

# Bisection Based Placement for the X Architecture

Satoshi Ono

Computer Science Department  
SUNY Binghamton  
Box 6000  
Binghamton NY 13902  
satoshi@binghamton.edu

Sameer Tilak

San Diego  
Supercomputer Center  
sameer@sdsc.edu

Patrick H. Madden

Computer Science Department  
SUNY Binghamton  
Box 6000  
Binghamton NY 13902  
pmadden@acm.org

## ABSTRACT

Rising interconnect delay and power consumption have motivated the investigation of alternative integrated circuit routing architectures. In particular, the X Architecture, which features preferred routing in diagonal directions, has gained a measure of industry support, and has even been validated at 65nm.

While there has been extensive study of Manhattan design methods, there are markedly fewer published results for non-Manhattan design. To help fill this gap, we study a patented placement method for the X Architecture; to our knowledge, there have been no prior published results for the method. Surprisingly, we find that the patented method in fact performs *worse* than traditional Manhattan methods – for both Manhattan and X routing metrics. We also present a theoretic formulation which explains why solution quality is degraded.

Many groups in industry are evaluating the merits of non-Manhattan routing architectures. By providing concrete experimental results, we hope to improve the accuracy of these evaluations.

## I. INTRODUCTION

As design sizes have increased over the years, interconnect delay has taken a progressively larger portion of total system delay. Interconnect also contributes to power consumption – long wires add switching capacitance and require large drivers. Power and delay are limiting performance of modern designs.

As a means to reduce interconnect lengths, the *X Routing Architecture*[18] was introduced. Under this routing model, wiring at 45-degree angles is allowed; this can result in lower wire lengths, and consequently lower delay and power consumption. The routing model has attracted some support from industry, and a few designs using the routing model have been produced.

It is clear that Steiner tree lengths for nets should improve with the introduction of diagonal wiring; an “X” based Steiner tree can never be longer than a Manhattan Steiner tree. The benefit *in practice*, however, can be difficult to measure. The X architecture consumes routing resources differently than Manhattan designs, and many tools in the design flow are not opti-

mized to support it.

Rectilinear design has been well studied by academia, and there are a number of established benchmarks for comparison of results; in this paper, our objective is to bring a similar level of rigor to non-Manhattan design. In particular, we consider a placement method for the X architecture patented by Teig and Ganley[17]. While the approach has been disclosed, to our knowledge there has been no published independent evaluation of the approach.

The patented method uses a recursive bisection framework, but can orient cut lines diagonally as a means to prefer diagonal wiring. In experiments performed with our implementation of this algorithm, we find a surprising result. While the method does result in placements that utilize an X routing architecture well, total wire lengths are in fact worse than “traditional” Manhattan-optimizing placement methods.

To be clear on this issue: from our experiments, we see no indication that a “diagonal cut line” bisection based placement method is superior to a traditional rectilinear bisection based placement method. Rather, we observe that diagonal cut lines in some sense put Manhattan routing architectures at a disadvantage.

Industry groups are actively evaluating non-Manhattan routing as a means to combat rising interconnect delay and power challenges. Our objective is to provide accurate and impartial experimental results. Design teams face a difficult decision when evaluating non-Manhattan routing architectures; while there appear to be some advantages, it is difficult to evaluate how great the advantages are, and if they outweigh disadvantages.

The remainder of this paper is organized as follows. We first describe routing metrics, and traditional placement techniques. Relatively little work has been published on non-Manhattan placement; we discuss this briefly. We next describe details of our implementation of the patented method, which uses diagonal cut lines to make better use of diagonal wiring resources. Experimental results are performed on a set of placement benchmarks commonly used by academic groups. The results may be somewhat unexpected; we present some observations on circuit structure and the shapes of regions being partitioned. These observations provide some intuition as to why




Fig. 1. The wire length required to connect a set of pins depends greatly on the routing architecture. For a three-pin net, we show (a) a rectilinear Minimum Spanning Tree, (b) a rectilinear Steiner Minimal Tree, and (c) an X Steiner Minimal Tree. With added routing directions, an X-based tree should never have higher length than a rectilinear tree.

this placement quality is not improved. We then conclude the paper with a discussion of future work.

## II. PREVIOUS WORK

In this section, we first discuss routing architectures, and then focus on placement techniques.

### A. Routing Architectures

For most of the history of integrated circuit design, wiring has been rectilinear, and is frequently gridded[13]. Non-rectilinear routing had been restricted to small jogs or transistor shaping within a cell[19, 12]. Non-Manhattan channel routing has also been investigated[16, 5]

The first proposal for the large scale use of non-Manhattan interconnect as a means to address delay and power was [9]. In this paper, two routing architectures were proposed; the “octilinear” approach has become better known as the “X” architecture, while the “hexagonal” approach has become known as the “Y” architecture. For the two routing architectures, wire length optimizing Steiner tree algorithms, a performance-driven required arrival time Steiner tree algorithm, and a global router were presented. Placements of standard cell benchmarks were evaluated, with the alternative architectures indicating potential wire length improvements.

Following this, an industry group, the X Initiative[20], has been organized to address research and fabrication challenges for octilinear non-Manhattan design. While there are a number of apparent advantages[18], only a few designs have been manufactured using the new routing model.

The “X Architecture” features lower interconnect metal layers in a traditional rectilinear configuration – this simplifies integration with existing tools. The upper interconnect layers are rotated diagonally, allowing shorter connections. Additionally, industrial routing tools developed for the X Architecture use a great deal of non-preferred direction routing; there are diagonal jogs on all layers, resulting in fewer vias and reduced wire lengths.

### B. Placement Techniques

As most routing architectures have been rectilinear, placement tools have been developed to take advantage of it. Traditional recursive bisection methods orient cut lines along the X and Y axis; we discuss the impact of this choice below.

Many analytic placement methods optimize horizontal and vertical wiring separately – this captures the wire length minimization objective for Manhattan designs, but is incorrect for non-Manhattan objectives. Annealing based methods are quite flexible and can support arbitrary objective functions; unfortunately, annealing does not scale as well as other methods, and high run times make it unattractive for large designs.

For non-Manhattan routing objectives, there are only a few placement results. In [4], the authors investigated both 45 and 60-degree routing metrics, using an annealing based placer with a Steiner-length optimization objective. When comparing to a Manhattan-length optimization, improvements ranging from 8.92% to 11.48% were observed. Unfortunately, this method does not appear to scale well; all designs considered in this work were small (less than 1500 nets), cells were modified to have unit area, and only five pins of high-degree nets were considered.

A hexagonal block floorplanner[10] has also been developed; this work focused primarily on efficiency of packing, and did not address wire length and routing architecture issues.

## III. BISECTION BASED PLACEMENT FOR THE X ARCHITECTURE

With traditional rectilinear design, it has been observed that placement tools produce results that are optimized for the Manhattan metric[9]. With random positioning of the pins on an interconnect net, one might expect that “X” routing would obtain a wire length savings of roughly 17% over Manhattan routing. In fact, [9] observed much less improvement. On average, the savings of X-based Steiner trees was only 10% over Manhattan Minimum Spanning Trees; relative to Manhattan Steiner trees, the savings was only around 4%.

Clearly, to gain the maximum advantage of a non-Manhattan routing architecture, the placement must be “tuned” to that architecture. The experiments of [4] indicate that this can be done; for industrial designs, however, the method must scale well. In this section, we discuss our implementation of the patented method by Teig and Ganley[17]; as it is based on recursive bisection, the method scales well.

Recursive bisection based placement is a classic approach[3]. In general, the method begins with a circuit netlist and a placement “region.” The netlist is repeatedly partitioned, while the placement region is repeatedly divided. Concepts such as terminal propagation[6] and cycling of partitions[7] are widely used; modern academic tools that are based on recursive bisection include *Capo*[11] and *feng shui*[2].

The location and orientation of cut lines during partitioning is known to impact results[15]. Using theory based on Rent’s Rule[14], a strategy of cut line sequencing was developed[21]. The contribution of [17] was the proposal to use diagonal cut lines.

### A. Motivation for Diagonal Cut Lines

With recursive bisection based placement, it has been observed that the lengths and orientations of interconnect wires are influenced by the direction of cut lines. The basic principle

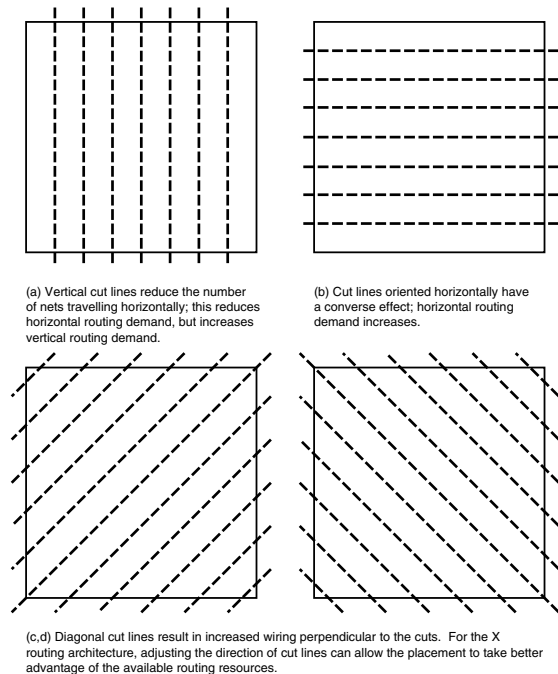


Fig. 2. The direction of cuts in bisection-based placement influences the routing demand. Early vertical cuts reduce horizontal demand (and increase vertical demand). The intention of diagonal cut lines is to prefer routing in non-Manhattan directions, taking better advantage of X Architecture routing resources.

is easy to understand: consider the examples shown in Figure 2(a) and (b). If one has early cut lines oriented vertically, only the nets cut in the first partitionings travel across the cut lines (horizontally). The number of nets cut in early passes is usually small compared to the total number of nets – and thus, the bulk of wiring in example (a) will be vertical. Wiring demand can be shifted either horizontally or vertically by adjusting the direction of cuts.

The motivation for using diagonal cut lines is to adjust routing demand such that better use is made of diagonal routing layers. The limited wire length improvement of X-based Steiner trees on Manhattan-driven placements makes it clear that many pins are either vertically or horizontally aligned – and thus get no benefit from diagonal resources.

### B. Implementation Details

We refer to our implementation of the patented method as *XPlace*; in many respects, it is a traditional recursive bisection placement tool. In our implementation, cuts are always vertical, horizontal, or at 45 degrees. We use a number of geometric techniques to simplify the management of space; with rectilinear bisection, regions are always rectangular. With diagonal cut lines, placement regions may be octagonal, requiring slightly more complex area and bisection computations. The area of a convex polygon can be computed easily. If the corner coordinates  $(x, y)$  are stored in counterclockwise order,




Fig. 3. Different cut sequences produced by our non-Manhattan placement tool. Command line options select from a traditional Manhattan approach, exclusive use of diagonal cut lines, or a mix (Manhattan followed by diagonal, or vice versa).

$$\begin{aligned}
 \text{Area} &= \frac{1}{2} \begin{vmatrix} x_0 & y_0 \\ x_1 & y_1 \\ \vdots & \vdots \\ x_n & y_n \\ x_0 & y_0 \end{vmatrix} \\
 &= ((x_0 \cdot y_1 + x_1 \cdot y_2 + \dots + x_n \cdot y_0) \\
 &\quad - (y_0 \cdot x_1 + y_1 \cdot x_2 + \dots + y_n \cdot x_0)) / 2
 \end{aligned}$$

When splitting a region with a bisection, the partitioner (in our implementation, *hMetis*[8]) returns two areas. To split an arbitrary octagonal region, we can use a closed-form computation.

### C. Cut Sequences

Our tool can use either alternating Manhattan cuts, alternating diagonal cuts, or a mix of both. One can estimate the number of “levels” of bisection required by taking the logarithm of the number of cells in the design. When switching from Manhattan to diagonal, we explored having the first 50% of cuts be oriented in a rectilinear fashion, and the remaining cuts oriented diagonally (we refer to this as “M+X”). We also used the converse, “X+M.”

## IV. EXPERIMENTAL RESULTS

To evaluate the merits of the patented method, we performed experiments with the “IBM version 2” benchmarks that are in wide use by academic research groups[1]. The naming scheme of the benchmarks may be confusing to those not familiar with them. Using an original set of 18 partitioning benchmarks, the authors of *Dragon* converted a subset for use in standard cell placement; we use all the converted benchmarks, and have not omitted any designs from this suite.

Benchmark	Our Tool				Results from [1]				
	M	X	M+X	X+M	KraftWerk	Capo8.6	Dragon2.23	feng shui 2.0	mPL 2.0
ibm01	.52	.67	.65	.59	.703	.549	.582	.524	.64
ibm02	1.53	1.83	1.81	1.67	2.15	1.59	1.58	1.47	1.61
ibm07	3.39	4.18	4.17	3.74	5.12	3.70	3.59	3.30	4.07
ibm08	3.73	4.78	4.77	4.10	4.66	3.84	3.82	3.66	4.25
ibm09	3.10	3.90	3.85	3.47	4.26	3.22	3.20	3.01	3.81
ibm10	5.76	7.48	7.26	6.31	7.61	6.15	6.02	5.67	6.61
ibm11	4.60	5.61	5.48	5.04	5.80	4.85	4.72	4.59	5.96
ibm12	8.04	9.79	9.48	8.67	10.41	8.58	8.58	7.75	9.44

TABLE I

MANHATTAN HALF-PERIMETER WIRE LENGTHS FOR PLACEMENTS PRODUCED USING MANHATTAN, X, AND MIXED CUTS. FOR REFERENCE, WE INCLUDE RESULTS FROM OTHER RECENT ACADEMIC TOOLS. THE PLACEMENTS PRODUCED BY OUR TOOL ARE DENSELY PACKED, AND LIKELY UNROUTABLE. THE RESULTS SHOW THAT OUR IMPLEMENTATION IS COMPARABLE TO OTHER RECENT WORK (WE NOTE THAT MANY OF THESE PLACEMENT TOOLS HAVE IMPROVED RESULTS, OR ARE FOCUSED ON SUCCESSFUL ROUTING, NOT SIMPLY WIRE LENGTH MINIMIZATION).

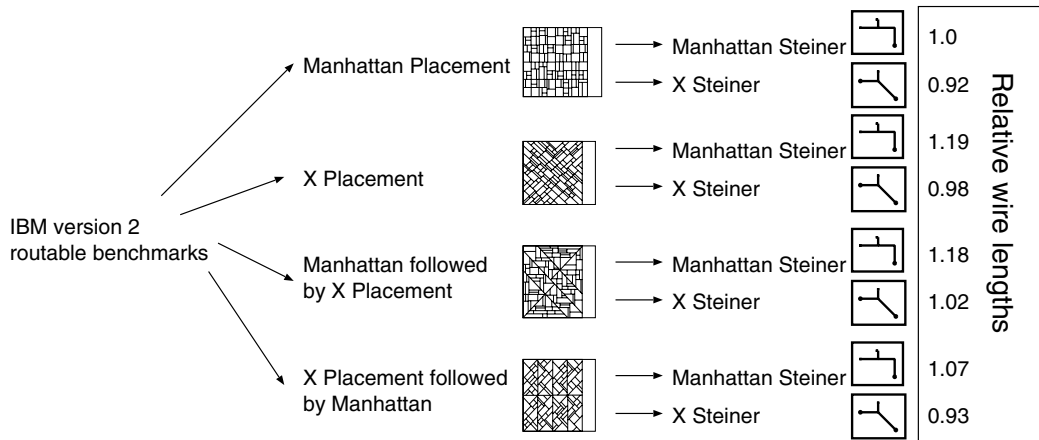


Fig. 4. An overview of the experiments performed; we used a common set of placement benchmarks, and tested a variety of cut sequences using our tool *XPlace*. We then compared wire lengths using both Manhattan and non-Manhattan Steiner tree metrics (wire lengths are higher than the half-perimeter results shown in Table I).

### A. Comparisons with Traditional Tools

First, to demonstrate that our placement tool is not a naive or fundamentally flawed implementation, we compare our results to those of other recent academic tools, using a rectilinear half-perimeter metric. Table I shows these results; the Manhattan cut sequences are competitive with other tools. When diagonal cut lines are employed, wire lengths increase—this is to be expected, as more wiring is oriented diagonally, and would have a higher length under the Manhattan metric.

Our placements are densely packed; we use the tool *feng shui 5.1* to perform legalization. This results in somewhat reduced wire lengths compared to other tools (which are optimizing routability as well as wire length).

### B. Evaluation of Steiner Tree Lengths

The experiments performed are illustrated in Figure 4; exact wire length results are shown in Table II. We explored different combinations of cut sequences, and evaluated wire lengths using both Manhattan and non-Manhattan Steiner trees. The top portion of the table contains Steiner wire lengths prior to legalization; the bottom shows results after legalization and detailed placement by *feng shui 5.1*.

### C. Interpretation of Results


From the experiments, it is clear that the use of diagonal cut lines changes the orientation of much of the circuit interconnect. For an “X” placement, the average wire length of an “X” Steiner trees is lower than that of a Manhattan Steiner tree, showing that the diagonal routing layers are more frequently used. The “X” placement has much higher wire length when routing is restricted to Manhattan directions – roughly 19% higher on average.

When compared to a traditional Manhattan placement, however, the advantages are not as compelling. Using “X” routing on a Manhattan placement gives an 8% improvement; with an “X” placement, the improvement is only 2%.


The lack of overall improvement can be tied to two issues. First, the diagonal cut lines are at best “rotating” the placement – as the “X” Steiner trees support 45-degree rotations transparently, one can hope for at best results equal to those obtained with Manhattan placement. We believe the wire length loss to be attributed to the shapes of regions generated during bisection; [21] showed that having regions that were excessively rectangular resulted in increased wire lengths, and the diagonal cut lines can have a similar effect. Figure 5 illustrates these points; when “X” Steiner trees are used, the wire lengths are at best equal if diagonal cut lines are used, and in many cases can be anticipated to be worse. While this was not our initial expectation, we believe that this analysis explains our experimental results.

## V. CONCLUSION

Circuit designs are pushed to the limit to maximize performance. To remain competitive in the semiconductor industry, no stone can remain unturned in the pursuit of faster, lower power devices. Interconnect wiring has been a bottleneck in



(a) If the entire placement region is rotated, Manhattan and diagonal cut lines will produce equivalent placements. If X routing is utilized, the tree lengths will be equal.



(b) Two regions with equal areas can have different distances to the opposite side of a partition. The dark shaded portion of the triangular region is “missing,” being replaced by equal lighter-shaded areas. The distance to the opposite side of the cut line is greater in the triangular configuration, resulting in higher wire length.

Fig. 5. (a) If one were to begin with a diamond-shaped placement region, using diagonal cut lines would result in a placement that was rotated at 45 degrees to a traditional Manhattan design; in all other respects, the placements would be identical. (b) The average distance between cells on opposite sides of a partition is greater if the region being partitioned is triangular.

design for many years – contributing significantly to both delay and power consumption.

While non-Manhattan wiring clearly can reduce interconnect lengths in principle, it is essential to know how much benefit it provides in realistic situations. Adoption of a new routing architecture complicates the design process; those in industry must examine carefully if the benefits outweigh the risks.

Our objective with this paper is to avoid being overly optimistic or pessimistic about the new routing architecture. To our knowledge, there have been no publically released experimental results for the method described in [17]; we have implemented this approach, and performed experiments on widely available benchmarks, so that our industry colleagues can make informed decisions. Our implementation is publically available and uses standard file formats, so that others may perform their own experiments.

From our experiments, one may expect an 8% reduction in Steiner tree lengths on Manhattan-based placements; this is comparable to the results reported earlier in [9]. Industry researchers are best suited to determine if 8% is “enough” to warrant the adoption of the new routing architecture. While the results of [4] are more encouraging, it is not clear if they extend to non-trivial circuit sizes.

As part of our current work, we are investigating alternative methods to optimize a circuit placement for non-Manhattan routing architectures. The 8% improvement is in some respects a conservative estimate; it is reasonable to expect that a properly tuned placement can achieve better results. We are also actively developing routing tools, and will report on the integrated set of tools when the work is complete.

Bench mark	Global Placement								Legalized Placement							
	Manhattan Steiner				X Steiner				Manhattan Steiner				X Steiner			
	M	X	M+X	X+M	M	X	M+X	X+M	M	X	M+X	X+M	M	X	M+X	X+M
ibm01	.620	.744	.736	.674	.570	.567	.630	.646	.622	.751	.718	.644	.567	.617	.633	.574
ibm02	1.82	2.04	2.05	1.92	1.68	1.73	1.80	1.70	1.82	1.98	2.01	1.86	1.67	1.72	1.77	1.68
ibm07	3.79	4.50	4.56	4.10	3.48	3.73	3.94	3.57	3.79	4.37	4.48	3.96	3.47	3.71	3.88	3.54
ibm08	4.46	5.41	5.48	4.77	4.09	4.50	4.73	4.18	4.44	5.26	5.39	4.62	4.07	4.47	4.68	4.14
ibm09	3.55	4.23	4.24	3.86	3.24	3.48	3.65	3.34	3.55	4.10	4.17	3.73	3.23	3.46	3.60	3.31
ibm10	6.54	8.09	7.97	7.00	5.96	6.68	6.87	6.08	6.54	7.87	7.85	6.79	5.93	6.66	6.79	6.02
ibm11	5.05	5.95	5.85	5.44	4.62	4.89	5.03	4.72	5.05	5.77	5.74	5.26	4.60	4.87	4.96	4.67
ibm12	9.07	10.6	10.4	9.57	8.28	8.77	8.95	8.36	9.07	10.3	10.3	9.34	8.25	8.74	8.87	8.31
avg.	1.0	1.19	1.18	1.07	0.92	0.98	1.02	0.93	1.0	1.15	1.16	1.04	0.91	0.98	1.01	0.93

TABLE II

COMPARISON OF TOTAL STEINER TREE LENGTHS FOR BOTH MANHATTAN AND X ROUTING METRICS, USING PLACEMENTS DERIVED FROM RECURSIVE BISECTION USING MANHATTAN, X, AND A MIX OF BOTH TYPES OF CUTS. X+M INDICATES MANHATTAN CUTS FOLLOWED BY DIAGONAL, WHILE M+X INDICATES THE REVERSE. LEGALIZATION AND DETAILED PLACEMENT WERE PERFORMED BY *feng shui 5.1*.

## REFERENCES

- [1] S. N. Adya, M. C. Yildiz, I. L. Markov, P. G. Villarrubia, P. N. Parakh, and P. H. Madden. Benchmarking for large-scale placement and beyond. *IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems*, 23(4):472–487, 2004.
- [2] A. R. Agnihotri, S. Ono, C. Li, M. C. Yildiz, A. Khatkhate, C.-K. Koh, and P. H. Madden. Mixed block placement via fractional cut recursive bisection. *IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems*, 24(5):748–761, 2005.
- [3] M. A. Breuer. A class of min-cut placement algorithms. In *Proc. Design Automation Conf*, pages 284–290, 1977.
- [4] H. Chen, C.-K. Cheng, A. B. Kahng, I. Mandoiu, and Q. Wang. Estimation of wirelength reduction for  $\lambda$ -geometry vs. Manhattan placement and routing. In *Proc. System Level Interconnect Prediction Workshop*, pages 71–77, 2003.
- [5] S. Das and B. Bhattacharya. Channel routing in manhattan-diagonal model. *Int'l Conf. on VLSI Design*, pages 43–48, 1996.
- [6] A. E. Dunlop and B. W. Kernighan. A procedure for placement of standard-cell VLSI circuits. *IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems*, CAD-4(1):92–98, January 1985.
- [7] D. J.-H. Huang and A. B. Kahng. Partitioning based standard cell global placement with an exact objective. In *Proc. Int. Symp. on Physical Design*, pages 18–25, 1997.
- [8] G. Karypis, R. Aggarwal, V. Kumar, and S. Shekhar. Multilevel hypergraph partitioning: Application in VLSI domain. In *Proc. Design Automation Conf*, pages 526–529, 1997.
- [9] C.-K. Koh and P. H. Madden. Manhattan or non-Manhattan? a study of alternative VLSI routing architectures. In *Proc. Great Lakes Symposium on VLSI*, pages 47–52, 2000.
- [10] J. Li, T. Yan, B. Yang, J. Yu, and C. Li. A packing algorithm for non-Manhattan hexagon/triangle placement design by using an adaptive O-tree representation. In *Proc. Design Automation Conf*, pages 646–651, 2004.
- [11] I. L. Markov. Seeing the forest and the trees: Steiner wirelength optimization in placement. In *Proc. Int. Symp. on Physical Design*, page to appear, 2006.
- [12] D. Marple, M. Smulders, and H. Hegen. Tailor: A layout system based on trapezoidal corner stitching. *IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems*, 9(1):66–90, 1990.
- [13] C. Mead and L. Conway. *Introduction to VLSI Systems*. Addison-Wesley, 1993.
- [14] C. E. Radke. A justification of, and an improvement on, a useful rule for predicting circuit-to-pin ratios. In *Proc. Design Automation Conf*, pages 257–267, 1969.
- [15] Kazuhiro Takahashi, Kazuo Nakajima, Masayuki Terai, and Koji Sato. Min-cut placement with global objective functions for large scale sea-of-gates arrays. *IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems*, 14(4):434–446, April 1995.
- [16] X. Tan and X. Song. Hexagonal three-layer channel routing. *Information Processing Letters*, 55:223–228, 1995.
- [17] S. Teig and J. L. Ganley. US patent 6,848,091: Partitioning placement method and apparatus, 2002.
- [18] S. L. Teig. The X Architecture: not your father's diagonal wiring. In *Proc. System Level Interconnect Prediction Workshop*, pages 33–38, 2002.
- [19] J. Waterkamp, R. Wicke, R. Bruck, M. Reinhardt, and G. Schrammeck. Technology tracking of non manhattan VLSI layout. In *Proc. Design Automation Conf*, pages 296–301, 1989.
- [20] X Initiative. Home page. <http://www.xinitiative.org>.
- [21] M. C. Yildiz and P. H. Madden. Improved cut sequences for partitioning based placement. In *Proc. Design Automation Conf*, pages 776–779, 2001.