

OPTIMALITY AND STABILITY STUDY OF TIMING-DRIVEN PLACEMENT ALGORITHMS

Jason Cong, Michail Romesis, Min Xie

Computer Science Department
University of California, Los Angeles
{cong,michail,xie}@cs.ucla.edu
Technical Report #030030

ABSTRACT

This work studies the optimality and stability of timing driven placement algorithms. To our knowledge, it is the first study of this kind. The contributions of this work include two parts: 1) We develop an algorithm for generating synthetic examples with known optimal delay for timing driven placement (T-PEKO). The examples generated by our algorithm can closely match the characteristics of real circuits. 2) Using these synthetic examples with known optimal solutions, we studied the optimality of several timing-driven placement algorithms for FPGAs by comparing their solutions with the optimal solutions, and their stability by varying the number of longest paths in the examples. For examples with a single longest path, the delay produced by the algorithms is from 10% to 18% longer than the optima on the average, and from 34% to 53% longer in the worst case. Furthermore, their solution quality deteriorates as the number of longest paths increases. For examples with more than 5 longest paths, their delay is from 23% to 35% longer than the optima on the average, and is from 41% to 48% longer in the worst case.

1. INTRODUCTION

Placement is one of the most important steps in the post-RTL synthesis process as it directly defines the interconnects, which have now become the bottleneck in circuit and system performance in DSM technologies. The placement problem has been studied extensively in the past 30 years. However, a recent study shows that existing placement solutions are surprisingly far from optimal. Using a set of constructed placement examples that match many industrial circuit characteristics with known optimal wirelength (PEKO), the study shows that the results of leading placement tools from both industry and academia are 70% to 150% away from the optimal solutions on those examples [1]. An extension of PEKO was presented [2], where new examples called PEKU (Placement Examples with Known Upper bounds) were created by inserting a user-specified percentage of non-local nets into a PEKO circuit. By relaxing the optimality constraint on a subset of connections, PEKU more accurately emulates real circuits in terms of wirelength distribution. Experiments showed that for PEKU benchmarks, state-of-the-art placers can be far away from the upper bound. In the extreme case, where each circuit consists of global connections only (G-PEKU benchmarks), existing tools can be 41% to 102% away in the worst case.

The examples with known optimal wirelength enable quantitative analysis about how far away state-of-the-art placement algorithms are from the optimal solutions. However, wirelength is not

the sole objective in circuit placement. In the era of DSM technology, the most important goal of placement is performance (delay) optimization. It is important that we extend the optimality study to timing-driven placement algorithms.

Existing timing-driven placement algorithms can be divided into two categories, net-based and path-based. Path-based algorithms [3, 4, 5] try to directly minimize the longest path delay. Since they maintain an accurate timing view during optimization, their complexity is usually high. Net-based algorithms [6, 7, 8, 9] first transform timing constraints into either length constraints or weights on individual nets. The information is then fed to a weighted wirelength minimization based placement engine to obtain a new placement with better timing. This process usually goes through a few iterations until no improvement can be made, or a certain iteration limit has been reached. Compared with path-based algorithms, net-based algorithms have lower complexity.

There are several works on generating timing-driven placement examples [10, 11]. However, none of them satisfy our need, since their optimal solutions are unknown. In this paper, we present an algorithm for generating timing-driven placement examples with known optimal delay under a simplified delay model (T-PEKO). These examples can closely match the characteristics of real circuits. Using these examples with known optimal delays, we studied the optimality of several timing-driven placement algorithms for FPGAs from commercial and academic tools by comparing their solutions to the optimal solutions, and their stability by varying the number of longest paths in the examples. We chose FPGA placement since it gives the flexibility to specify our delay model and cell library. To our knowledge, this is the first study of this kind. Experimental results for the academic tools show that for examples with a single longest path, the delay produced by the algorithms is from 10% to 18% longer than the optima on average, and from 34% to 53% longer in the worst case. Furthermore, their solution quality deteriorates as the number of longest paths increases. For examples with more than 5 longest paths, their delay is from 23% to 35% longer than the optima on average, and is from 41% to 48% longer in the worst case. The performance of the commercial tool that targets the Xilinx Virtex architecture is much better. The difference there from the optima in terms of delay ranges from 4% to 21%.

The rest of this paper is organized as follows: Section 2 presents the T-PEKO algorithm for the construction of timing-driven placement examples with a known optimal solution. Section 3 presents the comparison of the placement results for the T-PEKO suite produced by state-of-the-art, timing-driven placement algorithms with the optimal solutions. Section 4 presents conclusions and future




Figure 1: Graph of a basic logic element. It consists of a lookup table (LUT), a flip-flop and a multiplexer

work.

2. CONSTRUCTION OF PLACEMENT EXAMPLES WITH KNOWN TIMING OPTIMAL SOLUTION

2.1. Discussion of the FPGA Architecture and the Delay Model

We first present the architecture of the FPGA device that we assume during the construction of the examples. Each logic block (CLB) consists of two basic logic elements (BLE). One BLE is shown in Figure 1, and it contains a K -input LUT, a flip-flop and a multiplexer. The flip-flop's input is connected to the output of the LUT. The multiplexer selects the output of the LUT or the flip-flop.

Several delay models have been proposed to calculate the performance of a circuit. The most popular is the Elmore delay model [12]. Recent studies [13] have shown that under optimal buffer insertion, sizing and wire sizing, the delay of a wire is approximately linear to its length. For this reason, in this paper we use a linear delay model which can be summarized as follows:

- (i) The delay inside any LUT is a constant d_g , while any other delay inside a BLE is assumed to be zero.
- (ii) The delay of any interconnect $d_i(A, B)$ between two BLEs A and B is given by the following formula: $d_i(A, B) = \text{dist}(A, B) * d_u$, where $\text{dist}(A, B)$ is the Manhattan distance of BLE A from BLE B , while d_u is the constant delay between two adjacent BLEs (Manhattan distance equal to 1).

In reality, the BLEs of FPGA devices are more complex and include more connections, as will be shown in the Xilinx experiment section. The delay model also can be more complicated. However, our methodology is generic in that it is applicable as long as the interconnect delay between two adjacent nodes is always smaller than any delays between non-adjacent nodes and d_g, d_u are constants. Furthermore, it can be applied to ASICs as well, especially to standard-cell row-based architectures.

2.2. The T-PEKO Algorithm

Our methodology for the construction of the timing-optimal benchmarks works as follows: The first step is to obtain a placement solution of an existing combinational or sequential circuit. The second step is to perform timing analysis to find the longest path in the circuit using our delay model. Let d_k be the delay of the longest path, and x, y be the number of rows and columns of the device, respectively. The algorithm perturbs the netlist by inserting a path p_{opt} that connects $r + 1$ adjacent nodes, where r is given by $r = \max(\lceil d_k / (d_g + d_u) \rceil, \lceil ((x + y) * d_u + d_g) / (d_g + d_u) \rceil)$ (see Figure 2). Since the new netlist is the result of a perturbation of the original netlist, the smaller the perturbation, the stronger the similarities it has with the original circuit.

Before we present in detail the construction of the path p_{opt} , we denote some terms here:

- We call a netlist *valid* if: 1) It has no combinational loops, 2) It has no dangling BLEs, i.e., BLEs with at least one input (output) and no outputs (inputs), and 3) Each BLE has at most K inputs and 1 output. If a netlist is not valid, it is called *invalid*.

- *Static timing analysis* [14] constructs a *timing graph* whose vertices correspond to the pins of the circuit. The *timing edges* that connect the vertices of this graph are constructed in two ways: 1) Each net is converted into a set of directed edges that connect each source of the net to all sinks of the net. 2) Each LUT is represented by a set of *intracellular* edges that connect all the inputs of the LUT to its output.

- Each LUT is assigned a number called the *level* of the LUT, such that the following property is satisfied:

Property (1): For every timing edge of the timing graph of the circuit originating from the output pin of an LUT a to an input pin of another LUT b , we have $\text{Level}(a) < \text{Level}(b)$.

It is easy to see that if this property is satisfied, there are no combinational loops in the circuit. If some timing edges violate the above property, we can guarantee that by removing them the circuit is free of combinational loops. Note that this property does not have to hold for timing edges between pins of the same LUT, or between pins of a LUT and a flip-flop.

- A flip-flop f is *disconnected* if the multiplexer of the BLE selects the output of the LUT l , otherwise it is *connected*. In the remainder of this paper, when we mention that the status of a flip-flop is changed from connected to disconnected, we perform the following changes to the netlist: The flip-flop is removed from the netlist and all the fanouts of the flip-flop become fanouts of the LUT except for the ones that cause a violation of Property (1). Similarly, when we mention that the status of a flip-flop is changed from disconnected to connected the following changes are performed: A new net is added to the netlist from the output of the LUT to the input of the flip-flop, and all the previous fanouts of the LUT become fanouts of the flip-flop.

Before the construction of the path p_{opt} these initial steps take place:

- (i) The original mapped netlist is placed on the FPGA device.
- (ii) Static timing analysis is performed on the placed circuit. The longest path delay d_k is computed as well as the integer r according to the formula :

$$r = \max(\lceil d_k / (d_g + d_u) \rceil, \lceil ((x + y) * d_u + d_g) / (d_g + d_u) \rceil).$$
Every LUT is assigned to a level equal to the highest arrival time among its pins.

The construction of the path p_{opt} is as follows:

- (i) A BLE is selected at a corner of the device as the first node of the path. If the flip-flop of the BLE is disconnected its status is changed to connected. A new timing edge (if it does not exist already) is added¹ from the output pin of that flip-flop to an input pin of an adjacent LUT, which we

¹The insertion of a timing edge on the timing graph from the output pin s of LUT a to an input pin t of LUT b corresponds to the following changes on the netlist: If pin s is used and n is the corresponding net, add pin t to the sinks of n . If s is not used, create a new 2-pin net with s as its source and t as its sink.

easy to see that this additional connection does not violate Property (1).

- (iii) *A flip-flop becomes connected*: This modification does not change a valid netlist to invalid because: 1) There is no BLE with more inputs or outputs than before, 2) No dangling nodes are created, 3) No combinational loops are created.
- (iv) *A flip-flop becomes disconnected*: When a flip-flop becomes disconnected, its fanouts become fanouts of the corresponding LUT, except for the ones that cause a violation of Property (1). If any dangling nodes are created, they are connected to an IO pad or a BLE. It is easy to see that no other violations occur, so a valid netlist remains valid in this case.

Assuming that the original netlist is valid, and since all the intermediate changes do not change a valid netlist to invalid, the final netlist created by the T-PEKO algorithm is valid.

After the construction of p_{opt} , the algorithm iteratively performs timing analysis to identify paths that are longer than p_{opt} . The interconnect timing edges of these paths that do not belong to p_{opt} are removed. In this process it is possible to insert new timing edges, if dangling nodes are created. At the same time however, the corresponding flip-flop of the dangling node gets connected. As a result, the maximum possible delay of any path that includes these new edges is: $d_{max} = (x + y) * d_u + d_g$, since $(x + y) * d_u$ is the maximum possible delay of a net according to our timing model. Since at every iteration we insert edges only for the dangling nodes and we remove edges that are not in the p_{opt} , eventually the longest path of the circuit will be either p_{opt} , or another path that includes the new edges with a maximum delay of d_{max} . But we have that the delay TD of p_{opt} is $r * (d_g + d_u)$ and that $r \geq \lceil ((x + y) * d_u + d_g) / (d_u + d_g) \rceil$, so we have that $TD \geq d_{max}$. This proves that after a finite number of iterations, p_{opt} becomes the longest path of the circuit.

2.3. Increasing the Difficulty of the T-PEKO examples

In this work we study not only the optimality of the timing-driven placement algorithms, but also their stability for circuits with different characteristics. Ideally an algorithm is expected to perform well on various kinds of circuits. For this stability study we introduce two parameters for the construction of the circuits that control their difficulty for a placer, including:

- (i) *The number of longest paths*: The algorithm can create a user-specified number of disjoint longest paths. Assume that this number is M , that the delay of the critical path of the original circuit is d_t , and that integer r is computed as before. We create a path p_{opt} according to the same methodology as described earlier with the only difference that it connects $M * r + 1$ adjacent BLEs, instead of $r + 1$. The total delay of that path will be $TD = M * r * (d_g + d_u) \geq M * d_t$. Along this path at equal distances, we connect $M - 1$ flip-flops. As a result, the initial longest path is replaced by M paths, each one with a delay greater than or equal to d_t . Note that after this change, some other paths in the circuit might become longer, and they will be removed according to the same procedure as the one described in the previous subsection. In the end, these M paths will be the longest of the circuit.
- (ii) *The number of edges that connect longest paths*: To increase the degree of path sharing, T-PEKO will create some




Figure 3: Bridge construction. A bridge will be inserted between s and t .

nets to connect BLEs located on different longest paths. Figure 3 provides an example. AB and CD are two longest paths constructed as described in the previous paragraph. E is a BLE along the path AB , F is a BLE along the path CD . T-PEKO will connect E 's output s with one of F 's unused inputs t . This corresponds to inserting a timing edge between s and t in the timing graph. We call the newly added timing edge a *bridge*, denoted as $b(s, t)$. The following theorem guarantees that the netlist remains valid after this operation.

*Theorem 2: The netlist after inserting $b(s, t)$ is valid if $d(s) + dist(E, F) * d_u + d_g \leq d(t')$. Here, $d(s)$ is the arrival time of s , t' is the output pin of F , $d(t')$ is the arrival time of t' , and $dist(E, F)$ is the Manhattan distance of E from F . The longest paths in the original netlist remain the longest after $b(s, t)$ is inserted.*

Proof: It is obvious that the IO constraints on the BLEs will still be satisfied. If the netlist has combinational loops after the insertion of $b(s, t)$, $b(s, t)$ should be on a loop in the timing graph. Therefore, there exists a path from t' to s before inserting $b(s, t)$, indicating $d(t') < d(s)$. Contradiction.

If a path exists with a delay longer than TD after inserting $b(s, t)$, $b(s, t)$ should be on this path. However, since $d(s) + dist(E, F) * d_u + d_g \leq d(t')$, $b(s, t)$ will not increase the arrival time of any vertex in the timing graph. Therefore, a vertex exists in the timing graph with a delay higher than TD before inserting $b(s, t)$. Contradiction.

2.4. Extension to the Xilinx Architecture

The previously described algorithm targets our simplified model and FPGA architecture. With some modifications, T-PEKO is extended to create placement examples constructed for commercial tools. More specifically, in this subsection we describe how we created examples for the Xilinx Virtex architecture. In this architecture a CLB (configurable logic block) contains two logic cells, and each cell contains 2 LUTs (see Figure 4). Due to the interconnect architecture of Virtex [15], it is not guaranteed that the interconnection delay between non-adjacent nodes is shorter than the delay between two adjacent nodes.

To address this problem the artificial path we constructed for this architecture was slightly different from the general case of the previous section. The path must first visit all four nodes of a CLB before moving to an adjacent CLB. Figure 5 shows an example of two artificial paths that we created for a Xilinx Virtex device. These paths share the same CLBs in the middle row, but the first




Figure 4: A Virtex CLB contains 4 LUTs in 2 slices. Picture taken from the web site of Xilinx




Figure 5: An example of the artificial path on a Xilinx Virtex device. A box represents an LUT, while the dashed lines show the borders between different CLBs. The path traverses all the LUTs of a CLB, before moving to an adjacent CLB.

path moves to the CLBs of the upper row, while the other path moves to the CLBs of the bottom row. For our Xilinx experiments, we used this technique to create multiple paths. Similar changes must be performed when working on other FPGA architectures.

One additional problem when working with commercial tools is that the delay model is no longer known. It is true that delay tables can be extracted³, but they are not 100% accurate. Therefore, the timing analysis we perform is an approximation. It is not guaranteed that the artificial path is the longest path in the circuit. Still, we can consider the delay of that path as a tight upper bound of the optimal delay of the circuit.

In the experimental results section we are investigating the performance of the Xilinx place and route tool PAR on the T-PEKO examples tailored for the Virtex architecture.

3. EXPERIMENTAL RESULTS

We have implemented T-PEKO on a Sun Blade 1000 using C++. To generate the initial placement configurations needed by T-PEKO, we ran VPR [16] on 20 MCNC benchmarks using its timing driven

³We extracted the delay tables in the Virtex architecture as follows: We constructed nets connecting 2 LUTs. One LUT was fixed at a corner of the chip. The other was moved to every location on the chip. We filled the delay table with the delays reported by the Xilinx timing analysis tool in every case.

mode. The placement results were then fed into T-PEKO and perturbed. We varied M from 1 to 5 and generated 100 circuits. The maximum number of inputs and outputs on each BLE is 6 and 1 respectively. As for the delay parameters, $d_g = 1$ and $d_u = 1$. When possible, a maximum of 50 bridges were inserted between the M longest paths. The circuits were grouped as the T-PEKO suite and can be downloaded from [17].

Table 1 gives the characteristics of T-PEKO in terms of the number of CLBs, PIs, POs, flip-flops and nets. The column “Orig” shows the name of the original MCNC circuit from which the initial placement configuration is derived. The columns for $M = 0$ show the characteristics of the original MCNC circuits. Column “Opt” gives the optimal delay under our simplified delay model. For the same initial placement configuration, the optimal delay does not change for any value of $M > 0$ (of course, we do not know the optimal delay for $M = 0$). The perturbed circuits are very close to the original ones in these aspects for most cases. The circuits that were initially combinational were transformed into sequential after the insertion of flip-flops (40 in the worst case). The circuits are given in the format specified in [18]. Each circuit has a .net file describing the netlist of each circuit. It also has a .arch file specifying the combinational delay of each LUT, and the number of IOs for each CLB. To guarantee a fair comparison, we generated a .pad file for each circuit, which gives the pad locations extracted from the optimal solutions by our construction.

For our optimality and stability study, we experimented with two state-of-the-art FPGA placement algorithms, including:

- VPR [16], a well-known FPGA placement and routing package widely used for FPGA architecture evaluation [19]. Its optimization engine is based on simulated annealing. It combines connection-based and path-based timing-analysis. The cost function it uses trades off between wirelength and critical path delay. We used VPR v.4.3 downloaded from [20] in our experiment.
- PATH [21], the latest FPGA placement algorithm which presents a significant enhancement to VPR in timing optimization. It takes into consideration the path sharing effect. PATH introduces a new net weighting algorithm based on the concept of path-counting. We used PATH v.1.0 in our experiment.

One complication of FPGA architecture is that the delay between two BLEs depends not only on their Manhattan distance, but also the routing segments that connect the BLEs. Therefore, both algorithms use a preliminary routing procedure before placement to determine the delay between BLEs. To accommodate our simplified delay model, we modified the delay computation in each algorithm, so that the delay between BLEs is always the Manhattan distance between them multiplied by d_u . This change, in effect, makes our study of these algorithms independent of the FPGA architecture and their routing procedure. In our experiment, we set the tradeoff parameter of wirelength vs. delay to be 0.5, as suggested by [16]. Changing the value of this parameter to favor the critical path delay minimization did not seem to improve the final results.

For each circuit of T-PEKO, we run each algorithm 5 times. The results are summarized in Table 2 and Figure 6. The average difference between each algorithm’s result and the optimal solution is listed. For completeness, the best results for every circuit are reported. From the results, we make the following observa-

Table 1: Characteristics of the TPeko suite. Column “Orig” gives the initial circuit from which the perturbed circuits are derived. $M = 0$ corresponds to the characteristics of the original circuit. The perturbed circuits are very close to the original circuits in the number of CLBs, PIs, POs, flip-flops and nets. Column “Opt” gives the optimal delay for each circuit. It is the same for circuits derived from the same initial placement for $M \geq 1$.

Ckt	Orig	Opt	M = 0					M = 1					M = 3					M = 5				
			CLB	PI	PO	FF	NET	CLB	PI	PO	FF	NET	CLB	PI	PO	FF	NET	CLB	PI	PO	FF	NET
TPeko01	tseng	88	1047	51	122	385	1098	1056	51	105	371	1107	1059	51	85	341	1128	1060	51	81	326	1162
TPeko02	ex5p	108	1064	8	63	0	1072	1067	8	57	5	1075	1068	8	56	6	1102	1068	8	50	32	1127
TPeko03	apex4	106	1261	9	19	0	1271	1275	9	22	25	1284	1287	9	17	24	1310	1287	9	18	38	1347
TPeko04	dsip	116	1370	228	197	224	1598	1424	228	192	226	1652	1537	228	189	228	1789	1653	228	189	230	1932
TPeko05	misex3	92	1397	14	14	0	1411	1415	14	34	40	1429	1427	14	15	19	1461	1431	14	11	23	1496
TPeko06	diffeq	84	1497	63	39	377	1560	1502	63	26	366	1565	1511	63	19	354	1596	1512	63	18	343	1626
TPeko07	alu4	100	1522	14	8	0	1536	1535	14	8	2	1549	1546	14	8	6	1579	1561	14	8	11	1626
TPeko08	des	132	1591	256	245	0	1847	1604	255	206	10	1859	1686	254	194	16	1971	1772	254	191	17	2077
TPeko09	bigkey	124	1707	228	197	224	1935	1766	228	197	226	1994	1793	224	188	239	2036	1801	223	166	227	2075
TPeko10	seq	122	1750	41	35	0	1791	1754	41	40	18	1795	1756	41	38	27	1814	1756	41	28	24	1848
TPeko11	apex2	128	1878	38	3	0	1916	1894	38	8	13	1932	1912	38	5	17	1977	1913	38	8	25	2002
TPeko12	s298	166	1931	3	6	8	1934	1934	3	6	11	1937	1934	3	6	15	1986	1934	3	6	42	1988
TPeko13	frisc	170	3556	19	116	886	3575	3568	19	113	896	3587	3576	19	109	893	3628	3578	19	105	876	3648
TPeko14	elliptic	146	3604	130	114	1122	3734	3616	130	81	1099	3746	3628	130	68	1054	3789	3638	130	63	1027	3819
TPeko15	spla	184	3690	16	46	0	3706	3698	16	50	14	3714	3706	16	47	25	3773	3706	16	53	35	3773
TPeko16	pdcc	240	4575	16	40	0	4591	4595	16	55	30	4611	4607	16	48	35	4667	4607	16	39	27	4674
TPeko17	ex1010	290	4598	10	10	0	4608	4610	10	21	22	4620	4612	10	10	14	4673	4612	10	10	21	4673
TPeko18	s38417	164	6406	28	106	1463	6434	6417	28	96	1477	6445	6434	28	94	1485	6501	6441	28	88	1435	6520
TPeko19	s38584	164	6435	37	304	1260	6484	6457	37	262	1250	6494	6476	37	236	1236	6552	6487	37	233	1211	6575
TPeko20	clma	328	8382	61	82	33	8444	8393	60	60	128	8453	8402	60	57	117	8513	8408	60	53	120	8519

tions:

- For $M = 1$, the delay produced by the algorithms is from 10% to 18% longer than the optima of T-PEKO on average, and from 34% to 53% longer in the worst case.
- The solution quality of both algorithms deteriorates as M increases. For $M = 5$, the gap between their solutions and the optima is from 23% to 35% on average, and from 41% to 48% in the worst case.
- PATH outperforms VPR in all cases. The best results from PATH are on average 4% worse than the optima when $M = 1$, and 18% worse when $M = 5$.

Figure 7 shows the optimal configuration of TPeko20 with $M = 5$ and the results generated by both VPR and PATH. The nodes on the longest paths by our construction are colored in black in each solution. Furthermore, the critical timing edges in each solution are also colored in black. It can be seen that these nodes are indeed on the longest paths of both VPR and PATH’s results. However, the delay produced by both algorithms is far away from the optimal. Note that besides the longest path created by T-PEKO, there exist some other paths with the same delay, that include nets from the original circuit. Figure 7 shows several such paths in the optimal solution.

Using the method described in the previous section, we extended our study to the Xilinx placement engine, PAR, and constructed 15 synthetic circuits for it. The version we experimented is Release 5.1.03i - PAR F.26. To guarantee that PAR can find the minimum possible delay in this experiment, we set a loose delay constraint at the beginning and gradually tighten it until PAR can




Figure 6: Divergence vs M . The divergence from the optima is increasing with M .

no longer find a solution satisfying this constraint. For comparison, we used PAR to do routing on our constructed solutions and quoted the delay reported by its timing analysis tool. This value served as an upper bound to the optimal delay for our constructed circuit.

Table 3 gives the experimental results on these circuits. The first few columns give the circuit characteristics. The upper bound of the optimal delay by our construction is given in the column “UB,” the result by PAR is given in the column “PAR.” On average, the delay generated by PAR is 8.7% worse than our constructed solutions, and is 21.4% worse in the worst case. Com-




Figure 7: Three solutions for TPeko20. The nodes on the longest paths by our construction are colored in black. The timing edges on critical paths in each solution are colored in black, too. It can be seen these nodes are indeed on the longest paths in both VPR and PATH's results. However, the delay produced by both algorithms are far away from the optima. Note that besides the longest paths created by T-PEKO, there exist other paths with the same delay, that include nets from the original circuit. Several of them are shown in the optimal solution.

Table 2: Experimental results by VPR and PATH on the TPEKO suite. M corresponds to the number of initial longest paths. Average and minimum divergence from the optima by VPR and PATH is listed.

Circuit	Opt	M = 1				M = 3				M = 5			
		VPR		PATH		VPR		PATH		VPR		PATH	
		Avg	Best	Avg	Best	Avg	Best	Avg	Best	Avg	Best	Avg	Best
TPeko01	88	4%	0%	7%	0%	21%	17%	17%	15%	35%	27%	25%	20%
TPeko02	108	6%	1%	4%	1%	23%	15%	17%	13%	38%	31%	16%	14%
TPeko03	106	12%	4%	6%	0%	30%	23%	12%	8%	27%	24%	15%	10%
TPeko04	116	7%	0%	0%	0%	10%	7%	5%	3%	18%	14%	15%	9%
TPeko05	92	22%	7%	10%	0%	34%	25%	15%	5%	34%	24%	16%	11%
TPeko06	84	10%	5%	5%	1%	26%	20%	19%	8%	34%	20%	17%	14%
TPeko07	100	22%	10%	4%	0%	25%	20%	11%	7%	33%	27%	17%	13%
TPeko08	132	53%	44%	18%	7%	52%	47%	30%	24%	40%	20%	33%	27%
TPeko09	124	9%	2%	0%	0%	19%	14%	24%	15%	29%	28%	21%	18%
TPeko10	122	11%	6%	7%	1%	27%	25%	18%	12%	33%	29%	15%	13%
TPeko11	128	15%	10%	8%	0%	26%	19%	10%	6%	40%	33%	15%	12%
TPeko12	166	11%	4%	4%	1%	40%	16%	13%	11%	39%	15%	16%	14%
TPeko13	170	31%	25%	15%	8%	39%	22%	35%	19%	35%	29%	32%	22%
TPeko14	146	12%	8%	2%	1%	26%	21%	21%	12%	32%	26%	31%	25%
TPeko15	184	15%	11%	7%	1%	25%	22%	17%	10%	36%	33%	23%	18%
TPeko16	240	17%	7%	13%	3%	61%	55%	14%	12%	36%	25%	19%	13%
TPeko17	290	24%	17%	10%	2%	24%	15%	17%	12%	30%	21%	26%	21%
TPeko18	164	32%	15%	19%	5%	33%	15%	30%	20%	29%	22%	30%	16%
TPeko19	164	17%	13%	34%	26%	46%	25%	45%	38%	48%	41%	41%	37%
TPeko20	328	32%	25%	24%	14%	49%	37%	25%	19%	47%	30%	38%	31%
Avg.		18%	11%	10%	4%	32%	23%	20%	13%	35%	26%	23%	18%

pared with our experiment with VPR and PATH, the divergence here is much smaller. One possible reason is that the delay between two elements on a Virtex chip is not monotone with regard to their Manhattan distance and the PAR takes full advantage of the Virtex routing architecture to reduce the delay of long connections. Still, there is room for improvement that is not negligible.

4. CONCLUSIONS AND FUTURE WORK

This work studied the optimality and stability of timing-driven placement algorithms. We developed an algorithm for generating

Table 3: Experimental results on Xilinx PAR.

Circuit	Orig	chip	package	IOB	Slice	Net	UB (ns)	PAR (ns)	diff
TPeko01	ex5p	xcv50	bg256	73	573	1130	70.0	80.4	15.0%
TPeko02	tseng	xcv200	fg456	176	582	1182	78.2	90.5	15.7%
TPeko03	apex4	xcv50	bg256	33	669	1302	62.7	67.5	7.7%
TPeko04	misex3	xcv50	bg256	31	760	1415	54.7	56.9	4.1%
TPeko05	alu4	xcv50	bg256	27	766	1537	49.0	53.5	9.2%
TPeko06	dsip	xcv600	fg680	435	832	1690	78.2	95.0	21.4%
TPeko07	seq	xcv100	bg256	79	941	1832	64.6	69.9	8.3%
TPeko08	apex2	xcv100	bg256	49	968	1940	60.7	64.4	6.0%
TPeko09	s298	xcv100	bg256	16	1199	1964	64.8	67.9	4.7%
TPeko10	frisc	xcv200	fg456	152	1824	3562	57.4	61.3	6.9%
TPeko11	spla	xcv200	fg456	62	1961	3731	61.8	65.6	6.2%
TPeko12	elliptic	xcv200	fg456	248	1985	3766	62.7	69.4	10.6%
TPeko13	pdc	xcv200	fg456	59	2350	4593	50.5	53.5	6.0%
TPeko14	ex1010	xcv200	fg456	29	2350	4614	53.1	55.6	4.6%
TPeko15	clma	xcv600	fg680	130	4704	8463	64.3	67.3	4.7%
Avg.									8.7%

synthetic examples with known optimal delay for timing driven placement (T-PEKO). The synthetic examples generated by our algorithm can closely match the characteristics of real circuits. Using these synthetic examples with known optimal solutions, we studied the optimality of several timing-driven placement algorithms by comparing their solutions to the optimal solutions, and their stability by varying the number of longest paths in the examples. The results produced by the algorithms could be as far as 54% away from the optimal for our most difficult examples. The results seem to suggest that timing-driven placement algorithms, both net-based and path-based, have much room for improvement. Similar experiments for commercial FPGA architectures showed a smaller, but not negligible gap.

Future work includes the generation of similar placement examples that study the performance of placement algorithms for other objectives such as routability and power on both ASIC and FPGA designs.

5. ACKNOWLEDGEMENTS

This work is partially supported by the Semiconductor Research Corporation under Contract 98-TJ-686, partially supported by the National Science Foundation under Grant CCR0096383, and partially supported by DARPA/GSRC under contract number SA2211-23106. The authors would like to thank Dr. T. Kong for providing the latest version of PATH for our experiments. They would also like to thank Dr. R. Jayaraman for his valuable suggestions regarding the Xilinx experiments.

6. REFERENCES

- [1] C. Chang, J. Cong, and M. Xie, "Optimality and scalability study of existing placement algorithms," in *Proc. Asia South Pacific Design Automation Conference*, pp. 621 – 627, 2003.
- [2] J. Cong, M. Romesis, and M. Xie, "Optimality, scalability and stability study of partitioning and placement algorithms," in *Proc. International Symposium on Physical Design*, pp. 88 – 94, 2003.
- [3] M. Jackson and E. S. Kuh, "Performance-driven placement of cell based IC's," in *Proc. Design Automation Conf*, pp. 370–375, 1989.
- [4] A. Srinivasan, K. Chaudhary, and E. S. Kuh, "RITUAL: A performance driven placement for small-cell ICs," in *Proc. Int. Conf. on Computer Aided Design*, pp. 48–51, 1991.
- [5] T. Hamada, C. K. Cheng, and P. M. Chau, "Prime: a timing-driven placement tool using a piecewise linear resistive network approach," in *Proc. Design Automation Conf*, pp. 531–536, 1993.
- [6] A. E. Dunlop, V. D. Agrawal, D. N. Deutsch, M. F. Jukl, P. Kozak, and M. Wiesel, "Chip layout optimization using critical path weighting," in *Proc. Design Automation Conf*, pp. 133–136, 1984.
- [7] R. Nair, C. L. Berman, P. Hauge, and E. J. Yoffa, "Generation of performance constraints for layout," *IEEE Trans. on Computer-Aided Design*, vol. 8, no. 8, pp. 860–874, 1989.
- [8] R. S. Tsay and J. Koehl, "An analytic net weighting approach for performance optimization in circuit placement," in *Proc. Design Automation Conf*, pp. 620–625, 1991.
- [9] H. Eisenmann and F. M. Johannes, "Generic global placement and floorplanning," in *Proc. Design Automation Conf*, pp. 269–274, 1998.
- [10] M. Hutton, J. P. Grossman, J. Rose, and D. Corneil, "Characterization and parameterized random generation of digital circuits," in *Proc. Design Automation Conf*, pp. 94–99, ACM Press, 1996.
- [11] P. Verplaetse, D. Stroobandt, and J. Van Campenhout, "Synthetic benchmark circuits for timing-driven physical design applications," in *Proc. International Conference on VLSI*, pp. 31–37, CSREA Press, 2002.
- [12] W. C. Elmore, "The Transient Response of Damped Linear Networks," *Journal of Applied Physics*, vol. 19, pp. 55 –63, 1948.
- [13] J. Cong, and D. Pan, "Interconnect delay estimation models for synthesis and design planning," in *Asia Pacific Design Automation Conference*, pp. 97–100, 1999.
- [14] R. Hitchcock, G. Smith, and D. Cheng, "Timing Analysis of Computer Hardware," *IBM J. Res. Develop.*, vol. 26, pp. 100–108, 1982.
- [15] Xilinx Inc., *Virtex 2.5V FPGA Complete Data Sheet (all four Modules)*.
- [16] A. Marquardt, V. Betz, and J. Rose, "Timing-driven placement for FPGAs," in *Proc. of the ACM/SIGDA international symposium on Field programmable gate arrays*, pp. 203–213, ACM Press, 2000.
- [17] "<http://cadlab.cs.ucla.edu/~pubbench/tpeko.htm>,"
- [18] V. Betz and J. Rose, "VPR: A New Packing, Placement and Routing Tool for FPGA Research," in *Proc. of Seventh International Workshop on Field-Programmable Logic and Applications*, pp. 213–222, 1997.
- [19] V. Betz, J. Rose, and A. Marquardt, *Architecture and CAD for Deep-Submicron FPGAs*. Kluwer Academic Publishers, 1999.
- [20] "<http://www.eecg.toronto.edu/~vaghn/vpr/vpr.html>,"
- [21] T. Kong, "A novel net weighting algorithm for timing-driven placement," in *Proc. Intl. Conf. on Computer-Aided Design*, pp. 172 – 176, 2002.