[22] A. B. Kahng and G. Robins, "A New Class of Iterative Steiner Tree Heuristics with Good Performance", *IEEE Transactions on CAD* 11(7), July 1992, pp. 893-902.

[23] S. Khuller, B. Raghavachari and N. Young, "Balancing Minimum Spanning and Shortest Path Trees", *Proc. ACM/SIAM Symp. on Discrete Algorithms*, January 1993, to appear.

[24] S. Kim, R. M. Owens and M. J. Irwin, "Experiments with a Performance Driven Module Generator", *Proc. ACM/IEEE Design Automation Conf.*, 1992, pp. 687-690.

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

[26] I. Lin and D. H. C. Du, "Performance-Driven Constructive Placement", *Proc. ACM/IEEE Design Automation Conf.*, 1990, pp. 103-106.

[27] M. Marek-Sadowska and S. Lin, "Timing Driven Placement", *Proc. IEEE Intl. Conf. on Computer-Aided Design*, 1989, pp. 94-97.

[28] S. Prasitjutrakul and W. J. Kubitz, "A Timing-Driven Global Router for Custom Chip Design", *Proc. IEEE Intl. Conf. on Computer-Aided Design*, 1990, pp. 48-51.

[29] A. Prim, "Shortest Connecting Networks and Some Generalizations", *Bell System Tech. J.* 36 (1957), pp. 1389-1401.

[30] S. K. Rao, P. Sadayappan, F. K. Hwang and P. W. Shor, "The Rectilinear Steiner Arborescence Problem", *Algorithmica* 7 (1992), pp. 277-288.

[31] G. Robins, "On Optimal Interconnections", Ph.D. thesis (*technical report* CSD TR-920024), CS Department, University of California, Los Angeles, June 1992.

[32] J. Rubinstein, P. Penfield, and M. A. Horowitz, "Signal Delay in RC Tree Networks", *IEEE Trans. on CAD* 2(3) (1983), pp. 202-211.

[33] A. Srinivasan, K. Chaudhary and E. S. Kuh, "RITUAL: A Performance Driven Placement Algorithm for Small-Cell ICs", *Proc. IEEE Intl. Conf. on Computer-Aided Design*, 1991, pp. 48-51.

[34] S. Sutanthavibul and E. Shragowitz, "Adaptive Timing-Driven Layout for High Speed VLSI", *Proc. ACM/IEEE Design Automation Conf.*, 1990, pp. 90-95.

[35] S. Teig, R. L. Smith and J. Seaton, "Timing Driven Layout of Cell-Based ICs", *VLSI Systems Design*, May 1986, pp. 63-73.

[36] R. S. Tsay, "Exact Zero Skew", *Proc. IEEE Intl. Conference on Computer-Aided Design*, 1991, pp. 336-339.

[37] D. Zhou, F. P. Preparata and S. M. Kang, "Interconnection Delay in Very High-speed VLSI", *IEEE Trans. on Circuits and Systems* 38(7), 1991.

[38] D. Zhou, S. Su, F. Tsui, D. S. Gao and J. Cong, "Analysis of Trees of Transmission Lines", *technical report* UCLA CSD-920010.

References

[1] C. J. Alpert, T. C. Hu, J. H. Huang and A. B. Kahng, "A Direct Combination of the Prim and Dijkstra Constructions for Improved Performance-Driven Global Routing", *technical report* CSD-920051, UCLA Department of Computer Science, 1992.

[2] T. G. Andrews, ed., *Methods of Psychology*, New York, John Wiley, 1948.

[3] B. Awerbuch, A. Baratz and D. Peleg, "Cost-Sensitive Analysis of Communication Protocols", *Proc. ACM Symp. on Principles of Distributed Computing*, 1990, pp. 177-187.

[4] Barrera, T., J. Griffith, G. Robins and T. Zhang, "Narrowing the Gap: Near-Optimal Steiner Trees in Polynomial Time", *Proc. IEEE Intl. ASIC Conf.*, Rochester, Sept. 1993, pp. 87-90.

[5] K. D. Boese, J. Cong, A. B. Kahng, K. S. Leung and D. Zhou, "On High-Speed VLSI Interconnects: Analysis and Design" *Proc. Asia-Pacific Conf. on Circuits and Systems*, Sept. 1992, pp. 35-40.

[6] K.D. Boese, A. B. Kahng and G. Robins, "High-Performance Routing Trees with Identified Critical Sinks", *Proc. ACM/IEEE Design Automation Conf.*, June 1993, pp. 182-187.

[7] K. D. Boese, A. B. Kahng, B. A. McCoy and G. Robins, "Fidelity and Near-Optimality of Elmore-Based Routing Constructions", *Proc. IEEE Intl. Conf. on Computer Design*, October 1993, pp. 81-84.

[8] D.-S. Chen and M. Sarrafzadeh, "A Wire-Length Minimization Algorithm for Single-Layer Layouts", *Proc. IEEE Intl. Conference on Computer-Aided Design*, 1992, pp. 390-393.

[9] J. P. Cohoon and L. J. Randall, "Critical Net Routing", *Proc. IEEE Intl. Conf. on Computer Design*, 1991, pp. 174-177.

[10] J. Cong, A. B. Kahng, G. Robins, M. Sarrafzadeh, and C. K. Wong, "Provably Good Performance-Driven Global Routing", *IEEE Trans. on CAD* 11(6), June 1992, pp. 739-752.

[11] J. Cong, K.-S. Leung and D. Zhou, "Performance-Driven Interconnect Design Based on Distributed RC Delay Model", *Proc. ACM/IEEE Design Automation Conf.*, 1993, pp. 606-611.

[12] E. W. Dijkstra, "A Note on Two Problems in Connection With Graphs", *Numerische Mathematik* 1(1959), pp. 269-271.

[13] W. E. Donath, R. J. Norman, B. K. Agrawal, S. E. Bello, S. Y. Han, J. M. Kurtzberg, P. Lowy and R. I. McMillan, "Timing Driven Placement Using Complete Path Delays", *Proc. ACM/IEEE Design Automation Conf.*, 1990, pp. 84-89.

[14] A. E. Dunlop, V. D. Agrawal, D. N. Deutsh, M. F. Jukl, P. Kozak and M. Wiesel, "Chip Layout Optimization Using Critical Path Weighting", *Proc. ACM/IEEE Design Automation Conf.*, 1984, pp. 133-136.

[15] W. C. Elmore, "The Transient Response of Damped Linear Network with Particular Regard to Wide-band Amplifiers", *J. Applied Physics* 19 (1948), pp. 55-63.

[16] S. Even, *Graph Algorithms*, Potomac, MD, Computer Science Press, 1979.

[17] M. Hanan, "On Steiner's Problem with Rectilinear Distance", *SIAM J. Appl. Math.*, 14 (1966), pp. 255-265.

[18] P. S. Hauge, R. Nair and E. J. Yoffa, "Circuit Placement for Predictable Performance", *Proc. IEEE Intl. Conf. on Computer-Aided Design*, 1987, pp. 88-91.

[19] J.-M. Ho, G. Vijayan and C. K. Wong, "New Algorithms for the Rectilinear Steiner Tree Problem", *IEEE Transactions on Computer-Aided Design*, 9(2), 1990, pp. 185-193.

[20] M. A. B. Jackson and E. S. Kuh, "Estimating and Optimizing RC Interconnect Delay During Physical Design", *Proc. IEEE Intl. Conf. on Circuits and Systems*, 1990, pp. 869-871.

[21] M. A. B. Jackson, E. S. Kuh, and M. Marek-Sadowska, "Timing-Driven Routing for Building Block Layout", *Proc. IEEE International Symposium on Circuits and Systems*, pp. 518-519, 1987.

figure) must be on maximal segments with entry point q; it is easy to see that q is the closest connection between p, $Pin(a)$, and $Pin(b)$.

In Figure 22(b)-(d), edge (p, q) is a a straight edge. Let M be the MS containing (p, q), and let M' be the MS perpendicular to M with entry point q.

- In Figure 22(b), edge (q, a) is L-shaped and edge (q, b) is on the MS M'. By Lemma B4, M' must contain a sink, which will be contained in subtree T_b. Thus, $Pin(b)$ (n_2 in the Figure) is located on M'. Node a is the entry point for two branches perpendicular branches containing sinks (by Lemma B4); $Pin(a)$ is chosen arbitrarily from one of these branches (Line 3 in Figure 19). In Figure 22, either $Pin(a) = n_1$ or $Pin(a) = n_1'$; thus, it can be seen from the Figure that q is the closest connection between p, $Pin(a)$, and $Pin(b)$.

- In Figure 22(c), M' is the union of two branches. One of these branches contains a sink (by Lemma B4); without loss of generality, let this be the branch containing edge (q, a), with $Pin(a) = n_1$ in M'. Let B be the branch containing edge (q, b). If $Pin(b)$ is on B, then q will be the closest connection between p, $Pin(a)$ and $Pin(b)$. Otherwise, according to Lemma B2 we must have that b is the entry point to a far branch off of M'. Hence, if $Pin(b)$ is not on B, the b–$Pin(b)$ path in T^* contains only edges on far branches (by the criteria in Lines 8-10 in Figure 19; see $n_2 = Pin(b)$ in Figure 22(c)). Thus, $Pin(b)$ is contained in the upper-right quadrant relative to q in the Figure, and q is the closest connection between p, $Pin(a)$, and $Pin(b)$.

- Finally, consider the configuration in Figure 22(d). Here, MS M' is a branch of M containing node a and sink $Pin(a)$. Suppose that M' is a far branch; if $Pin(b)$ is not on MS M, then there must be a near branch off of M somewhere below q in T^* (otherwise, we could reduce all delays by shifting the entire half segment of M below q toward a). Let B_j be the near branch below q closest to q. Either sink $Pin(b)$ is on B_j, or the q_j–$Pin(b)$ path in T^* consists only of edges in B_j or far branches. In either case, $Pin(b)$ ($= n_2$ in the Figure) is contained in the lower-right quadrant relative to q. If M' is a near branch, an analogous argument again shows that $Pin(b)$ is in q's lower-right quadrant. Thus, q is the closest connection between p, $Pin(a)$ and $Pin(b)$. □

Except for redundancies and pruning of sub-optimal trees, BB-SORT-C searches over all possible ways to construct a Steiner tree sequentially, such that each sink is added by a closest connection to some edge in the current tree. Thus, we have:

Theorem B1: For any positive linear combination of sink delays, $f = \sum_{i=1}^{k} \alpha_i \cdot t(n_i)$, $\alpha_i > 0 \ \forall i$, algorithm BB-SORT-C returns a Steiner tree T^* which minimizes f. □

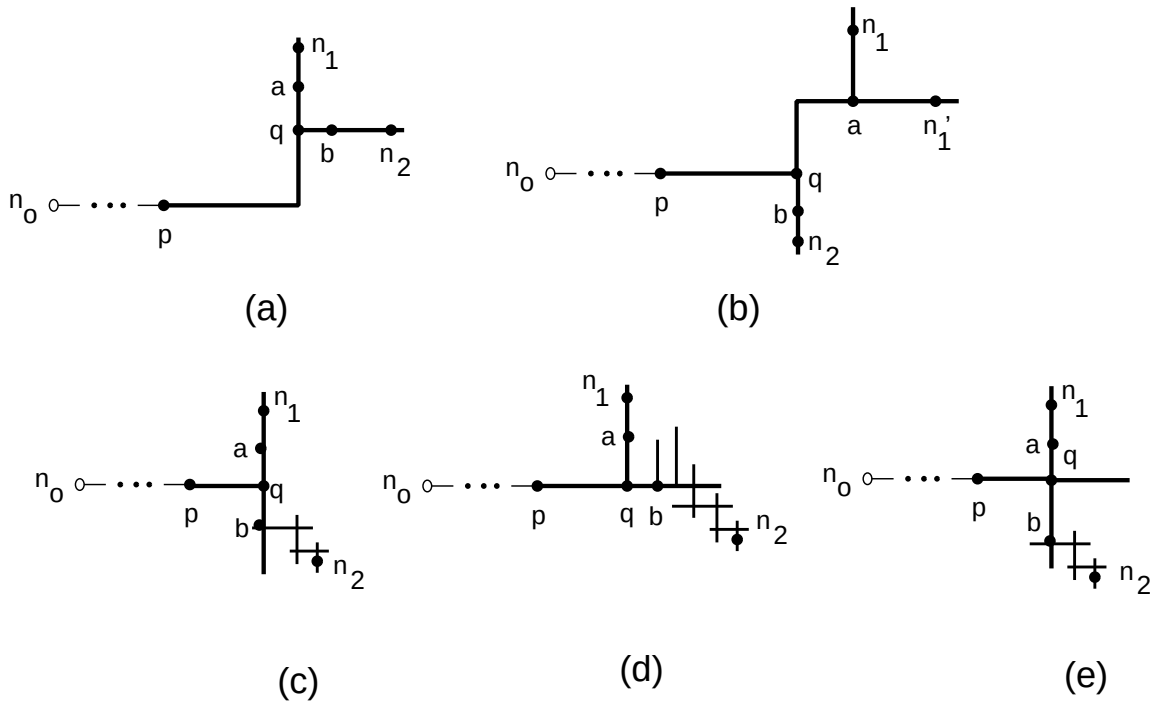

Figure 22: Five possible topologies at any Steiner node q in T^*. Each diagram shows two sinks n_1 and n_2 below node q in the tree, such that q is the closest connection between n_1, n_2 and q's parent p.

Proof: Part (i) of the Lemma is true since the construction of Figure 20 removes exactly one sink during each pass through Lines 3 to 9.

To show (ii), let p be the parent of the node q at Line 3 in Figure 20. The first case is when q is a sink or a degree-4 Steiner node in T_i (as in Figure 22(e)). In this case, edge (p, q) will remain in tree T_{i-1}. If (p, q) is L-shaped, we must have a connection as in Figure 22(a), where the two children of q are eventually replaced by sinks on the maximal segments with entry point q (i.e., n_1 and n_2 in the Figure). Both of these sinks have closest connections to (p, q) at q. If (p, q) is a straight edge, let M be the MS containing (p, q), and let a be a child of q in T^*. The sink $Pin(a)$ is assigned in the Figure 19 template such that the q–$Pin(a)$ path in T^* will contain only edges in M, edges in branches off of M, or edges in a sequence of far branches off of branches of M. (For example, consider the paths from q to sinks n_1 and n_2 in Figure 22(c)-(e).) Thus, $Pin(a)$ and p cannot be on the same side of a line that passes through q and is perpendicular to M. Consequently, q will be the closest connection between edge (p, q) and $Pin(a)$.

The second case is when q is a degree-3 Steiner node in T_i. Let a and b be the children of q in T^* such that $Pin(a)$ and $Pin(b)$ are q's children in T_i. Without loss of generality, we assume that $Pin(q) = Pin(a)$ and $n_i = Pin(b)$. We must show that q is located at the closest connection between nodes p, $Pin(a)$, and $Pin(b)$. There are four possible configurations for connections at q, as shown in parts (a)-(d) of Figure 22.

- In Figure 22(a), edge (p, q) is L-shaped and both $Pin(a)$ and $Pin(b)$ (denoted by n_1 and n_2 in the

T^* **Decomposition Procedure**
Input: Optimal delay tree T^*
Output: Sequence of sinks $n_1, \ldots n_k$ used to construct T^* using only closest connections of each n_i to T_{i-1}
1. $i = k$
2. **Repeat** until $i == 0$
3. Find a node $q \in T_i$ whose children are all leaves
4. **If** $q == n_0$ set c to be any child of n_0 in T^*
5. **Else if** q has degree 4
6. Set c to be the child of q in T^* on the MS containing edge $(parent(q), q)$
7. **Else** Set c to be a child of q such that $Pin(c) \neq Pin(q)$
8. $n_i = Pin(c)$
9. $T_{i-1} = T_i \backslash n_i$
10. $i = i - 1$

Figure 20: Procedure to determine a sequence of sinks $n_1, \ldots, n_k$ which can be used to construct T^* by a sequence of closest connections from n_i to tree T_{i-1}.

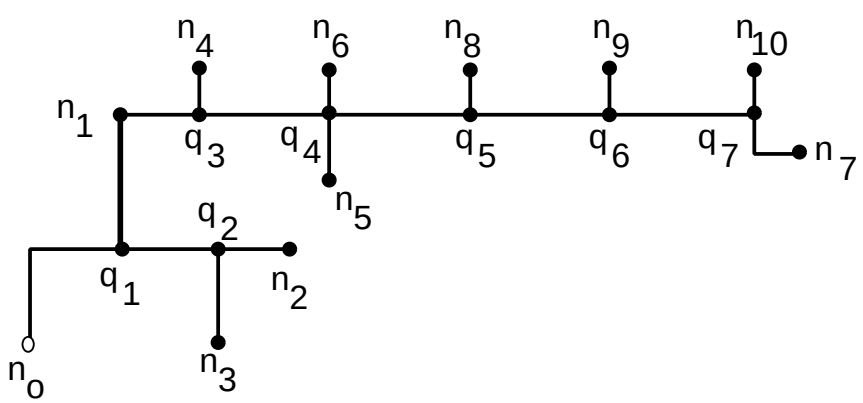

Figure 21: Example of the order in which pins are "peeled" from an optimal-delay Steiner tree T^*. (Sinks n_i are peeled from T^* in reverse order of their subscripts.)

Pin Assignments for Example Tree		
Steiner node q	$Pin(q)$	Reason for Assignment
q_1	n_1	(n_0, q_1) is L-shaped (Line 4)
q_2	n_2	n_2 is on MS containing (q_1, q_2) (Line 6)
q_3	n_4	n_1 is the entry point to MS containing (n_1, q_3) (Line 8)
q_4	n_5	far branch at q_3 and near branch at q_4 (Line 15)
q_5	$Pin(q_6)$	q_4 has degree 4 (Line 11)
q_6	$Pin(q_7)$	far branch at q_5 and no near branch at q_6 (Line 16)
q_7	n_7	far branch at q_6 and near branch at q_7 (Line 15)

Table 12: Pin assignment to Steiner nodes in the example of Figure 21. Line numbers refer to the corresponding line in the Pin assignment procedure of Figure 19.

if q is the source or a sink, then $Pin(q) = q$; and (ii) if q is a Steiner node, then $Pin(q)$ is chosen according to the template given in Figure 19.

$Pin(q)$ **Assignment Procedure**
Input: Optimal delay tree T^* Steiner node $q \in T^*$ such that $Pin(w)$ has been assigned for each $w \in T^*_q$, $w \neq q$
Output: $Pin(q)$
1. $p = parent(q)$ in T^* 2. **If** edge (p, q) is a straight edge 3. Set M to be the MS containing (p, q) 4. **If** edge (p, q) is L-shaped 5. Set c arbitrarily to be one of q's two children /* (q has exactly two children, by Lemma B1) */ 6. **Else if** T_q contains a sink on M 7. Set c to be the child of q on M 8. **Else if** p is the entry point to M 9. Set B to be the far branch of M at q /* (such a B exists by Lemma B2) */ 10. Set c to be the child of q on B 11. **Else if** p has degree 4 12. Set c arbitrarily to be one of q's children 13. **Else if** there is a near (far) branch of M at p 14. **If** there is a far (near) branch B of M at q 15. Set c to be the child of q on B 16. **Else** Set c to be the child of q on M 17. $Pin(q) = Pin(a)$

Figure 19: Criteria used to associate a sink $Pin(q)$ with each Steiner node q in the optimal-delay tree T^*. The assignment determines which sink in T_q will remain in the tree when q is removed from the current tree while "peeling off" sinks from T^*.

After $Pin(q)$ has been assigned, we can apply the rules described in Figure 20 to peel off sinks, thus determining the sequence in which sinks can be added to construct T^*. Note that node p in Line 3 of Figure 20 must exist since T_i is finite and has no cycles. Figure 21 gives an example of the decomposition procedure applied to an eleven-pin net. Sinks in the figure are labeled in reverse order of how the decomposition procedure might peel them off from the tree. (Other orders are possible because the decomposition procedure is not completely deterministic.) Table 12 shows how $Pin(q)$ was assigned for each Steiner node in Figure 21.

We now show that the procedure of Figure 20 gives a sequential decomposition of the optimal-delay tree T^*, such that each T_i is constructed by connecting sink n_i to tree T_{i-1} by a closest connection to some edge in T_{i-1}.

Lemma B5: There exists a sequence of subtrees $T_0 = \{n_0\}, T_1, T_2, \ldots, T_k = T^*$ such that for each i, $1 \leq i \leq k$, (i) there is a sink $n_i \in T_i$ such that $T_{i-1} = T_i \backslash n_i$, and (ii) either n_i is connected to n_0, or n_i makes a closest connection in T_i to some edge in T_{i-1}.

for $t(n_i)$ is

$$t(n_i) = (1-d) * r^2 + K_1 * r + K_0 \tag{9}$$

where K_1 and K_0 are constants. From Lemma B3, we know that $d \geq 1$, implying that $t(n_i)$ is a concave function of r. Similarly, delay to a sink n_j along a far branch b_j off of M will be equal to

$$-d * r^2 + K_1' r + K_0' \tag{10}$$

where again K_1' and K_0' are constants; this too is a concave function of r. Finally, delay to any sink whose source-sink path does not contain an edge in M will be linear in r, and thus also a concave function. Since any linear combination of functions that are each concave on a given interval will also be concave on that interval, f is concave in r and is minimized at one of its extreme values, i.e., at $r = 0$ or $r = x_2 - x_1$.

Thus, M may be moved so that it contains a new node, say p_i. If p_i is a sink, the lemma is proved. If p_i is a Steiner node, then because it has degree > 2, there must be a vertical MS incident to p_i, and this vertical MS must contain a sink since M is the lowest maximal segment not containing a sink. Hence, if p_i is a Steiner node, the shifted M will also contain a sink. □

A direct corollary of Lemma B4 is that all Steiner nodes in the Elmore-optimal Steiner tree are contained in the Hanan grid:

Corollary: Let X be the set of x-coordinates for all pins in N, and let Y be the set of y-coordinates in N. Then if (x, y) is the location of a Steiner node in T^*, $x \in X$ and $y \in Y$. □

Thus, only a finite number of possible Steiner point locations need to be considered. Hanan's original theorem may be viewed as a special case of this Corollary with the driver on-resistance $r_d \to \infty$.

B.4. Decomposition Theorem for T^*

To prove that BB-SORT-C will return the optimal-delay tree T^*, we show that T^* can be constructed by starting with a tree T_0 containing only n_0, then successively adding a sequence of sinks n_i, $1 \leq i \leq k$, each of which yields a tree T_i by making a closest connection to some edge in the current tree T_{i-1}. We show that such a sequence of trees exists by starting with $T^* = T_k$ and $i = k$, then "peeling off" an n_i at each iteration such that n_i was joined by a closest connection in T_i to some edge in $T_{i-1} = T_i \backslash n_i$.

At each step, we find an interior node $q \in T_i$ whose children are all leaves. Each of these leaves must be a sink, because all low-degree Steiner nodes (i.e., with degree < 3) are removed from $T_{i+1} \backslash n_{i+1}$. We choose one of q's leaves to be the n_i that is peeled off, and set $T_{i-1} = T_i \backslash n_i$. The choice of which leaf should be peeled is guided by the function $Pin(q)$, which specifies one of q's children that should *not* be peeled off from q. Thus, when q is removed as a low-degree Steiner node, the edge between q and its parent is replaced with an edge between $Pin(q)$ and q's parent. More formally, we define $Pin(q)$ for each node $q \in T^*$ as follows: (i)

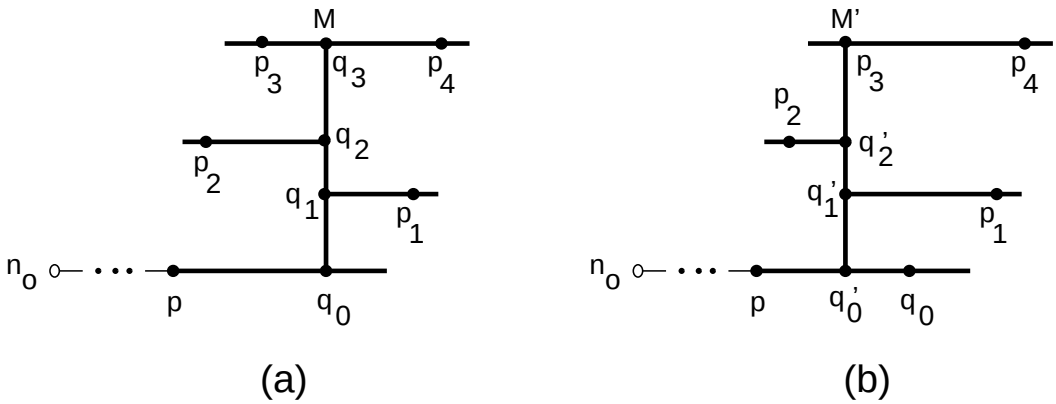

Figure 17: Proof of Lemma B3: (a) Example where $|Near(M)| = |Far(M)|$ for maximal segment M. M can be shifted to M', as shown in (b), to reduce delay at all sinks in the subtree $T^*_{q_0}$.

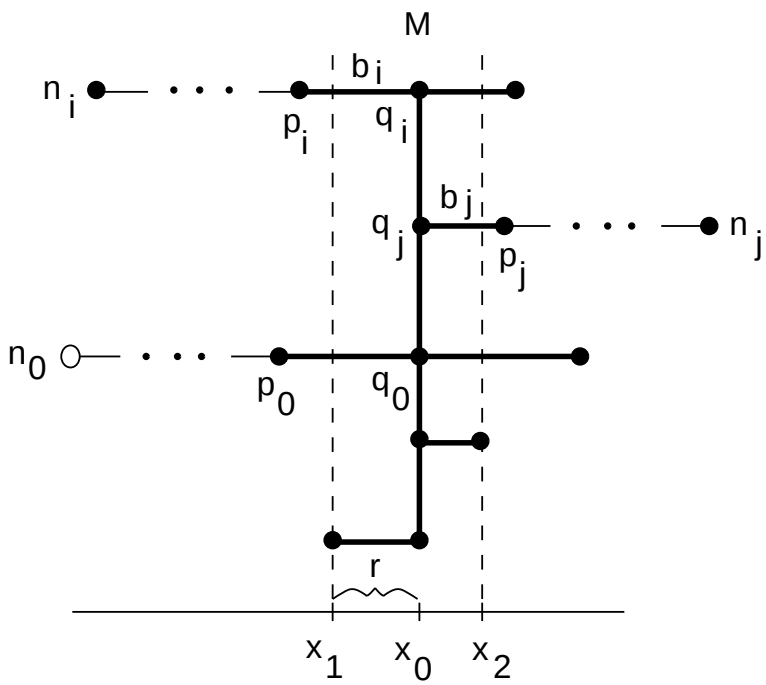

Figure 18: Proof of Lemma B4: because the objective function f is concave in r over the interval $0 \leq r \leq x_2 - x_1$, f will be minimized when the maximal segment M passes through either the gridline at x_1 or the gridline at x_2.

in M. Let the variable r, $0 \leq r \leq x_2 - x_1$, denote the position of M between the x-coordinates x_1 and x_2. We will show that minimizing the delay function f implies that either $r = 0$ or $r = x_2 - x_1$.

Let $d = Far(M) - Near(M)$. Consider the delay to some sink n_i located along a near branch b_i off of M which has entry point q_i. (In general, we let q_j denote the entry point to branch b_j.) Delay $t(n_i)$ is quadratic in r only along the edge (p_0, q_0) and along the edge (q_i, p_i), where p_i is the child of q_i on b_i. To be precise, the delay due to (p_0, q_0) is equal to $r * (r/2 - d * r + K)$, where K is some constant; the delay due to (q_i, p_i) is equal to $r * (r/2 + K') + K''$, where K' and K'' are again constants. Therefore, the equation

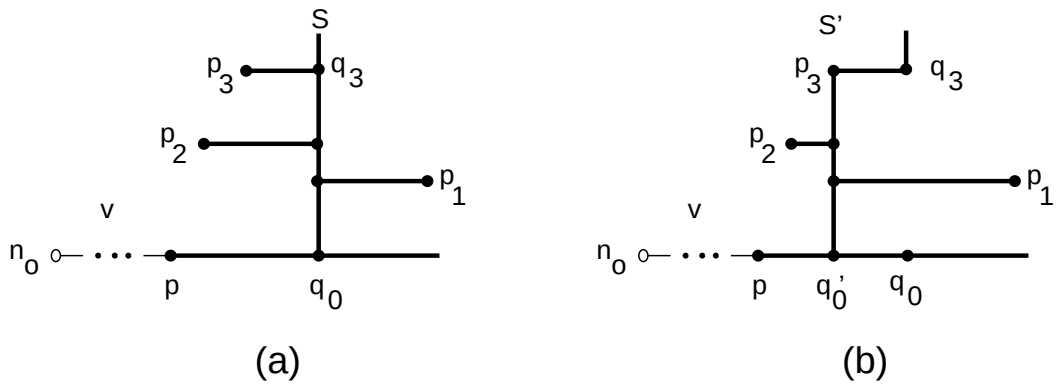

Figure 16: Proof of Lemma B2: Example (a) with $|Near(S)| > |Far(S)|$ for a segment S between q_0 and q_3; (b) shows how S can be shifted to S' to reduce delay to all sinks in $T^*_{q_0}$ and leave delay unchanged at all other sinks.

We can now use Lemma B2 to show that if an MS M has as many near branches as far branches, then it can be shifted so as to reduce delay to some sinks and leave delay to others unchanged. This is because shifting M toward the source will not increase total wirelength and will decrease some source-sink pathlengths and and shift capacitance along some source-sink paths closer to source. Intuitively, this is the proof behind our next lemma.

Lemma B3: In the optimal tree T^*, let M be an MS not containing n_0. Then $|Far(M)| > |Near(M)|$.

Proof: By Lemma B2, $|Far(M)| \geq |Near(M)|$. Suppose that the exact equality $|Far(M)| = |Near(M)|$ holds. Lemma B2 then implies that each endpoint of M has a near branch incident to it as in Figure 17(a) (otherwise, M would contain a subsegment S with all but one endpoint of M and having $|Near(S)| > |Far(S)|$.) In Figure 17(b) we show how M can be shifted toward the source without increasing total wirelength, while reducing source-sink pathlengths to nodes on the near branches of M and shifting capacitance toward the source for nodes on the far branches of M. Consequently, moving M will reduce delay to all sinks in $T^*_{q_0}$ and leave delay to all other sinks unchanged (or reduced if the shift reduces total wirelength), thereby contradicting the optimality of T^*. □

Lemma B4: In the optimal tree T^*, any maximal segment must contain either the source or a sink.

Proof: (See Figure 18.) Let M be a lowest maximal segment in T^* which does not contain either the source or a sink, i.e., every MS that is topologically below M contains a sink. Let q_0 be the entry point to M and let p_0 be the parent node of q_0 in T^*. Consider the possibility of shifting M, either toward the source or away from the source, without passing over any node in $T^*_{q_0}$ which is not in M. Without loss of generality, assume that M is a vertical segment with x-coordinate x_0, and with the near side of M having $x < x_0$. Let $x_1 < x_0$ be the be the closest value to x_0 on the near side of M such that shifting M to $x = x_1$ would cause M to intersect a node that is in $T^*_{q_0}$ but not in M. Similarly, let $x_2 > x_0$ be the closest value to x_0 on the far side of M such that shifting M to $x = x_2$ would cause M to intersect some node that is in $T^*_{q_0}$ but not

children are located on the far side of M (i.e., it is not the entry point to a larger far branch). For any segment S, we use $Near(S)$ (resp. $Far(S)$) to denote the set of near (resp. far) branches off of the maximal segment containing S. Figure 15 gives an example of an MS M with endpoints p_1 and p_2, entry point p_0, and four branches, including near branch b_1, far branch b_2, and a far branch consisting only of sink n_3.

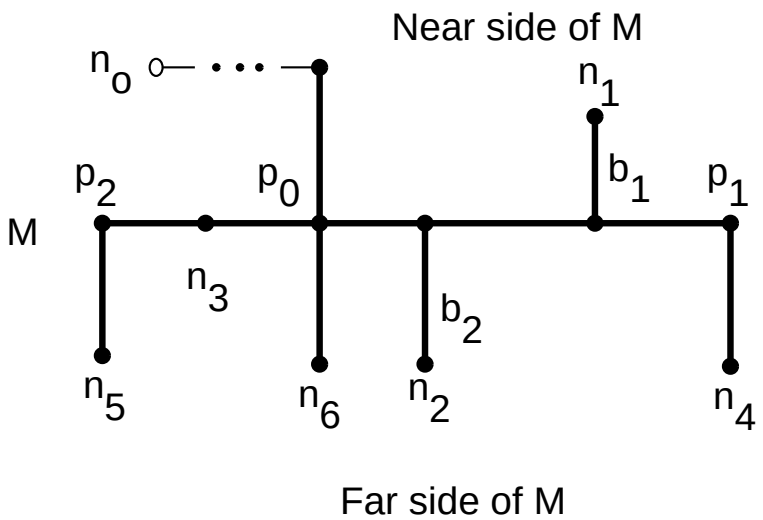

Figure 15: Example of a maximal segment M with entry point p_0, one near branch b_1, and four far branches, including b_2. Note that by definition, n_3 forms a far branch with no edges. Also, edge (p_0, n_6) does *not* form a far branch off of M because p_0 is not an entry point to the MS containing (p_0, n_6).

Lemmas B2 and B3 establish some properties that must hold for any maximal segment in T^*. Lemma B4 then uses these properties to show that each maximal segment in T^* will have a sink located on it. An immediate corollary of Lemma B4 is a generalization of the classic result of Hanan [17] to the Elmore delay objective. (Hanan's original theorem may be viewed as a special case of this Corollary with the driver on-resistance $r_d \rightarrow \infty$.)

Lemma B2: In the optimal tree T^*, let q_0 be the entry point to a maximal segment M not containing n_0. Let S be any segment contained in M and having q_0 as an endpoint. Then $|Far(S)| \geq |Near(S)|$.

Proof: By contradiction. Let S be the smallest segment in M with q_0 as an endpoint so that $Near(S) > Far(S)$. Then a portion of T^* between n_0 and q' looks like Figure 16(a). Label the branches $b_1, \ldots, b_j$ in order from entry point q_0. Figure 16(b) shows how we can shift segment S topologically toward the source; this effectively shifts wire from each near branch to a far branch with is topologically closer to the source (i.e., with a smaller label). Shifting S does not affect tree cost[22], and source-sink path lengths will be unchanged to all sinks except those connected to the tree through branches in $Near(S)$, which will have reduced source-sink path lengths. Consequently, the shift will decrease delay to all sinks below q_0 in T^* and leave delay to all other sinks unchanged, contradicting the optimality of T^*. □

[22] Unless q_0 is an endpoint for the MS containing edge (p, q_0), in which case tree cost will decrease.

can only be minimized when $x = 0$ or $x = c$. If $x = 0$ then either the node x is located at the source (and the lemma is satisfied) or $x \in D$, which is impossible. If $x = c$, then it is easy to see that $x = c$ is located at the closest connection between any edge e on the n_0—n_i path to which a is connected. □

Straightforward corollaries of Lemma B1 include: (i) that any non-source node in the optimal delay tree T^* must have degree ≤ 4, and (ii) that the possible configurations of edges incident to a Steiner node $q \in T^*$ are restricted to the five configurations shown in Figure 22. Note that Lemma B1 by itself is not sufficient to prove that BB-SORT-C will return the optimal delay tree. For example, if T^* connects a four-pin net into an "H" with two degree-3 Steiner nodes q_1 and q_2 (see Figure 14), then the parent of each non-source node v is connected by a closest connection to $T^* \backslash v$. However, T^* cannot be constructed by BB-SORT-C since the "H" cannot be formed by adding the three sinks sequentially by closest connections to the growing tree.

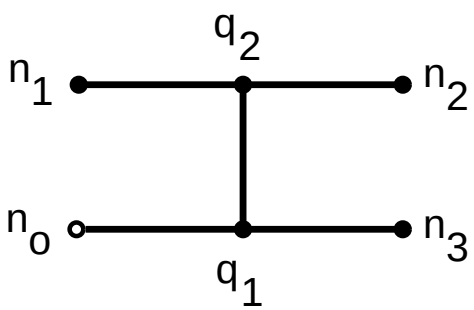

Figure 14: Example of a routing tree T which cannot be constructed by algorithm BB-SORT-C, but which satisfies the condition that each non-source node $v \in T$ makes a closest connection to each incident edge in $T \backslash v$.

B.3. Hanan Grid Proof for Steiner Nodes in T^*

We root any routing tree T at the source n_0 and for any node $v \in T$ define T_v to be the subtree of T rooted at v. We define a *segment* to be a contiguous set of straight edges in T which are either all horizontal or all vertical; a *maximal segment* (MS) is a segment not properly contained in any other segment. Let M be an MS in T. The node in M closest to n_0 on a source-sink path containing M is called the *entry point* to M. A segment containing all points in M to one side of a node v located on M is a *half segment* with respect to v, and a half segment with respect to the entry point of M is called a *branch*. A branch b is called a *branch off of* MS M' if M' contains b's entry point and is perpendicular to b. Note that any given segment, M, will divide the plane into two half-planes. If M does not contain n_0, then the half-plane containing the edge between M's entry point and its parent is called the *near side* of M (because it is "nearer" to the source), and the other half-plane is called the *far side* of M. (If M contains the source, the near and far sides of M can be labeled arbitrarily.) Branches off of M that are located on its near (resp. far) side are called *near* (resp. *far*) *branches*. In addition, a *sink* located on M is defined to be a *far branch* off of M if none of its

facts from elementary real analysis: (1) any concave function defined over an interval on the real line will be minimized at one of the boundary points of the interval; (2) multiplying any concave function by a positive constant produces another concave function; (3) the sum of any concave functions is also concave; and (4) any quadratic function of x having a negative coefficient for x^2 is concave in terms of x.

Consider the contribution made by the edge (x, a) to Elmore delay at various sinks n_j. First, consider the case of $n_j \in T_a$. Delay $t(n_j)$ is the sum of four functions: f_1 = delay from n_0 to p; f_2 = delay from p to x due to capacitance in $T^* \backslash b$; f_3 = delay from p to x due to capacitance in edge (b, q) and T_b; and f_4 = delay from x to n_j. Simple application of the Elmore formula for these four functions gives

$$f_1 = K_0 + K_1(a - x) \tag{2}$$

$$f_2 = x * (\frac{x}{2} + a - x + C_a + c - x + C_c) \tag{3}$$

$$f_3 = x * (b - q + C_b) \quad \text{if } x \leq q \tag{4}$$

$$f_3 = q * (b - q + C_b) \quad \text{if } x \geq q \tag{5}$$

$$f_4 = (a - x) * (\frac{a - x}{2} + C_a) + K_4 \tag{6}$$

where K_0, K_1, and K_4 are constants. To be precise, K_0 is the sum of resistance/capacitance products along the n_0—p path; K_1 is the sum of resistances from n_0 to p; and K_4 is the delay from a to n_j. We see that function f_1 is linear in x, while f_2 and f_4 are quadratic in x. The equation for $f_2 + f_4$ has a negative coefficient for x^2, and so $f_2 + f_4$ is concave. Function f_3 is linear and increasing for $x \leq q$ and remains constant for $x \geq q$; thus, f_3 is also concave in x. Consequently, $t(n_j) = f_1 + f_2 + f_4 + f_3$ is concave in x.

If $n_j \in T_c$ then f_1, f_2, and f_3 are identical to the case of $n_j \in T_a$. Function f_4 equals $(c-x)*(\frac{c-x}{2}+C_c)$. Again, $f_2 + f_4$ is concave in x, which implies that $t(n_j)$ is concave in x.

If $n_j \in T_b$ or $n_j = q$, then we can express Elmore delay to n_j in terms of three functions f_1, f_2 and f_3. The definitions of f_1 and f_2 are the same as for $n_j \in T_a$, and f_3 gives the delay from p to n_j due to capacitance in $T^* \backslash a$. The equation for f_1 is identical to that for $n_j \in T_a$, while f_3 is a constant in terms of x. Hence, f_1 and f_3 are concave. For f_2, we have

$$f_2 = x * (a - x + C_a) \quad \text{if } x \leq q \tag{7}$$

$$= q * (a - x + C_a) \quad \text{if } x \geq q \tag{8}$$

We recall that any continuous, piece-wise differential function of a real variable is concave as long as its first derivative is monotone decreasing. It is clear that this property holds for f_2, except possibly at $x = q$. Let f_2' be the derivative of f_2. Then for $x < q$, $f_2'(x) = a - 2x + C_a$, and for $x > q$, $f_2'(x) = -q$. Substituting q for x in these equations, we see that f_2' is indeed decreasing at $x = q$. Consequently, f_2 and $t(n_j)$ are concave in x.

Delay to any other sink in T^* is linear (and thus also concave) in x. Because f is a non-negative linear combination of concave functions, it is also concave. Therefore, over the interval $0 \leq x \leq c$ we know that f

an edge $e \in T \backslash v$ to which v is connected, then v is said to make a *closest connection* to e in T.

B.2. Proof of Closest Connections in T^*

Lemma B1: Let x be the parent of node $a \in T^*$, $a \neq n_0$. Then either $x = n_0$, or else x is located at the closest connection between a and each edge in $T^* \backslash a$ that is incident to a in T^*.

Proof: Consider a sink $n_i \in T^* \backslash a$ such that x is located on the n_0—n_i path in T^*. Assume that each L-shaped edge is oriented so that its interior points are as close to a as possible. We partition the continuous set of points along the n_0—n_i path into three sets, I, D, and $\{n_0\}$ (here, n_0 represents the point location of the source). If starting from a point p in the n_0—n_i path and moving along the path toward the source by a very small amount increases (decreases) the distance to a, then $p \in I$ ($p \in D$). Suppose $x \in D$. Then moving x closer to the source by a small amount along the n_0—n_i path will reduce both the overall tree cost and the path lengths to all sinks in T_a. Consequently, Elmore delay to all sinks will improve. Since T^* is optimal, this is impossible, and it must be that $x \notin D$.

Suppose then that $x \in I$. Let p and c be the endpoints of the maximal contiguous subset of I which contains x, with p topologically closer than c to n_0 (see Figure 13; note that $p \notin I$, but that there are points in I arbitrarily close to p). We will show that x must be located at c. We assume that there is exactly one node q between p and c in $T^* \backslash a$, and that q has degree 3 as shown in Figure 13. Our argument is easily extended when more nodes are present between p and c, and so we restrict our attention to the case shown in Figure 13.

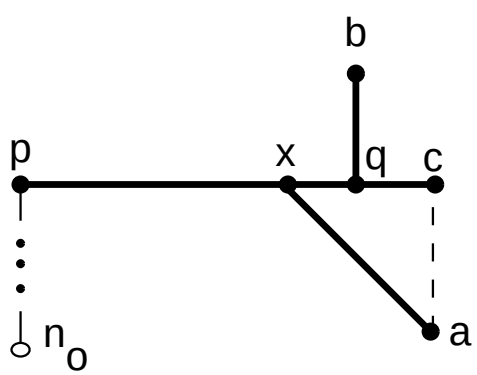

Figure 13: Proof of Lemma B1: we show that if x, the parent of a in T^*, is located on the maximal contiguous subset of I between locations p and c, then x must be located at c.

We will show that the delay function f is concave in terms of x for $0 \leq x \leq c$. For convenience, we overload x, a, b, c, and q to also represent the respective path lengths from p to these nodes or locations. Even though c is not necessarily a node in T^*, we will use T_c to represent the subtree of T^* below location c. Finally, recall that C_a, C_b, and C_c represent the tree capacitance in subtrees T_a, T_b, and T_c, respectively. We will show that the delay function f is concave in terms of x for $0 \leq x \leq c$. Our proof will invoke four

each node v' added to the queue in Line 6 of the template, and the total number of calls to Clean_Up is $O(k^2)$. $\square$

Appendix B: Optimal Steiner ERTs

For minimum-cost Steiner trees, the classic result of Hanan [17] restricts the choice of Steiner nodes to at most $k \cdot (k+1)$ points (the "Hanan grid") and enables finite branch-and-bound methods to determine optimal solutions. Here, we prove an analogous result for trees which optimize any linear combination of sink Elmore delays. Like Hanan, we show that any tree containing a Steiner node which is not a vertex in the Hanan grid can have its edges and Steiner nodes shifted to lie on the Hanan grid. However, we do not shift edges in the same way as Hanan (the edge shifts he uses can be suboptimal in terms of Elmore delay). Indeed, the result of, e.g., Lemma B1 below is obvious when minimizing tree cost, but requires a fairly involved proof when minimizing Elmore delay. Our development of the Hanan grid result becomes complete with the proof of Lemma B4 below. In Lemma B5, we extend our result to show that the branch-and-bound SORT-C method described in Section 4.4 returns the optimal delay Steiner tree.[21]

B.1. Definitions

We assume that all delays are defined in terms of Elmore delay. We seek to characterize the *optimal* Steiner tree over N, denoted by T^*, which minimizes the linear combination of sink delays $f = \sum_{i=1}^{k} \alpha_i \cdot t(n_i)$, with each $\alpha_i > 0$. (The case of some $\alpha_i = 0$ is effectively handled by setting these α_i to a small positive value.) We assume that T^* contains no Steiner nodes with degree < 3. For convenience, we normalize time and distance so that unit wire resistance and unit wire capacitance are both equal to one. We also consider a tree to be defined as a set of nodes and edges, so that the notations $v \in T$ for node v and $e \in T$ for edge e are well defined. An edge that is completely vertical or horizontal is called a *straight edge*; any other edge is called an *L-shaped edge*.

The *closest connection* between three nodes is the location of the single Steiner node in a minimum-cost Steiner tree over the three nodes. This location is unique and has coordinates given by the medians of the x- and the y- coordinates of the three nodes (if the minimum-cost Steiner tree is a chain, then the closest connection is the middle node). The *closest connection* between a node v and an edge e is the closest connection between v and the two endpoints of e. Assume that a Steiner tree T over N is rooted at n_0. We define $T \backslash v$ to be the tree induced by removing node v and all of its descendants from T, and then removing all degree-2 Steiner nodes from the resulting tree. We say that node $v \in T$ is *connected to* an edge $e \in T \backslash v$ if its parent node in T is located on edge e. If $parent(v)$ is located at the closest connection between v and

[21] The following clarifications should be made about our results. First, we allow the source pin n_0 to have degree > 4, which is in general physically impossible, but can be approximated by merging wires close to the source. Second, the optimal delay Steiner trees will not always be planar, as this is not required by our definition of an optimal-delay tree.

create an input tree that forces $\Omega(k^2)$ runtime [6] for GSR. However, GSR in practice seems to exhibit close to linear-time complexity, since multiple calls to procedure Clean_Up occur for very few nodes.

Theorem A2: For any input tree, (i) GSR returns a tree containing no V's and no U's, and (ii) GSR runs in $O(n^2)$ time in the worst case.

Proof: (i) Since GSR checks for V's and U's at each node in the tree, the output tree will contain a V or U only if GSR creates one at a node that has already been traversed. A new V or U can be produced at a node v only if the $n_0 - v$ path length is increased (which is impossible by Theorem 1) or if nodes are removed from the $n_0 - v$ path.

Removing a V at Line 8 will not introduce a new V or U at v_2 (in Figure 7), because the $n_0 - v_2$ path length is unchanged and a new Steiner point w_1 is added to this path. Removing a V will not introduce a V at v_3 either, because $v_1 w_1 v_3$ is not a V. A U may remain at v_3 after removing the V, but this will be detected later at Line 9.

Removing a U at v_4 in Figure 8 can only introduce a new V or U at w_2, v_4, or one of their descendants, because all other nodes have unchanged source-sink path lengths and no fewer Steiner nodes on their source-sink paths. The subroutine Clean_Up checks for V's and U's at w_2 and v_4, and recursive calls to Clean_Up will eventually terminate because a new V or U can be introduced only by reducing the number of nodes on the n_0–v_4 path.

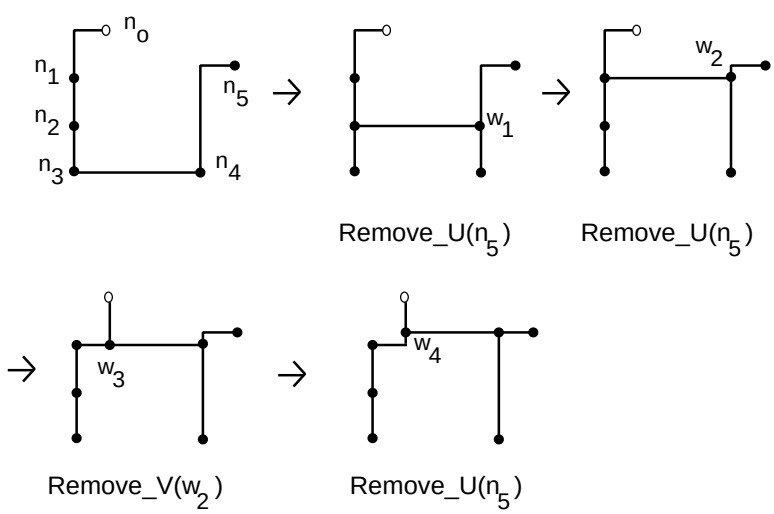

Figure 12: Example in which removing a U at n_5 requires two subsequent U-removals and a V-removal to complete the Clean_Up procedure.

Figure 12 shows how Clean_Up can require several recursive calls before terminating. However, for any node v', a call to Remove_U(v') will introduce a new V or U at v' or $parent(v')$ only if it reduces the number of nodes on the n_0–v' path. Because any Steiner tree connecting $k+1$ points can contain at most $2k$ nodes in total, there are $O(k)$ nodes on the n_0–v' path. Hence, at most $O(k)$ calls can be made to Clean_Up for

(iii) Assuming constant technology parameters[19], removing a U or a V can affect Elmore delay along a source-sink path in only three ways: a) changing the length of the path; b) changing tree capacitances along the path (i.e., increasing the wirelength of branches off from the path); and c) shifting tree capacitances along the path (changing where branches connect to the path). Removing a V will reduce some path lengths, reduce tree capacitances, and shift tree capacitances closer to the source, thereby reducing Elmore delay to all pins in the tree. Removing a U reduces path length to node v_4 in Figure 8 and shifts tree capacitance closer to the source for nodes v_2, v_3, and v_4. (For v_3, the capacitance that met the n_0-v_3 path at v_3 now meets the path at w_1 and w_2.) Removing 1-degree Steiner nodes reduces total wirelength, and thus reduces Elmore delay, while removal of 2-degree Steiner nodes leaves Elmore delay unchanged. Thus, GSR will not increase Elmore delay to any sink in the net. □

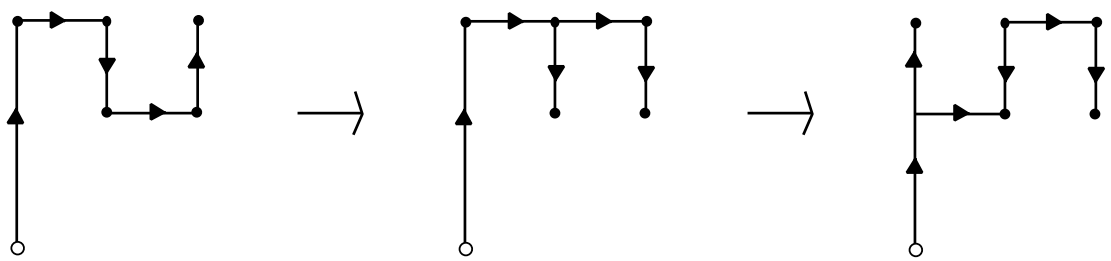

Figure 10: An example of a net with a source and five sinks for which processing the U's in a bottom-up order returns a tree with one remaining U.

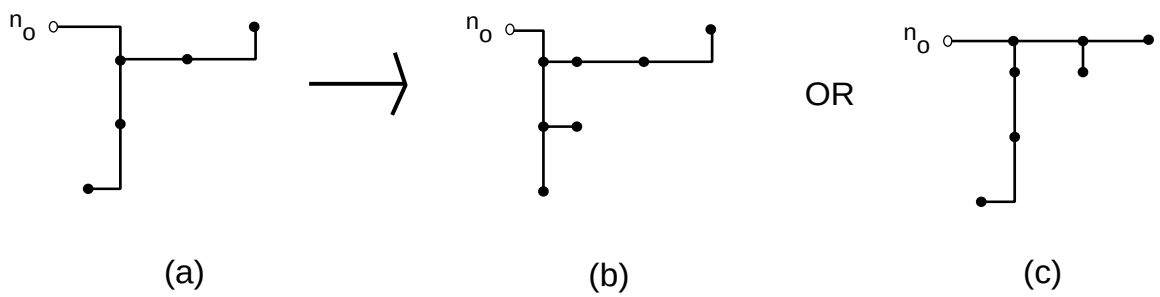

Figure 11: The GSR algorithm with input (a) can produce either tree (b) or tree (c), depending on the order in which the U's are processed.

We note that the order in which U's are removed from the tree is important. If the U's are processed in a bottom-up instead of a top-down order, then new U's can be introduced and the resulting tree may not have all of its U's removed, as can be seen in the example in Figure 10.[20] Furthermore, two different top-down orderings can produce different output trees (although both will have no remaining U's). An example is shown in Figure 11.

We now prove that GSR will remove all V's and all U's from any input tree, and that its worst-case time complexity is quadratic. Note that we have constructed a class of nets for which the 1-Steiner heuristic will

[19] I.e., including unit wire resistance, unit wire capacitance, driver resistance, and sink loading capacitances.

[20] By "bottom-up" we mean that each node is processed *after* all of its children in the tree, while a "top-down" ordering implies that each node is traversed *before* any of its children.

```
GSR Algorithm
Input: Steiner tree T with source n_0
Output: Steiner tree T with all U's removed
Remove all Steiner nodes of degree ≤ 2 from T;
Q ← {n_0};
While Q ≠ ∅
    v ← Dequeue(Q);
    For each node v' ∈ children(v) do
        Q ← Enqueue(v');
        If there is a V located at v'
            Remove_V(v')
        If there is a U located at v'
            Remove_U(v')
            Clean_Up(v')
Remove all Steiner nodes of degree ≤ 2 from T;

Subroutine Clean_Up(node: v')
If there is a V located at parent(v')
    Call Remove_V(parent(v'))
If there is a U located at v'
    Call Remove_U(v')
    Call Clean_Up(v')
Else
    If there is a U located at parent(v')
        Call Remove_U(parent(v'))
        Call Clean_Up(parent(v'))
```

Figure 9: Pseudo-code for the Global Slack Removal (GSR) algorithm. Local variables include a queue Q and nodes v and v'. We use $children(v)$ to denote the set nodes with children of v when the tree is rooted at n_0; $parent(v)$ denotes the parent of v in the rooted tree. The subroutine Remove_V(v) removes a V located at v as in Figure 7 and Remove_U(v) removes a U located at v as in Figure 8.

We now show that the tree returned by GSR dominates the input tree in terms of total tree cost, path length from the source to each sink, and Elmore delay at each sink. Let $cost(T)$ denote the cost of routing tree T.

Theorem A1: Given any tree T as input, GSR will return a tree T' such that (i) $cost(T') \leq cost(T)$; (ii) for each $i > 0$, the n_0-n_i path length in T' is less than or equal to the n_0-n_i path length in T; and (iii) the Elmore delay $t_{ED}(n_i)$ at each n_i in T' is less than or equal to the Elmore delay $t_{ED}(n_i)$ in T.

Proof: (i) Removing a V reduces cost in the routing tree; removing a U as in Figure 8 leaves tree cost unchanged; and by the triangle inequality the removal of a low-degree Steiner point will either reduce cost or leave it unchanged. These are the only operations in GSR that change the structure of the tree.

(ii) Removing a V does not affect source-sink path lengths; removing a U reduces the source-sink path length to the fourth node in the U (v_4 in Figure 8) and all of its descendants, and leaves all other source-sink path lengths unchanged; removing low-degree Steiner nodes does not affect source-sink path lengths.

tree to refer to any tree that can be output by the 1-Steiner algorithm.

Definition: A V is a subpath of *three* consecutive nodes on a root-leaf path in a routing tree such that the combined edge cost along the subpath is greater than the distance between its two end points (e.g., path v_1-v_3 in Figure 7(a)).

Definition: A U is a subpath of *four* consecutive nodes on a root-leaf path with edge cost greater than the distance between its end points (e.g., path v_1-v_4 in Figure 8(a)).

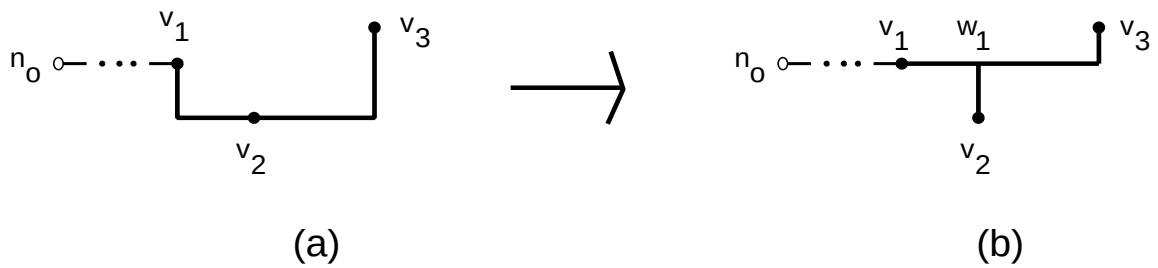

Figure 7: Removing a single "V" in the GSR Algorithm.

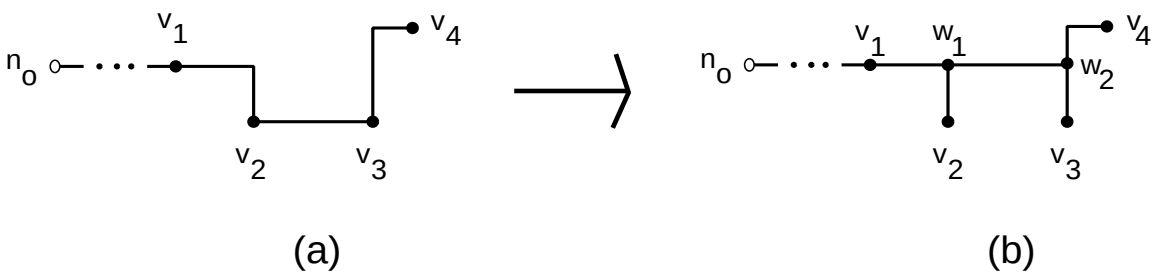

Figure 8: Removing a single "U" in the GSR Algorithm.

Note that the nodes in a V or a U can be either Steiner nodes or pins. A V can be removed from a routing tree by introducing a Steiner node which eliminates the overlap between the two adjacent edges, as in Figure 7(b). It is easy to see that, if a U (say $v_1v_2v_3v_4$) does not contain any V's, then its middle edge (v_2, v_3) must be either completely horizontal or vertical. Consequently, a U containing no V's can be removed by moving the middle edge and adding up to two new Steiner nodes as in Figure 8(b).

Figure 9 describes the GSR algorithm for removing V's and U's from any Steiner tree. We define a U (or V) to be *located at* node v if v is the node in U (or V) furthest topologically from the source. Three clarifying points should be noted. (i) GSR utilizes a queue Q which can be implemented arbitrarily as long as each node in the tree is processed before its children. In practice, a simple depth-first ordering suffices. (ii) The procedure Remove_U is invoked only for U's that do not contain any V's. Thus, it is executed as in Figure 8. (iii) All *low-degree* Steiner nodes of degree ≤ 2 are clearly superfluous and are removed since more U's can be found if these low-degree Steiner nodes are removed at the outset. Because each removal of a U can introduce additional low-degree Steiner nodes, they are removed again at the end of the algorithm.

delay objective to connect the remaining sinks). The SERT heuristic can also be applied before critical path information becomes available (an alternative is the ERT spanning tree heuristic, which has lower time complexity but does not introduce Steiner nodes). For nets on non-critical paths, minimizing wire length can take precedence over minimizing delay, hence traditional minimum-cost Steiner heuristics such as 1-Steiner [22] will are likely to be preferred.

Our heuristics which optimize Elmore delay directly are near-optimal in terms of SPICE-computed delay: we showed that our methods give Elmore delay that is nearly optimal, and we also showed that Elmore delay-optimal trees have nearly optimal SPICE delay. For spanning trees with five pins, we estimate that the optimal tree according to Elmore delay will be between 3% and 10% above SPICE-optimal, depending on the technology. Given that our SERT-C heuristic is between 0% and 5% above optimal in terms of Elmore delay for five-pin nets, we estimate that the SPICE delay suboptimality of our SERT-C heuristic ranges from 3% for MCM to about 12% for 0.5 μm and about 15% for 1.2 and 2.0 μm CMOS IC technologies.

Current work addresses interesting extensions of the CS-Steiner and ERT approaches to incorporate wiresizing, address general-cell layout with arbitrary routing region costs, and exploit the inherent parallelizability of our approaches. Finally, we leave as an open problem the reduction in time complexity of the ERT constructions.

6 Acknowledgements

We are grateful to the authors of [38] for use of their Two-Pole simulator code. Professor Ernest S. Kuh and Dr. Kamal Chaudhary provided many helpful comments regarding modeling and experimental methodology. Mr. Ashok Vittal of UC Santa Barbara gave many helpful comments and criticisms on the proofs in an earlier draft, and we also thank the anonymous referees for many detailed comments. Professors Jim Aylor and Michael Shur, and Mr. Sudhakar Muddu, all provided invaluable assistance with SPICE simulation methodology. Part of this work (ABK) was performed during a Spring 1993 sabbatical visit to UC Berkeley; support from NSF MIP-9117328 and the hospitality of Professor Ernest S. Kuh and his research group is gratefully acknowledged.

Appendix A: Global Slack Removal

Recall from Section 3.1 that *Global Slack Removal* (GSR) is an efficient post-processing enhancement to the CS-Steiner approach. The worst-case complexity of GSR is $O(k^2)$, although we believe the average-case complexity to be very close to $O(k)$. (GSR) is a linear-time postprocessing enhancement to the CS-Steiner approach. GSR shifts edges in the 1-Steiner output to maximize the monotonicity of all source-sink paths without any increase in total tree cost or Elmore delay to any sink. In what follows, we use the term *1-Steiner*

	IC1				IC2			
	$\|N\| = 5$		$\|N\| = 7$		$\|N\| = 5$		$\|N\| = 7$	
	delay	cost	delay	cost	delay	cost	delay	cost
SORT-C	1.0	1.111	1.0	1.112	1.0	1.161	1.0	1.158
SERT-C	1.042	1.046	1.083	1.047	1.049	1.120	1.114	1.106
(Std Err)	(.004)		(.006)		(.006)		(.009)	
1-Steiner	1.117	1.0	1.200	1.0	1.228	1.0	1.362	1.0
	IC3				MCM			
	$\|N\| = 5$		$\|N\| = 7$		$\|N\| = 5$		$\|N\| = 7$	
	delay	cost	delay	cost	delay	cost	delay	cost
SORT-C	1.0	1.175	1.0	1.165	1.0	1.296	1.0	1.262
SERT-C	1.046	1.140	1.112	1.112	1.000	1.296	1.001	1.256
(Std Err)	(.006)		(.010)		(.000)		(.0001)	
1-Steiner	1.275	1.0	1.429	1.0	1.455	1.0	1.634	1.0

Table 11: Empirical study of average sub-optimality of the SERT-C heuristic in terms of Elmore delay, using IC1, IC2, IC3 and MCM parameters. Simulations were run on 200 random nets for each net size. Delay is normalized to BB-SORT-C delay and tree cost is normalized to 1-Steiner cost. Standard errors for average SERT-C delays are shown in parentheses.

5 Conclusions

We have addressed a *critical-sink routing tree* (CSRT) formulation which arises when critical-path information becomes available during the timing-driven layout process. Two new classes of CSRT constructions are proposed: (i) the CS-Steiner method, which modifies a minimum Steiner tree to accommodate an identified critical sink, and (ii) the SERT-C method, which begins with a connection from the source to the critical sink and then grows a tree so as to minimize the increase in Elmore delay to the critical sink. Each of these algorithms is efficient, and offers very significant performance improvements over existing performance-driven routing tree constructions. We note that the greedy "Elmore routing tree" (ERT) approach underlying the SERT-C algorithm seems quite powerful. In particular, the approach encompasses a "generic" SERT Steiner router which outperforms all previous performance-driven routing algorithms in the literature. The ERT approach is also the first to *consistently*, and directly, optimize the Elmore delay formula itself, rather than an objective which heuristically abstracts Elmore delay. Since Elmore routing trees are efficiently computed, our approaches may lead to basic new utilities that can be integrated within existing performance-driven global routing codes. Assessments of the near-optimality of our Steiner constructions have led to a new characterization of Elmore-optimal Steiner trees, and to a new decomposition theorem for minimum-cost and minimum-Elmore delay Steiner trees; both of these results are of independent interest.

Which of our routing heuristics is most useful will depend on the application. The CS-Steiner heuristics H0 and HBest yield the smallest delay values for a single critical sink, but have high time complexity which may make them impractical for repeated wiring of large nets. Our SERT-C heuristic has time complexity of only $\Theta(k^2 logk)$ and is readily extended to the case of nets with multiple critical sinks (e.g., first apply SERT with its min-max delay objective to the critical sinks, then apply SERT-C with a weighted average

mines *optimal* trees with respect to any given linear combination of Elmore delays to critical sinks. We also present an entirely new "peeling decomposition" of any optimal Elmore delay Steiner tree into a sequence of subtrees, each of which adds a sink by a "closest connection" to some edge in the previous tree.

When the driver resistance r_d is very large, the optimal Elmore delay tree is a minimum-cost Steiner tree (recall Equation (1); also see [5]). As a consequence, our results extend very naturally to the well-studied problem of minimum-cost Steiner tree construction, and the restriction of Elmore-optimal Steiner nodes to the Hanan grid both generalizes and extends Hanan's original results. (As Hanan did for minimum-cost Steiner trees, we prove that every Steiner node in an Elmore-optimal tree is connected to one sink by a horizontal segment of edges, and to another sink by a vertical segment of edges. However, our techniques (Lemmas B1 - B4 in Appendix B) are much more powerful in order to address the optimality of the Steiner tree with respect to Elmore delay.) Our peeling decomposition and its extension to minimum-cost Steiner trees are of independent interest since they provide both a new characterization of, and a new means of generating, such trees.

Based on the results of Appendix B, we achieve a simple modification to our BBORT method which finds an optimal Steiner routing tree for any linear combination of Elmore delays to critical sinks. Rather than considering connections from each sink n_j outside the current tree to each sink n_i inside the tree as in BBORT, the *Branch-and-Bound* method for *Steiner Optimal Routing Trees with Critical Sinks* (BB-SORT-C) considers connections from n_j to each edge created when n_i was added to the tree. In other words, each node n_i already contained in T is replaced as a possible connection point by each of the edges created when n_i was added to the tree earlier. Again we use branch-and-bound pruning to reduce the complexity of the search.[17]

Table 11 compares Elmore delay for trees constructed by the SERT-C algorithm with optimal Elmore delay trees found by BB-SORT-C for each of the four technologies. The size of nets used in the comparison is limited to nets with six sinks (i.e., seven pins) because of the exponential time complexity of BB-SORT-C. For nets with seven pins, our results show that SERT-C achieves Elmore delay that is on average within 11.1% of optimal for the IC1 technology; results for IC2, IC3, and MCM parameters are very similar. The table also gives average tree costs for our constructions and the standard error of our estimate for the ratio between SERT-C and SORT-C delays.[18] We see that the SERT-C algorithm does not perform as well as the ERT algorithm in terms of nearness to optimality for the types of delay measures we have considered. Nevertheless, our results provide strong guidance for future efforts in performance-driven routing: even if future work improves the near-optimality of critical sink routing constructions, Table 11 shows that any future improvement in Elmore delay will be at most from 8% to 12% for nets with up to seven pins.

[17] Because we consider connections to up to three edges for each sink in the growing tree, our BB-SORT-C will introduce some redundancies in the tree topologies; we check for possible redundancies and prune the search at each redundant tree we find.

[18] Average running times for each 5-pin net (in seconds) are 0.006 (BB-SORT-C), 0.0004 (SERT-C), 0.0014 (SERT), and 0.0012 (1-Steiner). Average running times for 7-pin nets are 0.44 (BB-SORT-C), 0.0007 (SERT-C), 0.0055 (SERT) and 0.0035 (1-Steiner).

	IC1				IC2			
	$\|N\| = 5$		$\|N\| = 7$		$\|N\| = 5$		$\|N\| = 7$	
	delay	cost	delay	cost	delay	cost	delay	cost
ORT	1.0	1.103	1.0	1.133	1.0	1.140	1.0	1.175
ERT	1.007	1.104	1.017	1.142	1.010	1.159	1.022	1.215
(Std Err)	(.0015)		(.0021)		(.0017)		(.0022)	
SPT	1.085	1.290	1.130	1.395	1.058	1.290	1.096	1.395
MST	1.169	1.0	1.282	1.0	1.272	1.0	1.451	1.0

	IC3				MCM			
	$\|N\| = 5$		$\|N\| = 7$		$\|N\| = 5$		$\|N\| = 7$	
	delay	cost	delay	cost	delay	cost	delay	cost
ORT	1.0	1.146	1.0	1.190	1.0	1.432	1.0	1.547
ERT	1.011	1.172	1.027	1.252	1.009	1.585	1.024	1.892
(Std Err)	(.0018)		(.0025)		(.001)		(.0008)	
SPT	1.054	1.290	1.091	1.395	1.089	1.290	1.161	1.395
MST	1.311	1.0	1.499	1.0	1.894	1.0	2.457	1.0

Table 10: Elmore delays and wirelengths of various constructions using IC1, IC2, IC3 and MCM parameters. Simulations were run on 200 random nets for each net size. Tree cost is normalized to MST cost and delays are normalized to ORT delay. Standard errors for ERT delay are shown in parentheses.

of 0.21%, indicating a 95% confidence interval between 1.3% and 2.1% (i.e., an interval of within two times the standard error of the average).

Technology IC3 gives our worst results in terms of the optimality of ERTs. For the IC3 parameters and 7-pin nets, ERT gives an average value within 2.7% of ORT with a 95% confidence interval of between 2.2% and 3.2%. For MCM parameters, the Elmore-based ERT constructions are also very close to optimal: on average, they are within 2.4% of ORT delay for 7-pin nets. Finally, our tables compare the delays of the SPT construction with those of the ERT and MST solutions; the SPT outperforms the MST, but not the ERT, in terms of Elmore delay.

4.4 Elmore-Optimality of Steiner Tree Constructions

We have shown that our spanning tree constructions are nearly optimal when when we optimize the maximum Elmore delay over all sinks in the net. Because Steiner constructions give lower delay values than spanning trees in general, we close this section with a similar comparison for our SERT-C and SERT Steiner constructions. At first, this comparison appears very complicated because there are infinitely many possible locations for Steiner nodes. Indeed, while it is well-known that the result of Hanan [17] restricts the choice of Steiner nodes in a minimum-cost Steiner tree to at most $k \cdot (k+1)$ points, no such characterization has been established for a Steiner tree with optimal *Elmore* delay. In Appendix B, we present new theoretical results which restrict the choice of Steiner nodes in Elmore-optimal trees to exactly the same finite "Hanan grid" that contains the Steiner nodes of minimum-cost trees. This allows a finite algorithm which deter-

error of the average difference between ERT and ORT delays is equivalent to the standard error of average ERT delay.

avoids testing any topology more than once by requiring that sinks be added in the order of a breadth-first traversal of the tree (if two sinks are connected to the same parent node, then the sink with smaller index must be added to the tree first). It is easy to verify that according to this convention, any tree topology will imply a unique ordering of the sinks. Consequently, although BBORT tries all possible orderings of sinks, it calculates delay at most once for each tree topology. In Figure 6, lines 7 through 9 comprise the core of the branch-and-bound methodology used: if the delay in the current tree T' is greater than or equal to t_{min} (the current best-known delay for a complete tree), then procedure Add_Edges terminates and the algorithm backtracks. Otherwise, if T' is a complete tree, then t_{min} is set to the delay of T', or if T' is a partial tree, then Add_Edges recursively adds more edges to T'.

BBORT Algorithm
Input: signal net N with source $n_0 \in N$ **Output:** optimal-delay tree T_{opt} over N
1. $T = (V, E) = (\{n_0\}, \emptyset)$ 2. $t_{\min} = \infty$ 3. Call Add_Edges(T) 4. Output T_{opt}
Procedure Add_Edges(Tree: $T = (V, E)$) 5. **While** there exist $v \in V$ and $u \notin V$ such that $T' = (V \cup \{u\}, E \cup \{(u, v)\})$ is a new tree topology **Do** 6. Compute tree delay $t(T')$ 7. **If** $t(T') < t_{\min}$ **Then** 8. **If** $\|T'\| = \|N\|$ **Then** $T_{opt} = T'$; $t_{\min} = t(T')$ 9. **Else** Call Add_Edges(T')

Figure 6: Branch-and-Bound Optimal Routing Tree (BBORT) algorithm (recursive implementation).

To track all of the above simulation results, we have run BBORT trials on random sets of 200 nets for each of several net sizes. Our inputs are evaluated using the same four sets of technology parameters discussed previously. Table 10 compares Elmore delays of the BBORT and ERT constructions, as well as of the minimum spanning tree (MST) and shortest path tree (SPT) constructions, for each of the four technologies.[15] Delay for each tree is normalized to the ORT delay of the same net. Tree costs are similarly normalized to the MST cost of each net.

In the table, we see that ERTs over seven pins in the IC1 technology have an average maximum Elmore delay only 1.7% greater than optimal, while MSTs have average Elmore delay 28.2% greater than optimal. For smaller nets, ERTs are even better: for nets with five pins, ERT delays are only 0.7% above optimal on average, while MSTs are 16.9% above optimal. Our confidence in the average difference computed between ERTs and ORTs is very high: for instance, the 1.7% difference obtained for 7 pins has a standard error[16]

[15] The SPT construction is the tree which minimizes cost subject to each source/sink path having minimum length, i.e., it is a Steiner arborescence, or A-tree [11, 30].

[16] As used here, the term *standard error* is defined as follows. For a random variable X, let $\hat{X} = \sum_{i=1}^{n} X_i$ be an estimator for the expected value of X. The standard error of $\hat{X}$ is an estimate of its standard deviation over multiple sample sets, and is equal to the standard deviation of X divided by $\sqrt{n}$. Because delays are recorded as ratios to the ORT delay, the standard

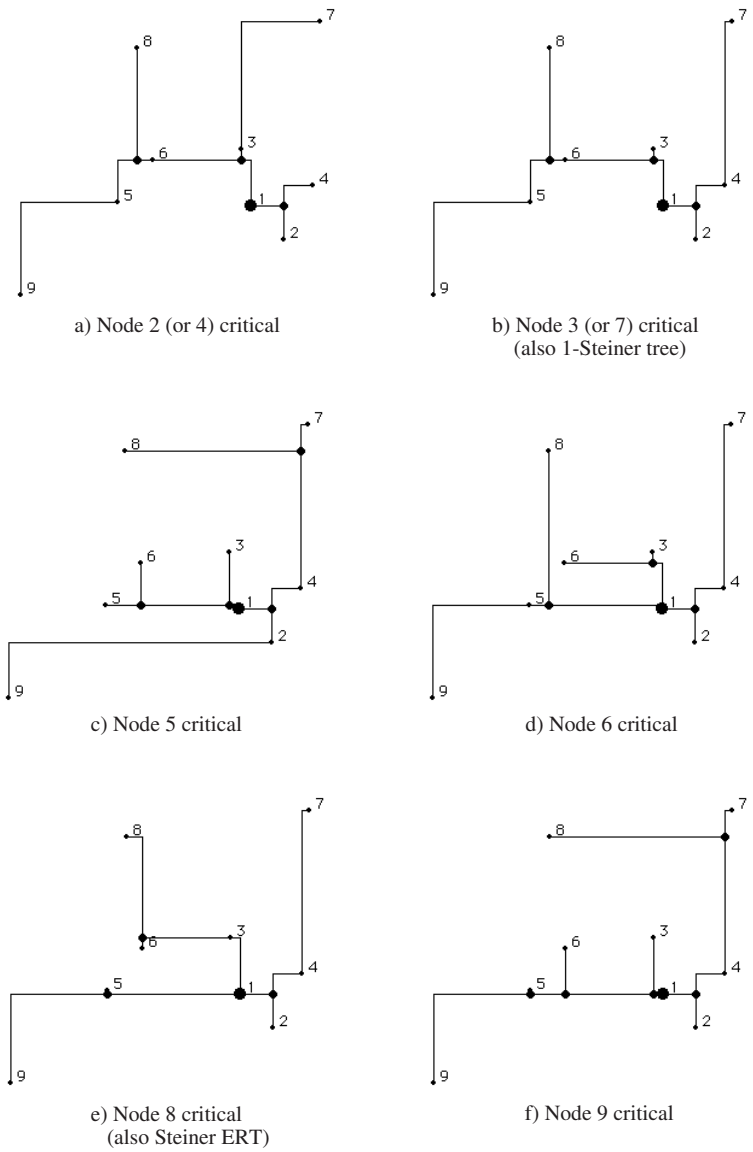

Figure 5: SERT-C tree constructions for a single 9-pin net, showing variation of solution with choice of critical sink n_c.

4.3 Elmore-Optimality of Spanning Tree Constructions

We have seen that the ERT constructions yield greatly improved signal delay when compared to previous methods. An obvious question is whether we still need to seek methods that better minimize Elmore delay. Thus, we have implemented a branch-and-bound algorithm which finds *optimal* generic routing trees according to Elmore delay. Starting with a trivial tree containing only the source pin, we incrementally add one edge at a time to the growing tree and evaluate the maximum sink delay. If this value exceeds the maximum sink delay in any *complete* candidate tree seen so far, we prune the search and backtrack to select a different edge at the previous step. A recursive implementation of this *Branch-and-Bound Optimal Routing Tree* (BBORT) search is shown in Figure 6. BBORT attempts to add sinks in all possible orders, but

and H0), and it does not require any simulator calls as does HBest.

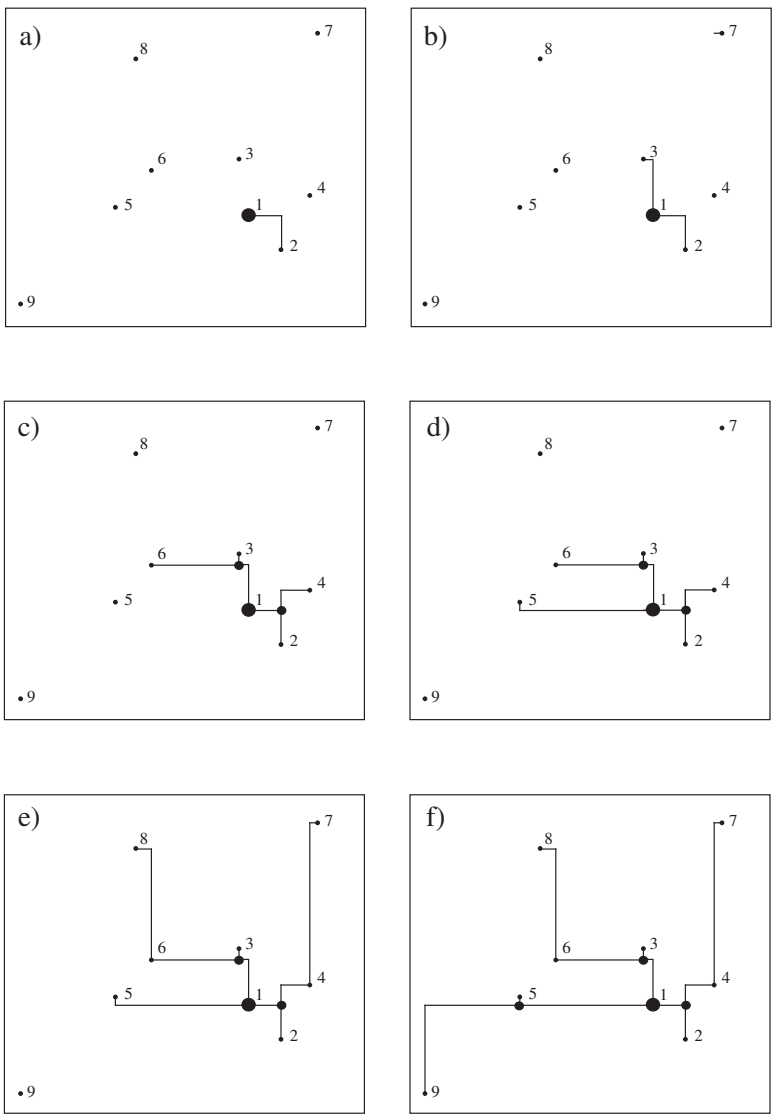

Figure 4: Example of the progressive SERT Steiner tree construction for a 9-terminal net using IC2 parameters. The source pin is labeled 1, and sinks are numbered in order of distance from the source.

Figures 4 and 5 illustrate the SERT and SERT-C algorithms for a 9-pin signal net using the IC2 technology parameters. Figure 4 shows the progressive growth of the SERT construction. Figure 5 contains the trees produced by SERT-C for the various choices of critical node. The tree constructed when n_c is node 3 or node 7 is also the 1-Steiner tree, and the tree constructed when n_c is node 8 is the same as the generic SERT result.

		IC1		IC2	
		$\|N\| = 5$	$\|N\| = 9$	$\|N\| = 5$	$\|N\| = 9$
Crit. Sink Delay	MST	0.645 ns	0.984 ns	0.395 ns	0.609 ns
	AHHK	.904	.837	.863	.770
	ERT	.879	.837	.804	.741
	1Stein	0.549 ns	0.848 ns	0.331 ns	0.520 ns
	SERT	.967	.884	.921	.806
	SERT-C	.947	.847	.870	.735
Max Delay	MST	0.758 ns	1.213 ns	0.485 ns	0.792 ns
	AHHK	.876	.805	.835	.747
	ERT	.855	.790	.786	.699
	1Stein	0.627 ns	1.028 ns	0.393 ns	0.664 ns
	SERT	.955	.853	.919	.780
	SERT-C	.970	.914	.962	.892
Ave WL	MST	1.64 cm	2.43 cm	1.64 cm	2.43 cm
	AHHK	1.16	1.09	1.16	1.09
	ERT	1.10	1.15	1.18	1.25
	1Stein	1.48 cm	2.18 cm	1.48 cm	2.18 cm
	SERT	1.06	1.09	1.11	1.18
	SERT-C	1.06	1.06	1.15	1.11

		IC3		MCM	
		$\|N\| = 5$	$\|N\| = 9$	$\|N\| = 5$	$\|N\| = 9$
Crit. Sink Delay	MST	0.262 ns	0.403 ns	2.82 ns	4.80 ns
	AHHK	.885	.749	.777	.608
	ERT	.782	.702	.472	.329
	1Stein	0.218 ns	0.342 ns	2.31 ns	4.09 ns
	SERT	.908	.781	.584	.384
	SERT-C	.839	.693	.567	.340
Max Delay	MST	0.326 ns	0.533 ns	3.86 ns	7.05 ns
	AHHK	.822	.730	.759	.632
	ERT	.764	.668	.544	.399
	1Stein	0.262 ns	0.444 ns	3.06 ns	5.92 ns
	SERT	.908	.759	.699	.481
	SERT-C	.954	.892	.859	.846
Ave WL	MST	1.64 cm	2.43 cm	16.4 cm	24.3 cm
	AHHK	1.16	1.09	1.04	1.07
	ERT	1.19	1.27	1.61	2.15
	1Stein	1.48 cm	2.18 cm	14.8 cm	21.8 cm
	SERT	1.13	1.22	1.66	2.27
	SERT-C	1.16	1.14	1.28	1.22

Table 9: Two-Pole simulation results for Elmore routing tree variants. Spanning ERT constructions are compared with MST and AHHK; Steiner SERT and SERT-C constructions are compared with 1-Steiner. All choices of critical sink are random, and all results are averaged over 100 random nets. MST and 1-Steiner results are reported in the physical units (nanoseconds or centimeters) while other results are reported as ratios to the corresponding MST or 1-Steiner results.

Since maximum sink delays still decrease, it is likely that overall skew in the routing tree will be reduced even when we treat the critical-sink formulation. Finally, we note that the SERT-C router produces very similar delays and costs compared to the HBest and H0 variants of CS-Steiner discussed in the previous subsection. However, SERT-C is more practical than HBest or H0 since it runs in $O(k^2 \log k)$ time (versus the $O(k^3)$ complexity of the best practical implementation of the 1-Steiner heuristic that is called by HBest

n_i as the critical sink.

Variants H0 and HBest significantly reduce delay to the critical sink, particularly in larger nets and for MCM interconnect technology where output driver and wire resistances are low. In other words, the simple strategy of connecting the critical node via a path with low branching factor is very successful for these cases. Of course, this strategy will produce larger routing cost.[13]

4.2 Elmore Routing Trees

We constructed Elmore routing trees for the same sets of random inputs used in the CS-Steiner experiments. Delay simulation results, again obtained using the Two-Pole simulator, are presented in the upper parts of Table 9. For comparison, the table includes data for the minimum spanning tree and AHHK tree [1] constructions.

Our results show that even as generic net-dependent routers, the ERT methods we propose are highly effective, beyond their relative efficiency and ease of implementation. For nets with 9 sinks, the spanning tree ERT construction reduces critical sink delay versus the MST construction by 16%, 26%, and 30% in the respective IC technologies and by 67% in the MCM technology. ERT also improves upon AHHK for most of the technologies, with reductions of 0% (IC1), 4% (IC2), 6% (IC3), and 46% (MCM). These results are particularly impressive because our AHHK data follows the experimental methodology in [1], which generates output trees for 21 different values of the c parameter and then chooses the best tree found for each signal net instance.[14]

The Steiner ERT variant also performs well as a generic high-performance router. For 9-pin nets, SERT improves critical sink delay versus the 1-Steiner routing by 19% and 62% for the IC2 and MCM technologies, respectively. The percentage reductions in maximum delay are somewhat greater for the IC technologies, but somewhat smaller for MCM interconnects. It should be noted that for the MCM technology, the ERT and SERT constructions tend to be star-like, producing tree costs significantly higher than those of the 1-Steiner construction. In practice, when delay is not an overriding concern, the user may recapture wirelength by simulating a larger output driver resistance.

Finally, even more significant reductions in delay can be achieved when a critical sink has been identified per the original CSRT formulation. The SERT-C algorithm improves over the SERT results by an *additional* reduction in delay at the critical sink of 5%, 7% and 6% for the three IC technologies, and 8% for MCM. Identification of a critical sink has clear advantages in terms of tree cost, particularly for MCM routing: the SERT-C trees have much less cost than the SERT outputs, while still improving the delay to the critical sink.

[13] Highly "star-like" topologies can possibly introduce other difficulties such as crossing wires, nodes with degree > 4, and capacitive coupling effects; these are not modeled by either SPICE or the Two-Pole simulator. Note that HBest uses calls to the Two-Pole simulator its delay analysis for candidate connections; see the definition of HBest in Section 3.1.

[14] According to [1], AHHK already achieves strong improvements over such other recent methods as shallow-light routing [10] or Steiner arborescences [11] when measured by the same Two-Pole simulation methodology. However, it should also be noted that delay reductions in practice will probably not attain exactly these magnitudes, partly because our modeling methodology cannot capture all of the device characteristics and delay effects related to the geometric embedding of our topologies.

4 Experimental Results

4.1 CS-Steiner Trees

		IC1		IC2	
		$\|N\| = 5$	$\|N\| = 9$	$\|N\| = 5$	$\|N\| = 9$
Critical Sink Delay	1Stein	0.549 ns	0.848 ns	0.331 ns	0.520 ns
	1Stein+GSR	.978	.964	.970	.954
	H0+GSR	.980	.824	.876	.700
	H1+GSR	.960	.883	.934	.827
	HBest+GSR	.929	.817	.867	.721
Ave WL	1Stein	1.48 cm	2.18 cm	1.48 cm	2.18 cm
	H0+GSR	1.29	1.22	1.29	1.22
	H1+GSR	1.04	1.06	1.04	1.06
	HBest+GSR	1.07	1.11	1.10	1.12

		IC3		MCM	
		$\|N\| = 5$	$\|N\| = 9$	$\|N\| = 5$	$\|N\| = 9$
Critical Sink Delay	1Stein	0.218 ns	0.342 ns	2.31 ns	4.09 ns
	1Stein+GSR	.968	.950	.952	.927
	H0+GSR	.849	.664	.550	.333
	H1+GSR	.922	.810	.857	.665
	HBest+GSR	.844	.693	.593	.340
Ave WL	1Stein	1.48 cm	2.18 cm	14.8 cm	21.8 cm
	H0+GSR	1.29	1.22	1.29	1.22
	H1+GSR	1.04	1.06	1.04	1.06
	HBest+GSR	1.11	1.12	1.22	1.21

Table 8: Two-Pole simulation results comparing CS-Steiner trees with 1-Steiner heuristic trees. Each entry corresponds to an average over delay computations for random critical sinks in each of 100 different random signal nets. 1-Steiner results are reported in the physical units (nanoseconds or centimeters) while other results are reported as ratios to the corresponding 1-Steiner results. Note that 1-Steiner and 1-Steiner plus GSR always produced nearly identical average costs.

We implemented each of the CS-Steiner variants H0, H1 and HBest, along with the 1-Steiner algorithm [22], using C on a Sun SPARC1 ELC workstation, and ran these algorithms on random 4- and 8-sink inputs.[12] We also applied our GSR post-processing algorithm (denoted as +GSR) to 1-Steiner and each of the CS-Steiner variants. Our inputs correspond to the four distinct technologies described in Table 1.

Table 8 gives delay and tree cost (WL) results and comparisons. The delays at all sink nodes correspond to 50% rise times estimated using the Two-Pole simulator [37] [38]. Each entry in Table 8 represents an average taken over every sink node in 50 random point sets. We emphasize that the 1-Steiner algorithm (or the BRBC, AHHK, etc. methods), being net-oriented, will return the same tree for a given sink set no matter which sink happens to be critical; the delays at the sinks n_i are in some sense "generic". In contrast, each of the three CS-Steiner variants can return a different tree for each choice of critical sink in the same net. Thus, for each variant we report the delay at n_i in the *specific* tree corresponding to identification of

[12] Results for 16-sink inputs have been reported in preliminary form, e.g., [6]. While such large inputs magnify the effect of our new methods, in practice most signal nets will be within the size range that we now discuss.

strategies are only motivated by Elmore delay and whose solution quality may therefore be more sensitive to the technology, the input instance, and the choice of parameters.

Time complexities for our ERT variants are analyzed as follows.

Observation 1: The SERT-C algorithm can be implemented in $O(k^2 \log k)$ time.

Proof: The effect on delay $t_{ED}(n_c)$ of inserting a new edge (u, w) into T arises only in the C_j terms in Equation (1), and is an *additive constant* no matter when (u, w) is added into the tree. Initially, we compute the best connection from each non-critical sink to the tree containing only edge (n_0, n_c). For each new sink added, at most three new edges will be inserted into the tree. In constant time, we can calculate the effects of connections from a given sink outside T to these three new edges (all previously computed effects remain unchanged and need not be recomputed). We can insert the new delay effects into a priority queue for each $u \notin V$ in $O(\log k)$ time and also retrieve the current minimum-cost connection for v in $O(\log k)$ time. Thus, each pass through the **while** loop of Figure 3 can be accomplished in $O(k \log k)$ time, giving an overall time complexity of $O(k^2 \log k)$. □

Observation 2: The ERT spanning tree algorithm can be implemented in $O(k^3)$ time, assuming constant unit wire resistance, unit wire capacitance, and sink capacitances.

Proof: The result follows from a simple observation: if a new tree edge incident to sink $u \in V$ (Line 3 of Figure 3) minimizes the maximum Elmore delay $max_i t_{ED}(n_i)$, it must connect u to the sink $v \notin V$ that is closest to u. Thus, at each pass through the **while** loop, we simply compute the shortest "outside connection" for each node in V, i.e., every possible u, in $O(k^2)$ time. We then add each of the $O(k)$ shortest outside connections to T in turn. Evaluating the Elmore delays at all sinks in each of the resulting trees requires $O(k)$ time per tree. Hence, each pass through the **while** loop requires $O(k^2)$ time, and this yields the $O(k^3)$ complexity result.[11] □

In practice, the complexity of the ERT algorithm will be transparent to the user, since k is typically small (e.g., our runtimes for the problem sizes discussed here are $O(10^{-2})$ seconds on Sun SPARC1 hardware; see also the Footnote 18 below). We know of no implementation of the SERT algorithm that is faster than $O(k^4)$. Intuitively, the difficulty is that (i) in Line 3 we must always consider $\Theta(k^2)$ Steiner connections, and (ii) the connection which minimizes $max_i t_{ED}(n_i)$ in Line 3 may not be the best one from the "perspective" of any individual sink in N or edge in T. Thus, we currently have a rather interesting situation where the CSRT problem formulation leads to an algorithm (SERT-C) that enjoys nearly quadratic speedup over the generic Steiner computation (SERT).

[11] Again, we note the fundamental difference between the ERT approach and the method of [28]: while [28] must add the single sink that is closest to the existing tree, the ERT algorithm identifies both a new sink and its connection such that Elmore delay is minimized.

net is spanned by the growing tree.[10] Note that greedy approach of the ERT algorithm can be generalized to any delay model by applying the appropriate estimator in Line 3 of Figure 3.

ERT Algorithm
Input: signal net N with source $n_0 \in N$
Output: routing tree T over N
1. $T = (V, E) = (\{n_0\}, \emptyset)$
2. **While** $\|V\| < \|N\|$ **do**
3. Find $u \in V$ and $v \notin V$ which minimize the maximum Elmore delay from n_0 to any sink in the tree $(V \cup \{v\}, E \cup \{(u, v)\})$
4. $V = V \cup \{v\}$
5. $E = E \cup \{(u, v)\}$
6. **Output** resulting spanning tree $T = (V, E)$

Figure 3: The ERT Algorithm: direct incorporation of the Elmore delay formula into a heuristic routing tree construction.

We apply the ERT approach to Steiner routing by allowing the new pin to connect to an *edge* (or the source) of the existing tree, possibly inducing a Steiner node on this edge at the point that is closest to the new pin. In this way, the number of ways a pin outside the current tree can be added at each iteration is at most the number of edges in the current tree plus one (i.e., a connection to the source). Note that the orientation of each "L-shaped" edge remains flexible until a Steiner node is placed on it.

- For *generic* performance-driven routing, our *Steiner Elmore routing tree* (SERT) algorithm iteratively finds $u \notin V$, $(v, v') \in E$, so that connecting u to the closest point on edge (v, v') minimizes the maximum source-sink Elmore delay in the resulting tree.
- To address *critical-sink* routing, our *Steiner Elmore routing tree with identified critical sink* (SERT-C) algorithm begins with a tree containing the single edge (n_0, n_c) in Line 1 of Figure 3, then continues as in the SERT algorithm, except that we minimize $t_{ED}(n_c)$ rather than the maximum delay to all sinks.

While CS-Steiner began with a minimum-cost Steiner tree and heuristically perturbed it to improve $t(n_c)$, SERT-C uses the opposite approach of starting with the required n_0-n_c connection and growing the routing tree while keeping $t_{ED}(n_c)$ as small as possible. Again, we note that SERT-C offers a consistent, *direct* incorporation of Elmore delay within its construction, in contrast to heuristics whose objectives or

[10] Our approach should be distinguished from the method of Prasitjutrakul and Kubitz [28] cited above, wherein A* heuristic search and the actual Elmore delay formula are used in a performance-driven routing tree construction. Like our method, [28] grows a routing tree over a net N starting from the source n_0; they perform A* search of a routing graph (e.g., in building-block design) to find the Elmore delay-optimal Steiner connection from the existing tree to a new sink. However, *the choice of this new sink is forced:* the algorithm always adds the sink that is closest (by Manhattan distance) to the existing tree, and thus falls into the standard pitfall of ignoring the underlying delay criterion. The effect of this difference is apparent in the ERT ordering of added nodes in Figure 4 of Section 4 below. Indeed, the method of [28] can yield Elmore delays substantially larger than those of ERT: given a very tall, "hairpin"-like version of Figure 1a with many sinks very closely spaced along the entire hairpin path, [28] forces the sinks to be added into the tree according to the path order (starting from the source n_0 at the lower left), yielding an obviously poor solution.

H1: The direct connection in Line 2 consists of the shortest possible wire that can join n_c to T_0, subject to the monotone path constraint.

HBest: Accomplish Line 2 by trying all shortest connections from n_c to edges in T_0, as well as from n_c to n_0; perform timing analysis on each of these routing trees, and return the tree with lowest delay at n_c.

The complexity of these variants is dominated by the construction of T_0 in Line 1 (or possibly by the simulator calls in HBest).

We enhance the CS-Steiner construction via an efficient *Global Slack Removal* (GSR) postprocessing algorithm. GSR [6] is similar to the method developed independently by Chen and Sarrafzadeh in [8], which also removes "U's" from interconnections. However, the objective of GSR is not to reduce tree cost (which is already effectively minimized by the 1-Steiner algorithm) but rather to maximize the monotonicity of all source-sink paths and reduce Elmore delay to all sinks. GSR accomplishes this without increasing overall tree cost. For expository reasons, we defer formal description of GSR, along with its proofs of correctness, to Appendix A.

3.2 Elmore Routing Trees

From the discussion of Section 2.2, we see that current routing objectives such as minimum tree cost, bounded tree radius, or prescribed cost-radius balance have often been *motivated* by the Elmore model. However, such objectives are abstractions: they do not directly optimize Elmore delay. Thus, the effectiveness of a given objective often depends on the prevailing technology, on the particular distribution of sink locations for a given signal net, and on the user's ability to find the parameter value (e.g., ϵ in the BRBC algorithm [10], or c in the AHHK algorithm [1]) which will yield a good solution for the particular input.

In this subsection, we depart from the abstraction inherent in "minimum cost" or "bounded radius" objectives, and propose a new greedy *Elmore routing tree* (ERT) approach which optimizes Elmore delay *directly* as the routing tree is constructed. The ERT approach is efficient, since Elmore delay at all nodes of a routing tree can be evaluated in linear time (see Footnote 1 above). Based on the performance results in Section 4 for both critical-sink and "generic" performance-driven routing formulations, we believe that the ERT approach, which we have embodied as the ERT, SERT and SERT-C algorithms described below, offers a basic new tool for VLSI routing.

The basic ERT approach is embodied in our Elmore routing tree (ERT) algorithm[9] for spanning trees (Figure 3), which is analogous to Prim's minimum spanning tree construction [29]: starting with a trivial tree containing only the source, we iteratively find a pin n_i in the tree and a sink n_j outside the tree so that adding edge (n_i, n_j) yields a tree with minimum Elmore delay. The construction terminates when the entire

[9] Note that "ERT approach" refers to our basic concept of optimizing Elmore delay directly via a greedy heuristic. In contrast, the "ERT algorithm" is simply one of many possible implementations of the ERT approach: specifically, it is a greedy spanning tree construction.

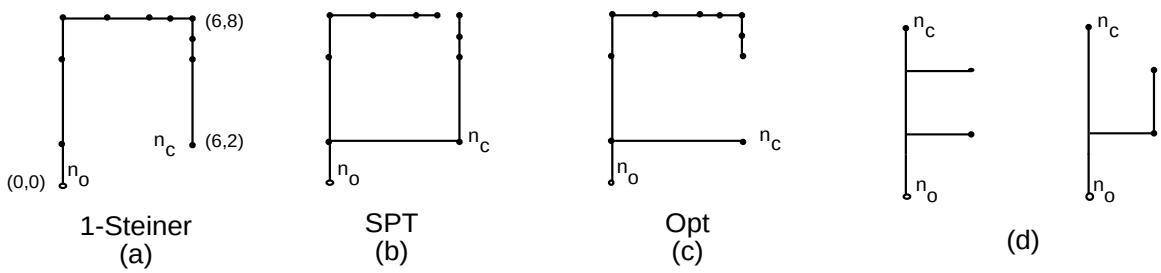

Figure 1: Parts (a)-(c): optimal Steiner tree (cost 2.0 cm, $t(n_c) = 3.34$ ns); minimum cost shortest-paths tree (cost 2.5 cm, $t(n_c) = 2.26$ ns); and optimal-delay tree (cost 2.2 cm, $t(n_c) = 1.67$ ns) for the same sink set. Coordinates shown are in mm, and the 1.2μ IC2 technology parameters (Table 1) were used with the Two-Pole simulator and 90% rise time delay criterion. Part (d): two distinct minimum-cost SPT solutions for a set of three sinks.

3 Two Classes of CSRT Heuristics

3.1 The CS-Steiner Approach

Given the observations above, we may characterize the optimal CSRT solution in Figure 1(c) as one which minimizes total tree cost, *subject to the path from n_0 to n_c being monotone* (i.e., of minimum possible length). This simultaneous consideration of radius and cost parameters recalls the motivations in [1] [9] [10], but here the tradeoff is formulated with respect to the critical sink n_c. We thus obtain our *CS-Steiner* heuristic for the CSRT problem (Figure 2).

CS-Steiner Algorithm
Input: signal net N; source $n_0 \in N$; identified critical sink $n_c \in N$
Output: heuristic CSRT solution T
1. Construct heuristic minimum-cost tree T_0 over $N - n_c$. 2. Form T by adding a *direct connection* from n_c to T_0, i.e., such that the n_0-n_c path in T is monotone.

Figure 2: The CS-Steiner heuristic.

The idea behind CS-Steiner is simple: construct a minimum-cost Steiner routing tree as usual, then "fix" the tree to reflect an identified critical sink. Since the algorithm template is quite general, we have examined a number of CS-Steiner variants. All of our variants use the 1-Steiner heuristic of Kahng and Robins [22] to construct the initial tree T_0 in Line 1. Section 4 reports results for the following three variants:[8]

H0: The direct connection in Line 2 consists of a single wire from n_c to n_0.

[8] We also studied two additional variants. **Variant H2** modifies Line 1 of CS-Steiner so that the initial heuristic tree T_0 is constructed over the entire net N. H2 then deletes the edge which lies directly above n_c when we root T_0 at n_0, and rejoins (the component containing) n_c to (the component containing) n_0 using a shortest possible wire from n_c, as in variant H1. **Variant H3** performs Lines 1 and 2 simultaneously by executing the 1-Steiner algorithm subject to a "maintaining monotone feasibility" constraint. In other words, we iteratively choose a Steiner point which minimizes the sum of the tree cost and the cost of any needed direct connection from n_c to n_0. The direct connection from n_c requires that there exist a monotone path through the "bounding boxes" of the edges in the path to n_0. Intuitively, this favors initial choice of Steiner nodes along some monotone path from n_0 and n_c, since such nodes will most rapidly reduce the marginal cost of adding the direct n_c-n_0 connection. The H2 and H3 variants yielded delays that were inferior to those of H0, H1 and HBest.

critical-sink and maximum sink delay criteria.

2.2 Intuitions from Elmore Delay

Because of its fidelity to SPICE-computed delay, Elmore delay is a good performance objective for constructing high-performance routing trees. Furthermore, the simplicity of the Elmore delay formula (1) allows us to intuit heuristics which effectively minimize delay.

Since r_{e_v} and c_{e_v} are usually proportional to the length of edge e_v, we see that $t_{ED}(n_i)$ has a quadratic relationship to the length of the n_0-n_i path, suggesting a min-radius criterion. However, the C_j term implies that Elmore delay is also linear in the total edge length of the tree which lies outside the n_0-n_i path, suggesting a min-cost criterion. The relative size of the driver resistance r_d heavily influences the optimal routing topology: if r_d is large, the optimal routing tree (ORT) is a minimum cost tree; as r_d decreases, the ORT tends to resemble a "star" topology. The size of r_d relative to unit wire resistance is a "resistance ratio" [5] that captures the technology vis-a-vis routing tree design. Values of the resistance ratio are larger for current-generation CMOS, but tend to decrease in MCM substrate and some submicron CMOS IC interconnects (Table 1).

In Figure 1, we show a signal net N with identified critical sink n_c, along with three routing trees: (a) the 1-Steiner tree, (b) a minimum-cost SPT, and (c) the optimal CSRT with respect to critical sink n_c. Based on this example, the example of Figure 1(d), and Equation (1), we make the following observations.

- The minimum cost solution (a) has large delay to the critical sink n_c due to the long source-sink path.
- However, requiring a monotone path to *every* sink, as in the SPT (b) or a Steiner arborescence [11, 30], can result in large tree capacitance which again leads to large delay at n_c.
- The optimal CSRT construction (c) illustrates the dependence of routing topology on the choice of critical sink, and reflects both the minimum-cost and the SPT solutions.
- Finally, Equation (1) implies that the number of Steiner points in the n_0-n_c path should be minimized, and the Steiner points "shifted" toward n_0 (i.e., branches off of the n_0-n_c path should occur as close to the source as possible). Figure 1(d) shows two trees which are both shortest-path trees and minimum Steiner trees, yet the rightmost tree has less signal delay at n_c.

	IC1		IC2		IC3		MCM	
Rank	CS	Max	CS	Max	CS	Max	CS	Max
1	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
2	1.006	1.049	1.003	1.046	1.002	1.050	1.001	1.044
3	1.011	1.087	1.005	1.088	1.005	1.089	1.001	1.108
4	1.014	1.120	1.006	1.128	1.006	1.133	1.002	1.189
5	1.016	1.154	1.007	1.153	1.006	1.158	1.003	1.250
6	1.017	1.180	1.007	1.184	1.006	1.191	1.004	1.297
7	1.026	1.201	1.012	1.215	1.007	1.221	1.005	1.363
8	1.040	1.227	1.021	1.243	1.014	1.247	1.005	1.419
9	1.074	1.253	1.046	1.273	1.036	1.279	1.014	1.474
10	1.160	1.282	1.138	1.311	1.120	1.322	1.047	1.531
11	1.180	1.306	1.155	1.336	1.134	1.345	1.049	1.594
12	1.224	1.330	1.207	1.371	1.182	1.380	1.058	1.652
13	1.246	1.353	1.218	1.399	1.191	1.417	1.060	1.713
14	1.288	1.387	1.254	1.436	1.225	1.449	1.064	1.763
15	1.306	1.417	1.269	1.468	1.233	1.483	1.066	1.831
16	1.327	1.436	1.309	1.495	1.283	1.515	1.103	1.886
17	1.351	1.491	1.344	1.572	1.326	1.595	1.427	1.987
18	1.380	1.517	1.376	1.600	1.354	1.629	1.431	2.039
19	1.417	1.554	1.427	1.641	1.413	1.672	1.475	2.079
124	8.04	5.46	10.36	6.51	10.81	6.83	18.34	10.41

Table 6: Average SPICE delay ratios for the top 20 topologies ranked according to SPICE for $|N| = 5$. Values are averaged over 50 random nets and normalized to the average delay of the best topology. Also included is the average ratio for the *worst* topology (rank 124).

Table 6 implies approximate suboptimality ranging between 0.6% for MCM and 2.4% for IC3.

	Critical Sink Delay		Maximum Delay	
Technology	$\|N\| = 4$	$\|N\| = 5$	$\|N\| = 4$	$\|N\| = 5$
IC1	1.029	1.099	1.009	1.001
IC2	1.039	1.096	1.005	1.002
IC3	1.038	1.078	1.013	1.002
MCM	1.019	1.031	1.001	1.001

Table 7: Average SPICE suboptimality of Elmore delay as measured by the ratio between the average SPICE delays of the Elmore-optimal and SPICE-optimal topologies.

A more direct measure of the suboptimality of Elmore delay is to compare SPICE delays of the Elmore-optimal and SPICE-optimal topologies. Table 7 shows averages of this measure of suboptimality for both the critical sink and the maximum delay criteria. For critical sink delay and $|N| = 5$, the average SPICE suboptimality of the Elmore-optimal topology is between 3.1% for MCM and 9.9% for IC1. Moreover, the Elmore-optimal topologies are closer to SPICE-optimal for the IC3 and MCM technologies, which have lower resistance ratios. We believe that the estimates of Elmore suboptimality for critical-sink delay in Table 7 are larger than those inferred above from Tables 4-6, due to the convexity of the relationship between SPICE rank and average SPICE delay (cf. Table 6). For maximum delay, Table 7 indicates that minimizing Elmore delay very nearly minimizes SPICE delay, with suboptimality of between 0.1% and 0.2% for the optimal Elmore topology. Thus, while the *accuracy* of Elmore delay has many dependencies on technology and is particularly weak for critical-sink delay, we find that the *fidelity* of Elmore delay is quite strong for both the

	Topologies	Linear vs SPICE		Elmore vs SPICE	
		$\|N\| = 4$	$\|N\| = 5$	$\|N\| = 4$	$\|N\| = 5$
IC1	Best	2.30	16.3	0.54	5.9
	5 Best	2.52	18.1	1.02	7.2
	All	2.43	17.0	0.92	8.0
IC2	Best	2.52	19.4	0.58	6.4
	5 Best	2.66	20.2	0.99	7.2
	All	2.44	16.9	0.94	7.9
IC3	Best	2.60	19.8	0.58	5.6
	5 Best	2.68	20.9	0.93	6.5
	All	2.43	16.5	0.93	7.7
MCM	Best	3.04	24.6	0.72	5.1
	5 Best	2.81	24.4	0.89	4.7
	All	2.33	15.7	0.89	7.1

Table 4: Average difference in rankings of topologies, in terms of 50% delay to a given random **critical sink** in each net. The sample consists of 50 random nets of each cardinality, and the 50% rise time delay criterion was used. The total number of topologies for each net is $4^{(4-2)} = 16$ for $|N| = 4$, and $5^{(5-2)} = 125$ for $|N| = 5$.

	Topologies	Linear vs SPICE		Elmore vs SPICE	
		$\|N\| = 4$	$\|N\| = 5$	$\|N\| = 4$	$\|N\| = 5$
IC1	Best	0.50	2.06	0.38	0.10
	5 Best	0.66	2.78	0.71	0.47
	All	0.94	7.74	0.65	1.39
IC2	Best	0.40	2.26	0.16	0.20
	5 Best	0.68	2.61	0.51	0.53
	All	0.87	7.02	0.43	1.24
IC3	Best	0.64	2.40	0.48	0.20
	5 Best	0.87	2.59	0.52	0.44
	All	1.04	6.96	0.60	1.22
MCM	Best	0.70	4.56	0.14	0.08
	5 Best	0.71	3.15	0.11	0.22
	All	1.02	7.01	0.16	0.86

Table 5: Average difference in rankings of topologies in terms of **maximum** delay over all sinks. The sample consists of 50 random nets of each cardinality, and the 50% rise time delay criterion was used.

have similarly established the fidelity of Elmore delay for circuit design: they plotted Elmore- versus SPICE-computed delays for a suite of 209 different place/route solutions of the same ripple-carry adder circuit, and also found a very high correlation between the two delay measures.

Table 6 shows the average increase in SPICE delay from optimal for the 20 top-ranking topologies, i.e., the 20 lowest SPICE delays for $|N| = 5$. For IC2, the average distance of 6.4 rank positions for the optimal critical sink Elmore topology implies an expected difference of approximately 1.6% in actual SPICE-computed delay; for IC3 a distance of 5.6 rank positions implies approximately 0.7% delay suboptimality; and for MCM a difference of 5.1 rank positions implies 0.4% delay suboptimality. For maximum sink delay,

Accuracy of Linear, Elmore and Two-Pole Delay Estimates for Maximum Sink Delay					
		$\|N\| = 4$		$\|N\| = 7$	
	Delay Ratio	average	std dev	average	std dev
IC1	SPICE/Linear	–	11.0%	–	11.4%
	SPICE/Elmore	0.79	1.9%	0.79	1.7%
	SPICE/2-Pole	1.39	1.7%	1.39	1.7%
	2-Pole/Elmore	0.568	0.34%	0.567	0.19%
IC2	SPICE/Linear	–	12.5%	–	12.9%
	SPICE/Elmore	0.83	4.1%	0.81	2.6%
	SPICE/2-Pole	1.44	3.6%	1.42	2.4%
	2-Pole/Elmore	0.572	0.81%	0.568	0.36%
IC3	SPICE/Linear	–	13.0%	–	13.6%
	SPICE/Elmore	0.87	6.0%	0.83	3.6%
	SPICE/2-Pole	1.51	5.1%	1.46	3.4%
	2-Pole/Elmore	0.574	1.32%	0.569	0.44%
MCM	SPICE/Linear	–	25.8%	–	23.3%
	SPICE/Elmore	0.79	2.3%	0.79	2.0%
	SPICE/2-Pole	1.39	2.1%	1.39	1.9%
	2-Pole/Elmore	0.568	0.96%	0.566	0.44%

Table 3: Accuracy of the Linear, Elmore and Two-Pole estimators for *maximum* sink delay. See previous table for explanatory notes.

solutions: we first rank *all* spanning tree topologies[6] by the given delay model, then rank the topologies again by SPICE delay, and then find the average of the absolute value of the difference between the two rankings for each topology. This measure of fidelity corresponds to a standard rank-ordering technique used in the social sciences [2]. We have run simulations to estimate this measure of fidelity for nets of size 4 and 5 using the linear and Elmore delay estimators and each of the four interconnect technologies[7]. (In Section 2.1.1 we showed that the ratio between Elmore delay and the Two-Pole estimator of [38] is very nearly constant. As would therefore be expected, fidelity values for the Two-Pole simulator are nearly identical to those for Elmore delay, and we do not report them here.)

Tables 4 and 5 show the fidelity to SPICE of the linear and Elmore delay estimators; the delay criterion is the 50% delay time to a given randomly-chosen critical sink in the net. We report the average difference in ranking over all topologies; the average rank difference for the topology which has lowest delay according to the estimator; and the average difference for the five topologies which have lowest delay according to the estimator. Our results show that Elmore delay has surprisingly high fidelity for the critical-sink delay criterion, and nearly perfect fidelity for the maximum sink delay criterion. For example, with 5-pin nets and IC3 technology parameters, optimal critical-sink topologies under Elmore delay averaged only 5.6 rank positions (out of 125) away from optimal according to SPICE, while the best topology for maximum Elmore delay averaged only 0.2 positions away from its "proper" rank using SPICE. Kim, Owens and Irwin [24]

[6] An early theorem of Cayley [16] implies that there are $|N|^{|N|-2}$ distinct spanning tree topologies for any given net N.

[7] Again, we use linear delay defined to be the source/sink pathlength. This definition leads to numerous ties between topologies; we break ties in favor of trees with lower total wirelength. Ties also occur with SPICE-computed delay because of the finite time step used; here we again break ties according to total wirelength.

Accuracy of Linear, Elmore and Two-Pole Delay Estimates for Critical-Sink Delay					
		$\|N\| = 4$		$\|N\| = 7$	
	Delay Ratio	average	std dev	average	std dev
IC1	SPICE/Linear†	–	28.4%	–	32.7%
	SPICE/Elmore	0.72	13.5%	0.69	15.4%
	SPICE/2-Pole	1.27	13.5%	1.23	15.4%
	2-Pole/Elmore	0.568	0.45%	0.566	0.22%
IC2	SPICE/Linear†	–	33.9%	–	38.8%
	SPICE/Elmore	0.74	16.1%	0.70	17.8%
	SPICE/2-Pole	1.30	15.9%	1.23	17.8%
	2-Pole/Elmore	0.572	0.92%	0.568	0.45%
IC3	SPICE/Linear†	–	34.9%	–	40.3%
	SPICE/Elmore	0.78	16.0%	0.72	17.8%
	SPICE/2-Pole	1.36	15.7%	1.27	17.9%
	2-Pole/Elmore	0.574	1.39%	0.571	0.80%
MCM	SPICE/Linear†	–	57.1%	–	61.6%
	SPICE/Elmore	0.69	20.5%	0.65	25.1%
	SPICE/2-Pole	1.20	20.8%	1.14	25.2%
	2-Pole/Elmore	0.568	0.96%	0.566	0.44%

Table 2: Accuracy of the Linear, Elmore and Two-Pole estimators for critical-sink delay. The table gives the average and standard deviation of the ratios between SPICE3e2 delay and estimated delay at a single random critical sink, averaged over 100 random nets. All nets are connected using MST constructions. Standard deviations are reported as a percent of the average ratio. (†) Linear delay is defined as the source/sink pathlength; because this is a distance rather than a time, we do not report a SPICE/Linear "ratio". However, we can report the standard deviation of this quotient, since it is independent of units.

and both the Two-Pole and Elmore estimators to be very consistent when measuring *maximum* sink delay, with standard deviations within 4% for all technologies on seven-pin nets. Thus, precomputed "correction factors" would seemingly compensate for the observed inaccuracy of these estimators. However, for delay at a random critical sink, the standard deviation of the accuracy ratio is consistently above 15%. This lesser consistency perhaps indicates that the traditional net-based performance objective is more "forgiving" of errors in the delay estimate than newer path-based delay objectives.[5]

2.1.2 Fidelity

A key observation is that precise accuracy is *not* really required of delay estimates used to construct routing trees. In practice, we only require that an estimator have a high degree of *fidelity* – i.e, an optimal or near-optimal solution according to the estimator should also be nearly optimal according to actual delay. To this end, we have defined a measure of fidelity vis-a-vis an exhaustive enumeration of all possible routing

[5] The small standard deviation of the accuracy values for maximum sink delay seems in part due to the rough similarity of the maximum source-sink distances over the examples studied. Note that, for example, with the MCM technology and $|N| = 7$, the average SPICE/Elmore ratio for a random sink is 0.65, whereas for the sink with greatest delay (generally furthest from the source) the ratio is 0.79. The critical-sink analysis in some sense better reveals this effect of sink distance from the source.

inductive effects; according to [5] and [38], its moment-based methodology is intermediate between SPICE and Elmore delay in both accuracy and computational efficiency.

2.1.1 Accuracy

In choosing a delay simulator, one traditionally measures *accuracy*, which may vary with the circuit technology and the specifics of a net (for instance, the number of pins it contains, or the size of its bounding box). Tables 2 and 3 indicate the accuracy of the linear, Elmore and Two-Pole models for each of the interconnect technologies described in Table 1. For each of the three estimators, the Tables give the average ratio of SPICE delay to the estimated delay, and also show the consistency of this ratio in terms of its standard deviation.[3] In Table 2, delay is calculated for a single random "critical" sink; in Table 3, delay is measured as the maximum delay at any sink in the net. For each net size, the results are averaged over 100 random nets with pin locations chosen from a uniform distribution over the routing area; each net is connected using the minimum cost spanning tree (MST) construction. We use MSTs rather than random tree topologies so that our comparisons will be for relatively good (albeit not necessarily optimal) routing solutions. (Observe that for a 7-pin net, finding the optimal-delay routing solution by exhaustive enumeration using SPICE is not computationally feasible.)

In all cases, the ratio of SPICE to Linear Delay has the largest standard deviation; this inaccuracy in the linear approximation is not surprising. It is also reasonable to expect poor "accuracy" of the Elmore and Two-Pole approximations with respect to SPICE, if only due to the somewhat ill-defined state of delay modeling and analysis noted in Footnote 3. Indeed, based on the average ratio of SPICE to Elmore delay or to the Two-Pole simulator, neither estimator seems particularly accurate: each is generally at least 20% away from SPICE on average at the critical sink.[4] Interestingly, Table 3 shows the ratio between SPICE

depth-first traversals: (1) to compute the delay along each edge and (2) to sum up the delays along each source-sink path. This fact is enabling to the efficient ERT methodology that we propose in Section 3.2.

[3] Again, we equate SPICE3e2 results with "actual delay". Our SPICE delay modeling uses constant unit resistance and capacitance values which vary with each interconnect technology. The root of the routing tree is driven by a resistor connected to the source; we remove driver attributes since we are concentrating on measuring delay within the interconnect. For the Two-Pole and SPICE simulators, every interconnect segment is broken into uniform segments, each at most 1/100th the length of the layout dimension, connected in series. To model sink loads, we use pure capacitive loads derived using minimum-size transistors. For all simulators, we have used the 50% rise time delay criterion, and we have measured both average sink delay and maximum sink delay. For the Two-Pole and SPICE simulators we have used time steps of 0.005 ns for the IC technologies and 0.05 ns for MCM.

We have found that our results are, qualitatively, independent of these detailed methodological choices (e.g., 50% rise time instead of 90% rise time as a delay criterion). However, it should be noted that many reasonable alternative simulation methodologies were possible. For instance, Elmore delay does not intrinsically correspond any delay time (it is simply the first moment of an impulse response), but can be said to correspond to a 63% delay criterion since $RC/2$ is the coefficient of s in the system transfer function $H(s)$ of a distributed RC line. On the other hand, the nature of the Two-Pole approximation makes it more suited to a 90% rise time criterion [38]. Other inconsistencies: while SPICE can model active devices as loads, the Two-Pole simulator can only handle "equivalent" sink capacitances; while SPICE and Two-Pole can model series inductance (for MCM interconnect), Elmore delay is solely a distributed RC model (and extracted inductance parameters are often not distributed with IC technology files) – indeed, the list of incomparable variables seems endless.

[4] Recall that we assigned a near-zero inductance value for the IC technologies, since inductance parameters were not available from MOSIS/MCNC, and since the Two-Pole simulator requires a non-zero inductance. We found that this does not change our results significantly. For example, if we increase the IC1 inductance parameter to 400 $fH/\mu m$, the average Two-Pole/Elmore ratio becomes 0.566 with a standard deviation of 0.51%. Elmore delay is independent of inductance since the first moment of the impulse response in a distributed RCL tree has no L term.

(i.e., signal delay is proportional to total tree capacitance, which is proportional to tree cost). In [9, 10, 34], the *linear* delay approximation is used; sink delays are thus proportional to source-sink path lengths, and a minimum-radius criterion is obtained.

Name	IC1	IC2	IC3	MCM
Technology	2.0 μm CMOS	1.2 μm CMOS	0.5 μm CMOS	MCM
r_d	164.0 Ω	212.1 Ω	270.0 Ω	25.0 Ω
unit wire resistance	0.033 $\Omega/\mu m$	0.073 $\Omega/\mu m$	0.112 $\Omega/\mu m$	0.008 $\Omega/\mu m$
unit wire capacitance	0.234 $fF/\mu m$	0.083 $fF/\mu m$	0.039 $fF/\mu m$	0.06 $fF/\mu m$
unit wire inductance	1x10^{-5} $fH/\mu m$	1x$10^{-5} fH/\mu m$	1x$10^{-5} fH/\mu m$	380 $fH/\mu m$
loading capacitance	5.7 fF	7.06 fF	1.0 fF	1000 fF
resistance ratio (x $10^6 \mu m$)	0.0050	0.0029	0.0024	0.0031
chip size	1x1 cm^2	1x1 cm^2	1x1 cm^2	10x10 cm^2

Table 1: Technology parameters for three CMOS IC technologies and one MCM technology. Parasitics and SPICE simulation decks for the IC1 and IC2 technologies are provided by MOSIS; IC3 parasitics are courtesy of MCNC; MCM interconnect parasitics are courtesy of Professor Wayne W.-M. Dai of UC Santa Cruz, and correspond to data provided by AT&T Microelectronics Division. The driver resistances (r_d) and sink loading capacitances are derived for minimum-size transistors. Note that inductance values for IC1 - IC3 are set to 1x$10^{-5} fH/\mu m$ (effectively zero) because they were not provided by MOSIS/MCNC, and because a non-zero inductance is required by the Two-Pole simulator (see Footnote 4 below).

Such simple delay approximations are known to be inaccurate as technology scales, e.g., smaller wire geometries imply that resistive effects of the interconnect become more dominant, particularly in relation to driver on-resistance (see the discussion below of "resistance ratio" effects, and note the four technology characterizations in Table 1). Furthermore, greater system speeds and layout areas may expose inductive effects on delay. Given these considerations, distributed RC delay approximations (e.g., that of Elmore [15]) or distributed RCL delay approximations (e.g., the "Two-Pole" simulator of Zhou et al. [38]) are of interest, since they are more accurate than linear or lumped-capacitance approximations while requiring less computation time than SPICE.

Elmore delay in an RC tree [15] [32] [36] is defined as follows. Given routing tree $T(N)$ rooted at the source n_0, let e_v denote the edge from node v to its parent in $T(N)$. The resistance and capacitance of edge e_v are denoted by r_{e_v} and c_{e_v}, respectively. Let T_v denote the subtree of T rooted at v, and let c_v denote the sink capacitance of v ($c_v = 0$ if v is a Steiner node). We use C_v to denote the *tree capacitance* of T_v, namely the sum of sink and edge capacitances in T_v. Using this notation, the Elmore delay along edge e_v is equal to $r_{e_v}(\frac{c_{e_v}}{2} + C_v)$. Let r_d denote the output driver resistance at the net's source. Then Elmore delay $t_{ED}(n_i)$ at sink n_i is:

$$t_{ED}(n_i) = r_d C_{n_0} \quad + \sum_{e_v \in path(n_0, n_i)} r_{e_v}(\frac{c_{e_v}}{2} + C_v) \tag{1}$$

Although Elmore delay has a compact definition and can be quickly computed[2], it does not capture all of the factors that account for delay. For example, the Two-Pole distributed RCL simulator [38] considers

[2] Elmore delay can be evaluated at *all* sinks in $O(k)$ time, as noted by Rubinstein et al. [32]. The calculation uses two

measure to guide the routing tree design, and derives motivating observations from analysis of the Elmore approximation for signal delay in distributed RC trees. Section 3 then presents our two main classes of CSRT algorithms. We first describe the *CS-Steiner* method, which perturbs an existing Steiner tree construction to account for the presence of identified critical sinks. We then propose an efficient class of *Elmore routing tree* (ERT) constructions which not only yield good CSRT solutions, but are also the first methods to optimize Elmore delay *directly* without any of the abstractions implicit in previous routing objectives. Section 3 also describes the extension of the ERT approach to net-dependent routing objectives. Experimental results are presented in Section 4, where we compare delays at critical sinks in our heuristic tree topologies with analogous delays obtained using the best-performing minimum Steiner tree heuristic [22] and the AHHK routing [1]. Our methods prove extremely effective, obtaining up to an *average* 69% reduction in signal delay to identified critical sinks in 8-sink nets. The ERT approach also yields *generic* high-performance routing trees when all sinks are equally critical: for 9-pin nets in 1.2μ CMOS IC (MCM) technology, we improve average sink delay by 19% (62%) and maximum delay by 22% (52%) over the minimum Steiner routing. We thus obtain a significant advance over the existing performance-driven routing tree constructions in the literature, including such recent works as [1] [10] [28]. Our results are complemented by a detailed analysis of the accuracy and *fidelity* of the Elmore delay approximation, and we furthermore provide *exact assessments versus optimal* for our heuristic tree constructions. To determine the latter data, we have developed a new theoretical characterization of Elmore-optimal routing trees, as well as a decomposition theorem for (Elmore-) optimal Steiner trees, which are of independent interest.

2 On Delay Approximations and Tree Design Objectives

For arbitrary signal nets N, the appropriate objective to use in *efficiently* constructing "high-performance routing trees" has not yet been established. In this section, we first consider necessary qualities for a delay approximation that is to be used in routing tree design. By studying both the relative accuracies and the relative *fidelities* of linear, distributed RC, distributed RCL, and SPICE-computed delay approximations, we demonstrate that the Elmore distributed RC delay approximation is of surprisingly high fidelity with respect to SPICE3e2. From Elmore's simple formula (i.e., the first moment of the impulse response in a distributed RC tree), we then develop revealing intuitions regarding the "correct" objective for critical-sink routing tree design.

2.1 Accuracy and Fidelity of Delay Approximations

Ideally, a routing algorithm will compute and optimize signal delays according to a detailed circuit simulation, such as that provided by SPICE. Since the computation times required by SPICE are prohibitive for routing tree construction, simpler delay approximations must be used. For example, the traditional minimum-cost Steiner tree objective, in addition to minimizing wiring area, corresponds to a *lumped-capacitance* model

delays. For example, Lin and Du [26] use a linear delay approximation so that their method updates the module placement to reduce the rectilinear distance between sources and critical sinks. Other path-oriented methodologies include those of Hauge et al. [18] and Teig et al. [35].

If a timing-critical path passes through a given net, the path-oriented approach can provide an explicit bound on delay at that net's critical sink. While the net-oriented approach may arguably provide only implicit routing constraints, it is still easy to identify critical sinks after the timing analysis has been performed, or *a priori* by finding paths in the design that contain more module delays. This reveals a "placement-routing mismatch": the performance-driven routing constructions reviewed above generally address *net-specific* objectives (min cost, min radius, cost-radius tradeoffs, etc.) and do not exploit the critical-path information that is available during iterative performance-driven layout. As a consequence, designers cannot realize the full benefit of high-quality timing-driven module placements. With this in mind, our work develops new high-performance routing tree constructions which *directly* exploit available critical-path timing information.

1.2 The Critical-Sink Routing Tree Problem

A *signal net* N consists of a set of pin locations $\{n_0, n_1, ..., n_k\}$ in the Manhattan plane, which are to be connected by a *routing tree* $T(N)$. We use n_0 to denote the *source*, with the n_i $(1 \leq i \leq k)$ denoting *sinks*. The *cost* of an edge e_{ij} in $T(N)$, denoted by d_{ij}, is the Manhattan distance between the endpoints n_i and n_j of the edge. The cost of the tree $T(N)$ is simply the sum of its edge costs. In a given routing tree $T(N)$, the signal delay between two terminals n_i and n_j is denoted by $t(n_i, n_j)$; the shorthand notation $t(n_i)$ indicates the delay from the source to the sink n_i. Finally, we allow each n_i to have an associated *criticality*, α_i, reflecting the timing information obtained during the performance-driven placement phase. Our goal is to construct a routing tree $T(N)$ which minimizes the weighted sum of sink delays:

Critical-Sink Routing Tree (CSRT) Problem: Given a signal net $N = \{n_0, n_1, ..., n_k\} \subset \Re^2$ with source n_0 and possibly varying sink criticalities $\alpha_i \geq 0$, $i = 1, \ldots, k$, construct a routing tree $T(N)$ such that $\sum_{i=1}^{k} \alpha_i \cdot t(n_i)$ is minimized.

This CSRT problem formulation is quite general, and easily captures traditional performance-driven routing tree objectives: (i) *average delay* to all sinks is minimized by using all $\alpha_i =$ some positive constant, then taking the L_1 sum of the weighted delays; and (ii) *maximum delay* to any sink is minimized by using all $\alpha_i =$ some positive constant, then taking the L_∞ sum of the weighted delays. In the discussion below, we will concentrate on the simple yet realistic case where *exactly one* critical sink, denoted by n_c, has been identified. In other words, we assume that $\alpha_c > 0$ and that all other $\alpha_i = 0$. Our methods may be generalized to the case where a small number of critical sinks is specified.

The remaining discussion is organized as follows. Section 2 discusses the appropriate choice of a delay

the former work also incorporates a hierarchy-based net ordering. For minimum Steiner tree routing, the 1-Steiner method [22] is the best-performing heuristic, and we therefore use it as a basis for comparison below.[1]

In [9], Cohoon and Randall proposed a heuristic which simultaneously considered both the *cost* (total edge length) and the *radius* (longest source-sink path length) of the routing tree. A more general formulation was given by Cong et al. [10], wherein a parameter ϵ guides the tradeoff between cost and radius minimization; the same authors in [10] proposed the "provably good" BRBC (bounded-radius, bounded-cost) algorithm, which affords both cost and radius simultaneously within *constant* factors of optimal. The BRBC method and works of Awerbuch et al. [3] and Khuller et al. [23] all achieve a smooth cost-radius tradeoff via the same basic idea: (i) make a depth-first traversal of the minimum spanning tree over the signal net, and (ii) if the accumulated path length from the source to some sink becomes too large, modify the tree to reduce that particular source-sink path length. The cost-radius tradeoff may also be viewed as one between competing minimum spanning tree (MST) (or minimum-cost Steiner tree) and shortest-path tree (SPT) constructions. Using this perspective, Alpert et al. [1] recently proposed the AHHK algorithm, which achieves a direct MST-SPT tradeoff. Finally, Cong et al. [11] have recently proposed the use of rectilinear Steiner arborescences [30], or A-trees; these are essentially minimum-cost SPTs with Steiner points allowed. The delay performance of the AHHK algorithm is superior to that of the BRBC or A-tree constructions [1], and thus below we use AHHK as another basis of comparison with our new methods.

1.1 Motivations For Critical-Sink Routing

In performance-driven layout for cell-based designs, timing-critical paths are determined by static timing analysis, and modules in these paths are then placed close together (see, e.g., [13, 18, 20, 26, 27, 34]). The static timing analysis thus *iteratively* drives changes within both the module placement and the global routing phases. Our contribution stems from carefully considering routing tree constructions within this overall performance-driven layout process.

In general, existing performance-driven placement algorithms may be classified as either *net-oriented* or *path-oriented*. *Net-oriented* placement typically uses centroid-connected star cost [33], probabilistic estimates of Steiner tree cost [20], minimum spanning tree cost [13] or the bounding box semiperimeter [27] to estimate wire capacitance and signal delay for a multi-terminal net. From this information, critical timing paths between primary inputs and primary outputs are computed, after which module placements are updated to reduce these "net-based" objectives for signal nets along the critical paths. By contrast, *path-oriented* placement considers delay between the source and a particular *critical sink* of a multi-terminal net. The critical sink is typically determined via timing analysis using known module delays and estimated path

[1] Recent studies by Barrera et al. [4], using the optimal Steiner code of J. Salowe, show that 1-Steiner is within 0.25% of optimal on average.

Near-Optimal Critical Sink Routing Tree Constructions*

Kenneth D. Boese, Andrew B. Kahng, Bernard A. McCoy†, and Gabriel Robins†

CS Dept., University of California at Los Angeles, Los Angeles, CA 90024-1596
† CS Dept., University of Virginia, Charlottesville, VA 22903-2442

Abstract

We present *critical-sink routing tree* (CSRT) constructions which exploit available critical-path information to yield high-performance routing trees. Our CS-Steiner and "Global Slack Removal" algorithms together modify traditional Steiner tree constructions to optimize signal delay at identified critical sinks. We further propose an iterative *Elmore routing tree* (ERT) construction which optimizes Elmore delay *directly*, as opposed to heuristically abstracting linear or Elmore delay as in previous approaches. Extensive timing simulations on industry IC and MCM interconnect parameters show that our methods yield trees that significantly improve (by averages of up to 67%) over minimum Steiner routings in terms of delays to identified critical sinks. ERTs also serve as generic high-performance routing trees when no critical sink is specified: for 8-sink nets in standard IC (MCM) technology, we improve average sink delay by 19% (62%) and maximum sink delay by 22% (52%) over the minimum Steiner routing. These approaches provide simple, basic advances over existing performance-driven routing tree constructions, including the recent works of [1, 10]. Our results are complemented by a detailed analysis of the accuracy and *fidelity* of the Elmore delay approximation; we also *exactly* assess the suboptimality of our heuristic tree constructions. In achieving the latter result, we develop a new characterization of Elmore-optimal routing trees, as well as a decomposition theorem for optimal Steiner trees, which are of independent interest.

1 Introduction

Due to the scaling of VLSI technology, interconnection delay has become a dominant concern in the design of complex, high-performance circuits [13, 34]. Performance-driven layout design has thus become an active area of research over the past several years. In this paper, we develop a new *critical-sink* problem formulation and new solutions for performance-driven routing tree design.

For a given signal net, the typical goal of performance-driven routing is to minimize average or maximum source-sink delay. Much early work implicitly equates optimal routing with minimum-cost Steiner routing. For example, Dunlop et al. [14] use static timing analysis to yield net priorities, so that the highest-priority nets may be routed by minimum Steiner trees, leaving lower-priority nets to subsequently encounter blockages. Jackson, Kuh, and Marek-Sadowska [21] and Prastjutrakul and Kubitz [28] have given approaches which are tuned to building-block layout and allow prescribed upper bounds on individual source-sink delays;

*Partial support for this work was provided by a GTE Graduate Fellowship and by NSF MIP-9110696, NSF Young Investigator Award MIP-9257982, ARO DAAK-70-92-K-0001, and ARO DAAL-03-92-G-0050. ABK was also supported by NSF MIP-9117328 during a Spring 1993 sabbatical visit to UC Berkeley.