Improving the Quadratic Objective Function in Module Placement *

Lars Hagen and Andrew B. Kahng

UCLA Department of Computer Science
Los Angeles, CA 90024-1596

Abstract

Traditional placement goals for cell-based designs involve minimizing either total wirelength or channel width; each of these metrics reflects final layout area. Typically, a net model is first used to transform the netlist hypergraph into a graph representation. A linear wirelength objective is then formulated, but then for technical reasons a quadratic form is actually optimized using relaxation or eigenvector methods. We address this seeming inconsistency and propose a simple transformation of the quadratic objective which recaptures the "linear" nature of the original minimum-wirelength objective. Computational results for a wide range of standard benchmarks show that this refinement gives very significant savings in both total wirelength and channel width for linear placement: these values are respectively reduced by an average of 7% and 18% over results obtained with the previous standard approach.

1 Preliminaries

Module placement has presented one of the most persistent challenges in IC layout synthesis. The true criterion for a good placement is efficient autoroutability, subject to performance and chip area constraints. Thus, placement models and objective functions are very hard to formulate. Many methods have been proposed, such as clustering, iterative improvement, row-assignment, recursive min-cut, force-directed, etc.; early techniques are surveyed by Hanan et al. [7], and Lengauer [9] gives a recent overview.

We may view the placement procedure as the assignment of n modules to n available slots such that a given cost function is minimized, i.e., we have a *quadratic assignment* formulation. In particular, to simplify our discussion we will address the problem of *linear placement*, where the modules are assigned to positions $1, \ldots, n$ on the one-dimensional integer lattice. The linear placement problem, which we will formalize below, captures a number of practical applications, e.g., placement within cell rows. Furthermore, as shown by Kurdahi and Sastry [8], two-dimensional placement in cell-based designs can be modeled as two interacting one-dimensional placement processes.

There are three main components to any module placement approach:

- First, a *net model* represents signal nets (hyperedges in the netlist hypergraph) by graph edges.
- Second, given the net model, an *objective function* is formulated for global minimization.
- Third, a (heuristic) *optimization method* must be used to minimize the objective function.

The results in this paper stem from the following simple observation: while standard net models and objective functions are aimed at minimizing wirelength, the accompanying optimizations actually minimize a completely different objective function which is *quadratic* in wirelength. Put another way, the traditional approach to placement algorithms becomes problematic in that one picks a net model which reflects a minimum wirelength objective, but then takes a "left turn" and instead minimizes the sum of squared wirelengths.

This inconsistency—between the objective function developed from the net model and the cost function actually optimized—is the subject of the present work. Our goal here is to determine an appropriate net model and objective function, and then be able to actually minimize a good approximation to this objective function. For a number of reasons, the linear wirelength objective is seen to be a good objective function. Furthermore, while technical reasons have historically motivated the substitution of the quadratic for the linear wirelength objective, we show that a simple change of weights in the quadratic form of the objective will "recapture" the effect of the original minimum-wirelength placement objective. Experimental results confirm that when our refinement is applied to traditional placement methods, very significant improvements result for both the total wirelength and maximum channel width metrics.

2 The Placement Process

Placement has traditionally been formulated as a quadratic assignment problem where modules must be placed into available slots. The slots are often assumed to be linearly ordered, e.g., for backplane formation or cell placement within assigned rows. As noted above, this is reasonably general since two-dimensional placement may be described by two simultaneous one-dimensional placement processes [8].

*This research was supported in part by NSF grant MIP-9110696

Linear Placement: Given a netlist graph $G = (V, E)$ with $|V| = n$ and connection weights c_{ij} between module pairs (v_i, v_j), map the modules onto distinct positions x_i taken from the set $\{1, \ldots, n\}$ such that an objective function dependent on the c_{ij} and the module positions is minimized.

2.1 Net Models

In the first phase of the placement algorithm, a net model is derived which yields a graph representation of the netlist hypergraph; the net model reflects the routing cost of the placed netlist and determines the c_{ij} values in the above formulation. Many net models have been proposed, including spanning paths, spanning cycles, spanning trees, star topologies, etc. Several models can suffer from nondeterministic asymmetry in the connection weights c_{ij}, i.e., not all adjacencies derived from a given k-pin net will be accorded the same significance. Furthermore, minimum spanning tree, centroid-based star (e.g., [10]), or other topologies are inherently dynamic, requiring recomputation with every change in the module placement (see [9] for a survey).

The most common net model is that of a weighted clique, where a k-pin net will induce $C(k, 2)$ edges among its k modules. The early survey by Hanan et al. [7] details several clique weighting variants which propose uniform weighting of the $C(k, 2)$ edges by such values as $\frac{2}{k}$, $\frac{1}{k}$, $\frac{1}{k-1}$, etc. Simple dimensional analysis shows that all of these net models are essentially identical in practice. Recent work has widely adopted a "standard" weighted clique model [9], wherein a k-pin net contributes $\frac{1}{k-1}$ to each of $C(k, 2)$ c_{ij} values. While the most obvious advantages of the clique model stem from its symmetry, additional intuition justifying this standard edge weighting can be developed as follows.

Since the goal of linear placement is to minimize the wire length contributed by each net, i.e., the net span in one dimension, the weighting of each net should be a function of the minimum possible wire length WL_{min} for each net. Moreover, the weighting should intuitively be inversely proportional to WL_{min} so that the objective function does not "try too hard" to place the k modules of a net into a span of fewer than k slots when such an arrangement is physically impossible.

In the discrete linear placement problem defined above, the minimum wirelength WL_{min} of a k-node net is $k-1$, and by the above intuition we should adopt a weighting function on the order of $\frac{1}{k-1}$. Thus, our heuristic picture of linear placement supports use of the present standard clique net model.

2.2 The Minimum Wirelength Objective

Certainly, the goal of module placement must encompass minimization of total wirelength. Each grid unit of wire requires an additional quantum of chip area which is dependent on the wiring pitch. Moreover, larger wirelength generally implies larger RC constants, which can adversely impact system performance and power requirements. Of course, layout area in cell-based designs is also determined by the number of wiring tracks used, i.e., the sum of channel widths is also an important objective function. Although minimum wirelength placement and minimum channel width placement are not perfectly correlated, we note that results of Adolphson and Hu [1] may be used to establish a probabilistic relationship between the two metrics. Indeed, our experimental results below show that when we use our new methods to improve wirelength, the channel width usually also improves. In view of these arguments, we choose to minimize a sum of wirelengths objective function; for linear placement, this corresponds to minimizing the sum of net spans. Such a conclusion is in agreement with researchers who have considered the tradeoffs between wirelength and squared-wirelength objectives. For example, Sigl et al. [10] recently concluded that "a linear objective function seems to reflect the actual wiring demands more accurately than the quadratic objective function".

2.3 Optimization Methods

Finally, it is well known that the Linear Placement formulation above is NP-complete when we minimize the objective function $\sum_{i,j} c_{ij}|x_i - x_j|$ [3]. By contrast, if we relax the slot constraint and introduce a quadratic form in the minimization, efficient numerical algorithms can be used to obtain a global optimum solution. More specifically, we represent the circuit netlist by the simple undirected graph $G = (V, E)$ with $|V| = n$ vertices $v_1, \ldots, v_n$. We use the $n \times n$ *connection matrix* $C = C(G)$, where $c_{vw} = 1$ if $(v, w) \in E$ and $c_{vw} = 0$ otherwise. If G has weighted edges, then c_{vw} is equal to the weight of $(v, w) \in E$, and by convention $c_{vv} = 0$ for all $v \in V$. If we let $d(v)$ denote the degree of node v (i.e., the sum of the weights of all edges incident to v), we obtain the $n \times n$ *diagonal matrix* D defined by $D_{ii} = d(v_i)$. As noted by, e.g., Hall [6], Cheng and Kuh [2], and Tsay et al. [12], the real eigenvector corresponding to the second smallest eigenvalue of Q gives a linear placement solution vector $\vec{x}$ which minimizes $\sum_{i,j} c_{ij}|x_i - x_j|^2$ (i.e. sum of squared wirelengths) subject to the constraint $|\vec{x}| = 1$.[1]

[1] Hall [6] noted that the eigenvectors of the matrix $Q = D - C$ solve the problem of finding the vector $\vec{x} = (x_1, x_2, \ldots, x_n)$ minimizing

$$z = \frac{1}{2}\sum_{i=1}^{n}\sum_{j=1}^{n} c_{ij}(x_i - x_j)^2$$

subject to the constraint $|\vec{x}| = (\vec{x}^T\vec{x})^{1/2} = 1$. Since $z = \vec{x}^T Q\vec{x}$, to minimize z we may form the Lagrangian

$$L = \vec{x}^T Q\vec{x} - \lambda(\vec{x}^T\vec{x} - 1).$$

Setting the first partial of L with respect to $\vec{x}$ equal to zero yields

$$2Q\vec{x} - 2\lambda\vec{x} = 0,$$

which can be rewritten as

$$(Q - \lambda I)\vec{x} = 0$$

Eigenvector solutions to the linear placement problem with quadratic objective function can then be found by either iterative relaxation methods [2] [12] [10] or sparse-matrix operator techniques such as the Lanczos method [5].

Some heuristic justifications for the squared wirelength objective have been put forth, notably that the metric reduces congestion since it tends to reduce the maximum wirelength of any net. However, it is by no means clear why so much effort has gone into optimizing the particular squared wirelength objective $\sum_{i,j} c_{ij}|x_i - x_j|^2$, which is not very reflective of the "true" cost function, $\sum_{i,j} c_{ij}|x_i - x_j|$. Indeed, much recent work simply points to the formulations used by Hall [6] and Cheng and Kuh [2], rather than presenting any rationale for the quadratic objective. As Lengauer notes in [9] (p. 317): "The main reason why quadratic wire-length estimation is so popular is that the method leads to quadratic cost functions that can be minimized easily .... Thus, the final motivation for using quadratic wire length is driven by methodical, not modeling, arguments."

It is interesting to note that very few researchers besides Lengauer have commented on this issue. Sigl et al. [10] claim to be the first to explicitly address the differences between the linear and quadratic objectives. As noted above, they find the linear objective to be more "accurate". However, in order to minimize the linear objective via a quadratic formulation, Sigl et al. use a dynamic net model which requires an iterative algorithm. Thus, the method in [10] alternately solves a quadratic program and updates the coefficients of the program until a convergence criterion is satisfied. By contrast, the next section proposes a heuristic which empirically allows direct computation of a good (in the sense of more closely reflecting the minimum wirelength objective) linear placement via, e.g., a Lanczos sparse-matrix code [4] [5].

3 Improving the Quadratic Objective Function

In this section, we briefly outline intuition leading to an enhancement of the traditional quadratic objective for linear placement. Recall from the discussion of Section 2 that we would like to minimize $z = \sum_{i,j} c_{ij}|x_i - x_j|$, i.e., the weighted sum of wirelengths. However, due to complexity reasons we prefer a quadratic objective function which is amenable to global optimization.

To minimize z, we could equivalently minimize the *square* of z if we hope to use a quadratic objective function. However, the expansion of z^2 is the highly complicated expression

$$\begin{aligned} c_{12}^2|x_1 - x_2|^2 + c_{13}^2|x_1 - x_3|^2 + \cdots \\ 2c_{12}c_{13}|x_1 - x_2||x_1 - x_3| + \cdots \end{aligned}$$

which can be rewritten as

$$z^2 = \sum_{i=1}^{n}\sum_{j=1}^{n} c_{ij}^2|x_i - x_j|^2 \;+\; \Phi \; (\equiv \mathit{mixed\ terms}).$$

Ignoring the sum of mixed terms Φ in this expression for z^2 leads us to a more natural quadratic minimization objective for linear placement:

$$z' = \sum_{i=1}^{n}\sum_{j=1}^{n} {c_{ij}}^2|x_i - x_j|^2$$

and we suspect that minimizing z' can be shown to be nearly equivalent to minimizing z^2, i.e., minimizing z, for a large class of problems. Certainly, the experimental results in the next section demonstrate that this simple modification leads to significant wirelength and channel width savings.

4 Results and Conclusions

Given the weighted clique model, we solve the associated linear placement problem using the standard eigenvector-based approach outlined in Section 2 above, using a Lanczos-based implementation reported in [5]. Given a graph representation G for the netlist, the eigenvector corresponding to the second-smallest eigenvalue of $Q(G) = D(G) - C(G)$ gives a linear ordering of the modules. We simply evaluate the sum of net spans (i.e., total wirelength) and the channel width induced by this linear ordering.

Tables 1 and 2 show results for linear placement of a number of MCNC benchmarks (the Primary and Test suites), as well as for additional industry netlists (three ILLIAC boards and two benchmarks obtained from Hughes Aircraft Co.) that were evaluated in [13] and [5]. Using the standard clique net model, where each k-pin net containing modules v_i and v_j contributes $\frac{1}{k-1}$ to the value c_{ij}, we obtain results that reflect the methods of [12], [2] and other recent work. However, when we square the final c_{ij} values and apply the same quadratic optimization techniques, we reduce both wirelength and maximum channel width by an average of 7% and 18%, respectively. Since each of these metrics is representative of overall layout area, we believe that these improvements have a great deal of practical significance.

These results motivate a number of interesting open issues. For example, it is possible that alternate functions of the c_{ij} (i.e., other than ${c_{ij}}^2$) can be applied to transform the standard quadratic objective.[2] As long

where I is the identity matrix. This gives an eigenvalue formulation for λ, and the eigenvectors of Q are the only non-trivial solutions for $\vec{x}$. The minimum eigenvalue 0 gives the uninteresting solution $\vec{x} = (1/\sqrt{n}, 1/\sqrt{n}, \ldots, 1/\sqrt{n})$, and hence the eigenvector corresponding to the second smallest eigenvalue λ is used. Again, note that the slot constraint is replaced by the requirement that $|\vec{x}| = 1$.

[2]It is interesting to note that very early work, e.g., that of Steinberg [11], proposed changing the exponent of the $|x_i - x_j|$ term in the objective function. However, no work has examined any alternate exponents for the c_{ij}.

Test problem	Number of elements	Wirelength		Gain (%)
		c_{ij}	c_{ij}^2	
IC67	67	2311	2204	4.63
IC116	115	4561	4790	-5.02
IC151	151	6809	6351	6.73
bm1	882	56739	52822	6.90
19ks	2844	569362	403835	29.07
Prim1	833	57991	52267	9.87
Prim2	3014	1029741	803172	22.00
Test02	1663	316837	377974	-19.30
Test03	1607	168994	130070	23.03
Test04	1515	175023	225495	-28.84
Test05	2595	573500	563515	1.74
Test06	1752	435829	294939	32.33

Table 1: Results for twelve industry benchmarks showing wirelength values when c_{ij} and c_{ij}^2 objective functions are used with an eigenvector method to yield the linear placement. Average wirelength improvement is 7%.

Test problem	Number of elements	Channel Width		Gain (%)
		c_{ij}	c_{ij}^2	
IC67	67	56	54	3.57
IC116	115	71	64	9.86
IC151	151	83	76	8.43
bm1	882	128	114	10.94
19ks	2844	470	299	36.38
Prim1	833	144	121	15.97
Prim2	3014	583	382	34.48
Test02	1663	364	389	-6.87
Test03	1607	166	130	21.69
Test04	1515	213	237	11.27
Test05	2595	458	354	22.71
Test06	1752	518	252	51.35

Table 2: Results for twelve industry benchmarks showing channel width values when c_{ij} and c_{ij}^2 objective functions are used with the eigenvector method to yield the linear placement. Average reduction in channel width is 18%. As with the wirelength results, note that the gains seem larger as problem size increases.

as fixed c_{ij} are used, the usual efficient quadratic optimization algorithms remain applicable. In fact, it is quite reasonable to envision a placement methodology which tests a number of alternate functions of the c_{ij}, then returns the best result.

In conclusion, we have retraced the traditional attack on linear placement of modules in cell-based layout. The basic observation is that a quadratic objective amenable to global minimization is usually substituted for the true linear wirelength objective; however, the standard quadratic form does not have a strong relationship to the original objective. Thus, we propose a simple modification of the weighting function used in the quadratic form of the objective. This allows global minimization of a function that more closely reflects the minimum wirelength goal. Indeed, experimental results confirm very significant savings in both wirelength and channel width. Since these metrics reflect layout area, we believe that our proposed modification will be highly useful within the context of existing placement methods.

References

[1] D. Adolphson and T.C. Hu, "Optimal Linear Ordering", *SIAM J. Appl. Math.* 25 (1973), pp. 403–423.

[2] C.K. Cheng and E.S. Kuh, "Module Placement Based on Resistive Network Optimization", *IEEE Trans. on CAD* (1984), pp. 218–225.

[3] M.R. Garey and D.S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, Freeman, 1979.

[4] G.H. Golub and C.F. Van Loan, *Matrix Computations*, Johns Hopkins, 1983.

[5] L. Hagen and A.B. Kahng, "Fast Spectral Methods for Ratio Cut Partitioning and Clustering", *Proc. IEEE Intl. Conf. on Computer-Aided Design*, 1991, pp. 10–13.

[6] K.M. Hall, "An r-dimensional Quadratic Placement Algorithm", *Management Science* 17(1970), pp. 219–229.

[7] M. Hanan, P.K. Wolff and B.J. Agule, "A Study of Placement Techniques", *J. Design Automation and Fault-Tolerant Computing* 2 (1978), pp. 28–61.

[8] F. Kurdahi and S. Sastry, "On the Applicability of the Weibull Model to Characterize Wire Length Distributions in Standard Cell Layouts", *Proc. European Design Automation Conf.*, 1991.

[9] T. Lengauer, *Combinatorial Algorithms for Integrated Circuit Layout*, Wiley-Teubner, 1990.

[10] G. Sigl, K. Doll and F.M. Johannes, "Analytical Placement: A Linear or a Quadratic Objective Function?", *Proc. ACM/IEEE Design Automation Conf.*, June 1991, pp. 427–432.

[11] L. Steinberg, "The Backboard Wiring Problem: A Placement Algorithm", *SIAM Review* 3(1) (1961), pp. 37–50.

[12] R.S. Tsay, E.S. Kuh and C. Hsu, "PROUD: A Fast Sea-of-Gates Placement Algorithm", *Proc. ACM/IEEE Design Automation Conf.*, 1988, pp. 318–323.

[13] Y.C. Wei and C.K. Cheng, "Ratio Cut Partitioning for Hierarchical Designs", *IEEE Trans. on CAD* 10(7) (1991), pp. 911–921.