

Optimal Dynamic Data Layouts for 2D FFT on 3D Memory Integrated FPGA

Ren Chen · Shreyas G. Singapura ·
Viktor K. Prasanna

Received: date / Accepted: date

Abstract FPGAs have been widely used for accelerating various applications. For many data intensive applications, the memory bandwidth limits the performance. 3D memories with through-silicon-via connections provide potential solutions to the latency and bandwidth limitations. In this paper, we revisit the classic 2D FFT problem to evaluate the performance of 3D memory integrated FPGA. To fully utilize the fine grained parallelism in 3D memory, data layouts which take into account the structure and organization of the memory are required. We propose dynamic data layouts for optimizing the performance of the 3D architecture. In 2D FFT, data is accessed in row major order in the first phase whereas, the data is accessed in column major order in the second phase. This column major order results in high memory latency and low bandwidth due to high row activation overhead of memory. Using the proposed dynamic data layouts, we improve memory access performance in the second phase without degrading the performance of the first phase. With parallelism employed in the third dimension of the memory, data parallelism can be increased to further improve the performance. We adopt a model based approach for 3D memory and we perform experiments on the FPGA to validate our analysis and evaluate the performance. Compared with the baseline architecture, our approach achieves up to $40\times$ peak memory bandwidth utilization

This material was supported by the NSF under Grant Number ACI-1339756

Ren Chen, Shreyas G. Singapura and Viktor K. Prasanna
University of Southern California, Los Angeles, CA 90089, USA

Ren Chen
E-mail: renchen@usc.edu

S. G. Singapura
E-mail: singapur@usc.edu

V. K. Prasanna
E-mail: prasanna@usc.edu

for column-wise FFT, thus resulting in approximately 97% improvement in throughput for the complete 2D FFT application.

1 Introduction

FPGAs have been used as accelerators for many applications such as Signal Processing, Image Processing, Packet classification etc [1–5]. The general purpose processors cannot keep up with the demands of these applications in terms of performance. Even with the high performance of FPGAs, meeting the throughput requirement of these applications is a challenging task. Most of the applications are data intensive and this translates to frequent accesses to the memory. The bottleneck in these cases is the low bandwidth and high latency of the memory.

3D memory has been widely studied in the research community with the high bandwidth and short latency access being the important parameters. 3D memories consist of stack of layers connected using Through Silicon Vias (TSVs) [6]. The high speed vertical TSVs along with the third dimension of memory result in short latencies and packs in large memory sizes compared to the conventional 2D memories. Although 3D memories are expected to provide $10\times$ bandwidth compared to 2D memory, this is subject to the ideal conditions. These include data layouts which reduce row activation overhead, high page hit rate for stride access, etc. These problems are similar to the issues in the conventional planar memories. But, employing the solutions in the context of 3D memory is not trivial due to the structure and organization of 3D memory.

In this paper, we target 2D FFT application on 3D memory integrated FPGA and evaluate its performance using throughput and latency as the metrics. 2D FFT is a data intensive application with stride memory access patterns. 2D FFT consists of two phases and the access patterns in the two phases require different data layouts to achieve high performance. The ideal data layout in the first phase is row major data layout whereas, the second phase requires a column major data layout. Therefore, a static data layout which improves the performance in one phase will lower the performance in the other phase and vice-versa. The main reasons for this low performance are high number of row activations and low page hit rate. Therefore, with a static data layout the true capability of 3D memory cannot be realized. We address this problem by extending our solution of dynamic data layouts [7] to 3D memory. The main contributions in this paper are:

1. Model the 2D FFT application on 3D memory integrated FPGA.
2. Develop optimal dynamic data layouts to optimize performance of 2D FFT on 3D memory.
3. Evaluation of optimized and baseline implementation with throughput and latency as the performance metrics.




Fig. 1: (a) 3D MI-FPGA Architecture (b) 3D Memory

2 Related Work

As the well-known simplest multidimensional FFT algorithm, the row-column algorithm has been commonly used to implement 2D FFT by performing a sequence of 1D FFTs [5, 8]. In this algorithm, input elements in an $N \times N$ array are stored in row-major order in the external memory such as DRAM. One major issue in the implementation of the 2D FFT architecture is the considerable delay caused by DRAM row activation which are mainly introduced by the strided memory access in the column wise 1D FFTs. To solve this problem, the authors in [3] propose a tiled data mapping method to improve the external memory bandwidth utilization. They logically divide the input $N \times N$ input array into $\frac{N}{k} \times \frac{N}{k}$ tiles and map the elements in each tile to consecutive memory locations. They conclude that the DRAM bandwidth utilization is maximized when the size of each tile is set to be the size of the DRAM row buffer. The authors in [9] extend the block data layout to 3D memory as well. However, this solution introduces non-trivial on-chip hardware resource cost for local transposition. Various traditional 2D memory based 2D FFT architectures achieving high throughput performance have been developed in [5, 10]. In [5], the authors propose a 2D decomposition algorithm which enables local 2D FFT on sub-blocks. In this way, the times of DRAM row activation is minimized. Vector radix 2D FFT in [10] presents a general structure theorem to construct a multi-dimensional FFT architecture by decomposing a large size problem into small size 2D FFTs. The external memory row activation overhead is not considered.

3D memory is expected to provide bandwidth higher than the 2D memory by an order of magnitude. There have been many works which have focused on this aspect of 3D memory. [11] implements matrix multiplication and 2D FFT on a Logic-in-Memory architecture. The architecture consists of a logic layer interleaved between two segments of memory layers to form a 3D architecture. The performance metrics are energy efficiency and bandwidth. In [12], the authors develop power efficient FFT on an architecture consisting of memory layers stacked on multiple FPGA layers. The authors focus on energy efficiency improvement while moving to a 3D architecture from a 2D architecture.

3 3D Memory Integrated FPGA (3D MI-FPGA)

3.1 Architecture Overview

Our model of 3D architecture consists of 3D memory integrated with FPGA interacting through TSVs. We extend our previous work on 3D architectures [13, 14]. Here, we provide a brief overview of 3D Memory Integrated FPGA (3D MI-FPGA). The architecture consists of three components: 3D memory, FPGA and TSVs. Fig. 1 illustrates the architecture of 3D memory integrated FPGA. The memory is composed of several layers (L_m) vertically stacked one above the other. Each of these layers is partitioned into several banks. Vaults (V) are defined as the group of banks (1, 2, 3, 4 in Fig. 1) across layers which share a set of interconnects (TSVs). This set of banks residing on one layer which belong to the same vault (B) is analogous to the number of banks in a chip in the 2D memory. These set of banks share the bus in 2D memory and they share the TSVs in the 3D memory architecture. This set of TSVs shared by the banks in a vault is denoted by N_{tsv} . Each vault has a dedicated memory controller which handles the memory accesses to that particular vault. These memory controllers form a separate layer in the memory. Different vaults can be activated at the same time as they do not share the TSVs. On the other hand, the banks in a given vault share the TSVs and the activation of these banks has to be pipelined or interleaved as in the case of 2D memory. Denoting by BW_{vault} the bandwidth of a vault, the peak bandwidth of 3D memory is $V \times BW_{vault}$. The FPGA architecture is similar to that of the conventional FPGA consisting reconfigurable logic, DSP blocks, on-chip memory (Block RAM and Distributed RAM) and memory controllers. The difference is that we model the FPGA to interact with the memory through the set of TSVs connecting the FPGA and the memory. These TSVs are between memory controllers on FPGA and those in the memory. FPGA accesses the data in the memory through the TSVs which are high speed, low latency vertical interconnects. The TSVs are characterized by the number of TSVs and latency of data transfer across them. These two parameters affect the amount of data that can be transferred between memory and FPGA in a given unit of time. Each TSV can transfer 1 bit of data at a time.

3.2 Timing Parameters

Bandwidth and latency of access to the 3D memory depend on a certain set of timing parameters and we discuss these in this section. Data in the 3D memory is stored in rows which combine to form a bank and a set of banks form a vault. Therefore, each row belongs to a specific bank and vault. When the memory is accessed, depending on the address, a specific row, bank and vault are activated. Therefore, although some of the parameters overlap with that of the 2D memory, certain additional parameters have to be defined taking

into account the architecture and different accesses possible in the context of a 3D memory. We model the 3D memory using the following parameters:

1. $t_{\text{diff-row}}$: minimum time required between issuing two successive activate commands to different rows in the same bank
2. $t_{\text{diff-bank}}$: minimum time required between successive activate commands to different rows in different banks in the same layer in the same vault
3. $t_{\text{in-row}}$: minimum time required between successive accesses to elements in the same row in the same bank
4. $t_{\text{in-vault}}$: minimum time required between accesses to different rows in different banks in different layers in the same vault

The values of the above parameters have a significant impact on latency and bandwidth of the 3D memory. In general, data from different vaults can be accessed simultaneously with zero delay between the accesses. Hence, a parameter such as $t_{\text{diff-vault}}$ is not defined. This is because, since vaults are completely independent and can be active at the same time, this parameter is equal to zero. Since the banks located in different layers but belonging to the same vault can be activated in a pipeline, this latency ($t_{\text{in-vault}}$) is lower than that of accessing data from banks belonging to the same layer and same vault. Other parameters are similar to the parameters of 2D memory. Therefore, accessing data from the same row in a bank ($t_{\text{in-row}}$) is faster than accessing data from two rows in different banks ($t_{\text{diff-bank}}$). The highest latency is seen when we access data from two different rows in the same bank in the same vault denoted by $t_{\text{diff-row}}$.

4 2D FFT Architecture

4.1 1D FFT Kernel

An N -point (floating-point) 1D FFT kernel is implemented by concatenating several basic components including radix block, data path permutation (DPP) unit, and twiddle factor computation (TFC) unit. The design of each architecture component relies on the FFT algorithm in use and will be described in this section. We applied the energy optimizations discussed in [4, 15, 16] to the components to reduce their energy consumption. The 1D FFT kernel supports processing continuous data streams to maximize design throughput and the memory bandwidth utilization.

4.1.1 Radix block

The radix block is used to perform a butterfly computation on the input samples. For example, the radix block for radix-4 FFT takes four input samples, performs the butterfly computation and then generates four results in parallel. Each radix block is composed of complex adders and subtractors. The structure of a radix block is determined by the FFT algorithm in use. Fig. 2a shows the structure of radix block for radix-4 FFT.

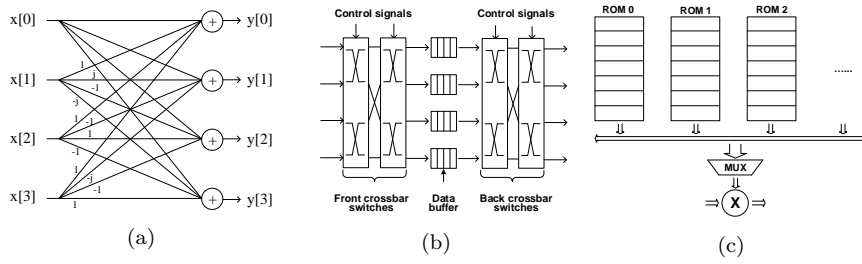


Fig. 2: (a) Radix-4 block (b) Datapath Permutation unit (c) Twiddle Factor Coefficient unit

4.1.2 Datapath Permutation (DPP) Unit

DPP unit is used for data permutation between butterfly computation stages in FFT. A DPP unit is composed of multiplexers and data buffers. In subsequent clock cycles, data from previous butterfly computation stage are first multiplexed and written into several data buffers. Each stored data element will be buffered for a fixed number of clock cycles and then read out. The size of each data buffer depends on the ordinal number of its present butterfly computation stage and the FFT problem size. Outputs from data buffers will also be multiplexed and fed into the next butterfly computation stage. Fig. 2b shows the DPP unit used for a radix-4 based FFT design. Each DPP unit consists of eight 4-to-1 multiplexers and four data buffers. In each cycle, a data buffer may be read and written simultaneously using different addresses.

4.1.3 Twiddle Factor Coefficient (TFC) Unit

A TFC unit consists of two parts: the TFC generation logic and the complex number multiplier. As shown in Fig. 2c, the TFC generation logic includes several lookup tables (functional ROMs) for storing twiddle factor coefficients, where the data read addresses will be updated with the control signals. The size of each lookup table is determined by the ordinal number of its present butterfly computation stage and the FFT problem size. Each lookup table can be implemented using a BRAM or distributed RAM (dist. RAM) on FPGA [17]. Each complex number multiplier consists of four real number multipliers and two real number adders/subtractors.

The proposed 2D FFT architecture is shown in Fig. 3, in which a Controlling Unit (CU) and a permutation network are introduced. The permutation network is developed based on our work in [1]. The CU is responsible for reconfiguring the permutation network to achieve the dynamic data layout.

4.2 Baseline Architecture

In the baseline architecture, when performing column-wise 1D FFTs, memory address is increased with a stride equal to N for a FFT of problem size $N \times$

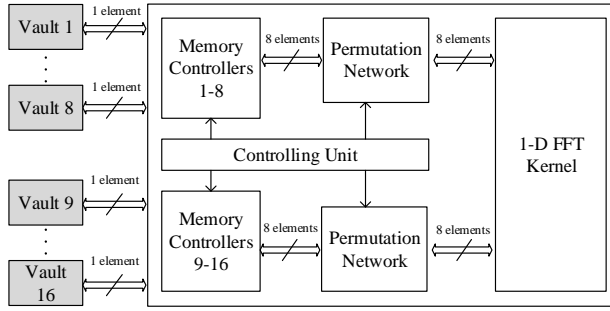


Fig. 3: 2D FFT Processor Architecture

N after each memory access. However, a minimum activate-to-activate delay ($t_{\text{diff-row}}$) exists when successively accessing two rows in the same bank or accessing two banks in the same vault ($t_{\text{diff-bank}}$ or $t_{\text{in-vault}}$). These delays result in a decline in 3D memory bandwidth utilization, and consequently the entire system throughput is affected.

4.3 Optimized Architecture

In the optimized architecture, the Controlling Unit (CU) is responsible for reconfiguring the permutation network dynamically to ensure data results of row-wise 1D FFTs are mapped onto the different vaults using the optimal dynamic data layout. Through this data remapping, activation of a different row will be only needed after several successive accesses on a row. This is in contrast to the baseline architecture, where every memory access results in a different row being activated. Thus, the impact of row activations on the entire system throughput is minimized. Furthermore, to reduce the overhead due to row activation, data inputs of several consecutive column-wise 1D FFTs are transferred from the 3D memory to the local memory on FPGA without waiting for the completion of the 1D FFT.

4.4 Optimal Dynamic Data Layouts

Our work in this paper is based on the dynamic data layouts (DDL) developed for the traditional 2D external memory in [7] where the data layout in memory is dynamically reorganized during computation. After reorganizations, non-unit stride accesses are converted to unit stride accesses, thereby reducing cache misses. The data layout is *optimal* from the performance point of view as it maximizes the memory bandwidth utilization. However, the data reorganization overhead regarding the latency and on-chip SRAM buffer consumption has not been considered. Figure 4 illustrate the key idea of our

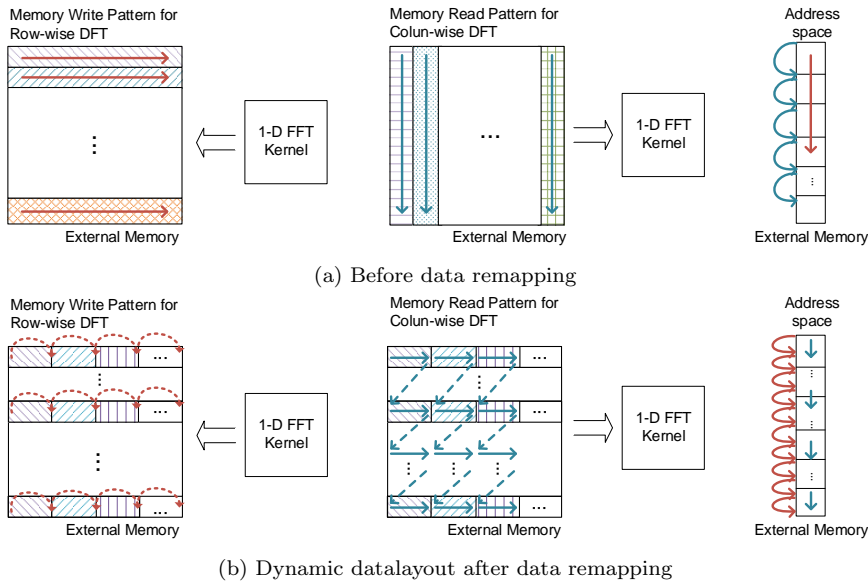


Fig. 4: Memory access patterns in 2-D FFT

proposed *optimal dynamic data layouts* in [18]. As shown in Figure 4, stride access is required for executing column-wise DFT without data remapping, as a result, the throughput for reading external memory declines significantly. By developing a data remapping approach, we proposed the *optimal dynamic data layouts* for 2D memory shown in Figure 4b such that peak memory bandwidth utilization is achieved with minimal data reorganization overhead. The data reorganization overhead is evaluated using the reorganizing latency and the on-chip buffer consumption. We further optimized our approach in [18] for the structure and organization of 3D memory based architecture. The data is distributed across all the vaults of the memory to increase the number of parallel accesses to the memory. The number of elements which can be accessed in one clock cycle of the FPGA depends on the data layout of the elements in the memory. This in turn affects the number of row activations and consequently the latency and throughput of the application. In the baseline architecture, row major order data layouts are employed. In our approach, instead of mapping results of row-wise FFTs to 3D memory in row major order, we employ block-based dynamic data layout, and the results are read block-by-block by the column-wise 1D FFT. The dynamic data layout is organized into blocks, each of size $w \times h$. w and h represent the width and height of the block, respectively. w is dynamically determined by the stride permutation to be performed in 1D FFTs. We denote the row buffer size in each bank of the 3D memory by s , the number of memory banks in each vault by b , the number of vaults

accessed in parallel by n_v . For a stride of value m , to achieve the *optimal dynamic data layout*, h is calculated based on the equation below:

$$h = \begin{cases} n_v & \text{if } 0 < m < sb \frac{t_{diff_bank}}{t_{diff_row}}; \\ n_v \cdot t_{diff_bank}/t_{in_row} & \text{if } sb \frac{t_{diff_bank}}{t_{diff_row}} \leq m < sb; \\ n_v \cdot t_{diff_row}/t_{in_row} & \text{if } m \geq sb. \end{cases} \quad (1)$$

The first condition in the Equation 1 represents the scenario where stride permutation results in accesses to the same row in a bank. This results in the lowest latency of t_{in_row} for each access. The second condition implies that the stride permutation causes different banks to be accessed and hence the latency for this access block of h is t_{diff_bank} . Lastly, the stride permutation is large enough so that the consecutive accesses are to the same bank and this is condition has the highest latency of t_{diff_row} . The value of h and hence the block size ($w = s/h$) varies with each of these conditions and results in the minimum number of row activations. It should be noted that we distribute the data across all banks in a vault and across all vaults. Since 3D memory allows for parallel access to multiple vaults, our approach maximizes the number of page hits and thereby extracts maximum bandwidth from the 3D memory.

4.5 Metrics of Evaluation

We evaluate the performance of 2D FFT on 3D memory integrated FPGA with respect to the metrics throughput and latency.

Throughput: defined as the maximum bandwidth of the memory supported by the application. It is measured in Giga Bytes per second (GB/s). Since our architecture is streaming data every cycle, the bandwidth at which memory operates determines the total execution time of the application. Therefore, higher the throughput, lower the execution time.

Latency: defined as the time elapsed between accessing first input from the memory and the time at which the first output is generated by the FFT kernel. We measure latency in the unit of ns. This penalty is paid just once and at the beginning of the processing. As we employ a streaming architecture, after the first output is generated, the subsequent outputs are produced in every cycle of operation of FPGA.

5 Experimental Results

The architecture parameters of the 3D memory used in model based performance analysis are 8 vaults (V), with 10 GB/s bandwidth per vault (BW_{vault}) and an overall 80 GB/s peak bandwidth ($V \times BW_{vault}$) for the entire 3D memory. The timing parameters of the 3D memory are: 5 ns (t_{in_row}), 8 ns (t_{diff_bank}) and 40 ns (t_{diff_row}). The 2D FFT architecture is implemented on state of the art FPGA Virtex 7 [17] and simulations were conducted using Xilinx Vivado 2014.2 development tools. Data parallelism of 32 is implemented for a data

Table 1: Throughput Comparison: Column-wise 1D FFT

	1024 × 1024 2D FFT	4096 × 4096 2D FFT	8192 × 8192 2D FFT
Throughput of column-wise FFT (Baseline)	0.8 GB/s	0.4 GB/s	0.4 GB/s
Peak bandwidth utilization	1.00%	0.5%	0.5%
Throughput of column-wise FFT (Optimized)	32 GB/s	25.6 GB/s	23.04 GB/s
Peak bandwidth utilization	40.0%	32.0%	28.8%

precision of 32 bits for each element. In order to give a thorough view of the performance of the 2D FFT implementation, we evaluate the system architecture combining the model based performance analysis of the 3D memory and the implementation of 2D FFT on FPGA.

Before evaluating the performance of the entire 2D FFT design, we estimate the throughput for the row-wise 1D FFT and column-wise 1D FFT separately. In the row-wise 1D FFT, the performance of the baseline architecture matches with that of the optimized architecture. This is attributed to the fact that the memory accesses of the baseline and optimized architectures are identical, i.e., sequential with the consecutive accesses to the same row. This results in the minimum number of row activations. In the second phase, i.e., column-wise 1D FFT, the difference in performance is significant. Table 1 shows the throughput performance of the 3D memory for baseline and optimized architectures. Table 1 shows that the column major access pattern and the row major data layout result in a very low throughput of a meager 0.8 GB/s for $N = 1024$ for the baseline architecture. With the optimized architecture of dynamic data layouts, for the same problem size, the 3D memory is able to sustain the bandwidth requirement of FFT kernel on FPGA and results in a throughput of 32 GB/s. It is also seen that the performance for column-wise FFT decreases with a larger problem size. This is because, with a problem size of $N = 1024$ and a row size of 16384 bits, 8 elements of 32 bits can be accessed in a time period of $t_{\text{diff-row}}$ from 8 vaults. With $N = 4096$ or $N = 8192$, only 4 elements can be accessed from 8 vaults in the same time period due to row major data layout. This problem is mitigated with the use of dynamic data layouts and the performance drop in optimized architecture is due to the drop in frequency of operation of FPGA. Through the proposed optimization, the peak bandwidth utilization is improved to 40.0%, 32.0%, and 28.8% for 1024×1024 , 4096×4096 and 8192×8192 size 2D FFTs respectively.

Table 2 presents the throughput and latency performance comparison between the baseline architecture and the optimized architecture for the complete 2D FFT application. The throughput of the baseline architecture is effectively halved due to the performance in column-wise 1D FFT phase. Using our proposed dynamic data layouts, the performance of the optimized architecture re-

Table 2: Performance Comparison: Complete 2D FFT application

FFT size	Baseline architecture			Optimized architecture			
	Throughput (GB/s)	Latency	Data Parallelism # elements	Throughput (GB/s)	Latency	Data Parallelism # elements	Performance improvement (throughput)
1024×1024	16.4	1.60 ms	32	32.0	524 μ s	32	95.1%
4096×4096	13.0	7.48 ms	32	25.6	2.4 ms	32	97.0%
8192×8192	11.7	145.4 ms	32	23.0	46.6 ms	32	96.6%

mains the same in both the phases and hence, the overall throughput remains high as shown in Table 2. The optimized architecture achieves a throughput of 32.0, 25.6 and 23.0 GB/s compared with the 16.4, 13.0, 11.7 GB/s performance of the baseline architecture for 1024×1024 , 4096×4096 and 8192×8192 problem sizes, respectively. The throughput performance is improved by 95.1%, 97.0%, 96.6% for 1024×1024 , 4096×4096 and 8192×8192 point 2D FFT, respectively. Similar analysis of the performance in terms of latency reveals that the latency of the optimized architecture is reduced by up to $3\times$ by using our proposed optimizations. Comparing the results of the throughput in the 1D FFT kernel and the entire 2D FFT architecture, we observe that the optimization for 3D memory access makes a significant contribution to the performance improvement. Moreover, the sustained throughput of the optimized 2D FFT architecture achieves up to 40% of the *peak memory bandwidth*, which is an upper bound on the performance of the chosen FFT algorithm and 3D system architecture.

6 CONCLUSION

In this paper, we proposed dynamic data layout optimizations to obtain a high throughput 2D FFT architecture on 3D memory integrated architecture. The proposed architecture achieves high throughput by maximizing and balancing the bandwidth between the external memory and FFT kernel on FPGA. By proposing the dynamic data layouts realized with the on-chip permutation network, the delay caused due to row activation overhead is highly reduced, thus leading to significant performance improvement. The experimental results comparing with the baseline architecture show that our implementation outperforms in throughput and latency. In the future, we plan to build a design framework targeted at throughput-oriented signal processing kernels, which enables automatic data layout optimizations addressing new 3D memory technologies.

References

1. Ren Chen and Viktor K. Prasanna. Energy and Memory Efficient Bitonic Sorting on FPGA. In *Proc. of ACM/SIGDA FPGA*, pages 45–54, 2015.

2. Ren Chen and Viktor K. Prasanna. Automatic generation of high throughput energy efficient streaming architectures for arbitrary fixed permutations. In *Proc. of IEEE Conference on Field Programmable Logic and Applications (FPL)*, pages 1–8. IEEE, 2015.
3. Berkin Akin, Peter A. Milder, Franz Franchetti, and James C. Hoe. Memory Bandwidth Efficient Two-Dimensional Fast Fourier Transform Algorithm and Implementation for Large Problem Sizes. In *Proc. of IEEE International Symposium on Field-Programmable Custom Computing Machines (FCCM '12)*,, pages 188–191, April 2012.
4. Ren Chen and Viktor K. Prasanna. Energy Efficient Parameterized FFT Architecture. In *Proc. of IEEE International Conference on FPL*, 2013.
5. Jung Sub Kim, Chi-Li Yu, Lanping Deng, Srinidhi Kestur, Vijaykrishnan Narayanan, and Chaitali Chakrabarti. FPGA Architecture for 2D Discrete Fourier Transform based on 2D decomposition for large-sized data. In *Proc. of IEEE Workshop on Signal Processing Systems*, pages 121–126, Oct 2009.
6. Hybrid Memory Cube Consortium. Hybrid Memory Cube Specification. http://hybridmemorycube.org/files/SiteDownloads/HMC_Specification%201_0.pdf.
7. Neungsoo Park and Viktor K. Prasanna. Dynamic Data Layouts for Cache-Conscious Implementation of a Class of Signal Transforms. *IEEE Transactions on Signal Processing*, 52(7):2120–2134, July 2004.
8. Wendi Wang, Bo Duan, Chunming Zhang, Peiheng Zhang, and Ninghui Sun. Accelerating 2D FFT with Non-Power-of-Two Problem Size on FPGA. In *Proc. of IEEE International Conference on Reconfigurable Computing and FPGAs (ReConFig '10)*, pages 208–213, Dec 2010.
9. Berkin Akin, Franz Franchetti, and James C. Hoe. Understanding the Design Space of Dram-optimized Hardware FFT Accelerators. In *Application-specific Systems, Architectures and Processors (ASAP), 2014 IEEE 25th International Conference on*, pages 248–255. IEEE, 2014.
10. Hong Ren Wu and Frank John Paoloni. The Structure of Vector Radix Fast Fourier Transforms. *IEEE Transactions on Acoustics, Speech and Signal Processing*, 37(9):1415–1424, Sep 1989.
11. Qiuling Zhu, Berkin Akin, H Ekin Sumbul, Fazle Sadi, James C Hoe, Larry Pileggi, and Franz Franchetti. A 3D-Stacked Logic-in-Memory Accelerator for Application-Specific Data Intensive Computing. In *Proc. of IEEE International Conference on 3D Systems Integration Conference (3DIC)*, pages 1–7. IEEE, 2013.
12. Peter Gadfort, Aravind Dasu, Ali Akoglu, Yoon Kah Leow, and Michael Fritze. A Power Efficient Reconfigurable System-in-Stack: 3D Integration of Accelerators, FPGAs, and DRAM. In *Proc. of IEEE International Conference on System-on-Chip Conference (SOCC)*, pages 11–16. IEEE, 2014.
13. Shreyas G. Singapura, Anand Panangadan, and Viktor K. Prasanna. Towards Performance Modeling of 3D Memory Integrated FPGA Architectures. In *Proc. of International Conference on Applied Reconfigurable Computing*. 2015.
14. Shreyas G. Singapura, Anand Panangadan, and Viktor K. Prasanna. Performance Modeling of Matrix Multiplication on 3D Memory Integrated FPGA. In *Proc. of 22nd Reconfigurable Architectures Workshop, IPDPDS*. 2015.
15. Ren Chen and Viktor K. Prasanna. Energy-efficient Architecture for Stride Permutation on Streaming Data. In *Proc. of IEEE Conference on ReConFig*, pages 1–7, 2013.
16. Ren Chen, Neungsoo Park, and Viktor K. Prasanna. High Throughput Energy Efficient Parallel FFT Architecture on FPGAs. In *Proc. of IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–6. IEEE, 2013.
17. Virtex-7 FPGA Family. <http://www.xilinx.com/products/virtex7>.
18. Ren Chen and Viktor K. Prasanna. DRAM Row Activation Energy Optimization for Stride Memory Access on FPGA-based Systems. In *Proc. of International Conference on Applied Reconfigurable Computing*. 2015.
19. Stefan Langemeyer, Peter Pirsch, and Holger Blume. Using SDRAMs for Two-Dimensional Accesses of Long $2^n \times 2^m$ -point FFTs and Transposing. In *Proc. of International Conference on Embedded Computer Systems (SAMOS '11)*, pages 242–248, July 2011.