A

Graph Minor Approach for Application Mapping on CGRAs

LIANG CHEN, National University of Singapore
TULIKA MITRA, National University of Singapore

Coarse-Grained Reconfigurable Arrays (CGRAs) exhibit high performance, improved flexibility, low cost, and power efficiency for various application domains. Compute-intensive loop kernels, which are perfect candidates to be executed on CGRAs, are usually mapped through modified modulo scheduling algorithms. These algorithms should be capable of performing both placement and routing. We formalize the CGRA mapping problem as a graph minor containment problem. We essentially test if the data flow graph representing the loop kernel is a minor of the modulo routing resource graph representing the CGRA resources and their interconnects. We design an exact graph minor testing approach that exploits the unique properties of both the data flow graph and the routing resource graph to significantly prune the search space. We introduce additional heuristic strategies that drastically improve the compilation time while still generating optimal or near-optimal mapping solutions. Experimental evaluation confirms the efficiency of our approach.

Categories and Subject Descriptors: C.1.4.2 [**Computer systems organization**]: Reconfigurable computing

General Terms: Algorithms, Performance

Additional Key Words and Phrases: Coarse grained reconfigurable arrays (CGRAs), Graph minor, Compilation

ACM Reference Format:
Liang Chen and Tulika Mitra, 2014. Graph Minor Approach for Application Mapping on CGRAs. *ACM Trans. Reconfig. Technol. Syst.* V, N, Article A (January 2014), 22 pages.
DOI : http://dx.doi.org/10.1145/0000000.0000000

1. INTRODUCTION

Coarse-Grained Reconfigurable Arrays (CGRAs) are promising alternatives between ASICs and FPGAs. Traditionally in embedded systems, compute intensive kernels of an application are implemented as ASICs, which have high efficiency but limited flexibility. Current generation embedded systems demand flexibility to support diverse applications. FPGAs provide high flexibility, but may suffer from low efficiency [Kuon and Rose 2007]. To bridge this gap, CGRA architectures, such as CHESS [Marshall et al. 1999], MorphoSys [Singh et al. 2000], ADRES [Mei et al. 2003a], DRAA [Lee et al. 2003], FloRA [Lee et al. 2009] etc., have been proposed. Typically these architectures arrange coarse-grained functional units (FUs) in a mesh structure. The FUs can be reconfigured by writing to a control (context) register on per cycle basis. Figure 1 shows a 4×4 CGRA with FUs connected in a mesh topology; each FU has a local register file and a configuration cache.

The compute-intensive loop kernels are perfect candidates to be mapped to CGRAs containing multiple FUs targeting high instruction-level parallelism. A number of CGRA mapping algorithms [Mei et al. 2003b; Bansal et al. 2003; Hatanaka and Bagherzadeh 2007; Friedman et al. 2009; De Sutter et al. 2008; Gnanaolivu et al. 2010] are inspired by compilation techniques for VLIW architectures as well as FPGA synthesis. For example, CGRA mapping algorithms adopt placement and routing techniques from FPGA synthesis domain and software pipelining based techniques such

This work was partially supported by Singapore Ministry of Education Academic Research Fund Tier 2 MOE2009-T2-1-033, MOE2012-T2-1-115, and Tier 1 T1-251RES1120.
Author's addresses: Liang Chen and Tulika Mitra, Department of Computer Science, School of Computing, National University of Singapore. Email: {chenliang,tulika}@comp.nus.edu.sg
Permission to make digital or hard copies of part or all 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 show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.
© 2014 ACM 1936-7406/2014/01-ARTA $15.00
DOI : http://dx.doi.org/10.1145/0000000.0000000

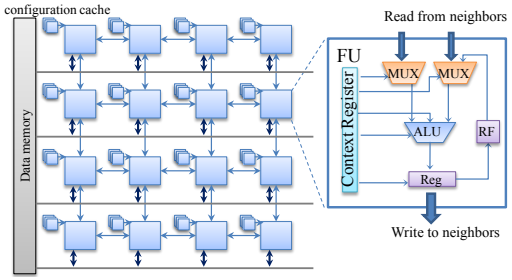

Fig. 1. A 4×4 CGRA

as modulo scheduling from VLIW compilation process. It is important to note that the inherent structure of the CGRAs is very different from both FPGAs and VLIW architectures. More concretely, the connectivity among the functional units in CGRAs is usually fixed unlike FPGAs where the interconnections can be reconfigured. Thus the mapping algorithms based on FPGA place and route techniques may find it challenging to identify feasible routing paths in fixed interconnect structure of CGRAs. Similarly, unlike VLIW architectures where all the FUs typically share a common register file, the FUs in most CGRAs have limited and explicit connections to the register files. Thus it is not prudent to perform register allocation as a post-processing step as is commonly done in VLIW scheduling. Instead, register allocation should be integrated in the early stage with scheduling (place and route) to achieve quality mapping.

In this work, we focus on developing an efficient CGRA mapping algorithm that generates high quality solution with fast compilation time. To first formalize the CGRA mapping problem, we notice that a number of recent works [Tuhin and Norvell 2008; Alle et al. 2008; Brenner et al. 2009; Gnanaolivu et al. 2010; 2011] follow subgraph homeomorphism [Fortune et al. 1980] formalization. The idea is to test if the data flow graph (DFG) representing the loop kernel is subgraph homeomorphic to the modulo routing resource graph (MRRG) representing the CGRA resources and their interconnects. Homeomorphism formulation allows subdivision of the DFG edges when being mapped onto the MRRG, i.e., a DFG edge can be mapped as a chain of edges (path) on the MRRG. Alternatively, additional vertices on a path consisting of a chain of edges on the MRRG can be smoothed out to create a single DFG edge. The additional nodes by sub-divisions model the routing of data from the source to the target FUs if they are not connected directly. However, subgraph homeomorphism requires the edge mappings to be node-disjoint (except at end points) or edge-disjoint [Fortune et al. 1980]. While subgraph homeomorphism provides an elegant formulation of the CGRA mapping problem, it excludes the possibility of sharing the routing nodes among single source multiple target edges [Park et al. 2008] (also called multi-net [De Sutter et al. 2008]) leading to possible wastage of precious routing resources.

Figure 2 illustrates the subgraph homeomorphism formulation. Figure 2(a) shows a simple DFG (for simplicity we have removed the loop back edge) being mapped onto a 2x2 homogeneous mesh CGRA shown in Figure 2(b). The DFG is homeomorphic to the subgraph of the MRRG shown in Figure 2(c) and thus the subgraph represents a valid mapping (again for simplicity we have removed additional nodes of the MRRG). In this homeomorphic mapping, edges (1,3) and (1,4) have been routed through three additional routing nodes marked by R. Notice that each routing node has degree 2 and has been added through edge subdivision (marked by dashed edges). Alternatively, the routing nodes in the MRRG subgraph can be smoothed out to obtain the original DFG. As mentioned earlier, by definition, edge subdivision cannot support route sharing.

In contrast, we model the CGRA mapping problem as graph minor containment problem which can explicitly model route sharing. A graph H is a minor of graph G if H can be obtained from a subgraph of G by a (possibly empty) sequence of edge contractions [Robertson and Seymour 1999]. In graph theory, an edge contraction removes an edge from a graph while simultaneously merging the two vertices it previously connected. In our context, we need to test if the DFG is a minor of the MRRG, where the edges to be contracted represent the routing paths in the MRRG. Unlike

edge subdivision (or its reverse operation smoothing), edge contractions are not restricted to simple paths. Thus graph minor formalism naturally allows for route sharing. Figure 2(d) shows a mapping under graph minor approach. It is a subgraph of the MRRG from which the DFG can be derived through two edge contractions as shown in Figure 2(e)-(f). In this example, we reduce the number of routing nodes from 3 (in subgraph homeomorphism mapping) to 2 (in graph minor mapping). While it is possible to support route sharing in [Park et al. 2008; De Sutter et al. 2008], we provide a formalization of the CGRA mapping problem under route sharing. This formalization enables us to design a customized exact graph minor testing approach that fully exploits the structure of the DFG and the CGRA interconnects to effectively navigate and prune the mapping alternatives.

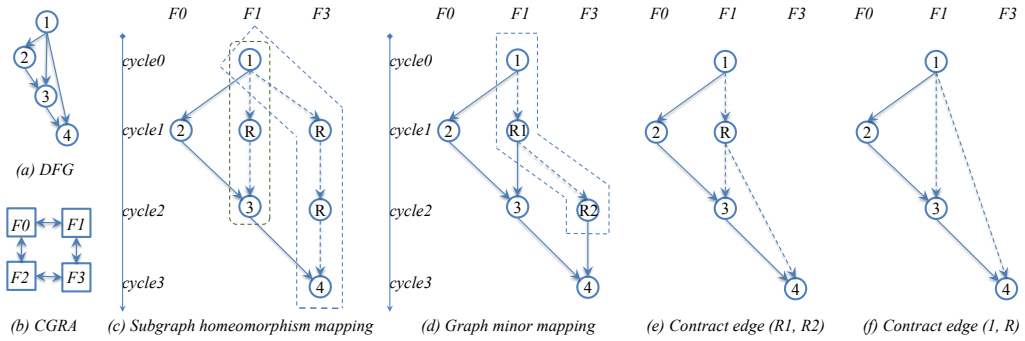

Fig. 2. Subgraph Homeomorphism versus Graph Minor formulation of CGRA mapping problem.

In parallel to our graph minor formalization [Chen and Mitra 2012] for CGRA mapping problem, [Hamzeh et al. 2012] proposed graph epimorphism formalization for the same problem. Their approach, called EPIMap, is quite elegant and models the novel concept of re-computation in addition to route sharing. Re-computation allows for the same operation to be performed on multiple FUs if it leads to better routing. In EPIMap approach, the DFG H is morphed into another graph H' (through introduction of routing/re-computation nodes and other transformations) such that there exists subgraph epimorphism from H' to H (many to one mapping of vertices from H' to H and adjacent vertices in H' map to adjacent vertices in H). Then EPImap attempts to find the maximal common subgraph (MCS) between H' and the MRRG graph G using standard MCS identification procedure. If the resulting MCS is isomorphic to H', then a valid mapping has been obtained; otherwise H is morphed differently in the next iteration and the process repeats.

The key difference with our approach is that while we develop a customized graph minor testing procedure that exploits structural properties of our graphs, EPIMap relies on off-the-shelf MCS identification algorithm. This can potentially lead to faster compilation time for graph minor approach. Both approaches introduce heuristics to manage the computational complexity; the transformation of the DFG as well as MCS identification require heuristics in EPIMap, while graph minor approach restricts the subgraph mapping choices. Thus, the quality of the solutions in both approaches depend on the loop kernel and the underlying CGRA architecture. On the other hand, the re-computation concept in EPIMap enables additional scheduling and routing options that can potentially generate better quality solutions for certain kernels. Finally, graph epimorphism and graph minor are quite unrelated concepts even though a detailed discussion on this topic is out of scope here. Instead, we provide quantitative comparison of the two approaches in Section 6.

The concrete contributions of this paper are as follows. We observe that the CGRA mapping problem in the presence of route sharing can be formulated as a graph minor containment problem. This allows us to develop a systematic and customized mapping algorithm that directly works on the input DFG and MRRG to explore the inherent structural properties of the two graphs during the mapping process. Experimental results confirm that our graph minor approach can achieve high quality schedules with minimal compilation time.

In the following, we will first discuss the existing approaches for CGRA mapping problem in Section 2. Backgrounds on modulo scheduling in CGRA mapping problem are provided in Section 3. We then formalize the CGRA mapping problem as a graph minor containment problem in Section 4. The proposed graph minor testing algorithm will be detailed in Section 5. Experimental evaluations comparing graph minor approach with different techniques are presented in Section 6.

2. RELATED WORK

Mapping a compute-intensive loop kernel of an application to CGRAs using modulo scheduling was first discussed in [Mei et al. 2003b]. In this simulated annealing based approach, the cost function is defined according to the number of over-occupied resources. The simulated annealing approach can have long convergence time, especially for large dataflow graphs. Routing through register files and register allocation problems are further explored in [De Sutter et al. 2008], which extends the work in [Mei et al. 2003b]. Register allocation is achieved by constraining the register usage during the simulated annealing place and route process. The imposed constraint is adopted from meeting graph [Eisenbeis et al. 1995] for solving loop cyclic register allocation in VLIW processors. In post routing phase, the registers are allocated by finding a Hamilton circuit in the meeting graph, which is solved as a traveling salesman problem [De Sutter et al. 2008]. This technique is specially designed for CGRAs with rotating register files. [Hatanaka and Bagherzadeh 2007] also follows the simulated annealing framework but aims at finding better cost functions for over-used resources. SPR [Friedman et al. 2009] is a mature CGRA mapping tool that successfully combines the VLIW style scheduler and FPGA placement and routing algorithms for CGRA application mapping. It consists of three individual steps namely scheduling, placement, and routing. The placement step of SPR also uses the simulated annealing approach.

List scheduling has been adopted in [Bansal et al. 2003], which analyzes priority assignment heuristics under different network traversal strategies and delay models. The heuristics utilize the interconnect information to ensure that data dependent operations can be mapped spatially close to each other. [Park et al. 2006] also gives priorities for operations and resources to obtain a quality schedule. The priorities are assigned according to the importance of routing from producer nodes to consumer nodes. This idea is further exploited in edge-centric modulo scheduling (EMS) [Park et al. 2008], where the primary objective is routing efficiency rather than operation assignments. The quality of a mapping using specific priorities highly depends on efficient heuristics for assigning these priority values to both operations and resources.

There are various approaches to CGRA mapping using techniques from graph theory domain. [Clark et al. 2006] integrates subgraph isomorphism algorithm to generate candidate mapping between a DFG and the resource graph of a coarse-grained accelerator. SPKM [Yoon et al. 2009] adopts the split and push technique [Di Battista et al. 1998] for planar graph drawing and focuses on spatial mappings for CGRAs. The mapping in SPKM starts from an initial drawing where all DFG nodes reside in the same group. One group represents a single functional unit. The group is then split into two and a set of nodes are pushed to the newly generated group. The split process continues till each group contains only one node, which represents a one-to-one mapping from DFG to the planar resource graph of CGRA.

A number of CGRA mapping approaches follow the subgraph homeomorphism formalizations including [Tuhin and Norvell 2008; Alle et al. 2008; Brenner et al. 2009; Gnanaolivu et al. 2010; 2011]. The mapping algorithm in [Tuhin and Norvell 2008] is adapted from MIRS [Zalamea et al. 2003], a modulo scheduler capable of instruction scheduling with register constraints. The adaptations for CGRA mapping include a cost function for routing and considerations for conditional branches. [Alle et al. 2008] partitions the DFG into substructures called HyperOps and these HyperOps are synthesized into hardware configurations. The synthesis is carried out through a homeomorphic transformation of the dependency graph of each HyperOp onto the resource graph. [Brenner et al. 2009] also formalizes the CGRA mapping as a subgraph homeomorphism problem. However, they consider general application kernels rather than loops. Particle swarm optimization is adopted for solving CGRA mapping problem in [Gnanaolivu et al. 2010; 2011]. The calculation

for fitness, which is used to move particles (DFG nodes) in particle swarm optimization, is specifically designed to optimize multiple objectives for routing. Compared to subgraph homeomorphism formalization, our graph minor approach takes advantage of sharing different data routes from one producer operation to multiple consumer operations.

EPIMap [Hamzeh et al. 2012] formalizes the CGRA mapping problem as an graph epimorphism problem with the additional feature of re-computations. The core of this approach consists of a subgraph isomorphism solver which finds the maximum common subgraph (MCS) [Levi 1973] between the DFG and the resource graph of CGRA. The idea is to transform the DFG iteratively by inserting dummy routing nodes or replicated operation nodes so that the routing requirements can be satisfied through the subgraph isomorphism solver. EPIMap can generate better scheduling results compared to EMS with similar compilation time. Most graph approaches solve a subset of the epimorphism problem defined in EPIMap. In our graph minor approach, rather than transforming the DFG, we directly explore the structural properties between the DFG and resource graph of CGRA during the mapping process.

3. MODULO SCHEDULING FOR CGRA

Given a loop from an application and a CGRA architecture, the goal of mapping is to generate a schedule such that the application throughput is maximized. The loop is represented as a data flow graph (DFG) where the nodes represent the operations and the edges represent the dependency among the operations. Figure 4(a) shows the DFG of a simple loop. Figure 4(b) shows a 2x2 CGRA consisting of four functional units (FUs) where the loop should be mapped to. The mapping problem consists of (a) scheduling the operations in space and time so as to satisfy the dependency constraints, and (b) explicit routing of the operands from the producers to the consumers.

3.1. CGRA Architecture

For simplicity of exposition, in the algorithm description we assume a homogeneous CGRA architecture with comprehensive FUs that can support all possible operations. However, our mapping approach can support diverse CGRA architectures through parameterizations. Our register file modeling approach can also support many different register file configurations such as NORF (architecture with no RF shown in Figure 3(a)), LRF (architecture with local shared RF shown in Figure 3(b)) and CRF (the architecture with central shared RF shown in Figure 3(c)). Heterogeneities for functional units are also supported in our framework. Experimental evaluations for different CGRA architectures will be presented in Section 6.

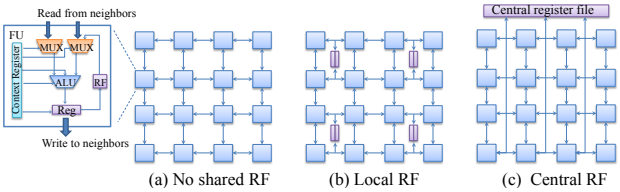

Fig. 3. 4×4 CGRAs with different register file configurations

3.2. Modulo Scheduling

Modulo scheduling is a software pipelining technique used to exploit instruction-level-parallelism in the loops by overlapping consecutive iterations [Rau 1994]. The schedule produced includes three phases: the prologue, the kernel, and the epilogue. The kernel corresponds to the steady state execution of the loop and comprises of operations from consecutive iterations. The schedule length of the kernel, which is also the interval between successive iterations, is called the initiation interval (II). If the number of loop iterations is high, then the execution time in the kernel is dominant compared to the prologue and the epilogue. Thus, the goal for modulo scheduling is to minimize the II value. Initially, the scheduler selects the minimal II (MII) value between resource-minimal II and

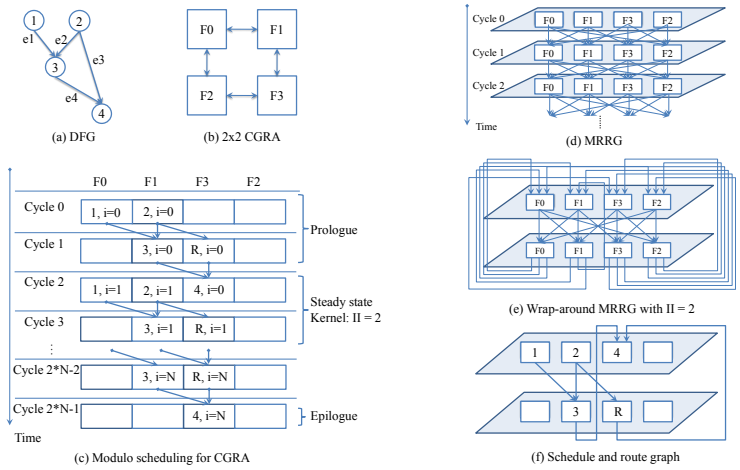

Fig. 4. Modeling of loop kernel mapping on CGRAs: An illustrative example.

recurrence-minimal II, and attempts to find a feasible schedule with that II value. If the scheduling fails, then the process is repeated with an increased II value.

Figure 4(c) shows the modulo-scheduled version of the loop in Figure 4(a) to the CGRA architecture in Figure 4(b) with prologue, kernel, and epilogue where II=2. Notice that *operation 4* from the ith iteration is executing in the same cycle with *operation 1* and *operation 2* from the $(i+1)^{th}$ iteration in the steady state. Also, we need to hold the output of *operation 2* in a routing node (R) till it gets consumed by *operation 4*. This explicit routing between FUs is what sets apart modulo scheduling in CGRAs from conventional modulo scheduling, where FUs are fully connected through the central register file (RF) and routing is guaranteed. In CGRAs, the modulo scheduler has to be aware of the underlying interconnect among the FUs and the RFs to route data.

3.3. Modulo Routing Resource Graph (MRRG)

Mei et al. [Mei et al. 2003b] defined a resource management graph for CGRA mapping, called Modulo Routing resource graph (MRRG), which has been used extensively in subsequent studies. In MRRG, the resources are presented in a time-space view. The nodes represent the ports of the FUs and the RFs, and the edges represent the connectivity among the ports. We adopt a simplified form of MRRG proposed in [Park et al. 2006] where a node corresponds to FU or RF rather than the ports. Our mapping technique integrates register allocation with scheduling. We model each RF as one node per cycle in the MRRG. The individual registers within RF are treated as identical elements and represented by the capacity of the RF as in *compact register file model* [De Sutter et al. 2008]. The usage of registers is tracked and constrained during the mapping procedure. The number of read and write ports per RF is also included as a constraint.

The MRRG is a directed graph G_{II} where *II* corresponds to the initiation interval. Given a graph G, we denote the vertex set and the edge set of G by $V(G)$ and $E(G)$, respectively. Each node $v \in V(G_{II})$ is a tuple *(n, t)*, where *n* refers to the resource (FU or RF) and *t* is the cycle. Let $e = (u, v) \in E(G_{II})$ be an edge where *u = (m, t)* and *v = (n, t+1)*. Then the edge *e* represents a connection from resource *m* in cycle *t* to resource *n* in cycle *t+1*. Generally, if resource *m* is connected to resource *n* in the CGRA, then node *u = (m, t)* is connected to node *v = (n, t+1)*, $t \geq 0$.

For example, Figure 4(d) shows the MRRG corresponding to the CGRA shown in Figure 4(b). The resources of the CGRA are replicated every cycle along the time axis, and the edges point forward in time. During modulo scheduling, when a node *v=(n, t)* in the MRRG becomes occupied, then all the nodes *v'=(n, t+k×II)* (where $k > 0$) are also marked occupied. For example, in the modulo schedule with II=2 shown in Figure 4(c), as F1 is occupied by *operation 2* in cycle 0, it is

also occupied by *operation 2* every $2 \times k$ cycle. In most CGRA mapping techniques, this modulo reservation for occupied resources is done through a modulo reservation table [Mei et al. 2003b].

3.4. MRRG with Wrap-around Edges

The goal of CGRA modulo scheduler is to generate II different configurations for the CGRA where each configuration corresponds to a particular cycle in the kernel. These configurations are stored in the configuration caches and provide configuration contexts to FUs every cycle. As these configurations are repeated every II cycles, the output from the resources involved in the last configuration cycle are consumed by the resources involved in the first configuration cycle. Thus instead of using MRRG where the time axis grows indefinitely till the steady state is achieved, we could restrict the time axis to the target II. We then need to add wrap around edges from the last cycle to the first cycle as shown in Figure 4(e) (similar graph is also used in [Friedman et al. 2009]). The modulo scheduled kernel in Figure 4(c) can now be simplified to the graph in Figure 4(f). We refer to this simplified graph as *schedule and route graph (SRG)*, which captures the scheduling plus routing information and is a subgraph of the MRRG. So instead of using a modulo reservation table, we can directly use MRRG with wrap around edges, which provides us an integrated view during mapping. *In the following, the term MRRG will be used to refer to MRRG with wrap around edges.*

4. CGRA MAPPING PROBLEM FORMALIZATION

We first present the formalization of the CGRA mapping problem in the form of subgraph isomorphism when no data routing is required and subgraph homeomorphism when routes are not shared. We then model the CGRA mapping as a graph minor problem [Robertson and Seymour 1999] between the DFG and the MRRG in the presence of route sharing. Meanwhile, we point out the necessary restrictions imposed in the formalization. We also provide the NP-completeness proof for the CGRA mapping problem under our graph minor formalization.

4.1. Subgraph Isomorphism and Subgraph Homeomorphism Mapping

Let H be a directed graph representing the DFG and G_{II} be a directed graph representing the MRRG with initiation interval II. We are looking for a mapping from the input graph *H* to the target graph *G*. In the ideal scenario of full connectivity among the FUs, all the data dependencies in the DFG can be mapped to direct edges in the MRRG. That is, for any edge $e = (u, v) \in E(H)$, there is an edge $e = (f(u), f(v)) \in E(G)$ where f represents the vertex mapping function from the DFG to the MRRG. This matches the definition of subgraph isomorphism in graph theory. Thus the CGRA application mapping problem can be solved using techniques for subgraph isomorphism from the graph theory domain [Ullmann 1976; Cordella et al. 2004].

In reality, data may need to be routed through a series of nodes rather than direct links. For example, the edges $(1, 3)$ and $(1, 4)$ in Figure 2(a) are routed through additional nodes. If an edge $e = (u, v) \in E(H)$ in the DFG can be mapped to a path from *f(u)* to *f(v)* in the MRRG *G*, it matches the subgraph homeomorphism definition [Fortune et al. 1980]. The subgraph homeomorphism techniques for CGRA mapping problem has been adopted in [Tuhin and Norvell 2008; Alle et al. 2008; Brenner et al. 2009; Gnanaolivu et al. 2010; 2011]. Subgraph homeomorphism, however, requires the edge mappings to be node-disjoint (or edge-disjoint), which means the nodes (or the edges) in the mapping paths for the edges carrying the same data cannot be shared.

4.2. Graph Minor

We now present graph minor [Robertson and Seymour 1999] based formulation of the application mapping problem on CGRAs with route sharing. In graph theory, an undirected graph *H* is called a minor of the graph *G* if *H* is isomorphic to a graph that can be obtained by zero or more edge contractions on a subgraph of *G*. An edge contraction is an operation that removes an edge from a graph while simultaneously merging together the two vertices it used to connect. More formally, a graph *H* is a minor of another graph *G* if a graph isomorphic to *H* can be obtained from *G* by

contracting some edges, deleting some edges, and deleting some isolated vertices. The order in which a sequence such operations are performed on G does not affect the resulting graph H.

A *model* of H in G is a mapping ϕ that assigns to every edge $e \in E(H)$ an edge $\phi(e) \in E(G)$, and to every vertex $v \in V(H)$ a non-empty connected tree subgraph $\phi(v) \subseteq G$ such that

(1) the graphs $\{\phi(v)|v \in V(H)\}$ are mutually vertex-disjoint and the edges $\{\phi(e)|e \in E(H)\}$ are pairwise distinct; and
(2) for $e = \{u, v\} \in E(H)$, the edge $\phi(e)$ connects subgraph $\phi(u)$ with subgraph $\phi(v)$.

H is isomorphic to a minor of G if and only if there exists a model of H in G [Adler et al. 2012].

4.3. Adaptation of Graph Minor for CGRA Mapping

We need to adapt and restrict the definition of graph minor. Graph minor is usually defined for undirected graphs. For directed graphs, the definition of edge contraction is similar to the undirected case [Robertson and Seymour 2004]. Figure 2(e)-(f) show examples of directed edge contractions.

We call the subgraph $M \subseteq G$ defined by the union of $\{\phi(v)|v \in V(H)\}$ and $\{\phi(e)|e \in E(H)\}$ as the schedule and route graph (SRG) of H in G. The SRG M is essentially the model of H in G. The edge set of M is partitioned into the *contraction edges* (the edges in $\{\phi(v)|v \in V(H)\}$) and the *minor edges* (the edges in $\{\phi(e)|e \in E(H)\}$). The minor edges support the data dependencies in the dataflow graph, while the contraction edges represent data routing through additional nodes. For example, in Figure 2(d), $\phi(1)$ is the subgraph inside the dashed region rooted at node 1. The dashed edges are the contraction edges, while the solid edges are the minor edges.

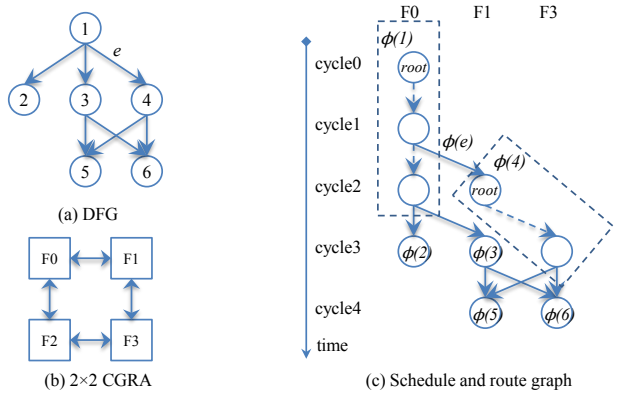

Fig. 5. Minor relationship between DFG and MRRG

Minor edge constraint. In graph minor definition, for $e = (u, v) \in E(H)$, the minor edge $\phi(e)$ connects $\phi(u)$ with $\phi(v)$. In other words, it is sufficient for $\phi(e)$ to connect any node in the subgraph $\phi(u)$ with any node in the subgraph $\phi(v)$. However, for our problem, we need to define one particular node in the subgraph $\phi(v)$ where the actual operation $\phi(v)$ takes place and it has to receive all the required inputs. The remaining nodes in $\phi(v)$ are used to route the result of the operation. More concretely, for our mapping, each subgraph $\phi(v) \subseteq G$ is a tree rooted at the node where the computation takes place. Let $root(\phi(v))$ be the root of the tree $\phi(v)$. Then we introduce the restriction that for $e = (u, v) \in E(H)$, the minor edge $\phi(e)$ connects $\phi(u)$ with $root(\phi(v))$. For example, the DFG in Figure 5(a) has an edge e that connects the DFG nodes 1 and 4, and it is mapped to a 2×2 CGRA shown in Figure 5(b). Then in the SRG, $\phi(1)$ has to connect to the root of $\phi(4)$ through a direct link $\phi(e)$ as shown in Figure 5(c).

Timing constraint. The wrap-around nature of the MRRG introduces another restriction. For SGR M to be a valid mapping, it has to satisfy the timing constraints as follows. For simplicity, let us first ignore the recurrence edges in the DFG. Then the DFG H is a directed acyclic graph. Let $u \in V(H)$ be a node in the DFG without any predecessor and $root(\phi(u)) = (m, t) \in M$

where $0 \leq t < II$ and M is the SRG, a subgraph of the MRRG. That is, u has been mapped to FU m in configuration t in the MRRG. We define the timestamp of u as $cycle(u) = t$, assuming u is executed in cycle t. Let $v \in V(H)$ be a DFG node with u as its predecessor node and $route(u, v)$ be the number of nodes (possibly zero) in the connecting path between $root(\phi(u))$ and $root(\phi(v))$ in the SRG M. For a mapping M to be valid, the following timing constraint, which ensures identical cycle along all input edges of v, must be satisfied for each internal DFG node v.

$$\forall u, u' \in pred(v): \quad cycle(u) + route(u, v) = cycle(u') + route(u', v)$$

We also define

$$\forall u \in pred(v): \quad cycle(v) = cycle(u) + route(u, v) + 1$$

where $pred(v)$ is the set of all predecessors of v in the DFG. Note that we are not doing modulo operation (w.r.t. II) while computing the cycle values. Figure 6 shows this timing computation. In the SRG, $root(\phi(2))$ is in cycle 0 and $root(\phi(3))$ is in cycle 1. However, $root(\phi(2))$ has to go through three routing nodes to reach $root(\phi(4))$; and $root(\phi(3))$ can directly pass the data to $root(\phi(4))$ in the next cycle. The timing constraint is then violated, leading to an invalid mapping.

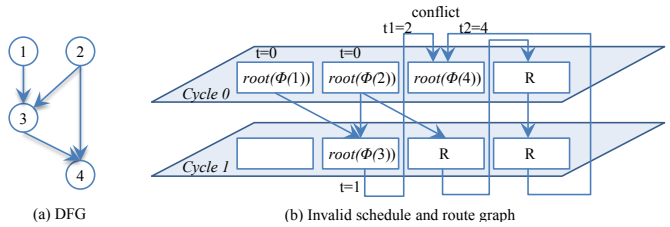

Fig. 6. Invalid mapping under timing constraint

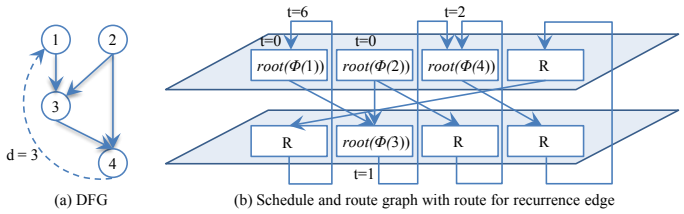

Fig. 7. Mapping with recurrence edge under timing constraint

For a recurrence edge $e = (u, v) \in V(H)$ in the DFG, we introduce additional timing constraint

$$route(u, v) = II \times d + cycle(v) - cycle(u) - 1$$

where d is the recurrence distance of e. Figure 7 shows how this timing constraint is used. Suppose $root(\phi(1))$ is executed in cycle 0; then it will receive the output of $root(\phi(4))$ 6 cycles (3 iterations) later. As $root(\phi(4))$ is executed in cycle 2 ($cycle(4) = 2$), the length of the route from $root(\phi(1))$ to $root(\phi(4))$ should be 2*3+0-2-1 = 3, which means the route contains three routing nodes. In fact, the timing constraint of normal edges are just special cases where distance d is equal to 0.

Attribute constraint. Each node in the DFG and the MRRG has an attribute that specifies the functionality of the node. For example, a node in the DFG can have memory operation as its attribute, while a node in the MRRG can have an attribute that signifies that it can support memory operations. Attribute constraint ensures that a DFG node is mapped to an MRRG tree subgraph whose root has a matching attribute. For example, the root of the tree subgraph for mapping a memory operation can only be a functional unit supporting memory accesses. However, other nodes in the tree subgraph can be any type of functional unit or register file.

Register file constraint. The mapping must ensure availability of register file read/write ports and capacity in the corresponding cycle if a link from/to the register file is used.

Restricted Graph Minor. We can now define application mapping on CGRAs as finding a valid subgraph (schedule and route graph) *M* of the MRRG such that the DFG can be obtained through repeated edge contractions of *M*. We call the DFG a restricted minor of the MRRG and the subgraph M represents the mapping. Alternatively, the DFG H is a minor of G if and only if there exists a model of H, represented by the schedule and route graph M, in G.

LEMMA 4.1.
The restricted graph minor problem for directed graphs is NP-complete.

PROOF. We first show that the restricted graph minor problem for directed graphs is in the set of NP. Given a mapping in the form of SRG $M \subseteq G$, we can check in polynomial time (a) the graphs $\{\phi(v)|v \in V(H)\}$ are mutually vertex-disjoint and the edges $\{\phi(e)|e \in E(H)\}$ are pairwise distinct, (b) for $e = (u, v) \in E(H)$, the edge $\phi(e)$ connects subgraph $\phi(u)$ with $root(\phi(v))$, and (c) the timing constraints as defined earlier are satisfied. That is DFG H is a minor of the G.

We now show that for general directed graphs, the restricted graph minor problem can be reduced to the Hamiltonian cycle problem, which is an NP-complete problem. The Hamiltonian cycle problem is to find a cycle in a directed graph *G* visiting each node exactly once. We can construct a graph *H* which is a directed cycle with $|V(G)|$ nodes. Finding the Hamiltonian cycle in *G* can now be reduced to finding a restricted graph minor between *H* and *G*. As $|V(G)| = |V(H)|$, each subgraph $\phi(v)$ can only consist of a single vertex and each edge mapping $\phi(e)$ where $e = (u, v) \in E(H)$ directly connects vertex $\phi(u)$ to vertex $\phi(v)$. This matches the exact definition of Hamiltonian cycle. Thus the restricted graph minor problem for directed graphs is NP-complete. □

5. GRAPH MINOR MAPPING ALGORITHM

Our solution for restricted graph minor containment problem is inspired by the tree search method (also called state space search) widely used to solve a variety of graph matching problems [Nilsson 1982]. The contribution of our solution is the introduction of customized and effective pruning constraints in the search method that exploit the inherent properties of the data flow graph and the CGRA architecture. We first present the exact restricted graph minor containment algorithm followed by description of additional strategies to accelerate the search process.

5.1. Algorithmic framework

Our goal is to map a DFG H to the CGRA architecture. Similar to the traditional modulo scheduling, we start with the minimum possible II, which is the maximum of the resource constrained II and the recurrence constrained II, that is, $II = \max(ResMII, recMII)$. Given this II value, we create the MRRG G_{II} corresponding to the CGRA architecture. If H is a minor of G_{II}, then the DFG can be mapped with initiation interval II. To check graph minor containment, we check if there exists a model or mapping of H in the form of a valid SRG $M \subseteq G_{II}$. If such SRG M does not exist, we increment the II value by one, create the MRRG corresponding to this new II value, and perform graph minor testing for this new MRRG. This process is repeated till we have generated an MRRG with sufficiently large value of II so that the DFG can satisfy the graph minor test. Algorithm 1 provides a high-level view of our mapping framework.

The core routine of the mapping algorithm *Minor()* performs graph minor testing. We consider all possible mapping between the DFG and the MRRG; thus our algorithm is guaranteed to generate a valid mapping if it exists. Clearly, the number of possible mappings between the DFG and the MRRG is exponential in the number of nodes of the DFG. That is, our search space is large. Our goal is to either (a) quickly identify a mapping such that the DFG passes the restricted minor test, or (b) establish that no such mapping exists. As mentioned earlier, we employ powerful pruning strategies to efficiently navigate this search space. We also carefully choose the order in which we

Algorithm 1: Graph Minor Mapping Algorithm

```
begin
    order_list := DFG_node_ordering(H);
    II := max(resMII, recMII);
    while do
        /*Create MRRG with II*/;
        G_II := Create_MRRG(G, II); M := ⊥;
        for all v ∈ V(H) and e ∈ E(H) do
            φ(v) := ⊥; φ(e) := ⊥;
        add all φ(v), φ(e) to M; /* empty mapping */
        if Minor(H, G_II, M) then
            return(M);
        II++;
Function Minor(H,G,M)
begin
    if no unmapped node in H then
        return(success);
    v := next unmapped node in H according to order_list;
    P := {p | p ∈ pred(v) ∧ φ(p) ≠ ⊥}; /*mapped predecessors of v */
    S := {s | s ∈ succ(v) ∧ φ(s) ≠ ⊥}; /*mapped successors of v */
    /*All candidate mappings are generated satisfying minor edge, timing, attribute, pruning constraints */
    Γ := min_map(v, P, S);
    for each φ(v) ∈ Γ do
        update M with φ(v);
        if Minor(H, G, M) then
            return(success);  /* mapping completed */
    if Γ = ⊥ then
        /* No feasible node mapping; expand predecessors */
        for each possible expansion do
            expand_map(v, P, M);
            /* attempt mapping v again */
            if Minor(H, G, M) then
                return(success);
    /* No node mapping; backtrack to the predecessor */
    return(failure);
```

attempt to map the nodes and the edges so as to achieve quick success in finding a valid mapping or substantial pruning that helps establish the absence of any valid mapping.

The procedure *Minor()* starts with an empty mapping. As mentioned earlier, restricted graph minor mapping for our problem requires mapping each vertex $v \in V(H)$ in the DFG to a tree $\phi(v) \subseteq G$ in the MRRG. Each edge $e = (u, v) \in E(H)$ is simply mapped to an edge $\phi(e) \in E(G)$ that connects some node in $\phi(u)$ to $root(\phi(v))$. Following this definition, we attempt to map the nodes one at a time in some pre-defined priority order, which will be detailed in Section 5.2.

There exist many possibilities to map a node $v \in H$ to a tree subgraph $\phi(v) \subseteq G$. However, the *min_map()* function in Algorithm 1 returns a set Γ of minimal valid mappings $\phi(v)$. Each minimal valid mapping contains minimal number of nodes and satisfies various constraints, including minor edge, timing, attribute and pruning constraints. The minor edge constraint ensures that all the edges connecting the mapped direct predecessors and successors of v can be mapped. More specifically, while mapping node v, we identify all its mapped direct predecessors P and successors S. We ensure that minor edge constraint can be satisfied between each node $p \in P$ and v as well as between v and each node $s \in S$. In other words, if node v has mapped direct successors, then we attempt to generate $\phi(v)$ containing additional routing nodes to ensure that $root(\phi(s))$, $s \in S$, can be reached from some node in $\phi(v)$. Meanwhile, $root(\phi(v))$ should be linked from every $\phi(p)$, $p \in P$. If node v does not have any mapped direct successor, $\phi(v)$ is generated containing a single node. Through *min_map()* function, edge mapping is automatically performed under minor edge constraint checking and we do not need to explicitly map the edges.

In addition, we check for timing constraint between v and its predecessors/successors to ensure that the data is routed correctly. Attribute compatibilities are checked between the DFG node v and the root of the candidate tree subgraph $root(\phi(v))$. If the target CGRA contains register files, the

register constraint is used to check for available ports and capacity. Finally, we also apply aggressive pruning constraints to eliminate mappings that are guaranteed to fail in the future.

If we get non-empty Γ for each node v, then we will eventually obtain a complete feasible solution. However, Γ could be empty if there is no minimal valid mappings. In this case, we have to explore more elaborate tree subgraph mappings for the candidate node v. This is done through *expand_map()* function. In *expand_map()* function, we add one extra node in $\phi(p)$ for each $p \in P$, which helps to enhance the routing path from $\phi(p)$ to $\phi(v)$. If we cannot map v even after all the possible expansions, then we backtrack and attempt a different mapping.

The mapping process continues till we have either mapped all the DFG nodes (i.e., the DFG is a restricted minor of the MRRG) or we have discovered that no such mapping is possible (i.e., the DFG is not a restricted minor of the MRRG) and we have to increment the II value.

5.2. DFG node ordering

An appropriate ordering of the DFG nodes during mapping is crucial to quickly find a feasible solution. We impose the constraint that the nodes along the critical path have higher priority, i.e., they appear earlier. This is because if the critical path cannot be mapped with the current II value, then we can terminate the search process and move on to the next II value.

In addition, we employ an ordering that helps us validate the timing constraints as discussed in Section 4.3. A node *v* is mapped only when at least one of its direct predecessor or successor has been mapped. That is *v* should appear in the ordering after at least one of its direct predecessor or successor nodes. The only exception is the first node in the ordering. The advantage of this ordering is that the timestamps $cycle(v)$ are generated appropriately for the nodes so that timing conflicts can be avoided early. When the DFG contains disjoint parts, a new timestamp is regenerated and propagated for every disjoint component during the mapping process.

Figure 8(b) shows a DFG and the ordering of the nodes through the arrow signs. We start with the input node 1 on the critical path. We proceed along the critical path to node 3 and node 4. Notice that we could not include node 2 after node 1 because none of its direct predecessors or successors would have appeared in the ordering by then. After node 4, we include node 2 in the ordering.

5.3. Mapping Example

Suppose we have a DFG as shown in Figure 8(b) and we are attempting to map it to a 2×2 CGRA array. Let us assume that we are currently considering II=2. For simplicity of exposition, we only draw the occupied edges in the MRRG. The entire mapping process is illustrated in Figures 8(c-g).

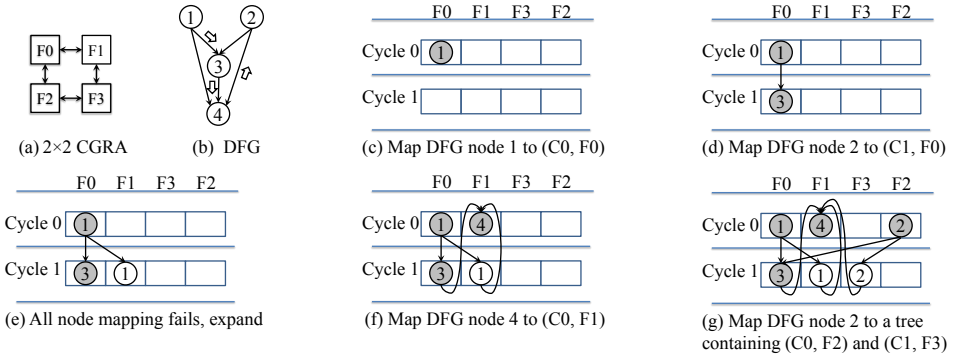

Fig. 8. An example of mapping process during the restricted graph minor test.

The process starts with mapping node 1. Node 1 is the initial node and it has no mapped direct successor. So the first tree subgraph generated by *min_map()* function contains only one node as shown in Fig 8(c): F1 in cycle 0 denoted as (C0, F0). Then we pick the next node in the priority list which is node 3. Again, this node has no mapped direct successors; so its tree mapping also contains

only one node. However, we need to make sure that $\phi(1)$ is directly connected with $root(\phi(3))$ according to the edge constraint imposed by the edge e = (1, 3) in DFG. Mapping node 3 to (C1, F0), as shown in Figure 8(d), can satisfy the constraint.

The next node in the priority list to be mapped is node 4. However, this time we fail to find any feasible node directly connected to the mapped direct predecessors $\phi(1)$ and $\phi(3)$. As mapping for node 4 fails, we expand its predecessor's mapping. An extra node (C1, F1) is added to $\phi(1)$ in Figure 8(e). Notice that to distinguish between root nodes and other nodes, the root nodes have been shadowed. Now node 4 can be mapped to (C0, F1) in Figure 8(f).

The final node in the list is node 2. This time, node 2 has two mapped successors, node 3 and node 4. Thus, we find a tree subgraph $\phi(2)$ containing (C0, F2) and (C1, F3) (see Figure 8(g)) that satisfy both the minor edge constraint (direct links to root nodes of $\phi(3)$ and $\phi(4)$) and the timing constraints at node 3 and 4. As all the nodes and the minor edges have been mapped successfully, DFG is a minor of $MRRG$ with $II = 2$.

5.4. Pruning constraints

Pruning constraints are important to reduce the compilation time. Pruning constraints look ahead and quickly identify if the current mapping can be extended to a successful final mapping. This lookahead helps to eliminate mappings that are guaranteed to fail in the future. Note that the pruning constraints do not affect the optimality of the solution.

Available resource constraint. This constraint simply checks that the number of available FUs of each type in the MRRG is larger than or equal to the number of unmapped DFG nodes of the same type. For example, the number of remaining available memory FUs must be larger or at least equal to the number of unmapped memory operations in the DFG. Global variables are used to record information about the available FUs and the unmapped DFG nodes and are updated every time the partial mapping changes. Thus both time and space complexity of this constraint are *O(1)*.

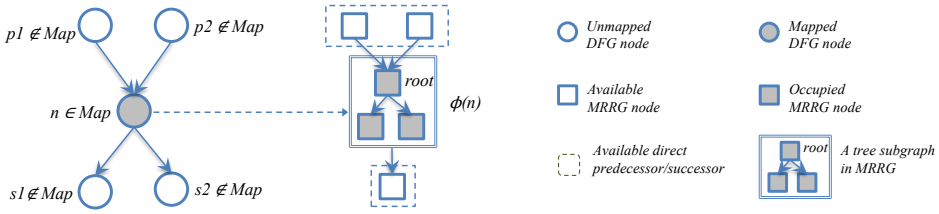

Fig. 9. Illustrations of degree pruning constraint.

Degree constraint. This constraint considers the local structures between the DFG H and the MRRG G. Let $\phi(n) \subseteq G$ be the tree subgraph representing the mapping of node $n \in V(H)$. The number of unmapped direct predecessors of n in the DFG must be smaller than or equal to the number of available direct predecessors of $root(\phi(n))$ in the MRRG.

On the other hand, if n has any unmapped direct successors, then the number of available direct successors of $\phi(n)$ must be at least one. This is because the data from $\phi(n)$ can be routed through any available outgoing node. For example in Figure 9, DFG node n is mapped to $\phi(n)$ in the MRRG. It has two unmapped direct predecessors and two unmapped direct successors. So $root(\phi(n))$ must have at least two available direct predecessors and there must be at least one available direct successor of $\phi(n)$ in the MRRG. Notice that the available direct successors of $\phi(n)$ are those available MRRG nodes directly connected from any node in $\phi(n)$.

The degree pruning constraint checks for all the DFG nodes in the current mapping. The time complexity for this pruning constraint is $O(cN)$, where N is the number of DFG nodes and c is the average number of producer nodes in $\phi(n)$ across all mapped DFG nodes n.

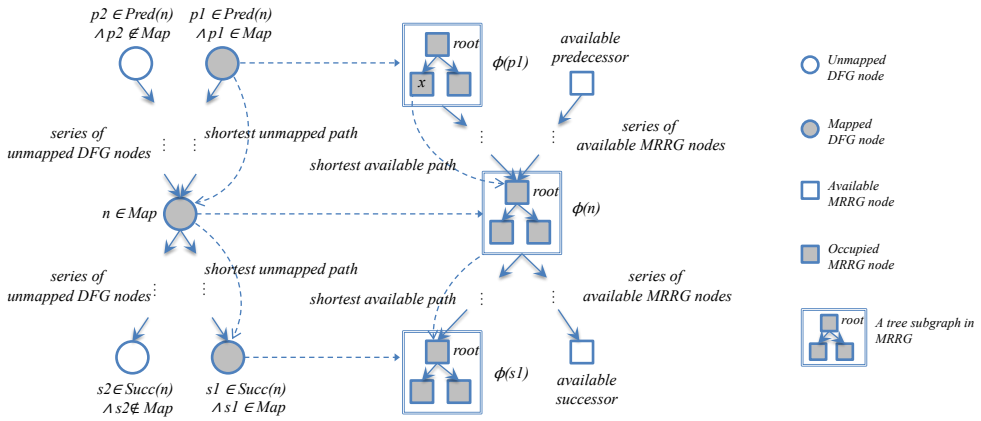

Fig. 10. Illustration of predecessor and successor constraints.

Predecessor and successor constraint. We further exploit structural patterns formed by each mapped DFG node n and its predecessors/successors as shown in Figure 10. We check the timing constraint inherently imposed by these patterns. We first calculate the shortest path lengths in both DFG and MRRG. The shortest paths defined here only consists of unmapped DFG nodes or available MRRG nodes except the two end nodes. For any mapped predecessor p of n, if p and n are connected through the shortest unmapped path $r = (p \rightsquigarrow n)$, then $\phi(p)$ and $\phi(n)$ should also be connected by a shortest available path $R = (x \rightsquigarrow root(\phi(n)))$, $x \in \phi(p)$, in MRRG. Thus we have $cycle(root(\phi(n)) - cycle(x) \geq max(length(R), length(r))$, which uses the fact that the timestamp differences must be at least equal to the length of the shortest path connecting the corresponding nodes either in the MRRG or in the DFG. Similar constraints are also applied to the patterns formed by n and its successors.

We also consider the relationships between a mapped DFG node n and its unmapped predecessors/successors. However, as these predecessors/successors have not been mapped yet, there is no explicit structural information to be used for pruning purpose. Instead, we calculate the number of available MRRG nodes those could be connected to $root(\phi(n))$ (or reached from $\phi(n)$) through available MRRG paths. The number must be at least equal to the number of unmapped predecessors (or successors) of n, which can be connected to (or from) n through unmapped DFG paths.

To obtain the reachability information in both the DFG and the MRRG during the mapping, two reachability matrices are built using an efficient algorithm by Italiano et al. [Italiano 1986]. The algorithm has a time complexity $O(K)$ with $O(K^2)$ space overhead, where K is the number of nodes in the input graph. Each element (u, v) in the matrix represents the shortest path length between the node u and node v. To build the reachability matrix for M MRRG nodes, the time complexity is $O(M^2)$. As the computation for reachability matrices is the most time consuming step, the overall time complexity for the pruning constraint is $O(M^2)$.

Feasibility constraint. In the final pruning constraint, we exploit the structural patterns of the unmapped DFG nodes. As shown in Figure 11, for each unmapped DFG node, we find all its mapped predecessors and successors reachable through unmapped paths. There must be at least one MRRG node that has the same connectivity to all the subgraphs the corresponding predecessors and successors have been mapped to. More specifically, let n is such an unmapped DFG node, p is a mapped predecessor of n and p is connected to n through an unmapped path. Then in the MRRG, there must be at least one available node m such that m could be connected from $\phi(p)$ through an available path. As this pruning constraint also depends on the reachability matrices, the complexity is $O(M^2)$.

5.5. Acceleration strategies

We now introduce additional strategies to further accelerate the compilation time. These strategies are integrated in the preprocessing step and the constraints in the algorithm infrastructure. All the strategies are designed in such a way that they do not impact the optimality of the mapping.

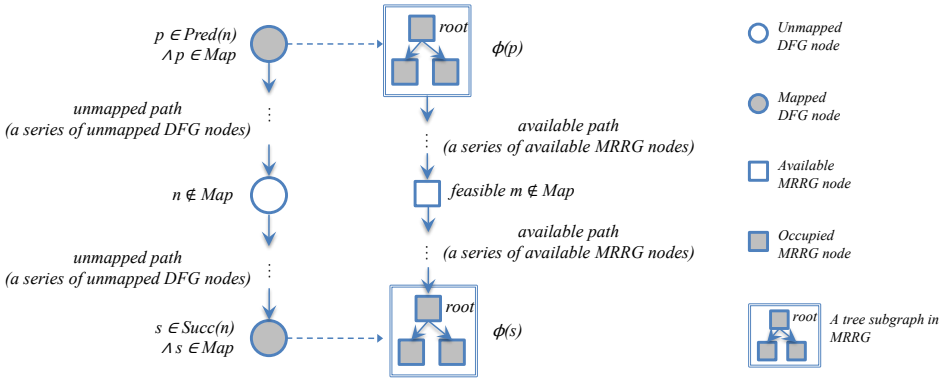

Fig. 11. Illustration of feasibility constraint.

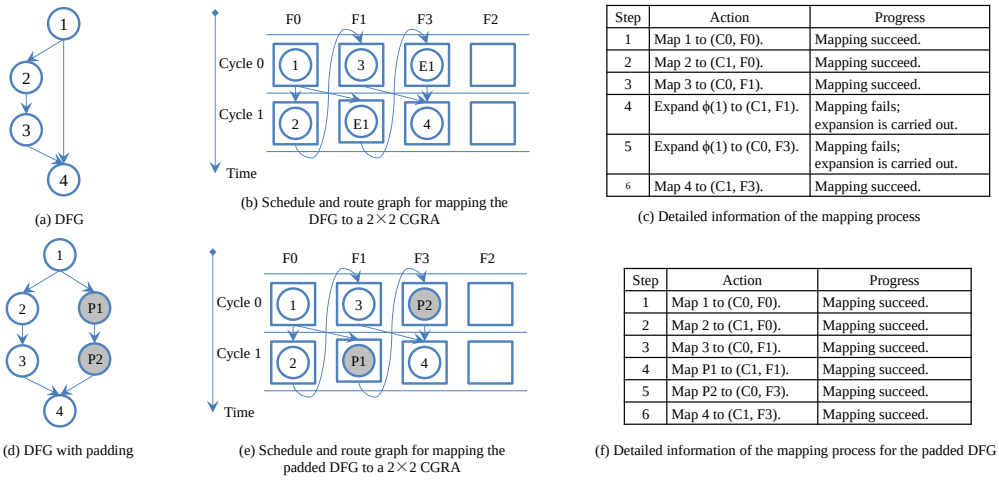

Step	Action	Progress
1	Map 1 to (C0, F0).	Mapping succeed.
2	Map 2 to (C1, F0).	Mapping succeed.
3	Map 3 to (C0, F1).	Mapping succeed.
4	Expand ϕ(1) to (C1, F1).	Mapping fails; expansion is carried out.
5	Expand ϕ(1) to (C0, F3).	Mapping fails; expansion is carried out.
6	Map 4 to (C1, F3).	Mapping succeed.

(c) Detailed information of the mapping process

Step	Action	Progress
1	Map 1 to (C0, F0).	Mapping succeed.
2	Map 2 to (C1, F0).	Mapping succeed.
3	Map 3 to (C0, F1).	Mapping succeed.
4	Map P1 to (C1, F1).	Mapping succeed.
5	Map P2 to (C0, F3).	Mapping succeed.
6	Map 4 to (C1, F3).	Mapping succeed.

(f) Detailed information of the mapping process for the padded DFG

Fig. 12. A motivating example for dummy node insertion.

5.5.1. Dummy nodes in the DFG. We introduce dummy nodes in the DFG during the preprocessing step. These dummy nodes are *only* used for routing, which means they can be mapped to non-computation nodes in the MRRG, e.g., register file nodes. Basically, the idea is based on the observation that expanding the tree mapping $\phi(v)$ for any node v is quite expensive. This is because $\phi(v)$ is expanded only after all attempts to map subsequent nodes have failed. Also the expansion is carried out incrementally, i.e., $\phi(v)$ is expanded one node at a time. The goal of introducing dummy nodes is to avoid the expansions as much as possible without affecting the quality of the solution.

Figure 12 shows an example of how dummy nodes can avoid expansion of node mapping. We want to map the DFG in Figure 12(a) to 2×2 CGRA. The mapping order is $1 \rightarrow 2 \rightarrow 3 \rightarrow 4$. The first three nodes 1, 2, and 3 can be mapped successfully. However, when we try to map node 4, the mapping attempt fails (Γ is empty) and we have to expand $\phi(1)$ twice in order to find the final feasible mapping for node 4. The final schedule and route graph is shown in Figure 12(b) with the expansion nodes for $\phi(1)$ denoted as $E1$. The detailed search process is also listed in Figure 12(c).

To avoid the mapping failures and expansions, we can add two dummy nodes $P1$ and $P2$, as shown in Figure 12(d). Suppose the mapping order for the new DFG is $1 \rightarrow 2 \rightarrow 3 \rightarrow P1 \rightarrow P2 \rightarrow 4$. After mapping the three nodes 1, 2 and 3, we will continue to map $P1$ and $P2$ without any failure. Finally, node 4 will be mapped successfully at the first attempt. The final schedule and route graph is shown in Figure 12(e) and the detailed mapping process is listed in Figure 12(f).

Clearly, dummy node insertion is useful in guiding the mapping process. So we add dummy nodes as part of DFG pre-processing step. We first assign scheduling levels to each DFG node using as soon as possible (*ASAP*) scheduling policy and as late as possible (*ALAP*) scheduling policy. The number of dummy nodes inserted to a DFG edge $e = (u, v) \in E(H)$ is equal to the difference between the *ASAP* level of v and the *ALAP* level of u. This is somewhat similar in concept to node balancing in [Hamzeh et al. 2012]. However, the difference is that we insert dummy nodes to accelerate the search process to obtain a feasible schedule. In the previous approach [Hamzeh et al. 2012], adding more balancing nodes is a requirement to obtain a valid schedule.

5.5.2. Fast implementation of pruning constraints. For large DFGs, the pruning constraints can increase the compilation time. The most expensive part is the reachability matrices computation. To reduce this overhead, we bypass updating the reachability matrix of the MRRG at each step. We do, however, generate the reachability information for the DFG statically in the beginning and for the MRRG at its generation step for each II value. We believe that the two static matrices provide limited but enough information for the pruning purposes. The static reachability matrices now record the reachability information between any two arbitrary nodes in the absence of any mapping, e.g., the element (x, y) in the MRRG matrix records the static shortest path length between nodes x and y. With only static reachability matrix, the pruning constraints have to be redesigned as follows.

Fast implementation of predecessor and successor constraints. Unlike the original constraint, the fast implementation only focuses on the structural patterns related to current mapping. Suppose the candidate DFG node n is mapped to $\phi(n)$ in the MRRG. For every mapped predecessor p of n, we can have the length value for the static shortest path $r_s = (p \rightsquigarrow n)$, from static DFG matrix. Let $R_S = (x \rightsquigarrow \phi(n))$, $x \in \phi(p)$, be the static shortest path between $\phi(p)$ and $\phi(n)$ in MRRG. x can be identified by checking the static MRRG matrix for all the nodes in $\phi(p)$. Utilizing the same fact used in the original constraint, we have $cycle(root(\phi(n))) - cycle(x) \geq max(length(R_S), length(r_s))$

Similarly, constraints are also imposed for the structural patterns formed by the candidate node and its mapped successors. The fast implementation reduces the runtime complexity from $O(M^2)$ to $O(cN)$ where c is the average number of nodes in $\phi(n)$ for each DFG node n.

Fast implementation of feasibility constraint. The basic idea for designing fast implementation of feasibility constraint is to consider the local effects of consuming one MRRG node for the remaining unmapped DFG nodes. Suppose the candidate MRRG node to be used for mapping is m, then the consumption will affect the potential mappings of those who also require m. If m is directly linked from any node in $\phi(p)$, p is a mapped DFG node, then the consumption of m can affect the mapping for the unmapped child $child_p$ of p. In other words, we need to ensure that apart from m, there is another available MRRG node m' that can be used to map $child_p$ satisfying certain timing constraints. For every mapped successor s of $child_p$, we can have the static shortest path $r_s = (child_p \rightsquigarrow s)$. Let R_S be the static shortest path connecting m' and $root(\phi(s))$, $R_S = (m' \rightsquigarrow root(\phi(s)))$. Following the same reasoning used before, we have

$$cycle(root(\phi(s))) - cycle(m') \geq max(length(R_S), length(r_s))$$

If m is a direct predecessor of the root node of $\phi(s')$, where s' is a mapped DFG node, similar constraints are used for the unmapped parent node of s'. The time complexity is also $O(cN)$.

5.6. Integration of Heuristics

Our modulo scheduling algorithm (Algorithm 1) can achieve the optimal II by definition. This is because it checks if the DFG is a minor of the MRRG for each value of II, starting with the minimum possible value. However, even with the pruning and acceleration strategies, the runtime of the optimal algorithm can be prohibitive when both the number of DFG nodes and the number of CGRA functional units are quite large. Therefore, we integrate some heuristics in the algorithm to speed up the search process. This may introduce sub-optimality, i.e., the search process may miss a valid mapping at lower II value even though it exists. But the compilation time improves significantly.

The first heuristic avoids backtracking between two unrelated nodes. In the optimal search process, if a node m cannot be mapped, then we backtrack to the node n which appears just before m in the DFG node ordering. However, node n may not be a predecessor or successor of node m in the DFG and hence may not be able to steer the search towards a successful mapping to m. Instead, we directly backtrack to the last predecessor or successor of node m in the ordering.

The second heuristic is motivated by the edge-centric mapping [Park et al. 2008]. In graph minor testing, instead of enumerating all possible tree subgraphs for node n, the procedure aims to find limited number of feasible subgraphs. The feasible subgraphs are chosen to be those with minimal number of nodes. After all the specified subgraphs have been explored, the node mapping fails.

The final heuristic makes it possible to escape from extensive subgraph expansions. We put a counter for each node mapping. The counter is increased every time an expansion is carried out. Once the counter reaches a pre-defined threshold value, we eliminate current mapping and backtrack to previous mappings. Our experimental evaluation reveals that this is the only heuristic that sometimes prevent us from reaching a feasible solution even if one exists.

6. EXPERIMENTAL EVALUATION

We now proceed to evaluate the quality and the efficiency of our mapping algorithm. We initially target a 4×4 CGRA with 2D mesh network architecture and no shared or central register file. The 4×4 array is the basic structure in many CGRA architectures and has been widely used to evaluate various mapping algorithms [Park et al. 2006; Park et al. 2008; Kwok and Wilton 2005; Hatanaka and Bagherzadeh 2007; Kim et al. 2011; Bansal et al. 2003]. For our initial experiments that compare against previous approaches, we assume each functional unit is comprehensive and is capable of handling any operation including memory operations. Later, we evaluate the versatility of graph minor mapping approach in supporting diverse CGRA architectures, such as heterogeneous functional units and various register file configurations. We also evaluate the scalability issue by mapping to 4×8, 8×8, 8×16 and 16×16 CGRAs.

Table I. Benchmark characterisitcs

Benchmark	*#ops*	*#MEM ops*	*#edges*	*Benchmark*	*#ops*	*#MEM ops*	*#edges*
SOR	17	6	11	***osmesa***	16	9	17
swim_cal1	59	23	39	***texture***	29	7	31
swim_cal2	62	26	44	***quantize***	21	8	24
sobel	27	7	34	***rgb2ycc***	41	15	44
lowpass	23	9	19	***rijndael***	32	13	35
laplace	20	8	16	***fft***	40	20	42
wavelet	12	4	6	***tiff2bw***	42	20	50
sjeng	36	13	21	***fdctfst***	59	16	80
scissor	12	4	13	***idctflt***	87	25	114

We select loop kernels from MiBench benchmark suite [Guthaus et al. 2001], SPEC2006 benchmark suite, and the benchmarks used in the EPIMap approach [Hamzeh et al. 2012]. Most of the benchmarks have an easily identifiable compute-intensive loop that performs the main functionality of the application and we select that loop for our experiments. For the few benchmarks with multiple loop kernels, we choose the representative one of them. *Rijndael* implements the AES standard where we choose the nested loop in its encryption subroutine. *Tiff2bw* converts a color TIFF image to greyscale image where we choose the nested loop in the first step that converts 16-bit color map to 8-bit. The benchmarks *Wavelet, Fdctfst, Idctfst* have multiple identical or similar loops and we choose one of them.

The DFGs for the loop kernels are generated from Trimaran [Chakrapani et al. 2005] back-end using Elcor intermediate representation [Aditya et al. 1998]. Benchmark characteristics are listed in Table I including the number of operations and the number of load/store operations. We assume that the memory operation includes both the address generation and the actual load/store operation.

Comparison with different techniques. There exist a number of approaches to CGRA mapping as presented in Section 2. We compare our graph minor approach (abbreviated as G-Minor here) with two previous techniques: simulated annealing based approaches and EPIMap [Hamzeh et al. 2012]. Simulated annealing (SA) based approaches [Mei et al. 2003b] are widely considered to provide high-quality mapping solutions with (possibly) longer compilation time. EMS, the edge-centric mapping approach [Park et al. 2008], provides significantly reduced compilation time with some degradation in the quality of the schedule compared to SA. As mentioned in Section 1, in parallel to G-Minor approach, [Hamzeh et al. 2012] have proposed graph epimorphism based mapping approach EPIMap that produces better quality solutions than EMS with similar compilation time. We compare G-Minor with EPIMap as it represents state-of-the-art CGRA mapping approach. For the comparison, we have re-implemented the EPIMap approach [Hamzeh et al. 2012] and the simulated annealing (SA) algorithm [Mei et al. 2003b] for 4×4 mesh CGRA with comprehensive functional units and no shared/central register file similar to the setup in [Hamzeh et al. 2012]. Our implementations of these two approaches allow route sharing. To demonstrate the benefits gained from using route sharing, we also create a subgraph homeomorphism mapping kernel. Moreover, we also integrate re-computation methodology introduced in EPIMap as a DFG pre-processing step in our G-Minor framework.

Figure 13 compares the scheduling quality for 18 benchmarks. The Y-axis represents the achieved II value. The first bar represents the minimal II value achievable considering only recurrence minimal and resource minimal II for each kernel. The remaining bars from left to right represent the II achieved for G-Minor, EPIMap, simulated annealing (SA), subgraph homeomorphism, and G-Minor with re-computation pre-processing (Rec-G-Minor) respectively.

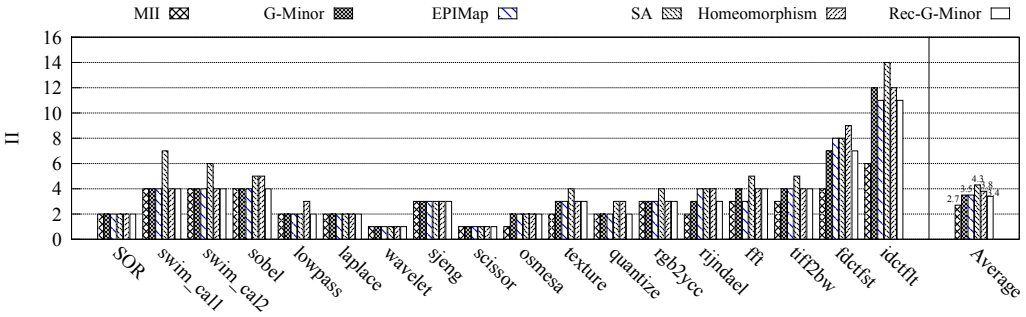

Fig. 13. Scheduling quality for G-Minor, EPIMap, SA, subgraph homeomorphism and G-Minor with re-computation

We first observe that the scheduling quality generated by EPIMap and G-Minor are quite similar. The achieved II value is different between the two for only 4 out of 18 benchmarks. For example, G-Minor produces better scheduling results for *rijndael* and *fdctfst*, while EPIMap performs better for *fft* and *idctflt*. Even for these benchmarks, the difference is only one cycle. The two reasons for the competitive results between G-Minor and EPIMap are the following. G-Minor exhaustively searches for minor with all routing possibilities, while EPIMap restricts the number of routing nodes. On the other hand, EPIMap provides extra choices for mapping the DFGs such as replication (or re-computation) for high fan-out nodes. An interesting possible future research direction would be to combine the relative strengths of G-Minor and EPIMap. We conduct preliminary evaluation by integrating re-computation with our G-Minor framework. It is shown in Figure 13 that in most cases, Rec-G-Minor can generate better scheduling results than G-Minor and EPIMap.

We observe that for a large subset of benchmarks (11 out of 18), both G-Minor and EPIMap achieve Minimal II (MII). SA, on the other hand, achieves minimal II value for 6 benchmarks. In general, G-Minor and EPIMap provide better schedules compared to SA. A possible reason is that in SA, a random operation is picked, replaced and routed in each step. It is inefficient in considering the placement and routing impacts among operations. This inefficiency gets worse when the routing

resources are limited such as in a 4×4 mesh CGRA. G-Minor and EPIMap, on the other hand, directly explore the structural properties of the graphs and hence the relationships among operations.

We carry out additional experiments to demonstrate the benefits of route sharing. We disable route sharing in our G-Minor algorithm to create a subgraph homeomorphism kernel. As shown in Figure 13, subgraph homeomorphism generates far worse schedules compared to G-Minor.

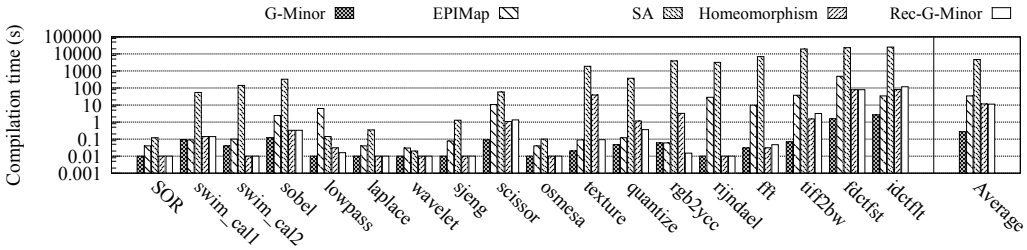

Fig. 14. Compilation time for G-Minor, EPIMap, SA, subgraph homeomorphism and G-Minor with re-computation

The runtime of the different approaches for all the benchmarks are shown in Figure 14, which is reported based on an Intel Quad-Core running at 2.83GHz with 3GB memory. It is well known that SA approaches require longer compilation time [Park et al. 2008] specially for large kernels. Similar compilation time has been reported in [Hamzeh et al. 2013]. G-Minor and EPIMap reduce compilation time significantly using more guided approach to mapping. The average compilation time for EPIMap is 34.26 sec, which is consistent with the timing reported in [Hamzeh et al. 2012]. G-Minor provides extremely fast compilation time of only 0.27 sec on an average. This is because the graph minor testing algorithm in G-Minor has been highly optimized using various pruning constraints and different acceleration strategies. EPIMap transforms the DFG and uses it as an input to an off-the-shelf maximal common subgraph (MCS) kernel [Levi 1973]. Thus the compilation time for EPIMap depends on the efficiency of the chosen MCS kernel. Besides, EPIMap might need to transform the DFG and repeat the MCS kernel computation multiple times when the mapping fails. This potentially leads to longer compilation time.

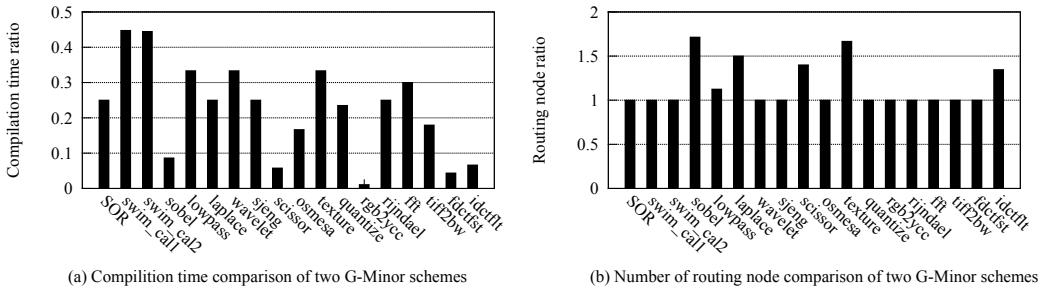

(a) Compilition time comparison of two G-Minor schemes

(b) Number of routing node comparison of two G-Minor schemes

Fig. 15. Experimental results for fast G-Minor scheme (with acceleration strategies) compared to slow G-Minor scheme.

Impact of acceleration strategies and heuristics. We evaluate reduction in compilation time using the acceleration strategies presented in Section 5.5. We compare compilation time for two different versions of G-Minor: the slow mode and the fast mode in Figure 15. The fast mode uses the acceleration strategies. Both modes achieve identical II for all the benchmarks because the acceleration strategies are designed such that they do not impact the quality of the solutions, but provide better guidance for the search process. In Figure 15(a), the compilation time of the fast mode is normalized w.r.t. the slow mode. The fast mode can effectively reduce the compilation time by more than 50%. The penalty for the fast mode is in the form of using more routing nodes. Figure 15(b) compares the number of routing nodes for the two schemes. The average ratio is around 1.15, which

means there are 15% extra routing nodes used in fast mode because the fast pruning constraints using static shortest path connectivity information can lead to more node expansions. The heuristics play crucial roles in achieving reasonable compilation time. In our experiments, 9 out of the 17 benchmarks will fail to return a feasible solution within 10 hours without the heuristics. Meanwhile, the II values of the remaining benchmarks match the results generated with heuristics.

Different CGRA configurations. As mentioned in Section 3.1, our approach can support different CGRA configurations. The experiment results for 4×4 CGRAs with different number of memory units and different register file configurations are shown in Figure 16. MxC denotes the availability of x columns of memory FUs in the array; and y is the number of registers in a register file. So an architectural configuration MxC-LRF-yR corresponds to an array with x columns of memory units and locally shared register files, each of which contains y registers. Each register file is associated with two read ports and one write port. The results indicate that memory units are the most critical resources. Adding more memory units brings substantial benefit by reducing the achieved II. However, adding more registers may not necessarily improve II. This is because the intelligent exploration of the search space can find mappings within limited routing resources. Adding more routing resources such as increasing the size of local/global register files can reduce the mapping efforts but could also end up with resource wastage. We notice that starting from M2C-LRF-1R configuration, increasing the number of registers and providing more connectivity through registers for routing do not reduce the value of the achieved II.

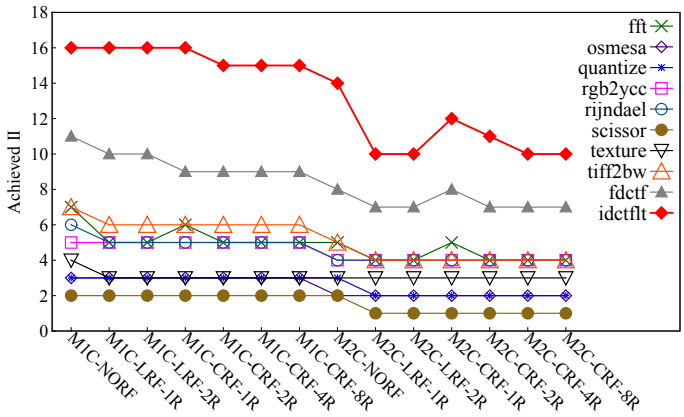

Fig. 16. Achieved II for different CGRA configurations.

Scalability. Our G-Minor fast mode can dramatically accelerate the compilation time. We test the scalability by configuring the size of NORF CGRA to 4×8, 8×8, 8×16 and 16×16 2D-mesh. To further stress the scalability, we generate 100 random DFGs where number of nodes is uniformly distributed in the range (0, 100]. We present the average compilation time for G-Minor and EPIMap with different CGRA sizes in Table II. The results confirm that G-Minor provides better scalability to map kernels on large CGRAs. We do not report compilation time for SA approaches as it takes too long to generate solutions for large CGRAs.

Table II. Compilation time for CGRAs with different sizes

	4×4 CGRA	4×8 CGRA	8×8 CGRA	8×16 CGRA	16×16 CGRA
Avg. compilation time (s) of G-Minor	0.23	0.61	1.51	3.12	7.08
Avg. compilation time (s) of EPIMap	54.78	570.72	837.92	1235.18	1385.27

7. CONCLUSIONS

We present an efficient technique for application mapping on CGRAs. We formalize the CGRA mapping problem as restricted graph minor containment of the data flow graph representing the computation kernel in the modulo routing resource graph representing the CGRA architecture. We design a customized and efficient graph minor search algorithm for our problem that employs aggressive pruning and acceleration strategies. We conduct extensive experimental evaluations of our approach and show that it achieves quality schedule with minimal compilation time.

REFERENCES

Shail Aditya, Vinod Kathail, and B Ramakrishna Rau. 1998. *Elcor's machine description system: Version 3.0*. Hewlett Packard Laboratories.

Isolde Adler, Frederic Dorn, Fedor V Fomin, Ignasi Sau, and Dimitrios M Thilikos. 2012. Fast minor testing in planar graphs. *Algorithmica* 64, 1 (2012), 69–84.

Mythri Alle, Keshavan Varadarajan, Reddy C Ramesh, Joseph Nimmy, Alexander Fell, Adarsha Rao, SK Nandy, and Ranjani Narayan. 2008. Synthesis of application accelerators on Runtime Reconfigurable Hardware. In *Proceedings of the 2008 international conference on Application-Specific Systems, Architectures and Processors (ASAP'08)*. IEEE, 13–18.

Nikhil Bansal, Sumit Gupta, Nikil Dutt, and Alexandru Nicolau. 2003. Analysis of the performance of coarse-grain reconfigurable architectures with different processing element configurations. In *the 2003 Workshop on Application Specific Processors, held in conjunction with the International Symposium on Microarchitecture (MICRO)*.

Janina A Brenner, Sándor P Fekete, and Jan C van der Veen. 2009. A minimization version of a directed subgraph homeomorphism problem. *Mathematical Methods of Operations Research* 69, 2 (2009), 281–296.

Lakshmi N Chakrapani, John Gyllenhaal, W Hwu Wen-mei, Scott A Mahlke, Krishna V Palem, and Rodric M Rabbah. 2005. Trimaran: An infrastructure for research in instruction-level parallelism. In *Languages and Compilers for High Performance Computing*. Springer, 32–41.

Liang Chen and Tulika Mitra. 2012. Graph minor approach for application mapping on CGRAs. In *Proceedings of the 2012 International Conference on Field-Programmable Technology (ICFPT'12)*. IEEE, 285–292.

Nathan Clark, Amir Hormati, Scott Mahlke, and Sami Yehia. 2006. Scalable subgraph mapping for acyclic computation accelerators. In *Proceedings of the 2006 international conference on Compilers, architecture and synthesis for embedded systems (CASES'06)*. ACM, 147–157.

Luigi P Cordella, Pasquale Foggia, Carlo Sansone, and Mario Vento. 2004. A (sub) graph isomorphism algorithm for matching large graphs. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 26, 10 (2004), 1367–1372.

Bjorn De Sutter, Paul Coene, Tom Vander Aa, and Bingfeng Mei. 2008. Placement-and-routing-based register allocation for coarse-grained reconfigurable arrays. In *Proceedings of the 2008 ACM SIGPLAN-SIGBED Conference on Languages, Compilers and Tools for Embedded System (LCTES'08)*. ACM, 151–160.

Giuseppe Di Battista, Maurizio Patrignani, and Francesco Vargiu. 1998. A split&push approach to 3D orthogonal drawing. In *Graph Drawing*. Springer, 87–101.

Christine Eisenbeis, Sylvain Lelait, and Bruno Marmol. 1995. The meeting graph: a new model for loop cyclic register allocation. In *Proceedings of the 1995 International Federation for Information Processing Working Group*. 264–267.

Steven Fortune, John Hopcroft, and James Wyllie. 1980. The directed subgraph homeomorphism problem. *Theoretical Computer Science* 10, 2 (1980), 111–121.

Stephen Friedman, Allan Carroll, Brian Van Essen, Benjamin Ylvisaker, Carl Ebeling, and Scott Hauck. 2009. SPR: an architecture-adaptive CGRA mapping tool. In *Proceedings of the 17th ACM/SIGDA International Symposium on Field Programmable Gate Arrays (FPGA'09)*. ACM, 191–200.

Rani Gnanaolivu, Theodore S Norvell, and Ramachandran Venkatesan. 2010. Mapping loops onto coarse-grained reconfigurable architectures using particle swarm optimization. In *Proceedings of the 2010 International Conference on Soft Computing and Pattern Recognition (SoCPaR'10)*. IEEE, 145–151.

Rani Gnanaolivu, Theodore S Norvell, and Ramachandran Venkatesan. 2011. Analysis of Inner-Loop Mapping onto Coarse-Grained Reconfigurable Architectures Using Hybrid Particle Swarm Optimization. *International Journal of Organizational and Collective Intelligence* 2, 2 (2011), 17–35.

Matthew R Guthaus, Jeffrey S Ringenberg, Dan Ernst, Todd M Austin, Trevor Mudge, and Richard B Brown. 2001. MiBench: A free, commercially representative embedded benchmark suite. In *the 2001 IEEE International Workshop on Workload Characterization*. IEEE, 3–14.

Mahdi Hamzeh, Aviral Shrivastava, and Sarma Vrudhula. 2012. EPIMap: using epimorphism to map applications on CGRAs. In *Proceedings of the 49th Annual Design Automation Conference (DAC'12)*. ACM, 1284–1291.

Mahdi Hamzeh, Aviral Shrivastava, and Sarma Vrudhula. 2013. REGIMap: Register-Aware Application Mapping on Coarse-Grained Reconfigurable Architectures (CGRAs). In *Proceedings of the 50th Annual Design Automation Conference (DAC'13)*. ACM, 18:1–18:10.

Akira Hatanaka and Nader Bagherzadeh. 2007. A modulo scheduling algorithm for a coarse-grain reconfigurable array template. In *Proceedings of the 21th International Parallel and Distributed Processing Symposium (IPDPS'07)*. IEEE, 1–8.

Giuseppe F. Italiano. 1986. Amortized efficiency of a path retrieval data structure. *Theoretical Computer Science* 48, 2-3 (1986), 273–281.

Yongjoo Kim, Jongeun Lee, Aviral Shrivastava, Jonghee W Yoon, Doosan Cho, and Yunheung Paek. 2011. High throughput data mapping for coarse-grained reconfigurable architectures. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 30, 11 (2011), 1599–1609.

Ian Kuon and Jonathan Rose. 2007. Measuring the gap between FPGAs and ASICs. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 26, 2 (2007), 203–215.

Zion Kwok and Steven JE Wilton. 2005. Register file architecture optimization in a coarse-grained reconfigurable architecture. In *Proceedings of the 13th International Symposium on Field-Programmable Custom Computing Machines (FCCM'05)*. IEEE, 35–44.

Dongwook Lee, Manhwee Jo, Kyuseung Han, and Kiyoung Choi. 2009. FloRA: Coarse-grained reconfigurable architecture with floating-point operation capability. In *Proceedings of the 2009 International Conference on Field-Programmable Technology (ICFPT'09)*. IEEE, 376–379.

Jong-eun Lee, Kiyoung Choi, and Nikil D Dutt. 2003. Compilation approach for coarse-grained reconfigurable architectures. *Design & Test of Computers, IEEE* 20, 1 (2003), 26–33.

Giorgio Levi. 1973. A note on the derivation of maximal common subgraphs of two directed or undirected graphs. *Calcolo* 9, 4 (1973), 341–352.

Alan Marshall, Tony Stansfield, Igor Kostarnov, Jean Vuillemin, and Brad Hutchings. 1999. A reconfigurable arithmetic array for multimedia applications. In *Proceedings of the 7th ACM/SIGDA international symposium on Field Programmable Gate Arrays (FPGA'99)*. ACM, 135–143.

Bingfeng Mei, Serge Vernalde, Diederik Verkest, Hugo De Man, and Rudy Lauwereins. 2003a. ADRES: An architecture with tightly coupled VLIW processor and coarse-grained reconfigurable matrix. In *Proceedings of the 13th International Conference on Field Programmable Logic and Application (FPL'03)*. Springer, 61–70.

Bingfeng Mei, Serge Vernalde, Diederik Verkest, Hugo De Man, and Rudy Lauwereins. 2003b. Exploiting loop-level parallelism on coarse-grained reconfigurable architectures using modulo scheduling. In *Proceedings of the 2003 Conference on Design, Automation and Test in Europe (DATE'03)*. IEEE, 296–301.

Nils J Nilsson. 1982. *Principles of artificial intelligence*. Springer.

Hyunchul Park, Kevin Fan, Manjunath Kudlur, and Scott Mahlke. 2006. Modulo graph embedding: mapping applications onto coarse-grained reconfigurable architectures. In *Proceedings of the 2006 International Conference on Compilers, Architecture and Synthesis for Embedded Systems (CASES'06)*. ACM, 136–146.

Hyunchul Park, Kevin Fan, Scott A Mahlke, Taewook Oh, Heeseok Kim, and Hong-seok Kim. 2008. Edge-centric modulo scheduling for coarse-grained reconfigurable architectures. In *Proceedings of the 17th International Conference on Parallel Architectures and Compilation Techniques (PACT'08)*. ACM, 166–176.

B Ramakrishna Rau. 1994. Iterative modulo scheduling: An algorithm for software pipelining loops. In *Proceedings of the 27th International Symposium on Microarchitecture (MICRO'94)*. ACM, 63–74.

Neil Robertson and Paul D Seymour. 1999. Graph minors. *Journal of Combinatorial Theory Series B* 77, 1 (1999), 162–210.

Neil Robertson and Paul D Seymour. 2004. Graph minors. XX. Wagner's conjecture. *Journal of Combinatorial Theory, Series B* 92, 2 (2004), 325–357.

Hartej Singh, Ming-Hau Lee, Guangming Lu, Fadi J Kurdahi, Nader Bagherzadeh, and Eliseu M Chaves Filho. 2000. MorphoSys: an integrated reconfigurable system for data-parallel and computation-intensive applications. *IEEE Trans. Comput.* 49, 5 (2000), 465–481.

M Tuhin and Theodore S Norvell. 2008. Compiling parallel applications to coarse-grained reconfigurable architectures. In *Proceedings of the 2008 Canadian Conference on Electrical and Computer Engineering (CCECE'08)*. IEEE, 1723–1728.

Julian R Ullmann. 1976. An algorithm for subgraph isomorphism. *J. ACM* 23, 1 (1976), 31–42.

Jonghee W Yoon, Aviral Shrivastava, Sanghyun Park, Minwook Ahn, and Yunheung Paek. 2009. A graph drawing based spatial mapping algorithm for coarse-grained reconfigurable architectures. *IEEE Transactions on Very Large Scale Integration Systems* 17, 11 (2009), 1565–1578.

Javier Zalamea, Josep Llosa, Eduard Ayguadé, and Mateo Valero. 2003. MIRS: Modulo scheduling with integrated register spilling. In *the 2003 Workshop on Languages and Compilers for Parallel Computing*. Springer, 239–253.