

Iris: Automatic Generation of Efficient Data Layouts for High Bandwidth Utilization

Stephanie Soldavini

Politecnico di Milano

Milan, Italy

stephanie.soldavini@polimi.it

Donatella Sciuto

Politecnico di Milano

Milan, Italy

donatella.sciuto@polimi.it

Christian Pilato

Politecnico di Milano

Milan, Italy

christian.pilato@polimi.it

ABSTRACT

Optimizing data movements is becoming one of the biggest challenges in heterogeneous computing to cope with data deluge and, consequently, big data applications. When creating specialized accelerators, modern high-level synthesis (HLS) tools are increasingly efficient in optimizing the computational aspects, but data transfers have not been adequately improved. To combat this, novel architectures such as High-Bandwidth Memory with wider data busses have been developed so that more data can be transferred in parallel. Designers must tailor their hardware/software interfaces to fully exploit the available bandwidth. HLS tools can automate this process, but the designer must follow strict coding-style rules. If the bus width is not evenly divisible by the data width (e.g., when using custom-precision data types) or if the arrays are not power-of-two length, the HLS-generated accelerator will likely not fully utilize the available bandwidth, demanding even more manual effort from the designer. We propose a methodology to automatically find and implement a data layout that, when streamed between memory and an accelerator, uses a higher percentage of the available bandwidth than a naive or HLS-optimized design. We borrow concepts from multiprocessor scheduling to achieve such high efficiency.

ACM Reference Format:

Stephanie Soldavini, Donatella Sciuto, and Christian Pilato. 2023. Iris: Automatic Generation of Efficient Data Layouts for High Bandwidth Utilization. In *28th Asia and South Pacific Design Automation Conference (ASPDAC '23)*, January 16–19, 2023, Tokyo, Japan. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3566097.3567892>

1 INTRODUCTION

Optimizing data transfers is one of biggest challenges in computing today [7, 20]. Many applications, particularly big data and machine learning (ML) algorithms, require huge amounts of data to be transferred and often this is an extreme bottleneck [16]. A lot of effort has been put into optimizing the computational aspects of these algorithms, particularly in the development and improvement of high-level synthesis (HLS) tools [9]. However, the speedup gained on the computation side has not been matched on the data-movement side and thus these applications cannot take full advantage of the

optimized accelerators. In an attempt to solve this problem, High-Bandwidth Memory (HBM) architectures with wide data busses are increasingly common. The Xilinx Alveo u280 has HBM with a maximum bandwidth of 460 GB/s and the Intel Stratix 10 MX has HBM with a maximum bandwidth of 409 GB/s. However, it is very difficult to realistically achieve these high bandwidths. Designers must put in a lot of manual effort to carefully ensure their design utilizes the full bus every single clock cycle [5]. In even the simplest designs, this data orchestration can be quite complex and resource intensive. In some cases, HLS tools can automatically fill the wide bus by unrolling the arrays, but only if the design meets stringent requirements. For instance, the bus width should be evenly divisible by the data width and the array length should be a power of two. Designers can manually make adjustments to meet these requirements or put in more effort to manually pack the bus. Even more effort is needed for a highly custom solution beyond packing equally sized data into evenly divided “lanes”. These highly custom designs are not uncommon, especially with custom-precision data types increasingly used in ML applications [15]. These arbitrarily-sized data prove difficult to fit onto a fixed-width bus.

We propose *Iris*, an algorithm for automatically finding a *data layout* (i.e., an organization of data in memory and in the bus lanes) such that, when streamed to an accelerator, maximizes the use of the available bandwidth. We borrow concepts from processor scheduling to solve this problem. Our contributions are:

- An algorithm to automatically find an efficient data layout;
- A methodology for generating a host-side function that creates a unified array of all the input data in the specified layout;
- A methodology for generating the accelerator-side, HLS-ready modules to convert such data layouts into streams for the kernels.

The automation of this process is useful not only for reducing manual designer effort, but also for rapid design-space exploration while tuning the width of custom-precision data types.

2 RELATED WORK ON HBM CHALLENGES

With the introduction of HBM architectures, designers should follow recommended guidelines to exploit the increased bandwidth. First, they should use a data width and clock frequency compatible with the architecture. For instance, the HBM in the Xilinx Alveo u280 platform operates at 450 MHz with a data width of 256 bits per channel, so the design should either use this frequency and width or 225 MHz with a 512 bit width. Additionally, transactions should be as large as possible to minimize the overhead per transaction [22].

Several works focused on maximizing the use of the channel bandwidth. The work in [21] analyzes database applications accelerated on FPGAs with HBM. They explicitly craft their designs to ensure that queries return parallelized data to use the full bandwidth.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ASPAC '23, January 16–19, 2023, Tokyo, Japan

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9783-4/23/01...\$15.00

<https://doi.org/10.1145/3566097.3567892>

Table 1: Summary of notation used in this work

m	There are m processors
j	There are j tasks
δ_j	The maximum number of processors task j can use at once
d_j	The due date of task j (time when j would ideally finish)
r_j	The release time of task j (earliest time a task can start)
C_j	Completion time of task j
C_{max}	Makespan, maximum completion time of all tasks
L_j	Lateness ($C_j - d_j$) of task j
L_{max}	Maximum lateness of all tasks

The work in [4] proposes HBM Connect, a customized crossbar for HBM access. They use a virtual HLS FIFO buffer to gather data for read and write operations. On the Alveo u280, they achieve up to 185 GB/s over 16 channels (where the ideal bandwidth would be 230 GB/s). The work in [11] proposes a novel sparse matrix-vector multiplier and an ILU0 preconditioned BiCGStab solver. They explicitly design their computation pipeline to accept a full 512-bit cache line. The main focus is to rearrange the data access (and therefore the compute kernel) to avoid transferring zero data. An accelerator for the single-source shortest path problem is proposed in [3]. Due to the random-access nature of graph problems, they focus on improving throughput and optimizing for HBM. The work in [6] proposes an automated methodology for HLS kernels to use more bandwidth than their naive code. Their methodology trades off between using more BRAM or achieving a greater speed.

Since bandwidth is valuable, minimizing the amount of data transferred is important. The deep CNN design in [19] aims at exploiting DRAM bandwidth, reducing the number of data transfers and avoiding stalls. The stream analytics engine in [14] reduces the amount of data to be accessed by putting a smaller, more regularly accessed portion of the data (pointers) into HBM and uses those to reduce DRAM accesses. The LLVM pass in [12] partitions data between DRAM and HBM. It could be extended to LLVM-based HLS tools. Deciding where to store data can help relieve some bandwidth congestion. Using a custom-width data format can reduce the total data by reducing the bit-width of each element. Custom precision was used in [1] to accelerate neural network training. This work targeted only CPU with no memory considerations. To the best of our knowledge, no prior work has focused on bandwidth optimization with custom data types.

3 PROBLEM FORMULATION AS SCHEDULING

Scheduling is a popular research problem to decide how jobs are assigned to resources to reduce the overall time to complete all activities while satisfying the constraints [18].

In our case, given a bus width (m) and a set of accelerator arrays, each with bitwidth (W_j), depth (D_j), and desired due date (d_j), we want a memory layout where data are packed most densely and the arrays arrive as close to their due dates as possible when transferred from memory to the accelerator. This can be viewed as a processor scheduling problem as follows: an m -bit wide bus is a multiprocessor system made of m identical processors and the data arrays are the tasks, j , with due dates d_j , and processing times $p_j = W_j \times D_j$. These “tasks” are preemptible, i.e., they can be scheduled discontinuously without incurring additional overhead. The “tasks” will be scheduled on multiple processors at once. The maximum number of bits an array can use on the bus at a time, or the

Table 2: Summary of additional symbols used in algorithms

p_j	Processing time (time units needed to execute) task j
p_{tot}	Total processing time for all tasks (i.e. C_{max} if $m = 1$)
l	Number of unique due dates d_j
d_{max}	Maximum (latest) due date of all d_j
W_j	Bitwidth of j
B_{eff}	Bandwidth efficiency
$h(j)$	Height (minimum possible execution time) of j
R_k	Set of tasks with release time r_k
β_j	Processors allocated to j
t	Current timestep being processed
τ	Length of interval being scheduled

maximum number of processors a task can use at once, is notated by δ_j . Arrays may be needed at different times in an accelerator. So each has a due date d_j , derived from the dataflow graph and the latencies of the nodes. To ensure the arrays can arrive as shortly after their due date as possible, we use the maximum lateness, $\gamma = L_{max} = \max_j (L_j)$, optimality criterion, where $L_j = C_j - d_j$ is the lateness of an array and C_j is the completion time, or last cycle the array is on the bus. All together, our problem is as follows: *in a system with m identical processors, we want to schedule preemptible tasks across several processors (where task j gains linear speedup by being scheduled on up to δ_j processors at once) such that each task is finished as soon after its due date d_j as possible.*

4 IRIS: OUR DATA LAYOUT ALGORITHM

We can find an $O(n^2)$ solution to an isomorphic problem of our formulation in [8]. This isomorphic problem uses release times r_j , or the time step when a task j is ready to begin execution, instead of due dates d_j , and optimizes the completion time, $C_{max} = \max_j (C_j)$ (also known as schedule length or makespan), instead of the maximum lateness L_{max} . This problem is described as follows: in a system with m identical processors, we want to schedule preemptible tasks with release time r_j across several processors to minimize the total schedule length (C_{max}). To convert between the two problems, each due date d_j is converted to a release time r_j by subtracting it from the maximum (latest) due date of all tasks such that $r_j = d_{max} - d_j$. Also, the solution schedule to the isomorphic problem should be read backward to find the solution to the original problem. In this way, tasks that originally have the latest due dates will have the earliest release times in the “backward” schedule. Fig. 1 shows how converting due dates d_j into release times r_j can yield the same schedule but reversed in time.

Algorithm 1.1 shows our proposed layout finding algorithm, called *Iris*. In Greek mythology, Iris is the messenger of the gods. Additional symbols used in the algorithms are summarized in Table 2. The algorithm in [8] is designed to optimize the completion time, C_{max} , given a set of release times, r_j , when tasks are available to be processed. The algorithm can be converted to optimizing for the minimum lateness, L_{max} , by changing a set of l due dates, $d_1 \leq d_2 \leq \dots \leq d_l$, to release times as follows: $r_j = d_{max} - d_j$. After scheduling using these new release times, the schedule can be read backward to optimize for L_{max} for input arrays using the due dates.

To adapt this algorithm for the **bus layout problem**, we modified it as follows. Instead of using a simple ratio, Iris uses the largest-remainder method (also known as the Hamilton Method) of apportionment to allocate processors to tasks [13]. This method ensures tasks are assigned whole numbers of “processors” (i.e., bus

Algorithm 1.1 Layout Algorithm

```

1:  $t := 0$ 
2: Group tasks with release time  $r_k$  in set  $R_k$ ,  $k = 1, \dots, l$ 
3: for  $k := 1$  to  $l$  do
4: Order  $R_k$  by nonincreasing values of  $h(j)$ 
5: while ( $r_{k+1} < t$ ) and ( $\exists_{j \in R_k} h(j) > 0$ ) do
6: FIND_CAPABILITIES( $R_k, \beta$ )
7: if  $\exists_{j, j+1 \in R_k} h(j) > h(j+1)$  then
    $\triangleright$  Shortest time before  $h(j), h(j+1)$  are equal
8:    $\tau' := \min \left\{ \frac{h(j) - h(j+1)}{\beta_j - \beta_{j+1}} : \frac{\beta_j}{\delta_j} \neq \frac{\beta_{j+1}}{\delta_{j+1}}, h(j) > h(j+1) \right\}$ 
9:   else
10:     $\tau' := \infty$ 
11:   end if
12:    $\tau'' := h(|R_k|)$   $\triangleright$  Time to earliest completion of any task
13:    $\tau := \min \{\tau', \tau'', r_{k+1} - t\}$   $\triangleright$  Interval is until next change
14:   For all  $j \in R_k$ , schedule  $j$  on  $\beta_j$  processors for the interval  $[t, t + \tau]$ 
15:    $h(j) := h(j) - \frac{\tau \beta_j}{\delta_j}$  for  $j \in R_k$   $\triangleright$  Subtract this proc. time from  $h$ 
16:    $t = t + \tau$   $\triangleright$  Update the timestep
17: end while
    $\triangleright$  Add any unfinished tasks to the next batch,  $R_{k+1}$ 
18: if  $\exists_{j \in R_k} h(j) > 0$  then  $R_{k+1} := R_{k+1} \cup \{j : j \in R_k, h(j) > 0\}$ 
19: end if
20: end for

```

Algorithm 1.2 Find_Capabilities Procedure

```

21: procedure FIND_CAPABILITIES( $X, \bar{\beta}$ )  $\triangleright X$  is a set of tasks
22:    $\bar{\beta} := \bar{0}$   $\triangleright \bar{\beta}$  is a vector of # processors allocated to each task  $j$ 
23:    $avail := m$   $\triangleright avail$  is the number of free processors
24:   while  $avail > 0$  and  $|X| > 0$  do
25:      $T :=$  set of the highest tasks in  $X$  with  $h(j) > 0$ 
26:     if  $\sum_{j \in T} \delta_j > avail$  then
27:        $\beta_j :=$  LRM_ALLOCATION( $T$ );  $avail := 0$ 
28:     else  $\triangleright$  Tasks in  $T$  can use at most  $avail$  processors
29:        $\beta_j := \delta_j$  for  $j \in T$ ;  $avail := avail - \sum_{j \in T} \delta_j$ 
30:     end if
31:      $X := X - T$   $\triangleright$  Remove scheduled tasks from the working set,  $X$ 
32:   end while
33: end procedure

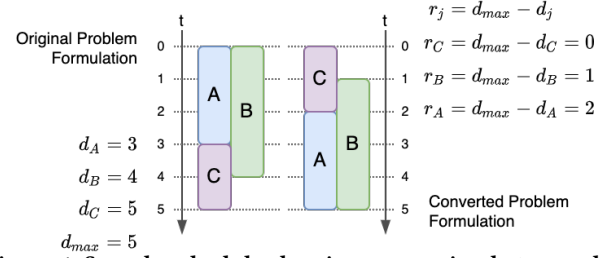
```

Algorithm 1.3 LRM_Allocation Procedure

```

34: procedure LRM_ALLOCATION( $T$ )
35:    $quota := \left( \sum_{j \in T} \delta_j \right) / avail$   $\triangleright$  Hare quota of processors
36:   for  $j \in T$  do
37:      $v_j := \frac{\delta_j}{quota}$   $\triangleright v_j$  is the processors requested per quota
38:      $\beta_j := \lfloor v_j / \delta_j \rfloor$   $\triangleright$  Assign  $\beta_j$  the largest multiple of  $\delta_j$  below  $v_j$ 
39:      $rem_j := v \bmod \delta_j$   $\triangleright$  Keep track of the remainder
40:      $avail := avail - \beta_j$ 
41:   end for
42:   Sort  $T$  by decreasing  $rem_j$ 
43:   for  $j \in T$  do
    $\triangleright$  If  $j$  fits in the remaining space, schedule it
44:   if  $avail > W_j$  then  $\beta_j := \beta_j + 1$ ;  $avail := avail - 1$ 
45:   end if
46:   if  $avail = 0$  then return  $\triangleright$  When there is no more space, done
47:   end if
48: end for
49: end procedure

```

**Figure 1:** Sample schedule showing conversion between due dates and release times.**Table 3:** Example set of inputs

Array	Width (W)	Depth (D)	Due Date (d)	Processing Time ($p = W \times D$)
A	2	5	2	10
B	3	5	6	15
C	4	3	3	12
D	5	4	6	20
E	6	2	3	12

Table 4: r_j , δ_j , and $h(j)$ for each array. Arrays sorted by non-decreasing d_j . $d_{max} = \max_j(d_j)$

Array	A	C	E	B	D
j	1	2	3	4	5
d_j	2	3	3	6	6
r_j	4	3	3	0	0
δ_j	8	8	6	6	5
$h(j)$	2	2	2	3	4

bit lanes). Also, regular multiprocessor tasks can be split arbitrarily, but array elements are indivisible. For instance, an array with 17-bit elements can use 17, 34, or 51 bits of a 64-bit bus (i.e., transferring one or more elements), but not 20 bits (i.e., transferring parts of the elements). To schedule indivisible elements, we modified the largest-remainder method to only allocate in multiples of the bitwidth (Line 38). The remainders can then be greater than one, but always less than the bitwidth of the element. Without this modification, the overhead for organizing the data with logical-shift and bitwise operations would be prohibitive.

A small example is presented here. Table 3 lists the characteristics of arrays. The total processing time, p_{tot} (total number of bits in all of the arrays) is 69. Ideally, $C_{max} \times m$ is as close to p_{tot} as possible to ensure there is the smallest amount of wasted bandwidth. Therefore, we compute **bandwidth efficiency** as:

$$B_{eff} = \frac{p_{tot}}{C_{max} \times m} \quad (1)$$

So, the ideal case is a value of 1 (or 100%), which means that the accelerator is fully utilizing the bandwidth.

A completely naive method would be to sort the arrays by increasing due date and place one element of each array into each 8-bit slot of memory. The resulting diagram is shown in Fig. 3. Array D would arrive 13 cycles after its due date of $d_D = 6$ ($L_{max} = 13$). The efficiency of this layout is $\frac{69}{19 \times 8} = 45.4\%$. An improvement would be to pack as many elements of an array as possible onto the bus at once. This homogeneous packing is more dense but still fairly naive. This layout is shown in Fig. 4. In this layout, $L_{max} = L_D = 7$ and the efficiency is $\frac{69}{13 \times 8} = 66.3\%$.

Listing 1 Sample host function for organizing arrays into the layout

```

void pack(int* A, int* B, int* C, int* D, int* E,
         unsigned char* out) {
    unsigned char curr;
    // 0 : C, B
    curr = ((*C++) & C_MASK) << (B_WIDTH + 1);
    curr |= ((*B++) & B_MASK) << (1);
    *out++ = curr;
    // 1 : D, B
    curr = ((*D++) & D_MASK) << (B_WIDTH);
    curr |= ((*B++) & B_MASK);
    *out++ = curr;
    // ...
    for (unsigned int t = 0; t < 2; t++) {
        // 7-8 : D, B
        curr = ((*D++) & D_MASK) << (B_WIDTH);
        curr |= ((*B++) & B_MASK);
        *out++ = curr;
    }
}

```

To convert this problem into one the algorithm can solve, the release times r should be computed from the set of due dates d as $r_j = d_{max} - d_j$. We show this and the computations for the maximum bits per cycle for an array, $\delta_j = \lfloor m/W_j \rfloor \times W_j$, and for the heights, $h(j) = p_j/\delta_j$ in Table 4. The set of unique release times is $r = \{0, 3, 4\}$. Using this set, the arrays must be sorted into groups, R_k , based on r_k (Line 2). Within each R_k , the arrays are ordered by nonincreasing values of $h(j)$: $R_0 = \{D, B\}$ ($r_j = 0$), $R_1 = \{C, E\}$ ($r_j = 3$), and $R_2 = \{A\}$ ($r_j = 4$). Then, each R_k is processed in order. The algorithm executed on the example arrays is shown in Fig. 2.

Each large box on the left side of Fig. 2 shows the current working group, R_k , and the current ready time, r_j , in the top left corner. Inside this box is each array currently ready to be processed. The curved arrows indicate the end of an iteration of the while loop (Line 5) of Algorithm 1.1. The reason for the value of τ in that iteration is listed on the arrow on the right side along with which array elements get placed into the layout at that interval. When the next R_k group is ready, the arrays which are not yet fully processed are added to the new group. At the end of the procedure, the final layout must be reversed to target L_{max} , as shown in Fig. 5. The latest arrays arrive only 3 cycles after their due dates ($L_{max} = 3$). The efficiency is now $\frac{69}{9 \times 8} = 95.8\%$, wasting only 3 bandwidth bits.

5 CODE GENERATION

Because all array details are statically known, we execute Iris during the compilation part to determine the data layouts and generate the necessary functions for decoding them into the accelerator.

Host-Side Organization. To transfer data from the CPU using the proposed layout, the host must aggregate the arrays into the layout efficiently. The procedure for organizing the data, given pointers to all of the input arrays and a pointer to allocated memory the size of the layout ($m \times C_{max}$), is as follows. We create each layout cycle using the machine-word-size of the host. For example, if the layout is for a 256-bit bus and the host uses a 64-bit word size, we organize the memory line in four adjacent `uint64` elements. The generator iterates over each array assigned to each cycle and logical-shift-left the next element of that array into the current word. When this word is full, it places it in its appropriate memory location and

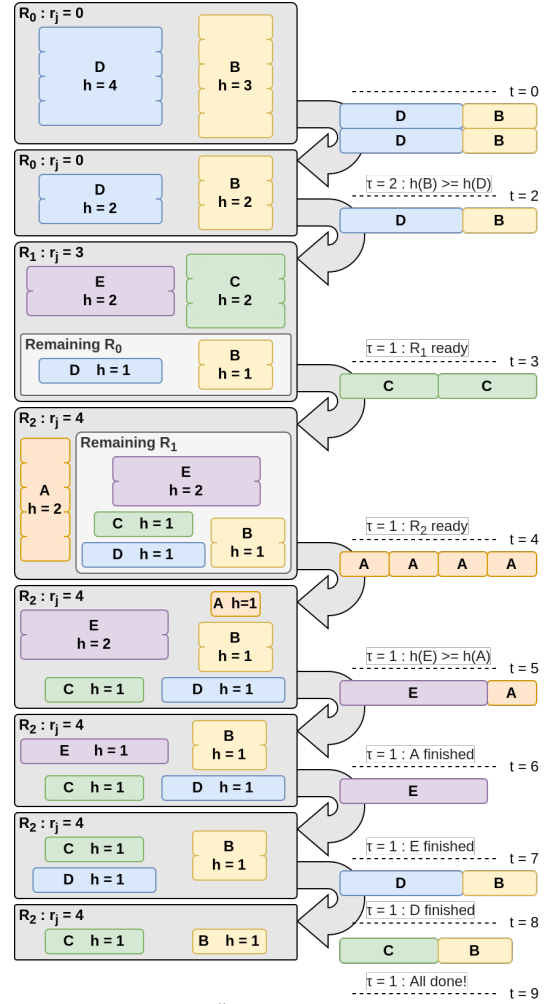


Figure 2: Our process of “scheduling” arrays into the layout.

starts the next one. After placing each array element, it increments its array pointer such that the next element will be inserted in the next steps. When an element spans across words, it shifts in the remaining bits to the top of the next word. The C function for the data organization of the example in Section 4 is shown in Listing 1. `X_WIDTH` and `X_MASK` constants represent the width of the array and a bitmask of the appropriate width, respectively. `(*X++)` will get the value at pointer `X`, and then post-increment it. In the case of $\tau > 1$ (e.g., in cycles 7-8), we use a for loop to create the same layout over several cycles. This simple function can be automatically generated from the given layout.

Accelerator-Side Decoding. Once the data are transferred into memory accessible by the accelerator (e.g., Xilinx HBM), the accelerator must interpret the data. We implement specialized modules to exchange data between memory and the appropriate streams.

The data-read module must have an initiation interval of 1 to maintain maximum bandwidth utilization. To achieve this, enough local memory ports must be available to store all data elements on the bus at once. For data elements that only appear once in any cycle of the layout, the stream interface or a private local memory (PLM) is sufficient. However, if two or more elements from the

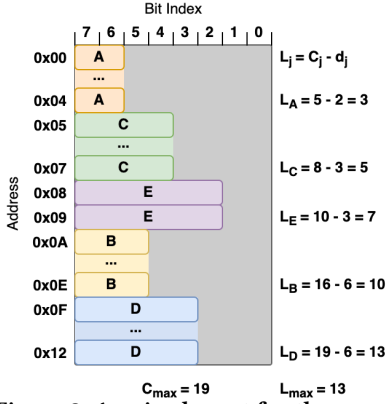


Figure 3: A naive layout for the example arrays.

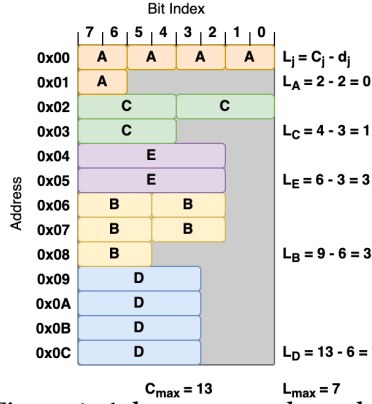


Figure 4: A homogeneously packed naive layout for the example arrays.

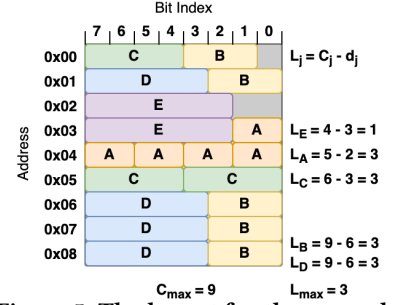


Figure 5: The layout for the example arrays generated by our method.

Listing 2 Sample HLS module for a data read module to decode the layout for an accelerator (Trimmed for brevity)

```

#define BUSWIDTH 8
#define A_FIFO_DEPTH 3
#define C_FIFO_DEPTH 1
void read_data(ap_uint<BUSWIDTH>* in_buf,
               hls::stream<ap_uint<A_WIDTH>>& dataA,
               // ...
               hls::stream<ap_uint<E_WIDTH>>& dataE) {
    ap_uint<BUSWIDTH> elem;
    dataA_t tmpA[A_FIFO_DEPTH];
    dataC_t tmpC[C_FIFO_DEPTH];
    for(unsigned int t = 0; t < 9; t++){
#pragma HLS pipeline II=1
        elem = in_buf[t];
        if (t == 0) {
            dataC << elem.range(7, 4);
            dataB << elem.range(3, 1);
        } else if (t == 1) {
            dataD << elem.range(7, 3);
            dataB << elem.range(2, 0);
            // ...
        } else if (t >= 7 && t <= 8) {
            dataA << tmpA[0];
            dataD << elem.range(7, 3);
            dataB << elem.range(2, 0);
        }
    }
}

```

same array are present in a single cycle, we need extra memories to temporarily store these elements to free up the bus quickly, rather than waiting for several cycles to read each element off the bus. For instance, if at most four elements from array A are on the bus in a single cycle, we need four write ports [17]. This can be implemented as a three-element shift-register where $A[i]$ is written straight to the destination, and $A[i+1]$, $A[i+2]$, and $A[i+3]$ are parallel-loaded into the shift-register, and each successive cycle has the next element written to the destination. However, if more elements of A are on the bus in these three cycles, additional depth might be needed in the shift-register. The maximum depth of the shift-register for an array is determined during layout creation by a running sum over each schedule interval.

Listing 2 shows a sample read module for the example layout written in Xilinx-style HLS code, using the HLS library for *arbitrary*

Table 5: Set of inputs for our test accelerators

Accelerator	Array	Width	Depth	Due Date (d)
Inv. Helmholtz	u	64	1331	333
	S	64	121	31
	D	64	1331	363
Matrix Multiplication	A	64	625	157
	B	64	625	157

Precision Types (`ap_uint`). This module sends each element to a stream for the appropriate array. Downstream dataflow modules can begin execution as soon as the first elements are sent. The constant X_WIDTH values are also the bitwidths of their respective arrays. The HLS tool estimates a latency of 11 clock cycles with only 29 flip-flops and 194 LUTs. For the naive read module (Fig. 3) HLS estimates a latency of 43 cycles and uses 54 flip-flops and 452 LUTs. Thus, we improve both latency and resource requirements.

6 EVALUATION

We implemented a prototype of Iris in Python which receives the input (e.g., bus bitwidth and array details) as a JSON file. This file can be automatically generated by reading array details from the kernel during HLS. To evaluate our method, we analyze several layouts generated for two real accelerators. In all cases, we use $m = 256$ to target the real bus width of the HBM on the Alveo u280.

Inverse Helmholtz. The work in [22] aims at deploying the Inverse Helmholtz operator on the Alveo u280. This operator is the building block of a computational fluid dynamics application. Due to the physical nature of the values, each array element uses 64 bits (double). In [22], the authors examine different strategies for optimizing the data transfers but using the *packed naive* approach for the HBM. Table 5 shows the depths and due dates of each array. d_S and d_u are simply the earliest time by which these arrays can feasibly be finished. D is needed later than u and S , so d_D is the earliest time by which u and S could both be feasibly finished by.

A naive layout following the pattern in Fig. 4 has an efficiency of 99.8% and $L_{max} = 364$. Our layout is 99.9% efficient, using one less cycle, and $L_{max} = 333$. Because these data widths are all evenly divisible into the bus-width, the metrics for our layout are nearly the same. However, we reduce the FIFO depth from 998 for u and D and 90 for S to 666 for u (-33%), 636 for D (-36%), and 30 for S (-67%). In the naive layout, four elements of each array are nearly always sent on

Table 6: Layout metrics with varied δ/W (Inv. Helmholtz)

		δ/W				
		Naive	4	3	2	1
Efficiency		99.8%	99.9%	98.8%	97.9%	51.1%
C_{max}		697	696	704	711	1361
L_{max}		364	333	341	348	998
FIFO	u	998	666	667	665	0
Depth	S	90	30	30	15	0
	D	998	636	631	620	0

Table 7: Layout metrics with varied W (Matrix Multiply)

(W_A, W_B)	(64, 64)		(33, 31)		(30, 19)	
	Naive	Iris	Naive	Iris	Naive	Iris
Efficiency	99.5%	99.8%	92.5%	98.9%	93.5%	97.3%
C_{max}	314	313	236	225	206	201
L_{max}	157	156	79	68	49	44
FIFO	A	468	312	535	467	546
Depth	B	468	312	546	478	576

the bus at a time. In our layout, instead, the three arrays are often interleaved together in the same cycle, relieving the contention pressure on the FIFOs. This improvement is important since BRAMs are usually a limiting factor for data-intensive applications.

We can even vary the maximum number of times an array can have elements on the bus in one cycle by reducing δ to a lower multiple of the bit-width. Table 6 summarizes the results when constraining the arrays as such. For $\delta/W > 1$, we slightly improve FIFO depth as δ/W decreases, along with slight efficiency and L_{max} degradation. When $\delta/W = 1$, the efficiency drops to 51.1% because there are only 3 arrays, so it is impossible to fill the entire bandwidth if they are all only allowed to have one element on the bus at a time. However, we eliminate the need for extra write-port FIFOs since only one element must be written to any array at a time. If a design is having difficulty due to area constraints, and not having data-transfer bottleneck issues, this layout may be useful.

Matrix Multiplication. We also test layouts for Matrix Multiplication, which is popular in many tensor-based applications [10, 2]. Its inputs are summarized in Table 5. The due dates for this application are both as soon as possible, as both inputs are needed at the same time. With $W = 64$ again, the naive layout and our layout perform nearly identically, with our layout only having slightly better L_{max} and FIFO depth. However, when we vary the bitwidths with custom precision data types, Iris achieves better results. Results for the matrix multiply layout are summarized in Table 7.

With custom precision data types, it is difficult to fit neatly into the bus width. Our layout algorithm determines how to better utilize the bandwidth without sacrificing performance. In the case of 64-bit data, the schedule length is reduced by one cycle, but the memory resources are reduced by 33%. For 33- and 31-bit widths, the schedule length, which directly correlates to data transfer time, is reduced by 5% and the overall FIFO memory resources are reduced by 13%. Finally, for the 30- and 19-bit width, the schedule length is reduced by 2% and the memory resources are reduced by 8%.

7 CONCLUSION

This work presented *Iris*, an algorithm designed to automatically create an efficient data layout that maximizes the use of the available

bandwidth. *Iris* was able to achieve higher bandwidth efficiency and lower lateness L_{max} for various accelerators. Also, the solutions created by *Iris* use fewer FPGA resources for the data read module, particularly in the case of the data FIFOs necessary to read from the bus every cycle. Also, as *Iris* is an automatic process, this relieves the designer of a huge manual effort and can even support rapid design space exploration when using custom data types.

ACKNOWLEDGMENTS

This work was partially funded by the EU Horizon 2020 Programme under grant agreement No 957269 (EVEREST).

REFERENCES

- [1] Grey Ballard, Jack Weissenberger, and Luoping Zhang. 2021. Accelerating neural network training using arbitrary precision approximating matrix multiplication algorithms. In *ICPP Workshops Article 16*, 1–8.
- [2] Tal Ben-Nun and Torsten Hoefer. 2019. Demystifying parallel and distributed deep learning: an in-depth concurrency analysis. *ACM Comput. Surv.*, 52, 4, Article 65, (Aug. 2019).
- [3] Yuze Chi, Licheng Guo, and Jason Cong. 2022. Accelerating SSSP for power-law graphs. In *FPGA*.
- [4] Young-kyu Choi, Yuze Chi, Weikang Qiao, Nikola Samardzic, et al. 2021. HBM connect: high-performance HLS interconnect for FPGA HBM. In *FPGA*, 116–126.
- [5] Young-kyu Choi, Yuze Chi, Jie Wang, Licheng Guo, et al. 2020. When HLS meets FPGA HBM: benchmarking and bandwidth optimization. (2020).
- [6] Jason Cong, Peng Wei, Cody Hao Yu, and Peipei Zhou. 2017. Bandwidth optimization through on-chip memory restructuring for HLS. In *DAC*.
- [7] William J. Dally, Yatish Turakhia, and Song Han. 2020. Domain-specific hardware accelerators. *Commun. ACM*, 63, 7, (June 2020), 48–57.
- [8] Maciej Drozdowski. 1996. Real-time scheduling of linear speedup parallel tasks. *Information Processing Letters*, 57, 1, 35–40.
- [9] Fabrizio Ferrandi, Vito Giovanni Castellana, Serena Curzel, Pietro Fezzardi, et al. 2021. Invited: bambu: an open-source research framework for the high-level synthesis of complex applications. In *DAC*.
- [10] Azzam Haidar, Stanimire Tomov, Jack Dongarra, and Nicholas J. Higham. 2018. Harnessing GPU tensor cores for fast FP16 arithmetic to speed up mixed-precision iterative refinement solvers. In *SC*.
- [11] Tom Hogervorst, Răzvan Nane, Giacomo Marchiori, Tong Dong Qiu, et al. 2021. Hardware acceleration of high-performance computational flow dynamics using high-bandwidth memory-enabled field-programmable gate arrays. *ACM TRETS*, 15, 2, Article 20, (Dec. 2021).
- [12] Dounia Khaldi and Barbara Chapman. 2016. Towards automatic hbm allocation using llvm: a case study with knights landing. In *LLVM-HPC*, 12–20.
- [13] Ulrich Kohler and Janina Zeh. 2012. Apportionment methods. *The Stata Journal*, 12, 3, 375–392.
- [14] Hongyu Miao, Myeongjae Jeon, Gennady Pekhimenko, Kathryn S. McKinley, et al. 2019. StreamBox-HBM: stream analytics on high bandwidth hybrid memory. In *ASPLOS*.
- [15] Mahdi Nazemi and Massoud Pedram. 2018. Deploying customized data representation and approximate computing in machine learning applications. In *ISLPED*.
- [16] Christian Pilato, Stanislav Bohm, Fabien Brocheton, Jeronimo Castrillon, et al. 2021. EVEREST: a design environment for extreme-scale big data analytics on heterogeneous platforms. In *DATE*, 1–6.
- [17] Christian Pilato, Paolo Mantovani, Giuseppe Di Guglielmo, and Luca P. Carloni. 2017. System-level optimization of accelerator local memory for heterogeneous systems-on-chip. *IEEE TCAD*, 36, 3, 435–448.
- [18] Sartaj K. Sahni. 1976. Algorithms for scheduling independent tasks. *J. ACM*, 23, 1, (Jan. 1976), 116–127.
- [19] Nimish Shah, Paragkumar Chaudhari, and Kuruvilla Varghese. 2018. Runtime programmable and memory bandwidth optimized fpga-based coprocessor for deep convolutional neural network. *IEEE TNNLS*, 29, 12.
- [20] John Shalf. 2020. The future of computing beyond Moore’s Law. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 378, 2166, (Jan. 2020).
- [21] Runbin Shi, Kaan Kara, Christoph Hagleitner, Dionysios Diamantopoulos, et al. 2021. Exploiting HBM on FPGAs for data processing. *ACM TRETS*, (Oct. 2021).
- [22] Stephanie Soldavini, Karl F. A. Friebe, Mattia Tibaldi, Gerald Hempel, et al. 2022. Automatic creation of high-bandwidth memory architectures from domain-specific languages: the case of computational fluid dynamics. *ACM TRETS*, (Sept. 2022).