SimPL: An Effective Placement Algorithm

Myung-Chul Kim, Dong-Jin Lee and Igor L. Markov
University of Michigan, Department of EECS, 2260 Hayward St., Ann Arbor, MI 48109-2121
{mckima, ejdjsy, imarkov}@eecs.umich.edu

In memoriam Frank Johannes

***Abstract*—We propose a self-contained, flat, quadratic global placer that is simpler than existing placers and easier to integrate into timing-closure flows. It maintains lower-bound and upper-bound placements that converge to a final solution. The upper-bound placement is produced by a novel *look-ahead legalization* algorithm. Our placer SimPL outperforms mPL6, FastPlace3, NTUPlace3 , APlace2 and Capo simultaneously in runtime and solution quality, running 7.10 times faster than mPL6 (when using a single thread) and reducing wirelength by 3% on the ISPD 2005 benchmark suite. More significant improvements are achieved on larger benchmarks. The new algorithm is amenable to parallelism, and we report empirical studies with SSE2 instructions and up to eight parallel threads.**

I. Introduction

Global placement currently remains at the core of physical design and is a gating factor for downstream optimizations during timing closure [3]. Despite impressive improvements reported by researchers [21] and industry software in the last five years, state-of-the-art algorithms and tools for placement suffer several key shortcomings which are becoming more pronounced at recent technology nodes. These shortcomings fall into four categories: (i) speed, (ii) solution quality, (iii) simplicity and integration with other optimizations, (iv) support for multi-threaded execution. We propose the SimPL algorithm that simultaneously improves results in the first three categories and lends itself naturally to thread-level and instruction-level parallelism on multicore CPUs.

State-of-the-art algorithms for global placement form two families: (i) *force-directed* quadratic placers, such as Kraftwerk2 [27], FastPlace3 [29] and RQL [30], and (ii) *non-convex optimization* techniques, such as APlace2 [16], NTUPlace3 [8] and mPL6 [7]. Force-directed quadratic algorithms model total net length by a quadratic function of cell locations and minimize it by solving a large sparse system of linear equations. To discourage cell overlap, *forces* are added pulling cells away from high-density areas. These forces are modeled by *pseudopins* and *pseudonets*, which extend the original quadratic function [14]. They are updated after each linear-system solve until iterations converge. Non-convex optimization models net length by more sophisticated differentiable functions with linear asymptotic behavior which are then minimized by advanced numerical analysis techniques [16]. Cell density is modeled by functional terms, which are more accurate than forces, but also require updates after each change to placement [8], [16]. Algorithms in both categories are used in the industry or closely resemble those in industry placers.

Tools based on non-convex optimization achieve the best results reported for academic implementations [8] and EDA vendor tools, but are significantly slower, which is problematic for modern flat SoC placement instances with tens of millions of movable objects. To scale the basic non-convex optimization framework, all tools in this family employ *netlist clustering* and *multilevel extensions*, sometimes at the cost of solution quality. Such multilevel placers perform many sequential steps, obstructing efficient parallelization. Moreover, clustering and refinement do not fully benefit from modern multicore CPUs. Due to their complexity, multilevel placers are also harder to maintain, improve, and combine with other physical-design techniques. In particular, clustered netlists complicate accurate static timing analysis, congestion maps and physical synthesis transformation, such as performance-driven buffering, gate sizing, fanin/fanout optimization, cloning, etc [3]. Hence, timing-closure flows often repeat global placement 3-4 times, alternating it with timing analysis, physical synthesis and congestion improvement.

State-of-the-art force-directed quadratic placers tend to run many times faster than non-convex optimization, but also use multilevel extensions in their most competitive configurations. Their solution quality is mixed. FastPlace3 underperforms mPL6, but the industry tool RQL closely related to FastPlace outperforms these two non-convex placers. Kraftwerk2 is the only competitive *flat* placer (i.e., it does not use clustering) and rivals other force-directed quadratic placers in speed. However, it lags behind in solution quality and poses several challenges, such as quickly solving Poisson's equation, ensuring the convergence of iterations and avoiding halos of unused space around macros. Our experience indicates that the performance

of Kraftwerk2 can be uneven, and stability can only be achieved with some loss of solution quality [18]. Several placers are described in the book [21] and journal papers [4], [8], [27].

Effective parallelization of CAD optimizations often requires redesign and simplication of entire algorithms to use fewer components, especially standard solvers, to avoid well-known limits to parallelism described by Amdahl's law. On the other hand, recent literature on parallel algorithms and GPGPU programming[1] often focuses on algorithms that are easier to parallelize, but are not the fastest or best-performing available [12], [15], [19]. Such results may be useful to illustrate specific parallelization techniques, but do not justify the need for parallelization. We believe that new EDA tool development should not solely focus on parallel processing, but rather on novel high-performance algorithms amenable to parallel processing.

In this work, we develop a new, self-contained technique for global placement that ranks as a flat partition-based and force-directed placement algorithm. It maintains lower-bound and upper-bound placements that converge to a final solution. The upper-bound placement is produced by a novel *look-ahead legalization* algorithm based on top-down geometric partitioning and non-linear scaling. Our implementation outperforms published placers simultaneously in solution quality and speed on standard benchmarks. The lower-bound placement is produced by solving a linear system with spreading forces. Our algorithm is simpler, and our attempts to improve overall results using additional modules and extensions from existing placers (such as *netlist clustering* [7], [16], [29], *iterative local refinement* (ILR) [29], and median-improvement (*BoxPlace*) [18]) were unsuccessful.

In the remainder of this paper, Section II describes the building blocks from which our algorithm was assembled. Section III introduces our key ideas and articulates our solution of the *force modulation* problem. The SimPL algorithm is presented in Section IV along with complexity analysis. Extensions and improvements are discussed in Section V, and empirical validation is described in Section VI. The use of parallelism is discussed in Section VII, and Section VIII summarizes our results.

II. Essential Concepts and Building Blocks

Circuit placement typically operates on a gate-level netlist, which consists of standard cells (NAND, NOR, MUX, half-adders, etc) and interconnect. Each standard cell has rectangular footprint with well-defined area. Some standard cells drive multiple other cells — such interconnects are captured by signal nets. Given a netlist $\mathcal{N} = (E, V)$ with nets E and nodes (cells) V, *global placement* seeks node locations (x_i, y_i) such that the area of nodes within any rectangular region does not exceed the area of (cell sites in) that region.[2] Some locations of cells may be given initially and fixed. The interconnect objective optimized by global placement is the Half-Perimeter WireLength (HPWL). For node locations $\vec{x} = \{x_i\}$ and $\vec{y} = \{y_i\}$, $\text{HPWL}_{\mathcal{N}}(\vec{x}, \vec{y})$= $\text{HPWL}_{\mathcal{N}}(\vec{x})$+$\text{HPWL}_{\mathcal{N}}(\vec{y})$, where

$$HPWL_{\mathcal{N}}(\vec{x}) = \Sigma_{e \in E}[\max_{i \in e} x_i - \min_{i \in e} x_i] \quad (1)$$

Efficient optimization algorithms often approximate $\text{HPWL}_{\mathcal{N}}$ by differentiable functions, as illustrated next.

Quadratic optimization. Consider a graph $\mathcal{G} = (E_{\mathcal{G}}, V)$ with edges $E_{\mathcal{G}}$, vertices V and edge weights $w_{ij} > 0$ for all edges $e_{ij} \in E_{\mathcal{G}}$. The *quadratic objective* $\Phi_{\mathcal{G}}$ is defined as

$$\Phi_{\mathcal{G}}(\vec{x}, \vec{y}) = \Sigma_{i,j} w_{i,j}[(x_i - x_j)^2 + (y_i - y_j)^2] \quad (2)$$

Its x & y components are cast in matrix form [4], [27]

$$\Phi_{\mathcal{G}}(\vec{x}) = \frac{1}{2}\vec{x}^T Q_x \vec{x} + \vec{c}_x^T \vec{x} + \text{const} \quad (3)$$

The Hessian matrix Q_x captures connections between pairs of movable vertices, while vector $\vec{c}_x$ captures connections between movable and fixed vertices [17, Section 4.3.2]. When Q_x is non-degenerate, $\Phi_{\mathcal{G}}(\vec{x})$ is a strictly convex function with a unique minimum, which can be found by solving the system of linear equations $Q_x\vec{x} = -\vec{c}_x$. Solutions can be quickly approximated by iterative Krylov-subspace techniques, such as the Conjugate Gradient (CG) method and its variants [26]. Since Q_x is symmetric positive definite, CG iterations provably minimize the residual norm. The convergence is monotonic [28], but its rate depends on the spectral properties of Q_x, which can be enhanced by *preconditioning*. In other words, we solve the equivalent system $P^{-1}Q_x = -P^{-1}\vec{c}_x$ for a nondegenerate matrix P, such that P^{-1} is an easy-to-compute approximation of Q_x^{-1}. Given that Q_x is diagonally dominant, we chose P to be its diagonal, also known as the *Jacobi preconditioner*. Our placement algorithm (Section IV-C) deliberately enhances diagonal dominance in Q_x.

The Bound2Bound net model [27]. To represent the HPWL objective by the quadratic objective, the netlist $\mathcal{N}$ is transformed in *two* graphs, $\mathcal{G}_x$ and $\mathcal{G}_y$, that preserve the node set V and represent each two-pin net by a single edge with weight $1/length$. Larger nets are decomposed depending on the relative placement of vertices — for each p-pin net, the *extreme* nodes (min and max) are connected to each other and to each

[1] GPGPU programming = General-Purpose programming on GPUs (Graphics Processing Units) [12], [15].

[2] In practice, this constraint is enforced for bins of a regular grid. The layout area is subdivided into equal, disjoint, small rectangles, and each rectangle limits total area of cells placed within.

internal node by edges, with the following weight

$$w_{x,ij}^{B2B} = \frac{1}{(p-1)|x_i - x_j|} \quad (4)$$

For example, 3-pin nets are decomposed into cliques with edge weight $1/2l$, where l is the length of a given edge. In general, this quadratic objective and the Bound2Bound (B2B) net decomposition capture the HPWL objective exactly, but only for the given placement. As locations change, the error may grow, necessitating multiple updates throughout the placement algorithm.

Most quadratic placers use the placement-independent star or clique decompositions, so as not to rebuild Q_x and Q_y many times [4], [29], [30]. Yet, the B2B model uses fewer edges than cliques ($p > 3$), avoids new variables used in stars, and is more accurate than both stars and cliques [27].

III. Key Ideas in Our Work

Analytic placement techniques first minimize a function of interconnect length, neglecting overlaps between standard cells, macros, etc. This initial step places many cells in densely populated regions, typically around the center of the layout. Cell locations are then gradually spread through a series of placement iterations, during which interconnect length slowly *increases*, converging to a final overlap-free placement (a small amount of overlap is often allowed and later resolved during detailed placement).

Our algorithm also starts with pure interconnect minimization, but its next step is unusual — most overlaps are removed using a fast *look-ahead legalizer* based on top-down geometric partitioning and non-linear scaling. Locations of movable objects in the legalized placement serve as *anchors* to coerce the initial locations into a configuration with less overlap, by adding pseudonets to baseline force-directed placement [14].

Each subsequent iteration of our algorithm produces (i) an almost-legal placement that *overestimates* the final result — through look-ahead legalization, and (ii) an illegal placement that *underestimates* the final result — through linear system solver. The wirelength gap between lower-bound and upper-bound placements helps monitor convergence (Section IV-C).

Solving the force-modulation problem. A key innovation in SimPL is the interaction between the lower-bound and the upper-bound placements — it ensures convergence to a no-overlap solution while optimizing interconnect length. It solves two well-known challenges in analytic placement: (1) finding directions in which to spread the locations (*force orientation*), and (2) determining the appropriate amount of spreading (*force modulation*) [18], [30]. This is unlike previous work, where spreading directions are typically based on *local information*, e.g., placers based on non-convex optimization use *gradient* information and require a large number of expensive iterations. Kraftwerk2 [27] orients spreading forces according to solutions of Poisson's equation, providing a global perspective and speeding up convergence. However, this approach does not solve the force-modulation problem, as articulated in [18].[3] The authors of RQL [30], which can be viewed as an improvement on FastPlace, revisit the force-modulation problem and address it by a somewhat *ad hoc* limit on the magnitude of spreading forces. In our work, the *look-ahead legalization algorithm* (Section IV-B), invoked at each iteration, determines both the direction and the magnitude of spreading forces. It is global in nature, accounts for fixed obstacles, and preserves relative placement to ensure interconnect optimization and convergence. Our placement algorithm does not require exotic components, such as a Poisson-equation solver used by Kraftwerk; our C++ implementation is self-contained.

Global placement with look-ahead. The legalized upper-bound placements constructed at every iteration of our placer can be viewed as *look-ahead* because they are used only temporarily and not refined directly. They pull cell locations in lower-bound placements not just away from dense regions, but also toward the regions where space is available. Such *area look-ahead* is particularly useful around fixed obstacles, where local information does not offer sufficient guidance. While not explored in this paper, similar *congestion look-ahead* and *timing look-ahead* based on legalized placements can be used to integrate our placement algorithm into modern timing-closure flows.

IV. Our Global Placement Algorithm

Our placement technique consists of three phases: initial placement, global placement iterations and post-global placement (Figure 1). Initial placement, described next, is mostly an exercise in judicious application of known components. Our main innovation is in the global placement phase. Post-global placement is straightforward, given current state of the art.

A. Initial Placement

Our initial-placement step is conceptually similar to those of other force-directed placers [27], [29], [30] — it entirely ignores cell areas and overlaps, so as to minimize a quadratic approximation of total interconnect length. We found that this step notably impacts the final result. Therefore, unlike FastPlace3 [29] and RQL [30], we use the more accurate Bound2Bound net

[3]The work in [18] performs force modulation with *line search* but is not currently competitive with state of the art.

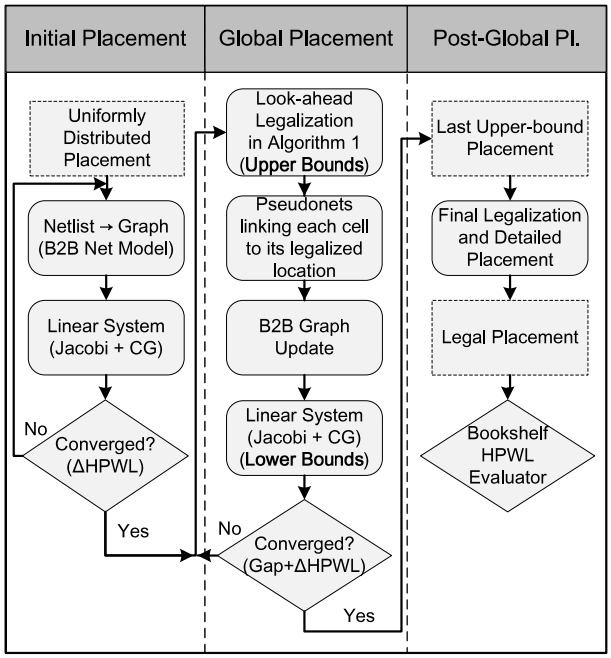

Fig. 1. The SimPL algorithm uses placement-dependent B2B net model, which is updated on every iteration. *Gap* refers to the difference between upper and lower bounds.

model from [27] reviewed in Section II. After the first quadratic solve, we rebuild the circuit graph because the B2B net model is placement-dependent. We then alternate quadratic solves and graph rebuilding until HPWL stops improving. In practice, this requires a small number of iterations (5-7), regardless of benchmark size, because the relative ordering of locations stabilizes quickly.

B. Look-ahead Legalization

Consider a set of cell locations with a significant amount of overlap as measured using bins of a regular grid. Look-ahead legalization changes the global positioning of those locations, seeking to remove most of the overlap (with respect to the grid) while preserving the relative ordering. This task can be formulated at different geometric scales by varying the grid. The quality of look-ahead legalization is measured by its impact on the entire placement flow. Our look-ahead legalization is based on top-down recursive geometric partitioning and non-linear scaling, as outlined in Algorithm 1. Cutlines C_c and C_B are chosen to be vertical at the top level and they alternate between horizontal and vertical directions with each successive level of top-down geometric partitioning.

Handling density constraints. For each grid bin of a given regular grid, we calculate the total area of contained cells A_c and the total available area of cell sites A_a. A bin is γ-*overfilled* if its *cell density* A_c/A_a exceeds given density limit $0 < \gamma < 1$. Adjacent γ-overfilled bins are clustered by Breadth-First Search

Algorithm 1 Look-ahead Legalization by Top-down Geometric Partitioning and Non-linear Scaling

Maximum allowed density γ, where $0 < \gamma < 1$
Current grid cell size
Floorplan with obstacles
Placement of cells
Queue of bin clusters $Q = \emptyset$

1: Identify γ-overfilled bins and cluster them // Fig. 2(a)
2: **foreach** cluster c **do**
3: Find a minimal rectangular region $R \supset c$ with density(R) $\leq \gamma$
4: $R.level$=1
5: Q.enqueue(R)
6: **while** !Q.empty() **do**
7: B=Q.dequeue()
8: **if** (Area(B) $<$ 4·grid cell size || $B.level \geq 10$) **then**
9: **continue**
10: M={movable cells in B}
11: **if** ($B.level$ % 2 == 0) **then** axis_direction D=HORIZ
12: **else** axis_direction D=VERT
13: C_c=D-aligned cutline to evenly split cell area in M
14: C_B=D-aligned cutline to evenly partition whitespace in B
15: (S_0, S_1)={two sub-regions of B created by cutline C_c}
16: M_0={movable cells in S_0}
17: M_1={movable cells in S_1}
18: (B_0, B_1)={two sub-regions of B created by cutline C_B}
19: Perform NON-LINEAR SCALING on $M_0 \perp$ to D in B_0
20: Perform NON-LINEAR SCALING on $M_1 \perp$ to D in B_1
21: $B_0.level$=$B_1.level$=$B.level$+1
22: Q.enqueue(B_0)
23: Q.enqueue(B_1)
24: **end while**
25: **end foreach**

(BFS), and look-ahead legalization is performed on such clusters. For each cluster, we find a minimal containing rectangular region with density $\leq \gamma$ (these regions can also be referred to as "clusters"). A key insight is that overlap removal in a region, which is filled to capacity, is more straightforward because the absence of whitespace leaves less flexibility for interconnect optimization.[4] If relative placement must be preserved, overlap can be reduced by means of x- and y-sorting with subsequent greedy packing. The

[4]In the presence of whitespace, the placer can move cells around without changing their relative ordering [2]. Removing whitespace suppresses this degree of freedom, giving fewer choices to the placer.

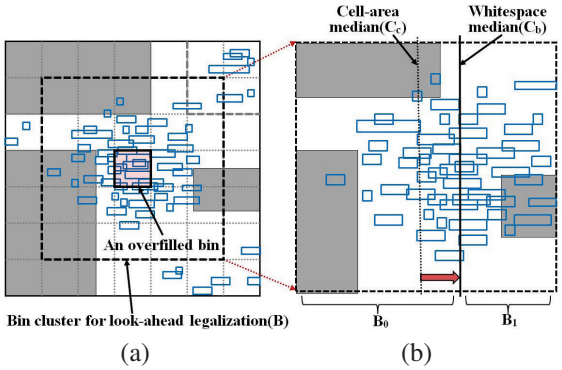

Fig. 2. Clustering of overfilled bins in Algorithm 1 and adjustment of cell-area to whitespace median by non-linear scaling (also see Figure 3). Movable cells are shown in blue, obstacles in solid gray.

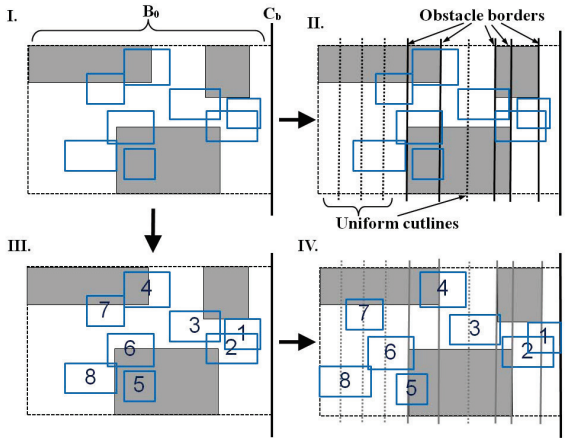

Fig. 3. Non-linear scaling in a region with obstacles (I): the formation of C_B-aligned stripes (II), cell sorting by distance from C_B (III), greedy cell positioning (IV).

next step, *non-linear scaling*, implements this intuition, but relies on cell-area cutline C_c chosen in Algorithm 1 and shifts it toward the median of available area C_B in the region, so as to equalize densities in the two sub-regions (Figure 2).

Non-linear scaling in one direction is illustrated in Figure 3, where a new region was created by a vertical cutline C_B during top-down geometric partitioning. This region is subdivided into vertical stripes parallel to C_B. First, cutlines are drawn along the boundaries of obstacles present in this region. Each vertical stripe created in this process is further subdivided (by up to 10 evenly distributed cutlines) if its available area exceeds $1/10$ of the region's available area. Movable cells in the corresponding sub-region created by C_c are then sorted by their distance from C_B and greedily packed into the stripes in that order. In other words, the cell furthest from the cutline is assigned to the furthest stripe. Each subsequent cell is assigned to the furthest stripe that is not filled yet.

For each stripe, we calculate the available site area A_a and consider the stripe filled when the area of assigned cells reaches γA_a. Cell locations within each stripe are linearly scaled from current locations (non-linearity arises from different scaling in different stripes).

Look-ahead legalization applies non-linear scaling in alternating directions, as illustrated in Figure 4 on one of ISPD 2005 benchmarks. Here, a region R is selected

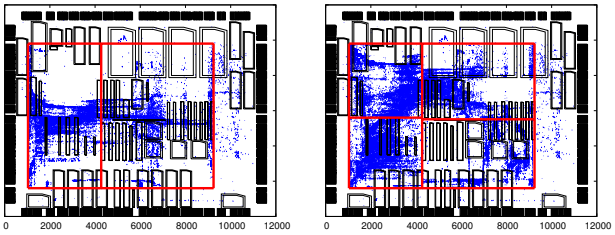

Fig. 4. Non-linear scaling after the first vertical cut and two subsequent horizontal cuts (ADAPTEC1) from intermediate steps between iterations 0 and 1 in Figure 7.

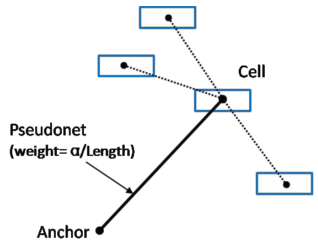

Fig. 5. An anchor with a pseudonet.

that contains overfilled bins, but is wide enough to facilitate overlap removal. R is first partitioned by a vertical cutline, after which non-linear scaling is applied in the two new sub-regions. Subsequently, look-ahead legalization (Algorithm 1) considers each sub-region individually and selects different horizontal cutlines. Four rounds of non-linear scaling follow, spreading cells over the region's expanse (Figure 4).

Despite a superficial similarity to cell-shifting in FastPlace [29], our non-linear scaling does not use cell locations to define bins/ranges, or map ranges onto a uniform grid.

Cutline shifting. Median-based cutlines are neither necessary nor sufficient for good solution quality. We therefore adopt a fast cutline positioning technique from [24]. On benchmarks whose obstacles cover $<20\%$ of total sites area, we find cutline positions C_c minimizing net cut for the top two levels of top-down geometric partitioning, with $<60\%$ of cell area per partition. We record the ratio ρ of cell areas in the two partitions and adjust the region's C_B cutline to the position that partitions the region's available area with the same ratio ρ. A related technique called ACG was developed at IBM in the context of min-cut placement, and their paper [2] describes relevant intuition.

C. Global Placement Iterations

Using legalized locations as anchors. Solving an unconstrained linear system results in a placement with significant amount of overlap. To pull cells away from their initial positions, we gradually perturb the linear system. As explained in Section IV-B, at each iteration of our global placement, top-down geometric partitioning and non-linear scaling generates a roughly legalized solution. We use these legalized locations as fixed, zero-area anchors connected to their corresponding cells in the lower-bound placement with artificial two-pin pseudonets. Furthermore, following the discussion in Section II, we note that connections to fixed locations do not increase the size of the Hessian matrix Q, and only contribute to its diagonal elements [17, Section 4.3.2]. This enhances diagonal dominance, condition number of $P^{-1}Q$, and the convergence rate of Jacobi-preconditioned CG.

In addition to weights given by the B2B net model on pseudonets, we control cell movement and iteration

convergence by multiplying each pseudonet weight by an additional factor $\alpha > 0$ computed as $\alpha = 0.01 \cdot (1 + \text{iterationNumber})$. At early iterations, small α values weaken spreading forces, giving greater significance to interconnect and more freedom to the linear system solver. As the relative ordering of cells stabilizes, increasing α values boost the pull toward the anchors and accelerate the convergence of lower bounds and upper bounds.

Grid resizing. To identify γ-overfilled bins, we overlay a uniform grid over the entire layout. The grid size is initially set to $S_{init} = 100 \times 100$ to accelerate the look-ahead legalization. However, in order to accurately capture the amount of overlap, the grid cell size decreases by $\beta = 1.06$ at each iteration of global placement until it reaches $4\times$ the average movable cell size.[5] Grid resizing also affects the clustering of γ-overfilled bins during look-ahead legalization (Section IV-B), effectively limiting the amount of cell movement and encouraging convergence at later iterations. A progression of global placement is annotated with HPWL values in Figure 7. The upper-bound placements on the right appear blocky in the first iteration, but gradually refine with grid resizing.

Convergence criteria. A convergence criterion similar to that in Section IV-A can be adopted in global placement. We alternate (1) look-ahead legalization, (2) updates to anchors and the B2B net model, and (3) solution of the linear system, until HPWL of solutions generated by look-ahead legalization stops improving. Unlike in the initial placement step, however, HPWL values of upper-bound solutions oscillate during the first 4-7 iterations, as illustrated in Figure 6. To prevent premature convergence, we monitor the gap between the lower and upper bounds. Global placement continues until (1) the gap is reduced to 25% of the gap in the 10th iteration and upper-bound solution stops improving or (2) the gap is smaller than 10% of the gap in the 10th iteration. On the ISPD 2005 benchmark suite, this convergence criterion entails 26-47 iterations of global placement. The final set of locations (global placement) is produced by the last look-ahead legalization as shown in Figure 1.

Convergence is guaranteed by the increasing weights of pseudonets. At each iteration, these pseudonets pull the lower-bound placement toward a legalized upper-bound placement. As the lower-bound placement becomes closer to a legal placement, it exhibits a decreasing amount of cell overlap. This, in turn, results in smaller cell displacements during look-ahead legalization. In the extreme, very high pseudonet weights force the lower-bound placement so close to the upper-bound placement, that look-ahead legalization does not change it, resulting in immediate convergence.[6] In practice, pseudonet weights are changed gradually to ensure low interconnect length. After the first few iterations, one typically observes monotonic convergence, as illustrated in Figure 6.

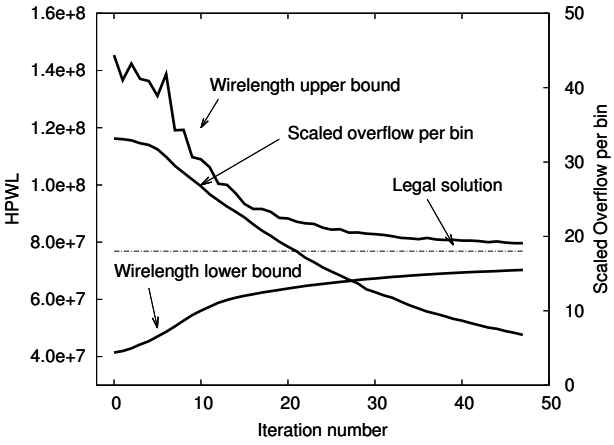

Fig. 6. Lower and upper bounds for HPWL, the scaled overflow per bin of the lower-bound placement at each iteration, and HPWL of the legal placement (ADAPTEC1)

D. Asymptotic Complexity Analysis

Modern placement algorithms are too complicated for asymptotic complexity analysis, but the bottlenecks of the SimPL algorithm yield to traditional analysis.

The runtime of global placement iterations is dominated by the Conjugate Gradient (CG) solver and look-ahead legalization. The complexity of each CG invocation is $O(m\sqrt{\kappa})$, where κ is the conditioning number of the matrix and m is the number of non-zero elements [28]. The number of non-zeros reflects the number of graph edges in the B2B model of the netlist. It grows linearly with the number of *pins* (cell-to-net connections) — a key size metric of a netlist. Another way to estimate the number of non-zeros is to observe that the average cell degree (the number of nets connected to a cell) is bounded by $d = 5$, or perhaps a slightly larger constant, for practical netlists.[7] Since $m \leq (d+1)n$ for n cells,[8] CG runs in $O(n\sqrt{\kappa})$ time.

Asymptotic runtime of look-ahead legalization is dominated by sorting cell locations by their x and y coordinates because non-linear scaling takes $O(n)$ time (several other linear-time steps take even less time in practice, therefore we do not discuss them). Given that look-ahead legalization operates on blocks of progressively smaller size, we can separately consider its

[5]This is similar to mesh sizing for finite-element methods in numerical analysis and especially in adaptive mesh refinement. Parameters can be optimized for benchmark suites using binary search. However, we have not tuned parameters to indiviual benchmarks.

[6]This convergence argument only assumes that look-ahead legalization does not change an upper-bound placement. It does not make any other assumptions about the look-ahead legalization algorithm or consistency of its results between iterations. Neither does it say anything about the quality of results.

[7]Even with large macros, whose number is limited by design area.

[8]Including diagonal matrix elements.

processing pass for the top-level blocks, then the pass for half-sized blocks, etc. Only $O(\log n)$ such passes are required for n cells. Each pass takes $O(n \log n)$ time because top-level blocks do not experience significant overlaps — in fact, each subsequent pass becomes faster because sorting is applied to smaller groups of cells. Hence, look-ahead legalization runs in $O(n \log^2 n)$ time.

We have observed that due to preconditioning, iteration counts in CG grow no faster than $\log n$, and each iteration takes linear time in n. Therefore one global placement iteration takes $O(n \log^2 n)$ time.

Empirically, SimPL requires $<$50 placement iterations, even for circuits with millions of cells. While the number of iterations might grow for larger circuits, this growth is very slow — possibly a polylog function of n. Empirical results in Section VI show that SimPL's advantage in runtime and solution quality over its closest competitor (FastPlace3) increases on larger netlists. Min-cut placement (Capo) exhibits asymptotic complexity $O(n \log^2 n)$, but lags behind SimPL in runtime and quality.

Space complexity of our algorithms is linear in the size of the input, and our implementations require a modest amount of memory.

V. Extensions and Improvements

The algorithm in Section IV can be improved in terms of runtime and solution quality. However, some of our attempts at improvement were unsuccessful. We report them here to warn the reader about their futility.

A. Selecting Windows for Look-ahead Legalization

During early global iterations, most movable cells of the lower-bound placement reside near the center of the layout region (Figure 7). In such cases, there is usually one expanded minimal rectangular region (cluster) that will encompass most of γ-overfilled bins. However, as global iterations progress, γ-overfilled bins will be scattered around the layout region, and multiple clusters of bins may exist. In our implementation, we process γ-overfilled bins in the decreasing order of density. Each expansion stops when the cluster's density drops to γ or the cluster abuts the boundaries of previously processed clusters. This strategy may generate incompletely expanded clusters, especially in mid-stages of global placement iterations. However, as the densest bins are processed first, the number of regions with peak density is guaranteed to decrease at every iteration except when the peak density itself decreases. At each iteration of global placement, look-ahead legalization is repeated up to ten times with increasing grid cell sizes until maximal density is decreased below γ.

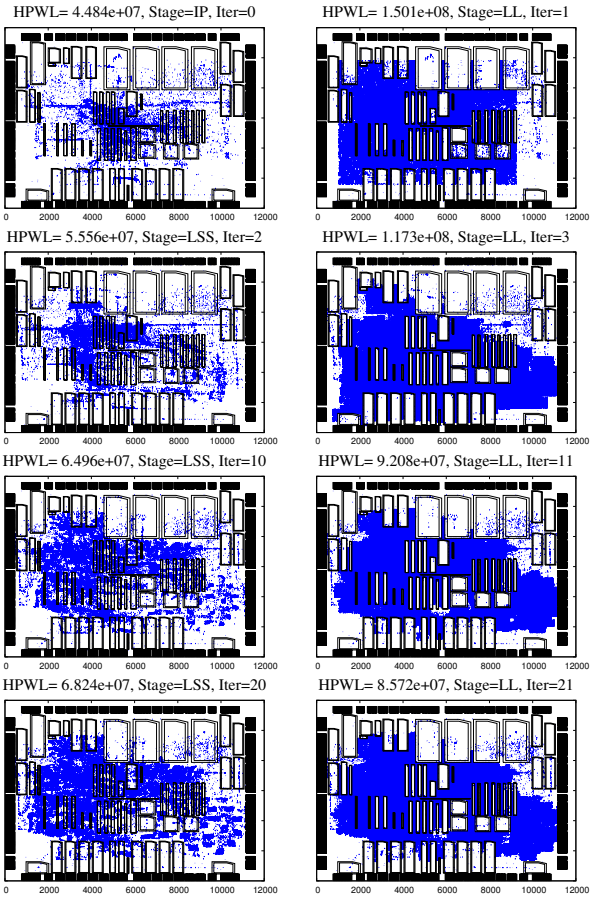

Fig. 7. A progression of global placement snapshots from different iterations and algorithm steps (adaptec1). IP=Initial Placement, LL=Look-ahead Legalization, LSS=Linear System Solver. Left-side placements show lower bounds and right-side placements show upper bounds.

B. Improving Asymptotic Runtime Complexity of Look-ahead Legalization

As explained in Section IV-D, asymptotic runtime of look-ahead legalization is largely determined by sorting cells by positions in directions perpendicular to cutlines. This sorting occurs at each level of top-down geometric partitioning. One way to improve asymptotic runtime complexity of look-ahead legalization is to invoke sorting less often, given that look-ahead legalization is to preserve the relative ordering among cells. Instead of sorting cells in each sub-region, we first establish two cell arrays sorted by x-coordinates and y-coordinates, respectively. At the second level of top-down geometric partitioning, two sub-regions inherit corresponding cells in-order from two sorted arrays of cells. In this way, if sorting is performed once at the top level of geometric partitioning, sorting at all successive levels can be replaced by selecting appropriate cells belong to current region in-order from two sorted arrays of higher level. This improves asymptotic runtime complexity of look-ahead legalization from $O(n \log^2 n)$ time to $O(n \log n)$.

Benchmark size (#cells)	APLACE2.0		CAPO10.5		FASTPLACE3.0		MPL6		NTUPLACE3		SIMPL	
	HPWL	Time	HPWL	Time	HPWL	Time	HPWL	Time	HPWL	Time	HPWL	Time
AD1 211K	78.35	35.02	88.14	25.95	78.16	2.50	77.93	18.36	81.82	8.20	**76.87**	**2.47**
AD2 255K	95.70	50.57	100.25	36.06	93.56	3.66	92.04	19.91	**88.79**	7.57	90.37	**3.40**
AD3 452K	218.52	119.53	276.80	78.19	213.85	8.48	214.16	58.92	214.83	15.62	**206.38**	**6.68**
AD4 496K	209.28	131.57	231.30	79.32	198.17	7.10	193.89	55.95	195.93	16.18	**186.00**	**5.88**
BB1 278K	100.02	44.91	110.92	41.78	96.32	3.77	96.80	22.82	98.41	13.22	**95.85**	**3.47**
BB2 558K	153.75	100.96	162.81	80.55	154.91	9.62	152.34	61.55	151.55	26.17	**143.56**	**7.58**
BB3 1.10M	411.59	209.24	405.40	182.94	365.59	21.59	344.10	85.23	360.07	51.08	**336.19**	**13.02**
BB4 2.18M	871.29	489.05	1016.19	567.15	834.19	40.93	829.44	189.83	866.43	115.06	**796.78**	**37.37**
Geomean	**1.09**×	**15.34**×	**1.20**×	**12.17**×	**1.05**×	**1.20**×	**1.03**×	**7.10**×	**1.05**×	**3.05**×	**1.00**×	**1.00**×

TABLE I
LEGAL HPWL (×10E6) AND TOTAL RUNTIME (MINUTES) COMPARISON ON THE ISPD 2005 BENCHMARK SUITE. EACH PLACER RAN AS A SINGLE THREAD ON A 3.2GHZ LINUX WORKSTATION. HPWL WAS COMPUTED BY THE GSRC BOOKSHELF EVALUATOR [1]. FULL NAMES OF BENCHMARKS ARE ABBREVIATED: "AD" FOR "ADAPTEC" AND "BB" FOR "BIGBLUE".

C. Unsuccessful Attempts at Improvement

Compared to other placement algorithms, SimPL uses a very modest set of interconnect optimizations. Therefore, we experimented with adding to SimPL several algorithms that were reported essential to the performance of other placers.

1. Our first attempt was to use netlist clustering to extend SimPL into a multilevel algorithm [7], [16], [29]. To this end, we implemented BestChoice clustering [20] used in FastPlace3 and were able to match its performance observed in FastPlace3 logs. This accelerated the initial CG solve in SimPL by about $2\times$, with essentially the same quality of results, but unclustering increased the amount of cell overlaps, and the refinement techniques that we tried were either ineffective or too time-consuming.

2. In our second attempt, we implemented *iterative local refinement* (ILR) [29], which is a stage of FastPlace-global where it spends 40-50% of its runtime. ILR is a simple move-based algorithm that post-processes results of quadratic placement by relocating cells to nearby grid bins, while keeping track of both HPWL and cell density. ILR did improve the results of our early prototypes, but adding it to SimPL does not improve final results. We believe that our *look-ahead legalization* algorithm provides sufficient density control with a moderate increase in HPWL. We also tried, unsuccessfully, the median-improvement (*BoxPlace*) algorithm from [18], which moves single cells to their HPWL-optimal locations, while considering adjacent cells fixed.

3. In a third attempt, we evaluated *ad hoc force modulation* used in RQL [30] that neglects 10% strongest forces. Sweeping the range from 1% to 10% did not reveal any improvement in our experiments.

4. In our fourth attempt at improvement, we reordered vertices in the netlist to improve memory locality for each invocation of CG. This technique is often applied to the matrices of linear systems and is known to reduce cache misses and runtime. We implemented the Reverse Cuthill-McKee (RCM) reordering, which is standard in numerical analysis. The locality of nets has significantly improved. However, CG did not run faster on any of our benchmarks — the default ordering in our benchmarks was already good enough.

In summary, we obtain state-of-the-art results without extensions reported essential to other placers (FastPlace3 [29], FDP [18], and RQL [30]). We have also experimented with several preconditioners for CG, but found the simplest of them — the diagonal (Jacobi) preconditioner — to work best in our application.

VI. EMPIRICAL VALIDATION

Our implementation was written in C++ and compiled with g++ 4.4.0. Unless indicated otherwise, benchmark runs were performed on an Intel Core i7 Quad CPU Q660 Linux workstation running at 3.2GHz, using only one CPU core. We compared SimPL to other academic placers on the ISPD 2005 placement contest benchmark suite with target density γ=1.0. Focusing on global placement, we delegate final legalization (into rows and sites) and detailed placement to FastPlace-DP [22], but post-process it by a greedy cell-flipping algorithm from Capo [6]. HPWL of solutions produced by each placer is computed by the GSRC Bookshelf Evaluator [1].

A. Analysis of Our Implementation

The SimPL global placer is a stand-alone tool that includes I/O, initial placement and global placement iterations. Living up to its name, it consists of fewer than 5,000 lines of C++ code and relies only on standard C++ libraries. There are four command-line parameters that affect performance — two for grid resizing (initial and step), and two for pseudonet weighting (initial and step). In all experiments we used default values described in Section IV.

Running in a single thread, SimPL completes the entire ISPD 2005 benchmark suite in 1 hour 18 minutes, placing the largest benchmark, BIGBLUE4 (2.18M cells), in 38 minutes using 2.1GB of memory. We report the runtime breakdown on BIGBLUE4 according to Figure 1, excluding 1.4% runtime for I/O.

Initial placement takes 5.0% of total runtime, of which 3.7% is spent in CG, and 1.3% in building B2B net models and sparse matrices for CG. **Global**

placement iterations take 47.4%, of which 19% is in the CG solver, and 9.9% is in sparse matrix construction and B2B net modeling. Inserting pseudonets takes 0.8%, and look-ahead legalization 17.7%. **Post-global placement** takes 46.2%, predominantly in detailed placement. Greedy orientation improvement and HPWL evaluation were almost instantaneous.

B. Comparisons to State-of-the-art Placers

We compared SimPL to other placers whose binaries are available to us.

We run each available placer,[9] including SimPL, in *default mode* and show results in Table I. The HPWL results reported by APlace2 [16], Capo10.5 [6], [25] and mPL6 [7] were confirmed by the GSRC Bookshelf Evaluator. However, FastPlace3 [29] reported lower HPWL by 0.25% to 0.96%. For consistency, we report the readings of the GSRC Bookshelf evaluator.

SimPL found placements with the lowest HPWL for seven out of eight circuits in the ISPD 2005 benchmark suite (no parameter tuning to specific benchmarks was employed). On average, SimPL obtains wirelength improvement of 7.73%, 16.47%, 4.38%, 2.98%, and 4.48% versus APlace2, Capo10.5, FastPlace3, mPL6, and NTUPlace3 respectively. SimPL was also the fastest among the placers on all eight circuits. It is 7.01 times faster than mPL6, which appears to be the strongest pre-existing placer. SimPL is 1.20 times faster than FastPlace3, which has been the fastest academic placer so far.

While we managed to obtain almost all best-performing academic placers in binaries, RQL reportedly outperforms mPL6 in HPWL by a small amount [30]. Comparing our HWPL results to numbers in [30], we observe five wins for SimPL and three losses. RQL is 3.1 times faster than mPL6, making it more than twice as slow as SimPL.

C. Scalability Study

To demonstrate SimPL's scalability to larger netlists, we generated variants of ISPD 2005 benchmarks with netlists that are twice as big with the same area utilization. In such a double-sized benchmark, each movable cell is split in two cells of smaller size, and each connection to the original cell is inherited by one of the split cells. Additionally, the two split cells are connected by a new two-pin net (Figure 8).

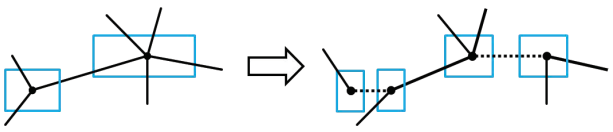

Fig. 8. Generation of double-sized netlists.

[9]The KraftWerk2 binary we obtained did not run on our system.

We compared SimPL to FastPlace3, mPL6, and NTUPlace3 on the double-sized benchmark suite and show results in Table II. mPL6 could not finish bigblue4. For bigblue3, FastPlace-DP was unable to completely legalize solutions produced by FastPlace3-global, hence we post-processed FastPlace-DP with Capo10.5's legalizer.

SimPL was the fastest among the placers on all eight circuits. It is 8.96 times faster than mPL6, and 1.49 times faster than FastPlace3. SimPL also found placements with the lowest HPWL for six out of eight circuits in the double-sized ISPD 2005 benchmark suite (no parameter tuning to specific benchmarks was employed). Comparing results in Table II to those in Table I, we observe that our placer has greater advantage on larger benchmarks. Furthermore, our runtime comparisons include detailed placement, but if SimPL is compared to FastPlace3-global without detailed placement, the average speed-up increases to 1.82 times from 1.58 times.

Compared to other placers, our implemantation uses a modest amount of memory — 1.65 times and 2.39 times less than mPL6 and NTUPlace3 repsectively, and 1.61 times more than FastPlace3. SimPL is using more memory than FastPlace3 when it constructs sparse matrices based on the Bound2Bound net model.

VII. Speeding up Placement Using Parallelism

Further speed-up is possible for SimPL on workstations with multicore CPUs.

A. Algorithmic Details

Runtime bottlenecks in the sequential variant of the SimPL algorithm (Section VI-A) — updates to the B2B net model and the CG solver — can be parallelized. Given that the B2B net model is separable, we process the x and y cases in parallel. When more than two cores are available, we split the nets of the netlist into equal groups that can be processed by multiple threads. To parallelize the CG solver, we applied a coarse-grain *row partitioning* [13] scheme to the Hessian Matrix Q, where different blocks of rows are assigned to different threads using OpenMP [11]. A critical kernel operation in CG is the Sparse Matrix-Vector multiply (SpMxV). Memory bandwidth is a known performance bottleneck in a uniprocessor environment [10], and its impact is likely to aggravate when multiple cores access the main memory through a common bus. We reduce memory bandwidth demand of SpMxV by using the *CSR* (Compressed Sparse Row) [26] memory layout for the Hessian matrix Q.

In addition to thread-level parallelism, our implementation makes use of streaming SIMD extensions

Ckts	FastPlace3.0			mPL6			NTUplace3			SimPL		
	HPWL	Time	Memory	HPWL	Time	Memory	HPWL	Time	Memory	HPWL	Time	Memory
adlx2	80.30	3.67	**0.20**	79.11	24.95	0.43	80.20	19.6	0.72	**77.02**	**3.05**	0.36
ad2x2	98.88	6.06	**0.23**	93.64	38.00	0.82	**91.56**	17.6	0.86	92.41	**4.11**	0.36
ad3x2	258.71	13.47	**0.42**	232.87	88.71	1.11	225.32	44.7	1.62	**215.56**	**7.58**	0.71
ad4x2	219.35	11.54	**0.44**	206.24	85.72	1.16	197.90	39.6	1.76	**193.18**	**7.25**	0.69
bb1x2	97.93	5.68	**0.25**	100.37	30.35	0.53	99.33	22.0	0.93	**96.39**	**4.71**	0.43
bb2x2	164.74	12.13	**0.49**	159.24	79.84	1.22	154.47	44.2	1.94	**148.43**	**9.02**	0.77
bb3x2	515.61	49.89	**0.93**	395.26	172.96	3.38	**386.65**	154.9	3.82	403.40	**22.24**	1.44
bb4x2	865.30	56.36	**1.94**	fail	fail	fail	866.78	267.9	7.86	**854.64**	**42.43**	2.91
Geomean	**1.10×**	**1.49×**	**0.62×**	**1.04×**	**8.96×**	**1.65×**	**1.02×**	**4.40×**	**2.39×**	**1.00×**	**1.00×**	**1.00×**

TABLE II
LEGAL HPWL (×10E6), TOTAL RUNTIME (MINUTES), AND PEAK MEMORY USAGE (GIGABYTES) COMPARISON ON THE DOUBLE-SIZED ISPD 2005 BENCHMARK SUITE. THE FAILURE OF MPL6 IN BIGBLUE4X2 DOES NOT APPEAR TO BE CAUSED BY OUT-OF-MEMORY CONDITIONS.

level 2 (SSE2) [23] (through g++ intrinsics) that perform several floating-point operations at once. SSE2 instructions are extensively used in our CG solver. Since SSE2 instructions are available in most modern CPUs, we used them in the default mode evaluated in Table I and Table II. The overall speed-up due to parallelism varies between different hardware systems, as it depends on the relation between CPU speed and memory bandwidth.

```
// inner product of two float vectors x and y
float inner_product(vector<float>&x, vector<float>&y)
{
  float p_acc[4], inner_product=(float)0.;
  __m128 X, Y, acc = _mm_set_zero_ps();
  unsigned i;
  #pragma omp parallel for schedule(static)
  private(X,Y) lastprivate(i) reduction(+:acc)
  num_threads(NUM_CORES)
  for (i=0 ; i<=x.size()-4 ; i+=4)
  {
     X = _mm_load_ps(&x[i]);
     Y = _mm_load_ps(&y[i]);
     acc = _mm_add_ps(acc, _mm_mul_ps(X, Y));
  }
  _mm_store_ps(p_acc, acc);
  inner_product = p_acc[0]+p_acc[1]+p_acc[2]+p_acc[3];
  for ( ; i<x.size() ; i++)
        inner_product+=x[i]*y[i];
  return inner_product;
}
```

Listing 1. Sample code for OpenMP and SSE2 parallelization for the inner-product operation.

After we parallelized the main bottlenecks, we noticed that look-ahead legalization started consuming a significant fraction of overall runtime. Fortunately, top-down geometric partitioning and non-linear scaling are amenable to parallelization as well. Notably, top-down partitioning generates an increasing number of sub-tasks of similar sizes which can be solved in parallel. Let Q be the global queue of bin clusters, as defined in Algorithm 1, and each thread has a private queue of bin clusters Q_i. First, we statically assign initial bin clusters to available threads such that each thread has similar number of bin clusters to start. After each level of top-down geometric partitioning and non-linear scaling on such bin cluster, each thread generates two sub-clusters with similar numbers of cells. Then the thread t_i adds only one of two sub-clusters to its own cluster queue Q_i for the next level of top-down geometric partitioning and non-linear scaling, while the remainder is added to the global cluster queue Q. Whenever Q_i of a thread t_i becomes empty, the thread t_i dynamically retrieves clusters from the global cluster queue Q. The number of clusters to be retrieved N is given by

$$N = \max(Q.size()/N_{threads}, 1)$$

where $N_{threads}$ is the total number of threads.

B. Empirical Studies

As part of our empirical validation, we ran SimPL on an 8-core AMD-based system with four dual-core CPUs and 16GByte RAM. Each CPU was Opteron 880 processor running at 2.4GHz with 1024KB cache. Single-thread execution was compared to eight-thread execution as shown in Table III. Our combination of multi-threading and SIMD instruction-level parallelization was 1.6 times faster on average than parallelization based on multi-threading alone. Theoretically, using SIMD instruction-level parallelization may speed-up CG by at most four times. However, SIMD-based implementation of SpMxV only provided marginal speed-ups and was not worth the development effort. This is because irregular memory access patterns of SpMxV prohibit the aligned loading of values (*MOVAPS* or *_mm_load_ps* in Listing 1) to SSE registers. Nevertheless, SSE instructions were helpful in other parts of the code and contributed to the overall speed-up in global placement, as illustrated in Table IV.

We note that look-ahead legalization operates on large datasets, but performs little computation per datum, which limits its performance by memory bandwidth. The amount of work per thread is so small that the overhead of thread creation outweighed the benefits. As a result, this part of SimPL scales poorly to >4 threads on available hardware, although this is probably not a fundamental limitation of the algorithm.

The overall speed-ups in global placement runtimes are shown in Table IV. Solution quality did not appreciably change, but peak memory usage increased by 1.91 times whereas runtime of global placement iterations was reduced by 2.4 times on average. The speed-ups saturate for more than 4 threads as look-ahead

Ckts	1 CORE	2 THREADS				4 THREADS				8 THREADS			
	CG+SSE	CG	CG+SSE	B2B	T&N	CG	CG+SSE	B2B	T&N	CG	CG+SSE	B2B	T&N
AD1	1.37	1.88	2.21	1.45	1.41	2.03	2.87	1.64	1.59	1.92	3.28	1.91	1.48
AD2	1.61	1.77	2.09	1.50	1.53	2.12	3.01	2.04	2.17	2.06	3.22	1.98	1.40
AD3	1.57	1.76	2.20	1.48	1.62	1.88	3.17	1.62	2.25	2.00	3.55	1.79	1.66
AD4	1.50	1.65	2.07	1.51	1.56	1.81	3.03	1.58	2.17	1.77	3.33	1.72	1.36
BB1	1.57	2.03	2.11	1.27	1.71	2.02	3.14	1.66	2.93	2.05	3.70	1.78	2.93
BB2	1.62	2.07	2.24	1.48	1.49	1.72	2.97	1.68	1.89	1.79	3.50	1.74	1.58
BB3	1.54	1.53	2.25	1.60	1.32	1.68	3.04	1.64	2.04	1.81	3.30	1.85	1.32
BB4	2.01	2.63	3.04	2.01	1.59	2.71	4.48	2.02	2.12	2.76	5.12	2.18	1.68
GM	**1.59×**	**1.89×**	**2.26×**	**1.53×**	**1.52×**	**1.98×**	**3.18×**	**1.73×**	**2.12×**	**2.00×**	**3.59×**	**1.86×**	**1.62×**

TABLE III
SPEED-UP RATIOS FOR CONJUGATE GRADIENT (CG), B2B NET MODEL CONSTRUCTION (B2B), AND TOP-DOWN GEOMETRIC PARTITIONING AND NON-LINEAR SCALING (T&N) ON THE ISPD 2005 BENCHMARK SUITE. RUNTIMES ARE COMPARED TO SINGLE-THREADED EXECUTION WITHOUT SUPPORT OF SSE INSTRUCTIONS.

Ckts	2 THREADS		4 THREADS		8 THREADS	
	no SSE	SSE	no SSE	SSE	no SSE	SSE
AD1	1.70	1.71	1.76	2.03	1.71	2.23
AD2	1.75	1.73	1.91	2.43	1.90	2.35
AD3	1.59	1.72	1.79	2.30	1.81	2.40
AD4	1.55	1.65	1.75	2.24	1.67	2.26
BB1	1.75	1.67	2.17	2.56	2.18	2.67
BB2	1.70	1.72	1.67	2.22	1.66	2.37
BB3	1.49	1.75	1.71	2.28	1.65	2.28
BB4	1.94	2.12	2.01	2.55	2.03	2.69
GM	**1.68×**	**1.75×**	**1.84×**	**2.32×**	**1.82×**	**2.40×**

TABLE IV
SPEED-UP RATIOS FOR GLOBAL PLACEMENT ON THE ISPD 2005 BENCHMARK SUITE. RUNTIMES ARE COMPARED TO SINGLE-THREADED EXECUTION WITHOUT SUPPORT OF SSE INSTRUCTIONS.

legalization scales poorly. The initial placement stage was accelerated by about 3 times. While CG remained the runtime bottle neck of SimPL on 8 threads (36% of global placement), look-ahead legalization became a close second ($> 31\%$ of global placement).

VIII. CONCLUSIONS AND FUTURE WORK

In this work, we developed a new, flat, partition-based and force-directed quadratic global placer. Unlike other state-of-the-art placers, it is rather simple, and our self-contained implementation includes fewer than 5,000 lines of C++ code. The algorithm is iterative and maintains two placements — one computes a lower bound and one computes an upper bound on final wirelength. These two placements interact, ensuring stability and fast convergence of the algorithm. The upper-bound placement is produced by a new *look-ahead legalization* algorithm, based on top-down geometric partitioning and non-linear scaling, and converges to final cell locations. In contrast, all analytic algorithms we reviewed (both force-directed quadratic and non-convex) derive their final solution from a lower-bound placement.

The use of partition-based techniques in upper-bound placements offers a solution to the force-modulation problem [18], [30] and removes the need for the so-called *hold forces* used by several force-directed placers.[10] As discussed in Section III, upper-bound placements perform an *area look-ahead*[11] that is instrumental in the handling of layout obstacles. APlace2, NTUPlace3, mPL6, as well as some force-directed placers, model obstacles by additional smoothened *penalty* terms in the objective function. Not only such terms introduce extra work, but they also add imprecisions to modeling. For similar reasons, SimPL avoids netlist clustering used by other placers. We have implemented several other techniques essential to well-known placers, such as BoxPlace [18], ILR [29], and *ad hoc* force modulation [30], but they did not improve SimPL results.

SimPL is highly competitive on ISPD 2005 benchmarks where it outperforms every placer available to us in binary both by solution quality and runtime. SimPL's advantage in runtime and solution quality over FastPlace3 and mPL6 grows on larger netlists. However, its most compelling advantages over prior state of the art deal with practical uses of placement in modern timing-closure design flows: (1) the reduced complexity of SimPL allows for fast implementation, parallel processing, and effective software maintenance; (2) the upper-bound placements facilitate tighter integration of timing and congestion optimizations into the global placement process, improving the speed and quality of physical synthesis.

The SimPL algorithm saw rapid adoption since its first publication at ICCAD 2010. At the ISPD 2011 placement contest, the winning team successfully implemented SimPL without having access to our source code. To our knowledge, at least two major EDA vendors are now using similar placement algorithms, and our own work with a state-of-the-art industry placer quickly produced significant improvements that will be discussed in future publications.

[10]Hold forces are used to ensure that the current placement is a force equillibrium. Then move forces are added so that the placement can be improved. While this techniques is needed to ensure convergence of iterations, SimPL relies on anchors and pseudonets to ensure convergence.

[11]The concept of *area look-ahead* was proposed in [9] for block-packing by nested bisection, where it checks if a given bisection admits a legal block packing in each partition. Area look-ahead was not used in [9] to spread standard cells from dense regions.

The implementation of SimPL described in this paper is designed for standard-cell layouts and does not yet handle movable macro-blocks. Our recent industry experience suggests that the majority of modern large-scale placement instances in practice do not require this feature, as their macro-blocks are fixed. However, mixed-size placement is useful for some mixed-signal SoCs, and we are addressing it in our ongoing work.[12]

Attempting to explain theoretically the strong performance of our placement algorithm, we have noticed similarities to *primal-dual* algorithms for convex [31] and combinatorial [5] optimization. Primal-dual methods maintain lower and upper bounds, expressed by primal and dual solutions that eventually converge to an optimal feasible solution. The interpretation of duality as swapping the problem's constraints for the objective function [31] is also consistent with our algorithm — *look-ahead legalization corresponds to imposing a no-overlap constraint while relaxing the linear constraints that capture the global minimum of the quadratic wirelength objective*. The key to the success of primal-dual algorithms [5], [31] is the observation that alternating progress in *primal* and *dual* solutions, i.e., improving the cost of feasible solutions and tightening the constraints for low-cost solutions, typically leads to faster convergence compared to one-sided optimizations. This effect is empirically observed in Section VI where SimPL is compared to pre-existing placement algorithms, all of which are one-sided.

REFERENCES

[1] S. N. Adya, I. L. Markov, "Executable Placement Utilities," http://vlsicad.eecs.umich.edu/BK/PlaceUtils/

[2] C. J. Alpert, G.-J. Nam. P. G. Villarrubia, "Effective Free Space Management for Cut-based Placement via Analytical Constraint Generation," *IEEE TCAD* 22(10) 2003, pp. 1343-1353.

[3] C. J. Alpert et al., "Techniques for Fast Physical Synthesis," *Proc. IEEE* 95(3), 2007, pp. 573-599.

[4] U. Brenner, M. Struzyna, J. Vygen, "BonnPlace: Placement of Leading-Edge Chips by Advanced Combinatorial Algorithms," *IEEE TCAD* 27(9) 2008, pp.1607-20.

[5] N. Buchbinder, J. Naor, *The Design of Competitive Online Algorithms via a Primal-Dual Approach*, NOW Publishers, 2009.

[6] A. E. Caldwell, A. B. Kahng, I. L. Markov, "Can Recursive Bisection Alone Produce Routable Placements?" *DAC* 2000.

[7] T. F. Chan et al., "mPL6: Enhanced Multilevel Mixed-Size Placement," *ISPD* 2006, pp. 212-214.

[8] T.-C. Chen et al.,"NTUPlace3: An Analytical Placer for Large-Scale Mixed-Size Designs With Preplaced Blocks and Density Constraints," *IEEE TCAD* 27(7) 2008, pp.1228-1240.

[9] J. Cong, M. Romesis, J. R. Shinnerl, "Fast Floorplanning by Look-ahead Enabled Recursive Bipartitioning," *TCAD* 25(9), 2006, pp. 1719-1732.

[10] G. Goumas et al., "Understanding the Performance of Sparse Matrix-Vector Multiplication," *Euromicro Int'l Conf. on PDP* 2008, pp. 283-292.

[11] L. Dagum, R. Menon, "OpenMP: An Industry Standard API for Shared-memory Programming," *IEEE Computational Science and Enginnering*, 1998, pp. 46-55.

[12] M. Garland, "Sparse Matrix Computations on Many-core GPU's," *Proc. DAC* 2010, pp. 2-6.

[13] L. Hsu et al., "Exploring the Cache Design Space for Large Scale CMPs," *ACM SIGARCH Computer Architecture News* 2005, pp. 24-33.

[14] B. Hu, M. Marek-Sadowska, "FAR: Fixed-points Addition & Relaxation Based Placement," *ISPD* 2005, pp. 161-166.

[15] D. A. Jamsek, "Designing and Optimizing Compute Kernels on NVIDIA GPUs," *Proc. ASPDAC* 2009, pp 224-229.

[16] A. B. Kahng, Q. Wang, "A Faster Implementation of APlace," *ISPD* 2006, pp. 218-220.

[17] A. B. Kahng, J. Lienig, I. L. Markov, J. Hu, "VLSI Physical Design: from Graph Partitioning to Timing Closure," Springer 2011, 312 pages.

[18] A. A. Kennings, K. Vorwerk, "Force-Directed Methods for Generic Placement," *IEEE TCAD* 25(10), 2006, pp. 2076-2087.

[19] K. Madduri et al., "Parallel Shortest Path Algorithms for Solving Large-Scale Instances", The Shortest Path Problem: Ninth DIMACS Implementation Challenge, pp. 249-290, *DIMACS*, 2009.

[20] G.-J. Nam et al., "A Fast Hierarchical Quadratic Placement Algorithm," *IEEE TCAD* 25(4), 2006, pp.678-691.

[21] G.-J. Nam, J. Cong, "Modern Circuit Placement: Best Practices and Results," *Springer*, 2007.

[22] M. Pan, N. Viswanathan, C. Chu, "An Efficient & Effective Detailed Placement Algorithm," *ICCAD* 2005, pp. 48-55.

[23] S. K. Raman, V. Pentkovski, J. Keshava, "Implementing Streaming SIMD Extensions on the Pentium III Processor" *IEEE Micro* 20(4) 2000, pp. 47-57.

[24] J. A. Roy, I. L. Markov, "ECO-System: Embracing the Change in Placement," *IEEE TCAD* 26(12) 2007, pp. 2173-2185.

[25] J. A. Roy et al., "Capo: Robust and Scalable Open-source Min-cut Floorplacer," *ISPD* 2005, pp. 224-226.

[26] Y. Saad, "Iterative Methods for Sparse Linear Systems," *SIAM* 2003.

[27] P. Spindler, U. Schlichtmann, F. M. Johannes, "Kraftwerk2 - A Fast Force-Directed Quadratic Placement Approach Using an Accurate Net Model," *IEEE TCAD* 27(8) 2008, pp. 1398-1411.

[28] L. N. Trefethen, D. Bau "Numerical Linear Algebra," *SIAM* 1997, pp. 296-298.

[29] N. Viswanathan, M. Pan, C. Chu, "FastPlace 3.0: A Fast Multilevel Quadratic Placement Algorithm with Placement Congestion Control," *ASPDAC* 2007, pp. 135-140.

[30] N. Viswanathan et al., "RQL: Global Placement via Relaxed Quadratic Spreading and Linearization," *DAC* 2007, pp. 453-458.

[31] S. J. Wright, "Primal-Dual Interior-Point Methods," *SIAM* 1987, 309 pages.

[12] Academic placers typically introduce this feature in dedicated publications, rather than in the first publication describing the baseline algorithm.

Myung-Chul Kim received the B.S. degree in Electronic and Electrical Engineering from the Pohang University of Science and Technology (POSTECH), South Korea, in 2006 and the M.S. in Electrical Engineering at the University of Michigan, in 2009. Currently he is a Ph.D. candidate at the University of Michigan. His research interests include VLSI Physical Design Automation with emphasis on placement, routing and timing analysis. Myung-Chul Kim is the winner of the ISPD 2010 clock-network synthesis contest and the recipient of the IEEE/ACM William J. McCalla Best Paper Award at ICCAD 2010. In 2011, he worked as a research intern at IBM Research.

Dong-Jin Lee was born in Gwangju, South Korea. He received the B.S. degree in electrical engineering from Seoul National University, South Korea, in 2006 and the M.S. and Ph.D. degrees in Electrical Engineering from University of Michigan, Ann Arbor, in 2008 and 2011 respectively. His research interests include EDA clock network synthesis, placement and routing algorithms. He is the winner of the ISPD 2009 and 2010 clock-newtwork synthesis contests and the recipient of the IEEE/ACM William J. McCalla Best Paper Award at ICCAD 2010. In 2009 and 2010, he worked as a research intern at Texas Instruments.

Igor L. Markov is an associate professor of Electrical Engineering and Computer Science at the University of Michigan. He received his Ph.D. in Computer Science from UCLA. He is a member of the Executive Board of ACM SIGDA, Editorial Board member of the Communications of ACM and IEEE Design & Test, as well as several ACM and IEEE Transactions. He chaired tracks at DAC, ICCAD, ICCD, DATE and GLSVLSI. Prof. Markov researches computers that make computers. He has co-authored three books and more than 180 refereed publications, some of which were honored by the best-paper awards at the Design Automation and Test in Europe Conference (DATE), the Int'l Symposium on Physical Design (ISPD) and IEEE Trans. on Computer-Aided Design. During the 2011 redesign of the ACM Computing Classification System, Prof. Markov lead the effort on the Hardware tree. Prof. Markov is the recipient of a DAC Fellowship, an ACM SIGDA Outstanding New Faculty award, an NSF CAREER award, an IBM Partnership Award, a Microsoft A. Richard Newton Breakthrough Research Award, and the inaugural IEEE CEDA Early Career Award.