

CRIMSON: Compute-Intensive Loop Acceleration by Randomized Iterative Modulo Scheduling and Optimized Mapping on CGRAs

Mahesh Balasubramanian[✉], Member, IEEE, and Aviral Shrivastava, Member, IEEE

Abstract—Coarse-grain reconfigurable arrays (CGRAs) are emerging accelerators that promise low-power acceleration of compute-intensive loops in applications. The acceleration achieved by CGRA relies on the efficient mapping of the compute-intensive loops by the CGRA compiler, onto the CGRA architecture. The CGRA mapping problem, being NP-complete, is performed in a two-step process, namely, scheduling and mapping. The scheduling algorithm allocates timeslots to the nodes of the data flow graph, and the mapping algorithm maps the scheduled nodes onto the processing elements of the CGRA. On a mapping failure, the initiation interval (II) is increased and a new schedule is obtained for the increased II. Most previous mapping techniques use the iterative modulo scheduling (IMS) algorithm to find a schedule for a given II. Since IMS generates a resource-constrained as-soon-as-possible (ASAP) scheduling, even with increased II, it tends to generate a similar schedule that is not mappable. Therefore, IMS does not explore the schedule space effectively. To address these issues, this article proposes CRIMSON, compute-intensive loop acceleration by randomized IMS and optimized mapping technique that generates random modulo schedules by exploring the schedule space, thereby creating different modulo schedules at a given and increased II. CRIMSON also employs a novel conservative test after scheduling to prune valid schedules that are not mappable. From our study conducted on the top 24 performance-critical loops (run for more than 7% of application time) from MiBench, Rodinia, and Parboil, we found that previous state-of-the-art approaches that use IMS, such as RAMP and GraphMinor could not map five and seven loops, respectively, on a 4×4 CGRA, whereas CRIMSON was able to map them all. For loops mapped by the previous approaches, CRIMSON achieved a comparable II.

Index Terms—Coarse-grained reconfigurable arrays (CGRAs), compiler, modulo scheduling, randomized scheduling.

I. INTRODUCTION

COMPUTING demands in human society continue to climb. Today there are numerous devices that collect, process, and communicate data from multiple sources, such

Manuscript received August 6, 2020; accepted August 31, 2020. Date of publication September 7, 2020; date of current version October 27, 2020. This work was supported in part by the National Science Foundation under Grant CSN 1525855 and Grant CCF 1723476 CAPA, and in part by the NSF/Intel Joint Research Center for Computer Assisted Programming for Heterogeneous Architectures. This article was recommended by Associate Editor P. Pande. (*Corresponding author: Mahesh Balasubramanian.*)

Mahesh Balasubramanian is with the School of Computing Informatics Decision and Systems Engineering, Arizona State University, Tempe, AZ 85287 USA (e-mail: mbalasu2@asu.edu).

Aviral Shrivastava is with the School of Computing Informatics, Decision and Systems Engineering, Arizona State University, Tempe, AZ 85048 USA. Digital Object Identifier 10.1109/TCAD.2020.3022015

as the Internet, cyber-physical and autonomous systems, sensor networks, etc., [1]. Extracting intelligent and actionable information from all these data—whether or not done by machine learning—is extremely compute-intensive, and often times limited by power, thermal, and other resource constraints [2]. Efficiency in the execution of these functionalities can be achieved by using application-specific integrated circuits (ASICs). However, they suffer from high production costs, and they quickly become obsolete as applications and algorithms evolve. Another promising alternative is field programmable gate arrays or FPGAs, but they lose efficiency in providing bit-level configurability, which is essential for their primary purpose—prototyping [3]. Coarse-grained reconfigurable architectures or CGRAs provide a very good middle ground with coarse-grain configurability (word and arithmetic operator-level), without much loss in power-efficiency when compared to ASICs [4]. As a result, there is a renewed surge in the application of CGRAs for compute-intensive workloads, including machine learning, embedded systems, and vision functionalities [5]–[7].

As shown in Fig. 1, CGRAs are simply an array of processing elements (PEs) arranged in a 2-D grid. The PEs are just bare arithmetic logic units (ALUs) that can receive inputs from the neighboring PEs, from the *Data Memory*, and its own small set of registers. Every cycle, the PEs receive an instruction from the *Instruction Memory*, and write the results to the output buffer, local register file, and/or the Data Memory. CGRA-based execution is highly parallel (16 operations can be executed simultaneously on a 4×4 CGRA) and power-efficient because instructions are in the predecoded form. There is no extensive pipeline for instructions to go-through before and after execution, and the PEs can exchange operands directly rather than going through the register files. Some of the early works on CGRA architecture include ADRES [3], PADDI [8], Kressarray [9], MATRIX [10], Morphosys [11], and Remarc [12]. ADRES [3] which showed CGRAs to be promising power-efficient accelerators with power efficiency of 60 Giga Operations per Watt (GOps/W) using a 32 nm technology.

The most common way to use CGRAs is to employ them as co-processors to CPU cores or processors, to speed up and power-efficiently execute compute-intensive applications—similar to GPUs. The execution of compute-intensive loops in the application can then be “offloaded” onto these CGRA co-processors, while the rest of the application can still execute




Fig. 1. Typical CGRA architecture consisting of 4×4 PEs connected in a 2-D mesh. Every cycle, each PE gets an instruction from the Instruction Memory, and can operate on the outputs of neighboring PEs and/or the Data Memory.

on the CPU. This heterogeneous computing paradigm requires compiler support to map compute-intensive loops of the application onto the PEs of the CGRA. Since the execution of a loop happens by software-pipelining on a CGRA, the objective of mapping is to layout the nodes of the data flow graph (DFG) onto a graph of the CGRA extended in time, so that the nodes can exchange the operands through the interconnection among the PEs and achieve correct and efficient execution. The repetition interval of the mapping (the time at which the next iteration of the loop can start) is called the initiation interval (II) and is the metric that determines the quality of mapping. Many techniques have been proposed to solve NP-complete [13] mapping problem of CGRAs efficiently [13]–[20]. Most of the newer methods work in these four steps: 1) create the DFG of the loop, and estimate the minimal II; 2) create the CGRA graph unrolled II times; 3) schedule the nodes of the loop onto the CGRA graph; and 4) map the nodes onto the PEs at their scheduled timeslots such that the dependencies among the nodes can be transferred through the connectivity among PEs. In case a valid mapping is not found, the II is increased, and steps from 2) onward are executed again. This process is repeated until a valid mapping is found. A mapping failure can occur in the fourth step due to the limited connectivity among the PEs of the CGRA, and because of the need to map new routing nodes. Routing nodes occur when dependent operations are scheduled in noncontiguous timeslots. In this case, the operands need to be routed from the PE on which the source operand is mapped, to the PE on which the destination operation is mapped. This is commonly referred to as the routing problem. One solution is to route the operands through the PEs in the intermediate timeslots. Since routing and mapping attempts often fail, existing CGRA mapping techniques have heavily focused on solving the problem encountered in the mapping and routing step. For example, [16], [17] route dependencies via PEs, [19] routes dependencies through the registers inside the PEs, [18] overlaps the routing paths carrying the same value, and [13] uses recomputation as an alternative to routing. MEMMap [21] routes dependent operations via Data Memory by adding store and load nodes. RAMP [20] proposes a heuristic to explore all the different routing options. However, all the previous approaches use the same iterative

TABLE I
ON EVALUATING 24 APPLICATIONS OF THE TOP THREE BENCHMARK SUITES ON A 4×4 CGRA, WE FIND THAT IMS-BASED RAMP WAS UNABLE TO MAP 5 OF THE LOOPS AND IMS-BASED GRAPHMINOR WAS UNABLE TO MAP 7 OF THE LOOPS. THE “X” IN THE TABLE DENOTES AN II WAS NOT OBTAINED EVEN AT A MAXIMUM II OF 50. THE MII IN THE TABLE DENOTES THE MINIMUM II, WHICH IS THE MAXIMUM OF EITHER RESMII OR RECMII

Suites	Loops	4x4		
		MII	RAMP-II	GraphMinor-II
MiBench	bitcount	3	3	6
	susan	2	3	3
	sha	3	3	3
	jpeg1	3	X	X
	jpeg2	2	X	X
Rodinia	kmeans1	2	2	2
	kmeans2	2	2	2
	kmeans3	2	2	2
	kmeans4	2	2	2
	kmeans5	2	2	2
	lud1	2	2	2
	lud2	2	2	2
	b+tree	2	2	2
	streamcluster	2	2	2
	nw	2	2	2
	BFS	2	2	2
	hotspot3D	5	X	X
Parboil	backprop	5	X	X
	spmv	3	3	3
	histo	2	2	2
	sad1	2	2	X
	sad2	2	2	2
	sad3	2	2	X
	stencil	4	X	X

modulo scheduling (IMS) [22] to find a valid schedule—and therein lies the problem.

The problem with IMS is that it generates a resource-constrained, as soon as possible (ASAP) schedule of nodes onto the CGRA PEs. When a mapping is not found, the traditional mapping techniques increase the II, and return to the scheduling step. The generated schedule does not change much, even when more resources are added toward the bottom of the CGRA graph. The resource-constrained ASAP schedule will be almost identical to the one obtained before, and the extra resources are not used! As a result, the mapping algorithm keeps on exploring the schedule space with the same schedule, and often no mapping can be found, even after huge increases in the II. Table I shows the evaluation of the 24 performance-critical loops from MiBench, Rodinia, and Parboil on a 4×4 CGRA, while being executed on the state-of-the-art IMS-based mapping algorithms, GraphMinor [18], and RAMP [20]. We can see that state-of-the-art RAMP was unable to find a valid mapping for five loops and GraphMinor was unable to find a valid mapping for seven loops on evaluation up to a maximum II = 50. One major observation was that, when these previous algorithms find a mapping, they achieve a very good II, but when the mapping fails, they are unable to map the loops even with II increments up to 50. For example, in loop *jpeg1*, while the minII was 3, both the techniques were unable to map the loop, even when the II was increased to 50.

Thus, the main problem in IMS is the absence of randomness in the scheduling algorithm. As a result, even when

the II is increased, the same schedule is generated without obtaining a valid mapping. Hence, this creates a need for an enhanced scheduling algorithm that explores the schedule space to increase the mappability of the compute-intensive loops. A more detailed explanation with a motivating example is given in Section IV. In this article, we propose compute-intensive loop acceleration by randomized IMS and optimized mapping on CGRA (CRIMSON). Instead of just using the resource constrained ASAP (RC_ASAP) schedule, CRIMSON generates both the RC_ASAP and resource constrained ASAP (RC_ALAP) schedules for all the nodes of DFG, similar to the concept of mobility used in high-level synthesis (HLS) [23]. CRIMSON then chooses a random time between RC_ASAP and RC_ALAP as the scheduling time for each node. As a result, every time a “new” schedule is obtained, CRIMSON is able to effectively explore the schedule space. CRIMSON also incorporates a novel conservative feasibility test after the scheduling step to check the mappability of the obtained schedule. This conservative test makes sure that the generated schedule will be mappable even after the addition of the new routing nodes, thereby rendering feasibility by quickly weeding out some unmappable schedules, and saving time. Among the 24 performance-critical (that account for more than 7% of execution time of the application) loops from MiBench, Rodinia, and Parboil, our approach CRIMSON was able to map all the loops for various CGRA sizes ranging from 4×4 to 8×8 . Our approach CRIMSON achieved a comparable II for the loops which were mappable by RAMP.

II. BACKGROUND AND TERMINOLOGY

CGRA compilers in general first create the DFG $D = (V, E)$ of a compute-intensive loop, where V refers to the nodes of the loop and E refers to the edges (data dependencies between nodes) in the graph. The constructed DFG is then software pipelined using IMS [22], where each node is assigned a schedule-time at which it should be executed.

Fig. 2(a) shows the DFG of a loop, and Fig. 2(b) shows the target CGRA architecture. The schedule of the DFG nodes are shown in Fig. 2(c), considering the resource and the recurrence cycle constraints. After scheduling, the nodes are then mapped onto the PEs of CGRA such that the dependent operands can be routed from the PE on which the source operation is mapped to the PE on which the destination operation is mapped through either registers, memory, or paths in the CGRA graph. A register can be used to route operands when the dependent operation is mapped to the same PE as the source operation. Memory can be used to route operands, but that requires inserting additional load and store instructions. A path is a sequence of edges and nodes in the CGRA graph that connect two PEs. In the simplest case, a path is just a single edge.

For simplicity, the mapping shown in Fig. 2(d) uses only edges to route dependencies. In this mapping, node a of iteration i (shown in dark color) is mapped to PE2 at time T , nodes b , and c are mapped to PEs, PE1 and PE2, respectively, at $T+1$. Similarly, nodes e, f of i th iteration are mapped in PE1 and PE2, respectively, at $T+2$. Node g of i th iteration




Fig. 2. (a) DFG of an application loop. (b) 1×2 CGRA target architecture. (c) IMS schedule of nodes of DFG. The x -axis is the modulo time. (d) Mapping of the scheduled nodes on the time-extended CGRA (TEC).

is mapped at PE1 at $T+3$. It can also be observed that a and g are mapped at T and $T+3$ in PE1 and PE2. Modulo schedule repeats itself every II cycles, in this case $II = 3$. The node g at T (shown in gray) is from $(i-1)$ th iterations. Likewise, the node a mapped at $T+3$ is from $(i+1)$ th iteration (shown in green). Based on the schedule, which considers the recurrences while scheduling, mapping a in PE2 satisfies the recurrence constraint of $f \rightarrow a$. i.e., the value of f at i th iteration can be routed to a at $(i+1)$ th iteration. In modulo scheduling, the interval in which successive instructions can begin execution is called the II [22]. II is considered as the performance metric for DFG mapping onto CGRA, as the total number of cycles required to execute a loop will be proportional to the II.

III. RELATED WORKS

CGRAs have been a luring accelerator option owing to their high performance and high power-efficiency. The ADRES CGRA [3] demonstrated to operate at 60 GOPs/W. The high power-efficiency of CGRA is due to instructions being in predecoded format with no long pipelines before and after execution, and the fact that PEs can exchange operands directly without going through a centralized register file. CGRAs rely on the compiler to map loops onto the PEs. Some CGRA application mapping techniques use generic algorithms [24]–[26] like genetic algorithms or simulated annealing [14], [27], [28] to explore the various possible mappings and come up with a valid one. While these genetic algorithms and simulated annealing come with inherent randomness, these methods take exorbitantly long times to find a valid mapping, since they have no conception of the DFG and CGRA graph structures.

Some of the older application-specific compilation techniques like DRESC [14] attempt to solve the scheduling and mapping problems together in one shot. However, this is inefficient, since these algorithms may spend a lot of time exploring a mapping, when even the schedule is infeasible. A valid schedule is a prerequisite of a valid mapping, and since scheduling is quite quick [22], it makes sense to first find a valid schedule, and then explore mapping solutions only for those schedules. As a result, most modern approaches separate the scheduling and mapping steps. When a mapping attempt fails due to limited connectivity or additional routing requirements, the II is increased, and a new schedule for




Fig. 3. (a) DFG of an application loop. (b) 2×2 CGRA target architecture. (c) Column 1 shows the nodes in the DFG and column 2 shows an IMS schedule for the nodes at $II = MII = 3$. (d) Mapping algorithm tries to map the nodes scheduled, but fails due to additional routing nodes “r” required to route nodes *f* and *i*. Failure to find a valid mapping, the II is increased to 4 and IMS is called again to schedule the nodes based on the workflow given in Fig. 4. (e) IMS schedule for an increased II ($II = 4$). (f) Even at an increased II , the mapping algorithm cannot find a valid mapping due to resource constraint at T_{i+1} which is not resolved at $II = 4$ and will not be resolved in any further increase in II .

this increased II is obtained followed by another attempt on mapping. This scheduling and mapping is repeated until a valid mapping is obtained.

Since mapping is harder, previous works concentrate on solving the mapping and routing issues. EPIMap [13] uses recomputation of some nodes to solve the routing problem. REGIMap [19] uses register file in the PEs to route the dependent operations, whereas MEMMap [21] uses the Data Memory to route the dependencies. More recent techniques like, RAMP [20] presents a heuristic to choose among a variety of routing options to try for unmapped nodes, CASCADE [29], on the other hand, increases data throughput by decoupling the memory accesses and the execution. Even though all these techniques have different mapping and routing strategies, they use the same scheduling algorithm, namely—IMS [22] proposed by Rau, for VLIW architectures, uses RC_ASAP approach to schedule the nodes of DFG. The problem is that, even when II is increased, IMS generates the same schedule, and is unable to explore the newly created scheduling space created by increased II .

Instead of taking a conventional ASAP/ALAP scheduling approach, EMS [16] proposes an alternative approach, where the nodes of the recurrent cycles are lifted or lowered on the time axis by assigning stages. These stages can consist of multiple schedule times. When an operation stage is reassigned based on the placement of its predecessor, all the dependent operation stages are also reassigned. However, while the EMS schedule is not ASAP, but it is still not randomized. As a result, the generated schedule for a higher II is very similar (if not the same) to the generated schedule at lower II s. HyCube [30] proposes a mapping technique for a highly connected CGRA that uses multihop multicast path system to communicate data in a single-cycle. In addition, HyCube’s interconnect crossbar switch is a part of the ISA, which makes it power-efficient. Like DRESC [14] approach of integrated scheduling, placement and routing (P&R), HyCube’s *ScheduleAndRoute* schedules and performs P&R in one shot. This faces the same issues as DRESC discussed above. Evidently, HyCube’s single-cycle communication may provide better II , but at the

cost of scalability. Since, the interconnect crossbar selection is a part of HyCube’s instruction set, for higher CGRA sizes HyCube’s instruction becomes longer.

The main contribution of this article is a random IMS algorithm to effectively search the scheduling space and an enhanced application mapping workflow to efficiently find a valid mapping of loops.

IV. MOTIVATING EXAMPLE

Let us consider the DFG of loop to be mapped on a 2×2 CGRA, shown in Fig. 3(a) and (b), respectively. Previous state-of-the-art techniques like RAMP, get a schedule from IMS [22] before mapping the nodes. IMS starts by computing the resource constrained minimum II (ResMII) and recurrence constrained minimum II (RecMII) from the DFG and the architecture description. For the given example in Fig. 3, total nodes = 9 and total resources available = 4. The minimum II (MII) is the maximum of RecMII and ResMII. Therefore for the above example, $MII = \text{ResMII} = \lceil 9/4 \rceil = 3$. After computing the MII, IMS sets the priorities for each node. Priority is a number assigned to each node, which is utilized during scheduling. Based on the height of the node, from the given DFG, the deepest node is given the least priority using depth-first search strategy. For the loop DFG given in Fig. 3(a), node *e* gets priority 0, nodes *d*, and *i* get priority 1, nodes *b*, *c*, *g*, *h* get priority 2 and finally *a* and *f* get priority 3. The nodes with higher priority number are scheduled first with earliest start time. The modulo scheduling starts with $II = MII$ for scheduling the nodes. The CGRA is time-extended, II times and a modulo resource table (MRT) is maintained to check for resource overuse for each timeslot. While trying to schedule each node, resource conflicts are checked. If there is a resource conflict a higher schedule time is tried. For the example DFG, the $II = MII = 3$. Nodes *a* and *f* are scheduled at modulo time 0 (0%3). Nodes *b*, *c*, *g*, and *h* are scheduled at modulo time 1 (1%3) without any resource constraint because there are 4 resources (PEs) at each modulo time. Nodes *d* and *i* are scheduled at modulo time 2 (2%3). Finally, *e* is scheduled at




Fig. 4. Overview of scheduling and mapping workflow of previous techniques.

modulo time 0 (3%). The IMS schedule of nodes [shown in column 1 Fig. 3(c)] at $II = 3$ is shown in Fig. 3(c) column 2.

With this prescribed schedule, mapping algorithms start to map the nodes, but eventually find that a routing node needs to be added to route operation f and i . Due to the unavailability of PEs in that timeslot a routing node cannot be added, as shown in Fig. 3(d). At this juncture, the mapping algorithm increases the II in an effort to find a schedule that is mappable. On increasing the II from 3 to 4, the IMS algorithm is invoked again to get a schedule. Since the priority calculation of IMS is DFG-based, all the nodes get the same priority. Now, IMS algorithm starts to schedule nodes based on the priorities for each node. Nodes a and f are scheduled at modulo time 0 (0%). Nodes b , c , g , and h are scheduled at modulo time 1 (1%). Nodes d and i are scheduled at modulo time 2 (2%) and e is scheduled at modulo time 3 (3%). The IMS schedule for $II = 4$ is shown in Fig. 3(e) column 2. Again, on failure to map, the mapping algorithm increases the II to 5. IMS repeats the process of assigning priorities to the nodes and as seen in $II = 4$, the priorities do not change. Nodes a and f are scheduled at modulo time 0 (0%). Nodes b , c , g , and h are scheduled at modulo time 1 (1%). Nodes d and i are scheduled at modulo time 2 (2%) and finally e is scheduled at modulo time 3 (3%). On comparing the schedules obtained for $II = 3$, $II = 4$, and $II = 5$, it can be seen that only node e has a different schedule time (from $II = 3$ to $II = 4$) and rest of the nodes have the same schedule. Hence, with IMS, it can be seen that an increase in the II does not correspond to a change in modulo schedule time of the nodes.

The algorithm keeps trying to find a valid mapping at higher II even when there is a mapping failure at a given modulo schedule. This process keeps on repeating endlessly. In the workflow of the previous techniques, as shown in Fig. 4, after finding the MinII and obtaining an IMS schedule, the mapping of the nodes begin assuming that the schedule is mappable. There are no mechanism to statically and systematically find the feasibility of the obtained schedule, which results in an infinite loop between the scheduling and the mapping stages.

V. CRIMSON: EFFICIENTLY ACCELERATE LOOPS BY RANDOMIZED ITERATIVE MODULO SCHEDULING AND OPTIMIZED MAPPING

A. Overview

To alleviate the challenges posed by IMS and the previous mapping algorithms, CRIMSON randomizes the schedule time




Fig. 5. Overview of CRIMSON workflow, with addition of RC_ASAP and RC_ALAP computation, randomized scheduling algorithm, and a feasibility test (shaded blocks in the image are proposed by this article).

of each node of the DFG by choosing a time between RC_ASAP and RC_ALAP. Additionally, CRIMSON proposes a change to the previous mapping algorithm workflow Fig. 4 by performing a feasibility test before the actual mapping.

Fig. 5 shows the modification to the traditional IMS-based workflow shown in Fig. 4. CRIMSON modifies the IMS-based mapping workflow by adding RC_ASAP and RC_ALAP computation steps before finding a random schedule. The “Create Randomized Schedule” block uses Algorithms 1 and 2 to find a random modulo schedule time. On a failure to find a schedule, “Create Randomized IMS” block is invoked λ times before increasing the II. When a random modulo schedule is obtained, the feasibility test statically analyzes if the obtained random schedule honors the resource constraints when routing nodes are added. If a schedule is found to be infeasible due to possible resource overuse, a different modulo schedule is tried for the same II. If the random schedule obtained is valid and feasible, then the mapping algorithm is called to add routing nodes and map the scheduled DFG onto the CGRA architecture.

B. Computing Resource-Constrained ASAP and Resource-Constrained ALAP

Algorithm 1 shows the CRIMSON’s randomized IMS. Lines 1 and 2 finds the RC_ASAP from the strongly connected components (SCCs)¹ of the DFG. The RC_ASAP is computed in line 3 of Algorithm 1 as a top-down, depth-first search approach, from the nodes that do not have any incoming edges in the current iteration. After computation of RC_ASAP, RC_ALAP is computed, starting from the nodes that do not have any outgoing edges in the current iteration and in a bottom-up (reverse), depth-first search manner, in line 4 of Algorithm 1.

C. Randomized Scheduling Algorithm

After computing RC_ASAP and RC_ALAP, Algorithm 1 line 5 populates the unscheduled array whereas line 6 sets a boolean scheduled operation to false for all the nodes, which is used in Algorithm 2. For all the unscheduled sorted nodes in the array, a random modulo timeslot is picked by honoring the resource constraints maintained by MRT, in line 10 of Algorithm 1.

¹Getting the list of SCCs ensures that the nodes in recurrence-cycles are scheduled first using *Sort_SCC()* function in line 5.

Algorithm 1: *Rand_Iterative_Mod_Schedule* (Input DFG D , CGRA CA, Input II)

```

1  $D' \leftarrow D$ ;
2  $SCCs \leftarrow Find\_List\_of\_SCCs(D')$ ;
3  $Find\_RC\_ASAP(II, SCcs, CA)$ ;
4  $Find\_RC\_ALAP(II, SCcs, CA)$ ;
5  $unscheduled \leftarrow Sort\_SCcs(SCcs)$ ;
6  $Set\_Scheduled\_op\_false(unscheduled)$ ;
7  $iter \leftarrow 0$ ;
8 while  $unscheduled\_size > 0 \ \& \ iter < threshold$  do
9    $operation \leftarrow unscheduled[0]$ ;
10   $TimeSlot \leftarrow$ 
11     $Find\_Random\_ModuloTime(operation, CA)$ ;
12    if ( $schedule(nodes, TimeSlot)$ ) then
13       $scheduled \leftarrow nodes$ ;
14    else
15      return failure;
16     $unscheduled \leftarrow Subtract(unscheduled, scheduled)$ ;
17     $iter++$ ;
18 if ( $iter == threshold \ \& \ unscheduled\_size > 0$ ) then
19  return failure;
return success;

```

The *schedule()* function in line 11 of Algorithm 1, schedules the node at chosen random timeslot. This *schedule* function sets the schedule time of the current operation and consecutively displaces the nodes that have resource conflicts. Previously scheduled nodes having a dependence conflicts with the current operation are also displaced after updating the RC_ASAP and RC_ALAP based on the current schedule operation. The displaced nodes are added to queue of unscheduled nodes. Similar to the *BudgetRatio* in IMS [22], the *iter* is a high value. On a failure to find a schedule, either due to unscheduled nodes lines 13 and 14 or if the *iter* value is greater than a threshold (lines 17 and 18), Algorithm 1 is invoked again. This is repeated λ times before increasing the II , in an attempt to find a valid schedule. This λ value is not reset for a particular II and used to control the failure due to unmappable schedule or a failure in the mapping step.

Algorithm 2 is called by CRIMSON's randomized iterative modulo schedule (*Rand_Iterative_Mod_Schedule*) Algorithm 1 line 10, to find a random timeslot between RC_ASAP and RC_ALAP. The RC_ASAP and RC_ALAP for a given operation is retrieved in lines 1 and 2 of Algorithm 2. Then, an array of timeslots is constructed using the *op_ASAP* and *op_ALAP*, line 4 of Algorithm 2. The array holds all the timeslots from *op_ASAP* with an increasing value of 1 until *op_ALAP*. If *op_ASAP* is equal to *op_ALAP* then the array size is one with either ASAP or the ALAP time. Each timeslot from the randomized array is checked for the resource constraint using MRT. The first valid timeslot is returned as the modulo schedule time for the operation. Due to the resource conflict if a valid timeslot is not present, there are two things to handle, 1) a timeslot for the operation should be chosen and 2) an already scheduled operation from that timeslot should be displaced. Concern 1) is handled in lines 13–17

Algorithm 2: *Find_Random_ModuloTime* (Operation op , CGRA CA)

```

1  $op\_ASAP \leftarrow get\_RC\_ASAP(op)$ ;
2  $op\_ALAP \leftarrow get\_RC\_ALAP(op)$ ;
3  $sched\_slot \leftarrow \emptyset$ ;
4  $timeslots \leftarrow get\_all\_timeslots(op\_ASAP, op\_ALAP)$ ;
5  $Randomize(timeslots)$ ;
6 while ( $sched\_slot == \emptyset \ \& \ timeslots\_size > 0$ ) do
7    $currTime \leftarrow timeslots[0]$ ;
8   if ( $ResourceConflict(op, currTime, CA)$ ) then
9      $timeslots \leftarrow Subtract(currTime, timeslots)$ ;
10    continue;
11   else
12      $sched\_time \leftarrow currTime$  ;
13 if ( $sched\_slot == \emptyset$ ) then
14   if ( $!Scheduled[op] \ ||$ 
15      $op\_ASAP > Prev\_Sched\_Time[op]$  ) then
16      $sched\_slot \leftarrow op\_ASAP$ ;
17   else
18      $sched\_slot \leftarrow Prev\_Sched\_Time[op] + 1$ ;
18 return  $sched\_slot$ ;

```

of Algorithm 2 where if the nodes has not been scheduled previously, *op_ASAP* is chosen as the schedule, else the previous schedule time of the operation is found and the modulo schedule time is computed using line 17. Concern 2) is addressed in the *schedule()* function in Algorithm 1 line 11, explained earlier. The methods addressing these concerns are similar to IMS implementation.

D. Novel Feasibility Test

Given a valid schedule, it may not be possible to map it because of two main reasons: 1) limited connectivity among the PE nodes and 2) the need to map the extra routing nodes that will be created as a result of scheduling. In a valid schedule dependent operations may be scheduled in noncontiguous timeslots. When this is the case, the operands need to be routed from the PE on which the source operand is mapped, to the PE on which the destination operation is mapped. The operands can be routed using a string of consecutive CGRA interconnections and PEs. These PEs are referred to as routing PEs, and the operation that is mapped on these PEs (just forward the operand from input to output) is called a routing operation. Because of the addition of these routing nodes, the generated schedule may not be mappable. Previous techniques assume that the schedule is mappable and spend a lot of time searching for a mapping when none is available. In order to avoid wasting time in exploring unmappable schedules, CRIMSON adds a conservative feasibility test to prune schedules that can be proven to be unmappable.

The feasibility test examines the random schedule produced, and for each routing resource that will be added in the future, it estimates the resource usage, considering path-sharing [18]. The feasibility test checks if the total number of unique nodes including the routing nodes per timeslot is less than or equal to the number of PEs in that timeslot.

$\text{schedule_nodes}_i + \text{routing_nodes}_i \leq PEs_i$, where i is the modulo timeslot. This feasibility check is performed for all the II timeslots. The mapping algorithm is invoked only for schedules that are feasible, unlike the previous approaches such as RAMP [20], where the mapping algorithm is invoked even for infeasible schedules. Since the time complexity of such mapping algorithms is high [time complexity of RAMP is $\mathcal{O}(N^8)$, where $N = n * m$, and “ n ” is the total nodes in the loop DFG, and “ m ” is the size of the CGRA], invoking them for infeasible schedules is counter productive. The feasibility test reduces the overhead incurred by the mapping algorithm by pruning the infeasible schedules.

E. Determining the λ Value

With every failure in the feasibility test a new schedule is obtained for a given II. The number of times a schedule is obtained for a given II is controlled by the λ value. The scheduling space that can be explored for a given II is calculated by the product of the total nodes in the DFG, the size of the CGRA, and the II, given in (1). A brute force exploration of the schedule space is time consuming. Lower λ values may increase the II prematurely, by superficial exploration of schedule space, whereas higher λ values increase the compilation time, due to elaborate exploration of the schedule space. Due to the randomness in the scheduling algorithm, a feasible schedule may be obtained faster by chance even for a higher λ value. The λ value is computed using

$$\lambda = \text{exploration_factor} \times n \times m \times \text{II} \quad (1)$$

where, “ n ” is the total number of nodes in the loop DFG, “ m ” is the size of the CGRA, and *exploration_factor* is the percentage of the schedule space that is to be explored. The *exploration_factor* is a user defined parameter. II is also one of the parameters that determines the λ value in (1), which means that a new λ is computed for each II. When the II is increased, the scheduling space is also increased therefore the scope of exploration gets broadened. A detailed discussion on the effects of *exploration_factor* on the scheduling time and II is given in Section VI-E.

F. Running Example

Fig. 6 shows the working of CRIMSON’s randomized iterative modulo schedule algorithm for the DFG and CGRA architecture shown in Fig. 6(a) and (b).² Instead of assigning a priority based on height like IMS, each node in DFG is assigned two times namely, RC_ASAP and RC_ALAP, which constitutes a good lower and upper bound for scheduling [16]. Similar to IMS, CRIMSON maintains an MRT to check for resource overuse during RC_ASAP and RC_ALAP assignment. The RC_ASAP is calculated from the nodes that does not have any incoming edges in the current iteration. These nodes are allotted RC_ASAP time as 0, which means, that the earliest start time of these nodes is at time 0. Based on the outgoing nodes from these start nodes and the delay of each operation, the RC_ASAP of consecutive nodes are

²The DFG and the architecture is the same as the motivation example Fig. 3(a) and (b).




Fig. 6. (a) DFG of the motivation example. (b) 2×2 CGRA architecture. (c) For each node of the DFG, RC_ASAP (column 2) and resource constrained ALAP (column 3) is first calculated. Then a random schedule time between RC_ASAP and RC_ALAP is chosen for each node. A valid randomized modulo schedule is shown in column 4. (d) With CRIMSON schedule a valid mapping is achieved by the mapping algorithm at II = 3.

computed in a depth-first manner (similar to IMS priority calculation). For the DFG in analysis, nodes a and f are assigned the RC_ASAP time as 0. Nodes b , c , g , and h are assigned RC_ASAP time as 1. Nodes d and i are assigned RC_ASAP time 2 and node e is assigned RC_ASAP time 3. The RC_ASAP times of each node is shown in Fig. 6(c) column 2. Next, starting from the last nodes of the DFG, i.e., nodes without any outgoing nodes in the current iteration, the nodes are assigned RC_ALAP in a reverse depth-first search manner, using $RC_{ALAP} = RC_{ASAP} + II - 1$. This ensures that $RC_{ALAP} \geq RC_{ASAP}$. For the given DFG, e is assigned RC_ALAP time 5, node h is assigned 3. Nodes d and i are assigned RC_ALAP time 4. Nodes b , c , and g are assigned RC_ALAP time 3. Finally a and f are assigned RC_ALAP time 2. The RC_ALAP times of each node is shown in Fig. 6(c) column 3.

After computing the RC_ASAP and RC_ALAP, CRIMSON chooses a random time between RC_ASAP and RC_ALAP, to schedule the nodes. Like IMS, CRIMSON maintains an MRT to check for resource overuse in each II modulo timeslot. After checking for resource constraints the modulo schedule time is chosen for each node. This randomization of modulo schedule time creates flexibility of movement for the nodes, which explores different modulo schedule spaces, thereby increasing the chances of finding a valid mapping by the mapping algorithm. A randomized modulo schedule for the example DFG is shown in Fig. 6(c) column 4, and a valid mapping for the scheduled nodes is shown in Fig. 6(d) at II = 3. The loop that was previously unmappable due to the restrictive scheduling of IMS Fig. 3, is now mappable at II = 3 due to randomization in assigning modulo schedule time.

If we take a closer look at the RC_ASAP and RC_ALAP times shown in Fig. 6(c) columns 2 and 3, we can observe that there is a chance that the RC_ASAP may be the modulo schedule chosen for all the nodes, since assigning a modulo schedule time for the nodes from RC_ASAP and RC_ALAP is randomized. As seen in Fig. 3(d) and (e), this schedule is not mappable. Unless there is a change to the workflow, there is a chance that finding a schedule that is unmappable and increasing the II to get a schedule process is repeated. To take care of this issue, CRIMSON proposes changes to the previous IMS-based workflow by statistically computing the feasibility

TABLE II
BENCHMARK CHARACTERISTICS

Suites	Loops	#nodes	#mem. nodes	#edges
MiBench	bitcount	22	4	28
	susan	31	8	35
	sha	31	10	39
	jpeg1	43	10	48
	jpeg2	28	6	33
Rodinia	kmeans1	15	6	17
	kmeans2	16	6	17
	kmeans3	17	4	20
	kmeans4	16	4	19
	kmeans5	12	2	13
	lud1	21	4	24
	lud2	20	4	24
	b+tree	13	2	13
	streamcluster	16	4	19
	nw	20	6	21
	BFS	28	10	32
	hotspot3D	76	20	96
	backprop	39	16	44
Parboil	spmv	25	8	27
	histo	18	4	20
	sad1	25	4	30
	sad2	19	4	20
	sad3	12	4	12
	stencil	69	16	94

of the scheduled nodes, prior to the mapping of the nodes. This makes sure that if a schedule is not mappable, a different random schedule is tried again for the same II. The number of times the mapping is tried for a given II is controlled by a threshold factor λ . With induced randomization in mapping and changes to the workflow, CRIMSON is able to achieve mapping of the application loops that were previously unmappable by IMS-based mapping techniques.

VI. EXPERIMENTAL RESULTS

A. Setup

Benchmarks We profiled top three of the widely used benchmark suites namely, MiBench [31], Rodinia [32], and Parboil [33]. The top performance-critical, nonvectorizable loops³ were chosen for the experiments. Loops that could not be compiled or the loops that were memory bound were not considered. Experiments were designed to consider only innermost loops so that a direct comparison with IMS can be made. These benchmarks depict a wide variety of applications from security, telecomm, etc., to parallel, high-performance computing (HPC) loops like sparse matrix-vector product (spmv). These loops on average across all the benchmark loops, corresponds to $\geq 50\%$ of the total application execution time.

Compilation: For selecting the loops from the application and converting the loops to the corresponding DFG, we used CCF [34]—CGRA compilation framework (LLVM 4.0 [35]-based). On top of the existing framework, to effectively compile the loops with control-dependencies (If-Then-Else structures), we implemented partial predication [36] as an LLVM pass, to convert the control-dependencies into data dependencies. Partial predication [36] can efficiently handle loops with nested if-else structures. The loop characteristics are shown in Table II including the number of nodes in the

³Maximum up to 5 loops per benchmark, with each contributing $>7\%$ of the execution time of the application when executed with standard inputs that are shipped with the benchmark suites.

DFG (only computing nodes are included and constants that can be passed in the immediate field of the ISA are excluded) and number of memory (load/store) nodes. CCF framework [34] produces DFG of the loop with separate address generation and actual load/store functionality. Furthermore, during the addition of routing resources after scheduling, we have implemented path-sharing technique proposed in GraphMinor [18]. Path-sharing can reduce the redundant routing nodes added. We implemented CRIMSON as a pass in the CCF framework including the λ value computation and the feasibility test. We also implemented the IMS-based state-of-the-art RAMP [20] and GraphMinor [18] as a pass in CCF. As observed in Table I, RAMP has demonstrated equal or better results when compared to GraphMinor. Hence, we compare CRIMSON against RAMP. We compiled the applications of the benchmark suites using optimization level 3 to avoid including loops that can be vectorized by compiler optimizations. We considered 2-D torus mesh CGRA of sizes 4×4 , 5×5 , 6×6 , 7×7 , and 8×8 .

B. CRIMSON Is Able to Schedule and Map Loops That Could Not Be Mapped Using RAMP

From Table III, we can infer that for loops, jpeg1, jpeg2, hotspot3D, backprop, and stencil, IMS-based state-of-the-art heuristic RAMP, was not able to find a valid mapping for a 4×4 CGRA (denoted by “X” in Table III). From the motivating example Fig. 3, IMS produces almost the same modulo schedule time for most of the nodes for any increase in II. CRIMSON, on the other hand, facilitates the exploration of different modulo scheduling times for nodes of the DFG, resulting in a valid mapping. It is observed that even at a lower CGRA size 4×4 , CRIMSON was able to map these particular loops. From Table III, when running on RAMP, loops that were not mappable on a 4×4 CGRA, were eventually mapped when allocated enough resources. For example, stencil which was unmappable by RAMP on a 4×4 CGRA was mapped on a 5×5 CGRA due to allocation of additional resources. Therefore, it can be said that the motivating example can also be mapped when allocated enough resources. From the motivating example, if Fig. 3(b) CGRA architecture was a 3×3 CGRA, then the IMS-based mapping algorithm would have used the extra resources provided to route the operation r . But this conclusion was not applicable to all the loops, meaning, loops, such as hotspot3D and jpeg2 were unable to find a valid mapping even when additional resources were allocated. RAMP was not able to achieve a mapping even at 8×8 CGRA for hotspot3D whereas RAMP was not able to achieve a mapping till 6×6 for jpeg2. While RAMP is able to map most of the loops at a higher CGRA size, CRIMSON with effective randomized modulo scheduling was able to map all the loops at size 4×4 . Additionally, for sad1 and sad3 loops, for which GraphMinor was not able to find a mapping, CRIMSON was able to achieve a mapping at 4×4 CGRA size.

C. CRIMSON Has Nearly Identical II for Loops That Could Be Mapped Using RAMP

From Table III, we can observe that for loops mapped using RAMP, the II obtained from CRIMSON was comparable to

TABLE III

COMPREHENSIVE TABLE SHOWING THE MII AND II ACHIEVED BY RAMP, AN EVALUATED IMS-BASED HEURISTIC, AND CRIMSON (CRIM.) FOR 24 BENCHMARK APPLICATION LOOPS FROM THREE MAJOR BENCHMARK SUITES AT 0.005 *exploration_factor*. THE “X” IN THE TABLE DENOTES THAT THERE WAS NO MAPPING OBTAINED FROM RAMP FOR AN INCREASING II UP TO 50. MII IN THE TABLE DENOTES THE MINIMUM II, WHICH IS THE MAXIMUM OF EITHER RESMII OR RECMII

Suites	Loops	4x4			5x5			6x6			7x7			8x8		
		MII	RAMP	CRIM.												
MiBench	bitcount	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3
	susan	2	3	4	2	2	2	2	2	2	2	2	2	2	2	2
	sha	3	3	4	2	X	3	2	3	2	2	2	3	2	2	4
	jpeg1	3	X	6	2	X	4	2	2	2	2	2	2	2	2	2
	jpeg2	2	X	5	2	X	3	2	X	2	2	2	2	2	2	2
Rodinia	kmeans1	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
	kmeans2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
	kmeans3	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
	kmeans4	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
	kmeans5	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
	lud1	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
	lud2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
	b+tree	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
	streamcluster	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
	nw	3	3	3	2	2	2	2	2	2	2	2	2	2	2	2
Parboil	BFS	2	2	3	2	3	3	2	2	3	2	2	2	2	2	3
	hotspot3D	5	X	10	4	X	7	4	X	7	3	X	6	3	X	4
	backprop	5	X	7	4	4	4	3	3	3	3	3	3	3	3	4
	spmv	3	3	3	2	2	2	2	2	2	2	2	2	2	2	2
	histo	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
	sad1	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
Parboil	sad2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
	sad3	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
	stencil	4	X	6	3	4	5	3	3	3	3	3	4	2	2	2




Fig. 7. Scheduling time comparison of CRIMSON and IMS.

RAMP across five different CGRA sizes ranging from 4×4 to 8×8 . We can see an occasional spike in the II in CRIMSON for *susan* at 4×4 and *stencil* on 5×5 , which is due to premature II increase by CRIMSON based on the λ value. To emphasize, λ is the maximum number of randomized schedules that are explored at the same II. A new schedule may be requested: 1) on a failure to find a randomized schedule; 2) on a failure of the feasibility test; or 3) a failure to map. The λ value is not reset for a given II. After exhausting the λ limit, the II is increased and a new RC_ASAP and RC_ALAP is computed along with a new λ value. The λ value is computed by 1 for each II. The λ value is determined by the user defined *exploration_factor*, which is the percentage of schedule space to that should be explored. If the *exploration_factor* is set too low, less modulo schedules are explored per II, thereby making it difficult to obtain a valid mapping and increasing the II prematurely. If the *exploration_factor* is set too high the time

to obtain a valid schedule/mapping increases, which negatively affects the compilation time of CRIMSON. Table III comprehensively conveys that CRIMSON has a nearly identical performance compared to RAMP for all the loops across different CGRA architectures that RAMP was able to map and CRIMSON is better than RAMP by mapping the five loops that were not mappable by RAMP and seven loops that were not mappable by GraphMinor on a 4×4 CGRA. The II obtained from CRIMSON is not always equal to or better than state-of-the-art RAMP and is dependent on the λ value.

D. Scheduling Time Comparison Between CRIMSON and IMS

The scheduling time for IMS [22] and CRIMSON are shown in Fig. 7, which is reported based on the execution of both the algorithms on Intel-i7 running at 2.8 GHz with 16 GB memory. As shown in Fig. 7, the x-axis is the scheduling time,




Fig. 8. Scheduling time versus II tradeoff trend for stencil.




Fig. 9. Scheduling time versus II tradeoff trend for hotspot3D.

i.e., time to obtain a valid schedule that is mappable, in μ s (microseconds) and the y-axis corresponds to the benchmark loops. The 19 benchmarks shown in Fig. 7 are those in which a mappable schedule was obtained by IMS. From Fig. 7, we can see that the scheduling time of CRIMSON is slightly higher than that of IMS. This is due to the additional computation of RC_ASAP and RC_ALAP, and the feasibility test (Fig. 5). For the loops shown, the *exploration_factor* was kept at 0.005.

E. Tradeoff Analysis Between Scheduling Time and II at Different λ Values

From 1, we can see that the λ value depends on the *exploration_factor*. This factor is defined as the percentage of modulo schedule space to be explored when there is an infeasible schedule or a mapping failure. The *exploration_factor* was changed from 0.5% (0.005) to 10% (0.1) and the corresponding scheduling time and II were recorded. The scheduling time numbers are recorded from executing CRIMSON on Intel-i7 running at 2.8 GHz and 16 GB memory and the compilation was performed for a 4×4 CGRA. A 4×4 CGRA was chosen because the II obtained by CRIMSON was much greater than the MII and the effect of λ can be shown clearly. In Figs. 8 and 9, the left y-axis (primary axis) denotes the CRIMSON scheduling time, in seconds, and the right y-axis (secondary axis) denotes the II obtained. The x-axis denotes the different *exploration_factors*. From (1) it is to be noted that as the *exploration_factor* increases, the λ value increases. From Figs. 8 and 9, it is evident that as *exploration_factor* increases the CRIMSON scheduling time increases, due to elaborate exploration of the schedule space at a given II. For lower value of the *exploration_factor*, superficial exploration of modulo schedule space prematurely increases the II but at lower scheduling time. We can also note from Fig. 9 at 0.1 that the above statement is not always true. At 0.1 the II decreases with the decrease in the scheduling time because a feasible

and a mappable schedule was obtained earlier in the modulo schedule space exploration due to the innate randomness of the CRIMSON scheduling algorithm.

VII. CONCLUSION

This article presented some of the major challenges encountered in the state-of-the-art mapping techniques with respect to scheduling and mapping of compute-intensive loops onto the CGRA. The previous mapping techniques use IMS scheduling that rarely showed a change in the modulo schedules for increased II, which obstructed the mapping algorithm to map the application loops onto the CGRA architecture. Additionally, previous mapping techniques assumed that the obtained IMS schedule is mappable and started to map the scheduled nodes. On a failure to map, due to the limited connectivity of the PEs or addition of routing nodes, the mapping algorithms increase the II and call IMS again to get a schedule that almost never changes. To mitigate these challenges, this article introduced CRIMSON, that comprehensively modeled RC_ASAP and RC_ALAP, picking a random modulo schedule time between these upper and lower boundaries. CRIMSON generated different schedules, thereby exploring different schedule spaces, on each invocation for a given or increased II. CRIMSON also introduced a novel feasibility test that pruned schedules that are unmappable. On evaluating the top 24 performance-critical loops from MiBench, Rodinia, and Parboil, CRIMSON was able to map 5 application loops that were unmappable by RAMP and 7 application loops that were unmappable by GraphMinor. The II achieved by CRIMSON was comparable to the II achieved by RAMP for the application loops that were mappable by RAMP.

REFERENCES

- [1] L. Zheng *et al.*, “Technologies, applications, and governance in the Internet of Things,” *Internet of Things-Global Technological and Societal Trends* (From Smart Environments and Spaces to Green ICT). Aalborg, Denmark : River Publ., 2011.
- [2] A. Crotty, A. Galakatos, K. Dursun, T. Kraska, U. Çetintemel, and S. B. Zdonik, “Tupleware: ‘Big’ data, big analytics, small clusters,” in *Proc. Conf. Innovat. Data Syst. Res. (CIDR)*, 2015.
- [3] F. Bouwens, M. Berekovic, B. De Sutter, and G. Gaydadjiev, “Architecture enhancements for the ADRES coarse-grained reconfigurable array,” in *Proc. Int. Conf. High Perform. Embedded Archit. Compilers*, 2008, pp. 66–81.
- [4] C. Nicol, “A coarse grain reconfigurable array (CGRA) for statically scheduled data flow computing,” *Wave Comput.*, Santa Clara, CA, USA, White Paper, 2017.
- [5] Y.-H. Chen, T. Krishna, J. S. Emer, and V. Sze, “Eyeriss: An energy-efficient reconfigurable accelerator for deep convolutional neural networks,” *IEEE J. Solid-State Circuits*, vol. 52, no. 1, pp. 127–138, Jan. 2017.
- [6] Y.-H. Chen, J. Emer, and V. Sze, “Eyeriss: A spatial architecture for energy-efficient dataflow for convolutional neural networks,” *ACM SIGARCH Comput. Archit. News*, vol. 44, no. 3, pp. 367–379, 2016.
- [7] X. Fan, H. Li, W. Cao, and L. Wang, “DT-CGRA: Dual-track coarse-grained reconfigurable architecture for stream applications,” in *Proc. IEEE 26th Int. Conf. Field Program. Logic Appl. (FPL)*, Lausanne, Switzerland, 2016, pp. 1–9.
- [8] D. C. Chen, “Programmable arithmetic devices for high speed digital signal processing,” Ph.D. dissertation, Dept. Elect. Eng. Comput. Sci., Univ. California, Berkeley, Ca, USA, 1992.
- [9] R. W. Hartenstein and R. Kress, “A datapath synthesis system for the reconfigurable datapath architecture,” in *Proc. ASP-DAC’95/CHDL’95/VLSI’95 With EDA Technofair*, Chiba, Japan, 1995, pp. 479–484.

- [10] E. Mirsky and A. DeHon, "MATRIX: A reconfigurable computing architecture with configurable instruction distribution and deployable resources," in *Proc. IEEE Symp. FPGAs Custom Comput. Mach. (FCCM)*, vol. 96. Napa Valley, CA, USA, 1996, pp. 17–19.
- [11] H. Singh, M.-H. Lee, G. Lu, F. J. Kurdahi, N. Bagherzadeh, and E. M. C. Filho, "MorphoSys: An integrated reconfigurable system for data-parallel and computation-intensive applications," *IEEE Trans. Comput.*, vol. 49, no. 5, pp. 465–481, May 2000.
- [12] T. Miyamori and K. Olukotun, "REMARC: Reconfigurable multimedia array coprocessor," *IEICE Trans. Inf. Syst.*, vol. 82, no. 2, pp. 389–397, 1999.
- [13] M. Hamzeh, A. Shrivastava, and S. Vrudhula, "EPIMap: Using epimorphism to map applications on CGRAs," in *Proc. 49th Annu. Design Autom. Conf.*, San Francisco, CA, USA, 2012, pp. 1284–1291.
- [14] B. Mei, M. Berekovic, and J. Y. Mignolet, "ADRES & DRESC: Architecture and compiler for coarse-grain reconfigurable processors," in *Fine-and Coarse-Grain Reconfigurable Computing*. Dordrecht, The Netherlands: Springer, 2007, pp. 255–297.
- [15] H. Park, K. Fan, M. Kudlur, and S. Mahlke, "Modulo graph embedding: Mapping applications onto coarse-grained reconfigurable architectures," in *Proc. Int. Conf. Compilers Archit. Synth. Embedded Syst.*, 2006, pp. 136–146.
- [16] H. Park, K. Fan, S. A. Mahlke, T. Oh, H. Kim, and H.-S. Kim, "Edge-centric modulo scheduling for coarse-grained reconfigurable architectures," in *Proc. 17th Int. Conf. Parallel Archit. Compilation Techn.*, Toronto, ON, Canada, 2008, pp. 166–176.
- [17] T. Oh, B. Egger, H. Park, and S. Mahlke, "Recurrence cycle aware modulo scheduling for coarse-grained reconfigurable architectures," *ACM SIGPLAN Notices*, vol. 44, pp. 21–30, Jun. 2009.
- [18] L. Chen and T. Mitra, "Graph minor approach for application mapping on CGRAs," *ACM Trans. Reconfig. Technol. Syst.*, vol. 7, no. 3, p. 21, 2014.
- [19] M. Hamzeh, A. Shrivastava, and S. Vrudhula, "REGIMap: Register-aware application mapping on coarse-grained reconfigurable architectures (CGRAs)," in *Proc. 50th Annu. Design Autom. Conf.*, Austin, TX, USA, 2013, p. 18.
- [20] S. Dave, M. Balasubramanian, and A. Shrivastava, "RAMP: Resource-aware mapping for CGRAs," in *Proc. 55th Annu. Design Autom. Conf. (DAC)*, San Francisco, CA, USA, 2018, pp. 1–6.
- [21] S. Yin, X. Yao, D. Liu, L. Liu, and S. Wei, "Memory-aware loop mapping on coarse-grained reconfigurable architectures," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 24, no. 5, pp. 1895–1908, May 2016.
- [22] B. R. Rau, "Iterative modulo scheduling," *Int. J. Parallel Program.*, vol. 24, no. 1, pp. 3–64, 1996.
- [23] B. Cheng, H. Wang, S. Yang, D. Niu, and Y. Jin, "A novel testability-oriented data path scheduling scheme in high-level synthesis," *Tsinghua Sci. Technol.*, vol. 12, pp. 134–138, Jul. 2007.
- [24] L. Zhou, D. Liu, M. Tang, and H. Liu, "Mapping loops onto coarse-grained reconfigurable array using genetic algorithm," in *Proc. 8th Int. Conf. Bio-Inspired Comput. Theor. Appl. (BIC-TA)*, 2013, pp. 801–808.
- [25] D. Liu, S. Yin, L. Liu, and S. Wei, "Polyhedral model based mapping optimization of loop nests for CGRAs," in *Proc. 50th ACM/EDAC/IEEE Design Autom. Conf. (DAC)*, Austin, TX, USA, 2013, pp. 1–8.
- [26] G. Lee, K. Choi, and N. D. Dutt, "Mapping multi-domain applications onto coarse-grained reconfigurable architectures," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 30, no. 5, pp. 637–650, May 2011.
- [27] A. Hatanaka and N. Bagherzadeh, "A modulo scheduling algorithm for a coarse-grain reconfigurable array template," in *Proc. IEEE Int. Parallel Distrib. Process. Symp. (IPDPS)*, Rome, Italy, 2007, pp. 1–8.
- [28] G. Ansaloni, K. Tanimura, L. Pozzi, and N. Dutt, "Integrated kernel partitioning and scheduling for coarse-grained reconfigurable arrays," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 31, no. 12, pp. 1803–1816, Dec. 2012.
- [29] D. Wijerathne, Z. Li, M. Karunaratne, A. Pathania, and T. Mitra, "CASCADE: High throughput data streaming via decoupled access-execute CGRA," *ACM Trans. Embedded Comput. Syst.*, vol. 18, no. 5s, pp. 1–26, 2019.
- [30] M. Karunaratne, A. K. Mohite, T. Mitra, and L.-S. Peh, "HyCUBE: A CGRA with reconfigurable single-cycle multi-hop interconnect," in *Proc. 54th Annu. Design Autom. Conf.*, Austin, TX, USA, 2017, pp. 1–6.
- [31] M. Guthaus, L. S. Ringenberg, D. Ernst, T. M. Austin, T. Mudge, and B. R. Brown, "MiBench: A free, commercially representative embedded benchmark suite," in *Proc. Workload Characterization Int. World Wide Web Workshop (WWC)*, 2001, pp. 3–14.
- [32] S. Che *et al.*, "Rodinia: A benchmark suite for heterogeneous computing," in *Proc. IEEE Int. Symp. Workload Characterization (IISWC)*, Austin, TX, USA, 2009, pp. 44–54.
- [33] J. A. Stratton *et al.*, "Parboil: A revised benchmark suite for scientific and commercial throughput computing," *Center Rel. High Perform. Comput.*, vol. 127, pp. 1–11, Mar. 2012.
- [34] S. Dave and A. Shrivastava, "CCF: A CGRA compilation framework," 2018.
- [35] C. Lattner and V. Adve, "LLVM: A compilation framework for lifelong program analysis & transformation," in *Proc. Int. Symp. Code Gener. Optim. (CGO)*, San Jose, CA, USA, 2004, pp. 75–86.
- [36] K. Han, J. Ahn, and K. Choi, "Power-efficient predication techniques for acceleration of control flow execution on CGRA," *ACM Trans. Archit. Code Optim.*, vol. 10, no. 2, p. 8, 2013.



Mahesh Balasubramanian (Member, IEEE) received the bachelor's degree from Anna University, Chennai, India, in 2011, and the master's degree in electrical engineering from the University of Texas at San Antonio, San Antonio, TX, USA, in 2013. He is currently pursuing the Ph.D. degree with Arizona State University, Tempe, AZ, USA.

His research interests lie in area of parallel computing and co-processor accelerators like CGRAs particularly for general purpose and HPC applications.

Mr. Balasubramanian is a recipient of the CIDSE Doctoral Fellowship, the Engineering Grad Fellowship, and the Ferdinand A. Stanchi Fellowship at Arizona State University.



Aviral Shrivastava (Member, IEEE) received the bachelor's degree in computer science and engineering from the Indian Institute of Technology, New Delhi, India, in 1999, and the master's and Ph.D. degrees in computer science and engineering from the University of California at Irvine, Irvine, CA, USA, in 2002 and 2006, respectively.

He is a Professor with the School of Computing Informatics and Decision Systems Engineering, Arizona State University, Tempe, AZ, USA, where he has established and heads the "Make Programming Simple" Lab. His research is funded by NSF, DOE, NIST, and several industries, including Microsoft, Raytheon Missile Systems, Intel, and Nvidia. His research lies in the broad area of "Software for Embedded and Cyber-Physical Systems." More specifically, his interested in topics around: 1) compilers and microarchitectures for heterogeneous and many-core computing; 2) protecting computation from soft errors; and 3) precise timing for cyber-physical systems.

Prof. Shrivastava was a recipient of the NSF CAREER Award in 2011 and the Outstanding Junior Researcher in CSE at ASU in 2012. His works have received several best paper nominations, including at DAC 2017, and a Best Student Paper Award at VLSI 2016. His students have received outstanding Ph.D. student award in CSE at ASU in 2017 and the Outstanding M.S. Student Award in CSE at ASU in 2012 and 2010. He is currently serving as the Deputy Editor-in-Chief of IEEE EMBEDDED SYSTEMS LETTERS, and an Associate Editor for ACM TRANSACTIONS EMBEDDED COMPUTING SYSTEMS, IEEE TRANSACTIONS ON COMPUTER-AIDED DESIGN, International Journal on Parallel Processing (Springer), and Design Automation for Embedded Systems (Springer). He has served as the Program Chair of CODES+ISSS 2017 and 2018, and LCTES 2019. He is currently the Virtual Conference Chair of ESWEEK 2020 and the Track Chair for RTSS 2020. He serves on the organizing and program committees of several premier embedded system conferences, including DAC, ICCAD, ISLPED, ESWEEK, and LCTES.