

CSL *COORDINATED SCIENCE LABORATORY*

ON THE ORDERING OF CONNECTIONS FOR AUTOMATIC WIRE ROUTING

LUTHER C. ABEL

UNIVERSITY OF ILLINOIS – URBANA, ILLINOIS

"THIS DOCUMENT HAS BEEN APPROVED FOR PUBLIC RELEASE AND SALE; ITS DISTRIBUTION IS UNLIMITED."

ON THE ORDERING OF CONNECTIONS FOR AUTOMATIC WIRE ROUTING

by

Luther C. Abel

This work was supported in part by the Joint Services Electronics Program (U. S. Army, U. S. Navy and U. S. Air Force) under Contract DAAB-07-67-C-0199, in part by the National Science Foundation under Grant NSF GK-15459, in part by the Army Research Office - Durham under Contract DAHC 04-72-C-0001, and in part by the Rome Air Development Center under Contract USAF 30(602)-4144.

Reproduction in whole or in part is permitted for any purpose of the United States Government.

This document has been approved for public release and sale; its distribution is unlimited.

ON THE ORDERING OF CONNECTIONS FOR AUTOMATIC WIRE ROUTING

Luther C. Abel

Abstract

Most wire routing programs utilize a max-e-running technique to route one connection at a time. Once routed, a wire cannot be moved even if it is subsequently discovered to interfere with the successful completion of other connections. The order in which the desired connections are presented to the routing algorithm has therefore been thought to be of critical importance. Experimental evidence is presented herein, however, to show that the performance of a router, when measured in terms of the total of the minimum (or ideal) lengths of the connections successfully completed, is in fact independent of the order in which connections are attempted.

INTRODUCTION

Virtually every automatic printed circuit card wire routing program in existence today utilizes Lee's Algorithm [1] or some other embodiment of Moore's maze-running technique [2]. Although these routines have a great advantage in providing a systematic search for a path between two points, they also have a fundamental shortcoming in that they route precisely one wire at a time, providing no feedback or anticipation in the routing process to avoid conflict between wires or to assure that some early wire routing will not prevent successful completion of some later connection.

Because of this blindness the order in which a set of connections (assuming the set is a priori completely defined) is presented to a wire routing program would seem to be of crucial importance to the successful routing of a printed circuit board. Although the topic of connection ordering has been vigorously debated at design automation conferences and workshops, no comparative study has actually been published.

Connection lengths are typically measured in terms of the rectilinear or "Manhattan" distance between the terminals to be joined ($l = \Delta x + \Delta y$). Two of the most common methods for interconnection ordering are in ascending order of length and descending order of length. Proponents of the former argue that it is easier to route a long wire around a short one than vice versa; supporters of the latter counter that longer wires are more difficult to lay out and hence should be attempted first. In the descriptions of a few reported systems [3,4], the authors declare that shortest connections must be routed first but give no justification; another author [5] concludes

on the basis of just two observations that "sorting by net size [wire length] does not improve layout."

When multilayer printed circuit boards are used for wiring, connections are frequently first assigned to individual layers by slope classes, that is, the sectors (say, octants) of a circle into which a line drawn between the two points to be joined would fall. For connections within one slope class, another possible ordering is based on the magnitude of the component of a connection's length perpendicular to the sector axis; such an ordering is an attempt to measure the degree of interference a wire presents to other wires within the sector, all of which should be roughly parallel to the sector axis.

DESCRIPTION OF EXPERIMENT

Wiring lists for thirty-eight different types of printed circuit boards from the ILLIAC IV Processing Element were used as a data base for these experiments. Up to twenty 16-pin dual-inline integrated circuit packages can be accommodated on a board in five rows of four packages each, with the packages parallel to the 100-pin connector located at one edge of the board. Integrated circuit package placement data were included in the board wiring lists. For convenience in discussing board layouts, it will be assumed that a cartesian coordinate system is superimposed on the board, with the Y-axis parallel to the edge with the connector.

Wiring was performed in an area approximately four inches square (10 cm by 10 cm) on a 50-mil (1.3 mm) grid. It was assumed that a pin connection pad could be formed entirely within one grid square. Signal sets

were interconnected using daisy-chain wiring (i.e., each signal connection point had at most two wires emanating from it), with the signal source constrained to be at one end of the chain.¹

Two printed circuit card layers were used for signal interconnections in these experiments. Power and ground were assumed to be supplied by other layers and are ignored in this work. A simple slope-class heuristic rationale was used to assign a connection to one of the two layers: all connections whose X-component of length exceeded their Y-component ($\Delta x > \Delta y$) were assigned to Layer 1, all others were assigned to Layer 2. The number of connections assigned to each layer for routing was approximately equal, but the connections on Layer 1 were on the average slightly longer, leading to a roughly 3:2 ratio of total minimum length of the connections attempted. Each required pin-pair interconnection was considered to be a separate wiring problem, divorced from all reference to its electrically common fellow connections (i.e., signal networks were decomposed into sets of independent pin-pair interconnection problems).

The wire routing program for these experiments was a straightforward implementation of Lee's Algorithm [1]. No attempt was made to speed up the program with heuristic "feeler" or "runner" techniques. Since the routing of long wires immediately adjacent to a row of pins tended to sharply curtail the

1. The size of the cards actually used in ILLIAC IV is slightly greater and wiring is more complex because the emitter-coupled logic family used in ILLIAC IV requires external pulldown and terminating resistors which were ignored for these experiments. Layouts of the actual cards were created manually; plated through holes were permitted at fixed locations to provide inter-layer stitching. A photograph of a typical card is included in [6].

successful routing of any subsequent connections to those pins, two heuristics were incorporated into the router to avoid this problem: Wires which joined two points in the same row of pins (i.e., connections with $\Delta x = 0$) were forced to swing away from the row of pins and, if possible, remain at least two grid squares away. Wires joining two points not in the same row were (again, if possible) forced to follow a Z-shaped path after departing from the terminals in the X-direction, avoiding the L-shaped path and concomitant blocking of a row of pins which might otherwise be chosen.

Inter-layer transitions by means of plated-through holes (vias) were not permitted. If a connection could not be completely routed on the layer to which it was assigned, it was simply abandoned and counted as a failure.

Wire routings were attempted for each layer of each board five times, using the following five connection ordering methods:

- a) Shortest connection first (abbreviated by the letter "S");
- b) Longest connection first (L);
- c) Random ordering (R);
- d) Connections with the shortest X-component of length first, with shortest overall length used for tie-breaking (X);
- e) Connections with the shortest Y-component of length first, again with shortest overall length used for tie-breaking (Y).

In any of the ordering methods, if two connections ranked equally in the selection process (e.g., were of the same length for the "S" or "L" methods), the one to be attempted first was chosen arbitrarily.

Statistics recorded for each run include the number of connections attempted and the number successfully completed, the length of wire used in routing the successful connections, the sum of the rectilinear distances

between points connected and the overall sum for all attempted connections. Those latter correspond to the minimum length of wire with which the connections might have been wired (ignoring the increase due to topologically unavoidable detours).

EXPERIMENTAL RESULTS

In order to meaningfully compare ordering methods, there must be an unbiased means of presenting our experimental results relating to the performance of the router with each of the ordering methods.

Certainly some connections are easier to route than others. If the number of connections successfully completed were simply counted and used as a performance criterion, the measure of the quality of an ordering method would be biased in favor of those methods which caused the router to attempt the easiest connections first. Specifically, since geometric arguments and the physical realities of board design indicate (at least to a first order approximation) that the ease of routing a wire is inversely proportional to its length, tabulating the number of completed connections would undoubtedly rank the shortest-first ordering method as best. Moreover, such a criterion would not give an adequate measure of how difficult it would be to route the remaining uncompleted connections were extra layers or manual intervention in the task permitted.

The relationship of the difficulty of routing to wire length suggests another measure: the total length of the wiring successfully routed. With this criterion, successful completion of one long, difficult wire should rank equally with an equivalent total length of short, easy wires. But evaluating ordering methods by the actual length of wire used for the

successfully routed connections is still less than perfect since it would bias the rating in favor of some ordering which inherently causes a large number of wires to be completed in greater than minimum length.

The real task of a router is to specify the layout of wires joining a set of terminals which, with ideal routing, would all be connected with wires of minimum length (actually, in minimum length plus some unavoidable detours whose minimum length can be a priori determined). Therefore the most unbiased measure of the performance of a router and/or the effect of a connection ordering method is the sum of the minimum or ideal lengths of the connections which were successfully routed.

This measure may be used directly, or it may be normalized to the size of the board being routed by dividing this length by the maximum length of wire which can possibly be put on the board. Typically, both of these lengths are measured in terms of routing grid squares, one square being the fundamental quantum of length. The boards used in these experiments had approximately 5800 available squares; wire lengths (wire densities) are hereafter expressed as percentages of this quantity.

The statistical distributions of this total minimum wire lengths for all attempted wire routings (i.e., $\sum (\Delta x + \Delta y)$ for all connections on a layer) for the 38 card types are shown² in Figure 1. The total ideal lengths of all connections ranged from 3% to 55%, with a mean length of approximately 36% for Layer 1 and 24% for Layer 2.

Figure 2 shows the distribution of the sums of the minimum lengths of successfully completed connections for the shortest-first, longest-

2. All curves have been smoothed by introducing some statistical "fuzziness" into each data point. As the occurrence of each point was plotted, specified fractional values were added to adjacent columns of the histogram.

first and random ordering methods. Figure 3 shows these ideal length distributions for the minimum-X-component-first and minimum-Y-component-first orderings, with the shortest-first distribution repeated for comparison. For these latter curves, recall that the connections on Layer 1 fall into the octants of the circle adjacent to the X-axis, while those on Layer 2 occur in the octants nearest the Y-axis. Hence a minimum-Y-component-first ordering for Layer 1 attempts earliest connections having minimum component of length perpendicular to the sector axis, while the minimum-X-component-first ordering for the same layer represents a particularly perverse ordering. For Layer 2, the minimum-X-component-first ordering similarly routes earliest those wires with the minimum off-axis component. Principal characteristics of all five distributions are summarized in Table 1.

These data represent our principal experimental result: when the summation of the distances between the points successfully connected by a router is used as a measure (and we contend this is the most unbiased comparison criterion), the order in which the connections are attempted has little if any effect on router performance.

True, there are some small differences between ordering methods; in these experiments ordering based on a connection's component of length perpendicular to the sector axis for a layer cumulatively performed slightly better than the others, with shortest-first ordering ranking second. But no ordering proved conclusively best for all 38 cards. In fact, each of the methods performed best for at least one card. Moreover, simply varying the arbitrary choices made when connections ranked equally in the ordering process changed individual results by the same percentages as the differences between

the cumulative results. Also, if the set of Layer 1 routings and the set of Layer 2 routings are regarded as separate experiments (as essentially they are), different rankings for the methods are observed. We therefore feel that the lack of strong differences between the methods is far more significant than the fact that small differences do exist.

Figures 4 and 5 show the distributions of the actual length of wire used for the successfully completed connections, and Table 2 summarizes these characteristics. A strong correlation between these ideal length distributions is observed, with an average of 1.5 times as much wire as the minimum actually required for the connections. But there are individual anomalies; for example, the ratio of actual to ideal wire lengths is consistently higher for random ordering.

The question of whether one ordering method might demonstrate a substantial, consistent superiority for some specific limited range of attempted wiring densities was investigated by plotting the completion rate (measured in ideal length) versus the total length of connections attempted for the various methods; no such prepotency was observed. Moreover, a large scatter was observed in the completion rates for sets of boards requiring essentially the same amount of wire to be routed. This leads us to believe that complex geometric and wire-interrelationship factors considerably influence the successful routing of a set of connections. Conversely, the completion rate was almost completely independent of the density of the attempted wiring.

Akers [7] and others have reported that routers perform as if there existed within them a fundamental barrier to complete routing such that the

ideal length of connections successfully routed simply cannot exceed a fixed fraction of the available board area. Stated another way, any set of connections whose total minimum length exceeds a given fraction of the board area is certainly doomed to contain wires which will not be successfully routed. The data of Figures 2 and 3, with their extremely sharp cutoff at 22-25%, strongly support this assertion.

The value of the cutoff point seems dependent on board geometry and the exact details of the routing algorithm such as the ordering method used and the method of choosing among a multiplicity of shortest paths. In addition to the small variations observed as the connection ordering method is changed, other experiments we have conducted indicate that the cutoff point is shifted by other geometric and algorithm variations, but remains just as sharp.

CONCLUSIONS

Experimental evidence of a statistically meaningful nature has been presented to show that the performance of a maze-running type router, when measured by the total ideal length of connections successfully routed, is independent of the order in which connections are attempted. If the creator of a computer-aided design system feels compelled to include an ordering method, then ordering based on a connection's component of length perpendicular to the sector axis for a layer would probably perform over a long term in a marginally superior manner. The existence of a sharp fundamental limit to the acceptable board wiring density for a router has also been demonstrated.

Although the router used in these experiments did not permit inter-layer transitions and the inclusion of such a facility would undoubtedly

increase the maximum wiring density which could be achieved, we feel the qualitative result concerning uniformity of performance independent of connection ordering would not change.

Further research must be done on the relationship of board geometry and the geometric properties of package placement to successful routing. Why, for example, did completion rates for boards of equal attempted wiring density vary by an almost 2:1 ratio in our experiments? Additional investigation is also needed into the algorithmic properties of routers to try to overcome their one-wire-at-a-time blindness. Methods for choosing the optimal path when a choice of routings exist seem crucial to preventing early wire routings from blocking later ones, and the blocking properties of wire forced outside their minimum-distance rectangles on the Manhattan grid requires examination.

Two additional factors may have a direct impact on the successful routing of a board: the spatial uniformity of attempted wiring density (that is, the presence or absence of "hot spots" through which many wires must pass) and the total component of ideal wire length perpendicular to the major wiring axis for a layer. If a correlation between these factors and the wireability of a board does indeed exist, then they perhaps can be used as package placement criteria, either in addition to or instead of the traditional placement objective of minimum wire length [8].

Ordering Method

Wiring Densities (percent of board area)

	Layer 1			Layer 2		
	MIN.	MEAN	MAX.	MIN.	MEAN	MAX.
Shortest - first	7	17	26	3	16	24
Longest -first	9	18	24	3	14	22
Random	9	17	26	3	14	22
Minimum X-componet first	7	17	25	3	16	26
Minimum Y-component first	11	20	29	3	15	24

Table 1. Distribution of ideal lengths of successfully completed connections.

Ordering Method

Wiring Densities (percent of board area)

	Layer 1			Layer 2		
	MIN.	MEAN	MAX.	MIN.	MEAN	MAX.
Shortest - first	13	28	46	4	23	40
Longest - first	17	29	44	3	21	36
Random	18	30	44	3	23	42
Minimum X-component first	14	28	49	4	24	45
Minimum Y-component first	15	30	46	4	23	41

Table 2. Distribution of actual wire lengths for successfully completed connections.

REFERENCES

- [1.] Lee, C. Y., "An algorithm for path connections and its applications," IRE Trans. Electron Comput., Vol. EC-10, September 1961, pp. 346-365.
- [2.] Moore, E. F., "Shortest path through a maze," in Annals of the Computation Laboratory of Harvard University, Harvard University Press, Cambridge, Mass., Vol. 30, pp. 285-292, 1959.
- [3.] Freeman, M., et al, "Multilayer printed wiring--computer aided design," Proc. Fourth Share/ACM/IEEE Design Automation Workshop, 1967, pp. 16-1 -- 16-28.
- [4.] Heiss, S., "A path connection algorithm for multi-layer boards," Proc. Fifth Share/ACM/IEEE Design Automation Workshop, 1968, pp. 6-1 -- 6-14.
- [5.] Moore, R. C., "Packaging flat pack intergrated circuits for earth satellites," Eighth International Circuit Packaging Symposium Record (Advances in Electronic Circuit Packaging, Vol. 8), 1967, Section 4/4, pp. 1-9.
- [6.] Slotnick, D. L., "The fatest computer," Scientific American, Vol. 224, February 1971, pp. 76-87.
- [7.] Akers, S. B., as reported in the minutes of the IEEE Circuits Standards Committee meeting held at the Naval Postgraduate School, Monterey, California, J. Bordogna, Secretary, May 1969, p. 19.
- [8.] Hanan, M. and Kurtzberg, J. M., "A review of the placement and quadratic assignment problems," IBM Corporation Research Report RC3046, April 1970.




Figure 1. Distribution of ideal total lengths of all connections
(a) Layer 1.




Figure 1. (continued)
(b) Layer 2.




Figure 2. Distribution of ideal total lengths of successfully completed connections for shortest-first, longest-first, and random orderings.
(a) Layer 1.




Figure 2. (continued)

(b) Layer 2.




Figure 3. Distribution of ideal total lengths of successfully completed connections for minimum-X-component-first, minimum-Y-component-first.

(a) Layer 1 ($\Delta x > \Delta y$ for connections on this layer).




Figure 3. (continued)

(b) Layer 2 ($\Delta y \geq \Delta x$ for connections on this layer).




Figure 4. Distribution of actual total lengths for successfully routed wires for shortest-first, longest-first and random orderings.
(a) Layer 1.




Figure 4. (continued)

(b) Layer 2.




Figure 5. Distribution of actual total lengths for successfully routed wires for minimum-X-component-first, minimum-Y-component-first, and shortest-first orderings.
(a) Layer 1.




Figure 5. (continued)

(b) Layer 2.

DOCUMENT CONTROL DATA - R & D

(Security classification of title, body of abstract and indexing annotation must be entered when the overall report is classified)

1. ORIGINATING ACTIVITY (Corporate author) Coordinated Science Laboratory University of Illinois Urbana, Illinois		2a. REPORT SECURITY CLASSIFICATION UNCLASSIFIED	
		2b. GROUP	
3. REPORT TITLE ON THE ORDERING OF CONNECTIONS FOR AUTOMATIC WIRE ROUTING			
4. DESCRIPTIVE NOTES (Type of report and inclusive dates)			
5. AUTHOR(S) (First name, middle initial, last name) Luther C. Abel			
6. REPORT DATE November 24, 1971		7a. TOTAL NO. OF PAGES 23	7b. NO. OF REFS 8
8a. CONTRACT OR GRANT NO. DAAB-07-67-C-0199; NSF GK-15459		9a. ORIGINATOR'S REPORT NUMBER(S) R-537	
b. PROJECT NO. DAHC 04-72-C-0001 USAF 30(602)4144		9b. OTHER REPORT NO(S) (Any other numbers that may be assigned this report) DILU ENG 71-2240	
c.			
d.			
10. DISTRIBUTION STATEMENT This document has been approved for public release and sale; its distribution is unlimited.			
11. SUPPLEMENTARY NOTES		12. SPONSORING MILITARY ACTIVITY Joint Services Electronics Program through U. S. Army Electronics Command; Army Research - Durham; Rome Air Dev. Center	
13. ABSTRACT Most wire routing programs utilize a maze-running technique to route one connection at a time. Once routed, a wire cannot be moved even if it is subsequently discovered to interfere with the successful completion of other connections. The order in which the desired connections are presented to the routing algorithm has therefore been thought to be of critical importance. Experimental evidence is presented herein, however, to show that the performance of a router, when measured in terms of the total of the minimum (or ideal) lengths of the connections successfully completed, is in fact independent of the order in which connections are attempted.			

KEY WORDS

LINK A

LINK B

LINK C

ROLE

WT

ROLE

WT

ROLE

WT

Computer-aided design

Computer design automation

Automatic design

Printed circuit

Wiring