Increasing FPGA Accelerators Memory Bandwidth with a Burst-Friendly Memory Layout

Corentin Ferry*†, Tomofumi Yuki*, Steven Derrien* and Sanjay Rajopadhye†
*Univ Rennes, CNRS, IRISA, Inria. Rennes, France.
†*Department of Computer Science*, *Colorado State University*. Fort Collins, Colorado, USA

arXiv:2202.05933v2 [cs.AR] 23 Feb 2022

***Abstract*—Offloading compute-intensive kernels to hardware accelerators relies on the large degree of parallelism offered by these platforms. However, the effective bandwidth of the memory interface often causes a bottleneck, hindering the accelerator's effective performance. Techniques enabling data reuse, such as tiling, lower the pressure on memory traffic but still often leave the accelerators I/O-bound. A further increase in effective bandwidth is possible by using burst rather than element-wise accesses, provided the data is contiguous in memory.**

In this paper, we propose a memory allocation technique, and provide a proof-of-concept source-to-source compiler pass, that enables such burst transfers by modifying the data layout in external memory. We assess how this technique pushes up the memory throughput, leaving room for exploiting additional parallelism, for a minimal logic overhead.

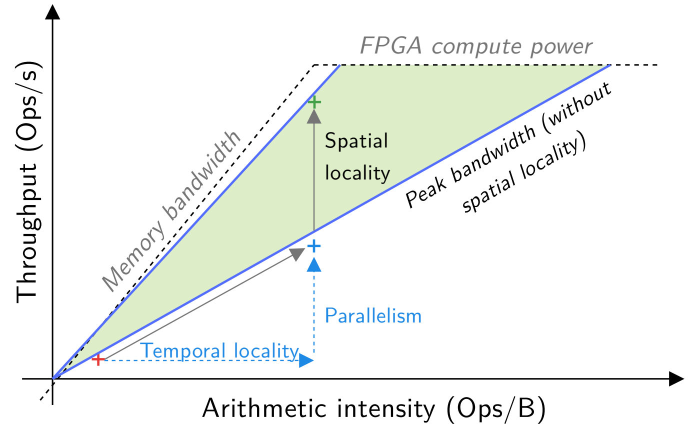

Fig. 1: Improving both temporal and spatial locality is needed to get a higher compute throughput. Limited amounts of on-chip memory means that temporal locality cannot be improved beyond a certain limit, therefore using spatial locality to increase memory bandwidth becomes necessary.

I. Introduction

Hardware acceleration of compute-intensive mathematical kernels is one of the most efficient ways to improve their performance. Compute-intensive algorithms, such as image processing algorithms or neural networks, can benefit greatly from the use of the massive parallelism and low latency of application-specific hardware. There are however significant hurdles preventing widespread use of such hardware: aside from the manufacturing cost of an application-specific integrated circuit (ASIC), which a programmable chip such as a field-programmable gate array (FPGA) can mitigate to some extent, the development effort needed to get a decently performing accelerator is orders of magnitude greater than developing a GPU version of the same algorithm.

Design automation tools have the ability to create massively-parallel accelerators. The result has so much processing power that memory accesses are not fast enough to feed the accelerator, which is then called *memory-bound* (the opposite situation, where the processing power is the bottleneck, being *compute-bound*). Cong et al. [1] emphasize that FPGAs have as much computing power as GPUs, but not the memory bandwidth. One way to address this *memory wall* [2] is to increase the *arithmetic intensity*, by either computing more for the same memory traffic or transferring less for the same amount of computations. Even so, memory-bound accelerators are still commonplace. Random memory accesses are among the primary causes of the memory wall, that lowers the effective bandwidth as illustrated by Figure 1.

Favoring burst accesses over random accesses is a key to increasing the the effective bandwidth, and therefore the performance of the accelerator. Doing so implies transforming the accelerator's memory access pattern; to be profitable, such a transformation must not degrade the on-chip compute power.

In this work, we propose a transformation that improves spatial locality of programs with uniform memory accesses, and integrate it into a polyhedral compiler flow. Broad classes of actively studied programs fall into this category, notably convolutions, FDTD or matrix multiplication. The idea is to transform the off-chip data layout and get as many contiguous off-chip accesses as possible, while preserving the compute-oriented on-chip optimizations.

The contributions of this work are:

- A memory layout and access pattern that favor burst accesses,
- A proof-of-concept compiler pass that automatically applies this memory layout,
- An evaluation of the performance of this layout, that shows it can fully utilize the available bandwidth, without a significant increase in area.

Our work relies on well-known techniques, such as iteration space tiling [3], [4], and tools, such as the Integer Set Library (ISL) [5] to represent polyhedral sets and Pluto [6] to apply iteration space tiling and generate a schedule.

This paper is organized as follows: first, it introduces the notion of data locality and on-chip parallelism-enabling techniques onto which we rely (Sections II and III); then, it

describes the construction of our burst-friendly off-chip layout (Section IV). Finally, it provides an evaluation and comparison with the state of the art (Section VI).

II. Background

This section explains the issue of memory bandwidth under-utilization, and the challenges that one faces to improve its use.

A. Increasing compute performance makes the accelerator memory-bound

Fine-tuning programs to use architectural optimizations (or architectural features on CPUs/GPUs) makes the design memory-bound.

In order to get the maximum performance with a given architecture, a program needs to fully exploit its hardware compute resources. This typically means using parallel execution units. Vector units are designed specifically for this purpose on CPUs and GPUs, with dedicated register files. FPGA and ASIC designers have to build their own on-chip architecture where processing units and memories can be actively used at all times.

High-level synthesis tools also have the ability to perform loop unrolling, akin to automatic loop vectorization for CPUs which has been part of state-of-the-art compilers for years [7], in order to exhibit more parallel computations. Loop pipelining, also available as a loop transformation in HLS tools, increases the compute throughput by ensuring all functional units are kept busy at all times. Combinations of these techniques [8], [9] yield drastic performance increase over non-parallel implementations.

Increasing parallelism often also increases memory traffic: to perform more operations at a time, an accelerator may need more data. Increasing parallelism without increasing arithmetic intensity makes the accelerator memory-bound, thus there is a need for memory optimizations such as data compression [10].

B. Improving locality increases memory access performance

To get a better performance, we need to reduce pressure on memory, and to do that, we need to improve both temporal and spatial locality.

A general rule that applies to all memory systems regardless of their hierarchy is that the further data is from the compute engine, the longer it takes to access it. Faster memory cells such as registers, close to the compute engines, are expensive and therefore scarce. Numerous program transformations seek to minimize the latency incurred by memory accesses, by improving the program's behavior in two ways:

- Minimizing cache misses. This means using data available at closer cache levels in priority, and computing as much as possible with it. As caches contain the most recently accessed data, this is called improving **temporal locality**.
- Accessing data in the same order as it is stored in memory, thereby exploiting the burst facilities offered by the memory controller and DRAM chips, resulting in a higher throughput. This is called improving **spatial locality**.

The *roofline model* [2], illustrated on Figure 1 is a visual way to see the effects of locality improvements. The throughput a system can reach is bound when either a *memory roofline* or a *compute roofline* is hit. The former case incurs an under-utilization of the on-chip parallel compute resources, which is undesirable, where the latter incurs idle memory bus time, which is good in regards to energy, and means no further memory improvements are needed. Figure 1 shows that improving both locality metrics is needed to go from the former to the latter case. The next subsections will go over locality-improving techniques and why further spatial locality improvement is needed.

C. Tiling loops to improve temporal locality

Temporal locality is obtained by transforming loops, in particular tiling them.

A common transformation to improve temporal locality is called loop tiling. Its principle is to break the workload of a loop (its *iteration space*) into smaller, single-shaped packets called *tiles*. Such tiles are small enough so that the amount of data required for them to execute fits in a local memory or cache. Tiling therefore increases temporal locality and arithmetic intensity.

The impact of tiling on performance is such that finding the best tile size and shape for specific classes of programs, such as matrix multiplication [11] or convolution [9], is a trending topic in compiler research.

Given the performance gain permitted by tiling, our work will focus on improving spatial locality specifically for tiled loops.

D. Tiling enables overlapping communications from execution

In order to use two data layouts and gain transfer performance without losing compute power, we have to separate communications and computations. It is always possible to do so with tiled loops due to atomicity of a tile.

To sustain a high performance, all parts of the accelerator should stay active at all times, whether they are dedicated to computation or communication. Thanks to the atomic nature of tiles, it is possible by definition to read a tile's input ahead of execution, and write its output to memory after execution. It is therefore possible to split the workload into three tasks, each one processing a different tile: while a tile of iterations is being executed, the accelerator can prepare the execution of the next tile by reading its input data, and write the results of the previous tile.

The usual structure of a hardware accelerator for a tiled loop, shown in Figure 2, follows this read-execute-write template. Communications are overlapped with computation in the form of a *task-level pipeline*, that HLS tools such as Vivado and Catapult provide ways to achieve automatically.

A task-level pipeline significantly increases the overall throughput: instead of processing a new tile after each one has completed the read-execute-write tasks, a new tile starts

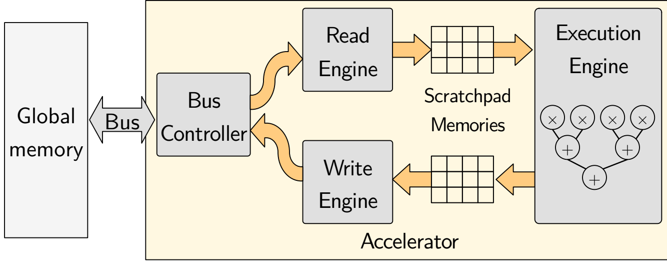

Fig. 2: Task-level structure of an accelerator: read, execution and writeback happen separately. *On-chip memory accesses are random, while off-chip accesses are contiguous.*

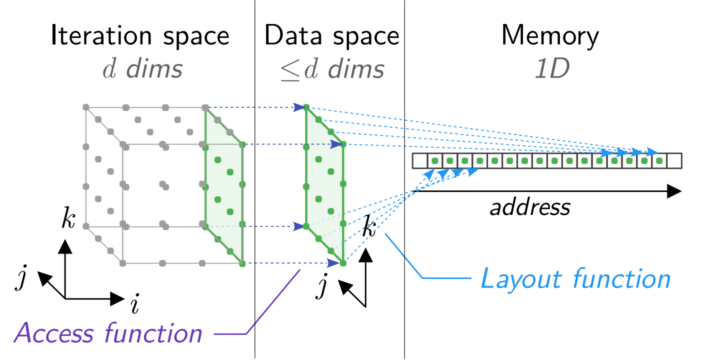

Fig. 3: Polyehdral iteration-to-memory flow: each memory access is the combination of an array access function and the layout function.

each time the readback of a tile completes, and at the same time, the execution and writeback of two other tiles complete.

E. Full bandwidth usage requires high spatial locality

Tiling alone does not bring spatial locality. The access pattern has to be transformed to achieve that.

Loop tiling, in general, increases temporal locality by having a tile's footprint fit in a cache or local memory. Traffic from / to global memory may still have a low throughput, keeping the program memory-bound. The memory access latency, and therefore the memory throughput, depends on the sequence of accessed addresses, which is called the memory access pattern.

The access pattern impacts a memory system's performance at several levels. At the DRAM chip level, switching rows and banks incurs a certain latency; reading physically contiguous data at consecutive addresses therefore performs better than accessing random locations. At the controller level, each transaction incurs an incompressible latency. Random accesses require one transaction each; burst accesses, which are reads or writes to several consecutive addresses, take in a single transaction. Making use of this feature gives a lower latency requires a contiguous access pattern.

The more contiguous accesses an access pattern has, the better its memory throughput is. Transforming a pattern this way is said to improve its spatial locality.

F. Spatial locality should be applied on flow-in and flow-out sets

The flow-in/out sets are the ones which need to be transformed to exhibit spatial locality.

Part of the data produced by a tile will be used later on by another tile. Given that tile size is usually determined to fill up on-chip memories, there is not enough on-chip memory to hold a tile's produced data for later use. Therefore this data needs to be stored in global memory. This specific set of data is called the *flow-out* set of the tile. Likewise, the *flow-in* set of a tile is those data produced earlier that need to be brought on chip before they are used.

Application-specific hardware accelerators usually feature scratchpad memories instead of caches. Copying the data in and out of scratchpads is up to the accelerator. The atomic nature of tiles allows splitting an accelerator's work into three discrete steps: read the flow-in data, perform the actual execution, and write the flow-out data to memory.

The flow-in and flow-out sets being the ones generating off-chip memory traffic, they are the ones for which we need to optimize the transfer from and to global memory.

G. Representing loops using the polyhedral model

We need a set representation of loops that supports transformations; the polyhedral model is a fit.

Loop transformations are more easily expressed as closed-form mappings of vector spaces, than as operations on loop nests. The usual program representation used for tiling is a compact, closed-form representation of a its control and data flows, called the polyhedral model.

Loops are mapped to a vector space called *iteration space* where each loop index is a dimension of this space. An affine function called *schedule* maps each iteration to an integer, giving an order of execution of iterations.

Constraints on the schedule are set by the *dependence pattern*, which is a set of affine functions mapping an iteration to those that it uses the result of: any iteration must be scheduled after those it depends on. This work focuses on programs with uniform dependencies; those are of the form $\vec{x} \mapsto \vec{x} + \vec{b}$ where $\vec{b}$ is a constant vector.

In the scope of this work, we will be performing affine operations on iteration vectors such as projection; it should be noted that tiling can be seen as the composition of affine operations on the iteration space.

H. Representing memory accesses using affine functions

Physical memory access functions can be decomposed into array access function + memory layout.

Every memory access performed by a program is made at an address that is computed at runtime from the loop indices using some *array access function*. Most programming languages support multi-dimensional arrays, making it possible to use multiple indices to access memory cells; however, memory is one-dimensional, so each multi-dimensional array must be mapped to a one-dimensional array in memory. Such a mapping is called a *memory layout*. Figure 3 shows the decomposition of memory accesses into the composition of array access functions and memory layout.

In the polyhedral model, access and memory layout functions are affine functions that respectively go from the iteration space to a data space, and from a data space to memory. There is one data space per array in the program, that mirrors an array, e.g. the array `int A[200][100]` is mapped to the set $\{A[i,j] : 0 \leqslant i < 200 \text{ and } 0 \leqslant j < 100\}$.

III. Related work

The high-level issue our work is addressing is the memory wall that appears on Figure 1, which is hindering the overall system performance. Techniques that aim to relieve it can be sorted into three main categories: those that increase the arithmetic intensity (right-wards on a roofline graph), those that increase performance (upwards on a roofline graph), and those that increase the effective bandwidth.

A. Increasing the arithmetic intensity

Increasing the AI reduces the pressure on memory in terms of volume of communication, be it with spatial locality / data reuse or with approximation.

Increasing the arithmetic intensity, per definition, means reducing the volume of data transferred per arithmetic operation. The main way to do that is to re-use the data already locally present on-chip or in a close level of the memory hierarchy.

Playing on the program's schedule to favor data reuse reduces the need for off-chip transfers. Loop tiling [3], [4] is the main technique in this aim. The primary target of tiling was to reduce the cache miss rate of CPUs, that have a cache hierarchy, but the technique also applies to software-programmable accelerators such as GPUs [12] , and thanks to high-level synthesis tools, it also applies to application-specific hardware accelerators as well [8]

Sharing on-chip resources to maximize their usage: On hardware accelerators, fine-tuning of on-chip memory allocation may eliminate off-chip accesses. Memory cells may be allocated multiple times, shared across tasks that do not interfere [13], increasing the usefulness of each memory cell, thereby reducing the amount of off-chip traffic.

B. Increasing the effective bandwidth

The root of the memory bandwidth issue is found in the massive parallelism the hardware can provide. When it is fully used, the memory latency may be higher than the compute latency, due to how the memory subsystem is designed and used. In this case, the design is said to be *memory-bound*. Such a limitation may be caused, for instance, by port contention, cache conflict misses, scalar accesses to global memory. All of these issues find their root in a sub-optimal memory access pattern, a sub-optimal physical placement (layout) of the data, or both.

Program and data transformations of several kinds can be used to adapt the access pattern to the data layout, or adapt the data layout itself to the program; proper use of the memory and its access interfaces result in an increase of the memory bandwidth to a value closer to the nominal memory chip bandwidth. Further increase of the *effective bandwidth*, which is the amount of useful data transferred per unit of time, may be achieved using data compression techniques.

Our wor is positioned among these techniques, but it is also essential that temporal locality and on-chip performance optimizations are applied to exhibit sufficient parallelism and create the demand for bandwidth.

1) Optimizing memory access pattern and layout: **Making an optimal use of bandwidth means have an on-chip data layout and access pattern free of port contention, and exhibit contiguity in off-chip accesses.**

Sub-optimal usage of memory bandwidth may be due to where the data is located in memory and the schedule of access requests. These can be tuned to alleviate certain issues they are at the root of.

a) Eliminating access conflicts: Access conflicts may occur in all memory architecture, when one wants to use a physical port or memory cell to convey multiple data at the same time. Mapping the data to other cells or addresses will resolve the conflicts. On FPGA chips, parallel memory access patterns may incur port contention. Such a phenomenon is common with on-chip memories, but may happen on multiple-port memory architectures such as high-bandwidth memory (HBM). Bank partitioning [14] is a key and widely used technique to bank conflict prevention, and is available in commercial high-level synthesis tools. Conflicts also happen in set-associative caches when multiple addresses share the same set and therefore the same physical memory cells in the cache. Appropriate array padding [15] reduces such conflict misses at the price of an increase in off-chip array size.

b) Exhibiting burst accesses by re-scheduling memory accesses: It may be possible to exhibit burst accesses by changing the access order of a given set of data addresses. If communications and computations are not separate [16], then the execution's schedule is also the memory access schedule; a loop transformation that maximizes DRAM row and burst use is found. Such a process changing the loop's execution order, it may break temporal locality. When communications and computations are executed separately and on-chip memories are used as scratchpads [8], the global memory access pattern does not need to match the on-chip access pattern. In this case, it is possible to improve memory bandwidth usage while keeping the same on-chip performance.

c) Re-allocating data in a burst-friendly layout: Regardless of the data layout, it is awlays possible to make burst accesses to memory, however the proportion of useful data accessed by such bursts may be low. It is possible to change the data layout to exhibit a high-usefulness burst access pattern. Data tiling [17] is one such technique; in the specific case of dense matrix product (GEMM), block matrix layout or *data tiling* as illustrated on Figure 4 (c) is known to have excellent spatial locality [18]. The dependence pattern of GEMM is such that the footprint of a tile can exactly fit a data tile, making it ideal layout for such an application. There is a trade-off between size and shape of data tiles and the usefulness of the burst accesses they permit, explored in works such as [19].

2) Increasing effective bandwidth by compressing data: **Compressing data increases the effective bandwidth, possibly at the cost of precision loss due to quantization.**

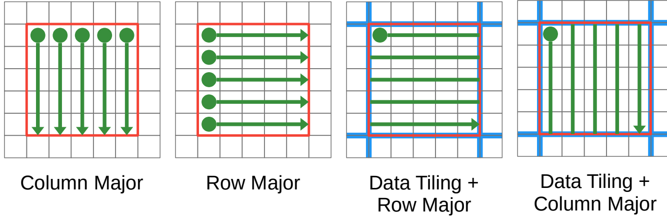

Fig. 4: Memory layouts (non-tiled and tiled) for a 2-dimensional data space, and access patterns for a tile. Each dot is the start of a burst access, that spans until the arrow tip.

Another class of solutions around the bandwidth issue is using compression, whether lossy or lossless. Ozturk et al. [19] created a dynamic lossless compression engine that acts like a cache, where local data is compressed before being sent out to memory, and decompressed when coming from memory. Compression combines two advantages: it saves both memory and bandwidth, at the cost of using extra cores or on-chip area to perform it. A major pitfall of compression combined with data tiling is that it requires to read or write a full tile even to access a single point from it.

Lossy compression enhances throughput as well, at the price of errors. Maier et al.'s perforation [20] is a form of lossy compression, as is the well-known JPEG algorithm. Nakahara et al. [10] use it on CNN inputs; their method is to compress the CNN input on the host using the lossy JPEG algorithm, decompress it on FPGA chip, and perform the inference there. As expected, accuracy goes down as the JPEG quality factor decreases, and there is a trade-off with the obtained speedup.

Sun et al. [21], to appear in 2022, tackle the same bandwidth problem as ours by using a mix of compression and data layout. Although this approach does not rely on polyhedral dependency analysis, it features the same base idea: group together data that is being used together.

C. Automatic synthesis of optimized hardware

Recent work in optimization of high-level synthesis has focused on memory allocation. Our work follows this trend.

There is a longtime sustained community interest in automating optimization of hardware design, including high-level polyhedral optimizations (as opposed to low-level RTL tuning). Compilers have been developed specifically to this aim.

Early work on polyhedral compilation was already targeting hardware design; Le Verge, Mauras and Quinton in 1991 [22] were using the Alpha language to automatically derive a systolic VHDL circuit from an Alpha polyhedral specification.

More recent work on HLS tools made it possible to explore tiling options for FPGA or ASIC accelerators (Pouchet et al., 2013 [8]) on various performance metrics such as latency, area, power consumption, memory bandwidth. Polyhedral transformations are done using specific tools, such as Pluto (Bondhugula et al., 2008 [6]) that automatically identifies tilable loop nests and applies tiling.

The most recent advances in tools bear a focus on data movement and memory issues. The SODA framework [23] automatically generates a dataflow-like pipeline structure with FIFO-ordered off-chip accesses; the data is automatically transformed have a specific allocation for this to work. This approach turns the data layout into a specific pre-determined layout, independent of the actual dependence pattern, whereas ours finds a data layout and an access pattern for each accelerator in function of the dependence pattern.

Xiang et al., 2022 [24] propose an approach that is complementary to ours, with an HLS code generator that infers the entire data movement and memory allocation, both off-chip and on-chip. This approach does not rely on fine-grain dependency analysis, and therefore has to keep the inner layout of each array; however, the location of each array in memory is carefully chosen so as to minimize the latency and maximize the bandwidth. Our approach is complementary: we make use of fine-grain dependency analysis to transform the inside of the arrays.

The next two sections explain in more detail the rationale and construction of our data layout, and the transformations a piece of code goes through to have it.

IV. Our method: Canonical Facet Allocation

This section details the core of our work, which principle is the following: by having different on-chip and off-chip data layouts, it is possible to have both on-chip parallelism (permitted by an adapted on-chip data layout) and low off-chip access latency (using a burst-friendly off-chip layout).

The next subsections explain the ideas that motivate our technique, and details how the off-chip data layout is built.

A. Trade-off between read and write contiguity

It is not possible to have a data layout that would guarantee at the same time and for all cases, that both the flow-in and flow-out of a tile can be accessed in a single transaction. This is due to the fact that, in the general case, results produced by one tile are consumed by more than one tile, as on Figure 5c, and each tile consumes results from multiple tiles as on Figure 5b.

Being able to read the whole flow-in data of a tile in a single burst transaction would require that a single contiguous region would be written to by multiple tiles. On Figure 5b, a contiguous flow-in region would be written to from seven distinct iteration tiles.

The same reasoning applies for the flow-out: if all the flow-out data from each tile was written to a single contiguous region of memory, the consumer tiles would pick their flow-in data from multiple regions and therefore make at least one access per producer tile. On Figure 5c, each colored region belongs to the flow-in of different tiles.

The two above scenarios are extreme cases, where either the flow-in of or flow-out of a tile is a single contiguous region (using a single burst), but the other way is scattered. There is therefore a trade-off between write contiguity and read contiguity. Canonical Facet Allocation takes the following stance on this trade-off:

- All write accesses are burst accesses (there are no element-wise writes),

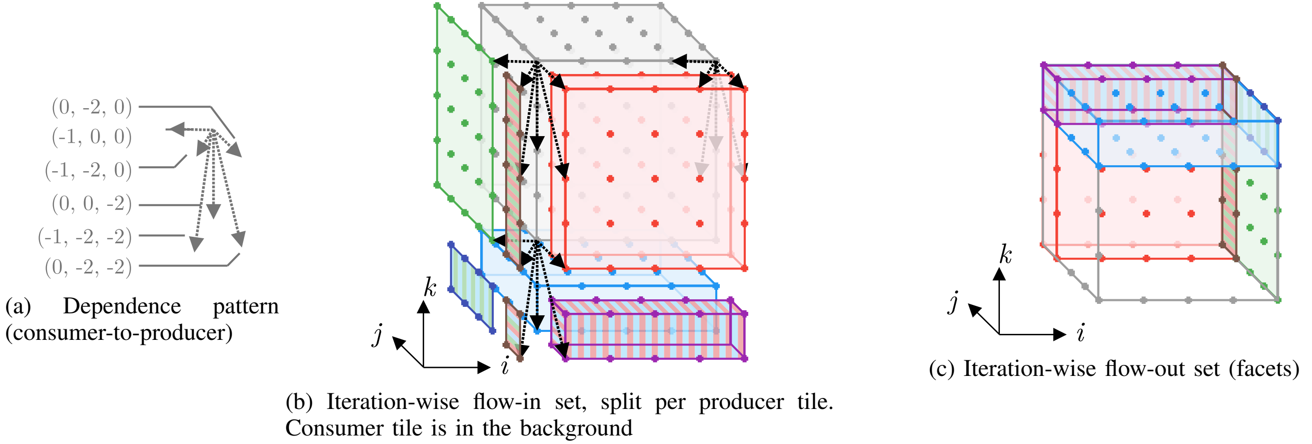

(a) Dependence pattern (consumer-to-producer)

(b) Iteration-wise flow-in set, split per producer tile. Consumer tile is in the background

(c) Iteration-wise flow-out set (facets)

Fig. 5: An instance of flow-in and flow-out sets. The flow-out set is the union of thicker versions of the tile faces (facets), while the flow-in set is composed of an union of either whole or partial facets. Some flow-in sets are adjacent in the iteration space; CFA reflects this adjacency in memory.

- Minimize the number of read transactions, whether they are element-wise or bursts.

The minimization objective provides no guarantee of absence of element-wise transactions, which sometimes may be unavoidable.

B. Separating off-chip and on-chip data layouts

It is, in general, not necessary that the off-chip data layout be the same as the on-chip layout. Indeed, both have a different objective: the on-chip layout should maximize access parallelism with low port contention, while the off-chip layout should allow minimal access latency and maximum throughput. We therefore propose to use two distinct layouts.

On-chip allocation has been the subject of many prior work, whether they are optimizing techniques or their automation; we assume it is already possible to find a suitable on-chip allocation to maximize parallel accesses. The core of our contribution is on off-chip allocation: we propose an off-chip allocation scheme adapted to tiled loop nests.

C. Principles of construction of off-chip allocation

We propose an allocation scheme that leverages multiple levels of contiguity. Relying on the shape of iteration space tiles, our allocation scheme offers contiguity for data produced by a tile (full-tile contiguity and intra-tile contiguity) as well as between adjacent tiles (inter-tile contiguity).

The data layout we propose honors these contiguity properties at the same time. The next subsections explain how we build it thanks to the combination of three techniques:

- multi-projection of each tile,
- data tiling,
- array dimension permutation.

D. Definitions

We use the following terminology in the next subsections:

Projection: We use a restricted definition of projection, considering only orthogonal projections. They are used to strip one dimension from a multi-dimensional space. For instance, the projection p_k such that

$$p_k(i_0, i_1, \ldots, i_k, \ldots, i_d) = (i_0, i_1, \ldots, i_{k-1}, i_{k+1}, \ldots, i_d)$$

removes the k-th dimension from an d-dimensional space.

Tile: A tile is an hyperrectangle, subset of the iteration space. Its size is $t_1 \times \ldots t_d$ in general; in the 3-dimensional example of Figure 5, tile size is respectively t_i, t_j, t_k along the i, j, k axis. A tile has coordinates $i_1, \ldots, i_d$ along each axis; in the example, these are ii, jj and kk.

First-level neighbor: A first-level neighbor of a tile is a neighbor that is reached with a move along a single canonical axis. For instance, $(ii, jj+1, kk)$ is a first-level neighbor of (ii, jj, kk), but $(ii+1, jj+1, kk)$ is not.

Second-level neighbor: A second-level neighbor of a tile is a neighbor such that is reached with a move along exactly two distinct canonical axes. For instance, $(ii-1, jj+1, kk)$ is a second-level neighbor of (ii, jj, kk), but $(ii+1, jj+1, kk+1)$ and $(ii+1, jj, kk)$ are not.

k-th level neighbor: By extension of the above two definitions, a k-th level neighbor of a tile is a neighboring tile that is reached with a move along exactly k canonical axes.

We will use the following notations:

- $E \subset \text{vect}\,(\vec{e_1}, \ldots, \vec{e_d})$: Iteration space of d dimensions
$$E = \{\vec{x} = (x_1, \ldots, x_d) : 0 \leqslant x_1 < N_1, \ldots, 0 \leqslant x_d < N_d\}$$
- $\vec{B}_1, \ldots, \vec{B}_p$: dependence vectors, such that rectangular tiling is legal. We assume all dependence vectors are backwards in all dimensions: $\forall i, j : \vec{B}_i \cdot \vec{e_j} \leqslant 0$.
- $N_1 \times \cdots \times N_d$: iteration space size
- $t_1, \ldots, t_d \in \mathbf{N}^*$: tile sizes

E. Hypotheses

We make the following hypotheses:

Uniform dependencies: It is, in general, impossible to predict the shape of the flow-in set of a tile when dependencies are not uniform. With uniform dependencies, the flow-in set of all tiles has the same shape, and the same goes for the

flow-out set. Many programs under active research fall under this hypothesis, ranging from stencils to convolutions.

Also, this hypothesis does not imply that all memory accesses are uniform - this only applies to accesses to read-write arrays (those that hold intermediate results).

Rectangular tiling: We assume tiles are rectangular in all dimensions; using polyhedral tools, it is notably possible to change the iteration space basis so that rectangular tiling becomes legal. Therefore, before applying CFA, we expect such a pre-processing to have been done if necessary. This does not imply a lack of generality.

Non-sparse data: It does not make sense to apply CFA on highly sparse data. The data layout of CFA is dense per construction, so that uniform dependencies yield uniform memory accesses. Using it with sparse data would lead to a significant amount of avoidable redundant data transfers. Sparse data should use a sparse representation instead.

F. Contiguity along multiple spatial directions: multi-projection

The first aspect of Canonical Facet Allocation is to make multiple projections of the iteration space. This makes it possible to have contiguity in single-dimensional memory along multiple spatial directions. The next two paragraphs explain why and how CFA achieves this.

1) Rationale: To get the best spatial locality, data should be laid out in memory the same way as the producer iterations are in the iteration space. Given that memory is abstracted to a one-dimensional space, in which it is not possible to have a multi-dimensional layout in memory. However, it is also true that:

- Not all data produced within an iteration tile needs to go through global memory - only flow-in and flow-out data need to,
- Only select neighbors of those iterations in flow-in / flow-out sets also belong to these sets.

An ideal memory allocation in terms of contiguity would be similar to a mathematical continuous function: two neighboring flow-in / flow-out iterations in *any* direction would also be neighbors in memory. Since the memory is one-dimensional, there is no such allocation. There is at most one direction per data space where this property can be verified, which is its direction of contiguity. It is however possible to use multiple data spaces, each one with a different direction of contiguity.

Because each flow-in / flow-out data point only needs specific directions of contiguity to exhibit spatial locality, we carefully choose the data that belongs to each data space. Visually, on Figure 5c, one would like to cut the flow-in / out data into three pieces, one per face of the cube that has flow-out data. This is what multi-projection achieves.

Multi-projection consists in creating multiple data spaces from a single iteration space, each data space corresponding to some projection of the iteration space. In the example of Figure 5c, there is a data space for each of the canonical hyperplanes (i, j), (i, k) and (j, k).

The next two paragraphs state how each projected data space is created, first with an example, then in the general case.

The idea is to make data spaces that are as thick as the dependence pattern "plunges" into the neighboring tiles.

2) Construction example: We start with a visual way to construct the data spaces, from a 3-dimensional iteration space, from Figure 5. The idea to construct the projected data space is, for every canonical hyperplane:

1) To determine how thick the facet (and consequently the data space) should be, and
2) To create a function that maps iteration coordinates to data space coordinates.

On Figure 5c, the part of the flow-out parallel to hyperplane (i, j), in light blue, has a thickness of 2; otherwise said, consumer tiles will need the result of the two uppermost (i, j) planes ($k \in \{3, 4\}$). The dual way to see it is the flow-in (Figure 5b): when moving the dependence pattern along the bottom plane of a consumer tile, the iterations this plane depends on (part of which the blue slab below the consumer tile) are located two planes below it.

The thickness of each data space is calculated using the dependence pattern: it is the maximum length of every dependence vector along the normal vector to the hyperplane we are projecting on. In the example of Figure 5, the dependence pattern is on Figure 5a. For hyperplane (j, k), it is the maximum absolute value of the component along the i axis of every dependence vector. We determine it to be 1: only the rightmost (j, k) plane of iterations is needed. Therefore, the data space created from the (j, k) hyperplane will be a simple two-dimensional array, that may be declared as `Dtype facet_i[`N_j`][`N_k`]`. We name it "facet" because it will ultimately contain data tiles holding entire facets.

To give coordinates to the iterations' results in the data space, we use a simple projection of the rightmost face of a tile:

$$p_i(i, j, k) = (j, k)$$

which domain is that rightmost face of each tile. As the tile size is 5, we get:

$$D(p_i) = \{(i, j, k) : i \equiv 4 \mod 5\}$$

Note that, although we do not do it at this point, the above projection function can be translated to code. Generating code of this function in its domain would result in a `for` loop that browses the rightmost face of each tile (in this example, tile size is 5 in each dimension): `for(`$j = 5jj$ `to` $5jj{+}4$`){ for(`$k = 5kk$ `to` $5kk{+}4$`){ facet_i[`j`][`k`] = iteration_result[4][`j`][`k`]}}`.

For hyperplane (i, j), the maximum absolute value of the component along the k axis of every dependence vector is 2. Therefore, the data space to be created will consist of the two uppermost (i, j) planes of every tile.

To create the mapping between iteration and data spaces, we use a modified projection, called a modulo projection: instead of getting rid of the component along the k axis, this projection replaces the k axis by $k \mod 2$:

$$p_k(i, j, k) = (i, j, k \bmod 2)$$

The domain is the two uppermost planes:

$$D(p_k) = \{(i, j, k) : 3 \leqslant k \bmod 5 \leqslant 4\}$$

3) General case: In the general case, i.e. for d-dimensional spaces, we determine the thickness of each facet (i.e. each data space) the same way as above. The thickness of each facet is determined by the longest dependence vector along the direction normal to that face. Assuming dependence vectors $\vec{B}_1, \ldots, \vec{B}_p$, the thickness of facet normal to $\vec{e_k}$ is given by:

$$w_k = \max_{q \in [\![1,p]\!]} \left| \vec{e_k} \cdot \vec{B}_q \right|$$

A proof that this is sufficient to hold all flow-in/flow-out data is located in annex.

We then take canonical *modulo* projections to map the iterations to the data space:

$$p_k(x_1, \ldots, x_d) = (x_1, \ldots, x_{k-1}, x_k \mod w_k, x_{k+1}, \ldots, x_d)$$

that are defined over the w_k last planes of a tile:

$$D(p_k) = \{(x_1, \ldots, x_d) : t_k - w_k \leqslant x_k \mod t_k \leqslant t_k - 1\}$$

4) Single-assignment scheme: The projections introduced in this section cause facets from tiles along a canonical axis to overwrite each other's data (all tiles along a canonical axis share the same memory locations for the data from the hyperplane normal to that axis). For instance, with the example of Figure 5c, all tiles along the i axis share the same projection of the (j, k) plane in the `facet_i` array.

The consequence of this overwriting is the possible introduction of memory dependences that may render the schedule of tiles illegal (data from a tile could be overwritten while it is still needed). In order not to break the schedule's legality, we choose to replicate the data spaces so as to make each data space a *single-assignment* space. This means each iteration tile will write in a memory space distinct from that of any other tiles.

Single-assignment spaces are created by introducing an extra dimension in the facet arrays that corresponds to the direction normal to the projection hyperplane. Continuing with Figure 5c example, the facet corresponding to (j, k) hyperplane is augmented with a dimension on the i axis, with so many entries as there are tiles on the i axis. The corresponding array becomes `facet_i`[N_i/5][N_j][N_k].

G. Flow-in from first-level neighbors: Full-tile contiguity

We want to ensure that each facet from each tile is written in the form of a burst access, which we call *full-tile contiguity*. Data tiling gives us this level of contiguity. The idea is to mirror the iteration space tiles on the data space; data tiling then allows to write each projected data space (facet) in a single burst access.

This transformation will improve both write and read performance: Flow-out facets are almost entirely used by the tile they are immediately adjacent to.

The next two subsections give an example and the general way to apply data tiling in the projected data spaces.

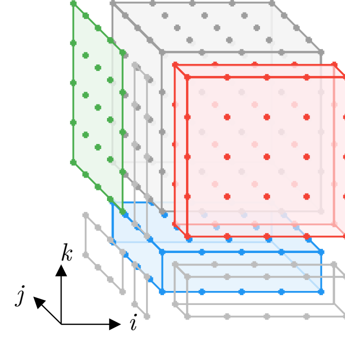

Fig. 6: Flow-in facets from first-level neighbors. Each facet is a data tile, i.e. a contiguous block of memory.

1) Example: We can continue with the example of Figure 5. In this example, we have built data spaces for each facet:

- `facet_i`[N_i/5][N_j][N_k]
- `facet_j`[N_j/5][N_i][N_k][2]
- `facet_k`[N_k/5][N_i][N_j][2]

We would like that the data from each facet of each tile to be written with a single burst access. For hyperplane (j, k), which is written to `facet_i`, we would like the destination space of projection p_i to be contiguous in memory, so that the code generated by the projection makes a single burst. We call this *full-tile contiguity*.

The destination space (image) of p_i is: $M_i = \{(i, j) : 5ii \leqslant i < 5(ii + 1) \wedge 5jj \leqslant j < 5(jj + 1)\}$, which corresponds to a flattened version of the iteration tile (ii, jj) onto the data space. The rectangular subset of the array `facet_i` that corresponds to M_i needs to be contiguous in memory.

The array `facet_i` will therefore be split into tiles of size 5×5, resulting in the following array:

`facet_i`[N_i/5][N_j/5][N_k/5][5][5]

2) General case: To mirror iteration tiles on the data space, we apply data tiling with the same tile sizes as the iteration space: the projection normal to $\vec{e_k}$ is therefore tiled with dimensions $(t_1, \ldots, t_{k-1}, t_{k+1}, \ldots, t_d)$.

Tiling the i-th dimension of an array with size t consists in replacing the i-th dimension by two dimensions, namely the quotient and the remainder of the Euclidean division of the original dimension by t. This operation gives $2(d-1)$-dimensional data spaces from a $(d-1)$-dimensional projection of the iteration space. All the quotient dimensions are moved first, and the remainder dimensions are moved last. We call the quotient dimensions the ***outer* dimensions**, and the remainder dimensions the ***inner* dimensions**.

Tiling a 3-dimensional data space A with tile sizes (t_i, t_j, t_k) would for instance be done with the following mapping from the original 3-dimensional to the tiled 6-dimensional array:

$$A[i][j][k] \mapsto A' \left[\frac{i}{t_i}\right] \left[\frac{j}{t_j}\right] \left[\frac{k}{t_k}\right] [i \% t_i][j \% t_j][k \% t_k]$$

This transformation is composed with the projection function of the previous section to give the actual allocation.

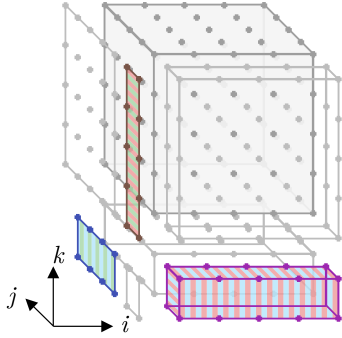

Fig. 7: Flow-in iterations from second-level neighbors of an iteration tile. CFA places these in memory next to a data tile already read to access them as “extensions”.

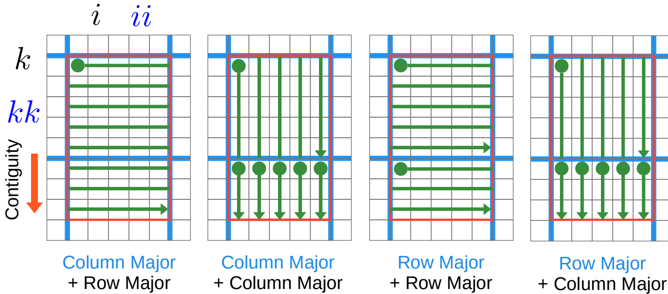

Fig. 8: The four possible data layouts for a 2-dimensional array in tiled data layout (inter-tile + full-tile layouts), and access pattern of a footprint that spans across data tiles. We seek inter-tile contiguity along direction k, which only column + row-major layout allows.

Data tiling provides a contiguity guarantee if the entire facet is read. As on Figure 7, we need to read only a subset of a facet from a second-level or third-level neighbor. Finer tweaking of the data layout is necessary so that these subsets are contiguous. The next subsections deal with these cases and brings in additional levels of contiguity.

H. Flow-in from second-level neighbors: Inter-tile contiguity

CFA mirrors contiguity across iteration tiles into contiguity across data tiles. This is called ”facet extensions”, as a burst access reading from data tile (holding a facet’s data) can be extended to fetch a neighboring tile’s data.

In the general case, part of the flow-in iterations of a tile are located in its second-level neighbors, such as those as shown in Figure 7.

We applied multi-projection and data tiling so that two neighboring tiles in memory are necessarily first-level neighbors in the iteration space. Contiguous reads from a first-level to a second-level iteration neighbor are possible, since a second-level neighbor is a first-level neighbor of a first-level neighbor. Such contiguity enables a cross-tile border read in a single burst access, as is shown for a two-dimensional array on Figure 8. We call this level of contiguity *inter-tile contiguity*.

This can be obtained by swapping the dimensions of the facet arrays. The next paragraphs explain how.

1) Example: We take back the example of Figure 5, of which the part of the flow-in data in second-level neighbors is shown in Figure 7. The producer iterations are part of two facets at the same time. We therefore have to:

1) Choose a direction of contiguity for each facet, and
2) Select the right facet to read each extension from.

If we call $E_{(ii,jj-1,kk-1)}$ the set coming from tile $(ii, jj-1, kk-1)$ (purple slab on Figure 7), we notice that this slab is a subset of both the (i, j) and (i, k) hyperplanes, so we can read it from both `facet_j` and `facet_k` arrays.

We choose that the direction of contiguity for the data space projected from the (i, k) hyperplane to be the k axis, and to read the extension $E_{(ii,jj-1,kk-1)}$ from `facet_j`. Figure 8 shows the four possible layouts for the `facet_j`: only column-major as a data tile layout and row-major as an intra-tile layout will allow reading $E_{(ii,jj-1,kk-1)}$ as a contiguous extension of the (i, k) facet of tile $(ii, jj-1, k)$. We choose this layout, and swap the dimensions of the `facet_j` array to match it. Column-major inter-tile layout means that dimensions are ordered this way: `kk`, `ii`; and row-major intra-tile layout gives the `i`, `k` order.

Accessing the `facet_j` array is therefore done with `facet_j[jj][ii][kk][k][i][2]`.

The result is that the purple and red slabs of Figure 5b can be read in a single, merged burst, from `facet_j`: they are contiguous along the k axis.

The same process can be repeated with the other facets.

2) General case: A tiled data space, as on Figure 8, has two dimensions corresponding to each canonical axis (one for tile coordinates, and one for intra-tile coordinates). By convention, if i is an axis, let’s designate `ii` the tile coordinate along that axis, and `i` the coordinate of a cell along the i axis inside a tile (this convention is followed on Figure 8).

The direction of inter-tile contiguity for a facet has to be chosen among those axes that are projected.

For a given projection, to make tiles along the i axis contiguous, then the `ii` dimension is moved as the last of the outer dimensions to be enumerated. Enabling contiguous reads is granted by promoting the `i` dimension the first of the inner dimensions to be enumerated.

Once the direction of inter-tile contiguity is picked, those parts of the flow-in sets that have been made contiguous can be merged to be read in a single burst.

I. Flow-in from third-level neighbors: Intra-tile contiguity

The subsets of flow-in coming from third-level neighbors, as is the set on Figure 9, are in general not contiguous in memory to data already accessed from first- or second-level neighbors. Still, this set may be read as a single contiguous burst. This section deals with the specific case of a three-dimensional iteration space.

In a 3-dimensional iteration space, there is only one flow-in set from a single third-level neighbor. It has a constant number of points, only a function of the dependence pattern. We call S_3 this subset.

For instance, on Figure 9, S_3 is the subset of iterations coming from tile $(ii-1, jj-1, kk-1)$ with $i = 4$, $j \in \{3, 4\}$

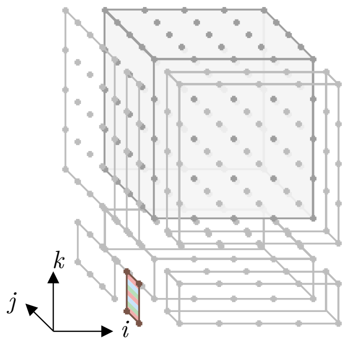

Fig. 9: Flow-in from third-level neighbor (pattern of Figure 5. CFA places these four points contiguously in memory within a data tile (intra-tile contiguity).

and $k \in \{3, 4\}$. This subset needs to be contiguous within one of the three projected facets.

It is possible to make S_3's data contiguous in memory by changing the order of the inner dimensions of `facet_k` (projection of (i, j) planes). Since this facet only contains iterations with $k \in \{3, 4\}$, we make S_3 contiguous using the order of dimensions kk, jj, ii, i, j, k. Then, the subset of Figure 9 is contiguous within `facet_k`: for any i, the points $(i, 3, 3)$, $(i, 3, 4)$, $(i, 4, 3)$ and $(i, 4, 4)$ are consecutive in memory. We can thus fetch them in a single burst.

The final layout of our arrays is therefore:

- `facet_i[ii][jj][kk][j][k]`
- `facet_j[jj][ii][kk][k][i][j % 2]`
- `facet_k[kk][jj][ii][i][j][k % 2]`

J. Case of k-th level neighbors

Although full, inter and intra-tile contiguity are always possible for first, second and third-level neighbors in a 3-dimensional space, it may not be in a 4 or higher-dimensional iteration space: the number of k-th level neighbors of a tile is C_k^d, which is higher than d, the maximum number of projections (i.e. of directions of contiguity).

V. Code generation

As part of our work, we have written a proof-of-concept source-to-source compiler pass that takes a C program corresponding to a tile and transforms it so that it uses CFA for global memory accesses.

Transforming the program to use CFA involves two main steps: analyzing the program's dependencies and calculating the facets, then transforming the program so that it uses the facets. Our proof-of-concept compiler pass sits within a compiler framework that targets high-level synthesis engines. Therefore, the code produced by our pass is turned into an FPGA accelerator using a synthesis tool such as Vitis HLS or Catapult.

This section explains how the program is transformed to include off-chip accesses to facets.

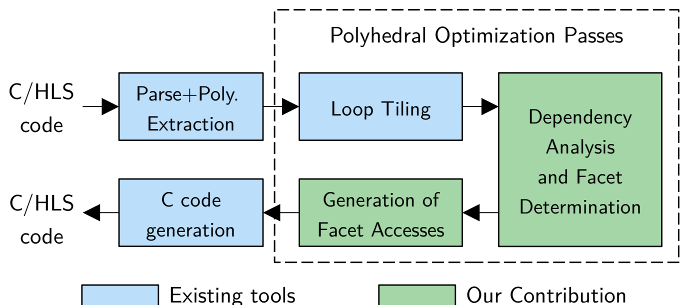

Fig. 10: Polyhedral compiler flow including the CFA pass (shown broken into two).

A. Overview: compiler pass

The flow of our compiler pass is as follows: it takes the the polyhedral representation of a program as input, under the hypotheses stated in the previous section. It determines what the facets are, under the form of sets of points. It then generates loops to scan these sets of points contiguously (copy-in / copy-out code), and wraps the tile's code with this copy-in and copy-out code. The copy-in/out code accesses global memory in CFA layout and turns it into the original program's layout for fast on-chip access.

The following subsections explain the transformations that are applied to the code and how copy-in/out code is generated.

B. Determining the facets and their layout

The technique described in Section IV is applied in order to determine the facets and their layout. Existing tools are used to do this: tiling is done using Pluto [6]; the Integer Set Library (ISL) is used to represent the iteration and data spaces, as well as the affine relations between them.

C. Copy-in/out code generation

CFA itself is only an allocation scheme for the data, which means it states *where* each datum is placed, not in *which order* the data is read or written. This subsection explains how we can make an access pattern that benefits from the contiguity properties of CFA, by issuing long burst accesses.

1) Making a contiguous flow-in/flow-out access pattern: We need to generate an access pattern that will take out the flow-in data from facets, and conversely, write the flow-out data into facets. This access pattern is supposed to be as contiguous as possible, and we expect the longest possible burst accesses.

Flow-out facets are entirely written with one transaction per facet - all the flow-out data is contained.

The intersection of the actual flow-in set with each facet may not be exactly a rectangle, and is perhaps not contiguous inside that facet. For this reason, a rectangular over-approximation of the set of accessed data is taken, like on Figure 11. That superset may span across facets from multiple iteration tiles - in this case, cross-tile contiguity ensures that a single transaction can bring all the data on-chip.

Given that a bounding box of a subset is contained in the bounding box of the containing set, doing so incurs less

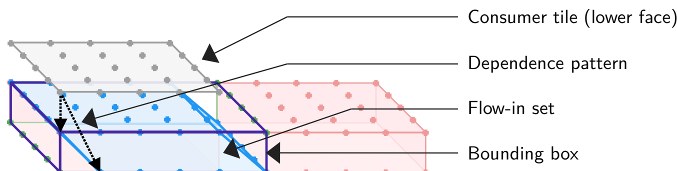

Fig. 11: Taking a rectangular over-approximation of the actual flow-in set incurs redundant reads. Contiguous data (facets) are in red: an over-approximation may be necessary to preserve access contiguity.

```
copy:
float* offChipAddr = &facet_0[l][n][m][0] + 0;
for(int I = 0; I <= 9215; I = I + 1) {
  #pragma HLS PIPELINE II=1
  int c6 = I / 96
  int c7 = I
  *offChipAddr = A_local[(96*n + c7 + 0)
  offChipAddr = offChipAddr + 1;
}
```

Fig. 12: Merger of two loops from which a burst access is inferred. on-chip addresses are random, while off-chip addresses are consecutive.

overhead than making a rectangular bounding box of the whole flow-in data.

Correctness: For flow-out data, due to the choice of using a tile-wise single-assignment allocation (no two different tiles write in the same memory cells), this over-approximation poses no correctness issues. For flow-in accesses, we must ensure this over-approximation doesn't break correctness (conflicts caused by two iterations sharing the same on-chip cell, while one is not part of the flow-in data). Filtering out the unneeded data coming from the bounding box is therefore necessary by adding a guard to the copy-in code.

2) Generating Burst-Capable Code: The code to be generated to fetch and write flow-in/flow-out sets has to trigger burst accesses on its memory controller. We generate synthesizable high-level synthesis (HLS) code, where bursts are inferred from `for` loops. The following conditions are sufficient for a burst to be inferred:

- The number of addresses accessed is explicit (e.g., the trip count of the copy loop nest is constant),
- If the target array size is known, the memory addresses are contained within its bounds,
- The addresses accessed are be consecutive (e.g., with a pointer increment),
- If applicable, the copy loop is pipelined with an initiation interval of 1 (assuming on-chip arrays are partitioned in such a way that this is possible).

Using the rectangular over-approximation of facet accesses, we can generate rectangular loop nests, that are coalesced into a single loop to force the inference of a single burst instead of a series of shorter bursts. as on Figure 12.

The copy code features two address generators, for off-chip and on-chip data. Off-chip address calculation is straightforward: when accesses are contiguous, one simply needs to provide the HLS tool with a pointer that starts at the beginning of the memory region to be accessed, and increment it. On-chip address calculation is performed at each cycle.

The resulting code for each facet is as on Figure 12, and is sufficient to get burst accesses using a commercial HLS tool like Vitis HLS.

D. Generating HLS code

The final step in CFA code generation is to generate a three-step coarse-grain pipeline:

- The first step reads the flow-in data from global memory in CFA allocation, and turns it into local allocation,

```
void toplevel(int i, int j, int k, float* facetIJ,
  float* facetIK, float* facetJK) {
#pragma HLS INTERFACE m_axi port=facetIJ
#pragma HLS INTERFACE m_axi port=facetIK
#pragma HLS INTERFACE m_axi port=facetJK
#pragma HLS DATAFLOW
    float buf1[TS*TS];
    float buf2[TS*TS];
    read(i, j, k, facetIJ, facetIK, facetJK, buf1);
    execute(i, j, k, buf1, buf2);
    write(i, j, k, buf2, facetIJ, facetIK, facetJK);
}
```

Fig. 13: Top-level function, assuming tile size is TS, local memories are $\texttt{TS}^2$ big and the iteration space is $\texttt{TS}^3$. This is the case of iterative stencils.

- The second step is tile execution,
- The last step is writeback from the accelerator to global memory.

This is implemented under the form of a function, as in Figure 13. Note the DATAFLOW pragma, which is used by the HLS engine to generate a coarse-grain pipeline where the three functions `read`, `execute` and `write` are executed in parallel, each function for a different tile.

VI. Evaluation

This section evaluates Canonical Facet Allocation by answering three questions with respect to the state of the art:

- Does Canonical Facet Allocation use all the available memory bandwidth?
- What is the improvement in effective bandwidth of CFA?
- Is there an area overhead due to using CFA versus single-layout allocations?

A. Experimental protocol

Experiments were carried out on a variety of uniform-dependence benchmarks listed in Table I. Benchmarks such as iterative stencils update an array in place, and differ by the dimensions of the iteration space and shape of the dependence pattern.

The platform used is a *Xilinx Zynq ZC706* board, including an *xc7z045ffg900-2* FPGA. The test accelerators only comprise read and write parts, the structure being that of Figure 14. Every baseline has been tested by connecting it to a single

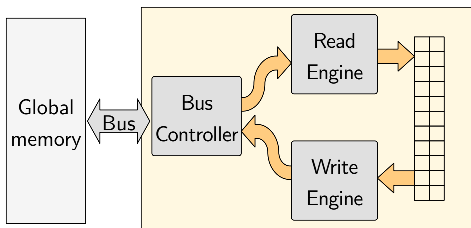

Fig. 14: Memory-bound accelerator used for benchmarking: only the read and write engines are implemented.

Dependence pattern	Nb of deps	Tile Sizes	Equivalent Application
jacobi2d5p	5	$16^3 \rightarrow 128^3$	Laplace equation
jacobi2d9p	9	$16^3 \rightarrow 128^3$	3×3 convolution
jacobi2d9p-gol	9	$16^3 \rightarrow 128^3$	2nd-order finite difference
gaussian	25	$4 \times 16^2 \rightarrow 4 \times 128^2$	5×5 Gaussian Blur
smith-waterman -3seq	7	$16^3 \rightarrow 128^3$	Alignment of 3 sequences

TABLE I: Benchmarks used for testing CFA. Equivalent applications have the same computational dependence pattern as the benchmark and would show similar performance.

AXI high-performance port mapped to DRAM (port HP0); the frequency of every design is 100.00 MHz, the AXI bus is 64-bit wide, and the data type transferred over the bus is 64-bit double-precision IEEE floating point numbers.

1) Baselines: We considered the following baselines for comparison with CFA:

- **Original Layout** (as done by Bayliss et al. [16]): a best-effort burst access pattern is determined under the original allocation. This access pattern does not issue any redundant reads, possibly at the expense of contiguity.
- **Bounding Box** (as done by Pouchet et al. [8]): a rectangular bounding box around the flow-in and flow-out data is taken so as to exhibit burst transfers; part of the bounding box is unused and redundantly transferred.
- **Data Tiling** (as done by Ozturk et al. [19]): data tiling is applied to the original arrays, and any tile that is accessed is entirely transferred. The reported value corresponds to the best performing tile size that is less or equal to the iteration tile size.

The original layout baseline introduces no transfer redundancy but has the smallest amount of and the shortest burst transfers. The two other baselines are tradeoffs between burst usefulness and bandwidth use: using a bounding box or data tiling result in using only long burst accesses at the price of transfer redundancy.

Each benchmark is tested against a variety of tile sizes, with 1:1, 1.5:1 and 2:1 ratios.

B. Results and discussion

1) Raw bandwidth: The raw bandwidth shown on Figure 15 indicates that the CFA layout and access pattern can reach close to 100% of the bus' maximum bandwidth, whereas other baselines exhibit high redundancy overhead (especially with the bounding box).

The high efficiency of our approach is mainly explained by three points:

- The small number of burst transfers per tile (4 in the case of 3-dimensional tiles). The data tiling approach uses a single burst access per tile, and the original layout and bounding box approaches issue multiple burst requests.
- The length of these burst transfers: one facet for the CFA approach,
- The ability of Vitis HLS to use burst access overlapping, which hides latency for long bursts even when they are decomposed into smaller burst accesses.

2) Effective bandwidth: The effective bandwidth assesses the usefulness of the transferred data: it only counts the data transferred that is actually useful for the application. Data transferred then ignored is consuming bus time, thus lowering the effective bandwidth. Figure 15 shows the effective bandwidth with colors, the difference with raw bandwidth being in grey. Two observations can be made:

- For the considered tile sizes, CFA is able to bring the effective bandwidth close to 100% of the bus bandwidth, which other allocations will not achieve.
- CFA is efficient even with small tile sizes. The `gaussian` benchmark, tiled with a small size in time (4) and larger spatial sizes (up to 128×128), shows that CFA exceeds 80% of the bus bandwidth for tile sizes above $4 \times 64 \times 64$.

The high usefulness of CFA (low difference between effective and raw bandwidth) is due to the choice of projections in CFA, which yields minimal redundancy.

3) Area cost: We analyze two distinct cost metrics specific to FPGA designs: the computational resources (Slice and DSP), and the storage resources (Block RAM).

a) Computational resources: **The cost of CFA itself in terms of hardware is the address generators. These are small, as about 95% of the logic area on our test platform remains available for compute engines.**

Regardless of the off-chip and on-chip allocations, the read and write engines take up a small fraction of the available logic resources. Figure 16 aggregates the area occupied by all baselines other than CFA for all tile sizes we have tried, and positions CFA. It can be observed that with tile sizes ranging from 16^3 to 128^3, except `gaussian` which tile sizes range from 4×16^2 to 4×128^2, designs occupy between 2 and 5% of the total slice area, and 0 to 3% of the total available DSP resources on an XC7Z045 FPGA chip. Canonical Facet Allocation does not show a significantly different slice occupancy than other baselines. For all benchmarks except `jacobi2d5p`, CFA requires some DSP blocks for some tile sizes, which are used to compute off-chip base addresses, but never needs more than 40 out of 900 (4%) of the DSP resources. The same observation holds for the `gaussian` benchmark, where the baseline using the original allocation takes between 26 and 30 DSP blocks.

Two conclusions can be drawn:

- Address generators are small units in terms of area, regardless of the allocation,

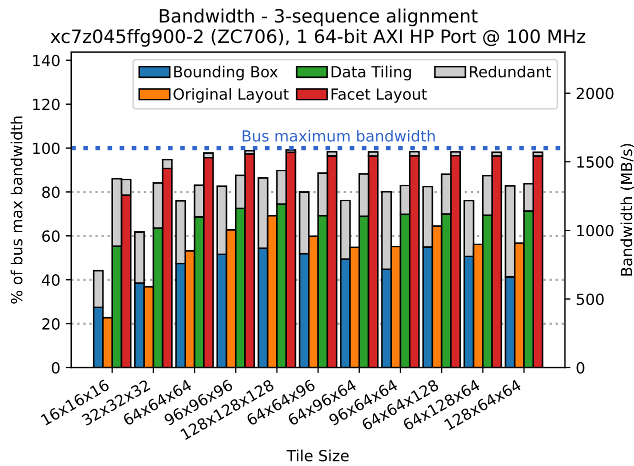

(a) `sw3d`

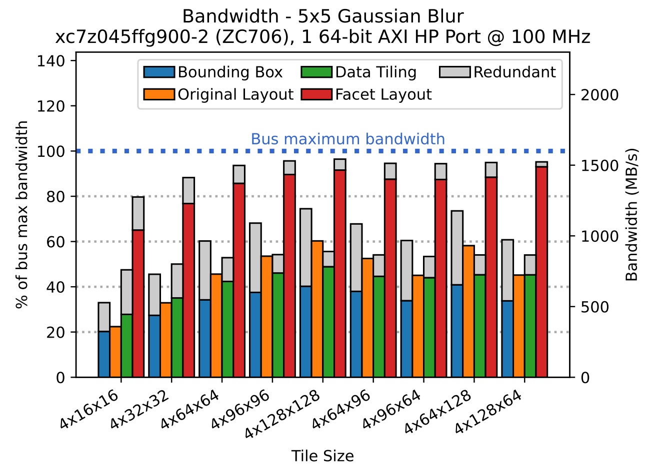

(b) `gaussian`

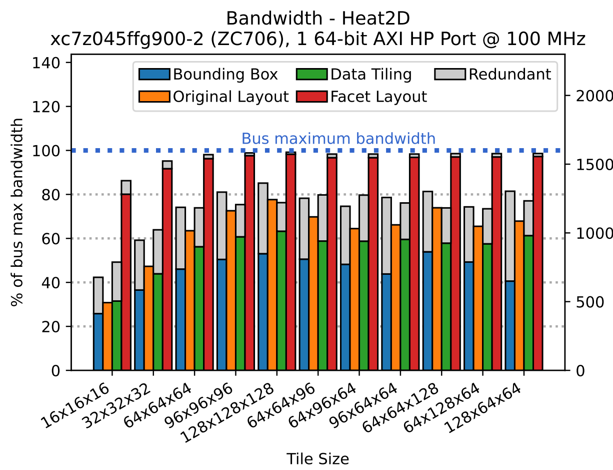

(c) `jacobi2d5p`

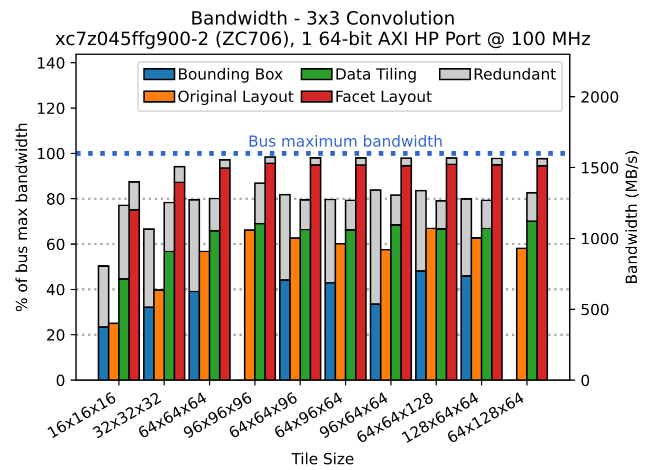

(d) `jacobi2d9p`

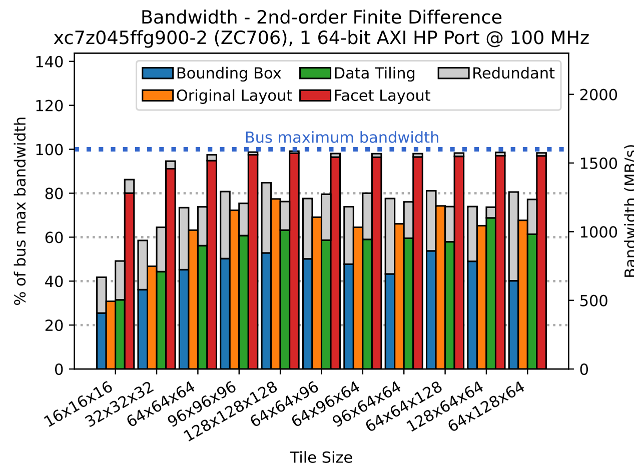

(e) `jacobi2d9p-gol`

Fig. 15: Bandwidth for all baselines, per benchmark. CFA is efficient even for tiles where one dimension is much smaller than the others. Bounding box is the technique used by Pouchet et al. [8], Data Tiling is used by Ozturk et al. [19], and the original layout is the best-effort baseline as in Bayliss et al. [16].

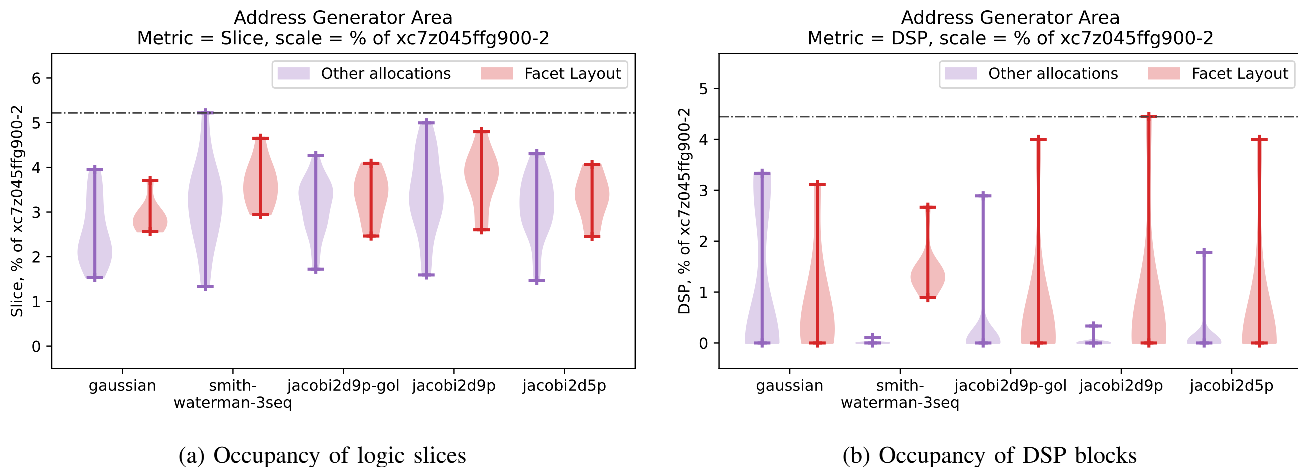

(a) Occupancy of logic slices

(b) Occupancy of DSP blocks

Fig. 16: Area occupied by CFA and all other baselines (aggregated), as a percentage of the available area of xc7z045ffg900-2. The vertical lines span from minimum to maximum.

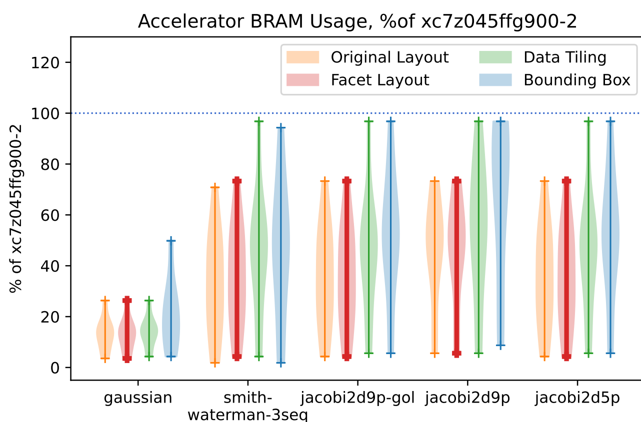

Fig. 17: BRAMs occupied by all baselines (in color), for each benchmark, as a proportion of BRAMs available in xc7z045ffg900-2.

- CFA does not exhibit a different area pattern than

It should additionnally be noted that the experiments were carried out without any compute logic on the accelerators; given the small size of the memory access modules, the synthesis tool did not have to significantly optimize the design for area.

b) Block RAM usage: **Using CFA does not change the on-chip allocation, therefore using CFA does not significantly increase the BRAM cost of a design.**

In an FPGA design, Block RAM resources used for on-chip data storage are shared between multiple actors. Even when the compute actor is not implemented, the memories needed to hold all the data on chip in and out of the memory actor must be present. Therefore, all designs do occupy a significant proportion, up to 95% of the available on-chip Block RAM, as Figure 17 shows. BRAM was, indeed, the factor limiting tile size - the larger the tile, the more data needs to fit into on-chip memories.

We can observe on Figure 17 that the distribution of on-chip memories using CFA and the original allocation is the same, with an exception (`smith-waterman-3seq`). This is due to the fact that CFA does not change the on-chip allocation, which is defining the amount of on-chip memory needed. The BRAM overhead for bounding box and data tiling baselines is mainly due to two facts: for the bounding box baseline, writing a superset of the tile footprint implies that the values written while not modified have been read and held on chip.

VII. Conclusion

Our work provides an answer to the under-utilization of memory bandwidth observed in many instances where an FPGA or ASIC accelerator is developed. The insufficient observed memory bandwidth results in stalls, preventing the full exploitation of the on-chip parallelism.

Memory allocation can be the cause of a significant under-utilization of the available bandwidth, due to the high latency of element-wise accesses. By introducing a data layout in memory for programs with uniform dependencies that is burst-friendly, we are able to overcome the under-utilization of the memory bandwidth and use it to its full extent. The method can be automated: we have developed a proof-of-concept procedure and compiler pass that implements the main code generation feature of CFA, which is the on-chip and off-chip address generator.

To further increase the benefits of CFA, the machine model we have considered may be extended to multi-port memory accesses, such as high-bandwidth memory, and distributed memories. In such architectures, there are multiple data ports; to benefit from all their bandwidth, one has to find an adequate repartition of data over each memory port to balance accesses.

References

[1] J. Cong, Z. Fang, M. Lo, H. Wang, J. Xu, and S. Zhang, “Understanding performance differences of FPGAs and GPUs,” in *2018 IEEE 26th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*, IEEE, apr 2018.

[2] S. Williams, A. Waterman, and D. Patterson, “Roofline: An insightful visual performance model for multicore architectures,” *Communications of the ACM*, vol. 52, p. 65, apr 2009.

[3] M. E. Wolf and M. S. Lam, “A data locality optimizing algorithm,” *ACM SIGPLAN Notices*, vol. 26, pp. 30–44, jun 1991.

[4] F. Irigoin and R. Triolet, “Supernode partitioning,” in *Proceedings of the 15th ACM SIGPLAN-SIGACT symposium on Principles of programming languages - POPL '88*, ACM Press, 1988.

[5] S. Verdoolaege, “isl: An integer set library for the polyhedral model,” in *Mathematical Software – ICMS 2010*, pp. 299–302, Springer Berlin Heidelberg, 2010.

[6] U. Bondhugula, A. Hartono, J. Ramanujam, and P. Sadayappan, “A practical automatic polyhedral parallelizer and locality optimizer,” in *Proceedings of the 29th ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI ’08, (New York, NY, USA), p. 101–113, Association for Computing Machinery, 2008.

[7] A. J. C. Bik, M. Girkar, P. M. Grey, and X. Tian, “Automatic intra-register vectorization for the intel architecture,” *International Journal of Parallel Programming*, vol. 30, no. 2, p. 65–98, 2002.

[8] L.-N. Pouchet, P. Zhang, P. Sadayappan, and J. Cong, “Polyhedral-based data reuse optimization for configurable computing,” in *Proceedings of the ACM/SIGDA international symposium on Field programmable gate arrays - FPGA '13*, ACM Press, 2013.

[9] C. Zhang, P. Li, G. Sun, Y. Guan, B. Xiao, and J. Cong, “Optimizing fpga-based accelerator design for deep convolutional neural networks,” in *Proceedings of the 2015 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*, FPGA ’15, (New York, NY, USA), p. 161–170, Association for Computing Machinery, 2015.

[10] H. Nakahara, Z. Que, and W. Luk, “High-throughput convolutional neural network on an FPGA by customized JPEG compression,” in *2020 IEEE 28th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*, IEEE, may 2020.

[11] R. Clint Whaley, A. Petitet, and J. J. Dongarra, “Automated empirical optimizations of software and the atlas project,” *Parallel Computing*, vol. 27, no. 1, pp. 3–35, 2001. New Trends in High Performance Computing.

[12] 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*, ICS ’12, (New York, NY, USA), p. 311–320, Association for Computing Machinery, 2012.

[13] X. Wei, Y. Liang, and J. Cong, “Overcoming data transfer bottlenecks in FPGA-based DNN accelerators via layer conscious memory management,” in *Proceedings of the 56th Annual Design Automation Conference 2019*, ACM, jun 2019.

[14] J. Cong, W. Jiang, B. Liu, and Y. Zou, “Automatic memory partitioning and scheduling for throughput and power optimization,” *ACM Transactions on Design Automation of Electronic Systems, Vol. 16, No. 2, Article 1*, vol. 16, pp. 1–25, 2011.

[15] C. Hong, W. Bao, A. Cohen, S. Krishnamoorthy, L.-N. Pouchet, F. Rastello, J. Ramanujam, and P. Sadayappan, “Effective padding of multidimensional arrays to avoid cache conflict misses,” *SIGPLAN Not.*, vol. 51, p. 129–144, June 2016.

[16] S. Bayliss and G. A. Constantinides, “Optimizing SDRAM bandwidth for custom FPGA loop accelerators,” in *Proceedings of the ACM/SIGDA international symposium on Field Programmable Gate Arrays - FPGA '12*, ACM Press, 2012.

[17] M. D. Lam, E. E. Rothberg, and M. E. Wolf, “The cache performance and optimizations of blocked algorithms,” vol. 25, pp. 63–74, 1991.

[18] J. R. Herrero and J. J. Navarro, “Using non-canonical array layouts in dense matrix operations,” in *Applied Parallel Computing. State of the Art in Scientific Computing*, pp. 580–588, Springer Berlin Heidelberg, 2006.

[19] O. Ozturk, M. Kandemir, and M. Irwin, “Using data compression for increasing memory system utilization,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 28, pp. 901–914, jun 2009.

[20] D. Maier, B. Cosenza, and B. Juurlink, “Local memory-aware kernel perforation,” in *Proceedings of the 2018 International Symposium on Code Generation and Optimization*, ACM, feb 2018.

[21] G. Sun, S. Kang, and S.-W. Jun, “BurstZ+: Eliminating the communication bottleneck of scientific computing accelerators via accelerated compression,” *ACM Transactions on Reconfigurable Technology and Systems*, vol. 15, pp. 1–34, jun 2022.

[22] H. L. Verge, C. Mauras, and P. Quinton, “The ALPHA language and its use for the design of systolic arrays,” *Journal of VLSI signal processing systems for signal, image and video technology*, vol. 3, pp. 173–182, sep 1991.

[23] Y. Chi, J. Cong, P. Wei, and P. Zhou, “SODA: Stencil with Optimized Dataflow Architecture,” in *Proceedings of the International Conference on Computer-Aided Design*, ACM, nov 2018.

[24] S. Xiang, Y.-H. Lai, Y. Zhou, H. Chen, N. Zhang, D. Pal, and Z. Zhang, “HeteroFlow: Ac Accelerator Programming Model with Decoupled Data Placement for Software-Defined FPGAs,” in *International Symposium on Field-Programmable Gate Arrays (FPGA)*, 2022.

APPENDIX

A. Set-wise construction of facets

Reminder of the notations:

- $E \subset \text{vect}(\vec{e_1}, \ldots, \vec{e_d})$: Iteration space (rectangular subspace of $\mathbf{Z}^d$), with $(\vec{e_1}, \ldots, \vec{e_d})$ an orthonormal base

$$E = \{\vec{x} = (x_1, \ldots, x_d) : 0 \leqslant x_1 < N_1, \ldots, 0 \leqslant x_d < N_d\}$$

- d : number of dimensions of E
- $\vec{B_1}, \ldots, \vec{B_p}$: dependence vectors, such that rectangular tiling is legal. We assume all dependence vectors are backwards in all dimensions: $\forall i, j : \vec{B_i} \cdot \vec{e_j} \leqslant 0$.
- $N_1, \ldots, N_d$: iteration space size (assuming a rectangular iteration space)
- $t_1, \ldots, t_d \in \mathbf{N}$: tile sizes (assume N_i if no tiling on dimension i)

Let $k \in [\![1, d]\!]$, $\vec{e_k}$ the canonical vector of the k-th dimension. The k-th *face* is given by those iterations which k-th coordinate is equal to $t_k - 1$.

The depth of the k-th facet is:

$$w_k = \max_{q \in [\![1,p]\!]} \left| p_{\vec{e_k}}\left(\vec{B_q}\right)\right| = \max_{q \in [\![1,p]\!]} \left|\vec{e_k} \cdot \vec{B_q}\right|$$

Let a tile of iterations be

$$T_{i_1,\ldots,i_d} = \left\{\vec{x} = (x_1, \ldots, x_d) : \forall q \in [\![1, d]\!] : \left\lfloor \frac{x_q}{t_q} \right\rfloor = i_q \right\}$$

The k-th facet for tile T is the set of iterations given by:

$$S_k(T) = \{(x_1, \ldots, x_d) \in T : t_k - w_k \leqslant x_k \mod t_k\}$$

B. Proof that the flow-out of every tile is contained inside facets

In order to replace the existing off-chip memory accesses by facet accesses, all the flow-in iterations of a tile need to be contained inside facets.

Proposition 1: Flow-in iterations are contained inside facets.

Purpose: Proves that all the flow-in data can be retrieved from facets.

Proof:

The iteration-wise flow-in set of a tile is defined as those iterations which result is used by this tile but are executed in another tile. It can be written as:

$$\varphi_i(T) = \left\{\vec{y} \in E \setminus T : \exists j \in [\![1, p]\!] : \vec{y} - \vec{B_j} \in T\right\}$$

Let $T = T_{i_1,\ldots,i_d}$ be a tile of iterations, $\vec{y} = (y_1, \ldots, y_d) \in \varphi_i(T)$, and T' be the tile of iterations containing $\vec{y}$. Let $j \in [\![1, p]\!]$ such that $\vec{y} - \vec{B_j} \in T$. Obviously, $\vec{B_j}$ is non-null, so let $q \in [\![1, d]\!]$ such that $\vec{B_j} \cdot \vec{e_q} < 0$.

Let's first show that there exists a q such that $\vec{B_j} \cdot \vec{e_q} \neq 0$ and $\left\lfloor \frac{y_q}{t_q} \right\rfloor \neq i_q$: given that $\vec{B_j} \neq \vec{0}$, there exists q such that $\vec{B_j} \cdot \vec{e_q} \neq 0$; if any such q verifies $\left\lfloor \frac{y_q}{t_q} \right\rfloor = i_q$, then for all $r \in \{1, \ldots, d\}$, $\left\lfloor \frac{y_r}{t_r} \right\rfloor = i_r$, thus $\vec{y} \in T$, which is false.

Let q be such that $\vec{B_j} \cdot \vec{e_q} \neq 0$ and $\left\lfloor \frac{y_q}{t_q} \right\rfloor \neq i_q$. By definition, $w_q \geqslant \left|\vec{B_j} \cdot \vec{e_q}\right|$. We have:

$$\left\lfloor \frac{y_q}{t_q} \right\rfloor \neq i_q \quad \text{and} \quad \left\lfloor \frac{y_q - \vec{B_j} \cdot \vec{e_q}}{t_q} \right\rfloor = i_q$$

As $\vec{B_j} \cdot \vec{e_q} < 0$, $\left\lfloor \frac{y_q}{t_q} \right\rfloor < i_q$. We thus have

$$y_q \mod t_q - \vec{B_j} \cdot \vec{e_q} \geqslant t_q$$

and as $w_q \geqslant \vec{B_j} \cdot \vec{e_q}$,

$$y_q \mod t_q \geqslant t_q - w_q$$

which proves that $\vec{y} \in S_q(T')$. As a result, *the iteration flow-in of a tile is contained inside an union of facets.*

■