

Wirelength driven floorplacement for FPGA-based partial reconfigurable systems

A. Montone², M. D. Santambrogio^{1,2}, D Sciuto²

¹**Computer Science and Artificial Intelligence Laboratory**
Massachusetts Institute of Technology
santambr@mit.edu

²**Dipartimento di Elettronica e Informazione**
Politecnico di Milano
alessio.montone@dresd.org, santambr, sciuto@elet.polimi.it

Abstract—The proposed work aims at identifying groups of *Reconfigurable Functional Units* that are likely to be configured in the same chip area, identifying these areas based on resource requirements, device capabilities and wirelength. The proposed floorplacement framework, tailored for Xilinx Virtex 4 and 5 FPGAs, uses an objective function based on *external wirelength*, i.e., the estimated length of the nets connecting each *Reconfigurable Functional Unit* to the corresponding required chip *Input Output Blocks*. The proposed approach results, as also demonstrated in the experimental results section, in a shorter external wirelength (an average reduction of 50%) with respect to purely area-driven approaches and a highly increased probability of re-use of existing links (90% reduction can be obtained in the best case).

I. INTRODUCTION

Nowadays one of the most important design styles in VLSI is represented by Field Programmable Gate Arrays (FPGAs). The standard FPGA design flow starts from an RT-level description of the circuit (e.g. provided by a HDL language) and ends with a configuration file (bitstream) configuring the desired circuit on the device. Much effort has been already placed towards improving the different stages of the design flow, from logic synthesis to placement and routing. While these problems have been addressed directly by the FPGA vendors [1] as much as by academic works [2], the floorplanning automation for FPGAs is a current research topic, to face the new challenges provided by FPGAs.

FPGAs present two unique aspects with respect to traditional VLSI designs: resource heterogeneity and reconfigurability. This requires to take into consideration the resource heterogeneity, to allow each architectural module to be placed in an area region containing all the needed resources. On the other hand, reconfiguration allows the possibility to change the architecture or the application implemented on the FPGA without requiring any physical action on the device. In particular, in partial dynamic reconfigurability, an architecture may change at runtime (i.e., without having a disruption of the provided functionality) a subset of its modules in order to modify its own behavior. Reconfigurability introduces time as a new variable within the floorplanning formulation.

A floorplanner taking into account both of these aspects has to find, for each module¹, an area assignment according to target device's capabilities considering that the modules can be replaced later on due to the system needs using the reconfiguration capabilities of the target device. The term floorplacer has been first introduced in the context of traditional VLSI design [3], [4] to emphasize how recently developed algorithms for VLSI design automation face placement and floorplanning tasks concurrently. In a similar spirit, resource management for partially and dynamically reconfigurable FPGAs can benefit from this paradigm. In this paper, we propose to develop a floorplacer for such a device, which will identify groups of modules that are likely to be configured in the same rectangular area and, consequently, identify these chip areas according to the modules' requirements, device capabilities and design objectives. Particularly, in this work, we focus on optimizing the wirelength since the base floorplacement framework.

The paper is structured as follows: the target FPGA architecture is described in Section II, while Section III introduces the problem and definitions of relevant concepts. Section IV presents the proposed solution, Section V describes the related work and the experimental results are presented in Section VI. Finally, conclusions and future work are outlined in Section VII.

II. FPGA RECONFIGURATION AND TARGET ARCHITECTURE

This section introduces the reconfiguration process from the FPGA's physical point of view. Sub-section II-A presents the physical limits of the reconfiguration process, while sub-section II-B shows the design flow proposed by Xilinx for reconfigurable architectures.

A. Smallest Reconfigurable Region

The reconfiguration process described here relates to the latest generation of Xilinx FPGAs, i.e. Virtex 5 devices.

¹In this paper we use the term *module* to address an architectural component after technology mapping

According to their datasheets [5] all the devices of this family, share a common structure. The entire FPGA is made of programmable logic (commonly referred to as CLBs for this FPGA family) and periodically distributed Block RAMs (BRAMs), while all the other resources (such as DSP blocks) are placed along the vertical edges of the FPGA. User logic is implemented combining all these resources and connecting them using channels and switchboxes. The information about device configuration is described in a binary configuration file named *bitstream*: logical functions implemented by CLBs (i.e., the content of the Look-up tables implementing a logical function with 6 inputs and 1 output), BRAMs content, routing information (i.e., logical status of switches managing the interconnections), and so on. Such devices are divided in 4 rows having the same height. The height of the row expressed in CLBs is a device parameter. The smallest reconfigurable element is 1 row high and 1 CLB wide and is referred to as *frame*, each frame is addressed by a row number (from 0 to 3) and column number (expressed in CLBs). The bitstream follows the device topology, hence inside a configuration bitstream one can easily identify data configuring a specific frame.

As previously mentioned, the frame is the smallest FPGA area that can be configured independently of the others. Given a generic module h_{module} rows high and w_{module} CLBs wide, the smallest area that can be involved in the reconfiguration process is a rectangle with $h_{\text{small-area}}$ rows height and $w_{\text{small-area}}$ CLBs width such that:

- The height in rows is the smallest integer greater than the module's height in rows

$$h_{\text{small-area}} = \lceil h_{\text{module}} \rceil \quad [\text{rows}]$$

- The width in CLBs is the smallest integer greater than the module's width in CLBs

$$w_{\text{small-area}} = \lceil w_{\text{module}} \rceil \quad [\text{CLBs}]$$

For example, given a module requiring for its placement 1.5 rows height and 30 CLBs width, the smallest area that can be assigned to this module, considering the reconfiguration constraints, is a rectangle of height and width of 2 rows and 30 CLBs, respectively. Configuration and reconfiguration processes take place by writing the bitstream inside the FPGA's *configuration memory*. In the case of partial reconfiguration the bitstream carries information addressing the frames that are going to be replaced by the carried data.

B. Xilinx Early Access Partial Reconfiguration Design Flow

Due to the fact that Xilinx FPGA Virtex families are the target devices of this work, the Xilinx Early Access Partial Reconfiguration (EAPR) design flow [6] will be briefly introduced here. This flow is compatible with dynamic reconfiguration. Let a *Reconfigurable Functional Unit* (RFU) be a technologically mapped netlist implementing a required functionality and a *Reconfigurable Region* (RR) be a rectangular FPGA area where two or more RFUs are going to be placed and routed (at design time) and configured (at runtime) according to the application implemented in the reconfigurable system. EAPR

allows the definition of a set of non-overlapping RRs. All the *static logic* (i.e. all logic that will always remain configured, including the glue logic) is placed outside the reconfigurable regions. EAPR allows a more scalable communication infrastructure for partial dynamic reconfigurable architectures, since it allows a RFU to communicate with static logic regardless of the position of other RFUs. Consequently, the static logic can provide a communication infrastructure accessible through a number of interface at least equal to the number of RRs. Each RFU can communicate with other configured modules or with the static part by just hooking up to the hardware macros corresponding to the communication infrastructure interfaces. RRs are allowed to use routing resources intersecting and even crossing RRs (the use of routing resources crossing RRs is the most relevant case). Figure 1 (a) provides an example of a reconfigurable architecture developed with the EAPR design flow. Note how the reconfigurable regions are aligned to the grid defined by rows and CLBs according to EAPR requirements.

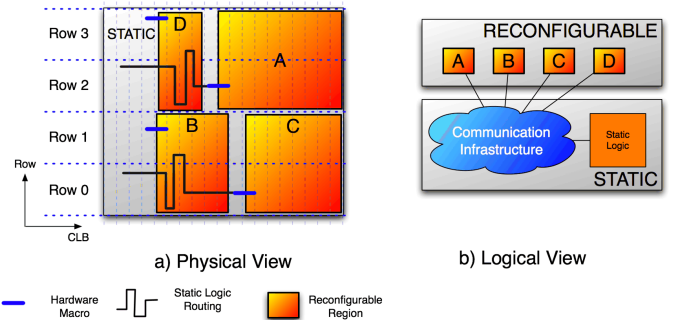


Figure 1. Target Reconfigurable Architecture based on EAPR: (a) physical and (b) logical views

According to the EAPR flow, hardware macros can be placed on the boundaries of RRs in order to define *pins* where RFUs can hook themselves. Such macros are made with pairs of CLBs, one side of the CLB pair is connected to a RR signal while the other is connected to a static logic signal. Previous design flows [7] required static logic *placed* and also *routed* outside RRs, while the EAPR relaxes this constraint.

C. Target Architecture

The target architecture considered in this work is based on the EAPR design flow and has a static part implementing the communication infrastructure providing a number of interfaces at least equal to the number of RRs (each RR may provide more than one link to the communication infrastructure). Each RFU can communicate with other configured modules or with the static part by just hooking up to the hardware macros corresponding to the communication infrastructure interfaces. The most general view of the target architecture is given in Figure 1: from a physical point of view RRs communicate with the static part through to EAPR stand-alone static-logic nets routing, while from a logical point of view the static logic is responsible for managing inter-module communication implementing different communication infrastructures (such

as bus-based, NoC-based, point to point, and so on). In conclusion, EAPR allows a more scalable communication infrastructure for partial dynamic reconfigurable architectures, since it allows a RFU to communicate with static logic regardless of the position of other RFUs. This approach simplifies the communication infrastructure design, while the frame size still needs to be considered during RRs floorplacement in order to prevent the configuration of a RFU inside a RR from interfering with other RFUs being executed on other parts of the FPGAs.

III. PROBLEM DESCRIPTION AND BASIC DEFINITIONS

Given a set of *Reconfigurable Functional Units*, RFUs, the resource- and configuration- aware floorplacement problem consists of finding, for each RFU, a chip area where it can be: placed and routed offline and configured at runtime. The entire process is subject to the constraint that each RFU must be associated with an area containing all the required types of resources for its functionality.

Let a RFU be a technologically mapped netlist implementing a required functionality, and a *Reconfigurable Region* (RR) be a rectangular FPGA area where two or more RFUs are going to be placed and routed (at design time) and configured (at runtime) according to the application implemented in the reconfigurable system. The goals of the floorplacer are to:

- 1) define the number of RRs and associate each RFU with one and only one RR (partitioning task).
- 2) find a position of RFUs inside the corresponding RRs (Temporal Floorplacement inside Reconfigurable Region),
- 3) find an area constraint for all the RRs inside the FPGA area (RR Floorplacement).

The set of input RFUs is represented by a *scheduled task graph*. Such a task graph is divided into time intervals named *static snapshot*. Each static snapshot is characterized by having the same set of configured and running RFUs, during its time interval. Given a RFU m , the function $\text{TIME}(m)$ returns the set of static snapshots containing m , i.e. static snapshots requiring that m is configured and running. According to these definitions, the scheduled task graph can be considered as a finite state automaton having static snapshots as states and the reconfiguration process as transitions.

Our proposed framework is tailored for Xilinx Virtex 4 and 5 FPGAs [5] that provide access to the off-chip logic through a set of FPGA pins dedicated for communication. The connections between the FPGA's internal logic and external pins are managed through three-state buffers named *Input Output Blocks* (IOBs). In Virtex 4 and 5 families such IOBs are placed along two columns placed around the center of the FPGA, as shown in Figure 2.

Our proposed approach specifically deals with the interconnection optimization. For FPGAs routing resources are often the limiting resources, hence, achieving better routability in a design may in fact determine the overall feasibility of a design.

Our floorplacement framework manages an objective function based on *external wirelength*, i.e., the estimated length

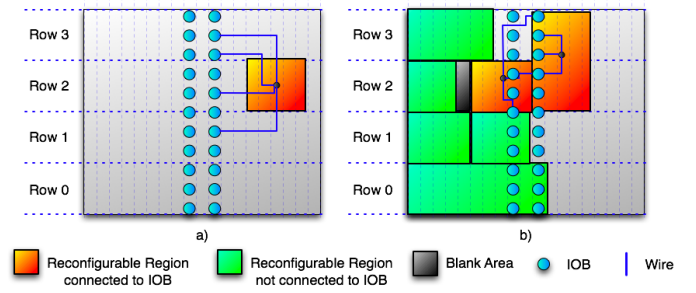


Figure 2. Example of RFU and illustration of the connections between the RR and the corresponding IOBs.

of the nets connecting each RFU to the corresponding IOBs. In order to formulate this objective function the distance between RFUs and IOBs need to be estimated. Experimental results proved that using the *Manhattan distance* between the center of the RFU and the position of the IOB provides a good approximation. The rationale of the wirelength-driven floorplacement is to constrain a group of RFUs characterized by being connected to a set of IOBs within the same area, thereby, creating a neighborhood. Hence, RFUs communicating with IOBs that are near each other, are kept together and constrained within the same area. This approach results in a shorter external wirelength and a highly increased probability of re-use of existing links. While the first outcome is intuitive, the second requires some elaboration. In a dynamically reconfigurable device a common approach is to allocate two main partitions on the device: the static part and the dynamically reconfigured part, where RFUs will be allocated. The communication infrastructure serving all dynamically inserted functionality is managed by the static part of the design and all the communication, both among RFUs and between RFUs and the static part, are performed by the communication infrastructure exposed through hardware macros. Also the communication between RFUs and IOB has to be managed by the static part of the design and similarly has to be exposed through hardware macros. If the floorplacement is aware of IOB positions, RFUs accessing to the same set of IOBs could be constrained within the same area and hardware macros may be re-used by the different RFUs. Consequently, the number of hardware macros that need to be provided by the static part of the design can be drastically lowered toward the theoretical limit of just one set of M hardware macros, i.e. just one for each IOB.

Each RFU of the input scheduled task graph can be annotated with information on the position of the IOBs or, simply, by a point C , named *centroid*, whose coordinates are given by the *arithmetic average* of the coordinates of the IOBs. The centroid represents the ideal position where an RFU should be positioned in order to minimize the external wirelength.

IV. THE FLOORPLACEMENT FRAMEWORK

In this section, we introduce our proposed framework [8], [9] for solving the floorplacement problem targeting the dynamically reconfigurable architecture and the reconfiguration technology described in the previous section. We will describe

the three algorithms defining our framework. The floorplacement process starts with a *Partitioning* step, where RFUs are first grouped into RRs according to two criteria, wirelength for external routing to IOBs and utilization of resources. In the second step, once the partitions have been computed, the position of each RFU inside the corresponding RR needs to be determined. This is performed by the *Temporal floorplacement* step inside RRs. Finally, the design is completed by placing the RRs on the FPGA using the *Reconfigurable Regions Floorplacement* step.

Algorithm 1 Partitioning into RRs

```

1: Buckets B;
2: for all Externally connected RFU  $r$  do
3:   B.add( $r$ );
4: end for
5: Wirelength-driven-partition(B);
6: Fix-existing-associations(B);
7: for all Remaining RFU  $q$  do
8:   B.add( $q$ );
9: end for
10: Resource-driven-partition(B);
  
```

A. Partitioning into RRs

Given N RRs, the N -RRs partitioning problem consists of finding a surjective binding $c_{m,n}$ of RFUs into RRs (i.e. each RFU has to be bound to one and only one RR and each RR has to contain at least one RFU). The algorithm firstly aims at grouping together externally connected RFUs (i.e., RFUs that are connected to IOBs) having nearest centroids and keeping RFUs with distant centroids in different RRs. We refer to this as the *wirelength-driven partitioning*. Secondly, the remaining RFUs are partitioned to minimize the variance of RRs' resource requirements along different static snapshots. In other words, for a given RR, the algorithm tries to keep the amount of resources needed by the RFUs configured and running inside the considered RR constant as much as possible across different static snapshots. We refer to this as the *resource-driven partitioning*.

The problem of wirelength aware partitioning of RFUs into RRs, can be reduced to the problem of clustering the corresponding centroids in a two dimensional space (i.e., chip area). Each identified cluster is associated with one and only one RR and the RFUs are partitioned into RRs according to the association between their corresponding centroids and clusters (a partition belongs to an RR if and only if its centroid belongs to the cluster associated with the RR). Once the wirelength-driven partitioning has been performed, the created partition is used as an initial solution by the resource-driven partitioner to partition the remaining RFUs that are not externally connected. This means that the surjective binding $c_{m,n}$ is no longer modified for externally connected RFUs.

While data-mining algorithms provide several tools to solve the clustering problem (e.g., k-means), they are tailored for large datasets. For our purposes, consisting of less than hundred RFUs per task graph, we adopted a simulated annealing

based approach. The annealer's data structure consists of a simple bucket having a set of RFU associated with one bucket. The RFUs are randomly moved from one bucket to another, and two RFUs (belonging to different buckets) can be swapped. Once the partitions have been computed, the position of each RFU inside the corresponding RR needs to be determined as follows.

B. Temporal Floorplacement inside Reconfigurable Regions

The aim of the Temporal Floorplacement inside Reconfigurable Regions (TFiRR) is to compute, for each RR, a set of height-width pairs describing rectangular areas where all RFUs bound to this RR can be successfully floorplaced. In this phase the final on-chip position of the rectangular area is not considered. For a target FPGA device that is divided by up to k rows for reconfiguration, the goal of this algorithm is to determine for each RR n a set of pairs

$$\Omega = \{ \langle n_h^1, n_w^1 \rangle, \langle n_h^2, n_w^2 \rangle, \dots, \langle n_h^k, n_w^k \rangle, \dots \}$$

such that actual placement of this RR on the device given as $A_n = \langle n_x, n_y, n_h^i, n_w^i \rangle \quad \forall n_x, n_y, i$, results in a feasible floorplacement independently of the final position n_x and n_y decided for this RR. Here, h , w , x and y stand respectively for the height, width and the two coordinates of the bottom-leftmost corner of the rectangular area.

Consequently, for each RR n , the set of height-width pairs Ω can be described by providing just four elements (due to technological constraints related to Xilinx Virtex 4 and 5 FPGAs that are divided in 4 rows)

$$\Omega = \{ \langle 1, n_w^1 \rangle, \langle 2, n_w^2 \rangle, \langle 3, n_w^3 \rangle, \langle 4, n_w^4 \rangle \}$$

where n_w^1 is the smallest width that RR n , floorplaced in 1 row, should have to feasibly host all the associated RFUs, n_w^2 is the smallest width that RR n (floorplaced in 2 rows) should have in order to feasibly host all the associated RFUs, n_w^3 and n_w^4 are the smallest widths that RR n , floorplaced in 3 and 4 rows respectively, should have in order to feasibly host all the associated RFUs.

The core of the TFiRR step is the computation of the pairs $\langle i, n_w^i \rangle$. In order to find, for a given height i , the minimum feasible width n_w^i , the algorithm has to check that every RFU can be successfully floorplaced inside the area described by the pair $\langle i, n_w^i \rangle$, i.e., for each RFU the algorithm has to provide a height, width, and a position within the RR n .

Such a problem is itself three-dimensional (i.e., two spatial dimensions and a temporal one). In order to simplify the problem the following assumption is introduced: *All RFUs' heights are equal to RRs' heights*. In such a way, the problem is reduced to a bidimensional packing problem having snapshot and the width as the only two coordinates. Given a RFU m , and a height i , the smallest position-independent width required by the RFU in order to be hosted inside an area of height equal to i rows, can be easily computed by taking into account the FPGA's resources periodic distribution.

The TFiRR algorithm works as follows:

For each RR n and each possible height $i \in \{1, 2, 3, 4\}$

- 1) Consider RFU m such that $c_{m,n} = 1$ (i.e., RFUs belonging to RR n), let RFUs' height be equal to RR's height, then, compute the minimum feasible width of RFUs m ;
- 2) Pack all the RFUs inside the RR in order to *minimize* the maximum width of RR n .

The packing of the RFUs can be performed with a zero temperature simulated annealing (ZT-SA) algorithm. For each static snapshot p an ordered list of RFUs m , such that $p \in \text{TIME}(m)$, is kept. The RFUs are ordered from the left-most to the right-most with respect to the RR's area. The following moves are applied:

- 1) *Randomly move an RFU.* Randomly pick an RFU and move to an integer position belonging to the interval $[0, \text{width}]$, where width is the current width of the RR.
- 2) *Randomly swap two concurrent RFUs.* Randomly pick two RFUs m' and m'' , being concurrently configured and running in at least one static snapshot, and swap their position inside RR.

To keep the floorplacement compact, each step of the annealer is followed by a *compression* function that computes, for each RFU, the left-most feasible solution preventing overlaps between RFUs.

The quality of this algorithm at first seems to be strictly related to the quality of the partition provided by the previous step. Our experiments show that TFiRR applied on non-partitioned task-graphs can reach the results of TFiRR applied to a partitioned task graph by increasing the number of iterations by at least two orders of magnitude. On the other hand, the difference between the objective functions remains fairly low. We observed degradations ranging between 1 – 5 %.

At the end of this second step, each resulting RR is annotated with a centroid whose coordinates are given by the arithmetic average of the corresponding coordinates of the RFUs associated with the considered RR. This identifies the ideal position where each RR should be placed in order to globally minimize the external wirelength of the associated RFUs. This formula for computing the centroid places more emphasis on the most heavily utilized IOBs. For example, if three RFUs are connected only to the USB interface, then all of them will have the same centroid C_{USB} that will occur three times in the set of centroids associated with the RR. In the third and final step, the centroids of the RRs will be used during the final floorplacement of RRs.

C. Reconfigurable Regions Floorplacement

The aim of this step is to define, for each RR n , an area

$$A_n = \langle n_x, n_y, n_h, n_w \rangle$$

The algorithm has to choose one $\langle n_w, n_h \rangle$ couple, for each RR n , out of the set provided by the TFiRR step. Furthermore, it has to determine the specific x and y positions on the FPGA area. According to classical floorplanning this task can be performed through simulated annealing. The RR Floorplacement algorithm is divided in two steps: the first




Figure 3. a) A variation of horizontal constraint graph used for floorplan representation b) Negative (N) and Positive (P) area slacks (the empty space on the second row is not included in P because it has a width greater than one of the empty areas in the upper row).

one floorplacing the RRs connected to IOBs (wirelength-driven RRs floorplacement) and the second one the remaining RRs (area-driven RRs floorplacement). The annealer's data structure consists of a Horizontal Constraint List (HCL) The first sub-step of this step is characterized by an objective function that must take directly into account the wirelength:

$$\Gamma = \left(\sum_{r \in \text{RR}} d(r, C(r)) \cdot \#\{RFU \in r\} \right)^f$$

where $d(r, C(r))$ represents the distance of the RR from its ideal position (centroid), $\#\{RFU \in r\}$ is the number of RFUs connected to IOBs belonging to the r -th RFU, while f is a positive number. If f is small, this indicates that the floorplacement is feasible and if it is large the floorplacement is not feasible. The goal of the floorplacement is to minimize Γ . Note that the objective function is weighted by the number of externally connected RFUs. This means that RRs containing more RFUs connected to the external world would benefit from a partitioning in the neighborhood of the centroid. Once the wirelength-driven RR floorplacement has been performed, the remaining RRs can be floorplaced by a purely area-driven algorithm. This second sub-step is guided by an objective function involving free area and feasibility of the final floorplacement. Given an RRs' floorplacement, the following quantities are defined: *Negative area slack* (N), i.e. the area of the floorplan crossing target device boundaries, and *positive area slack* (P), i.e. the greatest contiguous free area starting from the right-top most corner of the device and with non increasing width going bottom-ward. Figure 3 shows an example of such slacks. Given such slacks, the RR Floorplacement objective function is defined as follows (where $M \in \mathbb{N}$ and greater than the number of frames on the target FPGA area):

$$\Theta = P - M \cdot N$$

This objective function Θ is positive if the floorplacement is feasible (i.e. $N = 0$), otherwise it is negative (because $P < N$). The aim of the annealer is to maximize Θ , consequently, to provide a feasible floorplacement maximizing the contiguous FPGA area left free for static logic.

Given the HCL data structure, the following moves are defined:

- 1) *Randomly Swap two RRs:* randomly choose two RRs and swap their positions;




Figure 4. Limits of the floorplacement of big modules

- 2) *Move an RR to a randomly chosen position*: randomly pick an RR n and two coordinates $\langle \bar{x}, \bar{y} \rangle$ belonging to the device area.
- 3) *Span a randomly chosen RR over rows*: randomly choose an RFU, having height less than $\#ROWS$, and increase its height by 1 row.
- 4) *Unspan a randomly chosen RR*: randomly choose an RFU, having height greater than 1, and decrease its height by 1 row. It is the inverse of the span move.

It can be noticed how the floorplacement of a large number of small functional units is easier than the floorplacement of a small number of large functional units. A bad choice of floorplacement of a large functional unit during the early stages of the floorplacement is difficult to correct in the later steps, particularly when temperature decreases rapidly and each correcting move is likely to be rejected because it results in a worse objective function.

D. Identifying optimal number of partitions

The three steps described above define our floorplacement framework. The overall framework relies on the concept of partitioning, hence, some final remarks on how we control the granularity of these partitions are useful. In order to identify the most suitable number of partitions let us consider the maximum number of concurrently configured C_C RFUs, i.e. how many RFUs are present at most in one static snapshot in any partition. We define this quantity as follows:

$$C_C = \max \{n_{p,b} | \# \{RFU r | p \in \text{TIME}(r) \wedge r \in B_b\}\}$$

where $B_1, \dots, B_{\# \text{partitions}}$ represent the different output partitions. C_C can be considered as a good metric to describe the complexity of the entire partition. Let Γ defined as the global normalized variance in resource requirements, used to represent the heterogeneity of the partitions. We have observed experimentally that the number of partitions minimizing the product of Γ and C_C provides (in resource-driven partitioning algorithm) a good trade off between partition complexity and intra-partition variance.

V. RELATED WORKS

The problem of resource aware floorplanning tailored to FPGAs has been addressed in literature [8], [10]. Montone et

al. [8] proposed a floorplanning method that is only aware of resource requirements for the reconfigurable modules, taking into account the physical limits of the target FPGAs and the constraints imposed by the reconfiguration process. In this paper we extend such formulation to consider also the wirelength. Feng and Mehta [10] proposed an approach, where each architectural module has a list of required resources, hence each module has to be placed in an area region containing all the needed resources. This approach is based on a two steps algorithm: first the execution of Parquet [11] floorplanner with resource aware cost function consisting of a linear combination of Parquet's objective function and the amount of satisfied resource requirements. The result is then refined solving a flow maximization with cost minimization on a purposely built graph. While this approach is currently the state of the art in resource aware floorplanning for FPGAs, in the majority of cases, the result is not compatible with the FPGA reconfigurability process.

One of the earliest works dealing with reconfiguration aware floorplanning is [12]. An offline floorplanner is used to decide whether a required functionality should be implemented in hardware on a reconfigurable device or has to be executed via software on a general purpose processor. Different heuristics are proposed to minimize execution time, while reducing implemented modules' fragmentation across the device area.

The first definition of *temporal floorplanning* has been introduced in [13]. According to the author a temporal floorplanning consists of two phases: i) partitioning and sequencing design modules into design configurations (also called temporal partitioning or scheduling), and spatial positioning of design modules and wiring within the reconfigurable area, in the following named as reconfigurable region, of each configuration.

One of the most important contributions in three dimensional floorplanning is presented in [14], [15]. They introduce a full 3D floorplanner based on simulated annealing using a three dimensional transitive closure graph (TCG) and T-trees. They evaluate the time required by each module to communicate with RAM chips outside the FPGA in order to store results and read input, but they do not evaluate if a found floorplan, in particular its communication infrastructure, is feasible on a given device or not. Furthermore, they do not consider other resources than logic blocks.

A more recent work [16] faces the FPGA floorplanning problem considering partial dynamic reconfiguration, in terms of module reusability. They introduce a variation to the sequence pair representation [17] to represent the floorplan in different time frames. The aim of this work is to reduce the quantity of device area reconfigured between different time frames. Given two designs in two different time instants, the authors present a simulated-annealing based floorplanning able to reduce the reconfigured area between the two instants by exploiting re-use of already configured modules.

A comparison between previous works is provided in Table I. The most common limitations of prior work is in their lack of control over the feasibility of the resulting communication infrastructure. Furthermore, most of the related work do not

treat resource awareness as a primary goal. However, this is of utmost importance for making module allocation decisions for modern dynamically reconfigurable systems, which contain a variety of heterogeneous resources. Our proposed floorplacement method aims to address both these important problems, i.e., enabling an interconnect and resource aware design framework.

To the best of our knowledge, there are no prior works dealing with the minimization of wirelength between RFUs and IOBs. The problem of minimizing internal wirelength of RFUs is considered in the past [15], but this approach is based on an estimation of the internal wirelength related to RFUs' aspect ratio and cannot be easily extended to support external wirelength. Similarly, by applying the approach introduced in [16], a re-use of links can be obtained, but this approach aims only at re-using RFUs implementing the same functionality, while our approach aims at re-using links between different RFUs connected to the same set of IOBs.

VI. EXPERIMENTAL RESULTS

Our proposed approach for the wirelength and resource management has been validated both on randomly generated task graphs and real world applications from the domain of data processing for biomedical applications (i.e., collecting data from sensors, performing some preprocessing like FIR filtering, computing error detection codes and sending data through a network). This test set is composed of RFUs having a space requirement varying from few hundreds CLBs till several thousands (e.g., filtering functional units). Furthermore, the considered task graphs are generally composed of several dozens of RFUs and the criteria for their random generation are the same adopted in [8]. Similarly the RFUs' device dependent aspect-ratio is computed according to the description provided in [8]. The results introduced by our approach are application dependent, therefore we will describe three different metrics used to evaluate the results. For each metric, we will provide the description of the task graphs. (Table II shows a summary of the metrics and a summary of our comparison [8] that describes a floorplacement method that is only resource-aware, but does not consider the wirelength implications of resource management).

Table II
QUALITY METRICS AND THE SUMMARY OF THE VARIATION IN THESE METRICS COMPARED WITH EXISTING WORK [8].

Metric	Variation/Value
External Wirelength	reduction in external wirelength ranging between (90, 30)% compared to existing area-driven method
Links	reduction in number of links required ranging between (90, 0)% compared to existing area-driven approach
Blank Area	(5, 35)% of the final floorplacement

Our results confirm that introducing the wirelength awareness indeed improves the interconnect cost significantly. We observed a reduction of 90% in external wirelength in the best case and an average reduction of 50%. Task graphs with few externally connected RRs lead to the highest wirelength reduction. On the other hand, the reduction for external wirelength in task graphs with many externally connected RRs

with near centroids is less. In such task graphs only few RRs can be placed near their centroids and the other RRs have to be placed far away. The number of hardware macros provided by the static part can be reduced by 90% in the best case. Task graphs containing several RFUs accessing (in different static snapshots) the same set of IOBs (e.g. RFUs connected to the same external interface like USB) would particularly benefit from our approach, while task graphs having all the RFUs connected to a distinct sets of IOBs would not benefit from links reduction (this is the 0% reduction case referred in Table II).

Figure 5 shows how our proposed approach drastically reduces the number of required links. In this figure we also observe one weakness of our approach.

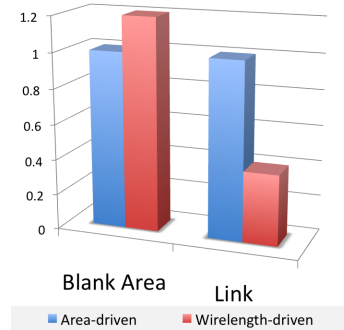


Figure 5. Comparison of the area and wirelength metrics among the area-driven and wirelength-driven approaches (normalized w.r.t. Area-driven approach).

The most relevant drawback of our approach is referred to as the *blank area* problem (i.e., the amount of area being surrounded by RRs but not assigned to any RR, see Figure 2). Let us consider a set of RRs, each one being externally connected and apply our approach several times, each time increasing the percentage of RRs that are considered (by the floorplacer) as attached to IOBs. During the first iteration no RR is considered as attached to IOBs, while during the last iteration all RRs are considered as attached to IOBs. Figure 6 plots the percentage of the final floorplacement that remains as blank area with respect to the percentage of RRs considered as attached to IOBs as a result of this experiment.




Figure 6. Example of wirelength reduction (normalized) and blank area left (percentage of the final floorplacement) for the floorplacement of only externally connected RRs plotted with respect to the percentage of RRs that are considered as externally connected by the algorithm.

When no RR is considered as externally connected a

Table I
COMPARISON AMONG PREVIOUS WORKS

Authors	Comm Infrastructure	Resources Aware	Reconfigurability Aware	Reusability Aware	Algorithm
Feng et al. [10]	No	Yes (high res. usage)	No	No	Sim. Annealing over seq. pairs + Flow Max over Flow Graph
Montone et al. [8]	No	Yes	Yes	For logic only	Sim. Annealing
Bazargan et al. [12]	No	No	Yes	No	Sim. Annealing over cubic modules (2d spatial, 1d temporal)
Yuh et al. [14], [15]	Limited, w/ High Overhead	No	Yes	No	Sim. Annealing over cubic modules using T-trees and TCG
Singhal et al. [16]	No	No	Yes	Yes	Sim. Annealing over seq. pairs

5% blank area is obtained (same as the purely area-driven approach). The peak is obtained when half of the RRs are managed by the wirelength-driven algorithm and the other half by the area-driven, in such a case the blank area may reach 30-35% of the final floorplacement. On the other hand, the blank area is generally divided in no more than three or four areas that are wide enough to be used by the static part of the design according to the EAPR design flow. Figure 6 shows also how the first externally connected RRs obtain a great wirelength improvement, while the latest one cannot obtain such improvement due to the previously introduced non-overlapping constraints.

Therefore, we observe that our approach can yield significant improvements in the cost of the communication infrastructure by trading-off a reasonable amount of blank area. Using our framework, the designer can choose between a) considering more RRs as externally connected, thereby, decreasing external wirelength or b) considering that beyond a certain point blank area overcomes benefits provided by the wirelength-driven approach. Hence, the designer may decide to floorplace in a wirelength-driven way only for the most relevant RRs leaving the others to be floorplaced by an area-minimizing approach. Finally, we observe that the blank area problem is not an issue for task graphs requiring most of the resources of the target FPGA (because the feasibility of the floorplacement requires as much area as possible to be used, hence blank area is reduced as a consequence) and for task graphs having few RFUs connected to IOBs (or many of them connected to the same IOBs). Finally, due the reuse of the Floorplacement Framework introduced in [8], [9], the execution time of the wirelength driven floorplacement is comparable with the one presented for the area driven floorplacer. Consequently, the proposed solution does not introduce any particular execution overhead.

VII. CONCLUSIONS AND FUTURE WORK

In this paper, we presented a resource- and configuration-aware floorplacement framework, tailored for Xilinx Virtex 4 and 5 FPGAs, using an objective function based on *external wirelength*. The proposed approach has achieved a shorter external wirelength and a highly increased probability of re-use of existing communication links. The reduction in wirelength ranges from 30% to 90% in comparison to a purely area-minimizing approach. Task graphs with few externally connected RRs lead to the biggest decrease, while external

wirelength in task graphs with many externally connected RRs show lower improvement. Our future work aims at directly addressing the blank area problem by modifying the objective function or floorplacing all the RRs, not connected to the IOBs, around the RRs connected to IOBs, trying to keep the RRs as close as possible in order to further reduce the blank area.

REFERENCES

- [1] *ISE 9.2i Manual*, Xilinx Incorporation, 2007.
- [2] V. Betz and J. Rose, *VPR: A new packing, placement and routing tool for FPGA research*. London, UK: Springer-Verlag, 1997.
- [3] J. Roy et al., "Min-cut floorplacement," *IEEE Transactions on CAD of IC and Systems*, vol. 25, no. 7, pp. 1313–1326, July 2006.
- [4] S. Adya, S. Chaturvedi, J. Roy, D. Papa, and I. Markov, "Unification of partitioning, placement and floorplanning," in *Computer Aided Design, 2004. ICCAD-2004. IEEE/ACM International Conference on*, Nov. 2004, pp. 550–557.
- [5] *Virtex 5 - Family Overview*, Xilinx Incorporation, Dec 2007.
- [6] *Early Access Partial Reconfiguration User Guide*, Xilinx Incorporation, 2006.
- [7] *Xilinx Application Note 290*, Xilinx Incorporation, December 2007.
- [8] A. Montone, F. Redaelli, M. D. Santambrogio, and S. Memik, "Reconfiguration-aware floorplacer for fpgas," *IEEE International Conference on ReConfigurable Computing and FPGAs*, 2008.
- [9] A. Montone, M. D. Santambrogio, and D. Sciuto, "A design workflow for the identification of area constraints in dynamic reconfigurable systems," *Electronic Design, Test and Applications, IEEE International Workshop on*, vol. 0, pp. 450–453, 2008.
- [10] Y. Feng and D. Mehta, "Heterogeneous floorplanning for FPGAs," in *VLSI Design, 2006. Held jointly with 5th International Conference on Embedded Systems and Design., 19th International Conference on*, Jan. 2006.
- [11] S. Adya and I. Markov, "Fixed-outline floorplanning: enabling hierarchical design," *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. 11, no. 6, pp. 1120–1135, Dec. 2003.
- [12] K. Bazargan, R. Kastner, and M. Sarrafzadeh, "3-d floorplanning: simulated annealing and greedy placement methods for reconfigurable computing systems," in *Rapid System Prototyping, 1999. IEEE International Workshop on*, Jul 1999, pp. 38–43.
- [13] M. Vasilko, "Dynasty: A temporal floorplanning based cad framework for dynamically reconfigurable logic systems," in *FPL '99: Proceedings of the 9th International Workshop on Field-Programmable Logic and Applications*. London, UK: Springer-Verlag, 1999, pp. 124–133.
- [14] P.-H. Yuh, C.-L. Yang, and Y.-W. Chang, "Temporal floorplanning using the t-tree formulation," in *Computer Aided Design, 2004. ICCAD-2004. IEEE/ACM International Conference on*, Nov. 2004, pp. 300–305.
- [15] —, "Temporal floorplanning using the three-dimensional transitive closure subgraph," *ACM Trans. on DAES*, vol. 12, no. 4, p. 37, 2007.
- [16] L. Singhal and E. Bozorgzadeh, "Multi-layer floorplanning on a sequence of reconfigurable designs," in *Field Programmable Logic and Applications, 2006. FPL '06. International Conference on*, Aug. 2006, pp. 1–8.
- [17] H. Murata et al., "Vlsi module placement based on rectangle-packing by the sequence-pair," *IEEE Transactions on CAD of Integrated Circuits and Systems*, vol. 15, no. 12, pp. 1518–1524, Dec 1996.