

ON THE BEHAVIOR OF CONGESTION MINIMIZATION DURING PLACEMENT

Maogang Wang

Majid Sarrafzadeh

Department of Electrical and Computer Engineering, Northwestern University, Evanston, IL 60208
mgwang,majid@ece.nwu.edu

ABSTRACT

Typical placement objectives involve reducing net-cut cost or minimizing wirelength. Congestion minimization is least understood, however, it models routability accurately. In this paper, we study the congestion minimization problem during placement. We introduce the notion of *consistent routing model* and promote its adoption by placement systems. First, we show that in this model the wirelength objective is indeed a good measure of congestion by establishing that a placement with minimum wirelength has minimum total congestion. We show that minimizing wirelength may (and in general, will) create locally congested regions. We demonstrate that most other congestion related objectives are ill behaved and they should only be used in a post processing step. We then propose several novel congestion minimization objectives. One in particular, called *overflow minimization with look-ahead*, performs very well and can be computed very efficiently in an incremental manner. At the end, we propose a post processing phase that further improves the congestion. By combining the overflow minimization with look-ahead and the post processing phase, we improve the congestion by more than 40% on the average.

1. INTRODUCTION

Automated cell placement for VLSI circuits has always been a key factor for achieving designs with optimized area usage, wiring congestion and timing behavior. As technology advances, the congestion problem becomes more and more important. With the advent of over-the-cell routing, the goal of every place and route methodology has been to utilize area to prevent spilling of routes into channels. It is this overflow of routes that accounts for an increase in area. The multiple routing layers have enough routing resources to route most wires as long as there are not too many wires congested in the same region. Excessive congestion will result in a local shortage of the routing resource.

Typical placement objectives involve reducing net-cut costs or minimizing wirelength. Because of its constructive nature, min-cut based strategies minimize the number of net crossings but fail to uniformly distribute them [7]. Congestion-driven placement based on multi-partitioning was proposed in [5]. It uses the actual congestion cost calculated from pre-computed Steiner trees to minimize the congestion of the chip, however, the number of partitions is limited due to the excessive computational load. The use of minimal wirelength as a metric to guide placement has been successful in achieving good placement. However, it only indirectly models congestion and the behavior of the router. Reducing the global wirelength helps reduce the

wiring demand globally, but does not prevent existing local congested spots. It is entirely feasible for a minimum wirelength solution to require more routing resources through a region than are available. Therefore, traditional placement schemes which are based solely on wirelength minimization [8, 3, 9, 1, 4, 2, 10] cannot adequately account for congestion.

The congestion problem in placement is not well studied. There are not many results on this problem [5, 6]. In this paper, we theoretically and experimentally study the congestion problem during placement. We first point out that minimizing wirelength is indeed equal to minimizing the average routing demand where the routing demand is the number of tracks needed for routed nets. We also give a theoretical distribution curve of the congestion on a layout. Then we focus on finding a good objective to effectively reduce the congestion in the final placement. The traditional wirelength objective is not a direct measure of the congestion. A good placement for wirelength may not be a good placement for congestion. Thus we are aiming at an objective which is at least better than the wirelength objective. Using the congestion cost directly as the objective is not effective. The congestion cost is a badly behaved objective function because it is not sensitive to placement moves.

We propose a new objective called *overflow with look-ahead* which is well behaved and is directly related to congestion. Experimental results prove that this new objective is effective. We also propose a flow-based post processing stage to further improve congestion. We get best congestion results by using this objective in the post processing stage. The placement produced by this new objective has on the average 40% less congestion than the best congestion results obtained by commonly used objective.

The rest of the paper is organized as follows: In Section 2, we formally define the congestion cost followed by some theoretical analysis on congestion and wirelength in a layout. In Section 3, we compare five different objectives to minimize the congestion. In Section 4 we propose an effective new objective to reduce congestion. The post processing stage is introduced in Section 5 and the conclusion is in Section 6.

2. THEORETICAL ANALYSIS

Intuitively speaking, congestion in a layout means too many nets are routed in local regions. In this paper we assume that we are given a netlist that consists of a set of cells connected by a collection of nets. Each net consists of a set of pins. Pins are to be assigned a geometric location on the layout surface in the placement process.

2.1. Definition of the Congestion Cost

The congestion cost is defined based on the global bin concept. We partition a given chip into several rectilinear regions, each of these regions is called a global bin or a bin for simplicity. Figure 1 shows an example. In Figure 1, we have $4 \times 4 = 16$ global bins.

The congestion is quantified as number of crossings between routed nets and global bin edges. Each global bin has two horizontal and two vertical edges surrounding it. Assume we have r rows and c columns of global bins. There are $(r + 1)$ rows of horizontal global bin edges and r rows of vertical edges as shown in Figure 1. Each row of horizontal edges has c edges and each row of vertical edges has $(c + 1)$ edges.




Figure 1. Layout of a circuit and global bins.

Given a placement, all the cells and pins have fixed positions on the chip. In order to get the congestion information, we need to estimate routing on the chip. We can use a “router” to route all the nets. This “router” is not necessarily a detail router. It can be a very simple global router or even a bounding box router. For each global edge, there are routed nets going across it. Therefore, for each global edge e , the routing demand of e , d_e , is defined as the number of the nets crossing e . The routing supply of a global edge e , s_e , is a fixed value which is a function of the length of the edge and technology parameters. A global edge e is congested if and only if the routing demand (number of the crossing nets) exceeds the routing supply of that edge ($d_e > s_e$). If a global edge e is congested, the overflow of e is defined as the exceeding amount of the routing demand over the routing supply of e . The overflow of e is zero if e is not congested. To describe overflow formally: overflow_e equals $d_e - s_e$ if $d_e > s_e$ and zero otherwise.

Using the above global bin and global edge notation, there are several parameters which can represent congestion in a placement. A.) The *total overflow* of a placement is defined as the summation of the overflow for all global edges. The amount of total overflow is the amount of total shortage of routing resource in the placement. Thus a placement with less total overflow is less congested. B.) Another measure of congestion is the number of *congested edges*. More congested edges means more regions are congested. While these two parameters both represent congestion in a placement, we feel that the total overflow is a better measure to describe congestion. The purpose of reducing congestion is

to provide better routability in the chip. Using the number of congested edges as the congestion cost does not serve this purpose very well. A placement with one highly congested edge is definitely worse than a placement with a couple of slightly congested edges. We will study both of these two congestion parameters in this paper.

2.2. Consistent Models for Wirelength and Congestion Estimation

Based on the above definition, the congestion cost is router-dependent. The specific value of the congestion cost for a given placement will be different for different routers. Figure 2 shows a 4-pin net in the global bin grid. In order to normalize the wirelength of the nets, we use the dimension of the global bin grid as the unit length. The width of a global bin is the unit length in the x direction and the height of a global bin is the unit length in the y direction. Given locations of all pins, there are a number of ways to route all the nets. For example, we can use the bounding box, the minimum spanning tree (MST) or the Steiner tree model to estimate the actual routing. A bounding box and a MST routing model are illustrated in Figure 2.

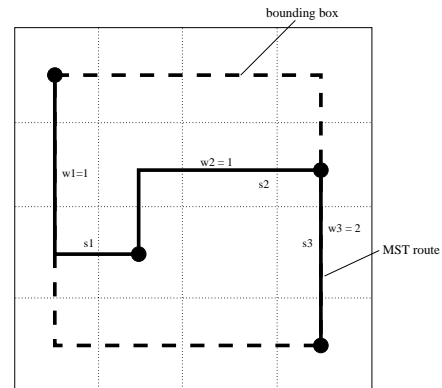


Figure 2. Consistent model.

The congestion is not independent of the wirelength cost. In order to find an effective algorithm to reduce the congestion, it is essential to understand the correlation between wirelength and congestion. As mentioned above, the calculation of wirelength and congestion are both based on the routing estimation of all the nets. After the routing estimation is done, the wirelength calculator calculates the total length of all the routes and the congestion calculator calculates the number of crossings between all the routes and all the global edges. Intuitively, the routing estimation method (router) used by the wirelength calculator and the congestion calculator should be the same in the placement stage. Different routing estimation method results in different degrees of accuracy and time complexity. For example, in most cases, the Steiner tree router should be a more accurate routing method than the bounding box router. It is inconsistent to use one routing method to measure wirelength and another method to measure congestion.

Usually algorithm designers use “consistent” routing models for estimation wirelength and congestion. Notice that the term *consistent* model here does not necessarily mean “the same” model for wirelength and congestion. Two different routing models sometimes also can have consistent results. As shown in Figure 2, for example, we can use the bounding box model to calculate the wirelength and the MST model to calculate the congestion. The routing model

assigns weight to each segment of the route. Two routing models are defined to be “consistent” if the total weighted length of the routes are the same. In Figure 2, all branches of the bounding box have a weight of 1, so the total weighted distance of the bounding box is 12. In the MST model, two branches have a weight of 1 and one branch has a weight of 2. This brings the total weighted length of the routes to be 12. Thus these two routing models, the uni-weighted bounding box model and the weighted MST model are consistent.

Formally, a model for net η consist of a set of segments, S_η . Each segment $s_i \in S_\eta$ has a length l_i and a weight w_i . Thus the total weighted length for net η , \mathcal{L}_η , is $\sum_{s_i \in S_\eta} w_i l_i$. And the total weighted length for all nets is $\sum_\eta \mathcal{L}_\eta = \sum_\eta \sum_{s_i \in S_\eta} w_i l_i$. If two different routing models have the same total weighted length for all nets, these two models are said to be consistent models.

2.3. Correlations between Wirelength and Congestion

Observation 1 *If we use consistent models to calculate wirelength and congestion, the total wirelength of a layout is equal to the total routing demand on all global edges, i.e., $\sum_\alpha l_\alpha = \sum_e d_e$, where l_α is the estimated length for net α and d_e is the routing demand for global bin edge e .*

Since we are using the dimension of the global bin as the unit length to measure the wirelength, each unit length wire will cross a global edge. Thus, each unit of the wirelength will contribute to one crossing between the wire and a global edge which is by definition one unit of the routing demand. If we use exactly the same routing model for calculating wirelength and congestion, the total units of wirelength will be equal to the total units of the routing demand. If we use consistent routing models for evaluating wirelength and congestion, by definition, the total wirelength is the same for two routing models. Thus the above claim is true for any consistent model.

This observation shows the underlying correlations between the wirelength cost and the congestion cost. When we minimize the wirelength cost, the total amount of routing demand is minimized. Thus the average routing demand on a global edge is minimized. Given a fixed amount of routing supply which is dependent on technology parameters, the less the routing demand is, the bigger chance we will get a low-congestion layout.

Observation 2 *The maximum routing demand is greater than or equal to the total wirelength divided by the number of global edges, that is, $\max(d_e) \geq \frac{\sum_\alpha l_\alpha}{\#edges}$ where l_α is the length of net α .*

Since the total wirelength is equal to the total routing demand according to Observation 1, the term $\frac{\sum_\alpha l_\alpha}{\#edges}$ is the average routing demand which is less than or equal to the maximum routing demand. Since the total wirelength is always known for a given placement, we can get a lower bound of the maximum routing demand on global edges without carrying out any congestion analysis.

2.4. Theoretical Analysis of Congestion Distribution

In this subsection, we will present theoretical analysis of distributions of the placement congestion. In order to study

the behavior of congestion, we need to make some assumptions. First, let us consider the x direction of the layout first. Assume we have n nets in the layout. They are labeled as $\alpha_1, \alpha_2, \dots, \alpha_n$. The length of the net α_i is l_i where $i = 1, 2, \dots, n$. Assume all the nets are two-pin nets. The left pin of the net α_i is located at x_i , the right pin of α_i is located at $x_i + l_i$. For simplicity, we assume that all the lengths are distributed on the layout randomly.

Assume the width of the chip is W . Based on the random distribution assumption, for net α_i , the left pin is uniformly distributed between segment $(0, W - l_i)$. This pin cannot be located between segment $(W - l_i, W)$ because otherwise there is not enough room to let this net be l_i long. Suppose a vertical global edge e is located at position x_e . We are going to calculate the probability that the net α_i intersects the global edge e . Since the length of α_i is l_i , obviously, α_i intersects e if and only if $x_i \leq x_e$ and $x_i + l_i \geq x_e$. Therefore, the probability that α_i intersects e is:

$$\begin{aligned} p_i(x_e) &= \frac{x_e}{W-l_i} && \text{if } 0 \leq x_e \leq l_i; \\ p_i(x_e) &= \frac{l_i}{W-l_i} && \text{if } l_i \leq x_e \leq W-l_i; \\ p_i(x_e) &= \frac{W-x_e}{W-l_i} && \text{if } W-l_i \leq x_e \leq W; \end{aligned}$$

In order to get the total expectation value of the number of the crossing nets at x_e , we need a statistical function $g(l)$. Function $g(l)$ gives the number of nets which has a length of l in the layout. This is a statistical function and can be obtained from placed benchmark circuits. Therefore, the expected number of wires crossing global edge e is $\sum_i p_i(x_e)g(l_i)$. The actual value of this function can be obtained by a numerical method.

The solid curve in Figure 3 shows the congestion distribution curve obtained from the benchmark circuit Primary1. The circuit was placed to minimize the wirelength. The x axis represents the x positions on the layout. The y axis is the routing demand at position x . For a random layout, theoretically, the most congested region is at the middle of the chip. The solid straight line in Figure 3 is the average congestion value at any position. Thus more than $\frac{4}{5}$ region of the layout has an above average congestion value. This is consistent with Observation 2.

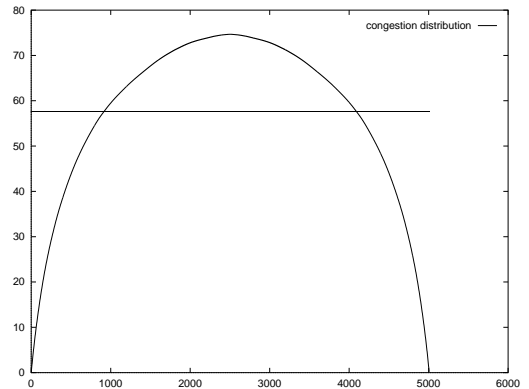


Figure 3. Theoretical congestion distribution on the x direction for Primary1.

The congestion distribution in y direction is the same as in x direction due to the symmetrical nature. Since the total congestion at any point in the chip is the summation of the congestion in x and y direction, the total congestion

distribution in a two-dimensional chip is a convex plane. Figure 4 shows an example.




Figure 4. Theoretical congestion distribution on a two-dimensional layout.

3. OBJECTIVE FUNCTIONS FOR CONGESTION MINIMIZATION

Our goal is to find a good placement with low congestion. This is an optimization problem. We need to set up an objective to optimize. In this section, we perform a series of experiments in order to determine what is a good objective to optimize in order to get a low-congestion layout.

An obvious choice will be the two congestion costs defined in Section 2.1. Since we have a clear definition of the congestion cost for a given placement, we can directly use this cost as the objective to minimize. We can use the total overflow or the number of congested edges as our objective. Besides this direct objective, we also have some other choices. Observation 1 in Section 2. shows that the wirelength cost is a reasonable objective to minimize the congestion cost. Thus the wirelength cost is also a candidate for an objective to minimize the congestion. Finally, we can put wirelength and congestion together to form a hybrid objective.

Before we run the actual tests to determine which objective is the best, let us perform some analysis first. We know that wirelength is indirectly correlated to congestion, so it would not give us the best result for congestion. The congestion objectives are direct measures of congestion. If we use an optimal optimization technique, we should be able to get a layout with the minimum congestion. However, since the placement problem is NP-hard, the optimization result highly depends on the properties of the objective function. A sensitive objective function will be easier for an optimization heuristic to find the global minimum.

The total overflow objective is not as sensitive as the wirelength objective. When we move a cell, the routing demand is changed on some global edges. However, if the routing demand before and after the change are both less than or equal to the routing supply of that edge, the overflow will not change. The number of congested edges is not a sensitive objective function either. In fact, this objective is less sensitive than the overflow objective. Unless the change in the routing demand happens around the routing supply value, the number of congested edges will not change. Therefore, the direct congestion cost may not be a very effective objective for iterative optimization techniques.

The hybrid objective which combines the wirelength and the congestion might be a reasonable objective to use. We

have two choices: combine wirelength and overflow together or combine wirelength and number of congested edges together. The wirelength part of the objective is sensitive to cell moves and is able to reduce the congestion globally. The congestion part of the objective controls the optimization technique to work towards the correct goal.

To summarize, we have the following five objectives to use to reduce the congestion in a placement:

- *WL*: Standard total wirelength objective (see Observation 1 in Section 2.3).
- *OF*: Total overflow in a placement. This is a direct measure of the congestion.
- *CE*: Total number of congested edges in a placement. This is also a direct measure of the congestion.
- *WLOF*: A hybrid objective which combines wirelength and overflow.
- *WLCE*: A hybrid objective which combines wirelength and number of congested edges.

In order to test these five objectives, we ran eight MCNC standard-cell benchmark circuits. The characteristics of these circuits are shown in Table 1. The size of the global bin grid is chosen so that each bin has roughly 10 – 50 cells.

TestCase	# Cells	# Nets	Global Bins
highway2	62	87	2×2
fract	125	163	3×3
Primary1	833	1266	8×8
Primary2	3014	3817	8×8
struct	1888	1920	8×8
biomed	6417	7052	16×16
avqs	21584	30038	20×20
avql	25114	33298	20×20

Table 1. Testing circuits information.

we can test the proposed objective with any placement heuristic. We have selected Simulated Annealing (SA). It is theoretically proved that given infinite amount of time, SA can get the global optimal result for any objective function. SA is widely used in VLSI CAD tools. The TimberWolf placement package [9] and the NRG placement tool [8] use simulated annealing and produces very good results on wirelength. Besides SA, other optimization techniques could be chosen as well. Results in this paper are obtained using NRG’s global placer. NRG has the same performance as TimberWolf. However, the objective of this paper is to show how to improve congestion of ANY placement result.

Table 2 shows the results for circuit biomed. Each column of Table 2 is corresponding to one of the five testing objectives. As we mentioned above, there are a couple of parameters which can be used to represent congestion. The first row of the table is the total overflow of the final placement. The second row is the number of congested edges and the third row is the maximum routing demand among all global edges.

For simplicity, we put the testing results for all eight benchmark circuits in one table, Table 3. Each row of the table corresponds to one benchmark circuits and each column corresponds to one testing objective. Since we believe that the total overflow is a more accurate measure for congestion, we put the total overflow at each entry of Table 3. Lower overflow value means better placement. As we can see from the table, the objective CE is obviously the

	WL	OF	CE	WLOF	WLCE
Orflw	69	426	58302	188	139
#ConEdge	20	393	480	31	21
MaxCon	68	66	634	94	78

Table 2. Comparison between different objectives for circuit *biomed*.

worst one. This is because CE is the least sensitive objective among all five objectives. Objective OF is not good either due to its insensitive nature. Best congestion results are obtained using one of the other three objectives. There is no objective which can produce best result in all eight benchmark circuits. The results show that WL is the winner on the average for all the eight circuits. However, we know that in practice the placement with minimal wirelength does not always satisfy the congestion constraint. Therefore, we need to find a new objective which can reduce the congestion more effectively than the wirelength objective.

TestCase	WL	OF	CE	WLOF	WLCE
highway2	12	8	72	6	16
fract	14	32	274	16	16
Primary1	24	97	4697	41	28
Primary2	33	105	16810	95	97
struct	49	184	12044	55	52
biomed	69	426	58302	188	139
avqs	336	1141	322948	392	315
avql	465	2254	269063	316	383

Table 3. Total overflow cost using five different objectives.

4. A NEW OBJECTIVE

We believe that an effective congestion objective should be sensitive to placement moves and directly related to the congestion cost. An objective function is used in evaluating a change of placement (a move) in an optimization algorithm. For any global edge e , the routing supply is s_e . Suppose the routing demand of e is d_e before a move and d'_e after the move. The direct overflow cost of this move will be $\max(d_e, s_e) - \max(d'_e, s_e)$. As we can see, if $d_e < s_e$ and $d'_e < s_e$, the cost of the move will be zero. However, if d_e or d'_e is close to s_e , i.e., $s_e - \delta \leq d_e, d'_e \leq s_e$ where δ is a small number, the change on d_e is still useful to evaluate the move. For example, an increase in d_e will result in a higher probability of changing the edge e from uncongested to congested in later moves; and a decrease in d_e will help the edge e stay uncongested in the future. On the other hand, if d_e and d'_e are both far less than s_e , i.e., $d_e, d'_e \leq s_e - \delta$, we do not care the change in d_e because the edge e will more likely remain uncongested in the near future. Based on this discussion, we propose a modified overflow objective, *overflow with look-ahead*. The cost of each move is $\max(d_e, s_e - \delta) - \min(d'_e, s_e - \delta)$ where δ is an adjustable parameter.

We test this overflow with look-ahead objective on the same eight benchmark circuits mentioned in the previous section. The results are summarized in Table 4. The second column is total overflow results obtained by the overflow with look-ahead objective. The third column is the best results obtained by the previous five objectives in Section

3. The fourth column is the percentage improvement of the new objective over the best of the old objectives.

TestCase	look-ahead OF	best old obj.	% imp.
highway2	2	6	66.7%
fract	12	14	14.2%
Primary1	2	24	91.6%
Primary2	11	26	57.6%
struct	24	52	53.8%
biomed	40	69	42.0%
avqs	827	315	-162%
avql	1412	316	-346%

Table 4. Congestion results using the modified overflow objective.

From Table 4, the overflow with look-ahead objective does not work well in all benchmark circuits. The reason is that this overflow with look-ahead objective is still not sensitive enough to let an optimization technique find a global optimal solution easily. Thus for large circuits like avqs and avql, the optimization algorithm gets stuck in a local minima. However, this objective works extremely well on small circuits. We have an average 54% improvement compare to the old objectives for six small circuits. The good performance of this new objective on small circuits suggests that this objective works well locally in a small region. Thus it should be effective in a post processing stage.

5. POST PROCESSING USING OVERFLOW MINIMIZATION WITH LOOK-AHEAD

We propose a two stage process to reduce the congestion in a layout. In the first stage, we use the wirelength as the objective to minimize the average congestion. After the first stage is done, we can perform post processing to further reduce the congestion. In the post processing stage, we use the overflow with look-ahead cost as the objective to minimize.

We implemented a greedy algorithm to do the post-processing overflow minimization. The greedy algorithm can test the effectiveness of the post-processing idea with look-ahead objective. In the implementation, we evaluate moving a cell or exchanging two cells using the modified overflow objective. Then we make this move or exchange if and only if it can give us a lower objective cost value.

We run simulated annealing using the wirelength objective as our first stage. The output placement from the first stage will be the input to the post processing stage. The quality of the annealing result is closely related to the number of searches made at each temperature. Thus we prepared two kinds of inputs for the post processing stage. The low quality input is obtained by running the first stage really fast. And the high quality input is obtained by a slow annealing procedure. Table 5 shows the post processing results on low quality inputs. The second column of the table is the total overflow before the post processing stage. The third column is the total overflow after the post processing stage. The fourth column is the percentage improvement on the total overflow by using the post processing stage. The average improvement is 5% and the final overflow value is not good compare to the results obtained by the old objectives. This results shows that the post processing stage using the overflow with look-ahead objective can only make local optimizations. If the input placement is not good,

the post processing stage will not be able to get a good low-congestion placement.

TestCase	before PP	after PP	% improv.
highway2	6	6	0%
fract	20	20	0%
Primary1	74	68	8.1%
Primary2	883	831	5.9%
struct	867	850	2.0%
biomed	1525	1169	23.3%
avqs	14658	14534	0.8%
avql	16354	16269	0.5%
ave.			5.1%

Table 5. PP results on low quality inputs.

Table 6 shows the results from the post processing stage with the high quality inputs. The *before PP* column in the table is the results before the post processing stage. The percentage improvement column is the improvement of the post processing stage comparing to the results before post processing. The table shows that post processing stage can significantly reduce the congestion cost if the input placement is good. We get an average 47.7% improvement comparing to the congestion results before post processing. It works well on both large and small circuits. Table 7 compares the post-processing results with the best results we obtained previously,

TestCase	before PP	after PP	% improv.
highway2	6	6	0%
fract	14	14	0%
Primary1	24	8	66.7%
Primary2	26	11	57.7%
struct	74	39	47.3%
biomed	69	6	91.3%
avqs	336	125	62.8%
avql	465	207	55.5%
ave.			47.7%

Table 6. PP results on high quality inputs.

The *best old* column is the best results obtained by the previous five objectives. The percentage improvement column is the improvement of using the post processing stage comparing to the best old results. This table shows that among all the congestion reduction methods studied in this paper, this post processing method using the overflow with look-ahead objective has the best results. The results of it are on the average 41.9% better than all studied methods. In experiments, we observed the fact that the post-processing stage will increase the total wirelength by about 5%. However, this increasement is not important in today's over-the-cell routing scenario.

6. CONCLUSION

In this paper, we studied the behavior of congestion minimization in placement. As shown both by theoretical and experimental results, the congestion cost is a poorly behaved function. If we use the congestion cost directly as the objective to minimize, we will not get a low congestion placement. Our theoretical analysis showed that there are some correlations between wirelength and congestion in a placement. Specifically, the total wirelength is equal to the total routing demand of a placement for consistent routing

TestCase	after PP	best old	% vs. old
highway2	6	6	0%
fract	14	14	0%
Primary1	8	24	66.7%
Primary2	11	26	57.7%
struct	39	52	25%
biomed	6	69	91.3%
avqs	125	315	60.3%
avql	207	316	34.5%
ave.			41.9%

Table 7. PP results compared to old results.

models. Therefore, minimizing wirelength is helpful in minimizing congestion globally. We proposed a new objective named overflow with look-ahead to reduce the congestion. This objective is experimentally proved to be more effective than any direct congestion objective. We get extremely good congestion results by using this new objective in a post processing stage. On the average, this method produces more 40% less congestion than other studied methods.

REFERENCES

- [1] J. Cong and S. K. Lim. "Multiway Partitioning with Pairwise Movement". In *International Conference on Computer-Aided Design*, pages 512–516, 1998.
- [2] H. Eisenmann and F. M. Johannes. "Generic Global Placement and Floorplanning". In *Design Automation Conference*, pages 269–274. IEEE/ACM, 1998.
- [3] D. Huang and A. B. Kahng. "Partitioning-based Standard-cell Global Placement with an Exact Objective". In *International Symposium on Physical Design*, pages 18–25. ACM, April 1997.
- [4] J. M. Kleinhans, G. Sigl, F. M. Johannes, and K. J. Antreich. "GORDIAN: VLSI Placement by Quadratic Programming and Slicing Optimization". *IEEE Transactions on Computer Aided Design*, 10(3):365–365, 1991.
- [5] S. Mayrhofer and U. Lauther. "Congestion-Driven Placement Using a New Multi-Partitioning Heuristic". In *International Conference on Computer-Aided Design*, pages 332–335. IEEE/ACM, November 1990.
- [6] P. N. Parakh, R. B. Brown, and K. A. Sakallah. "Congestion Driven Quadratic Placement". In *Design Automation Conference*, pages 275–278. IEEE/ACM, 1998.
- [7] Y. Saab. "A Fast Clustering-based Min-cut Placement Algorithm with Simulated-annealing Performance". *VLSI Design: An International Journal of Custom-Chip Design, Simulation, and Testing*, 5(1):37–48, 1996.
- [8] M. Sarrafzadeh and M. Wang. "NRG: Global and Detailed Placement". In *International Conference on Computer-Aided Design*. IEEE, November 1997.
- [9] C. Sechen. *VLSI Placement and Global Routing Using Simulated Annealing*. Kluwer, B. V., Deventer, The Netherlands, 1988.
- [10] K. Shahookar and P. Mazumder. "VLSI Cell Placement Techniques". *ACM Computing Surveys*, 23(2):143–220, June 1991.