

Parallel FPGA Router using Sub-Gradient method and Steiner tree

Rohit Agrawal, Chin Hau Hoo, Kapil Ahuja, and Akash Kumar

Abstract—In the FPGA (Field Programmable Gate Arrays) design flow, one of the most time-consuming step is the routing of nets. Therefore, there is a need to accelerate it. In [2], the authors have developed a Linear Programming (LP) based framework that parallelizes this routing process to achieve significant speedups (the resulting algorithm is termed as ParaLaR). However, this approach has certain weaknesses. Namely, the constraints violation by the solution and a local minima that could be improved. We address these two issues here.

In this paper, we use the LP framework of [2] and solve it using the Primal-Dual sub-gradient method that better exploits the problem properties. We also propose a better way to update the size of the step taken by this iterative algorithm. We perform experiments on a set of standard benchmarks, where we show that our algorithm outperforms the standard existing algorithms (VPR [1] and ParaLaR).

We achieve upto 22% improvement in the constraints violation and the standard metric of the minimum channel width when compared with ParaLaR (which is same as in VPR). We achieve about 20% savings in another standard metric of the total wire length (when compared with VPR), which is the same as for ParaLaR. Hence, our algorithm achieves minimum value for all the three parameters. Also, the critical path delay for our algorithm is almost same as compared to VPR and ParaLaR. On an average, we achieve relative speedups of 3 times when we run a parallel version of our algorithm using 4 threads.

Index Terms—FPGA, Lagrange relaxation multipliers, sub-gradient, Steiner tree.

I. INTRODUCTION

According to the Moore's law, the number of transistors in an integrated circuit is doubling approximately every two years. In the FPGA design flow, the routing of nets (which are a collection of two or more interconnected components) is one of the most time consuming step. Hence, there is a need to develop fast routing algorithms that tackle the problem of the increasing numbers of transistors per chip, and subsequently, the increased runtime of FPGA CAD tools. This can be achieved in two ways. First, by parallelizing the routing algorithms for hardware having multiple cores. However, the pathfinder algorithm [8], which is one of the most commonly used FPGA routing algorithm is intrinsically sequential. Hence, this approach seems inappropriate for parallelizing all types of FPGA routing algorithms.

Second, instead of compiling the entire design together, the users can partition their design, compile partitions progressively, and then assemble all the partitions to form the entire design. Some existing works have proposed this approach [9], [10]. However, the routing resources required by one partition

may be held by another partition, i.e. there is no guarantee to have balanced partitions. In other words, in this approach, there is a need to tackle the difficulties arising in sharing of routing resources.

The authors in [2] overcome the limitations of existing approaches by formulating the FPGA routing problem as an optimization problem. Here, the objective function is linear and the decision variables can only have binary values. Hence, the FPGA routing problem is converted to a LP minimization problem (LP is an optimization technique in which the objective function and the constraints are linear). In this LP, the dependencies that prevent the nets from being routed in parallel are examined and relaxed by using Lagrange relaxation multipliers. The relaxed LP is solved in a parallel manner by the sub-gradient method and the Steiner tree algorithm, which is called ParaLaR.

This parallelization gives significant speedups. However, in this approach, the sub-gradient method is used in a standard way that does not always gives feasible solution (i.e. some constraints are violated). Further, by this approach, although the metric of total wire length is reduced as compared to VPR [1] (which is another commonly used algorithm for FPGA routing), but the metric of minimum channel width needs to be further improved.

There are many variants of the sub-gradient method and a problem specific method gives better result. In this paper, we use the same framework as for ParaLaR, but use an adapted sub-gradient method. Our approach substantially solves the above two problems. That is, as compared to results in [2], the number of infeasible solutions and the minimum channel width requirement both reduce by about 22% (which is same as in VPR). As compared to VPR, we save about 20% in the total wire length as well, which is the same improvement as obtained in ParaLaR. Hence, our algorithm achieves minimum value for all the three parameters. The critical path delay for our algorithm is almost same as compared to VPR and ParaLaR. On an average, we achieve relative speedups of 3 times when we run a parallel version of our algorithm using 4 threads.

The rest of this paper is organized as follows: Section II describes the formulation of the FPGA routing as an optimization problem. Section III explains the implementation of our proposed approach. Section IV presents experimental results. Finally, Section V gives conclusions and discusses future work.

II. FORMULATION OF THE OPTIMIZATION PROBLEM

The routing problem in FPGA or electronic circuit design is a standard problem that is formulated as a weighted grid graph $G(V, E)$ of certain set of vertices V and edges E , where a cost

R. Agrawal and K. Ahuja are with Discipline of Computer Science and Engineering, Indian Institute of Technology Indore, India e-mail: kahuja@iiti.ac.in.

C. H. Hoo is with Department of Electrical and Computer Engineering, National University of Singapore, Singapore.

A. Kumar is with Centre for Advancing Electronics, Technische Universität Dresden, Germany

is associated with each edge. In this grid graph, there are three types of vertices; the net vertices, the Steiner vertices, and the other vertices. A net is represented as a set $N \subseteq V$, and set N consist of all net vertices. A Steiner vertex is not part of the net vertices but it is used to construct the net tree, which is the route of a net (i.e. a sub-tree T of the graph G). A net tree is also called a Steiner tree.

Fig. 1 shows an example of 4×4 grid graph. In this figure, the black color circles represent the net vertices; the gray color circles represent the Steiner vertices; and the white color circles are the other vertices. The horizontal and the vertical lines represent the edges (as above, these edges have a cost associated with them but that is not marked here). Two net trees are shown by dotted edges.

The number of nets and the set of vertices belonging to each net is given. The objective here is to find a route for each net such that the union of all the routes will minimize the total path cost of the graph G . The goal here is to also minimize the channel width requirement of each edge. Both these objectives are explained in detail below, after (1).

To achieve the above two objectives, the problem of routing of nets is formulated as a LP problem given as follows [2]:

$$\begin{aligned} & \underset{x_{e,i}}{\text{Minimize}} \sum_{i=1}^{N_{nets}} \sum_{e \in E} w_e x_{e,i}, \\ & \text{Subject to} \sum_{i=1}^{N_{nets}} x_{e,i} \leq W, \forall e \in E, \\ & A_i x_i = b_i, i = 1, 2, \dots, N_{nets} \\ & x_{e,i} = 0 \text{ or } 1. \end{aligned} \quad (1)$$

This optimization problem minimizes the total path cost of FPGA routing, where N_{nets} is the number of nets, E is the set of edges, w_e is the cost/ time delay associated with the edge e , $x_{e,i}$ is the decision variable that indicates whether an edge (routing channel) e is utilized by the net i (value 1) or not (value 0), x_i is the vector of all $x_{e,i}$ for net i that represents the i^{th} net's route tree, A_i is the node-arch incidence matrix, and b_i is the demand/ supply vector. The inequality constraints are the channel width constraints that restrict the number of nets utilizing an edge to a constant W (which is iteratively reduced as well; discussed later). The equality constraints guarantee that a valid route tree is formed for each net, and these are implicitly satisfied by our solution approach.

To find a feasible route for each net efficiently, the above LP should be parallelized. There are two main challenges here, which are discussed next.

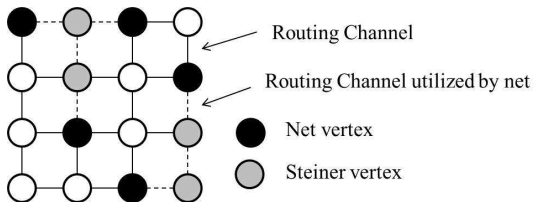


Fig. 1. A 4×4 grid graph.

A. The channel width constraints

The first challenge to parallelize the LP given in (1) is created by the channel width constraints. These constraints

introduce dependency in the parallelizing process, and therefore, should be eliminated or relaxed (see [2] for further details). The Lagrange relaxation [4] is a technique well suited for problems where the constraints can be relaxed by incorporating them into the objective function¹. For our problem, λ_e times the corresponding channel width constraints are added to the original objective function to obtain the modified LP. That is, we have the following [2]:

$$\underset{x_{e,i}, \lambda_e}{\text{Minimize}} \left(\sum_{i=1}^{N_{nets}} \sum_{e \in E} w_e x_{e,i} + \sum_{e \in E} \lambda_e \left(\sum_{i=1}^{N_{nets}} x_{e,i} - W \right) \right), \quad (2)$$

Subject to $A_i x_i = b_i, i = 1, 2, \dots, N_{nets}.$

$$x_{e,i} = 0 \text{ or } 1 \quad \text{and}$$

$$\lambda_e \geq 0.$$

This LP is independent of the channel width constraints, and hence, it can be solved in a parallel manner. After rearranging the objective function in (2), the above modified LP is given as [2]

$$\begin{aligned} & \underset{x_{e,i}, \lambda_e}{\text{Minimize}} \left(\sum_{i=1}^{N_{nets}} \sum_{e \in E} (w_e + \lambda_e) x_{e,i} - W \sum_{e \in E} \lambda_e \right), \quad (3) \\ & \text{Subject to} A_i x_i = b_i, i = 1, 2, \dots, N_{nets}. \\ & x_{e,i} = 0 \text{ or } 1 \quad \text{and} \\ & \lambda_e \geq 0. \end{aligned}$$

In (3), $(w_e + \lambda_e)$ is the new cost associated with the edge e .

B. The choice of decision variable

The second challenge to solve the LP given by (1) or the modified LP given by (3) is created by the decision variables $x_{e,i}$. These decision variables are restricted to take value either 0 or 1 (as earlier, if the edge e is utilized by the net i , then $x_{e,i} = 1$ else $x_{e,i} = 0$). Thus, this is a binary integer linear program (BILP), which is non-differentiable, and hence, cannot be solved by conventional methods such as the Simplex method [11], the interior point method [12], etc. Some methods to solve non-differentiable optimization problems include sub-gradient based methods [3], the approximation method [13], etc.

The sub-gradient based methods are commonly used algorithms to minimize non-differentiable convex functions $f(x)$. These are iterative in nature that update the variable x as $x^{k+1} = x^k - \alpha^k g^k$, where α^k and g^k are the step size and a sub-gradient of the objective function, respectively, at iteration k . In [2], the LP given in (3) is not solved directly by a sub-gradient based method but only the Lagrange relaxation multipliers are obtained by it. After this (i.e. after solving Lagrange relaxation multipliers), the minimum Steiner tree algorithm is used in a parallel manner for FPGA routing. Here, the decision variables $x_{e,i} \forall i \in 1, 2, \dots, N_{nets}$ can have only binary values. Just using a sub-gradient method will not always give binary solutions. Moreover, using a Steiner tree algorithm helps us in achieving feasible routing (equality constraints are implicitly satisfied).

¹If this is not possible, then a Lagrange heuristic can be developed [14], [15], [16].

III. IMPLEMENTATION

There are many variants of sub-gradient based methods such as the Projected sub-gradient method [3], the Primal-Dual sub-gradient method [5], the Conditional sub-gradient method [6], the Deflected sub-gradient method [6], etc. In [2], authors use the Projected sub-gradient method, where the Lagrange relaxation multipliers are calculated as $\lambda^{k+1} = \max(0, \lambda^k + \alpha^k h)$. Here, λ^k and λ^{k+1} are the Lagrange relaxation multipliers at the k^{th} and the $(k+1)^{th}$ iteration, respectively; and $h \in g^k$, i.e. a sub-gradient of the objective function given in (3) at the k^{th} iteration. Also, α^k denotes the size of the step taken in the direction of the sub-gradient at the k^{th} iterative step, and is updated as $\alpha^k = 0.01/(k+1)$. This approach satisfactorily parallelizes FPGA routing and gives better results over VPR [1], but there are many inequality constraints that are violated for some cases. Furthermore, the minimum channel width requirement needs to be improved further.

In the formulated LP given by (1), w_e is constant $\forall e \in E$. Therefore, minimizing the objective function (that is, the total path cost of FPGA routing) automatically minimizes the channel width (i.e. $\sum_{i=1}^{N_{nets}} x_{e,i}$). We start with a constant value of W , and then solve the optimization problem given by (3). This gives us the total path cost and the channel width (also the inequality constraints violation from (1), i.e. $\max(0, \sum_{i=1}^{N_{nets}} x_{e,i} - W)$). Next, we reduce the value of W and again follow the above steps to obtain a better local minima both for the total path cost and the channel width. If we obtain a reduced channel width, then the inequality constraints violation may reduce, increase or remain same. Usually, it decreases. Therefore, the above two problems are interlinked².

In our proposed work, we overcome the deficiencies of the existing approach of FPGA routing discussed in [2] by appropriately calculating the Lagrange relaxation multipliers and the corresponding step sizes. We implement three different variants of sub-gradient based methods, namely, the Projected sub-gradient method (as done in [2]), the Primal-Dual sub-gradient method, and the Deflected sb-gradient method. The difference among these variants include the different expression for the iterative update of the Lagrange relaxation multipliers and the corresponding step sizes.

A. Our algorithm

Next, we discuss our algorithm that better exploits the problem properties. We use the Primal-Dual sub-gradient method because the LP given in (3) is the Lagrange dual problem of the LP given in (1), and hence, this method is a natural fit here. That is, the Primal-Dual sub-gradient method is useful because of the way the Lagrange relaxation multipliers are updated. That is,

$$\lambda_e^{k+1} = \lambda_e^k + \alpha^k \max\left(0, \sum_{i=1}^{N_{nets}} x_{e,i} - W\right), \quad (4)$$

where $\sum_{i=1}^N x_{e,i} - W$ is a sub-gradient of the objective function at the k^{th} iteration (the partial derivative of the objective function in (3)), and $\lambda_e^0 = 0 \forall e \in E$ is the most general initial guess [4].

Let us now compare (4) with the update given in the Projected sub-gradient method (as discussed in the above paragraphs). For both the methods, if the inequality constraints are violated at the k^{th} iteration, then the Lagrange relaxation multiplier at the $(k+1)^{th}$ iteration is incremented by α^k times the sub-gradient of the objective function at the k^{th} iteration³. Otherwise, for the Primal-Dual sub-gradient method, the value of the Lagrange relaxation multiplier at the $(k+1)^{th}$ iteration is the same as the k^{th} iteration, while for the Projected sub-gradient method, it may change. In general, this works better because, if there is no constraints violation at the k^{th} iteration, then the Lagrange relaxation multiplier at the $(k+1)^{th}$ iteration should remain the same.

Next, we discuss the choice of the step size. If the step size is too small, then the algorithm would get stuck at the current point, and if it is too large, the algorithm may oscillate between any two non-optimal solutions. Hence, it is very important to select the step size appropriately. The choice of step size can be either constant in all the iterations or can be reduced in each successive iteration. In our proposed scheme, the computation of step size involves a combination of the iteration number as well as the norm of the *KKT* operator (Karush-Kuhn-Tucker operator) of the objective function at that particular iteration [6] (instead of using the iteration number only, as given in [2]). This ensures that the problem characteristic is used in the computation of the step size. That is,

$$\alpha^k = (1/k) / \|T^k\|_2,$$

where k is the iteration number, T^k is the *KKT* operator for the objective function of (3), and $\|T^k\|_2$ is the 2-norm of T^k .

The sub-gradient based methods are iterative algorithms, and hence, we need to check when to stop. There is no ideal stopping criterion for sub-gradient based methods, however, some possible measures that can be used [3] are discussed below (including our choice).

- If at an iteration k , $g^k \leq 0$ and $\lambda^k g^k = 0$, then we obtain the optimal point. Therefore, we stop here. However, this stopping criteria is achieved only if strong duality holds but, in case of our problem, there is weak duality⁴.
- Let at iteration k , f^* and f_{best}^k are the optimal value and the best possible value, respectively, of the objective function, then the sub-gradient iterations can be stopped when $|f_{best}^k - f^*| \leq \epsilon$ (where ϵ is a very small positive number). In this criteria, the optimal value of the objective function is required in advance, which we do not have.
- In diminishing step size, as discussed at the start of Section III, when the step size becomes too small, the sub-gradient method would get stuck at the current iteration, and hence, we can stop sub-gradient iterations. However, for our problem, there is no proper criteria for deciding the lower limit of the step size.

As any of the above stopping criteria do not suit us, we stop our algorithm after a sufficient and fixed number of iterations, as used in [2].

IV. EXPERIMENTAL RESULTS

Using the earlier algorithm, we perform experiments on a machine with single Intel(R) Xeon(R) CPU E5-1620 v3

³For the Primal-Dual sub-gradient method this is obvious. For the Projected sub-gradient method, λ^k , α^k , and h all would be positive.

⁴Detail of strong and weak duality can be found in [6].

²For further detail, please see Section 3 of [2].

running at 3.5 GHz and 64 GB of RAM. The operating system is Ubuntu 14.04 LTS, and the kernel version is 3.13.0-100. Our code is written in C++11 and compiled using GCC version 4.8.4, and the resulting compiled code is run using different number of threads. We compare our method with VPR [1] and ParaLaR [2]. For comparison purpose, VPR 7.0 from the Verilog-to-Routing (VTR) package and ParaLaR are compiled using the same GCC version. Some parameters in the input-output pad and the configuration logic blocks (CLBs) of VPR and ParaLaR are modified to run them identical to our model. We test on MCNC benchmark circuits [7], which range from small sized to large sized logic blocks. We use an upper limit of 50 for the number of iterations of the sub-gradient method.

In Table I, we compare the total path cost, the channel width, and the critical path delay (in nanoseconds) as obtained by our algorithm with VPR and ParaLaR. For sake of easy comparison, we call the total path cost as the total wire length here. These metrics are independent of the number of threads used, therefore, here we do not present the results for different number of threads (which is discussed in the following paragraph). If we look at the total wire length, then our algorithm gives average savings of 20.13% over VPR, which is the same as for ParaLaR. If we look at the channel width, then our algorithm gives 22.12% improvement over ParaLaR, which is the same as for VPR. Recall, the constraints violation is the difference of the channel width and the input W. Hence, this improves proportionally to the channel width improvement. Hence, our algorithm achieves minimum value for all the three parameters.

Unfortunately, minimizing the channel width and total wire length causes our algorithm to incur some extra cost (very less) in terms of critical path delay. We can see from Table I that this delay for our algorithm is on an average only 1.95% higher than that of VPR, and on an average 3.87% higher than that of ParaLaR. Hence, our algorithm incurs very little cost in terms of time.

Recall, the underlying goal of [2] and us (we improve the algorithm in [2]) is to efficiently parallelize the routing process. Hence, next we report results when using different number of threads in Table II. The benchmark dataset used is the same as discussed in the earlier paragraph (first column in Table II). Columns second through sixth give the absolute runtime and the remaining columns give the relative speedups. The speedup of execution with n threads is calculated as

$$Speedup = \frac{Execution\ time\ with\ n\ threads}{Execution\ time\ with\ 1\ thread}.$$

The symbols $1X$, $2X$, $3X$, $4X$ in Table II refers to the execution of our algorithm with 1, 2, 3 and 4 threads, respectively. We can observe from this table that our algorithm (when we run it using 1 thread) is 2.16 times faster than that of VPR. It can also be observed from this table that when we use our algorithm with 2 threads, on an average, speedups of upto 1.78 times are obtained over the single thread execution. Similarly, using 3 threads, on an average, speedups of upto 2.29 times, and when using 4 threads, on an average, speedups of upto 2.95 times are observed.

We also calculate the speedup of ParaLaR, and compare it with our proposed method. We achieve almost similar speedups, and hence, we do not report this data here. Thus, our

proposed method achieves similar parallelization as compared to ParaLaR.

V. CONCLUSIONS AND FUTURE WORK

In this work, we extend the work of [2] in proposing a more effective parallelized FPGA router. We use the LP framework of [2] and use the Primal-Dual sub-gradient method with better update of the Lagrange relaxation multipliers and the corresponding step sizes.

Experiments on the standard benchmarks show that using our algorithm gives improvements of upto 22% in the standard metric of the minimum channel width as compared to ParaLaR [2], our parent algorithm (which is the same as in VPR [1], another standard algorithm). This proportionally reduces the constraints violation, which was a problem in ParaLaR. We achieve the same total wire length (another standard metric) as ParaLaR. This is 20% better than the corresponding data for VPR. Hence, our algorithm achieves minimum value for all the three parameters. Our algorithm incurs very less extra timing cost in terms of the critical path delay, and executing it in parallel gives speedups of upto 3 times with 4 threads (over our serial implementation).

The Lagrange relaxation technique that we use, is not always guaranteed to satisfy the corresponding constraints (as observed in Sections III and IV). Hence, one future direction is to develop a Lagrange heuristic [14]–[16] specific to our problem to avoid this behavior. Another future direction involves experimenting on the Titan benchmark [17].

REFERENCES

- [1] V. Betz, J. Rose, "VPR: A new packing, placement and routing tool for FPGA research", in Proc. 7th Int. Workshop on Field-Programmable Logic and Applications (FPL), London, UK, September 1997, pp. 213-222.
- [2] C. H. Hoo, A. Kumar, Y. Ha, "PARALAR: A parallel FPGA router based on Lagrangian relaxation", in Proc. 25th Int. Conf. on Field-Programmable Logic and Applications (FPL), London, UK, September 2015, pp. 1-6.
- [3] S. Boyd, L. Xiao, A. Mutapcic, "Subgradient methods", Notes for EE392o (Stanford University), October 2003, pp. 1-21.
- [4] M. L. Fisher, "The Lagrangian relaxation method for solving integer programming problems", Management Science, vol. 27, no. 1, pp. 1-18, 1981.
- [5] S. Boyd, "Primal-Dual subgradient method", Notes for EE364b (Stanford University), Downloaded in December 2017, pp. 1-13.
- [6] B. Guta, "Subgradient optimization methods in integer programming with an application to a radiation therapy problem". PhD Thesis, Technische Universität Kaiserslautern, 2003.
- [7] S. Yang, "Logic synthesis and optimization benchmarks user guide: version 3.0", Microelectronics Center of North Carolina (MCNC), January 1991.
- [8] L. McMurchie, C. Ebeling, "PathFinder: A negotiation-based performance-driven router for FPGAs", in Proc. 3rd Int. Symposium on Field-Programmable Gate Arrays, Napa Valley, USA, February 1995, pp. 111-117.
- [9] L. A. F. Cabral, J. S. Aude, N. Maculan, "TDR: A distributed-memory parallel routing algorithm for FPGAs", in Proc. 12th Int. Conf. on Field-Programmable Logic and Applications (FPL), Montpellier, France, September 2002, pp. 263-270.
- [10] M. Gort, J. H. Anderson, "Accelerating FPGA routing through parallelization and engineering enhancements special section on PAR-CAD 2010", IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, vol. 31, no.1, pp. 61-74, 2102.
- [11] R. H. Bartels, G. H. Golub, "The simplex method of linear programming using LU decomposition", Communications of the ACM, vol. 12, no. 5, pp. 266-268, 1969.
- [12] I. J. Lustig, R. E. Marsten, D. F. Shanno, "Interior point methods for linear programming: Computational state of the art", ORSA Journal on Computing, vol. 6, no. 1, pp. 1-14, 1994.

TABLE I
COMPARISON OF QUALITY OF RESULTS BETWEEN OUR ALGORITHM, VPR, AND PARALaR.

Benchmark	Total wire length			Channel width			Critical path delay (ns)		
	Proposed	VPR [1]	ParaLaR [2]	Proposed	VPR [1]	ParaLaR [2]	Proposed	VPR [1]	ParaLaR [2]
Alu4	5087	6538	5087	35	38	48	7.01	7.43	6.89
Apex2	7927	10233	7928	48	50	59	7.90	8.23	7.50
Apex4	5652	7190	5650	49	48	61	7.50	7.25	6.38
Bigkey	4173	4711	4173	18	26	23	3.69	2.67	4.30
Clma	50310	62086	50328	81	76	104	15.90	15.13	14.95
Des	7050	8892	7047	30	30	41	5.63	5.83	5.32
Diffeq	4522	6299	4522	42	34	50	5.60	5.71	5.60
Dsip	4935	5952	4935	31	28	34	3.86	3.12	2.89
Elliptic	15198	19150	15202	58	54	79	9.30	9.02	9.25
Ex1010	23277	29474	23268	67	62	81	11.90	10.83	12.31
Ex5p	4921	6289	4921	50	50	66	6.40	7.90	7.15
Frisc	19668	24095	19659	71	62	89	12.54	11.67	11.22
Misex	5229	6789	5230	47	44	57	6.70	6.01	6.42
Pdc	30685	36803	30667	75	74	84	12.40	13.60	11.90
S298	5208	6610	5208	37	40	49	12.90	10.24	11.54
S38417	21705	28671	21707	50	48	86	8.05	8.50	8.47
Seq	7672	9691	7671	51	50	63	6.98	6.38	6.67
Spla	20404	25115	20402	66	64	86	11.34	12.50	10.61
Tseng	2436	3504	2436	34	38	47	5.27	5.70	5.27
Total	246059	308092	246041	940	916	1207	160.87	157.72	154.64
Average	12950.47	16215.37	12949.53	49.47	48.21	63.53	8.47	8.30	8.14

TABLE II
EXECUTION TIME (IN SECOND) OF VPR AND OUR ALGORITHM, AND SPEEDUP WHEN USING DIFFERENT NUMBER OF THREADS.

Benchmark	Execution time (s)					Speedup			
	VPR	1X	2X	3X	4X	1X vs VPR	2X vs 1X	3X vs 1X	4X vs 1X
Alu4	8.28	10.45	6.25	5.11	3.35	0.79	1.67	2.05	3.12
Apex2	8.58	32.49	17.15	11.53	8.52	0.26	1.89	2.82	3.81
Apex4	4.7	8.06	4.23	2.9	2.67	0.58	1.91	2.78	3.02
Bigkey	2.81	0.97	0.67	0.7	0.61	2.90	1.45	1.39	1.59
Clma	395.86	83.97	44.38	30.24	28.18	4.71	1.89	2.78	2.98
Des	9.57	2.68	1.59	1.31	0.96	3.57	1.69	2.05	2.79
Diffeq	6.54	2.43	1.43	1.27	0.87	2.69	1.7	1.91	2.79
Dsip	4.72	0.76	0.61	0.52	0.49	6.21	1.25	1.46	1.55
Elliptic	52.14	27.56	14.23	9.85	8.99	1.89	1.94	2.80	3.07
Ex1010	37.05	26.4	13.88	9.5	7.56	1.40	1.90	2.78	3.49
Ex5p	6.38	4.42	2.49	1.74	1.6	1.44	1.78	2.54	2.76
Frisc	56.86	9.75	5.53	3.88	3.33	5.83	1.76	2.51	2.93
Misex	5.55	18.74	9.74	8.95	5.59	0.30	1.92	2.09	3.35
Pdc	306.69	107.2	56.41	51.27	27.75	2.86	1.90	2.09	3.86
S298	8.34	6.72	3.61	3.38	2.51	1.24	1.86	1.99	2.68
S38417	26.13	10.76	6.04	4.28	3.86	2.43	1.78	2.51	2.79
Seq	10.67	21.25	11.19	7.5	5.8	0.5	1.90	2.83	3.66
Spla	54.79	156.51	79.81	65.84	49.79	0.35	1.96	2.38	3.13
Tseng	1.74	1.53	0.93	0.83	0.56	1.14	1.65	1.84	2.73
Total	1007.4	532.65	280.17	220.6	163.17				
Average	53.02	28.03	14.75	11.61	8.59	2.16	1.78	2.29	2.95

- [13] P. B. Dimitri, "Nondifferentiable optimization via approximation", in Balinski, M. L., Wolfe, P. (Ed.): Nondifferentiable Optimization (Springer, Berlin, Heidelberg, 1975), pp. 1-25.
- [14] O. G. Czibula, H. Gu, Y. Zinder, "A Lagrangian relaxation-based heuristic to solve large extended graph partitioning problems", in: Kaykobad, M., Petreschi, R. (Ed.) WALCOM: Algorithms and computation (Springer, Cham, 2016), pp. 327-338.
- [15] K. Holmberg, M. Joborn, K. Melin, "Lagrangian based heuristics for the multicommodity network flow problem with fixed costs on paths", European Journal of Operational Research, vol. 188, no. 1, pp. 101-108, July 2008.
- [16] S. Deleplanque, S. K. Sidhoum, A. Quilliot, "Lagrangian heuristic for a multi-plant lot-sizing problem with transfer and storage capacities", RAIRO-Operations Research, vol. 47, no. 4, pp. 429-443, 2103.
- [17] K. E. Murray, S. Whitty, S. Liu, J. Luu, V. Betz, "Titan: Enabling large and complex benchmarks in academic CAD", in Proc. 23rd Int. Conf. on Field-Programmable Logic and Applications (FPL), Porto, Portugal, September 2013, pp. 1-8.