

# TensorLib: A Spatial Accelerator Generation Framework for Tensor Algebra

Liancheng Jia<sup>1</sup>, Zizhang Luo<sup>1</sup>, Liqiang Lu<sup>1</sup>, and Yun Liang<sup>1,2</sup>

<sup>1</sup>Center for Energy-Efficient Computing and Applications, School of EECS, Peking University

<sup>2</sup>Pengcheng laboratory, China

{jlc, semiwaker, liqianglu, ericyun}@pku.edu.cn

**Abstract**—Tensor algebra finds applications in various domains, and these applications, especially when accelerated on spatial hardware accelerators, can deliver high performance and low power. Spatial hardware accelerator exhibits complex design space. Prior approaches based on manual implementation lead to low programming productivity, rendering thorough design space exploration impossible. In this paper, we propose TensorLib, a framework for generating spatial hardware accelerator for tensor algebra applications. TensorLib is motivated by the observation that, different dataflows share common hardware modules, which can be reused across different designs. To build such a framework, TensorLib first uses Space-Time Transformation to explore different dataflows, which can compactly represent the hardware dataflow using a simple transformation matrix. Next, we identify the common structures of different dataflows and build parameterized hardware module templates with Chisel. Our generation framework can select the needed hardware modules for each dataflow, connect the modules using a specified interconnection pattern, and automatically generate the complete hardware accelerator design. TensorLib remarkably improves the productivity for the development and optimization of spatial hardware architecture, providing a rich design space with trade-offs in performance, area, and power. Experiments show that TensorLib can automatically generate hardware designs with different dataflows and achieve 21% performance improvement on FPGA compared to the state-of-the-arts.

## I. INTRODUCTION

Tensor algebra is a prevalent tool and has been successfully applied to a broad range of applications such as machine learning, data analytics. Tensor algebra features different dimensions, sizes and computation patterns, which require special hardware acceleration. For example, the 2-D convolution is one of the most popular tensor operations in deep learning applications [10]. It involves a 4-D weight and 3-D input and requires to accumulate the partial sums in four dimensions [10]. MTTKRP is a widely used tensor operation for tensor factorization in recommendation systems, which takes one 3-D tensor and two matrices as inputs and generates a matrix. Due to the regular computation patterns, spatial hardware accelerators are commonly used for the acceleration of tensor algebra [3], [9], [16]–[18].

The majority of the spatial accelerator designs follow a hierarchical architecture [13]. As shown in Figure 1, spatial

Yun Liang is the corresponding author.

accelerators usually consist of an array of homogeneous processing elements (PEs), an on-chip network that connects PEs together, a shared scratchpad buffer, and a system controller. The array of PEs can provide huge parallelism, and the connection between PEs can exploit different types of data reuse. While most spatial accelerators adopt the same hierarchical architecture, the actual implementation of each design can vary a lot. Among the various design parameters, hardware dataflow plays the most important role as it determines how the tensor is computed and communicated between PEs. There exists a sufficient large design space of dataflows for the spatial hardware accelerator design. Initially, dataflow is categorized by specifying the tensors that are reused temporarily inside each PE. For example, [16] uses output stationary systolic array dataflow because the output tensor elements stay inside PE during execution. Similarly, [9] uses weight stationary dataflow and [3] uses row stationary dataflow.

The complex structure of spatial accelerator and its large design space leads to low productivity, rendering thorough design space exploration impossible. To improve the programming productivity, HLS tools have been used for accelerator designs, which supports hardware generation with software-style programming [4], [16]. Some recent works also design DSLs or other notations to represent the dataflow and hardware architecture of spatial accelerators [11]–[14]. The user can use DSL to express the architecture and the compiler generates hardware code in HLS. However, the auto-generated HLS code is hard to optimize, resulting in low performance. Recent advances in hardware programming introduce highly parameterized and modular design principles based on Scala and Python [1]. The high-level language can be equally expressive as Verilog which supports cycle-level RTL description, but they also support functional and object-oriented programming. The high-level programming features enable a large variety of hardware instances to be generated with the same hardware template, which is extremely useful for spatial accelerators with a variety of dataflows.

In this paper, we propose TensorLib, a framework for generating spatial accelerator for tensor algebra. We use Space-Time Transformation (STT) [2] as a means of expression to represent the hardware dataflow. STT maps loop instances to hardware spatially (coordinates in the PE array) and temporally (timestamp of execution) and it can cover the complete

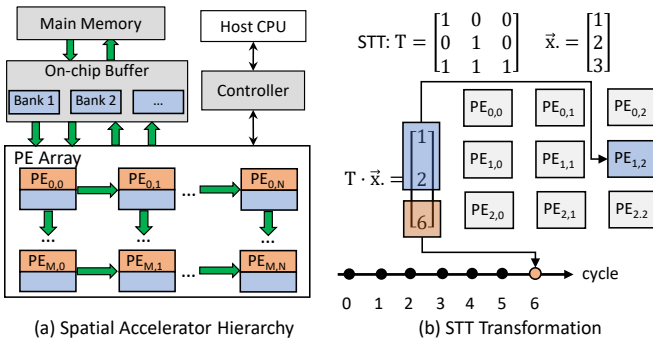


Fig. 1: Spatial accelerator and Space-Time Transformation

design space of spatial dataflows with a linear transformation. By inferring the pattern of data reuse using STT, we can directly represent various dataflows used in spatial accelerators including unicast, systolic, multicast, stationary, etc.

We observe that different hardware dataflows can share common hardware modules at each level of the hierarchical structure. In other words, a new dataflow can be derived from an existing dataflow by only modifying certain parts. Based on these findings, we develop a few basic hardware component templates for each level of hierarchy using Chisel [1], facilitating hardware reuse of different dataflows. When a dataflow is specified by STT, TensorLib can automatically select the templates and connect them together to build the hardware architecture. Some recent HLS frameworks also apply STT [4], [8], [11], but they are limited in both generability and performance. First, their proposals mainly target systolic architectures which do not cover the complete space of dataflows, and only support limited applications. Second, the tensor algorithm is tightly coupled with architecture in their frameworks, making it hard to optimize. We overcome these challenges by providing a thorough dataflow analysis based on STT and building highly parameterized hardware module templates. Overall, TensorLib makes the following contributions,

- We propose TensorLib, a framework to automatically generate various hardware dataflow implementation of spatial accelerators for tensor algebra applications.
- We develop a formal representation of spatial hardware dataflows using space-time transformation which can cover a comprehensive design space of dataflows.
- We integrate the STT-based dataflow analysis with Chisel-based hardware template to build a accelerator generator with rich design space, high productivity and good performance.

Experiment shows that TensorLib is able to generate a large number of dataflow architectures for various tensor applications. Compared to the state-of-the-art, Tensorlib achieves 21% throughput improvement on FPGA for matrix multiplication benchmark, and supports more dataflows and applications, which remarkably improves both performance and generality. The source code of Tensorlib is available at <https://www.github.com/kirliavc/tensorlib>.

## II. SPACE-TIME TRANSFORMATION BACKGROUND

Space-time transformation (STT) is a linear transformation that maps the tensor algebra to hardware from a spatial and temporal perspective [2]. STT can be used for tensor algebra whose computation can be described in perfect nested loops. The PE array can be viewed as a hypercube, and the execution of hardware can be identified as a space vector and a time scalar indicating where and when the computation takes place. STT transforms a point in the loop nest to the space-time vector in hardware execution using a matrix multiplication operation. For example, given a loop iteration in the loop nest  $\vec{x} = [i, j, \dots]^T$  and a transformation matrix  $T$ , the execution space and time can be calculated as follows,

$$\begin{bmatrix} \vec{p} \\ t \end{bmatrix} = T\vec{x} \quad (1)$$

where space vector  $\vec{p}$  means the PE coordinates inside the PE array and time scalar  $t$  means the time step of execution. Figure 1 (b) shows an example using matrix multiplication  $C[i, j]_+ = A[i, k] \times B[k, j]$  as the target tensor algebra. There are three iterator variables  $i, j$  and  $k$ . When  $i = 1, j = 2, k = 3$ , using equation 1 we can get  $\begin{bmatrix} \vec{p} \\ t \end{bmatrix} = (1, 2, 6)^T$ , which means  $A[1, 3] \times B[3, 2]$  takes place at PE (1, 2) at the sixth cycle. Since a PE can only perform one operation per cycle, matrix  $T$  must be a full-ranked matrix so that there is an one-to-one mapping between iterator space and space-time space.

## III. FRAMEWORK

Figure 2 presents the overview of TensorLib framework. The workflow can be divided into two steps: dataflow generation and hardware implementation generation. For the dataflow generation, TensorLib uses the tensor algebra described by a nested loop and an STT matrix ( $T$  in Equation (1)) as input. The framework first maps the nested loop into the spatial hardware PE array using STT. After that, by analyzing the recurrence of the same tensor element in different space-time vectors, TensorLib determines the reuse pattern and dataflow type of each tensor. The computation involves multiple tensors and STT generates different dataflows for each tensor. We categorize the dataflow for each tensor into five types: Unicast, Stationary, Systolic, Multicast (Reduction Tree for output) and 2D-reuse. Details are discussed in Section IV.

The second step is hardware generation. The framework generates the 3-level hierarchy of spatial accelerators in a bottom-up manner. It first uses the dataflow type of each tensor to select the internal modules of PE and connect them with the computation IP to form the PE structure. Next, it connects the PEs together with the generated patterns to form the PE array. Finally, it generates the memory modules with access pattern and external memory IPs, and the controller which provides control signals for both PE and memory ports. Details are discussed in section V.

## IV. DATAFLOW GENERATION

To represent a dataflow, it is critical to capture the movement of each tensor. Given a loop iteration  $\vec{x}$ , the tensor index

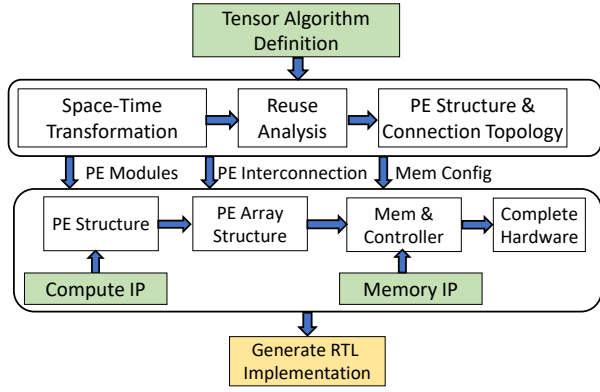


Fig. 2: The overview of TensorLib

accessed by  $\vec{x}$  can be expressed as  $\vec{I} = A\vec{x}$ , where  $A$  is the corresponding access matrix. The mapping between tensor index and loop iteration index can be 1-to-N so that the same tensor element is accessed by multiple loop iterations. By transforming the loop iterations to the space-time domain, the reuse of the same tensor element forms the *reuse hyperplane* in space-time space, which determines how the tensor moves (e.g. dataflow). By replacing the  $\vec{x}$  in Equation 1 with accessing matrix, we have,

$$AT^{-1} \begin{bmatrix} \vec{p} \\ t \end{bmatrix} = \vec{I} \quad (2)$$

Based on this, we can build a reuse subspace composed of all the space-time points that reuse the same tensor index. Next, based on the rank of the reuse subspace, we divide the reuse shape into three cases as shown in Table I. Firstly, if the rank is 0, the hyperplane is actually a single point, which means the tensor element only appears once in the whole computation without any reuse. The unicast dataflow requires large on-chip bandwidth since every PE reads data from the on-chip memory independently.

If the rank equals to 1, the hyperplane is a straight line, and the dataflow is determined by the direction of the line. We can calculate  $d\vec{p}$  and  $dt$  easily using the base vector of the hyperplane, as showed by Equation 3, where  $Eig$  is the first eigenvector,  $E$  is the identity matrix,  $(AT^{-1})^{-}$  is the pseudoinverse of  $AT^{-1}$ .

$$\begin{bmatrix} d\vec{p} \\ dt \end{bmatrix} = Eig(E - (AT^{-1})^{-}(AT^{-1})) \quad (3)$$

There are totally three sub-cases:

- $d\vec{p} = \vec{0}, dt \neq 0$ . The same tensor element appears in the same PE, but at different time steps, which means that the tensor is stationary and stays inside one PE.
- $d\vec{p} \neq \vec{0}, dt \neq 0$ . The tensor element appears in different PEs and at different time steps. This corresponds to systolic dataflow where data are delayed for one cycle before being sent out to other PE.
- $d\vec{p} \neq \vec{0}, dt = 0$ . The tensor element appears in different PEs but at same time step. The same data must be transformed to different PEs at the same time, which

TABLE I: Dataflow analysis with STT

Subspace Dimension	Shape	Space-Time Reuse Space	Tensor Dataflow
0	-		Unicast
1	$d\vec{p} = 0, dt \neq 0$		Stationary
	$d\vec{p} \neq 0, dt \neq 0$		Systolic
	$d\vec{p} \neq 0, dt = 0$		Multicast
2	t-axis Vertical		Broadcast
	t-axis Parallel		Multicast & Stationary
	t-axis Intersect		Systolic & Multicast

corresponds to a multicast dataflow. If it is a output tensor, the dataflow indicates that different partial results is generated simultaneously by different PEs, which requires a reduction tree to generate final results.

Finally, if the rank is 2, the reuse subspace is a 2D plane. In this case, the shape of reuse space is categorized based on the relationship with time axis. There are totally three sub-cases:

- Vertical to t-axis. The same tensor element is broadcasted to all the PEs in the array at the same cycle.
- Parallel to t-axis. The tensor element is firstly broadcasted to a group of PEs, and the element keeps stationary in the PE during the execution stage.
- Intersect with t-axis. The tensor element is firstly broadcasted to a group of registers, and then traverses between PE in a systolic style.

The PE array dataflow and interconnection pattern of 2D reuse subspace is similar to the 1D case because both of them are formed with 1D reuse patterns (systolic, multicast and stationary). Finally, typical tensor algebras such as 2D Convolution contain more than three loop nests. For a 2D PE array, we need to select three loops (2D space + 1 time) from the loop nest to map to space-time space. The remaining loops are executed sequentially which doesn't influence PE dataflow. When PE and memory sizes are determined, the loops are performed tiling to fit the hardware resources.

Using the example in Figure 1 (b), the accessing matrix for  $A[i, k]$  is  $\begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}$ . With Equation (3), we can find that the dimension of reuse subspace is 1 and  $\begin{bmatrix} d\vec{p} \\ dt \end{bmatrix} = (0, 1, 1)^T$ , which means tensor A uses the systolic dataflow with vertical direction.

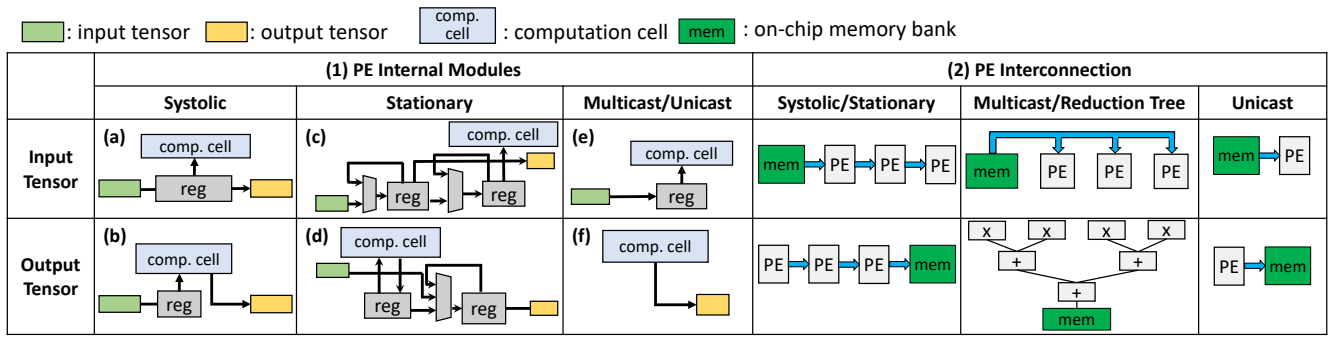


Fig. 3: The PE internal modules and interconnection modules for different dataflows

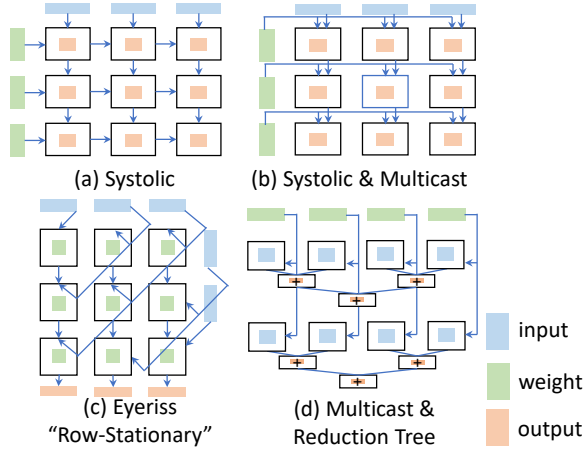


Fig. 4: Examples of PE interconnection for different dataflows

## V. HARDWARE GENERATION

After dataflow generation, the next step is to implement the architecture with the specific dataflow. The architecture of spatial accelerators can be separated into three levels: PE structure, PE-Array interconnection, on-chip shared buffer and controller. In this section, we discuss how to generate each level of hardware architecture for different dataflows.

### A. PE Generation

The PE structure consists of the computation cell and the internal modules that connect the cell with I/O ports. The computation cells are manually implemented IPs. The difference in dataflows mainly lies in the internal modules that connect the I/O ports and the computation cells. Since each tensor can use different dataflow and the internal modules for each tensor do not connect to each other, they can be implemented independently. As shown in Table 1, the dataflow of each tensor inside PE has three types: (a) systolic, (b) stationary and (c) multicast or reduction tree. Each tensor can be either input or output of the algorithm. The implementation is different for input and output tensor because the output tensor needs to read results from the computation cell, so there are 6 conditions in total. In TensorLib framework, we build the implementation of PE internal modules for each condition, and the circuit diagram is presented in Figure 3.

- Modules (a) and (b) are designed for systolic dataflows. Tensor elements always transfer to their neighboring PEs every cycle. The difference is that the output data is generated from the computation cell and the input data is transferred to the next PE directly.
- Modules (c) and (d) are designed for stationary tensors. Tensor elements stay inside PE during execution, but the data needs to be updated when the execution stage ends. It uses a double-buffer structure to enable the parallelism of computation and data communication. In (d), one register is used to update the result of the current stage, and the other is used to transfer the results of the last stage.
- Modules (e) and (f) are designed for multicast and unicast dataflows, where data are received and passed directly.

Different PE dataflow can be constructed by selecting the corresponding modules. For example, output stationary [16] dataflow contains two modules (a) (systolic, input) and one module (d) (stationary, output). Weight stationary [9] dataflow contains one (a), one (b) (systolic, output) and one (c) (stationary, input). Eyeriss [3] dataflow uses one multicast and one stationary module for input, and systolic module for output. The selected components are connected with the PE body and computation cell components to generate the complete PE structure.

### B. Interconnection and Memory Generation

Different dataflows also require different connection topologies between PEs. Figure 3 (2) shows the PE interconnection patterns for different dataflows. The systolic and stationary dataflows connect adjacent PEs together. The direction of systolic dataflow interconnection is determined by the reuse vector  $(d\vec{p}, dt)$ . The output of  $PE(x, y)$  connects to  $PE(x + dx, y + dy)$  after delaying  $dt$  cycle. For multicast input dataflow, the same input data is transferred from the on-chip buffer to PEs in a row directly at the same cycle. The output multicast dataflow is implemented with a reduction tree to perform a reduction on the output of PEs. For unicast dataflow, different PEs are totally independent and connect to the on-chip memory bank directly. The next step is on-chip memory generation. Each group of PEs that reuse the same tensor indexes is assigned with a particular memory bank. TensorLib automatically generates the memory module and connects to the PE array with the pattern shown in Figure





TABLE III: FPGA Performance Comparison on MM workload

Device	Susy [11]		PolySA [4]		TensorLib	
	Arria-10		VU9P		VU9P	
Workload	MM	Conv	MM	Conv	MM	Conv
LUT	40%	35%	49%	49%	68%	73%
DSP	93%	84%	89%	89%	75%	75%
BRAM	32%	30%	89%	71%	51%	73%
MHz	202	220	229	229	263	245
Gop/s	547	551	555	548	673	626




Fig. 6: Power and area result of different dataflow designs

INT16 datatype. We set the target frequency to 320MHz for ASIC synthesis. Figure 6 gives the area and energy performance of different dataflow architectures generated by TensorLib. Each point in the graph refers to one dataflow, and there are totally 148 points for GEMM and 33 points for depthwise-Conv2D. Experiments show that dataflow choice has a larger impact on energy consumption than area. The energy variation of GEMM range from 35mW to 63mW, which shows 1.8X difference, while the area has only 1.16X difference. Compared with other dataflows, dataflow with two multicast input (MMT, MMS) consumes more energy. However, reduction tree output dataflow doesn't cost too much energy, although they have similar STT-level representation. Dataflows with stationary tensor also consume more area and energy because of the control signals for stationary data.

### C. Comparison with Prior works

We also compare TensorLib with PolySA [4] and Susy [11]. PolySA and Susy can generate the hardware implementation of systolic arrays with a polyhedral model or STT. However, their designs are limited in performance and supported algorithms. Both Susy and PolySA only support systolic array dataflow, and they fail to generate hardware for algorithms that don't fit well in systolic architecture, such as Depthwise convolution. We synthesize our implementation with systolic array (KCX-STs) dataflow on MM and Conv2D workloads with FP32 datatype. The size of PE array in our FPGA implementation is  $10 \times 16$  and the vectorization degree in each PE is 8. The result is shown in Table III. Our synthesis result achieves 673 Gop/s throughput and 263 MHz frequency, which obtains 21% throughput and 15% frequency improvement compared with the state-of-the-art generators.

Recently, AutoSA [15] extends PolySA with I/O optimization and more tensor algorithms. It also uses Autobridge [7] to improve the frequency by assigning each hardware module to one particular FPGA slot and minimizing the slot-

crossing cost. Inspired by this, we also manually optimize the physical placement using Vivado toolchain. This will improve the frequency of MM design to 328 MHz on VU9P.

## VII. CONCLUSION

In this paper, we propose TensorLib, a framework for generating spatial accelerator for tensor algebra. TensorLib uses STT to analyze the tensor reuse behavior to generate the dataflow type and communication direction for each tensor. We build reusable hardware module templates for each dataflow and Tensorlib automatically select hardware modules to construct PE structure, PE interconnection, on-chip memory and controller for complete spatial accelerator. Experiment results show that TensorLib can generate various accelerators for tensor applications with different dataflows, and achieve better performance than state-of-the-art generators on FPGA.

## ACKNOWLEDGMENT

This work was supported in part by the Beijing Natural Science Foundation (No. JQ19014), Beijing Academy of Artificial Intelligence (BAAI), and Key-Area Research and Development Program of Guangdong Province (No. 2019B010155002).

## REFERENCES

- [1] Jonathan Bachrach et al. Chisel: constructing hardware in a scala embedded language. In *DAC*, 2012.
- [2] Donald G. Batus et al. Efficient exploration of nonuniform space-time transformations for optimal systolic array synthesis. In *ASAP 1993*.
- [3] Yu-Hsin Chen et al. Eyeriss: A spatial architecture for energy-efficient dataflow for convolutional neural networks. In *ISCA 2016*.
- [4] Jason Cong and Jie Wang. Polysa: polyhedral-based systolic array auto-compilation. In *ICCAD 2018*.
- [5] Zidong Du et al. Shidiannao: shifting vision processing closer to the sensor. In *ISCA 2015*.
- [6] Hasan Genc et al. Gemmini: An agile systolic array generator enabling systematic evaluations of deep-learning architectures. *arXiv*, 2019.
- [7] Licheng Guo et al. Autobridge: Coupling coarse-grained floorplanning and pipelining for high-frequency HLS design on multi-die fpgas. In *FPGA*, 2021.
- [8] Liancheng Jia et al. Generating systolic array accelerators with reusable blocks. *IEEE Micro*, 2020.
- [9] Norman P Jouppi et al. In-datacenter performance analysis of a tensor processing unit. In *ISCA*, 2017.
- [10] Alex Krizhevsky et al. Imagenet classification with deep convolutional neural networks. In *NIPS 2012*.
- [11] Yihsiang Lai et al. Susy: A programming model for productive construction of high-performance systolic arrays on fpgas. In *ICCAD 2020*.
- [12] Yihsiang Lai et al. Heterocl: A multi-paradigm programming infrastructure for software-defined reconfigurable computing. In *FPGA*, 2019.
- [13] Liqiang Lu et al. Tenet: A framework for modeling tensor dataflow based on relation-centric notation. In *ISCA*, 2021.
- [14] Nitish Kumar Srivastava et al. T2s-tensor: Productively generating high-performance spatial hardware for dense tensor computations. In *FCCM*, 2019.
- [15] Jie Wang et al. AutosA: A polyhedral compiler for high-performance systolic arrays on fpga. In *FPGA*, 2021.
- [16] Xuechao Wei et al. Automated systolic array architecture synthesis for high throughput CNN inference on fpgas. In *DAC*, 2017.
- [17] Qingcheng Xiao et al. Fcnlib: An efficient and flexible convolution algorithm library on fpgas. In *DAC*, 2020.
- [18] Qingcheng Xiao et al. Hasco: Towards agile hardware and software co-design for tensor computation. In *ISCA*, 2021.