and

MemRef dialects. The Affine dialect uses techniques from polyhedral compilation to provide a powerful abstraction for affine operations and analyses, such as dependence analysis and loop transformations. The Arithmetic dialect is intended to hold basic integer and floating-point mathematical operations, and the Complex dialect is intended to hold complex numbers creation and arithmetic operations. The MemRef dialect is intended to hold core memref creation and manipulation operations [3].

- **Affine Dialect:** We implement the computation-heavy part of the DSL in Affine dialect, by lowering from the tensor type that FFT dialect operates on to the MemRef type that is indexed via an affine loop-nest. Tensors represent an abstract value-typed sequence of data. By using tensor and tensor operations, we can increase the productivity of algorithm developers since it is similar to the notations used in mathematics. The MemRefs dialect, on the other hand, represents the lower level buffer access, builds a bridge to the actual computer memory.

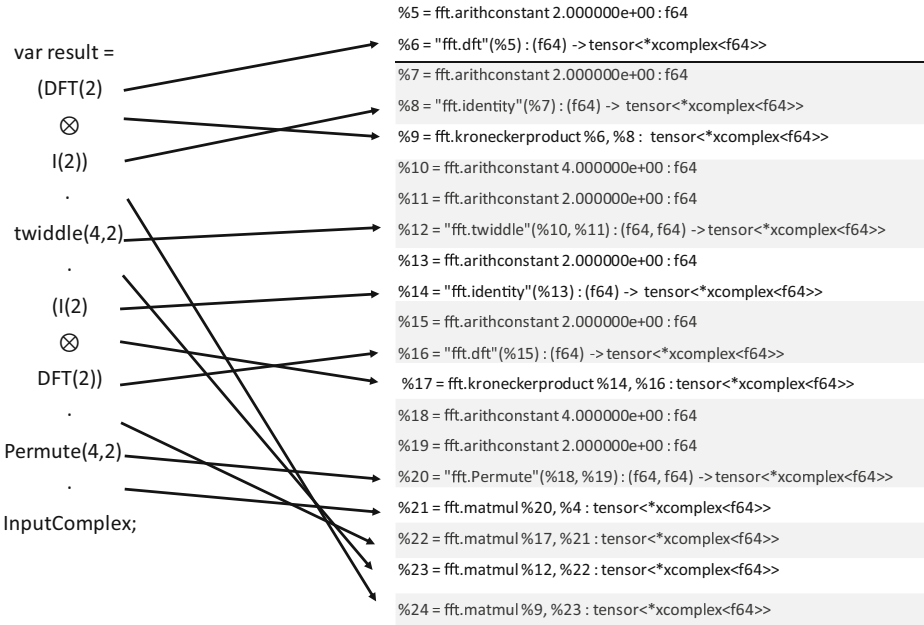


Fig. 2. Mapping from the Recursive FFT to MLIR.

To implement the operators, we allocate a chunk of memory for the output tensor, construct loops to compute each element of the output tensor, then store them to the corresponding index of the output memory. The scalarized tensor arithmetic operations are performed by corresponding operations in the Complex dialect. The lowering result of a matrix multiplication operator is shown in the Listing 1.3. We take advantage of the existing opti-

mizations in the Affine dialect, such as loop fusion, AffineScalarReplacement and AffineLoopInvariantCodeMotion. These optimization passes can help perform operator fusion, eliminate redundant load/store and hoists loop invariant operations out of Affine loops.

```

1 From:
2   %10 = fft.matmul %9, %3 : (tensor<4x4xcomplex<f64>>,
3   tensor<4x1xcomplex<f64>>) ->
4   tensor<4x1xcomplex<f64>>
5 To:
6   affine.for %arg0 = 0 to 4 {
7     affine.for %arg1 = 0 to 1 {
8       affine.for %arg2 = 0 to 4 {
9         %18 = affine.load %9[%arg0, %arg2] :
10        memref<4x4xcomplex<f64>>
11        %19 = affine.load %3[%arg2, %arg1] :
12        memref<4x1xcomplex<f64>>
13        %20 = complex.mul %18, %19 : complex<f64>
14        %21 = affine.load %2[%arg0, %arg1] :
15        memref<4x1xcomplex<f64>>
16        %22 = complex.add %21, %20 : complex<f64>
17        affine.store %22, %2[%arg0, %arg1] :
18        memref<4x1xcomplex<f64>>
19      }
20    }
21  }

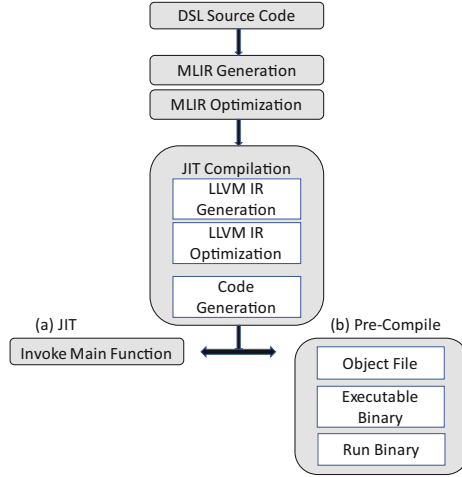
```

**Listing 1.3.** Affine Code Example for FFT.MatMul Operation

**4.2.3 Phase 3: Translation** There exist infrastructures in MLIR to perform a full conversion from the Affine, MemRef, and Complex dialects to the LLVM dialect. Then, we can emit the LLVM IR from the LLVM dialect.

**4.2.4 Phase 4: Code Generation** We set up a JIT compiler using the MLIR wrapper over LLVM OrcJit, and pass the optimization and debug flags to the JIT compiler. The pass manager is also populated by MLIR. Then, the JIT compiler will perform the LLVM’s middle-end optimization and code generation.





**Fig. 3.** Compilation modes.

#### 4.2.4.1 Ahead-Of-Time vs Just-In-Time Compilation

We support two types of compilation modes in FFTc: Ahead-of-Time (AOT) and Just-in-Time (JIT) compilation. The compilation modes can be seen in Fig. 3, where they share multiple components and are in line with similar compilation flows (e.g., in OpenCL’s Online/Offline compilation [1]). In short, both modes start by parsing the DSL source code and transforming/optimizing it using our MLIR intermediate representation. Next, we lower the MLIR down to LLVM IR. Once in LLVM IR, the two modes differ: using the JIT mode, we directly execute the main function of our compiled targets and exit afterward. The AOT mode, instead, transforms the LLVM IR representation to an object file, links with eventual standard libraries, and outputs a machine code binary file that can be invoked by the user.

Using either model has benefits and limitations. For example, the AOT mode can be faster and speed up the final execution significantly but has the limitation that the FFT size needs to be constant. The JIT model, on the other hand, is slower but allows the FFT size to be variable at runtime. In short, the AOT mode trades flexibility for performance, while the JIT mode honors flexibility over performance.

## 5 Experimental Setup

We evaluate our FFTc on the Kebnekaise supercomputer that is located at the HPC2N HPC center in Umeå, Sweden. Kebnekaise nodes have a dual-socket Intel Xeon Gold 6132 CPU, 192 GB of RAM. The operating system is Ubuntu 20.04.4 LTS. The version of LLVM we use to embed the FFTc is 15.0.0. We run the Ahead-of-Time compilation mode FFTc 1,000 times, and we calculate

error computing the standard deviation for 30 execution rounds. We developed a Python script to generate the recursive implementation of the Cooley-Tukey FFT algorithm, using our FFTc DSL. An example of the output program is shown in Listing 1.1. Albeit our script can generate different FFT algorithm implementations, in this paper, we only present the results of the recursive Cooley-Tukey algorithm.

## 6 Results

As first step of our evaluation, we verify the correctness of DSL implementation. We test different random input vectors with different sizes: the input sizes are the powers of two, from 32 to 1024. We employ complex numbers in double-precision. We compare the results with the NumPy’s FFT function, that is based on FFTW. The error is calculated as  $\frac{|result_{DSL} - result_{Numpy}|}{FFT_{size}}$ . The error is smaller than 1e-7 for each run.

For the next step, we evaluate the performance in the JIT mode. We measure the execution time of size 32 recursive FFT under JIT mode. The execution time is shown in Fig. 4. In the figure, the item Parser&MLIRGen stands for frontend compilation, ‘builtin.func’ stands for MLIR compilation pipeline, ‘Jit’ stands for both LLVM Jit compilation and running time. It is clear from analyzing the figure that the frontend takes a minor portion of the execution time. The MLIR pipeline takes the largest part of the execution time. Most of the time is spent in the optimization passes such as AffineLoopFusion and AffineScalarReplacement. We can choose whether to run these optimization passes or not by passing optimization flag to FFTc, currently there are O0/O2/O3 available. The Jit part takes much smaller portion compared with MLIR pipeline, under O3 optimization option for both LLVM middle-end compilation and code generation. In actual applications, FFT algorithms may run many times while only need to be compiled once, so the compilation time does not matter considerably. As a future plan, we intend to reduce the compilation time, such as multi-threading the compiler and remove redundant operations in Affine passes.

Under Pre-Compiled mode, we compare the FFTc pre-compiled binary with FFTW 3.3. We built FFTW with gcc compiler, enabled the SIMD instructions. The input size of the FFT are the powers of 2, we use single thread to run the program. The result is shown in Fig. 5, the standard deviation is shown as the black lines over bars.

We run four versions of FFT using FFTc: direct DFT implementation and Cooley-Tukey recursive FFT implementation with different optimization flags (O0/O2/O3). It is expected that the DFT performs much better than recursive implementations, because current implementation for FFT is computed through dense matrix multiplication, and to achieve the  $O(N \log N)$  complexity FFT must be sparse matrix computation. The workload of the currently developed recursive FFT is much larger than DFT. However, we intend to use the current solution to showcase the functionality of FFTc and are planning to rewrite the

### Total Execution Time Report

Wall Time / Seconds	Name
0.0034 (0.0%)	Parser & MLIRGen
0.0003 (0.0%)	Inliner
0.0000 (0.0%)	(A) CallGraph
0.0000 (0.0%)	'builtin.func' Pipeline
0.0002 (0.0%)	Canonicalizer
6.2268 (90.4%)	builtin.func' Pipeline
0.0001 (0.0%)	{anonymous}::ShapeInference
0.0001 (0.0%)	Canonicalizer
0.0000 (0.0%)	CSE
0.0000 (0.0%)	(A) DominanceInfo
0.0116 (0.2%)	{anonymous}::AffineToLLVMLoweringPass
0.0226 (0.4%)	Canonicalizer
0.0014 (0.0%)	CSE
0.0000 (0.0%)	(A) DominanceInfo
0.6238 (9.1%)	AffineLoopFusion
5.5622 (80.7%)	AffineScalarReplacement
0.0000 (0.0%)	(A) PostDominanceInfo
0.0000 (0.0%)	(A) DominanceInfo
0.0009 (0.0%)	AffineLoopInvariantCodeMotion
0.0384 (0.6%)	{anonymous}::FFTToLLVMLoweringPass
0.0000 (0.0%)	output
0.6154 (8.9%)	Jit
0.0057 (0.1%)	Rest
6.8903 (100%)	Total

Fig. 4. JIT Mode Performance for size 32 recursive FFT

computation in sparse form as a future work. The performance with optimization flag O3 is better than O2 and O0. The difference between O2 and O3 flag is that under O2, the AffineScalarReplacement pass will not be executed. For size




Fig. 5. FFTc Single Thread Performance Compared with FFTW

128 the O2 is slightly better than O3. Investigating the MLIR, the AffineScalar-Replacement performs memory access optimizations. In addition, there is also a similar optimization pass in LLVM pipeline. We plan to further investigate this issue in the future.

When comparing the performance between FFTc Cooley-Tukey code and FFTW, we note that here is still a significant performance gap. We believe that this gap can be attributed to (amongst others) the following reasons:

- The recursive factorized FFTs are computed through matrix-matrix multiplication where the matrices are not expressed as sparse matrices.
- We do not take full advantage of MLIR/LLVM infrastructure to generate high performance code. Examples of such a features are loop tiling, unrolling and jam and vectorization in the MLIR/LLVM pipeline.
- We do not support yet an autotuning mechanism, such as the FFTW *planner*, to decompose the FFT problem into simpler sub-problems, later solve the simpler sub-problems using codelets generated by genfft. Currently, our implementation is similar to genfft: for the FFTs with large-size input, the generated code is extremely large and introduces considerable compilation overhead.

## 7 Discussion and Conclusion

In this paper, we have introduced FFTc— an emerging, work-in-progress DSL for describing different FFTs variants. The goal of FFTc is to decouple algorithm description from hardware-specific details and ultimately provide higher productivity and better portability without sacrificing performance. To this end, we have chosen an abstract language representation that is not unlike the mathematical formulas we are used to describing FFTs. We show how such an abstract language design can be mapped down-to machine code by leveraging existing MLIR and LLVM infrastructure. The performance – while not a direct objective of this paper – of our DSL is not yet on par with state-of-the-art FFTW, but is never-the-less a good starting point to further build upon in future performance-focused studies, such as extending our compiler with support for OpenMP tasking or vectorization.

**Acknowledgement.** Funding for the work is received from the European High-Performance Computing Joint Undertaking (JU), Grant Agreement No. 3893 (IO-SEA). I want to thank Steven W. D. Chien (wdchien@kth.se) for his help with the proofread.

## References

1. Intel Online/Offline Compilation. <https://www.intel.com/programmable/technical-pdfs/683521.pdf>
2. Kaleidoscope: implementing a parser and AST. <https://llvm.org/docs/tutorial/MyFirstLanguageFrontend/LangImpl02.html>

3. MLIR dialects document. <https://mlir.llvm.org/docs/Dialects/>
4. MLIR toy language. <https://mlir.llvm.org/docs/Tutorials/Toy/Ch-1/>
5. Franchetti, F., al.: Discrete fourier transform on multicore. *IEEE Sig. Process. Mag.* **26**(6), 90–102 (2009)
6. Franchetti, F., al.: SPIRAL: extreme performance portability. *From High Level Specification High Performance Code* **106**(11), 1935–1968 (2018)
7. Frigo, M.: A fast fourier transform compiler. In: *ACM SIGPLAN 1999 Conference on Programming Language Design and Implementation*, pp. 169–180 (1999)
8. Frigo, M., Johnson, S.G.: The design and implementation of FFTW3. *Proc. IEEE* **93**(2), 216–231 (2005)
9. Lattner, C., Adve, V.: LLVM: a compilation framework for lifelong program analysis amp; transformation. In: *CGO 2004*, pp. 75–86 (2004)
10. Lattner, C., al.: MLIR: Scaling compiler infrastructure for domain specific computation. In: *2021 IEEE/ACM International Symposium on Code Generation and Optimization*, pp. 2–14 (2021)
11. Van Loan, C.: Computational frameworks for the fast Fourier transform. In: *SIAM* (1992)