

# Timing-Constrained FPGA Placement: A Force-Directed Formulation and Its Performance Evaluation

SRILATA RAMAN<sup>a</sup>, C. L. LIU<sup>b</sup> and LARRY G. JONES<sup>a</sup>

<sup>a</sup>Motorola Inc. Austin, TX, USA; <sup>b</sup>University of Illinois at Urbana-Champaign, Urbana, IL, USA

In this paper we present a simple but efficient timing-driven placement algorithm for FPGAs. The algorithm computes forces acting on a logic block in the FPGA to determine its relative location with respect to other blocks. The forces depend on the criticality of nets shared between the two blocks. Unlike other net-based approaches, timing constraints are incorporated directly into the force equations to guide the placement. Slot assignment is then used to move the blocks into valid slot locations on the FPGA chip. The assignment algorithm also makes use of the delay information of nets so that the final placement is able to meet the timing criteria specified for the circuit. The novelty of the approach lies in the formulation of the force equations and the manner in which weights of the nets are dynamically altered to influence the placement. Experiments conducted on industrial test circuits and MCNC circuits give very promising results and indicate that the algorithm succeeds in significantly reducing the maximum delay in the circuit. In addition, routability is not adversely affected and running time is low.

*Keywords:* Symmetrical FPGA, timing-oriented placement, force-based algorithm, slot assignment

## 1. INTRODUCTION

The emphasis in VLSI designs today is on high circuit performance which has become synonymous with high speed of the circuits. With recent advances in the fabrication technology, delays of interconnects play a dominant role in determining the speed of a chip. This is particularly true in the case of FPGA architectures in which the logic blocks have predetermined delays and the interconnects dictate the speed of the configured chip. One of the chief motivations for this work is to study timing-based layout algorithms designed specifically for FPGAs. Although a lot of effort has been devoted to technology mapping

of the circuit logic onto FPGAs, layout issue has not received comparable attention. In FPGAs, in contrast to gate-array and standard-cell technologies, layout involves dealing with a lot of constraints because of limited and preset routing resources. In addition, for high-speed designs, timing constraints must be incorporated into the layout tools for regulating the timing behavior of the circuits. Most of the work in layout for FPGAs has concentrated on detailed routing, and placement is carried out using techniques such as simulated annealing which can take up a lot of time. Our aim is to incorporate timing delays of interconnects at the design layout stage and determine how they influence the speed and routability of the FPGA.

For gate-array and standard-cell technologies, incorporating timing constraints in the placement algorithm has been studied extensively yielding good placement results [1,2,3,4]. Efforts in timing-driven placement algorithms can be streamlined into two categories: *net-based* and *path-based*. In a net-based algorithm, higher weights are assigned to nets lying on critical paths as in [1] or stringent upper bounds are imposed on the lengths of these nets as in [5]. The net-based approach helps to improve the timing behavior without incurring high computational overhead. A very good theoretical basis in support of this approach is provided in [6]. In a path-based algorithm, the delays of all nets lying on the critical paths are controlled simultaneously [7,8]. Most of these algorithms employ a mathematical programming approach to obtain timing-accurate solutions but are computationally intensive. However, work in the area of timing-driven layout for FPGAs has been scant, two recent papers being [9] and [10].

There are three key points in our algorithm that distinguish it from other net-based techniques. Most net-based algorithms use some form of net-weighting to assign weights to nets depending on their criticality. Often, the weights of the nets are selected from a prespecified range and it is left to the discretion of the user to specify the exact value for each net. Our algorithm makes a significant departure from this practice. We use specified allowable delays of nets to directly determine the weights of nets which are then used in the formulation to guide the placement. The second feature of our algorithm is that it ensures that non-critical nets are not adversely affected in the placement process. The third feature is an adaptive net weighting scheme in which the weights of the nets are altered dynamically in accordance with the path delay requirements to guide the subsequent placement.

The organization of the paper is as follows. In section 2, we describe our timing-constrained placement algorithm for symmetrical FPGAs. The factors that make FPGA layout problem challenging and the underlying timing model are described first, followed by details of each step of the algorithm. Section 3

presents experimental results and performance evaluation of the placement algorithm.

## 2. TIMING-CONSTRAINED FPGA PLACEMENT

The objective of timing-based placement is to optimize locations of logic blocks on the chip for maximum speed, density, and high percentage of routability. It is worthwhile to consider at this point the issues that distinguish FPGAs from other technologies thereby making it imperative to revisit and alter the existing CAD layout tools. Since FPGA chips are not customized, it is very critical to implement circuits on them with as little wastage of resources as possible. Placement takes the longest time to complete in layout. Timing behavior and routability can be influenced much more at the placement stage than at the routing stage. This is because the flexibility of routing is limited by an already existing arrangement of blocks. In gate array, standard-cell, and custom cell designs the prime objective is to minimize the routing area and the routing resources are not predetermined. The routing resources in FPGAs typically consist of pre-laid segments of interconnects of fixed lengths that can be extended using connection matrices. These matrices have capacitances associated with them and therefore contribute to the delay. As a result, the delay of a net is influenced by the number of segments and the number of connection matrices used by it. The CAD layout tools must take into account this effect in order to be efficient. Timing-driven placement in FPGAs is an art in tradeoffs. The main tradeoff is between a placement that meets the timing requirements and one that is routable for the FPGA architecture under consideration. The dominant and alterable part of the delays in FPGA layout comes from interconnects since the logic blocks have fixed delays. In contrast, in gate-array and standard-cell technologies, the delays in the circuit can also be controlled by redesigning the logic blocks or by transistor sizing.

We assume a generic symmetrical FPGA architecture in this paper similar to that in [9]. It has a two-dimensional array of logic blocks interconnected by vertical and horizontal routing channels. The logic blocks are the functional elements used to construct the circuit. The input-output blocks on the periphery of the chip provide interface with the logic not resident on the chip. Both these blocks are user-programmable.

The general flow of our algorithm is outlined in figure 1. The algorithm has three phases as indicated by the **while** loop of the figure. Before delving into the details of each of the phases, we describe the concepts underlying the timing model assumed in this paper.

## 2.1 Timing Issues

The terminology used in the rest of the paper is as follows. *Path delay* is the time required by a signal to propagate along a path between two input-output blocks or storage elements. A path is associated with two times—*latest required arrival time* and *actual arrival time*. The former refers to the latest time by which the signal should arrive at the path sink for it to be processed correctly. It serves as an upper bound on the path delay. The latter refers to the slowest time that the signal actually takes to arrive at the path sink. A path is deemed *critical* if its total delay is too close to the latest required arrival time. Nets lying on the critical paths are termed critical nets. For this paper, we consider the problem of handling long paths, that is, paths whose delays exceed the latest required ar-

rival times. We also assume that structural false paths have been eliminated and are not specified during the layout process.

The input to the timing-constrained placement algorithm is a circuit net-list that describes how the logic blocks are interconnected. The circuit has  $M$  logic blocks (LBs),  $N$  input-output blocks (IOs) that serve as the Primary Inputs (PIs) and Primary Outputs (POs), and pins located on the periphery of the logic blocks and the IO blocks. A net connects pins on the same or different blocks (logic or IO). The delays of logic blocks and IOs are assumed to be given and are treated as constants. In a logic design, the logic block and nets feeding it are associated with two delays—rising and falling. For simplification and without loss of generality, we consider only a single delay as in [1].

The nature and arrangement of the pre-laid interconnects constitute the routing architecture of the FPGA. We assume there are two kinds of interconnects available—uniform length grid segments (equal to the height or width of a logic block) joined together by connection matrices, and long segments that span the length and width of the FPGA chip. The long segments are typically reserved for routing global nets and nets with large fanout. If the length of a net is  $l_n$  and its delay is  $d_n$ , then if  $c$  is the capacitance,

$$d_n = c * f(l_n) + k_1 \quad (1)$$

Here,  $k_1$  is a constant and  $f(l_n)$  is a function of the length of the net. In case of FPGAs, the function  $f$  is governed by the number of segments and number of connection matrices used to route the net, and  $c$  is the sum of the capacitance contributed by a connection matrix and its incoming grid segment. According to equation (1), solving a placement problem with net delay constraints is same as solving a problem with constraints on the wire length of nets. Hence, we use the terms net delay and net length interchangeably in this paper.

The zero-slack algorithm [5] is used to obtain a set of upper bounds on the lengths of nets and these con-

```

begin
  identify critical paths of the circuit CP
  find upper bounds on all net lengths ;
  while timing constraints not met
    begin
      1. perform force-directed placement;
      2. perform slot assignment;
      3. check path delays and associated net delays;
         remove paths from CP that meet the timing criteria;
         reassign weights to nets in remaining paths;
    end
  end
end

```

FIGURE 1 Overall Structure of the Algorithm.

stitute timing constraints for the placement problem. Critical path information is obtained from a timing analysis done prior to placement. Zero-slack algorithm maximizes the range of wire lengths for each net and returns a set of *non-unique* bounds on the net lengths. If the timing constraints are to be met, all nets must have lengths less than or equal to their respective upper bounds. While the upper bounds over-constrain a lot of other net-based algorithms, we use the non-uniqueness of the upper bounds to our advantage, as will be demonstrated in the section on weight assignment.

## 2.2 A Force-Directed Placement Algorithm

The technique we adopt is based on the classical force-directed placement (FDP) method [11,12]. The method uses a mathematical formulation to model the forces among the logic blocks (LBs) and solves the resulting set of equations to determine the relative placements of the LBs. The idea is based upon a physical model in which if a signal net is shared between two LBs, there is an attractive force between them; otherwise, a repulsive force comes into play between them. The attractive force tries to bring the connected LBs closer together while the repulsive force tries to move apart the LBs that do not share any nets. In the classical approach, the main objective is to obtain a placement that minimizes the total wire length. Consequently, the forces are a function of the number of signal nets between the LBs or the connectivity. This method increases routability by placing strongly connected components close together but does not ensure that timing requirements on the critical paths are met. In the formulation detailed next, instead of using a connectivity-based formulation we model the force equations to reflect the timing delays of the nets and thereby ensure that timing constraints specified for the critical paths of the circuit are satisfied. We refer to this as timing-constrained force-directed placement (TFDP).

There are three key points in favor of the force-directed formulation that have prompted its use.

- The salient feature is that it takes a global perspective of the problem. This implies that the formulation takes into account constraints on all the blocks concurrently and hence attempts to find an arrangement that does justice to constraints on all the blocks. This is unlike other iterative methods that perturb a subset of the blocks at any time and use a cost function to evaluate the effect of the perturbation. This can adversely affect the blocks that are not perturbed if the cost function is not accurately modeled. This is particularly true of timing-based formulations since it is difficult to capture the effect of block perturbations on all path or net delays in a single cost function.
- The technique is relatively insensitive to initial placement, very much unlike other commonly used placement techniques such as the partitioning-based ones and those based on simulated annealing. This means that the initial random positioning of blocks has very little effect on the final placement. This is a very nice feature because it obviates the need to make several runs to get the best placement, as needs to be done in many iterative algorithms.
- The number of iterations required for convergence to the final placement is dependent more on the topology of the circuit rather than on the size of the circuit. The size of the circuit primarily affects the time required for each iteration only.

In addition to these, the timing-based formulation of our algorithm has distinct advantages over other net-based placement techniques. The other techniques rely on assigning weights to nets that are deemed critical after the timing analysis. The manner in which the weights are assigned is often left to the discretion of the user. For instance, typically, a range of weights (1 to 10) to be used for the nets is given. The non-critical nets are assigned a weight of 1 while the rest of the nets are assigned weights depending on how critical they are. However, it is very difficult to assess the criticality of each net in order to assign a suitable weight to it. Our algorithm obviates this shortcoming by using the upper bounds on the net lengths obtained after the timing analysis

directly as weights in the formulation, thereby eliminating the intermediate step of translating net delays into net weights. This not only ensures that the algorithm closely follows the actual timing constraints, but also avoids the burden of picking the right set of net weights. More details on the actual implementation are given in section 2.4. In order to better understand the formulation, we consider the circuit shown in figure 2(i). The square boxes are the LBs and the letters inside them designate their names. The 2-tuples below the LBs in figure 2(ii) give their physical coordinates on the chip. The critical path of the circuit is shown as the bold line from LBs  $a \rightarrow b \rightarrow c$ . All nets along this path are critical. The bounds on the lengths of the two net segments of the critical path are shown as  $L_a$  and  $L_c$ . In order to incorporate the timing constraints into the formulation, we model an attractive force between two LBs if they lie on the same path and share a net that is considered *critical*. Thus, referring to figure 2(ii), there is an attractive force  $Fatt_{ba}$  and  $Fatt_{bc}$  between LBs  $b$  and  $a$ , and between LBs  $b$  and  $c$ . The magnitude of this force depends on the allowable net length of the critical net that the two LBs share. Thus, if  $\Delta S_{ba}$  denotes the Manhattan distance between the two blocks  $b$  and  $a$ , then the attractive force between  $b$  and  $a$  is given by  $Fatt_{ba} = -\Delta S_{ba}/W_{ba}$ . Here,  $W_{ba}$  is the *weight* assigned to the critical net shared by the two LBs and is initially set to  $L_a$ . The intuition behind such a formulation is that smaller the allowable net length, greater will be the attraction between the blocks so that they are kept close to each other. Similarly, there is an attractive force between LBs  $b$  and  $c$  that is given by  $Fatt_{bc} = -\Delta S_{bc}/W_{bc}$ , with  $W_{bc} = L_c$  initially. More discussions on net

weights will be deferred until the section devoted specifically to weight assignment. While the attractive force takes care of bringing the LBs closer together as dictated by the allowable net lengths, a repulsive force is needed to move the LBs that are not on the same critical path away. To reflect this need, a repulsive force manifests itself between two LBs if they do not lie on the same critical path and hence do not share a critical net. The key point here is that a force of repulsion can exist between two blocks even when they have some connections in common as long as they do not lie on the same critical path. Thus, there is repulsive force,  $Frep_{bd} = R$ , between LBs  $b$  and  $d$ , and  $Frep_{be} = R$ , between LBs  $b$  and  $e$ ,  $R$  being the repulsion factor. In general, the repulsion factor between two LBs is 0 if they share a critical net; otherwise, it is given by  $R$  which is calculated as:

$$R = (1/TK)[\sum_{i=1}^M \sum_{j=1}^M A_{ij}]$$

Here,  $K$  is a constant (experimentally found to be between 1 and 2),  $M$  is the number of logic blocks, and  $A_{ij}$  is 1 if the LBs  $i$  and  $j$  share a critical net and 0 otherwise.  $T$  is the total number of  $A_{ij}$ 's that are 0. The total force on LB  $b$  is then given by the sum of the attractive and repulsive forces acting on it due to all the other LBs. The total force on each of the other blocks can be determined in a similar fashion. Since the information about criticality of the nets is embedded in the force equations, there is a tight coupling between the blocks that share critical nets. This favors our primary objective of controlling the delays on the critical paths. Once the force-formulation has been done, the algorithm tries to find positions for the LBs such that the forces acting on them are reduced to zero; in other words, the attractive and repulsive forces are balanced. This is done by iteratively allowing the blocks to move in accordance with the forces acting upon them until either the total forces on all the blocks fall below a tolerance level or the blocks do not move significantly from their current positions. The forces are resolved in the X and Y direc-

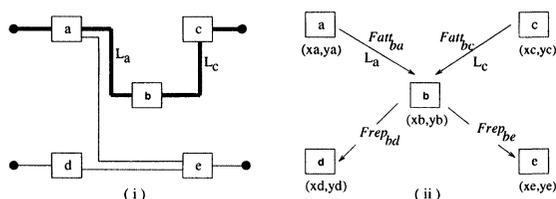


FIGURE 2 Calculation of forces:(i) Sample Circuit (ii) Forces on block b

tions for this purpose resulting in  $2M$  nonlinear equations that are solved using the quadratically convergent Modified Newton-Raphson technique. We would like to emphasize here that unlike other net-based approaches, we do not order the nets on the basis of their weights for placement purposes. The IO blocks are treated as fixed blocks during the course of the algorithm and their positions are predetermined. All the other blocks that have the logic mapped onto them are free to move anywhere within the boundary of the chip. In order to keep the LBs within the physical boundary of the chip, the center of gravity of the chip is taken to be the geometrical center of the chip and the force equations are augmented to reflect this requirement.

### 2.3 Slot Assignment

Once the relative location of the blocks has been obtained at the end of the first phase, the LBs have to be assigned to fixed positions, *slots*, on the FPGA chip since the solution from the first phase of the algorithm assumes a continuous plane for movement of the blocks. The slots are predetermined by the architecture of the FPGA under consideration. We use a linear assignment algorithm for performing the slot assignment. In this phase also, we make use of the timing information on the nets to position the blocks. The assignment algorithm employs an  $M \times S$  cost matrix,  $S$  represents number of slots on the FPGA such that  $M \leq S$ . Each element of the matrix represents the cost function value  $L_{ij}$  which is the cost of moving the block  $i$  from its current position to slot  $j$ . In most of the linear assignment algorithms the cost function is chosen to be the Euclidean distance between the current position of the block and the position of the slot. But this measure minimizes the movement of the blocks only. This has the effect of moving the blocks to the nearest available slots without regard to the effect on the delays of the nets that a block shares with other blocks. To ensure that the delay bounds on nets are preserved after the first phase, we use a cost function that reflects the timing constraints imposed on the nets. In our implementa-

tion of the assignment algorithm, the cost function is chosen to be the increase in net lengths of all nets that are connected to a block caused by a change in the position of the block. If there are  $t$  nets connected to block  $i$ , then

$$L_{ij} = \sum_{k=1}^t \Delta l_{ijk}$$

where,  $\Delta l_{ijk}$  represents the increase in length of net  $k$  due to shift in location of block  $i$  from its current position to slot  $j$ . The motivation behind choosing such a cost function is that we wish to reduce the increase in the net length and hence the delays on the nets as a result of moving the blocks.

### 2.4 Adaptive Weight Assignment

As stated earlier, the attractive force between two logic blocks  $i$  and  $j$  is biased by a weight  $W_{ij}$  which is dependent on the timing requirement of the critical path. An appropriate choice of weights is very important since it dictates how well the timing requirements will be met. It also affects the convergence of the force-directed placement algorithm since it affects the attractive forces. Although the zero-slack algorithm gives a set of upper bounds on the net lengths, which when satisfied, guarantee that the circuit will meet the timing constraints specified, it can also over-constrain the problem. To circumvent this, it is noteworthy here that the set of upper bounds on net lengths is not a unique set, which means that it is possible that some other combination of net lengths would satisfy the timing constraints just as well. Consequently, it is possible that while some net lengths along a critical path are in excess of the upper bounds specified for them, other nets along the same path have lengths well below their upper bounds, thereby ensuring that the timing requirement on the path is still satisfied. Hence, adhering strictly to the upper bounds without paying attention to the path delays should be avoided. With this intent, in our approach, the weights of the nets are not fixed, but dynamically altered during the course of the algorithm to reflect the current path delay requirements. We re-

fer to this approach as *adaptive net-weighting* and have found that it improves the performance considerably. The steps involved in adaptive weight assignment are now described.

- Start with an initial set of net weights that are the upper bounds on the lengths of the nets obtained from the zero-slack algorithm. Weight  $W_{ij}$  between LBs  $i$  and  $j$  is the minimum of all upper bounds on net lengths shared between the two LBs.

$$W_{ij} = \min_k \{U_{n_k}\} \quad (2)$$

where,  $U_{n_k}$  is the upper bound on the length of net  $k$  shared between LBs  $i$  and  $j$ . Perform the force-directed placement and slot assignment.

- The upper bound on a path length is the sum of the upper bounds of the lengths of nets along the path. This represents a path constraint and is given as input to the placement algorithm. Calculate the current length of a path as the sum of lengths of all nets along the path. Check for paths whose time delays and hence path lengths are in excess of the constraints set for them. Define *overshoot* as the difference between the current path length and the upper bound on the path length.
- Find nets along the path whose net lengths exceed the corresponding upper bounds. Let there be  $V$  such nets. Apportion the overshoot among the  $V$  nets by decrementing the respective weights by  $(overshoot/V)$ . Weights of other nets on the path remain same.

- Reevaluate  $W_{ij}$  as the minimum of the weights of all nets shared between LBs  $i$  and  $j$ . This has the effect of increasing the attractive force between the blocks whose nets have delays longer than desired without adversely affecting the other nets.

The first two phases of the algorithm—force-directed placement and slot assignment are repeated with the new set of weights. This is an iterative process which is carried out until there is no significant improvement in the path lengths that violate their respective upper bounds. We refer to each such iteration as a *pass* through the entire algorithm. In the end, limited and local pairwise exchange is carried out to reduce the delays on paths that do not meet the constraints, if any such paths are left. It needs mention here that when weights are decremented, they are not allowed to fall below a certain threshold to ensure that the attractive forces do not become numerically large. Another important point is that since the net weights are always decremented, it does not in any way affect the other paths of which these nets may be a part. This is in contrast to other net-based approaches in which, often, the non-critical paths become critical due to the manner in which net weights are assigned.

### 3. EXPERIMENTAL RESULTS

For the experiments, we tested eight circuits [14] that represent a good mix of industrial circuits and circuits chosen from the MCNC suite of benchmarks.

Table I Circuit Design Details

Ckt no	Design	# LBs	# IOs	# NETS	Array Size	Array Pads
1	top1	54	26	107	8 × 8	68
2	top8	64	54	141	8 × 8	68
3	freqcntr	66	64	156	10 × 8	84
4	f51m	42	16	50	8 × 8	68
5	firtgl	59	31	87	8 × 8	68
6	comp	66	35	98	10 × 10	68
7	9sym	95	10	104	10 × 10	68
8	terml	116	44	150	12 × 12	84

Table I summarizes the circuit details, the first three circuits in this table are industrial circuits and the remainder are MCNC benchmarks. We use the Xilinx 3000 series FPGAs [13] to implement the circuits. The columns on array size and array pads in the table give the description of the FPGA used. We map the circuits onto the FPGA architecture using the technology mapper in the Xilinx Design Environment. The logic blocks and IO blocks so configured are then placed using our placement algorithm. For detail routing, we use the delay-driven router that is a part of the Automatic Place and Route tool (APR version 3.13) available in the Xilinx Design environment. A point that needs mention here is that the placement tool in APR is based on the simulated annealing algorithm while the router uses the rip-up and retry algorithm. Critical nets are taken care of by assigning them higher weights in APR. The code for our algorithm is implemented in the C programming language and experiments are carried out on a Sun SPARC workstation. We assume that the positions of the IO blocks are given and kept fixed during the placement.

We use TFDP to denote results obtained by running our timing-constrained force-directed placement algorithm, FDP to denote results obtained by running the connectivity-driven version of the force-directed placement algorithm with timing constraints ignored, and XAPR to denote results obtained by running the APR tool of Xilinx with timing constraints specified. We now discuss the performance of our algorithm in terms of various metrics.

### 3.1 Performance Evaluation

*Path slack* is measured as the difference between the upper bound on the length (or delay) of a path and its actual length (or delay). It is desirable to have all paths with positive slacks if there are to be no timing violations in the circuit. *Path slack distribution* is a histogram of path slacks and gives an indication of the relative proportion of critical paths that have negative slacks. The histogram of figure 3 shows a typical path slack distribution obtained as a result of timing-driven placement for circuit 8 (term1). In figure 3(a), we compare two distributions corresponding to TFDP and FDP. The horizontal axis corresponds to the various ranges of path slacks (in length units) and the vertical axis gives the number of paths that fall in a particular slack range. It is clear from the distribution that imposition of timing constraints increases the number of paths with positive slacks in the case of TFDP when compared with FDP. In figure 3(b), we compare the distributions after the first pass and after the second pass of TFDP. Weight assignment, as described in section 2.4, precedes the second pass of TFDP. This histogram serves to demonstrate the usefulness and effectiveness of weight assignment. As can be seen in the figure, in the second pass, as a result of reassignment of net weights, there are far fewer nets in the negative slack region. After the third pass, there is no path with negative slack left and the circuit meets the timing constraints. Such a behavior is exhibited by all the test circuits and timing con-

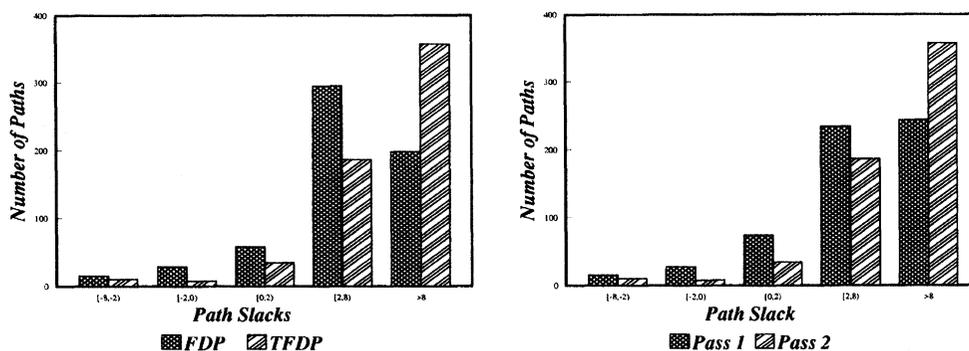


FIGURE 3 Slack Distribution: (A) FDP vs TFDP (B) Pass 1 vs Pass 2.

straints specified for them are met after at most three passes of the algorithm.

*Maximum (worst) delay* is measured as the delay corresponding to the longest path in the circuit after the completion of routing. This measure is important because it is the slowest path in the circuit that dictates how fast the circuit can operate. To evaluate the speed improvement in the circuit as a result of timing-driven placement, the maximum delay of any path in the circuit is shown in Table II. For this metric, we show three sets of values corresponding to those obtained from FDP, TFDP and XAPR. As for the improvement in speed due to imposition of timing constraints, it is clear that TFDP performs better than FDP for all the circuits, the maximum speed improvement being 20.7ns for circuit 1. This reaffirms the usefulness of timing-driven placement. Comparing the maximum delays obtained from TFDP versus those obtained from XAPR, we find that our algorithm exhibits superior performance for all the circuits. The maximum speed improvement is 50.2ns for circuit 6 ! This demonstrates that the delay formulation used in the force equations and the concept of weight assignment is very effective in meeting all the timing constraints as well as in increasing the speed of the circuit.

Having evaluated our algorithm with respect to timing performance, we now discuss the impact of timing-driven placement on metrics related to wireability. *Total Wire Length (TWL)* is an estimate of the sum of wire lengths of all the nets in the circuit. While it is difficult to establish a closed form relationship between the estimated wire length and the routability of a circuit layout, it is well known that

minimizing the total wire length helps to reduce the routing area. A single-trunk steiner tree is used to approximate the length of a net. Since timing-constrained placement is to minimize path delays in the circuit and not the TWL, the latter tends to be larger than in the case when the objective of the algorithm is to minimize TWL. However, it is not desirable to satisfy the delay requirements by compromising a great deal on total wire length. We measure the fractional increase in TWL due to timing-constraint consideration over the TWL obtained when the aim is to minimize TWL. The results for these are summarized in Table 2. The percentage increase in TWL in the case of TFDP is shown in the corresponding column of the table. As expected, TFDP results in an increase in wire length when compared to FDP in most cases. However, there are some interesting exceptions to this, such as circuits 5 and 8, which actually show a decrease in the TWL when compared to FDP ! One of the reasons for this behavior is that our delay formulation favors bringing blocks that share critical nets closer together. Consequently, there may be some critical nets which connect several logic blocks (or have large fanout) and cause all these blocks to be close to each other, thereby reducing the wire length.

A placement that meets all timing constraints but results in great deal of wire congestion is not desirable. Because of the limited interconnection resources of the FPGAs, *routability* assumes greater importance than in other technologies. The number of unrouted nets serves as a yardstick of the routability of the given placement. As can be seen in Table 2, the number of unrouted nets in placements obtained from FDP and TFDP are about the same in most cases.

Table II Experimental results

Ckt	Max Delay (ns)			% TWL Incr. TFDP	Unrouted Nets		CPU Time (sec)		
	FDP	TFDP	XAPR		FDP	TFDP	FDP	TFDP	XAPR
1	73.2	52.5	54.9	0.24	12	11	43	88	152
2	46.6	46.1	69.3	0.019	9	9	65	200	233
3	67.8	50.8	67.2	0.18	0	0	99	456	525
4	55.2	55.1	89	0.15	0	0	29	59	89
5	53.6	53.5	63.7	-0.04	3	3	64	129	149
6	54.3	43.2	93.4	0.33	0	0	63	192	316
7	75.9	74.9	93.4	0.01	3	3	131	160	568
8	91.7	83.9	97.7	-0.016	7	6	297	727	977

While in three of the test circuits, all nets are completely routed, in two circuits TFDP results in an additional unrouted net compared to FDP and in the remaining circuits wireability is not affected. This indicates that wireability degradation is extremely low in TFDP when compared with FDP.

The last important metric deals with the computational resource required for the timing-driven placement algorithm. Since FPGAs are geared towards fast design turnaround times, it is essential to ensure that the timing-driven placement algorithm does not become a computational bottleneck. In Table 2, TFDP requires more CPU time than FDP. This is because of the additional computation required in the weight assignment phase of our algorithm and also due to the additional passes required as a result to ensure that timing constraints are completely satisfied. TFDP takes at least twice as long as FDP as seen in circuit 4. However, our algorithm is faster than simulated annealing based placement algorithm used in XAPR. For circuit 7, TFDP takes 408s less than XAPR and for circuit 5, TFDP takes 20s less than XAPR.

From the above discussion, it is clear that our timing-constrained force-directed algorithm succeeds in attaining the primary objective of meeting the timing constraints specified for the circuit without seriously degrading the wireability of the resulting placement and without incurring excessive computational time. The results also demonstrate that reassigning weights to nets that lie on paths that do not meet timing constraints leads to very good timing results.

#### 4. CONCLUSIONS

In this paper, we have presented a timing-driven placement algorithm for FPGAs that uses forces on the logic blocks to determine their placement. The algorithm employs a novel dynamic weight assignment technique for the nets to ensure that timing constraints for the circuit are met. The algorithm is evaluated comprehensively on eight test circuits and the results demonstrate that the algorithm performs very well in meeting the timing constraints, leads to circuit

designs with higher speeds, and does not adversely affect the wireability or running time. All these factors are extremely important when implementing high-performance circuits using FPGAs.

#### Acknowledgments

We would like to thank the reviewers for the helpful suggestions. This research was partially supported by a Grant MIP-92-22408 from the National Science Foundation.

#### References

- [1] M. Burstein and M. N. Youssef, "Timing influenced layout design," *Proc. of 22nd ACM/IEEE Design Automation Conference*, pp. 124–130, 1985.
- [2] T. Gao and P. M. Vaidya and C. L. Liu, "A new performance driven placement algorithm," *Proc. of 1991 Intl. Conference on Computer Aided Design*, pp. 44–47, 1991.
- [3] W. E. Donath and et al., "Timing driven placement using complete path delays," in *Proc. of 27th ACM/IEEE Design Automation Conference*, pp. 84–89, 1990.
- [4] W. K. Luk, "A fast physical constraint generator for timing driven layout," *Proc. of 28th ACM/IEEE Design Automation Conference*, pp. 626–631, 1991.
- [5] R. S. Tsay and J. Koehl, "An analytic net weighting approach for performance optimization in circuit placement," *Proc. of 28th ACM/IEEE Design Automation Conference*, pp. 620–625, 1991.
- [6] R. S. Tsay and J. Koehl, "An analytic net weighting approach for performance optimization in circuit placement," *Proc. of 28th ACM/IEEE Design Automation Conference*, pp. 620–625, 1991.
- [7] A. Srinivasan, "An algorithm for performance-driven initial placement of small-cell ICs," in *Proc. of 28th ACM/IEEE Design Automation Conference*, pp. 636–639, 1991.
- [8] M. A. B. Jackson and E. S. Kuh, "Performance-driven placement of cell based ic's," *Proc. of 26th ACM/IEEE Design Automation Conference*, pp. 370–375, 1989.
- [9] S. Brown and J. Rose and Z. G. Vranesic, "A detailed router for Field-Programmable Gate Arrays," *IEEE Transactions on Computer-Aided Design*, vol. 11, pp. 620–628 May 1992.
- [10] S. K. Nag and K. Roy, "Iterative Wireability and Performance Improvement for FPGAs," in *Proc. 30th ACM/IEEE Design Automation Conference*, pp. 321–325, June 1993.
- [11] K. J. Antreich and F. M. Johannes and F. H. Kirsch, "A new approach for solving the placement problem using force models," *Proc. Intl. Symposium on Circuits and Systems*, pp. 481–486, 1982.
- [12] N. R. Quinn, Jr. and M. A. Breuer, "A force directed component placement procedure for printed circuit boards," *IEEE Transactions on Circuits and Systems*, pp. 377–388, June 1979.
- [13] Xilinx, *The programmable gate array data book*. Xilinx,

- [14] Srilata Raman, Timing-Constrained Layout Algorithms for Symmetrical FPGAs, Ph.D. Thesis, Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, Illinois. Inc., 1989.

### Authors' Biographies

Srilata Raman is a Research Staff Member in the Advanced Design Technology group in Motorola, Austin. She received the Ph.D. degree from the Department of Electrical and Computer Engineering at the University of Illinois at Urbana-Champaign, Urbana, Illinois. Her research interests include VLSI-CAD, Optimization algorithms, parallel processing for CAD, and incremental design techniques. She is a member of IEEE, Computer Society, and ACM.

C. L. Liu obtained his B.Sc. degree of Cheng Kung University in 1960 and his Sc.D. degree in 1962, all in Electrical Engineering at the Massachusetts Institute of Technology. He is currently a Professor of Computer Science at the University of Illinois at Ur-

bana-Champaign. His areas of research interest are Design and Analysis of Algorithms, Computer Aided Design of Integrated Circuits, and Combinatorial Mathematics.

Dr. Larry G. Jones received the BS and Ph.D. degrees in Computer Science from the Pennsylvania State University in 1978 and 1986, respectively. From 1978 to 1980 he worked for Data General as a systems programmer designing and implementing programming language compilers and compiler generators. On completion of his Ph.D. he joined the University of Illinois at Urbana-Champaign where he was an Assistant Professor of Computer Science. In 1992, he joined Motorola in Austin, Texas where he is currently managing the High Density Design Technology Group. He has publications in the areas of computer-aided design of integrated circuits, programming language implementation, and active-value database systems.



**Hindawi**

Submit your manuscripts at  
<http://www.hindawi.com>

