

Are Floorplan Representations Important In Digital Design?

Hayward H. Chan[†], Saurabh N. Adya[‡] and Igor L. Markov[†]

[†]The University of Michigan, Department of EECS, 1301 Beal Ave., Ann Arbor, MI 48109-2122

[‡]Synplicity Inc., 600 W. California Ave, Sunnyvale, CA 95054

{hhchan, imarkov}@umich.edu, saurabh@synplicity.com

ABSTRACT

Research in floorplanning and block-packing has generated a variety of data structures to represent spatial configurations of circuit modules. Much of this work focuses on the geometry of module shapes and seeks tighter packing, as well as improvements in the asymptotic worst-case complexity of algorithms for standard tasks. In this work we consider the implications of interconnect optimization on the value of floorplan representations and establish a framework for comparing different representations. By analyzing performance bottlenecks in block packing and properties of floorplan representations, we show that many of the mathematical results in floorplanning do not translate into better VLSI layouts. This is confirmed by extensive empirical data for stand-alone floorplanners and integrated applications.

Categories and Subject Descriptors: B.7.2 [Integrated Circuits]: Design Aids — *placement and routing*; G.4 [Mathematical Software]: Algorithm Design and Analysis; J.6 [Computer-Aided Engineering]: Computer-Aided Design.

General Terms: Algorithms, experimentation.

Keywords: Circuit layout, floorplanning, sequence pair, B*-tree.

1. INTRODUCTION

Floorplanning has traditionally been important in VLSI design because it determines the top-level spatial structure of a chip. Today automatic floorplanning is encouraged by the growing adoption of embedded memories and IP blocks in SoC designs. While human designers may have non-trivial insights into design objectives, they cannot analyze millions of possible spatial configurations.

A floorplan can be represented by the locations of the blocks, as in [10], but this complicates the generation of new overlap-free floorplans. We note that VLSI floorplanners typically rely on local search, especially simulated annealing, and spend most of their runtime to incrementally modify and evaluate candidate floorplans. Therefore, topological representations are more common as they guarantee that all encoded packings are overlap-free. Such a representation encodes relative positions among blocks in a way that is amenable to perturbation. For example, the sequence pair [16] captures a packing by a pair of permutations and can be modified in $O(1)$ time, shifting the computational burden to the retrieval of block locations. Different representations are often compared by the algorithmic complexity of their *evaluation*.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ISPD'05, April 3–6, 2005, San Francisco, California, USA.

Copyright 2005 ACM 1-59593-021-3/05/0004 ...\$5.00.

Table 1: Published results of area-minimization on MCNC benchmarks of different representations, except for several results for CBL and ECBL that are incorrect (INC), according to the authors. Results known to be optimal are boldfaced.

Floorplan representation	Published results				
	<i>apte</i>	<i>xerox</i>	<i>hp</i>	<i>ami33</i>	<i>ami49</i>
Optimal [4]	46.9	19.8	8.95	time-out	
Sequence pair [22]	46.9	19.8	8.95	1.21	36.8
TCG [13]	46.9	19.8	8.95	1.20	36.8
TCG-S [14]	46.9	19.8	8.95	1.19	36.4
O-tree [18]	46.9	20.2	9.16	1.24	37.7
B*-tree [5]	46.9	19.8	8.95	1.27	36.8
Corner-block list (CBL) [8]	INC	19.8	INC	1.18	36.1
ECBL [27]	INC	19.9	INC	1.19	36.7
Twin-binary sequence [26]	47.2	19.8	9.03	1.21	37.0
TBS extended [26]	47.4	19.8	9.02	1.19	36.9
Q-sequence [29]	46.9	19.9	9.03	1.19	36.8
ACG [28]	46.9	19.9	9.03	1.19	36.8

Slicing floorplans are those that can be recursively bisected by horizontal and vertical cut-lines down to single blocks. They can be encoded by slicing trees or Polish expressions [23]. The first encoding of an arbitrary floorplan — a sequence pair — appeared 10 years ago [16]. Today more than a dozen representations exist [5, 7, 8, 13, 14, 22, 26, 27, 28, 29]. Relations among them are studied in [25] and will be summarized in Section 2.

Competitive VLSI floorplanners frequently use simulated annealing because this allows modifying the objective function in applications. They differ in the choice of floorplan representation, which may affect runtime and solution quality. New floorplan representations are typically justified by improvements in the algorithmic complexity of evaluation, the type of encoded floorplans, the amount of redundancy in the encoding and the total number of encoded configurations. To this end, linear- or near-linear time evaluation algorithms are already known for a number of representations (sequence pair, B*-tree, etc), and would be difficult to improve upon if the location of each block is required (a possible exception is incremental evaluation). Many existing representations encode a sufficiently broad range of floorplans.

Significant empirical improvements are also unlikely because popular benchmarks are so small that many publications already report optimal solutions (see Table 1). The use of simulated annealing leads to particularly controversial reporting of results. For example, many papers [5, 13, 14, 29] only report best results out of an unknown number of independent runs. Since temperature schedules are rarely reported, it is conceivable to tune them to individual benchmarks. Needless to say, such results may be difficult to reproduce. For example, the performance of TCG and TCG-S is questioned in [28] since there are significant discrepancies between the experimental results and results reported in [13, 14]. Many representations [5, 22, 28, 29] have not been evaluated in the context

of interconnect optimization, and in other cases, wirelength optimization is shown in a token experiment without much discussion. For example, ACG in [28] is presented as “more suitable for interconnect plan” than existing representations, but no results on interconnect are reported. This state of confusion reminds of the critical analysis of VLSI placement literature in [15], and our work is motivated similarly. However, a major difference is that many results in floorplanning have been mathematically proven, therefore our focus is on relevance and applicability rather than correctness.

Larger-scale layout is a good context for justifying improvements in asymptotic complexity of optimization algorithms, and excellent area-packing results were reported recently for benchmarks with over 100 blocks [4, 12]. However, these results have little or no impact on interconnect minimization. The work in [2] suggests that in fixed-outline interconnect minimization [9], annealing is outperformed for over 100 blocks by a hybrid algorithm that uses min-cut partitioning where possible and resorts to annealing-based packing when necessary. Whether or not the choice of topological floorplan representations is significant for interconnect-driven floorplanning is one of the questions addressed in our work.

In this work we group existing floorplan representations into families of related data structures that have identical “solution spaces” and/or share similar evaluation algorithms. To avoid redundant comparisons, we then select (i) sequence pair [16] that is used in the well-known tool Parquet [1] and is related to TCG and TCG-S [14], as well as (ii) B*-tree [24], which is similar to O-tree [18] and for which a multi-level extension has been reported [12]. Several unrelated representations (e.g., mosaic floorplans) are excluded from our comparison because their applicability is more restricted. We replace sequence pair with B*-tree in the annealing-based floorplanner Parquet [1] so as to compare the two representations using exact same temperature schedule. Additionally, making fair and realistic comparisons requires several technical results, which we develop in this work. Empirical data show that the size of the solution space of a representation and worst-case complexity of its evaluation are not very relevant to actual performance. To scale our analysis to mixed-size placement, we embed floorplanning comparisons into the framework of *min-cut floorplacement* that additionally involves min-cut partitioning [2] and has been demonstrated on up to 220K movable objects. In this framework, the layout is first partitioned into subregions, minimizing the interconnect between them. Fixed-outline floorplanning is invoked in regions with relatively large macros in it. In our work floorplacement is used for benchmarking purposes as a practical application of floorplanning where both packing and interconnect optimization are significant. In terms of stand-alone evaluation of block packers, floorplacement generates a large number of diverse floorplanning instances allowing an unbiased comparison of different block packing techniques.

In the remaining part of the paper, Section 2 describes and compares existing floorplan representations. We outline our evaluation framework in Section 3, present and analyze experimental results in Section 4 and conclude with a discussion in Section 5.

2. FLOORPLAN REPRESENTATIONS

A key property of popular topological floorplan representations is that they capture at least one area-optimal solution. Other than that, their solution spaces may be quite different. The floorplan in Figure 3c cannot be captured by a B*-tree but can be captured by a sequence pair [16] and by a corner-block list [8]. Some representations may actually capture the exact same set of floorplans, such as sequence pair, TCG and TCG-S. Therefore, we distinguish several families of topological representations.

2.1 Families of representations

Sequence pair [16], TCG and TCG-S [14] are shown to be equivalent in [14] in the sense that they share the same $(n!)^2$ solution space and capture the exact same set of floorplans. Each sequence pair corresponds to one TCG and vice versa. TCG-S is a hybrid of TCG and sequence pair, which contains a constraint graph from TCG and a sequence from sequence pair. Since TCG and TCG-S capture the same set of floorplans and take $O(n^2)$ time to evaluate, we only evaluate sequence pair since (a) it is simpler, and (b) annealing moves takes less time to evaluate. Multiple sequence pairs can represent the same floorplan, as we will see in Section 2.2, and the exact number of general floorplans is in $O(n!2^{5n}/n^{4.5})$, which is much smaller than $(n!)^2$ [20]. This redundancy is typically viewed as a limitation of sequence pair, TCG and TCG-S. However, such an argument is only applicable to exhaustive-search algorithms (as in [4]), but not necessarily local search. The literature on Boolean Satisfiability suggests that redundancies in the solution space, e.g., symmetries [19], make local search more successful by increasing the number of paths leading to desirable configurations (i.e., increasing the basins of attraction of global optima).

Both O-tree [7] and B*-tree [5] use a single tree to represent a horizontally compacted packing (Fig.2), but differ in the bit-level implementation of the tree. O-tree uses a rooted-ordered tree with arbitrary vertex degrees, while B*-tree uses a binary tree. Therefore they share the same $O(n!2^{2n-2}/n^{1.5})$ solution space and capture the same set of floorplans. As pointed out in [5], the regularity of B*-tree makes it more attractive in annealing-based floorplanners.

Another family of floorplan representations captures only floorplans with zero wasted area (*mosaic floorplans*) [8] and includes the corner-block list, twin-binary sequence [25] and Q-sequence [29]. In practical applications with fixed-sized blocks (e.g., embedded memories, datapaths and IP blocks) it is extremely rare that no space is wasted. Therefore some mosaic representations are extended to handle general floorplans by inserting empty rooms. Area-optimality is guaranteed when enough rooms are inserted. ECBL [27], an extension of corner-block list, requires n^2 empty rooms to capture the area-optimal solutions, but this leads to a dramatic growth of the solution space and makes evaluation runtimes impractical. Extensions of twin binary sequence to capture area-optimal floorplans without redundancies [26] do not outperform published results for sequence pair in [22]. Mathematical properties of various floorplan representations are discussed in [25].

If the representations share the same solution space, one can equalize the move set, and thus the differences will only affect runtime (some moves may be faster or slower). Therefore, we do not consider TCG and extrapolate results for sequence pair to TCG. Similarly, we study B*-trees and extrapolate results to O-trees, with a potential caveat about runtime. On the other hand, the solution space defines an intrinsic bound on the expressiveness of a representation, and this bound may directly impact solution quality. The example in Figure 3c shows that B*-tree may not necessarily capture min-wirelength solutions that can be captured by sequence pair. Therefore, we compare solution spaces rather than specific details of individual floorplan representations.

Incidentally, the two representations we have chosen for comparisons, sequence pair and B*-tree, appear best-studied in the literature among non-slicing representations. Sequence pair has been extended to handle fixed blocks [17], arbitrary convex and concave rectilinear blocks [6] and soft blocks [1]. Similar extensions have been proposed for B*-tree [24] or can be easily extrapolated. Both representations have been used in multi-level or hierarchical floorplanning [1, 12] scaling to tens of thousand blocks.

2.2 Sequence pair

Unlike graph-based representations, a sequence pair [16] is a pair of permutations (orderings) of the N blocks. The two permutations capture geometric relations between each two blocks. Recall that since blocks cannot overlap, one of them must be to the left or below from the other, or both. In sequence pair

$$\langle \dots, a, \dots, b, \dots \rangle, \langle \dots, a, \dots, b, \dots \rangle \Rightarrow a \text{ is to the left of } b \quad (1)$$

$$\langle \dots, a, \dots, b, \dots \rangle, \langle \dots, b, \dots, a, \dots \rangle \Rightarrow a \text{ is above } b \quad (2)$$

Every two blocks constrain each other in either vertical or horizontal direction, and only these constraints are recorded. Therefore, placements produced from sequence pair must be aligned to given horizontal and vertical axes, e.g., $x = 0$ and $y = 0$. There are $n!^2$ sequence pairs for n blocks, but multiple sequence pairs may encode the same block placement, e.g., for three identical square blocks, both $\langle a, c, b \rangle, \langle c, a, b \rangle$ and $\langle a, c, b \rangle, \langle c, b, a \rangle$ encode the placement with a straight on top of c , and b aligned with c on the right (Fig.3a). We say that a block placement is “representable” (or “can be captured”) by a sequence pair *iff* there exists a sequence pair which encodes that placement. The use of sequence pair for area minimization is justified by the fact [16] that at least one minimal-area placement is representable.




Figure 1: Two sequence pairs with edges of the horizontal (dashed) and vertical (solid) constraint graphs. Transitive edges are omitted.

The original $O(n^2)$ -time evaluation algorithm from [16] has been considerably simplified in [21]. Another variant in [21] runs in time $O(n \log(n))$, and later work in [22] reduces runtime to $O(n \log(\log(n)))$ without affecting the resulting block locations. While O-trees [18] and corner block lists [8] can be evaluated in linear time, the difference in complexity is dwarfed by implementation tuning, e.g., the annealing schedule. The implementation in [22] outperforms many published results in terms of runtime and solution quality.

2.3 B*-tree

B*-tree represents a compacted packing by a binary tree, in which each node corresponds to a block (see Fig.2b). The root node represents the bottom-left block, which must exist if we compactify in that direction. For example, B_5 in Fig.2b is the bottom-left block. A left child is the lowest right neighbor of its parent and a right child is the lowest block above its parent that shares the same x -coordinate with its parent. In Fig.2b, B_7 and B_1 are the left and right children of B_6 respectively.

Given a B*-tree, block locations can be found by a depth-first traversal of the tree. After block A is placed at (x_A, y_A) , we consider its left child B and set $x_B = x_A + w_A$, where w_A is the width of A . Then y_B is the smallest non-negative value that avoids overlaps with previously placed blocks. After returning from recursion at block B , we consider the right child C of A : $x_C = x_B$, and y_C is smallest possible so as to avoid overlaps. This algorithm can be implemented in $O(n)$ time with the *contour* data-structure. The contour of a packing defines its upper outline (jagged) and can be implemented as a doubly-linked list of line segments (Fig.2c). When we put a new block on top of the contour at a certain x -coordinate, it takes amortized $O(1)$ time to determine its y -coordinate [5].




Figure 2: A packing represented by (a) an O-tree and (b) an equivalent B*-tree; the contour of the packing is shown in (c). Block B_5 is the root node. Thicker arrows link parents to their left children.




Figure 3: (a)-(b) Multiple B*-trees can represent the same packing; B*-tree does not capture packing (c), which is represented by the sequence pair $\langle a, b, c \rangle, \langle c, a, b \rangle$.




Figure 4: (a) shows the vertical B*-tree of the horizontal B*-tree in Fig.2b. (b) shows the reverse horizontal B*-tree of the B*-tree in Fig.2b. All blocks are packed to the top instead of the bottom. (c) shows the reverse vertical B*-tree of the vertical B*-tree in (a). All blocks are packed to the right. Blocks are linked with their left-child by boldfaced arrows.

All packings represented by B*-trees are necessarily compacted so that no single block can move down without creating overlaps. Therefore, some packings representable by sequence pair cannot be captured by B*-tree, as shown in Fig.3c. The same packing can be captured by multiple B*-trees (Fig.3). There are $O(n!2^{2n-2}/n^{1.5})$ B*-trees [5] packings, which is fewer than $n!^2$ possible sequence pair. B*-tree captures all compacted packings, and hence some area-optimal packings. However, it may capture none of interconnect-optimal packings.

3. OUR EVALUATION FRAMEWORK

Our evaluation framework is based on the open-source floorplanner Parquet-2.1 that implements simulated annealing and uses the sequence pair representation [1]. The temperature schedule can be easily extracted from its source code. Parquet supports both traditional outline-free min-area floorplanning and fixed-outline floorplanning [9] for any combination of hard and soft blocks, with optional interconnect optimization. We replace sequence pair with B*-tree in Parquet to facilitate fair comparisons with exact same temperature schedules. With some effort, we ensure that all features of Parquet are supported when B*-tree is used. Probabilities of trying and accepting different moves are as similar as possible, although the move sets for the two representations are slightly different. All our implementations are available in Parquet 3.1 [1].

3.1 Floorplanning

Outline-free floorplanning. The following moves are used with B*-tree, (i) swapping two blocks, (ii) rotating a block and (iii) moving a block to another position in the tree (specific to B*-tree). The following moves are used with sequence pair: (a) swapping two blocks in either sequence or both, (b) rotating a block and (c) deleting a block from one of the sequences and inserting it at a random position (specific to sequence pair). During interconnect

optimization, movements of single blocks are guided towards locations that optimize adjacent interconnect. For the two floorplan representations, block rotations and interconnect-driven moves are applied with the same probabilities. To investigate the effects of compaction on both area and interconnect optimization, we apply the compaction algorithm in Section 2.3 every five moves, and report floorplanning results with and without this feature.

Floorplan compaction for B*-tree. Compaction has been studied in [7, 11] with the goal to compress a given floorplan in both directions. Note that packings captured by B*-trees and O-trees are only guaranteed to be compacted in one direction (down). We developed the following simple compaction scheme. Given a (horizontal) B*-tree, we construct a vertical B*-tree that compacts all blocks to the left. For example the horizontal B*-tree in Fig.2b has vertical B*-tree in Fig.4a. From a vertical B*-tree, we construct a horizontal B*-tree and then iterate this process until all blocks are compacted downward and to the left. Our algorithm for “re-orienting” a B*-tree is derived from the B*-tree evaluation algorithm, which places blocks one by one, in a depth-first order. When we place a block B , we identify the block A (or the bottom edge) that prevents it from moving towards the bottom as its parent, and try to mark it as the left-child of A . If block A already has left-child C , we try to mark the new block B as the right-child of C . If C already has a right-child, we consider the right-child of C , proceed similarly until we reach a block C' with no right-child and mark new block B as the right-child of C' . In Fig.4a, for example, when we place the block B_2 , we identify the bottom edge as its parent, whose left-child is already set to be block B_5 . Therefore, we consider its right-child B_3 , which has no right-child, and set B_2 as its right-child. On the other hand, when block B_7 is placed, we identify block B_8 as its parent. Since block B_8 has no left-child, we set B_7 to be its left-child. B*-tree evaluation takes $O(n)$ time for n blocks, but the worst-case overhead for placing block B in our re-orientation algorithm is $O(n)$ due to the tree-traversal needed to find B 's parent. We can avoid this traversal by recording the last left-child of each block, so that next time when we add a block B as a left-child of block A , we can go directly to the last left-child of A . For example in Fig.4a, when we place the block B_2 , we read from the bottom-edge that B_3 is the last left-child of the bottom-edge. Therefore, we can consider B_3 directly, and add B_2 as the left-child of B_3 without tree traversal. After we place the new block B_2 , we update B_2 as the last left-child of the bottom-edge. This last optimization enables constructing the vertical B*-tree in $O(n)$ time.

Fixed-outline floorplanning. In the fixed-outline mode [9], Parquet uses slack-based moves from [1]. The x -slack of a block is the distance by which it can move with other blocks fixed. A key result in [1] shows that the width of a floorplan cannot be shortened unless some block with zero x -slack is moved, and similarly for y -slacks. The slack computation in [1] is specific to sequence pair, and in this work we develop slack computation for B*-tree. Conceptually, we construct a *reverse* B*-tree for a horizontal B*-tree that compacts blocks to the top, instead of the bottom (Fig.4b). The difference in y -coordinates of each block in these two horizontal B*-tree is its y -slack. To evaluate the x -slacks, we first construct the vertical B*-tree T that compacts blocks to the left. Then we construct the reverse vertical B*-tree T' from T that compacts blocks to the right (Fig.4c). The difference in x -coordinates of each block in T' and the original horizontal B*-tree is its x -slack. Note that slack computation for B*-tree takes $O(n)$ time. This allows us to make slack-based moves identical for B*-tree and sequence pair.

Optimizing soft blocks. Aspect ratios of soft blocks can be optimized based on their x and y slacks. For example, we reshape blocks with zero x -slack and positive y -slack by decreasing their

widths and increasing their heights. This reduces the number of critical paths in the x direction without creating new critical paths in the y direction, and may improve floorplan area. Again, we ensure that the soft-block moves for sequence pair and B*-tree are as similar as possible, and applied with the same probability. Empirical data suggest the same trends for floorplanning with hard blocks and for floorplanning with soft blocks. Therefore, we mainly report on hard-block floorplanning and illustrate the similarity by representative results with soft blocks.

3.2 Floorplacement

Min-cut floorplacement introduced in [2] integrates fixed-outline floorplanning into traditional min-cut placement to solve a more general layout problem, which includes cell placement, floorplanning, mixed-size placement and achieving routability. At every step of min-cut placement, either partitioning or wirelength-driven, fixed-outline floorplanning is invoked, depending on whether large blocks are present in a given layout region. If floorplanning fails to satisfy the fixed outline, an earlier partitioning decision is undone, adjacent regions are merged and the larger region is re-floorplanned to find a legal placement of the blocks. Empirically, this framework improves the scalability and quality of results for traditional wirelength-driven floorplanning. A major implication of min-cut floorplacement is that large layouts with blocks of different sizes may not require large-scale block packing. Indeed, min-cut partitioning may divide the core regions into smaller sub-region before block packing is involved. For a given layout region (bin), a block-packing instance is constructed as follows. All connections between modules in the bin and other modules are propagated to *fixed terminals* at the periphery of the bin. As the bin may contain numerous standard cells, the number of movable objects is reduced by clustering standard cells into soft placeable blocks using a simple bottom-up connectivity-based algorithm. Our implementation only clusters small cells but not large modules.

Preliminary results of floorplacement experiments reported in the next section indicated to us that wirelength evaluation in block packing consumed a significant portion of runtime. Therefore we made several modifications to the floorplacer described in [2], so as to reduce the number of nets in block-packing instances it generates. First, when formulating a block-packing instance, the modified floorplacer ignores all nets whose bounding boxes contain the fixed outline of the instance. Such nets do not affect the optimality of block-packing solutions, and removing them is a clear win. This is similar to removing *inessential nets* in min-cut placement [3]. The main difference is that min-cut partitioning affects only the x -span or the y -span of each net, depending on the direction of the outline, but floorplanning affects both. Another improvement to the floorplacer from [2] deals with soft blocks that consist of small cells clustered together. We noticed that in many cases a pair of such clusters would be connected by numerous two-pin nets that can be conglomerated. Therefore, we (i) extended the Parquet floorplanner used in [2] to use net weights, and (ii) developed a near-linear-time algorithm for conglomerating “parallel” two-pin nets. This algorithm creates an adjacency vector (expandable array) for every vertex, in which it stores indices of other vertices adjacent to the current vertex through two-pin nets. Each adjacency vector is sorted, after which repeated indices always appear next to each other and can be counted in linear time. Each group of identical indices is then represented by a single two-pin net whose weight is the multiplicity of the index. Empirically, these two improvements significantly reduce the number of nets considered during block packing, resulting in an overall 10% speed-up of min-cut floorplacement on the 18 IBM benchmarks (more so on larger benchmarks) without a noticeable impact on solution quality.

Table 2: Comparison of sequence pair and B*-tree in outline-free mode with hard blocks. Parquet is run on a 3.2GHz Linux workstation with 1GB RAM, and all data are averaged over 50 independent runs.

		area weight = 1.0		area weight = 0.6, wire weight = 0.4		time-per-move (area_only : area+wire)	
		dead-space % / HPWL (e3) / runtime (s) / time-per-move (ms)		estimated		actual	
ami33	sequence pair	9.54%	132 / 0.30s / 4.82ms	14.3%	76 / 1.77s / 29.7ms	1 : 8.28	1 : 6.16
	B*-tree (no compaction)	5.58%	146 / 0.33s / 7.64ms	15.8%	75 / 1.52s / 31.5ms		1 : 4.12
	B*-tree (with compaction)	5.38%	145 / 0.45s / 7.40ms	14.6%	74 / 1.62s / 31.3ms		1 : 4.23
ami49	sequence pair	9.12%	1855 / 0.66s / 7.24ms	14.1%	823 / 4.30s / 48.5ms	1 : 7.19	1 : 6.70
	B*-tree (no compaction)	5.26%	2176 / 0.68s / 10.7ms	16.2%	847 / 3.64s / 50.6ms		1 : 4.73
	B*-tree (with compaction)	4.91%	2120 / 0.95s / 10.8ms	15.1%	835 / 3.95s / 51.6ms		1 : 4.78
n100	sequence pair	9.18%	397 / 3.11s / 17.6ms	11.6%	322 / 16.6s / 93.9ms	1 : 5.48	1 : 5.33
	B*-tree (no compaction)	5.48%	468 / 2.66s / 20.3ms	11.3%	320 / 14.7s / 99.5ms		1 : 4.90
	B*-tree (with compaction)	5.18%	477 / 3.79s / 20.3ms	8.68%	322 / 15.6s / 98.6s		1 : 4.86
n200	sequence pair	10.2%	732 / 12.9s / 37.9ms	13.4%	589 / 76.6s / 219ms	1 : 3.95	1 : 5.78
	B*-tree (no compaction)	5.92%	931 / 10.7s / 39.4ms	12.0%	580 / 65.2s / 219ms		1 : 5.56
	B*-tree (with compaction)	5.68%	926 / 15.2s / 39.0ms	9.27%	592 / 71.2s / 223ms		1 : 5.72
n300	sequence pair	10.7%	954 / 32.1s / 64.3ms	14.1%	707 / 159s / 305ms	1 : 3.18	1 : 4.74
	B*-tree (no compaction)	6.27%	1294 / 23.7s / 57.1ms	11.9%	700 / 141s / 313ms		1 : 5.48
	B*-tree (with compaction)	5.73%	1360 / 33.3s / 56.4ms	9.34%	730 / 149s / 308ms		1 : 5.46

4. EMPIRICAL RESULTS AND ANALYSIS

We now compare sequence pair to B*-tree within simulated annealing in several different contexts. All experiments are performed on a 3.2GHz Pentium4 workstation with 1GB RAM running Linux.

4.1 Floorplanning

The first batch of results is on the MCNC and GSRC benchmarks (see Table 3), except for those with 11 blocks and smaller.

Outline-free floorplanning. In Table 2 we report average results over 50 runs, while best-seen dead-space is typically 1-2% smaller. Evidently B*-tree packs better than sequence pair, which is consistent with B*-tree’s capturing only horizontally compacted floorplans and with published results [5]. By profiling average time per move, we empirically determine that floorplan evaluation dominates in pure area-optimization mode, i.e., when the netlist is not considered. This is consistent with the emphasis on algorithm complexity in the floorplanning literature. For example, the linear-time algorithm for evaluating a B*-tree is considered faster than the naive $O(n^2)$ algorithm for sequence pair. However, as we see in Table 2, this asymptotic comparison has little value for GSRC and MCNC benchmarks with 200 blocks or less. The more sophisticated B*-tree evaluation is faster than sequence pair evaluation *only with over 200 blocks*. To this end, we recall that $O(n \log(n))$ -time and $O(n \log(\log(n)))$ -time evaluation algorithms exist for sequence pair. In another surprising result, repeated compaction of the floorplan does not lead to significantly better results even when only area is minimized, which suggests that floorplan compaction algorithms are only useful in post-processing.

Interconnect optimization. When we configure our annealer to optimize a linear combination of area and wirelength, B*-tree leads to smaller packings, while sequence pair achieves smaller wirelength (see Table 2). This is consistent with the observation that B*-tree does not capture interconnect-optimized packings that are not horizontally compacted. Table 2 shows that runtime is now dominated by wirelength evaluation, as time-per-move quadruples. Intuitively this makes sense because floorplanning instances usually have many more nets and pins than blocks (see Table 3). In this

context floorplan evaluation is not a runtime bottleneck in practice. Below, we compare the number of floating-point operations used by wirelength computation and that by floorplan evaluation, to show that *floorplan evaluation should not be a bottleneck in general*.

First, we optimistically count floating-point operations needed to compute the HPWL of a netlist for a given placement (we ignore assignments, etc). The location of each pin p is given by the location (b_x, b_y) of its module’s center, module dimensions (w, h) and the pin’s relative offset (f_x, f_y) from the center. Namely, $p_x = b_x + wf_x$ and similarly for p_y (here wf_x can be precomputed for hard blocks). Thus, each pin requires 4 operations to compute its location, and absolute locations of pads are given. To calculate the HPWL of a degree- d net, one finds the smallest and largest x and y pin coordinates, which takes $4d$ floating-point comparisons. A special-case computation for 2-pin nets uses only two comparisons, saving 6 comparisons out of 8. Three more operations compute the HPWL of a net as $(x_{max} - x_{min}) + (y_{max} - y_{min})$. Now assume that the netlist has p pins, q pads and N nets of which N_2 are two-pin nets and observe that $p + q = \sum_i d_i$. Hence the number of floating-point operations to compute the HPWL is $4p + 4(p + q) + 3N - 6N_2$.

Next, we consider the original $O(n^2)$ -time algorithm for evaluating sequence pairs [16]. It uses the horizontal and vertical constraint graphs with E_H and E_V edges respectively. Block locations are calculated by two depth-first traversals, which requires $(E_H + E_V)$ floating-point additions and as many comparisons. Each pair of blocks is either horizontally or vertically constrained, and each block is connected to the sources and sinks of the two constraint graphs (Fig.1). Therefore $E_H + E_V = \binom{n}{2} + 4n$, and the total is $2(E_H + E_V) = n(n + 7)$ floating-point operations. Table 3 estimates such operation counts for evaluating sequence pair and HPWL, per benchmark. It also counts actual floating-point operations in the faster algorithm that we use [21] (which does not build constraint graphs). Both algorithms have worst-case complexity $O(n^2)$, and the one from [16] always takes the same amount of time. However, plotting the data in Table 3 suggests a difference in asymptotic average-case behavior. Using more than 500 randomly-generated floorplans ranging from 16 to 16K blocks, we fit the performance of the faster algorithm to $cn^{1.3}$ with good accuracy. Further, the overhead of wirelength evaluation is shown in Table 2. Surprisingly, the actual overhead increases for larger benchmarks, contrary to what floating-point counts suggest. To this end, we observe that wirelength evaluation operates on much larger datasets than floorplan evaluation (see Table 3), and larger netlists may not fit in processor cache. Overall, floorplan evaluation is never a runtime bottleneck in our experiments.

Table 3: Attributes of MCNC and GSRC benchmarks

	# blks	# nets	# pads	# pins	# float ops: estimated actual		
					WL eval.	seq. pair eval.	
ami33	33	123	42	522	4089	1320	562
ami49	49	408	22	953	6914	2744	1117
n100	100	885	334	1873	14253	10700	3180
n200	200	1585	564	3599	28669	41400	9710
n300	300	1893	569	4358	34593	92100	15863

Table 4: Comparison of sequence pair and B*-tree in fixed-outline mode without HPWL optimization. Parquet is run on a 3.2GHz Linux workstation with 1GB RAM, and all data are averaged over 50 independent runs. All blocks are hard.

	time per move	10% dead-space		20% dead-space		
		aspect ratio = 1.0	aspect ratio = 2.0	aspect ratio = 1.0	aspect ratio = 2.0	
		success-rate % / average runtime				
<i>ami33</i>	sequence pair	10.4ms	18% / 0.50s	24% / 0.50s	100% / 0.14s	100% / 0.15s
	B*-tree (no compaction)	13.9ms	72% / 0.39s	68% / 0.45s	98% / 0.09s	100% / 0.12s
	B*-tree (with compaction)	13.9ms	72% / 0.40s	84% / 0.42s	96% / 0.10s	100% / 0.13s
<i>ami49</i>	sequence pair	20.4ms	58% / 0.92s	68% / 0.90s	96% / 0.34s	98% / 0.31s
	B*-tree (no compaction)	24.7ms	66% / 0.87s	62% / 1.01s	96% / 0.26s	98% / 0.32s
	B*-tree (with compaction)	24.5ms	84% / 0.74s	92% / 0.75s	96% / 0.25s	100% / 0.30s
<i>n100</i>	sequence pair	47.2ms	76% / 4.31s	60% / 4.81s	100% / 1.71s	100% / 1.69s
	B*-tree (no compaction)	53.3ms	74% / 3.97s	98% / 3.44s	100% / 1.28s	100% / 1.29s
	B*-tree (with compaction)	53.5ms	100% / 2.74s	100% / 2.63s	98% / 1.27s	100% / 1.30s
<i>n200</i>	sequence pair	175ms	82% / 20.5s	38% / 23.8s	100% / 8.43s	100% / 8.64s
	B*-tree (no compaction)	204ms	20% / 20.5%	62% / 19.2s	100% / 6.31s	100% / 6.78s
	B*-tree (with compaction)	205ms	94% / 13.9s	100% / 13.8s	100% / 5.90s	100% / 6.28s
<i>n300</i>	sequence pair	344ms	80% / 51.1s	30% / 59.9s	100% / 22.4s	100% / 23.0s
	B*-tree (no compaction)	442ms	2% / 47.6s	14% / 45.7s	100% / 17.3s	100% / 17.3s
	B*-tree (with compaction)	400ms	98% / 35.4s	100% / 33.6s	98% / 13.8s	100% / 14.5s

Fixed-outline floorplanning. Tables 4 and 5 show that moves are more expensive in fixed-outline floorplanning. In particular, the average time-per-move for B*-tree grows faster than what its linear-time evaluation algorithm suggests. We trace these effects to slack-based moves that involve computing the x - and y -slacks for all blocks, and then prioritizing blocks by slacks. For both representations, this involves two floorplan evaluations, two rounds of floating-point subtractions, etc. However, slack-based moves are important to ensure that the final floorplan fits into a given outline [1]. B*-tree typically achieves higher success rates than sequence pair, which is expected since B*-tree captures only horizontally-compacted packings. Since Parquet terminates soon after finding the first legal solution, higher success rates improve runtime.

When optimizing interconnect subject to a fixed-outline, one is trading off success rate and interconnect cost. More aggressive annealers explore solutions with smaller wirelength at the risk of not finding any legal floorplans that fit into the outline. Therefore, success rates in Table 5 are substantially lower than those in Table 4. Due to compaction, B*-tree packings are more likely to fit into the outline than those represented by sequence pair. Therefore one can view B*-tree as more conservative and may expect to achieve the same trade-off by increasing or decreasing the weight of the interconnect term in the objective function of simulated annealing. However, *some low-interconnect floorplans may elude B*-tree with any and all weight configurations*. As in the outline-free mode, wirelength evaluation dominates runtime, and periodic floorplan compaction does not improve results.

Optimizing soft blocks. The trends we reported so far carry over to soft blocks — Table 6 compares sequence pair and B*-tree on soft-block version of MCNC and GSRC benchmarks, with compaction moves disabled. As before, (i) B*-tree achieves higher success rates and takes less time on average, and (ii) wirelength evaluation considerably increases average time per move.

4.2 Floorplacement

To scale our previous comparisons up, we embed them into the min-cut floorplacement framework, where (i) the netlist and the layout area are first partitioned into smaller regions, and (ii) block packing is used only when macros occupy a significant portion of region. We work with the floorplacer Capo 9.0 [2], which uses the block packer Parquet. With minor modifications we can now use either sequence pair or B*-tree within Capo. Our experiments are performed with a suite of 18 publicly-available mixed-size placement benchmarks IBM-MSwPins [2] that have non-square blocks

Table 6: Comparison of sequence pair and B*-tree in fixed-outline mode. All data are averaged over 50 independent runs, and all blocks are soft. The outline has 20% whitespace and aspect ratio 1.

	Without HPWL optimization	
	sequence pair	B*-tree
	success rate % / runtime (s)	time-per-move (ms)
<i>ami33</i>	100% / 0.19s / 10.7ms	100% / 0.11s / 14.4ms
<i>ami49</i>	100% / 0.38s / 20.9ms	98% / 0.28s / 25.5ms
<i>n100</i>	100% / 2.11s / 47.4ms	100% / 6.89s / 131ms
<i>n200</i>	100% / 12.6s / 191ms	100% / 8.99s / 217ms
<i>n300</i>	100% / 65.4s / 343ms	100% / 49.7s / 464ms

	With HPWL optimization			
	sequence pair	B*-tree		
	success rate % / HPWL (e3)	runtime (s)	time-per-move (ms)	
<i>ami33</i>	98% / 84	1.73s / 36.6ms	82% / 74	1.35s / 40.2ms
<i>ami49</i>	100% / 880	4.44s / 64.3ms	80% / 974	3.06s / 69.9ms
<i>n100</i>	100% / 338	21.0s / 130ms	94% / 340	14.9s / 134ms
<i>n200</i>	100% / 638	124s / 363ms	86% / 647	104s / 417ms
<i>n300</i>	100% / 740	356s / 563ms	70% / 810	378s / 747ms

and non-trivial pin offsets. Final wirelengths and overall runtimes of mixed-size placement are shown in Table 7. Sequence pair outperforms B*-tree by $<1.0\%$ in wirelength and, marginally, in runtime.¹ The difference is insufficient to declare a winner as further tuning may improve results by 1-2%. However, the choice of floorplan representation appears to make little impact overall. Of course, this is partly due to the fact that final results of mixed-size placement also depend on other subsystems of Capo, such as several min-cut partitioners and end-case placers. We also note that regularly compacting B*-tree floorplans during annealing does not impact final results. The fraction of floorplacement runtime used by block packing is also reported, where we again see little difference.

Our next experiment aims to distinguish block-packing effects from floorplacement, and yet avoid overspecializing to MCNC and GSRC benchmarks. We collect a rich set of over 2000 block-packing instances generated during floorplacement and solve them with *stand-alone* sequence pair and B*-tree floorplanners. Recall that during floorplacement standard cells may be clustered into soft blocks, as described in Section 3.2. Table 8 suggests that our saved instances are fairly diverse, ranging from 1 or 2 blocks to more than 100 blocks, which is interesting for benchmarking purposes. While average block counts are small, fairly large block-packing

¹In min-cut placement, it is common to see that better optimization of interconnect at higher levels speeds up lower levels.

Table 5: Comparison of sequence pair and B*-tree in fixed-outline mode with HPWL optimization. Parquet is run on a 3.2GHz Linux workstation with 1GB RAM, and all data are averaged over 50 independent runs. All blocks are hard.

	time per move	10% dead-space		20% dead-space		
		aspect ratio = 1.0	aspect ratio = 2.0	aspect ratio = 1.0	aspect ratio = 2.0	
		success-rate % / HPWL / average runtime				
<i>ami33</i>	sequence pair	36.4ms	0% / - / 2.11s	2% / 80 / 2.08s	80% / 80 / 1.67s	86% / 83 / 1.60s
	B*-tree (no compaction)	39.9ms	2% / 79 / 1.88s	0% / - / 1.91s	58% / 76 / 1.32s	58% / 79 / 1.28s
	B*-tree (with compaction)	40.2ms	6% / 82 / 1.93s	2% / 82 / 1.99s	92% / 81 / 1.05s	92% / 84 / 1.11s
<i>ami49</i>	sequence pair	63.6ms	6% / 868 / 5.11s	4% / 855 / 5.09s	90% / 934 / 3.90s	78% / 942 / 4.10s
	B*-tree (no compaction)	67.4ms	0% / - / 4.58s	0% / - / 4.54s	52% / 926 / 3.34s	38% / 948 / 3.63s
	B*-tree (with compaction)	67.7ms	0% / - / 4.71s	0% / - / 4.76s	78% / 1022 / 2.45s	88% / 1057 / 2.85s
<i>n100</i>	sequence pair	132ms	38% / 340 / 20.5s	28% / 348 / 19.7s	100% / 338 / 15.3s	100% / 345 / 15.4s
	B*-tree (no compaction)	139ms	10% / 340 / 17.5s	2% / 334 / 18.0s	100% / 343 / 8.69s	100% / 342 / 8.14s
	B*-tree (with compaction)	139ms	100% / 398 / 2.74s	100% / 411 / 2.63s	98% / 404 / 1.27s	100% / 419 / 1.30s
<i>n200</i>	sequence pair	351ms	24% / 621 / 88.7s	0% / - / 97.5	100% / 614 / 68.0s	100% / 630 / 68.0s
	B*-tree (no compaction)	410ms	0% / - / 84.3s	0% / - / 82.4s	98% / 633 / 40.8s	100% / 642 / 41.1s
	B*-tree (with compaction)	410ms	94% / 729 / 13.9s	100% / 756 / 13.8s	100% / 745 / 5.90s	100% / 770 / 6.28s
<i>n300</i>	sequence pair	565ms	20% / 749 / 199s	2% / 805 / 209s	100% / 745 / 149s	100% / 771 / 150s
	B*-tree (no compaction)	719ms	0% / - / 177s	0% / - / 82.4s	98% / 633 / 40.8s	100% / 642 / 41.1s
	B*-tree (with compaction)	721ms	98% / 951 / 35.4s	100% / 996 / 33.6s	100% / 971 / 13.8s	100% / 1018 / 14.5s

Table 8: Comparison of sequence pair and B*-tree on clustered floorplanning instances generated by floorplacement on IBM-MSwPins benchmarks. ibm05 does not have any macros.

	# inst-ances	max / min # blocks	avg # blocks	% of successful instances	
				sequence pair	B*-tree
ibm01	74	47 / 1	16.8	68 / 91.9%	68 / 91.9%
ibm02	50	93 / 1	13.6	45 / 90.6%	45 / 90.0%
ibm03	27	69 / 1	18.6	18 / 66.7%	19 / 70.4%
ibm04	92	62 / 1	20.3	71 / 77.2%	74 / 80.4%
ibm06	68	49 / 1	20.3	53 / 77.9%	54 / 79.4%
ibm07	50	59 / 1	18.7	43 / 86.0%	41 / 82.0%
ibm08	79	78 / 1	17.9	62 / 78.5%	57 / 72.2%
ibm09	55	144 / 1	23.3	43 / 78.2%	43 / 78.2%
ibm10	303	65 / 1	16.2	242 / 79.9%	243 / 80.2%
ibm11	81	57 / 1	18.4	61 / 75.3%	59 / 71.8%
ibm12	206	213 / 1	17.2	188 / 91.3%	190 / 92.2%
ibm13	130	71 / 1	17.2	100 / 83.3%	98 / 81.7%
ibm14	286	93 / 1	18.0	219 / 76.6%	219 / 76.6%
ibm15	198	52 / 1	17.7	157 / 79.3%	154 / 77.8%
ibm16	154	172 / 1	23.1	123 / 79.9%	123 / 79.9%
ibm17	236	48 / 1	19.2	215 / 91.1%	214 / 90.7%
ibm18	110	67 / 1	21.1	85 / 77.3%	81 / 73.6%

instances appear once in a while — typically when a large block triggers floorplanning in a region with many small macros (Capo currently does not cluster macros, but this can be implemented if needed). Again, all results for sequence pair and B*-tree are close.

In Table 9 we break down the 286 instances generated from ibm14 by size to study the performance of sequence pair and B*-tree at different levels of physical hierarchy in floorplacement. One may suspect that sequence pair and B*-tree respond differently to clustered instances of different sizes, but Table 9 shows little difference at every min-cut level. This reinforces our earlier observation that sequence pair and B*-tree perform very similarly.

5. CONCLUSIONS

Floorplan representations have been the center of floorplanning research in the last decade, and their geometric properties have been extensively investigated. However, empirical evaluations have been restricted to a small set of benchmarks, and anecdotal evidence points to several versions of MCNC benchmarks that generate incompatible area and wirelength results. To this end, our work offers a comprehensive comparison of two well-researched floorplan representations in stand-alone block packing and in the context of min-cut floorplacement. This comparison is facilitated by several technical results presented in our work, such as the evaluation of floorplan slack in terms of B*-tree and the simplification

Table 9: Comparison of sequence pair and B*-tree on clustered floorplanning instances generated by floorplacement on the ibm14 benchmark.

min-cut level	# inst-ances	max / min # blocks	avg # blocks	# (%) of successful instances	
				sequence pair	B*-tree
5	2	95 / 49	72.0	2 / 100%	1 / 50.0%
6	0	- / -	-	0 / -	0 / -
7	4	56 / 5	39.8	4 / 100%	3 / 75.0%
8	6	38 / 14	25.7	4 / 66.7%	4 / 66.7%
9	55	49 / 2	23.4	47 / 85.5%	49 / 89.1%
10	87	49 / 2	17.7	63 / 78.8%	60 / 69.0%
11	107	49 / 1	15.6	79 / 73.8%	82 / 76.6%
12	22	28 / 1	8.6	17 / 77.3%	17 / 77.3%
13	3	7 / 1	3.3	3 / 100%	3 / 100%
total	286	95 / 1	18.0	219 / 76.6%	219 / 76.6%

of block-packing instances generated by min-cut floorplacement.

Our experiments strongly suggest that many theoretical results in block packing, traditionally formulated in the context of simulated annealing, have little relevance in digital design. For example, we show that a well-known algorithm for evaluating sequence pairs [21] runs in $\Theta(n^{1.3})$ time in practice, which is better than its estimated worst-case complexity $O(n^2)$. Given that practical block-packing instances rarely exceed 300 blocks, this algorithm empirically outperforms a linear-time algorithm for B*-tree, but the difference is dwarfed by wirelength evaluation. Other popular properties of representations, such as redundancies in the solution space, make little impact in the presence of wirelength optimization. Since congestion, delay and power usually take longer to compute than wirelength, interconnect evaluation in general is the bottleneck in annealing-based floorplanning. To this end, move-based incremental algorithms for interconnect evaluation can be very useful, but are difficult to design because even minor changes to relative block locations can significantly affect the structure of packed floorplans. Without incremental evaluation that is faster both asymptotically and empirically, new floorplan representations and fast evaluation algorithms seem irrelevant. Our conclusions are supported by extensive experiments on standard MCNC and GSRC benchmarks, as well as a rich set of block-packing instances derived from min-cut floorplacement of large mixed-size netlists.

Algorithms that are not based on annealing may be competitive in area packing [4], but are often impractical because their objective function cannot be modified in applications. More generally, while we have not considered all possible algorithmic and design contexts for justifying new floorplan representations, changing the current *status quo* on their utility may require new breakthroughs.

Table 7: Floorplacement using sequence pair and B*-tree respectively on IBM-MSwPins mixed-size benchmarks. Capo is run on a 3.2GHz Linux workstation with 1GB RAM. Best HPWL for each benchmark is boldfaced. ibm05 does not have any macros.

	sequence pair			B*-tree					
				with compaction			no compaction		
	HPWL (e6) / time (s) / % time used in floorplanning / no. successful			FP instances / no. FP instances					
ibm01	2.48 / 221	23.3%	71/80	2.51 / 230	29.0%	68/82	2.47 / 217	25.4%	78/86
ibm02	5.21 / 447	10.2%	42/45	5.64 / 468	21.1%	53/59	5.56 / 519	23.9%	49/54
ibm03	8.27 / 647	18.8%	40/45	8.15 / 580	13.2%	63/65	8.05 / 533	7.50%	24/29
ibm04	8.81 / 727	22.6%	93/99	8.84 / 738	20.7%	90/100	8.91 / 676	16.7%	84/96
ibm06	7.34 / 751	12.5%	50/56	7.36 / 758	12.6%	66/75	7.21 / 770	14.6%	53/57
ibm07	12.24 / 1091	10.5%	53/57	11.55 / 1039	6.48%	86/92	11.44 / 1072	8.97%	78/85
ibm08	14.57 / 1343	23.3%	49/59	14.83 / 1227	16.5%	67/78	14.88 / 1753	40.8%	63/75
ibm09	15.06 / 1302	10.1%	59/66	15.60 / 1348	11.8%	49/57	15.51 / 1376	13.5%	40/44
ibm10	34.30 / 2346	20.9%	265/298	34.71 / 1227	17.4%	254/293	33.27 / 2224	20.0%	263/294
ibm11	21.95 / 1814	7.74%	69/76	21.82 / 1348	10.5%	60/69	22.14 / 2009	16.4%	56/67
ibm12	41.23 / 2614	15.3%	165/180	43.44 / 2254	18.0%	195/218	40.69 / 2427	16.6%	220/246
ibm13	27.29 / 2411	12.2%	90/101	27.23 / 1867	15.1%	79/87	28.55 / 2565	16.5%	94/103
ibm14	40.79 / 4048	10.9%	250/275	41.02 / 4123	11.2%	216/245	41.21 / 4101	11.3%	257/285
ibm15	57.30 / 5532	10.3%	136/160	53.96 / 5218	7.95%	197/230	64.14 / 5799	8.98%	201/234
ibm16	65.41 / 5851	7.60%	138/148	64.49 / 5877	8.02%	117/134	64.42 / 5802	8.28%	121/142
ibm17	77.31 / 6010	3.26%	213/225	77.69 / 6096	5.42%	230/254	77.39 / 6134	4.14%	241/265
ibm18	48.81 / 5320	2.88%	61/75	48.85 / 5204	2.98%	72/77	49.45 / 5309	2.38%	94/100
Average	0% / 0%			+0.53% / -0.44%			+0.68% / +2.39%		

6. REFERENCES

- [1] S.N. Adya and I.L. Markov, "Fixed-outline Floorplanning: Enabling Hierarchical Design," *IEEE Trans. on VLSI* 11(6), pp.1120-35, 2003. <http://vlsicad.eecs.umich.edu/BK/parquet/>
- [2] S.N. Adya, S. Chaturvedi, J.A. Roy, D.A. Papa and I. L. Markov, "Unification of Partitioning, Floorplanning and Placement," *ICCAD 2004*, pp. 550-557.
- [3] A. E. Caldwell, A. B. Kahng, I. L. Markov, "Optimal Partitioners and End-case Placers for Standard-cell Layout," *IEEE Trans. on CAD* 19(11), pp. 1304-1314, 2000.
- [4] H. H. Chan, I. L. Markov, "Practical Slicing and Non-slicing Block-Packing without Simulated Annealing," *ACM/IEEE Great Lakes Symp. on VLSI 2004*, pp. 282-287.
- [5] Y.-C. Chang, Y.-W. Chang, G.-M. Wu and S.-W. Wu, "B*-trees: A New Representation for Non-Slicing Floorplans," *DAC 2000*, pp. 458-463.
- [6] K. Fujiyoshi and H. Murata, "Arbitrary Convex and Concave Rectilinear Block Packing Using Sequence Pair," *ISPD 1999*, pp. 103-110.
- [7] P.-N. Guo, C.-K. Cheng and T. Yoshimura, "An O-tree Representation of Non-Slicing Floorplan," *DAC '99*, pp. 268-273.
- [8] X. Hong et al., "Corner Block List: An Effective and Efficient Topological Representation of Non-Slicing Floorplan," *ICCAD 2000*, pp. 8-13.
- [9] A. B. Kahng, "Classical floorplanning harmful?" *ISPD 2000*, pp. 207-213.
- [10] R. E. Korf, "Optimal Rectangle Packing: New Results," *ICAPS 2004*, pp. 142-149.
- [11] M. Lai and D. Wong, "Slicing Tree Is a Complete Floorplan Representation," *DATE 2001*, pp. 228-232.
- [12] H.-C. Lee, Y.-W. Chang, J.-M. Hsu, and H. H. Yang, "Multilevel Floorplanning/Placement for Large-Scale Modules using B*-trees," *DAC 2003*, pp. 812 - 817.
- [13] J.-M. Lin and Y.-W. Chang, "TCG: A Transitive Closure Graph-Based Representation for Non-Slicing Floorplans," *DAC 2001*, pp. 764-769.
- [14] J.-M. Lin and Y.-W. Chang, "TCG-S: Orthogonal Coupling of P*-admissible Representations for General Floorplans," *DAC 2002*, pp. 842-847.
- [15] P. H. Madden, "Reporting of Standard Cell Placement Results," *IEEE Trans. on CAD* 21(2), Feb. 2002, pp. 240-247.
- [16] H. Murata, K. Fujiyoshi, S. Nakatake and Y. Kajitani, "VLSI Module Placement Based on Rectangle-Packing by the Sequence Pair," *IEEE Trans. on CAD* 15(12), pp. 1518-1524, 1996.
- [17] H. Murata and E. S. Kuh, "Sequence-Pair Based Placement Methods for Hard/Soft/Pre-placed Modules," *ISPD 1998*, pp. 167-172.
- [18] Y. Pang, C.-K. Cheng and T. Yoshimura, "An Enhanced Perturbing Algorithm for Floorplan Design Using the O-tree Representation," *ISPD 2000*, pp. 168-173.
- [19] S. Prestwich, "Supersymmetric Modelling for Local Search," *SymCon '02*, September 2002. <http://user.it.uu.se/~pierref/astra/SymCon02/>
- [20] Z. C. Shen and C.C.N. Chu, "Bounds on the Number of Slicing, Mosaic, and General Floorplans," *IEEE Trans. on CAD* 22(10), pp. 1354 - 1361.
- [21] X. Tang, R. Tian and D. F. Wong, "Fast Evaluation of Sequence Pair in Block Placement by Longest Common Subsequence Computation," *DATE 2000*, pp. 106-111.
- [22] X. Tang and D. F. Wong, "FAST-SP: A Fast Algorithm for Block Placement Based on Sequence Pair," *ASPDAC 2001*, pp. 521-526.
- [23] D. F. Wong and C. L. Liu, "A New Algorithm For Floorplan Design," *DAC 1986*, pp. 101-107.
- [24] G.-M. Wu and Y.-C. Chang and Y.-W. Chang, "Rectilinear block placement using B*-trees," *ACM Trans. on Design Autom. of Electronic Systems* 8(2), pp. 188-202, 2003.
- [25] B. Yao et al., "Floorplan Representations: Complexity and Connections," *ACM Trans. on Design Autom. of Electronic Systems* 8(1), pp. 55-80, 2003.
- [26] E.F.Y. Young, C.C.N. Chu and Z.C. Shen, "Twin Binary Sequences: A Nonredundant Representation for General Nonslicing Floorplan," *IEEE Trans. on CAD* 22(4), pp. 457-469, 2003.
- [27] S. Zhou, S. Dong, C.-K. Cheng and J. Gu, "ECBL: An Extended Corner Block List with Solution Space including Optimum Placement," *ISPD 2001*, pp. 150-155.
- [28] H. Zhou and J. Wang, "ACG-Adjacent Constraint Graph for General Floorplans," *ICCD 2004*.
- [29] C. Zhuang, Y. Kajitani, K. Sakanushi and L. Jin, "An Enhanced Q-Sequence Augmented with Empty-Room-Insertion and Parenthesis Trees," *DATE 2002*, pp. 61-68.




Figure 5: Sample placements produced by Capo 9.0 on ibm01 using (a) sequence pair and (b) B*-tree. The discrepancies in wirelength versus Table 7 represent variability in Capo results.