

Closing the Gap: Near-Optimal Steiner Trees in Polynomial Time

Jeff Griffith, Gabriel Robins, *Member, IEEE*, Jeffrey S. Salowe, and Tongtong Zhang

Abstract—The minimum rectilinear Steiner tree (MRST) problem arises in global routing and wiring estimation, as well as in many other areas. The MRST problem is known to be NP-hard, and the best performing MRST heuristic to date is the Iterated 1-Steiner (I1S) method recently proposed by Kahng and Robins. In this paper, we develop a straightforward, efficient implementation of I1S, achieving a speedup factor of three orders of magnitude over previous implementations. We also give a parallel implementation that achieves near-linear speedup on multiple processors. Several performance-improving enhancements enable us to obtain Steiner trees with average cost within 0.25% of optimal, and our methods produce optimal solutions in up to 90% of the cases for typical nets. We generalize I1S and its variants to three dimensions, as well as to the case where all the pins lie on k parallel planes, which arises in, e.g., multi-layer routing. Motivated by the goal of reducing the running times of our algorithms, we prove that any pointset in the Manhattan plane has a minimum spanning tree (MST) with maximum degree 4, and that in three-dimensional Manhattan space every pointset has an MST with maximum degree of 14 (the best previous upper bounds on the maximum MST degree in two and three dimensions are 6 and 26, respectively); these results are of independent theoretical interest and also settle an open problem in complexity theory.

I. INTRODUCTION

THE MINIMUM rectilinear Steiner tree problem is central to VLSI physical design phases such as global routing and wiring estimation, where we seek low-cost topologies to interconnect the pins of signal nets [30], [37]:

The Minimum Rectilinear Steiner Tree (MRST) problem: Given a set P of n points, find a set S of Steiner points such that the minimum rectilinear spanning tree (MST) over $P \cup S$ has minimum cost.

The cost of a tree edge is the Manhattan distance between its endpoints, and the cost of a tree is the sum of its edge costs. The terms *points* and *pins* are synonymous and are used interchangeably, depending on the context (a *net* is a set of pins). Fig. 1 shows an MST and an MRST for the same four-pin net.

Research on the MRST problem has been guided by several fundamental results. First, Hanan [21] has shown that there

Manuscript received July 27, 1993; revised May 27, 1994. This work was supported in part by NSF Young Investigator Award MIP-9457412 (awarded to G. Robins) and by NSF grants MIP-9107717 and CCR-9224789. This paper was recommended by Associate Editor M. Sarrafzadeh.

J. Griffith is with the Department of Computer Science, University of Tennessee, Knoxville, TN 37996-1301 USA.

G. Robins and T. Zhang are with the Department of Computer Science, University of Virginia, Charlottesville, VA 22903-2442 USA.

J. S. Salowe is with Qwestech, Inc., Falls Church, VA 22043 USA.

IEEE Log Number 9404037.




Fig. 1. A minimum spanning tree (left) and MRST (right) for a fixed net; hollow dots represent the original pointset P , while solid dots represent the set S of added Steiner points.




Fig. 2. Hanan's theorem: There always exists an MRST with Steiner points chosen from the intersections of all the horizontal and vertical lines passing through all the points.

always exists an MRST with Steiner points chosen from the intersections of all the horizontal and vertical lines passing through all the points in P (see Fig. 2), and this result was generalized by Snyder [42] to all higher dimensions. However, a second major result by Garey and Johnson [16] establishes that despite this restriction on the solution space, the MRST problem remains NP-complete; this has given rise to numerous heuristics as surveyed by Hwang, Richards, and Winter [27].

In solving intractable problems, we often seek provably good heuristics having bounded worst-case error from optimal. Thus, a third important result is the discovery by Hwang [25] that the rectilinear MST is a fairly good approximation to the MRST, with a worst-case performance ratio of $\frac{\text{cost}(\text{MST})}{\text{cost}(\text{MRST})} \leq \frac{3}{2}$. This implies that any MST-based strategy that improves upon an initial MST topology will also enjoy a performance ratio of at most $\frac{3}{2}$. This has prompted a large number of Steiner tree heuristics that resemble classic MST construction methods [23], [24], [26], [31], [32], all producing Steiner trees with average cost 7 to 9% smaller than the MST cost [27].

Unfortunately, all MST-based MRST constructions were recently shown by Kahng and Robins [29] to have a worst-case performance ratio of exactly $\frac{3}{2}$. This negative result has motivated research into alternate schemes for MRST approximation, with the best performing among these being the Iterated 1-Steiner (I1S) algorithm [28], [38]. I1S always

performs strictly better than $\frac{3}{2}$ times optimal,¹ and also performs very well in practice, achieving almost 11% average improvement over MST cost, which is on average less than 0.5% away from optimal [40]. The Iterated 1-Steiner method was generalized to arbitrary weighted graphs by Alexander and Robins [1], [2], and is thus a suitable basis of a practical global router that must handle congestion, obstacles, etc. [15].

The performance success of IIS was achieved at the expense of a high time complexity: although a more efficient variant of IIS can be implemented to run within time $O(n^2 \log n)$ [28], the computational geometric methods employed to achieve this time bound hide large constant factors and are also difficult to code. Thus, actual previous implementations of IIS typically use a more straightforward approach that requires time $O(n^4 \log n)$.

The first contribution of this paper is a practical implementation of IIS that runs within time $O(n^3)$. Our method is based on a dynamic minimum spanning tree update scheme and establishes the practicality and viability of the iterated 1-Steiner approach. For 100 points, our new implementation is about three orders of magnitude faster than the naive implementation, and the speedup over the naive implementation increases with the number of points. This dramatic improvement in speed enables for the first time the testing of IIS on nets containing several hundred pins.

Since a typical CAD environment consists of a network of workstations, exploiting the available large-grain parallelism provides a natural and effective means of improving the performance of CAD algorithms. With this in mind, a second contribution of our work is a parallel version of IIS that achieves high parallel speedups. Since Steiner tree construction is a computationally expensive part of global routing, our parallel implementation may be viewed as an important first step toward obtaining a "Steiner engine," i.e. an efficient tool for producing near-optimal Steiner trees.

Our third contribution entails several performance-improving enhancements to the IIS method. Our methods rely on an approach that deviates from pure greed and instead employs randomness in order to break ties during Steiner point selection. Although asymptotically no slower than the original IIS variants, our methods afford improved average performance. Extensive simulations indicate that for uniformly distributed random nets of up to 8 pins, the average performance of our enhanced IIS algorithm is only 0.25% away from optimal. Moreover, for 8-pin nets, our method produces the optimal Steiner tree in 90% of all instances. We also propose a method of improving performance at the expense of running time, allowing a smooth tradeoff between solution quality and efficiency.

¹Recently, Berman and Ramaiyer [7] and Foessmeier, Kaufmann, and Zelikovsky [6], [13] have extended the fundamental work of Zelikovsky [46], [47] to yield a method similar to IIS (specifically, to the "batched" IIS method described below) with performance ratio bounded by $\frac{11}{8}$; this work settles in the affirmative the longstanding open question of whether there exists a polynomial-time rectilinear Steiner tree heuristic with performance ratio strictly smaller than $\frac{3}{2}$ [25]. At the time of this writing, Berman, Foessmeier, Karpinski, Kaufmann, and Zelikovsky [13] further improved the performance bound of their polynomial-time rectilinear Steiner heuristic to $\frac{91}{45} = 1.271$.

Next, we generalize IIS and its variants to three dimensions, as well as to the intermediate case where all pins lie on k parallel planes. This formulation has several applications, including multi-layer routing [9], [20], [22] and the design of buildings [41]. Empirical testing suggests that this approach is effective for three-dimensional Steiner routing, yielding up to 15% average improvement over MST cost in three-dimensional Manhattan space.

Finally, in order to reduce the running time of the dynamic MST-maintenance component of our algorithms, we prove the following results under the Manhattan metric: 1) every two-dimensional pointset has an MST with maximum degree of at most 4; and 2) every three-dimensional pointset has an MST with maximum degree of at most 14 (the best previously known bounds for two and three dimensions were 6 and 26, respectively). Our results and algorithms on degree-bounded minimum spanning trees are of significant independent theoretical interest [39] and settle several open issues in complexity theory.²

The rest of the paper is organized as follows. In Section II we review the IIS method. Section III outlines our more efficient implementation of IIS, and Section IV describes a variant of IIS with enhanced performance. Section V generalizes IIS to three dimensions and proposes an efficient method for three-dimensional MST-maintenance. Section VI proves tight bounds for the maximum MST degree in two and three dimensions under the Manhattan metric and discusses the implications of our results to a problem in complexity theory. Section VII outlines the parallel implementation, and Section VIII presents extensive simulation results regarding performance, running times, and parallel speedups. We conclude in Section IX with directions for further research. A preliminary version of this work appeared in [3] and in [4].

II. REVIEW OF THE ITERATED 1-STEINER METHOD

We begin with a review of the Iterated 1-Steiner method of Kahng and Robins [28]. For two pointsets P and S , we define the MST savings of S with respect to P as $\Delta\text{MST}(P, S) = \text{cost}(\text{MST}(P)) - \text{cost}(\text{MST}(P \cup S))$. We use $H(P)$ to denote the set of Hanan Steiner point candidates (i.e., the intersections of all horizontal and vertical lines passing through points of P). For a pointset P , a 1-Steiner point $x \in H(P)$ maximizes $\Delta\text{MST}(P, \{x\}) > 0$. The IIS method repeatedly finds 1-Steiner points and includes them into S . The cost of the MST over $P \cup S$ will decrease with each added point, and the construction terminates when there is no x with $\Delta\text{MST}(P \cup S, \{x\}) > 0$. Although a Steiner tree may contain at most $n - 2$ Steiner points [18], IIS may add more than $n - 2$ Steiner points; therefore, at each step we eliminate any extraneous Steiner points having degree 2 or less in the MST over $P \cup S$. Fig. 3 illustrates a sample execution of IIS, and Fig. 4 describes the algorithm formally.

Although a single 1-Steiner point may be found in $O(n^2)$ time using complicated techniques from computational geom-

²Robins and Salowe [39] investigate the maximum MST degree for higher dimensions and other L_p norms, and relate the maximum MST degree to the so-called "Hadwiger" numbers.




Fig. 3. Execution of Iterated 1-Steiner (IIS) on a 4-pin net. Note that in step (d) a degree-2 Steiner point is formed, and is eliminated from the topology.

| |
|--|
| <p>Algorithm Iterated 1-Steiner (IIS) [28] Input: A set P of n points Output: A rectilinear Steiner tree over P $S = \emptyset$ While $T = \{x \in H(P) \Delta MST(P \cup S, \{x\}) > 0\} \neq \emptyset$ Do Find $x \in T$ with maximum $\Delta MST(P \cup S, \{x\})$ $S = S \cup \{x\}$ Remove from S points with degree ≤ 2 in $MST(P \cup S)$ Output $MST(P \cup S)$</p> |
|--|

Fig. 4. The Iterated 1-Steiner algorithm.

etry [17], [28], such methods suffer from large constants in their time complexities, and they are notoriously difficult to implement. Thus, a *batched* variant of IIS is usually favored that efficiently adds an entire set of “independent” Steiner points in a single *round*, thereby affording both practicality and reduced time complexity [28], [38].

Following Hanan’s result, for each candidate Steiner point $x \in H(P)$ the Batched 1-Steiner (BIS) variant computes the induced MST savings $\Delta MST(P, \{x\})$. Next, we select a maximal “independent” set of Steiner points, where the criterion for independence is that no candidate Steiner point is allowed to reduce the potential MST cost savings of any other candidate. More formally, a set S of Steiner points is *independent* if $\Delta MST(P, S) \geq \sum_{x \in S} \Delta MST(P, \{x\})$. The weight of set S is $\Delta MST(P, S)$, and our goal is to find an independent set of maximum weight; the Steiner points in that independent set are grouped together during a round of BIS.

Using a reduction from the NP-complete problem of finding a maximum independent set in a graph, it is easy to show that our maximization problem is NP-complete also, even if the independent sets obey the “inheritance property,” one of the axioms for matroids (see Cormen *et al.* [11]). Our independent sets do not necessarily obey the inheritance property or the exchange property for matroids. Nevertheless, we can use a greedy approximation, described in Fig. 5, to approximate the weight of a maximum independent set. Note that the greedy algorithm would find a maximum independent set if the subset system was a matroid. This approximation is efficient and seems to work well in practice; it would be interesting to prove nontrivial performance bounds.

| |
|--|
| <p>Algorithm Batched 1-Steiner (BIS) [28] Input: A set P of n points Output: A rectilinear Steiner tree over P While $T = \{x \in H(P) \Delta MST(P, \{x\}) > 0\} \neq \emptyset$ Do $S = \emptyset$ For $x \in T$ in order of non-increasing ΔMST Do If $\Delta MST(P \cup S, \{x\}) \geq \Delta MST(P, \{x\})$ Then $S = S \cup \{x\}$ $P = P \cup S$ Remove from P Steiner points with degree ≤ 2 in $MST(P)$ Output $MST(P)$</p> |
|--|

Fig. 5. The Batched 1-Steiner (BIS) algorithm.

Once an approximate maximum independent set S is determined, it is inserted into P and we iterate this process with P set to $P \cup S$ until we reach a round that fails to induce a Steiner point. The total time required by BIS is $O(n^4 \log n)$ per round (the number of rounds in practice is a small constant independent of net size, i.e., less than 3 on average [28]). The BIS algorithm is summarized in Fig. 5.

III. A NEW, FASTER IMPLEMENTATION

In speeding up the MST-savings computations, a key observation is that once we have computed an MST over the pointset P , the addition of a single new point x into P can only induce a small constant number of changes between $MST(P)$ and $MST(P \cup \{x\})$. This follows from the observation that each point can have at most eight neighbors in a rectilinear planar MST, i.e. at most one per octant [24]. Thus, to update an MST with respect to a newly added point x , it suffices to consider only the closest point to x in each of the eight plane octants with respect to x (below we show that for each point it suffices to examine at most four potential candidates for connection in the MST).

Thus, we have the following linear-time algorithm for dynamic MST maintenance: connect the new point x to each of its potential neighbors (i.e. the closest point to x in each of the $O(1)$ octants around x), then delete the longest edge on any resulting cycle. This dynamic MST maintenance scheme reduces the time complexity of each round of BIS from $O(n^4 \log n)$ to $O(n^3)$, a substantial savings. An execution example of this method is given in Fig. 6, and Fig. 7 describes it formally. Note that dynamic MST maintenance can also be achieved in sub-linear time [14], but such methods seem impractical due to their complicated description and large hidden constants. A similar method was also used by Yao [45] to obtain a sub-quadratic MST algorithm in higher dimensions, but no attempt was made to optimize the number of necessary regions, whereas we also optimize the number of regions.

We now show that only four regions suffice for dynamic MST maintenance, namely the four regions defined by the two lines oriented at 45 and -45 degrees (Fig. 8(a)); we call this the *diagonal partition*. We couch the discussion in general terms so that we can later extend our terminology and techniques to the three-dimensional case. We begin by defining the following key property for regions and partitions:

The Uniqueness Property: Given a point p , a region R has the *uniqueness property* with respect to p if for every pair of




Fig. 6. Dynamic MST maintenance: adding a point to an existing MST entails connecting the point to its closest neighbor in each octant and deleting the longest edge on each resulting cycle (the Euclidean metric has been used for clarity in this example).

| Dynamic MST Maintenance (DMSTM) |
|--|
| Input: A set P of n points, $MST(P)$, a new point x |
| Output: $MST(P \cup \{x\})$ |
| $T = MST(P)$ |
| For $i = 1$ to $\#regions$ do |
| Find in region $R_i(x)$ the point $p \in P$ closest to x |
| Add to T the edge (p, x) |
| If T contains a cycle Then remove from T the longest edge on the cycle |
| Output T |

Fig. 7. Linear-time dynamic MST maintenance.

points $u, w \in R$, either $\text{dist}(w, u) \leq \text{dist}(w, p)$ or $\text{dist}(u, w) \leq \text{dist}(u, p)$.

We use $\text{dist}(u, w)$ to denote the Manhattan distance between the two points u and w . A partition of space (i.e., into a finite set of mutually disjoint regions whose union covers the space) is said to have the uniqueness property if each of its regions has the uniqueness property. Clearly, any partition scheme having the uniqueness property can be used in dynamic MST maintenance, since each region having the uniqueness property can contain at most one candidate for connection in the MST. Naturally, simpler partition schemes (i.e., ones containing a smaller number of regions) are preferable to more complicated ones. We now prove that the diagonal partition has the uniqueness property.

Lemma 3.1: *Given a point p in the Manhattan plane, each region of the diagonal partition with respect to p has the uniqueness property.*

Proof: The two diagonal lines through p partition the plane into four disjoint regions R_1 through R_4 (see Fig. 8(a)). The boundary points between two neighboring regions may be arbitrarily assigned to either region. Consider one of the four regions, say R_1 , and let $u, w \in R_1$ (Fig. 8(b)). Assume without loss of generality that $\text{dist}(u, p) \leq \text{dist}(w, p)$ (otherwise swap u and w in this proof). Consider the diamond

D in R_1 with one corner at p , and with u on the boundary of D (see Fig. 8(c)). Let c be the center of D , so that c is equidistant from all points on the boundary of D , and let a ray starting at p and passing through w intersect the boundary of D at the point b . By the triangle inequality, $\text{dist}(w, u) \leq \text{dist}(w, b) + \text{dist}(b, c) + \text{dist}(c, u) = \text{dist}(w, b) + \text{dist}(b, c) + \text{dist}(c, p) = \text{dist}(w, p)$. Thus, w is not closer to p than it is to u and the region R_1 therefore has the uniqueness property. The other three regions are handled similarly. It follows that the diagonal partition has the uniqueness property. \square

A natural question is how to determine an optimal partitioning scheme for a given dimension and metric, i.e., finding a partition scheme that contains the smallest possible number of regions, yet still possesses the uniqueness property. The existence of pointsets in the Manhattan plane where the MST is forced to have degree 4 (i.e., the five points corresponding to the center and four corners of a diamond) establishes the optimality of the diagonal partition, in the sense that no partition of the Manhattan plane into less than *four* regions can have the uniqueness property. Section V addresses the problem of finding an optimal partition for three-dimensional Manhattan space.

We have shown above that in the Manhattan plane, the degree of any particular single MST node can be made to be 4 or less. But note that this does not imply that the degrees of *all* nodes can be made 4 or less *simultaneously*, since decreasing the degree of one node can increase the degree of neighboring nodes. Thus, it is not immediately obvious that in the Manhattan plane there always exists an MST with maximum degree 4; this requires additional proof as detailed in Section VI below.

IV. A VARIANT WITH IMPROVED PERFORMANCE

At each iteration, the basic IIS heuristic uses pure greed to select a 1-Steiner point, and this may unfortunately preclude additional savings in subsequent iterations. A similar phenomenon may occur due to tie-breaking among 1-Steiner candidates that induce equal savings. For example, in Fig. 9 we observe that an unfortunate choice for a 1-Steiner point can interfere with the savings of future potential 1-Steiner candidates, resulting in a suboptimal solution.

Empirical tests indicate that ties in MST savings for the various 1-Steiner point candidates occur very often. Therefore, in order to avoid breaking ties in ways that would preclude possible future savings, we propose the following scheme: when an MST savings tie occurs among a number of 1-Steiner candidates, rather than using a deterministic tie-breaking rule, we instead randomly select one of the 1-Steiner candidates and proceed with the execution. We then run this randomized variant of IIS m times on the same input and select the best solution (i.e., the least costly of the m trees produced), where m is an input parameter.

In order to further avoid the pitfalls of a purely greedy strategy (i.e., getting trapped in local minima), we also propose a mechanism that allows IIS to select a 1-Steiner candidate if its MST savings is within δ units from that of the best candidate, where δ is again an input parameter. This strategy




Fig. 8. The *diagonal partition* of the plane into four regions with respect to a point p (a) has the *uniqueness property*: for every two points u and w that lie in the same region (b), either $\text{dist}(w, u) \leq \text{dist}(w, p)$ or else $\text{dist}(u, w) \leq \text{dist}(u, p)$ (c).




Fig. 9. An unlucky tie-breaking choice for a 1-Steiner point may interfere with the savings of other potential 1-Steiner candidates (a). If P_1 is selected in the first iteration (b), then the MST savings of both P_2 and P_3 vanish during the second iteration, yielding a suboptimal tree of cost 7 (c); on the other hand, if P_2 is selected in the first iteration, then P_3 may be selected in the second iteration, yielding an optimal tree of cost 6 (d).

| |
|---|
| Algorithm Enhanced Iterated k-Steiner (EIkS) |
| Input: A set P of n points, parameters $\delta \geq 0$, $k \geq 1$, and $m \geq 1$ |
| Output: A rectilinear Steiner tree over P . |
| $T = \text{MST}(P)$ |
| Do m times |
| $S = \emptyset$ |
| While $C = \{X \subseteq H(P) \mid X \leq k, \Delta\text{MST}(P \cup S, X) > 0\} \neq \emptyset$ Do |
| Find $Y \in C$ with maximum $\Delta\text{MST}(P \cup S, Y)$ |
| Randomly select $Z \in C$ with $\Delta\text{MST}(P \cup S, Z) \geq \Delta\text{MST}(P \cup S, Y) - \delta$ |
| $S = S \cup Z$ |
| Remove from S points with degree ≤ 2 in $\text{MST}(P \cup S)$ |
| If $\text{cost}(\text{MST}(P \cup S)) < \text{cost}(T)$ Then $T = \text{MST}(P \cup S)$ |
| Output T |

Fig. 10. The Enhanced Iterated k -Steiner (EIkS) method.

would enable the acceptance of a slightly suboptimal choice (with respect to the best immediate possible savings), with the possibility of realizing greater savings in future iterations.

Finally, we note that performance may be further improved if instead of looking for individual 1-Steiner points, we search for *pairs* of Steiner candidates that offer maximum savings with respect to other candidates or pairs of candidates. For example, such an *Iterated 2-Steiner algorithm* (I2S) will optimally solve the pointset of Fig. 9. Combining these three techniques of 1) nondeterministic tie-breaking, 2) near-greedy search, and 3) k -Steiner selection, we obtain a new *Enhanced Iterated k -Steiner* (EIkS) algorithm, as shown in Fig. 10.

Note that the original IIS algorithm of Kahng and Robins [28] (see Fig. 4) is equivalent to our new EIkS algorithm with $k = 1, m = 1$, and $\delta = 0$. Our EIkS scheme can also be extended using a “noninterfering” criterion as in [28], to

yield an *enhanced batched k -Steiner* (EBkS) algorithm, where a maximal number of Steiner points are added during each round. The time complexity of EBkS is $O(m \cdot n^{2(k-1)} \cdot T(n))$, where $T(n)$ is the time complexity of B1S. For fixed m , EB1S runs asymptotically within the same time as B1S, namely $O(n^3)$ per round. The EBkS method improves the quality of the solutions at the expense of running time, allowing a smooth tradeoff between performance and efficiency. Although EBkS is guaranteed to always yield optimal solutions for $\leq k + 2$ pins, its time complexity increases exponentially with k ; thus, in order to remain within polynomial time, k must be fixed.

Finally, it was observed empirically that only a small fraction of the Hanan candidates have positive MST savings in a given round of B1S; moreover, only candidates with positive MST savings in an earlier round are likely to produce positive MST savings in subsequent rounds. Therefore, rather than examine the MST savings of all Hanan candidates in a given round, we may consider only the candidates that produced positive savings in the previous round. The empirical simulations described in Section VIII indicate that this strategy significantly reduces the time spent during a round, without degrading solution quality. We call this streamlined version of B1S the *modified batched 1-Steiner* (MB1S) algorithm.

V. STEINER ROUTING IN THREE DIMENSIONS

Three-dimensional packaging is emerging as a viable VLSI design technology [9], [20], [22]; however, most existing CAD routing tools and techniques still implicitly address two dimensions only. In contrast, the EIkS method readily generalizes to arbitrary dimensions. We distinguish between the general three-dimensional version of the Steiner problem and the less general (but more realistic) multi-layer formulation, where the points of P lie on L parallel planes. Note that the unrestricted three-dimensional version of the Steiner problem occurs in the limit when $L = \infty$, and the standard two-dimensional formulation is the case $L = 1$. The cost of routing between one layer and another (i.e., using vias) is likely to be substantially higher than staying on the same single layer (i.e., in terms of manufacturing expense, signal propagation delay, etc.), and this may be modeled by varying the distance between the layers. In Section VIII, we present simulation data for several combinations of values for L and n .

Our three-dimensional ElKS method may be implemented efficiently by using Snyder's [42] generalization of Hanan's theorem to higher dimensions. In particular, there always exists an optimal Steiner tree whose Steiner points are chosen from the $O(n^3)$ intersections of all orthogonal planes (i.e., planes parallel to the coordinate axis) passing through all points in P . The three-dimensional analog of Hwang's result suggests that the maximum MST/MRST ratio for three dimensions is at most $\frac{5}{3}$ (this is a consequence of a more general conjecture for higher dimensions [19]), although there is currently no known proof of this. An example consisting of six points located in the middles of the faces of a rectilinear cube establishes that $\frac{5}{3}$ is a lower bound for the worst MST/MRST performance ratio in three dimensions. Thus, we expect the average performance of our heuristics, expressed as percent improvement over MST, to be higher in three dimensions than it is in two dimensions; this is indeed confirmed by our experimental results in Section VIII. Also as expected, the average performance improves as the number of layers L increases.

As noted above, once we have computed an MST over a pointset P , the addition of a single new point p into P can only induce a small constant number of topological changes between $\text{MST}(P)$ and $\text{MST}(P \cup \{p\})$. This follows from the fact that in a fixed dimension, each point can have at most $O(1)$ neighbors in a rectilinear MST, as noted above in Section III. Thus, MST savings in three dimensions may be efficiently calculated by partitioning the space with respect to the new point p into $O(1)$ mutually disjoint regions $R_i(p)$ having the uniqueness property, namely that only the closest point to p in each region $R_i(p)$ may be connected to p in the $\text{MST}(P)$. This implies that in three-dimensional Manhattan space, linear time suffices to compute the MST savings of each 1-Steiner candidate.

Using similar arguments to those of Lemma 3.1, we can partition three-dimensional Manhattan space into 14 regions, with the uniqueness property holding for each region. Such a partition corresponds to the faces of a truncated cube (Fig. 11(a)), i.e., a solid obtained by chopping off the corners of a cube, yielding six square faces and eight equilateral triangle faces (Fig. 11(b)); this solid is known as a "cuboctahedron" [34]. The 14 solid regions of this partition are induced by the 14 faces of the cuboctahedron, namely the six pyramids with square cross-section (Fig. 11(c)) and the eight pyramids with triangular cross-section (Fig. 11(d)). Again, points located on region boundaries may be arbitrarily assigned to any of the adjacent regions. We call this particular partition of space the *cuboctahedral partition*, and we refer to the two types of induced regions as *square pyramids* and *triangular pyramids*, respectively. We now show that the cuboctahedral partition has the uniqueness property.

Theorem 5.1: *Given a point p in three-dimensional space under the Manhattan metric, each of the 14 regions of the cuboctahedral partition of space with respect to p has the uniqueness property.*

Proof: To show the uniqueness property for the square pyramids, consider one of the square pyramids R with respect to p (Fig. 11(c)), and let $u, w \in R$. Assume without loss of generality that $\text{dist}(u, p) \leq \text{dist}(w, p)$ (otherwise swap the

roles of u and w in this proof). Consider the locus of points $D \subset R$ that are distance $\text{dist}(u, p)$ from p . (Fig. 11(e)); D is the upper half of the boundary of an octahedron. Let c be the center of the octahedron determined by D , so that c is equidistant from all points of D . Let b be the intersection of the surface of D with a ray starting from p and passing through w . By the triangle inequality, $\text{dist}(w, u) \leq \text{dist}(w, b) + \text{dist}(b, c) + \text{dist}(c, u) = \text{dist}(w, b) + \text{dist}(b, c) + \text{dist}(c, p) = \text{dist}(w, p)$. Thus, w is not closer to p than it is to u , and therefore the region R has the uniqueness property. The other square pyramids are handled similarly.

To show the uniqueness property for the triangular pyramids, consider one of the triangular pyramids R with respect to p (Fig. 11(d)), and let $u, w \in R$. Assume without loss of generality that $\text{dist}(u, p) \leq \text{dist}(w, p)$ (otherwise swap the roles of u and w). Consider the locus of points D in R that are distance $\text{dist}(u, p)$ from p (Fig. 11(f)). Let b be the intersection of D with a ray starting from p and passing through w . By the triangle inequality, $\text{dist}(w, u) \leq \text{dist}(w, b) + \text{dist}(b, u) \leq \text{dist}(w, b) + \text{dist}(b, p) = \text{dist}(w, p)$. Thus, w is not closer to p than it is to u and therefore the region R has the uniqueness property. The other triangular pyramids are handled similarly. \square

We thus have the following:

Corollary 5.2: *Given a pointset P in three-dimensional Manhattan space and an additional new point p , there exists an MST over $P \cup \{p\}$ where p has degree of at most 14 in the MST.*

Proof: The cuboctahedral partition of space with respect to p yields 14 regions, each possessing the uniqueness property. This implies that for any two points inside a region, one is closer to the other rather than to the origin; thus, only one point inside each region is a viable candidate for an MST neighbor of p . Therefore, the degree of p (or any other single point) in the MST can be made 14 or less. \square

It is still an open question whether for three dimensions the cuboctahedral partition is optimal (i.e., whether the cuboctahedral partition has a minimum number of regions among all partitions having the uniqueness property); we conjecture that it is. On the other hand, we can show that 13 is a lower bound on the maximum MST degree in three-dimensional Manhattan Space:

Theorem 5.3: *There are three-dimensional pointsets for which the maximum degree of any MST is at least 13.*

Proof: Consider the pointset $P = \{(0, 0, 0), (\pm 100, 0, 0), (0, \pm 100, 0), (0, 0, \pm 100), (47, -4, 49), (-6, -49, 45), (-49, 8, 43), (-4, 47, -49), (-49, -6, -45), (8, -49, -43), (49, 49, 2)\}$.

The distance between every point and the origin is exactly 100 units, but the distance between any two non-origin points is strictly greater than 100 units. Therefore, the MST over P is unique with a star topology (i.e., all 13 points must connect to the origin in the MST), and thus the origin point has degree 13 in the MST. \square

A remaining open question is whether there exists a three-dimensional pointset where the maximum MST degree is forced to be as high as 14. Given that each point can connect to at most 14 neighbors in the MST, we obtain the following linear-time algorithm for dynamic MST maintenance




Fig. 11. A truncated cube (a)–(b) induces a three-dimensional cuboctahedral space partition where each region has the uniqueness property. The 14 induced regions consist of six square pyramids (c) and eight triangular pyramids (d). Using the triangle inequality, the uniqueness property may be shown to hold for each region (e)–(f).

(DMSTM) in three dimensions: connect the new point in turn to each of its ≤ 14 potential neighbors, then delete the longest edge on each resulting cycle. This is a three-dimensional analog of the two-dimensional method given in Figs. 6 and 7.

VI. ON THE MAXIMUM MST DEGREE

As noted above, Lemma 3.1 does not imply that in the Manhattan plane the maximum MST degree is 4, nor does Corollary 5.2 imply that the maximum MST degree in three

| Algorithm Parallel Iterated 1-Steiner (PIIS) | |
|---|---|
| Input: | A set P of n points (assuming p processors available) |
| Output: | A rectilinear Steiner tree over P |
| Master Process: | |
| $S = \emptyset$ | |
| Do Forever | |
| For $i = 1$ to p Do | |
| Send the i^{th} processor $P \cup S$ and $\lfloor \frac{ H(P) }{p} \rfloor$ different Hanan candidates from $H(P)$ | |
| For $i = 1$ to p Do | |
| Receive from the i^{th} processor the candidate h_i with the largest ΔMST | |
| Let x be the h_i with the largest $\Delta\text{MST}(P \cup S, \{h_i\})$ for all $1 \leq i \leq p$ | |
| If $\Delta\text{MST}(P \cup S, \{x\}) > 0$ Then | |
| $S = S \cup \{x\}$ | |
| Remove from S points with degree ≤ 2 in $\text{MST}(P \cup S)$ | |
| Else Output $\text{MST}(P \cup S)$ and Stop | |
| Slave Process: | |
| Receive from the master process $P \cup S$ and a set of Hanan candidates h_1, \dots, h_k | |
| Return the candidate h_i with the largest $\Delta\text{MST}(P \cup S, \{h_i\})$ | |

Fig. 12. The Parallel Iterated 1-Steiner algorithm. The master process sends each one of p slave processors $\frac{1}{p}$ of the Hanan candidates; a slave process then determines the candidate with the best ΔMST and returns it to the master process. Out of all the p candidates received, the master process then uses the candidate with the highest MST savings for inclusion into the pointset. To further increase the parallel execution speed, slower processes are sent smaller sets of candidates, in proportion to their observed speed (for clarity, this simple load-balancing scheme is not shown in the above template).

dimensions is 14, since reducing the MST degree of any one point may increase the MST degree of other points. It turns out, however, that ties for connection during MST construction may always be broken appropriately so as to keep the overall MST degree low. We begin by defining a strict version of the uniqueness property:

The Strict Uniqueness Property: Given a point p , a space region R has the *strict uniqueness property* with respect to p if for every pair of points $u, w \in R$, either $\text{dist}(w, u) < \text{dist}(w, p)$ or $\text{dist}(u, w) < \text{dist}(u, p)$.

Note that if a region has the strict uniqueness property it also has the (nonstrict) uniqueness property. Clearly each d -dimensional region satisfying the strict uniqueness property may contribute at most 1 to the maximum MST degree. We now prove that by breaking ties judiciously, the maximum MST degree can be made no larger than the number of d -dimensional regions in a partition having the strict uniqueness property:

Theorem 6.1: Given a partition of d -dimensional space into r regions, and given that only $r' \leq r$ of these regions are d -dimensional and have the strict uniqueness property (the rest of the $r - r'$ regions being lower-dimensional, and are not required to have the uniqueness property), then the maximum MST degree in this space is r' or less.

Proof: Given a pointset P , perturb the coordinates of each point by a tiny amount so that the lower-dimensional regions with respect to each point do not contain any other points. This is always possible to do, and yields a new perturbed pointset P' . Clearly, interpoint distances in P' only differ by a tiny amount from the corresponding interpoint distances in P , and therefore the cost of the MST's over P' and P can differ by only a similarly tiny amount, which can be made arbitrarily small. But the MST over P' has maximum degree r' , since only the r' d -dimensional regions of the partition are nonempty with respect to the points of P' . We now use the topology of the MST for P' to connect the corresponding points of P ; this would induce an MST over P having maximum degree r' . \square

Applications of this technique to 2 and 3 dimensions are immediate:

Corollary 6.2: Every pointset in the Manhattan plane has an MST with maximum degree 4.

Proof: Modify the diagonal partition into a *strict* diagonal partition having a total of eight regions: four two-dimensional open wedges (i.e., not containing any of their own boundary points), and four one-dimensional rays (i.e., the boundaries between the wedges). By arguments similar to those of Lemma 3.1, each of the open wedges possesses the strict uniqueness property, and thus by Theorem 6.1 points lying on the boundaries between wedges can be perturbed into the interiors of the wedges themselves. Thus, the maximum MST degree given such a strict diagonal partitioning scheme is 4. \square

The bound of 4 on the maximum MST degree in the Manhattan plane is tight, since there are examples that achieve this bound (e.g., the center and vertices of a diamond). Note that the best previously known upper bound for the maximum MST degree in the Manhattan plane was 6, as was stated without proof in [24].

Corollary 6.3: Every pointset in three-dimensional Manhattan space has an MST with maximum degree 14.

Proof: Modify the cuboctahedral partition into a *strict* cuboctahedral partition having a total of 38 regions: 14 three-dimensional open pyramids (i.e., 8 triangular pyramids and 6 square pyramids, each not containing any of their own boundary points), and 24 two-dimensional regions (i.e., corresponding to all the boundaries between the pyramids). By arguments identical to those of Theorem 5.1, each of the open pyramids possesses the strict uniqueness property, and thus by Theorem 6.1, points lying on the boundaries between the 14 pyramids can be perturbed into the interiors of the pyramids themselves. Thus, the maximum MST degree given such a strict cuboctahedral partition scheme is equal to the number of three-dimensional regions of the cuboctahedral partition, i.e. 14. \square

While Theorem 5.3 gives an example that illustrates that the maximum MST degree in three-dimensional Manhattan space can be as large as 13, it is still open whether there exist examples in three-dimensional Manhattan space that force an MST degree of 14. Note that the best previously known bound for the maximum MST degree in three-dimensional Manhattan space was 26, as implied by a result from the theory of sphere packing [12], [44].

Our results regarding MST bounds also settle some open questions in complexity theory, since it is known that the problem of finding a degree-bounded MST is NP-complete, even when the degree bound is fixed at 2 (yielding the Traveling Salesman Problem) or at 3, as shown by Papadimitriou and Vazirani [36]. Corollary 6.2 implies that the degree-bounded MST problem in the Manhattan plane is solvable in polynomial time when the degree bound is fixed at 4 or more, since we have shown how to efficiently find an MST that meets these maximum degree constraints; this was previously an open problem. Similarly, Corollary 6.3 implies that the degree-bounded MST problem in three-dimensional Manhattan space is solvable in polynomial time when the degree bound is fixed




Fig. 13. An example of the output of Batched 1-Steiner (B1S) on a random 300-point set (hollow dots). The Steiner points produced by B1S are denoted by dark solid dots.

at 14 or more (since we have shown how to efficiently find an ordinary MST that meets such maximum degree constraints). Monma and Suri [35] used a similar perturbation argument to prove that for any pointset in the Euclidean plane, there is an MST with maximum degree of 5.

VII. A PARALLEL IMPLEMENTATION

Since a typical CAD environment consists of a network of workstations, exploiting the available large-grain parallelism provides a natural and effective means of improving the performance of CAD algorithms. We therefore propose a parallel implementation of the Iterated 1-Steiner method that achieves high parallel speedups. The I1S algorithm is highly parallelizable because each one of p processors can compute independently and in parallel the MST savings of $O(\frac{n^2}{p})$ of the Steiner candidates. In our parallel implementation, all processors send their best candidate to a master processor, which selects the best of these candidates for inclusion into the pointset. This procedure is iterated until no improving candidates can be found (the B1S algorithm parallelizes similarly).

We used the Parallel Virtual Machine (PVM) system [5], [43] to control remote processors. PVM is a widely available software package that allows a heterogeneous network of parallel and serial computers to be used as a single computational resource. The PVM system consists of two parts, a daemon process, and a user library, providing mechanisms for initiating processes on other machines and for controlling synchronization and communication among processes. Using PVM as a framework for parallelism alleviates the need to hand-code the synchronization and communication protocols from scratch; for the specific details on how PVM manages the underlying system resources, see [5], [43].

Each of the processors in our parallel implementation is a SUN workstation, communicating over an Ethernet. The master processor sends to the p available processors equal-sized subsets of the Hanan candidate set $H(P)$; each slave process then determines from among the candidates that it received the one with the best Δ MST and returns the winner to the master process. Out of all the p candidates received, the master process then determines the candidate with the highest MST savings and includes it into the pointset. This parallel version of the Iterated 1-Steiner algorithm is shown in Fig. 12.

To further increase the parallel execution speed, we implemented a load-balancing scheme in order to mitigate any significant variations that may exist between the speeds of different processors (due to a heavy CPU load, the underlying architecture type, swapping behavior, etc.) To this end, the wall-clock response time of each processor is tracked. If any individual processor is determined to be considerably slower than the rest, it is henceforth given smaller tasks to perform (i.e., it is sent less Hanan candidates, in an amount proportional to its observed speed). If a processor does not complete a task within reasonable time, it is sent an abort message and its computation task is reassigned altogether to the fastest idle processor available. This prevents individual slow (or crashed) processors from seriously impeding the speed of the overall computation.

VIII. EXPERIMENTAL RESULTS

We have implemented both serial and parallel enhanced versions of the B1S algorithm, using C in the SUN workstation environment. Our code is available upon request. The serial B1S heuristic has been benchmarked on up to 10000 instances of each net size. The instances were generated randomly by independently choosing the coordinates of each point from a uniform distribution in a 10000×10000 grid; such instances are statistically similar to the pin locations of actual VLSI nets and are the standard testbed for Steiner tree heuristics [27]. As is the convention in the Steiner approximation literature [27], we evaluate the performance of our method by comparing the cost of our solutions to the MST cost over the same inputs. Performance results are summarized in Table I and are illustrated in Fig. 14(a)–14(c). B1S yields Steiner trees with cost averaging almost 11% less than the MST cost.³

We timed the execution of the serial and parallel versions of B1S, using both the naive $O(n^4 \log n)$ implementation [28] and our new $O(n^3)$ implementation, which incorporates the efficient, dynamic MST maintenance as described in Section III. We observe that the number of rounds required by B1S is always very small, and may therefore be considered a constant; for example, when $n = 300$, the average number of rounds is only about 2.5, and we never observed B1S to require more than 5 rounds.

The parallel implementation uses nine SUN 4/40 (IPC) workstations, with a SUN 4/75 (Sparc2) as the master processor. The improvements in running time are dramatic: for $n = 100$, our fast serial implementation is 247 times faster than the naive serial implementation (as measured by CPU time), and

³Recently, other Steiner heuristics with performance approaching that of I1S were proposed by Borah, Owens, and Irwin [8], Chao and Hsu [10], and by Lewis, Pong, and VanCleave [33].

TABLE I
 BATCHED 1-STEINER (B1S) STATISTICS: THE PERFORMANCE FIGURES DENOTE PERCENT IMPROVEMENT OVER MST COST. ALSO GIVEN ARE STATISTICS REGARDING THE NUMBER OF STEINER POINTS PRODUCED, THE NUMBER OF ROUNDS, AND THE NUMBER OF STEINER POINTS PRODUCED IN EACH ROUND

| n | Batched 1-Steiner (B1S) Performance Results | | | | | | | | | | | | |
|-----|---|-------|-------|---------------|--------|-----|------------------|------|-----|--------------------|-------|------|------|
| | Performance | | | Number of SPs | | | Number of rounds | | | Ave #SPs per round | | | |
| | Min | Ave | Max | Min | Ave | Max | Min | Ave | Max | 1 | 2 | 3 | 4 |
| 4 | 0.00 | 8.83 | 27.19 | 0 | 2.10 | 4 | 0 | 0.97 | 2 | 1.14 | 0.37 | 0.00 | 0.00 |
| 5 | 0.00 | 9.35 | 23.59 | 0 | 2.59 | 5 | 0 | 1.08 | 3 | 1.51 | 0.46 | 0.06 | 0.00 |
| 6 | 0.00 | 9.83 | 25.99 | 0 | 3.09 | 6 | 0 | 1.15 | 3 | 1.91 | 1.42 | 0.03 | 0.00 |
| 7 | 0.41 | 9.82 | 23.44 | 2 | 3.52 | 7 | 1 | 1.17 | 3 | 2.32 | 0.59 | 0.02 | 0.00 |
| 8 | 0.00 | 9.92 | 20.93 | 0 | 3.99 | 8 | 0 | 1.23 | 3 | 2.73 | 0.96 | 0.17 | 0.00 |
| 9 | 1.04 | 10.07 | 20.62 | 2 | 4.48 | 7 | 1 | 1.31 | 3 | 3.12 | 1.16 | 0.19 | 0.00 |
| 10 | 1.26 | 10.24 | 20.03 | 3 | 4.95 | 8 | 1 | 1.34 | 4 | 3.56 | 1.11 | 0.10 | 0.01 |
| 11 | 2.01 | 10.32 | 20.81 | 3 | 5.42 | 10 | 1 | 1.33 | 3 | 4.01 | 1.01 | 0.03 | 0.00 |
| 12 | 1.85 | 10.32 | 18.83 | 2 | 5.88 | 10 | 1 | 1.43 | 4 | 4.31 | 1.28 | 0.09 | 0.00 |
| 13 | 2.06 | 10.27 | 19.74 | 3 | 6.39 | 11 | 1 | 1.46 | 4 | 4.80 | 1.23 | 0.09 | 0.04 |
| 14 | 3.01 | 10.28 | 19.56 | 3 | 6.82 | 11 | 1 | 1.48 | 4 | 5.18 | 1.25 | 0.14 | 0.07 |
| 15 | 2.78 | 10.41 | 18.10 | 4 | 7.27 | 12 | 1 | 1.49 | 4 | 5.59 | 1.33 | 0.12 | 0.00 |
| 16 | 2.47 | 10.42 | 18.45 | 4 | 7.76 | 13 | 1 | 1.52 | 4 | 6.01 | 1.33 | 0.13 | 0.00 |
| 17 | 3.00 | 10.58 | 17.94 | 5 | 8.33 | 14 | 1 | 1.56 | 3 | 6.50 | 1.38 | 0.27 | 0.00 |
| 18 | 3.62 | 10.46 | 18.19 | 5 | 8.70 | 14 | 1 | 1.58 | 4 | 6.85 | 1.52 | 0.26 | 0.01 |
| 19 | 3.71 | 10.37 | 18.05 | 5 | 9.14 | 14 | 1 | 1.61 | 5 | 7.21 | 1.58 | 0.10 | 0.01 |
| 20 | 4.35 | 10.39 | 17.90 | 5 | 9.63 | 15 | 1 | 1.60 | 4 | 7.68 | 1.58 | 1.00 | 0.00 |
| 30 | 6.49 | 10.64 | 14.73 | 9 | 13.47 | 20 | 1 | 1.85 | 4 | 11.61 | 1.75 | 0.21 | 0.01 |
| 40 | 6.28 | 10.68 | 14.55 | 12 | 18.10 | 26 | 1 | 1.98 | 5 | 15.95 | 2.47 | 0.23 | 0.00 |
| 50 | 7.25 | 10.72 | 15.28 | 15 | 23.70 | 33 | 1 | 2.04 | 4 | 19.97 | 2.78 | 0.56 | 0.00 |
| 60 | 8.67 | 10.77 | 14.29 | 19 | 27.53 | 37 | 1 | 2.12 | 5 | 23.88 | 3.81 | 0.58 | 0.00 |
| 70 | 7.80 | 10.82 | 13.55 | 23 | 32.26 | 42 | 1 | 2.17 | 4 | 28.02 | 3.52 | 0.62 | 0.00 |
| 80 | 7.34 | 10.83 | 12.88 | 28 | 36.99 | 50 | 1 | 2.20 | 4 | 32.07 | 4.66 | 0.64 | 0.01 |
| 90 | 7.99 | 10.84 | 13.52 | 32 | 41.61 | 56 | 1 | 2.24 | 4 | 36.18 | 5.14 | 1.17 | 0.04 |
| 100 | 8.78 | 10.86 | 13.60 | 35 | 46.30 | 59 | 1 | 2.25 | 4 | 40.21 | 5.55 | 1.31 | 0.01 |
| 110 | 8.04 | 10.86 | 12.86 | 40 | 50.98 | 67 | 1 | 2.28 | 4 | 44.96 | 5.77 | 0.53 | 0.08 |
| 120 | 7.83 | 10.86 | 12.62 | 44 | 55.59 | 68 | 1 | 2.29 | 5 | 48.64 | 7.15 | 1.02 | 0.00 |
| 130 | 8.25 | 10.85 | 12.63 | 47 | 60.03 | 73 | 1 | 2.34 | 4 | 52.87 | 7.22 | 0.82 | 0.00 |
| 140 | 8.45 | 10.84 | 12.57 | 51 | 64.97 | 79 | 1 | 2.38 | 5 | 56.94 | 8.12 | 1.07 | 0.22 |
| 150 | 9.07 | 10.92 | 12.63 | 57 | 69.85 | 87 | 1 | 2.39 | 4 | 59.92 | 8.28 | 1.13 | 0.00 |
| 200 | 9.39 | 10.97 | 12.32 | 79 | 93.10 | 107 | 2 | 2.47 | 4 | 80.04 | 11.68 | 1.48 | 0.50 |
| 250 | 10.26 | 10.98 | 11.68 | 103 | 116.32 | 131 | 2 | 2.57 | 4 | 100.00 | 14.53 | 0.82 | 0.00 |
| 300 | 9.76 | 10.97 | 12.18 | 131 | 137.67 | 145 | 2 | 2.50 | 4 | 121.85 | 16.85 | 1.33 | 0.00 |

our parallel implementation running on 10 processors is 1163 times faster than the naive implementation (as measured by elapsed wall-clock time). It should be noted that for the serial implementation, the CPU time and the elapsed wall-clock time are essentially identical, since we report averages over hundreds of cases, and all of our algorithms are CPU-bound (i.e., swapping and I/O time is negligible). Detailed execution timings for various cardinalities are given in Table II, and are illustrated in Fig. 14(d). Note that most of the observed speedup over the naive implementation stems from the fast serial implementation using the dynamic MST update scheme (i.e., as described in Section III). The parallel implementation only adds a speedup factor somewhat less than the number of processors/workstations available (10 SUN's in our case). The serial speedup grows as a function of the net size, since our new serial implementation is asymptotically faster than the naive serial implementation.

As noted by Ganley and Cohoon [15], most nets in actual VLSI designs have a small number of pins (i.e., less than 10). It is therefore of particular interest to observe the behavior of our algorithms on small nets. Even for small pointsets, our new implementations are considerably faster than the previous, naive ones; for example, for $n = 5$, the new serial B1S is on average twice as fast as the naive implementation, while for

$n = 10$ it is 7 times as fast. Our parallel speedup increases with problem size, reaching about 7.2 for $n = 250$ running on 10 processors. This enables us to examine the asymptotic behavior of B1S for much larger pointsets than was previously possible: the output of Batched 1-Steiner (B1S) for a random pointset of size 300 is shown in Fig. 13.

We have also implemented the E1KS and the EBkS algorithms, and the following algorithms were tested side-by-side on the same inputs:

- **B1S**—The Batched 1-Steiner method, with the code further streamlined for speed;
- **MB1S**—The “modified” B1S variant that uses only winning candidates from previous rounds;
- **EB1S**—The enhanced version of B1S;
- **I2S**—The Iterated 2-Steiner algorithm;
- **EI2S**—The enhanced version of I2S;
- **META(B1S,MB1S)**—The *metaheuristic* over heuristics B1S and MB1S, i.e., the best solution found by either of these heuristics;
- **META(EB1S,I2S,EI2S)**—The metaheuristic over EB1S, I2S, and EI2S;
- **META(B1S,MB1S,EB1S,I2S,EI2S)**—The metaheuristic over B1S, MB1S, EB1S, I2S, and EI2S;
- **OPT**—The optimal Steiner tree algorithm [40].

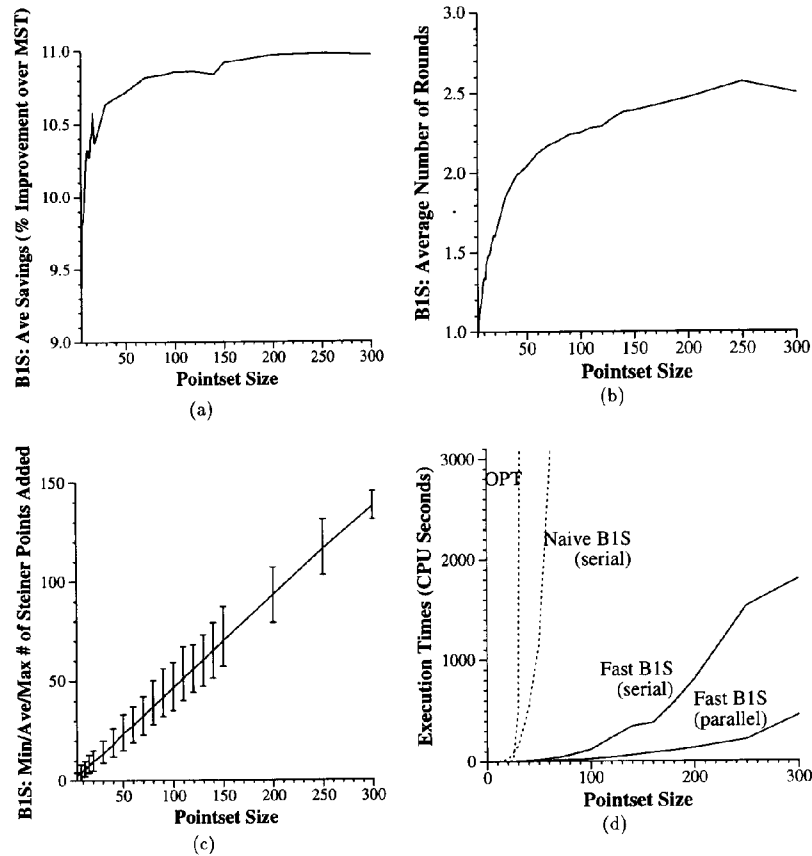


Fig. 14. (a) Average performance of B1S, shown as percent improvement over MST cost. (b) Average number of rounds for B1S. (c) Average number of Steiner points induced by B1S (vertical bars indicate the range of the minimum and maximum number of Steiner points added). (d) Average execution times (in elapsed wall-clock seconds) for the serial and parallel B1S, using both the naive MST implementation and our new incremental MST maintenance scheme, over a wide range of net sizes; note that for the serial versions, the elapsed wall-clock times are essentially the same as the CPU time, since the benchmarks represent average running times over hundreds of cases (the algorithms are CPU-bound, so that swapping and I/O have negligible effect on the running time).

Recall from the discussion at the end of Section IV that MB1S is a more efficient version of B1S, since it only examines a fraction of the Hanan candidates in a typical round (i.e., only the ones with positive MST savings in the previous round). OPT is the fastest known *optimal* rectilinear Steiner tree algorithm [40]. All variants have been benchmarked on up to 10000 random instances of each net size. Fig. 15(a) shows the performance comparison of MB1S, EI2S, and OPT, while Table III gives more detailed performance data. We observe that the average performance of our methods is extremely close to optimal: for $n = 8$, EI2S is on average only about 0.11% away from optimal and solutions are optimal in about 90% of the cases. Even for $n = 30$, MB1S is only about 0.30% away from optimal, and yields optimal solutions in about one quarter of all cases. Table IV tracks the percentage of cases where the various heuristics find the optimal solution, and this data is also depicted in Fig. 15(b).

Table V gives the average CPU times, in seconds, for each heuristic and net size. Our most time-efficient algorithm is MB1S, requiring an average of 0.009 CPU seconds per 8-pin net, and an average of 0.375 seconds per 30-pin net. Using

EIkS (or EBkS) with values of k greater than 2 improves the performance, but slows down the algorithm; it is easy to see that for arbitrary k , EIkS (EBkS) always yields optimal solutions for $\leq k + 2$ pins, but has time complexity greater by a factor of $n^{2(k-1)}$ than that of EB1S. This allows a smooth tradeoff between performance and efficiency. However, the performance of the EBkS algorithm with $k = 2$ is already so close to optimal, that in most applications increasing k further is not likely to justify the incurred time penalty.

In three dimensions, we observed that the average performance of EB1S approaches 15% improvement over MST cost and the performance increases with the number of planes L . It is not surprising that the average savings over MST cost in three dimensions is higher than it is in two dimensions, since the worst-case performance ratio in three dimensions is higher also (i.e., $\frac{5}{3}$ for three dimensions versus $\frac{3}{2}$ for two dimensions). Fig. 15(c) shows the performance of our method in three dimensions for various values of the number of parallel planes L , including the unrestricted three-dimensional case, corresponding to the limit when L approaches ∞ . Table VII gives statistics for three dimensions on the number of Steiner

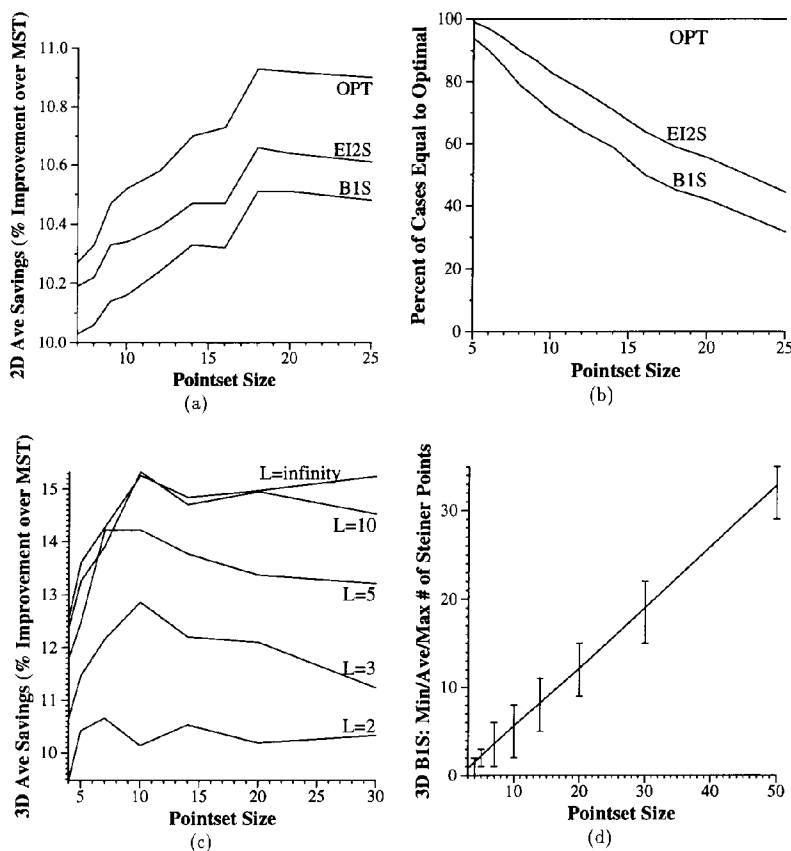


Fig. 15. (a) Average performance in two dimensions of EI2S, B1S, and OPT; note that EI2S is only 0.25% (or less) away from optimal. (b) Percentage of all cases when the heuristics find the optimal solution (note that EI2S yields optimal solutions a large percentage of the time). (c) Average performance of EB1S in three dimensions for various values of L = number of parallel planes. (d) Average number of Steiner points added by B1S in three dimensions for $L = \infty$.

points induced by B1S (see Fig. 15(d)), as well as on the number of rounds that occur in B1S before termination. As is the case in two dimensions, the number of rounds for B1S in three dimensions is on average very small. Table VI gives more detailed performance data. In all cases, the L parallel planes were uniformly spaced in the unit cube (i.e., they were separated by $\frac{G}{L}$ units, where $G = 10000$ is the gridsize). Unfortunately, the OPT algorithm of Salowe and Warme [40] does not easily generalize to higher dimensions; thus, we were not able to compare our three-dimensional version of EB1S to optimal.

IX. CONCLUSION

We have proposed enhanced serial and parallel implementations of the Batched 1-Steiner heuristic (B1S), achieving speeds of up to three orders of magnitude faster than previous implementations. Moreover, the speedup increases with the number of points. This has enabled the testing of B1S on several hundred points for the first time, and we observed that for such large pointsets B1S consistently improves 11% over MST cost.

Next, we enhanced B1S by using a near-greedy approach with random tie-breaking. Our method enjoys the same asymptotic time complexity as B1S, yet offers improved average performance. We also allow performance increase at the expense of running time, creating a smooth tradeoff between solution quality and computational efficiency. Extensive simulations indicate that for typical nets, the average performance of our methods is less than 0.25% away from optimal, and our solutions are actually optimal for up to 90% of uniformly distributed nets of typical sizes.

We generalized B1S and its variants to three dimensions, as well as to the case where all the pins lie on L parallel planes, which arises in, e.g., three-dimensional VLSI and multi-layer routing. Our methods are highly parallelizable and generalize to arbitrary weighted graphs; thus, they are suitable to support a multi-layer global router, where obstruction and congestion considerations affect routing. Since Steiner tree construction is a computationally expensive component of global routing, our techniques suggest the feasibility of a "Steiner engine" for efficiently computing near-optimal Steiner trees.

We reduced the running time of our algorithms through a dynamic MST-maintenance scheme, and we proved that

TABLE VII
STATISTICS REGARDING THE NUMBER OF STEINER POINTS INDUCED BY
BIS IN THE UNRESTRICTED THREE-DIMENSIONAL CASE (i.e.,
 $L = \infty$). ALSO GIVEN ARE THE NUMBER OF ROUNDS EXECUTED BY
BIS. SHOWN ARE THE MINIMUM, AVERAGE, AND MAXIMUM VALUES

| #SPs and #rounds for BIS in 3D | | | | | | |
|--------------------------------|------|-------|-----|---------|------|-----|
| n | #SPs | | | #rounds | | |
| | min | ave | max | min | ave | max |
| 3 | 0 | 0.90 | 1 | 1 | 1.90 | 2 |
| 4 | 0 | 1.53 | 2 | 1 | 2.18 | 3 |
| 5 | 1 | 2.25 | 3 | 2 | 2.33 | 4 |
| 7 | 1 | 3.63 | 6 | 2 | 2.60 | 5 |
| 10 | 2 | 5.63 | 8 | 2 | 2.92 | 7 |
| 14 | 5 | 8.22 | 11 | 2 | 3.18 | 6 |
| 20 | 9 | 12.14 | 15 | 2 | 3.26 | 5 |
| 30 | 15 | 18.99 | 22 | 3 | 3.53 | 6 |
| 50 | 29 | 32.86 | 35 | 3 | 4.14 | 6 |

ACKNOWLEDGMENT

We would like to thank Mr. Tim Barrera for his help with the coding, and the referees for their helpful comments.

REFERENCES


- [1] M. J. Alexander and G. Robins, "An architecture-independent unified approach to FPGA routing," Dept. of Computer Science, Univ. of Virginia, Tech. Rep. CS-93-51, Oct. 1993.
- [2] M. J. Alexander and G. Robins, "A new approach to FPGA routing based on multi-weighted graphs," in *Proc. ACM/SIGDA Int. Workshop on Field-Programmable Gate Arrays*, Berkeley, CA, Feb. 1994.
- [3] T. Barrera, J. Griffith, S. A. McKee, G. Robins, and T. Zhang, "Toward a Steiner engine: Enhanced serial and parallel implementations of the iterated 1-Steiner algorithm," in *Proc. Great Lakes Symp. VLSI*, Kalamazoo, MI, Mar. 1993, pp. 90-94.
- [4] T. Barrera, J. Griffith, G. Robins, and T. Zhang, "Narrowing the gap: Near-optimal Steiner trees in polynomial time," in *Proc. IEEE Int. ASIC Conf.*, Rochester, NY, Sept. 1993, pp. 87-90.
- [5] A. Beguelin, J. J. Dongarra, G. A. Geist, R. Manchek, and V. S. Sunderam, "A user's guide to PVM: Parallel virtual machine," Oak Ridge National Laboratory, Tech. Rep. ORNL/TR-11826, 1991.
- [6] P. Berman, U. Foessmeier, M. Karpinski, M. Kaufmann, and A. Z. Zelikovsky, "Approaching the 5/4-Approximation for rectilinear Steiner trees," Wilhelm Schickard-Institut für Informatik, Tech. Rep. WSI-94-06, 1994.
- [7] P. Berman and V. Ramaiyer, "Improved approximations for the Steiner tree problem," in *Proc. ACM/SIAM Symp. Discrete Algorithms*, San Francisco, CA, Jan. 1992, pp. 325-334.
- [8] M. Borah, R. M. Owens, and M. J. Irwin, "An edge-based heuristic for rectilinear Steiner trees," Dept. of Computer Science, Pennsylvania State Univ., Tech. Rep. CS-93-003, 1993.
- [9] D. Braun, J. Burns, S. Devadas, H. K. Ma, K. M. F. Romeo, and A. Sangiovanni-Vincentelli, "Chameleon: A new multi-layer channel router," in *Proc. ACM/IEEE Design Automation Conf.*, 1986, pp. 495-502.
- [10] T. H. Chao and Y. C. Hsu, "Rectilinear Steiner tree construction by local and global refinement," *IEEE Trans. Computer-Aided Design*, pp. 303-309, 1994.
- [11] T. H. Cormen, C. E. Leiserson, and R. Rivest, *Introduction to Algorithms*. Cambridge, MA: MIT Press, 1990.
- [12] J. T. Croft, K. J. Falconer and R. K. Guy, *Unsolved Problems in Geometry*. New York: Springer-Verlag, 1991.
- [13] U. Foessmeier, M. Kaufmann, and A. Z. Zelikovsky, "Fast approximation algorithms for the rectilinear Steiner tree problem," Wilhelm Schickard-Institut für Informatik, Tech. Rep. WSI-93-14, 1993.
- [14] G. N. Fredrickson, "Data structures for on-line updating of minimum spanning trees," *SIAM J. Comput.*, vol. 14, pp. 781-798, 1985.
- [15] J. L. Ganley and J. P. Cohoon, "Routing a multi-terminal critical net: Steiner tree construction in the presence of obstacles," in *Proc. IEEE Intl. Symp. Circuits Syst.*, London, England, May 1994.
- [16] M. Garey and D. S. Johnson, "The rectilinear Steiner problem is NP-complete," *SIAM J. Applied Math.*, vol. 32, pp. 826-834, 1977.
- [17] G. Georgakopoulos and C. H. Papadimitriou, "The 1-Steiner tree problem," *J. Algorithms*, vol. 8, pp. 122-130, 1987.
- [18] E. N. Gilbert and H. O. Pollak, "Steiner minimal trees," *SIAM J. Applied Math.*, vol. 16, pp. 1-29, 1968.
- [19] R. L. Graham and F. K. Hwang, "A remark on Steiner minimal trees," *Bull. Inst. Math. Acad. Sinica*, vol. 4, pp. 177-182, 1976.
- [20] A. Hanafusa, Y. Yamashita, and M. Yasuda, "Three-dimensional routing for multilayer ceramic printed circuit boards," in *Proc. IEEE Int. Conf. Computer-Aided Design*, Santa Clara, CA, Nov. 1990, pp. 386-389.
- [21] M. Hanan, "On Steiner's problem with rectilinear distance," *SIAM J. Applied Math.*, vol. 14, pp. 255-265, 1966.
- [22] A. C. Harter, *Three-Dimensional Integrated Circuit Layout*. New York: Cambridge University Press, 1991.
- [23] N. Hasan, G. Vijayan, and C. K. Wong, "A neighborhood improvement algorithm for rectilinear Steiner trees," in *Proc. IEEE Int. Symp. Circuits Syst.*, New Orleans, LA, 1990.
- [24] J.-M. Ho, G. Vijayan, and C. K. Wong, "New algorithms for the rectilinear Steiner tree problem," *IEEE Trans. Computer-Aided Design*, vol. 9, pp. 185-193, 1990.
- [25] F. K. Hwang, "On Steiner minimal trees with rectilinear distance," *SIAM J. Applied Math.*, vol. 30, pp. 104-114, 1976.
- [26] F. K. Hwang, "An $O(n \log n)$ algorithm for rectilinear minimal spanning trees," *J. ACM*, vol. 26, pp. 177-182, 1979.
- [27] F. K. Hwang, D. S. Richards, and P. Winter, *The Steiner Tree Problem*. North-Holland, 1992.
- [28] A. B. Kahng and G. Robins, "A new class of iterative Steiner tree heuristics with good performance," *IEEE Trans. Computer-Aided Design*, vol. 11, pp. 893-902, 1992.
- [29] A. B. Kahng and G. Robins, "On performance bounds for a class of rectilinear Steiner tree heuristics in arbitrary dimension," *IEEE Trans. Computer-Aided Design*, vol. 11, pp. 1462-1465, 1992.
- [30] A. B. Kahng and G. Robins, *On Optimal Interconnections for VLSI Layout*. Boston, MA: Kluwer Academic Publishers (in press).
- [31] J. H. Lee, N. K. Bose, and F. K. Hwang, "Use of Steiner's problem in sub-optimal routing in rectilinear metric," *IEEE Trans. Circuits Syst.*, vol. 23, pp. 470-476, 1976.
- [32] K. W. Lee and C. Sechen, "A new global router for row-based layout," in *Proc. IEEE Int. Conf. Computer-Aided Design*, Santa Clara, CA, Nov. 1990, pp. 180-183.
- [33] F. D. Lewis, W. C. Pong, and N. VanCleave, "Local improvement in Steiner trees," in *Proc. Great Lakes Symp. VLSI*, Kalamazoo, MI, Mar. 1993, pp. 105-106.
- [34] A. L. Loeb, *Space Structures: Their Harmony and Counterpoint*. New York: Birkhauser, 1991.
- [35] C. Monma and S. Suri, "Transitions in geometric minimum spanning trees," *Discrete & Computational Geometry*, vol. 8, pp. 265-293, 1992.
- [36] C. H. Papadimitriou and U. V. Vazirani, "On two geometric problems relating to the traveling salesman problem," *J. Algorithms*, vol. 5, pp. 231-246, 1984.
- [37] B. T. Preas and M. J. Lorenzetti, *Physical Design Automation of VLSI Systems*. Menlo Park, CA: Benjamin/Cummings, 1988.
- [38] G. Robins, "On optimal interconnections," Ph.D. Dissertation, CSD-TR-920024, Dept. of Computer Science, UCLA, 1992.
- [39] G. Robins and J. S. Salowe, "On the maximum degree of minimum spanning trees," in *Proc. ACM Symp. Computational Geometry*, Stony Brook, NY, June 1994, pp. 250-258.
- [40] J. S. Salowe and D. M. Warne, "An exact rectilinear Steiner tree algorithm," in *Proc. IEEE Int. Conf. Computer Design*, Cambridge, MA, Oct. 1993, pp. 472-475.
- [41] J. M. Smith and J. S. Liebman, "Steiner trees, Steiner circuits and the interference problem in building design," *Engineering Optimization*, vol. 4, pp. 15-36, 1979.
- [42] T. L. Snyder, "On the exact location of Steiner points in general dimension," *SIAM J. Comput.*, vol. 21, pp. 163-180, 1992.
- [43] V. S. Sunderam, "PVM: A framework for parallel distributed computing," *Concurrency: Practice and Experience*, vol. 2, pp. 315-339, 1990.
- [44] G. F. Tóth, "New results in the theory of packing and covering," *Convexity and its Applications*, 1983.
- [45] A. C. C. Yao, "On constructing minimum spanning trees in k -dimensional spaces and related problems," *SIAM J. Comput.*, vol. 11, pp. 721-736, 1982.
- [46] A. Z. Zelikovsky, "An 11/6 approximation algorithm for the network Steiner problem," *Algorithmica*, vol. 9, pp. 463-470, 1993.
- [47] A. Z. Zelikovsky, "A faster approximation algorithm for the Steiner tree problem in graphs," *Information Processing Letters*, vol. 46, pp. 79-83, 1993.



Jeff Griffith received his Bachelor's degree in mathematics in 1992 from the University of Virginia in Charlottesville. He is presently working on his Ph.D. in computer science at the University of Tennessee in Knoxville. His areas of interest include physical design for VLSI circuits, as well as scientific computing.



Jeffrey S. Salowe received the B.A. Degree in mathematics from the University of Virginia, Charlottesville, VA, in 1983, and the MS. and Ph.D. degrees in computer science from Rutgers University, New Brunswick, NJ, in 1985 and 1987, respectively. Currently, he is a consultant for Questech, Inc. His research interests include the analysis of algorithms, computational geometry, data structures, and graph algorithms.



Gabriel Robins (S'91-M'91) received the Ph.D. degree in 1992 from the UCLA Computer Science Department, where he won the Distinguished Teaching Award and held an IBM Graduate Fellowship. Currently, Dr. Robins is assistant professor of computer science at the University of Virginia in Charlottesville, where he recently won the NSF Young Investigator Award. His primary areas of research are VLSI CAD and geometric algorithms, with recent work focusing on performance-driven routing, Steiner tree heuristics, computational geometry, motion planning, pattern recognition, and computational biology. Dr. Robins is the author of a Distinguished Paper at the 1990 IEEE International Conference on Computer-Aided Design, and his Ph.D. dissertation was nominated for the ACM Distinguished Dissertation Award. He is a member of ACM, IEEE, MAA, and SIAM.



Tongtong Zhang studied at the University of Beijing, China for three years and received her Bachelor's degree in 1992 from Bridgewater College, Virginia. She is currently working on her Ph.D. in computer science at the University of Virginia in Charlottesville. Her areas of interest include physical design for VLSI circuits and computational biology.