

A Runtime FPGA Placement and Routing Using Low-Complexity Graph Traversal

RICARDO FERREIRA, LUCIANA ROCHA, ANDRÉ G. SANTOS, and JOSÉ A. M. NACIF,
Universidade Federal de Viçosa
STEPHAN WONG, TU Delft
LUIGI CARRO, Universidade Federal do Rio Grande do Sul

Dynamic Partial Reconfiguration (DPAr) enables efficient allocation of logic resources by adding new functionalities or by sharing and/or multiplexing resources over time. Placement and routing (P&R) is one of the most time-consuming steps in the DPAr flow. P&R are two independent NP-complete problems, and, even for medium size circuits, traditional P&R algorithms are not capable of placing and routing hardware modules at runtime. We propose a novel runtime P&R algorithm for Field-Programmable Gate Array (FPGA)-based designs. Our algorithm models the FPGA as an implicit graph with a direct correspondence to the target FPGA. The P&R is performed as a graph mapping problem by exploring the node locality during a depth-first traversal. We perform the P&R using a greedy heuristic that executes in polynomial time. Unlike state-of-the-art algorithms, our approach does not try similar solutions, thus allowing the P&R to execute in milliseconds. Our algorithm is also suitable for P&R in fragmented regions. We generate results for a manufacturer-independent virtual FPGA. Compared with the most popular P&R tool running the same benchmark suite, our algorithm is up to three orders of magnitude faster.

Categories and Subject Descriptors: B.7.2 [Integrated Circuits]: Design Aids—*Placement and routing*

General Terms: Design, Algorithms, Performance

Additional Key Words and Phrases: FPGA, Place and route, Run-time reconfiguration, Graph traversal

ACM Reference Format:

Ricardo Ferreira, Luciana Rocha, André G. Santos, José A. M. Nacif, Stephan Wong, and Luigi Carro. 2015. A run-time FPGA placement and routing using low complexity graph traversal. *ACM Trans. Reconfig. Technol. Syst.* 8, 2, Article 9 (March 2015), 16 pages.
DOI: <http://dx.doi.org/10.1145/2660775>

1. INTRODUCTION

Modern Field-Programmable Gate Arrays (FPGAs) provide support to perform on-the-fly reconfiguration while other untouched hardware areas continue execution. This feature, known as Dynamic Partial Reconfiguration (DPAr), allows efficient allocation of the FPGA resources by wisely scheduling hardware modules over time according to the application needs. DPAr benefits include higher occupancy and power efficiency.

This work was supported by TU Delft, Netherlands and the Brazilian Institutions: Science without Borders/CNPq, CAPES, FAPEMIG, UFV, UFRGS, Funarpos/FUNARBE, and Gapso.

Authors' addresses: R. Ferreira (corresponding author), L. Rocha, A. G. Santos, and J. A. M. Nacif, Universidade Federal de Viçosa, Viçosa, CEP 36570.000, Brazil; email: ricardo@ufv.br; S. Wong, EEMCS, Computer Engineering P.O. Box 5031, 2600 GA Delft, The Netherlands; email: j.s.s.m.wong@tudelft.nl; L. Carro, Instituto de Informática, Universidade Federal do Rio Grande do Sul, Av. Bento Gonçalves, 9500, CEP 91501-970 Po Box: 15064, Porto Alegre, Brazil; email: carro@inf.ufrgs.br.

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.

© 2015 ACM 1936-7406/2015/03-ART9 \$15.00

DOI: <http://dx.doi.org/10.1145/2660775>

In order to implement DPaR, a series of challenges should be addressed. These problems are being solved by different techniques [Papadimitriou et al. 2011; Gericota et al. 2003; Handa and Vemuri 2004; Sidiropoulos et al. 2012]. In this article, we focus on the hardware module Placement and Routing (P&R) problem. When P&R dynamically maps a new hardware module, a nonused region should be identified in the target architecture. This region should also provide both sufficient hardware resources and minimize the impact to upcoming DPaR procedures.

The P&R is a highly computationally intense task that is traditionally performed offline. The P&R of modern complex FPGAs commonly takes several hours to be performed. The most used approaches [Ludwin and Betz 2011] to address P&R include recursive partitioning, analytical placements, genetic algorithms, and simulated annealing. These approaches are not suitable for a runtime P&R algorithm because they are very slow. Recently, a Just-in-Time (JIT) framework [Sidiropoulos et al. 2012; Hübner et al. 2011] to support DPaR was introduced. This strategy reduces the fragmentation for hardware resources (preallocated FPGA area, etc.) in partial reconfiguration context. The configuration bitstream for a virtual FPGA is computed at runtime by performing technology packing and P&R. However, the P&R still takes a few seconds to complete, hence runtime P&R remained impractical.

In this article, we propose a novel P&R based on a graph mapping model that performs nearly three orders of magnitude faster in comparison to the state-of-the-art of P&R algorithms [Luu et al. 2011]. The main contributions of this article are (1) a novel polynomial P&R greedy heuristic based on graph mapping, (2) a local routing for FPGA by exploring the graph locality and by prioritizing the critical path, (3) an adaptable P&R for fragmented regions, and (4) a runtime P&R suitable for dynamic partial reconfiguration frameworks.

The remainder of this article is outlined as follows. Section 2 presents a novel implicit graph model used for the FPGA representation. Section 3 describes our mapping approach and how we explore the graph model to produce the P&R. In Section 4, we evaluate our P&R algorithm compared to the VPR tool [Luu et al. 2011]. Section 5 presents related work. Finally, we conclude in Section 6.

2. A FPGA GRAPH MODEL

A logic circuit has a direct correspondence to a graph, where the edges represent the wires and the nodes represent the logic functions or FPGA LUTs (after the technology mapping stage). An FPGA could also be modeled as a graph, and the P&R as a graph mapping problem from a source graph (circuit) to a target graph (FPGA). In this section, a novel implicit FPGA graph model is proposed.

Although several FPGA graph models have been proposed, our implicit model introduces new concepts. First, traditional graph structures use an adjacency/incidence list or matrix to store the graph. For the implicit model, it is not required to store which $LUT_{i,j}$ is connected to $switch_{i,j}$, since the connection is implicitly based on the index numbers. Second, the model is as close as possible to the FPGA components (LUTs, wires, and switch boxes). Since there are several commercial FPGA families, the proposed graph is based on a virtual FPGA as presented in Sidiropoulos et al. [2012] and Hübner et al. [2011]. Moreover, the virtual FPGA has several advantages: It is independent of the underlying hardware, it can be directly implemented on a traditional off-the-shelf FPGA device, and it allows partial configuration as shown in Sidiropoulos et al. [2012, 2013], even if the target device does not include partial reconfiguration infrastructure. Finally, since the algorithm is designed for a runtime implementation, the routing cost and availability is efficiently verified by extending the concept of local connections.

The FPGA graph supports three node types: LUTs, wire segments, and switch boxes, which are detailed in the following subsections.

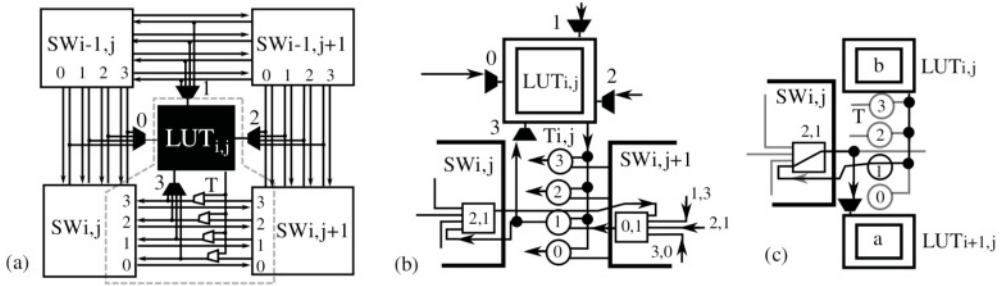


Fig. 1. (a) FPGA LUT and switches; (b) LUT output graph; and (c) local LUT connection.

2.1. LUT and Wire Nodes

Let us consider a homogeneous FPGA architecture with one four-input LUT per logic block and W channel tracks. Each LUT input is fully connected to all adjacent routing tracks, as depicted in Figure 1(a). The directions west, north, east, and south are labeled as 0, 1, 2, and 3, respectively. Each LUT will be represented as a $LUT_{i,j}$ node, where i, j represent the row and column within an FPGA, respectively. The LUT node includes four multiplexers (one for each input). Figure 1(a) depicts an example with four bidirectional channel tracks similar to the FPGA described in Hübner et al. [2011], where the LUT output is connected to the south bidirectional tracks from right to left. Each wire track multiplexer will be represented as a wire track node (T), and it will be labeled by i, j , and the wire track number. Figure 1(b) depicts four T nodes: $T_{i,j,3}$, $T_{i,j,2}$, $T_{i,j,1}$, $T_{i,j,0}$. The edges are not stored by the implicit graph.

2.2. Switch Box and the Set of Switch Nodes

Each switch box will be replaced by a set of Switch Nodes (SW), one for each output wire. The switch box is the key component in the FPGA design and routability. It is well known that the switch box needs a large number of configuration bits in comparison to track and LUT nodes. The proposed implementation is based on the Wilton switch [Wilton 1997]. All connections, through the wire tracks, should traverse at least one SW node. The in/out number for a switch box depends on the track channel width. Since each SW node implements a single switch output, there is a direct correspondence to a physical multiplexer. Each multiplexer receives four inputs tracks, one from each SW direction (west, north, ...). If there are C channel tracks, then one entire switch box will be represented by a set of $4C$ SW nodes. The label i, j, d, c identifies the SW node, where c is the wire track number and d is the direction (west, north, ...). Figure 1(b) depicts two SW nodes that are connected to the node $T_{i,j,1}$ and the $LUT_{i,j}$ input 3. The $SW_{i,j,2,1}$ is from the east side of switch box i, j , and $SW_{i,j,0,1}$ is from the west part of the switch box $i, j + 1$. The wire track $k = 1$ for the switch $SW_{i,j+1}$ in Figure 1(b) is implemented as a multiplexer with the following input tracks following the Wilton pattern [Wilton 1997]: $t_{i,j+1,0,k}$, $t_{i,j+1,1,\frac{w-k}{w}}$, $t_{i,j+1,2,k}$, and $t_{i,j+1,3,\frac{k-1}{w}}$. The SW nodes are also described by index numbers in the implicit representation.

3. GRAPH-BASED PLACEMENT AND ROUTING ALGORITHM

The proposed P&R algorithm is based on the Depth-First (DF) traversal. As already mentioned, the P&R is implemented by a graph mapping from the k -input LUT graph to the FPGA graph described in Section 2. At compile time, the synthesis tool is responsible for generating the input graph, which is represented by a k -input LUT graph. The graph is stored as an edge list in DF order. During the DF traversal, the descendant

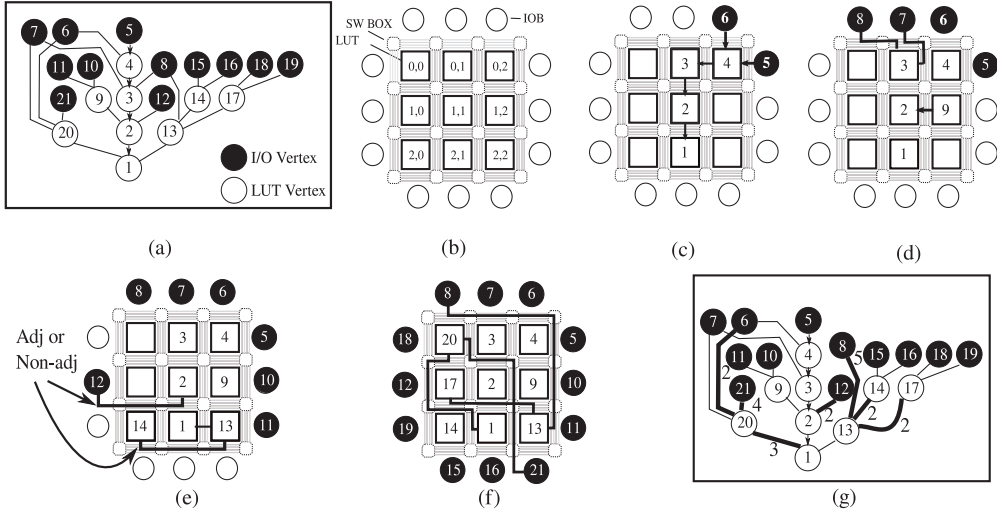


Fig. 2. (a) Input Graph; (b) FPGA; (c) first path; (d) partial P&R; (e) nonadjacent LUTs; (f) longest connections; and (g) Final cost.

vertices have been ordered by the depth the traversal can reach. Figure 2(a) depicts a simple graph, where vertex 1 has three descendants: 2, 13, and 20. The deepest the traversal can go is through vertex 2, hence it will be the first to be visited. The goal is to prioritize the critical path during the mapping.

For ease of explanation, Figure 2(b) depicts a simplified view of the FPGA graph, where T and SW nodes are not detailed, and they will be represented by wiring segments. The term *vertex* will be used for the k-input LUT graph and the term *node* for the FPGA graph. The Depth-First Placement and Routing (DFPR) algorithm maps a LUT vertex onto a LUT node and an input/output vertex onto an in/out node. The SW and T nodes are used for routing the edges. An edge from the k-input LUT graph is mapped in one or more SW and T nodes onto the target FPGA graph. The P&R aims to minimize the number of SW and T nodes by using a greedy approach.

The DFPR pseudocode is presented in Algorithm 1. Consider the DF traversal depicted in Figure 2, where the vertex numbers follow a DF order driven by the critical path. First, the edges in the path from 1 to 5 are visited and mapped as depicted in Figure 2(c). The LUTs are labeled by the row and column indexes. It is important to note that the FPGA graph is also traversed in DF order. Since all LUTs were free, the critical path was mapped in an optimal way by using local adjacent nodes. It is also important to remark that, in addition to the placement, all edges had also been routed with optimal cost.

The next edge is $6 \rightarrow 4$, since the vertex 4 is placed at $LUT_{0,2}$, then DFPR maps the vertex 6 into an adjacent I/O. Next, the edges $7 \rightarrow 3$ and $8 \rightarrow 3$ are also mapped with optimal cost. Figure 2(d) presents a snapshot in which the DFPR maps the edge $9 \rightarrow 2$ by placing the vertex 9 at $LUT_{1,2}$, adjacent to node $LUT_{1,1}$ where vertex 2 was placed.

The edge $13 \rightarrow 1$ is also placed and routed with minimal cost. However, vertex 14 could not be placed in an adjacent position next to vertex 13, as depicted in Figure 2(e), since vertex 13 is placed at the bottom-right corner, and there is no free node among its eight neighbors. Therefore, a nonadjacent LUT is used. As introduced in the following section, the $LUT_{2,0}$ could be considered adjacent to the $LUT_{2,2}$, even when the routing has a cost of two switch boxes.

ALGORITHM 1: Depth-First Placement and Routing (DFPR) Algorithm

Input: *Edges* = an edge list in depth-first order driven by the critical path
Output: Mapped FPGA graph

for each edge $e : b \rightarrow a$ in *Edges* **do**
 // Place and Routing First Fanout Edge ;
 if b is not Placed and there is a free LUT $L \in \text{Adj}(a)$ **then**
 Place b in L , and route $L \rightarrow \text{Lut}(a)$;
 else if b is not Placed **then**
 Place and Route b in the nearest free neighbour of a by doing breath-first search
 else
 Insert a in Fanout List of b
 end
end
 // Routing remaining edges ;
for each node x in Fanout List **do**
 Multicast Routing of all x 's Fanout
end

Figure 2(f) depicts the final mapping, where the longest connections are highlighted. Most edges are mapped with an optimal cost. Figure 2(g) depicts the original graph with cost labels. The bold lines are used to show the connections for which the cost is greater than 1. Normal lines have cost 1. Finally, for this example, the average cost per edge is 1.56 switch boxes per connection or mapped edge, which is a promising result for a greedy algorithm to solve an NP-complete problem based on a simple and single graph traversal. The optimal solution has an average cost per edge equal to 1.39.

The time complexity is polynomial because each edge is visited only once. Previous work on CGRA has already used a graph mapping approach based on DP traversal [Ferreira et al. 2007]. However, a circuit at LUT level has vertices with high fanout/fanin degrees. Moreover, the routing infrastructure, as well as the lower FPGA granularity at bit level, results in more complex mapping and modeling in comparison to CGRAs. The following sections describe the concept of adjacent and nonadjacent nodes and the multicast routing of the remaining edges.

3.1. Adjacent LUT Nodes

The DFPR algorithm is based on DF traversal in both graphs. Therefore, in an optimal scenario, two adjacent vertices $b \rightarrow a$ should be mapped in two adjacent LUT nodes $\text{LUT}_b \rightarrow \text{LUT}_a$. However, this assumption is not valid for all vertices since there are physical P&R constraints in the FPGA graph. Therefore, the vertex b could be placed far away from the vertex a because there are no available LUTs in vertex a 's neighborhood.

Initially, let us consider an edge $b \rightarrow a$. At least one track and one SW should be traversed in the FPGA graph, as depicted in Figure 1(c). This mapping example is the best local situation for an edge. However, it could be impossible even for adjacent row and/or column LUT. Therefore, we expand the concept of FPGA adjacent nodes. Figure 3(a) depicts the cost to connect $b \rightarrow a$ if b is placed on the top of a . The cost depends on both the position of b and the input border of a . The output will traverse 2, 1, 3, and 2 SWs to connect to the a inputs: west, north, east, and south, respectively.

Figure 3(b) depicts the cost if b is placed on the left of a , and Figure 3(c) depicts both the left and the top possibilities and costs. Instead of depicting the cost at the target node a inputs, we present the costs at the source node borders.

Assume that only the top and the left LUTs are free. The lowest cost option is to place b at the top LUT and to connect it by using the north input of a . If the north input

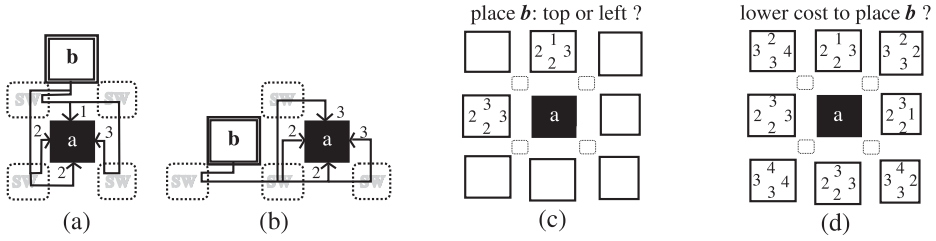


Fig. 3. (a) Top LUT routing cost; (b) left LUT; (c) top or left LUT; and (d) 8-neighborhood.

is unavailable, there are four options with cost 2: *top* \rightarrow *west*, *top* \rightarrow *south*, *left* \rightarrow *west*, and *left* \rightarrow *south*. If both west and south inputs are unavailable, the routing cost 3 will be checked, and there are two options: the top or the left LUT. Therefore, we propose to extend the concept of adjacent node to routing costs greater than 1. Figure 3(d) depicts the routing costs for an 8-neighborhood. If the routing cost 3 is considered, most nodes in Figure 2(f) are placed and routed as adjacent nodes, even the edges $14 \rightarrow 13$, $17 \rightarrow 13$, and $20 \rightarrow 1$, and only the I/O nodes 21 and 8 have a routing cost greater than 3. Although the example depicted in Figure 2 is simple, the target FPGA has the minimum square size, and all I/Os and LUTs are occupied.

3.2. Nonadjacent LUT Nodes

The example in Figure 2 is composed only by nine vertices and 12 input vertices, but Figure 2(f) depicts two cases where nonadjacent nodes should be routed. Let us consider now the input vertices. Some input vertices could be handled as adjacent nodes, such as vertices 5 and 6. However, edge $21 \rightarrow 20$ should be routed to a nonadjacent node since there is no place for the source vertex in the target node neighborhood. When the depth-limited depth-first search in adjacent nodes fails, we find the closest node by using an unlimited breadth-first search around $L_{i,j}$ in the FPGA graph. Subsequently, the source vertex is placed, and the routing algorithm is performed. Another situation occurs when the source node is already placed but it has multicast edges such as input vertices 6, 7, and 8 in Figure 2. Since these vertices have been already placed as adjacent nodes of 4 and 3, when the edge $6 \rightarrow 20$ is visited, per the second time, the multicast routing algorithm handles this connection.

3.3. Multicast Routing of Fanout Larger than One

The DFPR algorithm uses the concept of locality based on the first edge traversal. The algorithm starts from the outputs to the inputs. If a vertex is composed of more than one fanout, the vertex position is defined by the first visited edge. The remaining edges are inserted in a fanout list, as depicted in DFPR pseudocode (see Algorithm 1). After all edges have been visited, the multicast edges will be routed.

For multicast edges, we propose using a simple and greedy routing algorithm. Our algorithm is based on Network-on-Chip (NoC) XY routing [Dehyadgari et al. 2005; Lin et al. 1994]. The routing path starts from the source node, moving through switch boxes until target node at row i and column j is reached. We adapted the algorithm to an FPGA architecture considering switch box track channels and LUT side inputs.

Figure 4(a) depicts a simple source-to-target routing example. First, the algorithm tries to route in the row direction. However, since the output track in the row direction is already used, the algorithm routes in the column direction. The next two steps are accomplished in the row direction, and, finally, the routing reaches the target at the north input.

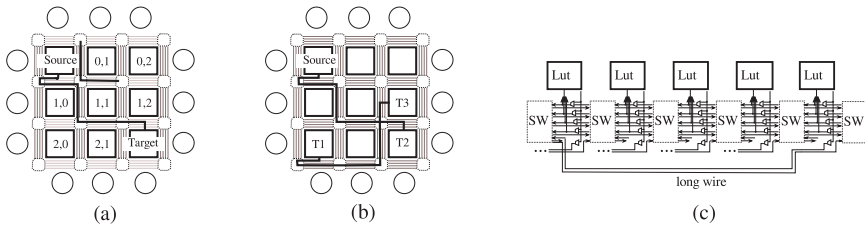


Fig. 4. (a) Row/Column routing; (b) multicast routing; and (c) one long wire.

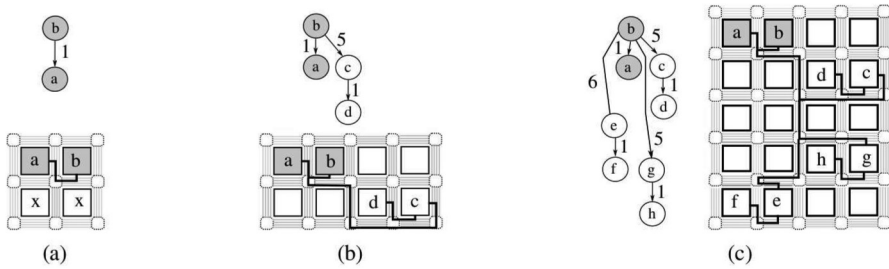


Fig. 5. (a) Fanout one; (b) fanout two; and (c) multiple fanout.

Figure 4(b) depicts a multicast routing to three target destinations: T_1 , T_2 , and T_3 . Multicast routing strategy reduces the overall switch cost by sharing routing resources. When using the multicast strategy, the routing should observe track directions. For example, the south track is connected from right to left. Therefore, since the T_1 input is in the south side, our algorithm uses one extra switch box to perform the routing.

The DFPR algorithm is based on DF order edge visit. When a source vertex is visited more than two times, the algorithm does not route its edges on the fly. A list of target vertices is created during the traversal for each multiple fanout source node. Finally, when all edges have been visited, all multicast edge lists are processed. In most circuit functions, control signals are connected to several destinations as a multicast signal. The routing cost of these signals strongly contributes to the total routing cost.

3.4. Fanout Locality and First Edge

Figure 5(a) depicts an edge $b \rightarrow a$ during the placement process. Since a is already visited and placed, the position of b will be adjacent to a if there is a free node. If there is no available node, the algorithm uses a nonadjacent node. However, if node b is composed by a double fanout, as depicted in Figure 5(b), the greedy approach used by the DFPR algorithm does not ensure that c will be placed close to b , even if there is an optimal position for b near a . The position of c is chosen based on the locality of its fanout. For this example, c is placed close to d . The routing path from b to c will traverse five switch boxes, and it uses six wire segments. Therefore, although the cost of $b \rightarrow a$ is 1, the cost of $b \rightarrow c$ is 5. The DFPR algorithm focuses only on the first visited edge. It is important to note that the first edge belongs to the original critical path. Since most digital circuits include a large number of single fanout nodes, the heuristic can significantly reduce the P&R execution time without increasing the critical path.

Now, we consider a simple example of the vertex with fanout 4, depicted in Figure 5(c). Vertex b includes the following fanout vertices: a , c , e , and g . The first edge $b \rightarrow a$ is used to place b , and g , e , and c are placed near their fanouts h , f , and d , respectively. Since b is a multicast node, as explained in Section 3.3, the multicast routing is performed after the final placement. A possible multicast routing is depicted in Figure 5(c), where

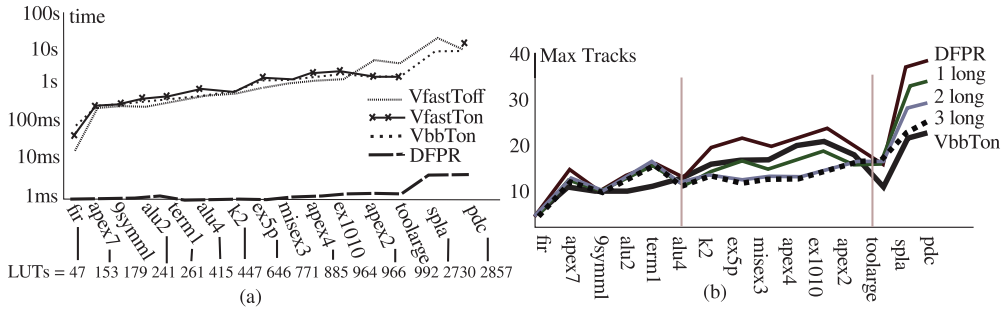


Fig. 6. (a) VPR and DFPR execution time in log scale; and (b) maximum number of tracks and long wires.

the mapped edges $b \rightarrow c$, $b \rightarrow e$, and $b \rightarrow g$ present a cost of 5, 6, and 5 switch boxes, respectively. Assuming that $b \rightarrow g$ is not in the original critical path, since its cost after the P&R is 5, this path could become the critical path. The total number of used switch boxes and wiring segments is computed by considering that some resources are shared by multicast connections. The four fanout mapped edges of b require 9 switch boxes and 13 wiring segments in total.

4. EXPERIMENTAL RESULTS

The DFPR algorithm is compared to VPR [Luu et al. 2011] by using the combinational MCNC benchmarks [MCNC 2010]. The target FPGA architecture has one 4-input LUT per logic block. Each routing channel contains 50 segment tracks, all tracks are fully connected to each adjacent LUT input, and all wires have one unit length (two adjacent switch boxes). Moreover, the LUT output is always connected to the south border, similar to the virtual FPGA proposed in [Sidiropoulos et al. 2012]. Our virtual FPGA is composed of an array of $N \times N$ LUTs to map a circuit with N^2 LUTs, which corresponds to the minimum square. We present and analyze the DFPR execution time, the maximum number of tracks, and the total wire length in Sections 4.1 and 4.2. Section 4.3 compares the greedy DF placement to three simulated annealing-based placements [Luu et al. 2011]. Section 4.4 presents an edge distribution analysis of the proposed DF placement. Finally, we evaluate the fragmentation in a partial reconfigurable scenario in Section 4.5.

4.1. Execution Time

The experiments were performed on an Intel i3 370M, 2.4GHz, 3MB L2 cache machine. The execution time is measured by using GNU gprof tools. The VPR is executed in three different modes: *VfastToff*, *VfastTon*, and *VbbTon*. *VfastToff* and *VfastTon* use the fast-mode based on user guide recommendations [Luu et al. 2011] by setting the parameter: *fast on*. *VfastToff* uses an additional parameter *timinganalysis off* to switch off the timing optimization, and *VfastTon* switches it on. We also propose another fast mode *VbbTon* with the following parameters: *innernum 1* *placedalgorithm boundingbox* *timinganalysis on*. In addition, we fix the channel length by using *routechainwidth 50*. Figure 6(a) depicts the execution time for considering a set of 15 MCNC benchmarks [MCNC 2010] ordered by the LUT size (depicted in the bottom line). The DFPR algorithm is on average $1950\times$ faster than VPR in *VbbTon* mode. *VfastToff* is the faster VPR mode for the small benchmarks; however, *VfastTon* and *VbbTon* performs better for the four largest benchmarks. Moreover, the *VfastTon* routing fails for *spla*. The DFPR is a greedy heuristic, and each LUT edge is only visited once in comparison to Simulated Annealing (SA)-based approaches that perform several tries. Therefore, our heuristic reduces the execution time by 2–3 orders of magnitude. For the evaluated

Table I. Placement Comparison to Baseline V_{CR} Placement: Execution Time, Wiring, and Critical Path

| Placement | V_{CR} baseline | | |
|------------------|--------------------------|-------------------|-----------------|
| | Executing Time (Speedup) | Critical Path (%) | Wire Length (%) |
| DF | 10218× | 26.4 % | 43.8% |
| $V_{bb}T_{on}$ | 14.4× | 27.0 % | 0.1% |
| $V_{fast}T_{on}$ | 8.4× | 1.8 % | 4.8 % |

benchmarks, the VPR performs from 10^5 up to 10^6 swap operations. The impact of swap operations, the wire length, and the critical path are analyzed in the following sections.

4.2. Maximum Track Number

We use 50 tracks in our experiment as proof of concept to evaluate the maximal track usage in a scenario with enough routing resources for both the VPR and DFPR routing. Although the DFPR could generate a routable solution, the maximal track occupation increases as a function of circuit size, as depicted in Figure 6(b). In commercial FPGAs, long wires are used to speed up connections to avoid channel congestion and large width tracks. Long wires are similar to single wires, except that each one spans two or more Switch Boxes (SW), imposing lower routing delays.

Figure 6(b) depicts the maximal track usage in presence of long wires that span 5 SW by considering the use of 1, 2, or 3 long wire tracks per channel. Figure 4(c) depicts an example of adding one long wire. There are three regions. The first region shows the small size benchmarks, which have less than 500 LUTs (Fir \rightarrow Alu4); the DFPR algorithm requires on average 17% more tracks than VPR $V_{bb}T_{on}$. The worst case requires 15 tracks, which is small compared to 50 available tracks. The middle region in Figure 6(b) depicts the medium-sized benchmarks from 500 up to 1,000 LUTs or 2,000–4,000 equivalent gates, where the DFPR algorithm uses on average 23.4% more tracks, where the maximum usage is 23 tracks for ex1010 benchmark. By adding one long wire, the DFPR presents a track usage equivalent to VPR. Finally, for the last two and large benchmarks (3,000 LUTs), the DFPR uses 76% more tracks than VPR. However by adding three long wires, the DFPR usage gets close to VPR usage. Therefore, using long wires is an effective technique to reduce the amount of routing resources with no P&R execution time degradation.

In addition to the track usage, our proposed DFPR P&R requires on average $3.27 \times$ more wiring segments than $V_{bb}T_{on}$. It is well known that the final routing quality depends on the placement choices. Therefore, in the following section, we analyze separately our greedy placement algorithm and its required resources.

4.3. Placement Evaluation

This section evaluates the tradeoff between execution time and placement quality. Our single-try DF placement is compared to three VPR SA placements [Luu et al. 2011]: $V_{bb}T_{on}$, $V_{fast}T_{on}$, and V_{CR} . The placements $V_{bb}T_{on}$ and $V_{fast}T_{on}$ are described in Section 4.1. In order to measure only the placement execution time, the *placeonly* VPR option is set. In addition, following the VPR reference manual [Luu et al. 2011], execution time/quality tradeoff could be explored by changing the *innernum* option. Therefore, we added the V_{CR} placement that optimizes the quality by using the default parameter: *-innernum 10*. Figure 7(a) depicts the placement execution time in log scale, and Table I summarizes the results. V_{CR} reaches the best results and therefore it is used as the baseline comparison. However, the V_{CR} disadvantage is the long execution time. Our DF placement is on average 10,218×

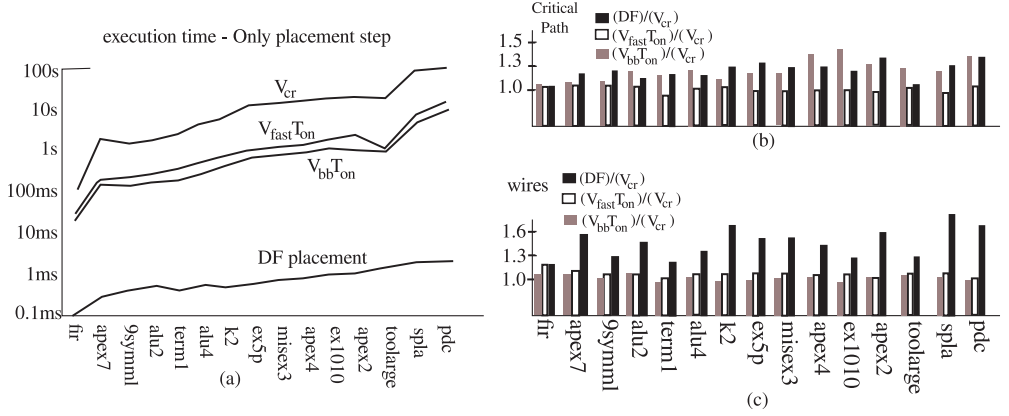


Fig. 7. (a) Placement execution time; (b) normalized critical path; and (c) normalized wire length.

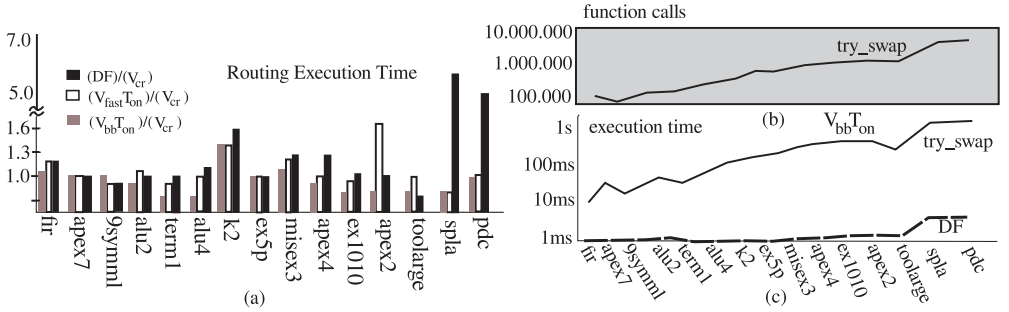


Fig. 8. (a) Normalized VPR routing execution time; (b) VPR “try_swap” function calls; and (c) VPR “try_swap” execution time.

To evaluate the critical path and the wire length, we propose to apply the same routing algorithm for all placements, including our DF algorithm. The VPR routing is set to the default parameters and 50 tracks. Regarding the critical path, the $V_{fast}T_{on}$ reaches a near-optimal solution in comparison to V_{CR} , and it is on average only 1.8% worse than V_{CR} . Figure 7(b) depicts the normalized critical path in comparison to the baseline V_{CR} . A value greater than one denotes a critical path increase. On average, the DF and $V_{bb}T_{on}$ are equivalent, and both slow down the critical path by a factor of 27%. However, the DF execution time is on average $710\times$ and $1,218\times$ faster than $V_{bb}T_{on}$ and $V_{fast}T_{on}$, respectively.

Regarding the wire length, Figure 7(c) depicts the normalized wire length in comparison to V_{CR} . $V_{bb}T_{on}$ and V_{CR} have almost the same wire length usage, whereas $V_{fast}T_{on}$ is about 4.8% worse with respect to the wire length. Our DF placement increases the wire length on average in 43.8%, since it is a single-try greedy heuristic that prioritizes the critical path placement. In order to provide better results, the VPR placements try multiple positions, increasing the execution time by at least $710\times$.

Figure 8(a) depicts the normalized execution time for the VPR routing algorithm for all placements. For the *pdc* and *spla* benchmarks, as detailed in the next section, the VPR routing requires more computational efforts for our DF placement since it only explores the locality of single fanout nodes.

Figure 8(b) and (c) depicts the number of calls and the execution time of VPR “try_swap” function in comparison to DF algorithm. The “try_swap” is the most

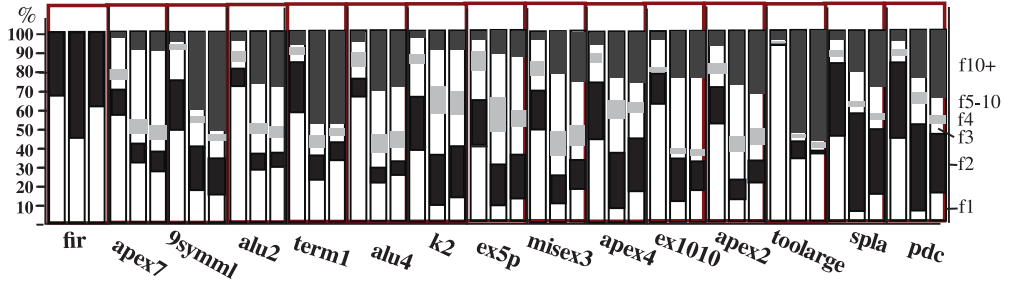


Fig. 9. Original fanout nodes and mapped edge distribution for DF and V_{CR} placements.

time-consuming function in SA placements, as observed in Sidiropoulos et al. [2013]. We measure the “try_swap” for the $V_{bb}T_{on}$, which is the fastest SA placement. This function is called 115,097 times for *9symml* and 5,243,379 times for *pdc*, in comparison to our DF placement that tries just one position per node.

In summary, the results suggest that $V_{bb}T_{on}$ is the best SA placement regarding the wire length and execution time tradeoffs, $V_{fast}T_{on}$ is the best SA placement regarding the critical path and execution time tradeoffs, and our proposed DF substantially reduces the execution time with acceptable quality for a single try runtime placement. Finally, DF placement results are very promising, although there is still room for improvement in the placement quality tradeoffs. The following section discusses the first steps to understand the differences between DF and SA placements by investigating the mapped edges.

4.4. Mapped Edge Analysis

Since placement quality directly impacts routing cost, this section analyzes the wire occupancy for the mapped edges as a function of the node fanout to better understand the results of DF and V_{CR} placements. Suppose f_i to be a class of nodes or vertices for which the fanout is i . Suppose $e_{first} = b \rightarrow a$ to be the first visited edge of b . As already mentioned, the DFPR performs a simultaneous traversal in both original and FPGA graphs. The placement of b is defined when e_{first} is visited during the DF traversal. Moreover, b is placed as near as possible to a to optimize the cost of the edge e_{first} . For instance, suppose x has fanout 2: $x \rightarrow y$, and $x \rightarrow z$. If only y is placed near x , the edge $x \rightarrow y$ could be mapped by using one switch, whereas the edge $x \rightarrow z$ uses five switches. On average, the mapped edges of x uses three switches.

The DF placement prioritizes the f_1 nodes. Figure 9 depicts a three-bar graph for each benchmark. The first bar depicts the vertex distribution in the original graph as a function of f_i . For instance, the *fir* benchmark is composed by 68% of f_1 vertex and 32% of f_2 vertex. The second bar depicts the mapped occupation for f_1 and f_2 after the DF placement followed by the VPR routing. The f_1 nodes use 45% of the mapped wires, and f_2 nodes use 55%. The third bar depicts the f_i mapped distribution for the V_{CR} , where the f_1 class occupies 63% of the total wires, and the f_2 uses 37%.

Although a digital circuit has on average a fanout 3, the distribution is not uniform. The majority of nodes are f_1 or f_2 . For example, the benchmarks *apex7*, *9symml*, and *alu2* have $f_1 + f_2 = 71\%$, 74% , and 81% , respectively. Figure 9 depicts the f_i in order with the following colors: f_1 = white, f_2 = black, f_3 = white, f_4 = gray, f_{5-10} = white, f_{10+} = black. The class f_{5-10} represents f_5, f_6, \dots, f_{10} , and the class f_{10+} represents all f_i such that $i > 10$. The total wire cost is in general dominated by the high fanout vertices, even though they correspond to a small fraction in the original graph. Considering the *term1* benchmark, although the high fanout f_{10+} corresponds

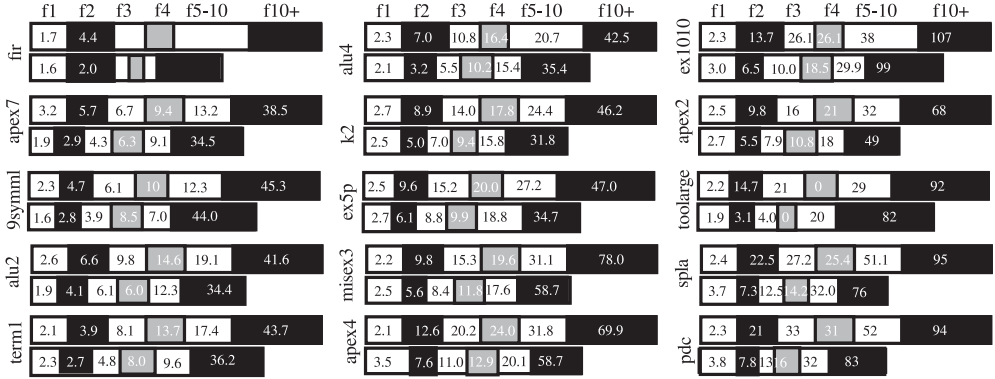


Fig. 10. Average switch utilization as a function of f_i classes for DF and V_{CR} placements.

to 6.5% of the input vertices, the mapped edges will consume 51% of the total wires for both our DF placement and the V_{CR} placement.

It is possible to verify that the DF algorithm prioritizes f_1 vertices. Considering the *spla* benchmark, f_1 represents 44.1% of all vertices, and it consumes only 5.8% of the total wires of *spla*. However, f_2 consumes a significant amount of wires resources for large benchmarks. For instance, f_2 represents 39.9% of all vertices and will consume 48.2% of the total wires. Therefore, f_2 contributes significantly to DF wire length. Although for the V_{CR} placement, f_2 is 27.8% of the total wires, f_1 will consume 15.2% of the total wires of *spla*. Since the DF placement is on average 10,000 faster than V_{CR} placement, novel approaches derived from DF could be proposed by exploring f_2 and f_{10+} edges/vertices to reduce the wire length and/or critical path at the cost of more execution time based on the locality of these fanout classes.

In addition to the f_i wire occupancy depicted in Figure 9, Figure 10 shows the average number of switches per class. For instance, the f_1 of *term1* consumes on average 2.1 and 2.3 switches per mapped edge for the DF and the V_{CR} placements, respectively. For f_2 , the average is 3.9 and 2.7, which means $\frac{3.9}{2} = 1.95$ and $\frac{2.7}{2} = 1.35$ switches per mapped edge. For the large benchmark *pdc*, DFPR optimizes the f_1 and uses 2.3 switches on average in comparison to 3.8 used by the V_{CR} . However, f_2 is mapped on 21 switches on average by the DF placement. Since the e_{first} is probably mapped by using two switches, the second edge will consume on average 19 switches, which increases the wire length. In addition, edges that were not in the critical path could contribute to the final critical path. Despite that, as already mentioned, the DF placement achieves very low execution time, around milliseconds, and local strategies could optimize the resource utilization without increasing the execution time.

4.5. Fragmentation

Consider a function set $W = A, B, \dots$ to be used during the execution of one reconfigurable application. If the bitstream for each function w_i is generated at compile time, the FPGA could have a poor resource utilization. Moreover, it could be impossible to place all functions, even if there is enough routing and LUT resources. Figure 11(a) depicts an example where function D cannot be placed due to fragmentation. Despite the two empty areas (white parts) containing enough resources, it is not possible to place D since it is generated as a single rectangular region.

To evaluate the potential of the DF algorithm on fragmented areas, consider two scenarios as depicted in Figure 11(b). Use the black color for the area that has been already allocated. In the first scenario (top), we have two occupied disjoint areas,

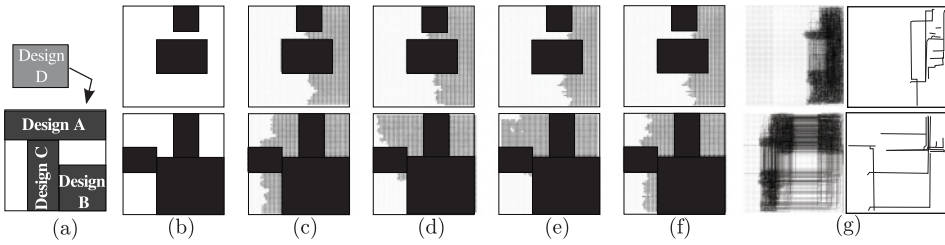


Fig. 11. (a) Fragmentation, (b) initial scenario, (c) Apex2, (d) Apex4, (e) Too_large, (f) Ex1010, (g) routing and critical path.

Table II. Wiring and Critical Path for Fragmented Regions

| Bench | Scenario 1 | | | Scenario 2 | | |
|-----------|------------|---------------|-------|------------|---------------|-------|
| | Wires | Critical Path | Ratio | Wires | Critical Path | Ratio |
| too_large | 7846 | 3.08 | 1.04 | 7328 | 3.52 | 1.19 |
| apex2 | 13328 | 1.67 | 1.43 | 15151 | 1.72 | 1.48 |
| apex4 | 14002 | 1.62 | 1.42 | 17426 | 1.77 | 1.55 |
| ex1010 | 13922 | 1.91 | 1.73 | 17713 | 1.86 | 1.69 |
| average | | | 1.41 | | | 1.47 |

whereas in the second scenario (bottom), most of the area has been already allocated. There are three disjoint free areas in the second scenario. For both scenarios, four MCNC benchmark circuits have been mapped with around 1,000 LUTs onto a target FPGA array of 60×60 LUTs with channel width 50. Figure 11(c) depicts the area where the function *Apex2* is placed (in gray). Figures 11(d), 11(e), and 11(f) present the area for *Apex4*, *Too_large*, and *Ex1010*.

The DF algorithm executes the placement in less than 2 milliseconds, even in fragmented areas. Since it is a graph-based approach, it is easily adaptable to fragmented regions. The DF output is routed by using VPR to obtain the critical path. Figure 11(g) depicts the VPR [Luu et al. 2011] screenshot for the *Ex1010* benchmark routing and critical path. Most of wiring segments are concentrated on the placement region. Moreover, wiring segments are also used to route through disjoint areas. Table II shows the critical path and total wiring segments for both scenarios. The column Ratio depicts critical path length increases compared with optimal critical path. The optimal critical path is obtained by using VPR at timing driven mode, which is mapped on a square contiguous area, as shown in Section 4.3. On average, the critical path increases 41% and 47% for scenarios 1 and 2, respectively. On average, the routing resource measured as wiring segments increases 57% and 79% for scenarios 1 and 2, respectively.

5. RELATED WORK

It is well known that FPGA P&R are NP-complete problems [Donath 1980], and these problems are handled separately and solved sequentially one after the other. Due to their complexity, several heuristics have been proposed in the past three decades.

The placement heuristics could be classified into three categories: SA, partitioning, and analytical. The first and most popular approach is based on SA [Luu et al. 2011; Betz and Rose 1997; Lin and Wawrzynek 2010; Ludwin and Betz 2011]. The SA-based placement reaches a high-quality solution regarding the critical path and/or the total wire length. Nevertheless, the SA approach has a long execution time. Several SA-based

placements [Ludwin and Betz 2011; Lin and Wawrzynek 2010; Wu and McElvain 2012; Sidiropoulos et al. 2013] have been proposed, and most of them use the VPR framework [Luu et al. 2011; Betz and Rose 1997] to perform a baseline comparison. In Lin and Wawrzynek [2010], a dynamically adaptive stochastic tunneling algorithm to avoid the freezing problem in the SA approach is proposed. However, only a marginal reduction of 18.3% in runtime is obtained over VPR. A parallel approach for SA has been proposed in Ludwin and Betz [2011]. The algorithm evaluates multiple SA moves in parallel, and the experimental results achieve an average speedup ranging from $2\times$ up to $7\times$ in comparison to VPR. In Wu and McElvain [2012], a modified SA algorithm has been presented by using low temperatures, which results in a speedup factor of $3\times$ in comparison to a traditional SA algorithm with random movements without critical path and wire length degradation.

The second placement approach is based on partitioning techniques [Maidee et al. 2005]. First, the circuit is partitioned, and then the overlap regions are removed; finally, the SA algorithm is used to perform the local placement. Although the execution time is reduced, the quality of results is worse regarding the critical path and total wire length.

The third category is analytical placement [Xu et al. 2011; Lin et al. 2013]. Today, some commercial tools have replaced the SA based approach with analytical approaches [Lin et al. 2013], which is a scalable technique to handle high-capacity devices. An analytical approach based on a near-linear net model was proposed in Xu et al. [2011] that is $5\times$ faster than VPR fast mode. Moreover, this placement obtains a 9% reduction in critical path delay. The routing is performed by using VPR router. The reported CPU time is in the range of seconds for the MCNC benchmarks. Recently, an analytical placement proposed in Lin et al. [2013] achieves a speedup factor of $7\times$ compared to VPR, with an average reduction of 7% in the critical path without total wire length degradation.

The analytical and partitioning-based approaches perform two placements: a global placement of large blocks followed by a detailed placement. The global placement could generate overlapped regions that will be removed using a legalization step. Then, an SA-based detailed placement is done. The DFPR algorithm differs from the previous one, and it can be included in none of the three placement categories. Moreover, the SA technique is not used, and the placement is performed in a single step.

Regarding the routing step, a new approach is presented in Gort and Anderson [2011]. When combined with a low-cost architecture change, this new approach results in a 34% reduction in router runtime, at the cost of a 3% area overhead. The routing time is around 1 second for the MCNC benchmarks.

All previous works are offline approaches at compile time. Moreover, recent work targets scalable approaches for large circuits, since the P&R times of high-density FPGAs (more than 100K LUTs) can easily reach a full day. Unfortunately, the SA approach is not suitable for P&R of this magnitude, and the analytical/partitioned approaches are used. The DFPR algorithm proposed here focuses on partial reconfiguration frameworks, where it should dynamically replace only a part of a large circuit. However, even for medium-sized circuits like MCNC Benchmark, the previous work execution time, in seconds, is not suitable for runtime.

Regarding runtime P&R, a similar work [Sidiropoulos et al. 2012, 2013] proposes a JIT P&R that is $7.34\times$ faster than VPR. Moreover, the DFPR algorithm and the JIT approach lead to significantly lower fragmentation of hardware resources at the LUT level. However, whereas JIT P&R time is around seconds, the DFPR algorithm executes in a few milliseconds, thanks to its data structure and simplicity, thus allowing its usage in runtime applications.

6. CONCLUSIONS AND FUTURE WORK

A novel P&R approach is proposed in this article, one that has demonstrated that a low-complexity P&R algorithm based on a graph traversal can achieve a huge reduction in execution time in comparison to traditional SA-based approaches. The algorithm is well-suited to be used at runtime in a partial reconfiguration framework, whereas traditional P&R approaches are not suitable for runtime because they suffer from high overheads. Furthermore, the P&R is modeled as a graph mapping problem. The input graph is the circuit to be mapped, and the output graph is the FPGA. The P&R algorithm performs a DF graph traversal over the input graph, resulting in the final mapping. Unlike traditional offline algorithms, our approach does not try similar solutions, allowing the P&R to execute in milliseconds. We presented results that validate the proposed solution, reaching up to three orders of magnitude speedup compared with the state-of-art VPR tool [Luu et al. 2011]. Beyond the algorithm, this work distills several contributions for implementing P&R approaches: (1) To the best of our knowledge, the DFPR algorithm is the first single-try placement approach that produces routable solutions for the MCNC benchmarks. (2) Instead of using traditional graph data structures, we propose an implicit representation based on index numbers. (3) We developed a new placement algorithm based on first-edge locality. (4) We expand the adjacent node zone. (5) We show a fanout distribution analysis and its impact on the final routing cost.

This work demonstrates that a greedy P&R strategy significantly reduces the execution time compared to the state-of-the-art of P&R on homogeneous virtual FPGA with a single LUT CLB. However, modern FPGA devices include heterogeneous resources (multipliers and BRAM) as well as complex CLB with multiple LUTs and customized interconnections. Future work will address new challenges to extend the proposed graph traversal approach to these heterogeneous FPGA devices.

REFERENCES

- V. Betz and J. Rose. 1997. VPR: A new packing, placement and routing tool for FPGA research. In *International Conference on Field Programmable Logic and Applications (FPL'97)*. Springer-Verlag, Berlin, 213–222.
- M. Dehyadgari, M. Nickray, A. Afzali-Kusha, and Z. Navabi. 2005. Evaluation of pseudo adaptive XY routing using an object oriented model for NOC. In *International Conference on Microelectronics*. IEEE, 204–208.
- W. E. Donath. 1980. Complexity theory and design automation. In *Design Automation Conference*. ACM, New York, NY, 412–419.
- R. Ferreira, A. Garcia, T. Teixeira, and J. M. P. Cardoso. 2007. A polynomial placement algorithm for data driven coarse-grained reconfigurable architectures. In *ISVLSI*. IEEE, 61–66.
- M. G. Gericota, G. R. Alves, M. L. Silva, and J. M. Ferreira. 2003. Run-time management of logic resources on reconfigurable systems. In *Design, Automation and Test Conference (DATE'03)*. ACM/IEEE, 974–979.
- M. Gort and J. H. Anderson. 2011. Reducing FPGA router run-time through algorithm and architecture. In *International Conference on Field Programmable Logic and Applications (FPL'11)*. IEEE, 336–342.
- M. Handa and R. Vemuri. 2004. An efficient algorithm for finding empty space for online FPGA placement. In *Design Automation Conference (DAC'04)*. ACM/IEEE, 960–965.
- M. Hübner, P. Figuli, R. Girardey, D. Soudris, K. Siozios, and J. Becker. 2011. A heterogeneous multicore system on chip with run-time reconfigurable virtual FPGA architecture. In *Workshops and PhD Forum (IPDPSW)*. IEEE, 143–149.
- M. Lin and J. Wawrzynek. 2010. Improving FPGA placement with dynamically adaptive stochastic tunneling. *IEEE Transactions on CAD of Integrated Circuits and Systems* 29, 12 (2010), 1858–1869.
- T. Lin, P. Banerjee, and Y. Chang. 2013. An efficient and effective analytical placer for FPGAs. In *Design Automation Conference (DAC'13)*. ACM/IEEE, Article 10, 6 pages.
- X. Lin, P. K. McKinley, and L. M. Ni. 1994. Deadlock-free multicast wormhole routing in 2-D mesh multi-computers. *IEEE Transactions on Parallel and Distributed Systems* 5 (1994), 793–804.

- A. Ludwin and V. Betz. 2011. Efficient and deterministic parallel placement for FPGAs. *ACM Transactions on Design Automation of Electronic Systems* 16, 3, Article 22 (2011), 23 pages.
- J. Luu, I. Kuon, P. Jamieson, T. Campbell, A. Ye, W. M. Fang, K. Kent, and J. Rose. 2011. VPR 5.0: FPGA CAD and architecture exploration tools with single-driver routing, heterogeneity and process scaling. *ACM Transactions on Reconfigurable Technology and Systems* 4, 4, Article 32 (Dec. 2011), 23 pages.
- P. Maidee, C. Ababei, and K. Bazargan. 2005. Timing-driven partitioning-based placement for island style FPGAs. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 24, 3 (2005), 395–406.
- MCNC. 2010. BLIF Benchmark Suit. Retrieved from <http://cadlab.cs.ucla.edu/~kirill/>.
- K. Papadimitriou, A. Dollas, and S. Hauck. 2011. Performance of partial reconfiguration in FPGA systems: A survey and a cost model. *ACM Transactions on Reconfigurable Technology and Systems (TRETs)* 4, 4 (2011), 36.
- H. Sidiropoulos, K. Siozios, P. Figuli, D. Soudris, and M. Hubner. 2012. On supporting efficient partial reconfiguration with just-in-time compilation. In *PhD Forum (IPDPSW), IEEE*. IEEE, 328–335.
- H. Sidiropoulos, K. Siozios, P. Figuli, D. Soudris, M. Hübner, and J. Becker. 2013. JITPR: A framework for supporting fast application's implementation onto FPGAs. *ACM Transactions on Reconfigurable Technology and Systems* 6, 2, Article 7 (Aug. 2013), 12 pages.
- Steven J. E. Wilton. 1997. *Architectures and Algorithms for Field-Programmable Gate Arrays with Embedded Memory*. Ph.D. Dissertation. University of Toronto.
- Q. Wu and K. S. McElvain. 2012. A fast discrete placement algorithm for FPGAs. In *Proceedings of the ACM/SIGDA International Symposium on Field Programmable Gate Arrays (FPGA'12)*. ACM, New York, NY, 115–118.
- M. Xu, G. Grewal, and S. Areibi. 2011. Starplace: A new analytic method for FPGA placement. *Integration, the VLSI Journal* 44, 3 (2011), 192–204.

Received December 2013; revised July 2014; accepted July 2014