Steve Meyer

Pragmatic C Software 220 Montgomery Street, Suite 925 San Francisco, CA 94104

Abstract: Since mathematically optimal layout algorithms seem unattainable, layout needs to become a more experimental science. This paper advocates the use of the bench mark circuits in controlled experiments. The bench mark layout environment, open problems solvable by experiment, and the relation of layout experiments to scientific methodology are discussed. The paper concludes by showing that physical cell synthesis by computer program is unpromising since it is no easier than the general artificial intelligence problem.

1. Introduction

In a perfect world, there would exist provably efficient layout algorithms. Unfortunately, in reality most layout algorithms are NP complete [3]. There is even a lack of consensus concerning what makes one circuit layout superior to another. Circuit area is probably most important, but among numerous conflicting criteria, electrical characteristics, congestion, via number, timing, and power distribution must be considered. Progress in such unstructured problem domains can often be facilitated by means of scientific experiments. The various sets of benchmark circuits compiled for the physical design workshops [2] [18] provide such an experimental opportunity. The purpose of this paper is to advocate use of the bench mark circuits in more systematic experiments, to give open problems that can be answered by means of bench mark experiments, and finally to illustrate limitations of the experimental method by showing that layout synthesis of macros cell descriptions into mask patterns is impossible since it is no easier than the general artificial intelligence problem. Since solution of artificial intelligence problems is unlikely in the foreseeable future, there is no need for experimentation.

2. The Bench Mark Layout Environment

The bench mark layout environment has facilitated growth in placement and global routing because it retains the essential aspects of layout problems while omitting the numerous details required by a commercial circuit design environment. The YAL language [24] has proven effective in coding net lists, cell libraries, and placements because it is tailored to the algorithm development environment (see [15] for a discussion of problems of generalized coding schemes), and because its level of abstraction has proven to be correct. The omission of various technology dependent factors such as cell electrical function, pin pair wire segment decomposition, via representation, and mask geometric details has not caused problems. The macro cell library includes enough different cell types to allow nearly any circuit to be coded, yet avoids the complexity of the hundreds of different macro cells that would be required by a commercial system.

The inclusion of cell blockages and requirements for correct feed through handling has allowed progress in various global routing problems that were rarely dealt with before the advent of the bench mark circuits [10].

The mere existence of bench mark circuits is valuable since within the IC system design area far too many results have previously used only anecdotal evidence (see for example [8], pp. 26-28, 30). It is true that various layout algorithms were compared before the bench mark circuits were available (see [4] [5] [6] for example), but the pre bench mark work compared published algorithms as implemented for one particular layout system. When results for a technique did not match the algorithm originator's claims, the discrepancy could usually not be explained. But even this primitive comparison methodology proved valuable since it led to the commercial systems ([6] [13] [18]) that even today have still layed out the vast majority of ASICs. At least at LSI Logic during development in the 1980s of the original algorithms, availability of the current bench mark comparisons would have been valuable since placement and routing were different computer programs for which different corporate entities were responsible. This made total layout system comparisons somewhat problematic since both the placement program and global router had to take advantage of a different program's effectively black box algorithm.

3. Current Bench Mark Usage and Problems

Bench mark circuits are currently primarily used as test cases during layout program development, for generating experimental results when new algorithms are described (see recent layout papers in DAC or ICCAD proceedings), and to hold competitions to choose the best (or few best) layout systems at meetings of this workshop. These applications are interesting but they do not provide much help in guiding a semiconductor company in developing a layout system tailored to its commercial product lines. This product specific customization or parameter setting is currently required by even the most advanced academic algorithms (see for example [22], p. 44). Also, Commercial systems always have product type specific requirements, and eventual production volume. Knowing that one program produces less area in the abstracted bench mark environment is not of much use when developing a commercial system.

Even within the bench mark environment, comparison problems have arisen during the previous workshops. Some systems used the UTMC router [20] in the 1987 and 1988 row based style workshops while others used their own router. Even for systems that used the UTMC router, there were comparison problems. The UTMC router inserts one grid wide feed through cells where needed. Some of the placements required this feature, but some were made worse by it. Some placements achieved small area but required more feed through than the UTMC router was able to add. Those placements would probably, but not definitely, require more area after feed through addition. Various placements used a substrate size dictated by I/O pad geometry. These placements required larger area than was required by those which ignored I/O pads. The area determined from the circumference needed for I/O pads was much larger than the area required by internal cells. The bench mark system locates all macro cell pins on the second metal layer so that unconnected pins block the vertical feed through track. Systems designed to take advantage of free vertical tracks caused by unconnected pins might produce poorer results than they otherwise would. See [7, p. 127] and [25, fi gure 5] for more recent comparison problems.

Comparisons according to fi nal area also have various methodological problems. It is possible for a system using an inferior algorithm but a better implementation to produce less area than what seems to be a better algorithm for which the implementation or choice of approximations is problematic. If an algorithm works poorly, it is currently not possible to isolate the reason. It could be a implementation mistake, a hidden background factor that caused the algorithm to be specialized to the original layout system, or an algorithm problem. Without controlled experimentation it is not possible to determine the reason for the better or worse results. Finally, by competitively comparing area, a promising new algorithm may not continue to be developed because the early versions do not initially produce competitively small area. It is generally considered a bad idea in research and development to put all effort into developing one approach to the exclusion of all alternative approaches.

4. Proposed Bench Mark Usage

This paper proposes that the bench marks be used for scientific experiments for which all aspects of the layout program are controlled except the one under study. The results are then reported using the normal scientific method. By presenting results in this manner, it allows new algorithms to be developed since the particular aspect for which they are superior can be presented. It also allows commercial layout system implementers to much better evaluate the value of various algorithms independent of implementation quality. It may even allow general algorithmic questions to be answered which have importance beyond IC layout.

5. Relation to Scientific Methodology

It is possible to perform experiments in the bench mark system that previous work in the methodology of science has identified as having importance for scientific growth. The most widely known condition first identified by Professor Kuhn requires the ability to solve puzzles [12]. There must be a way to decide within one approach (sometimes called a research program) if one technique is superior to another. For example, it should be possible to decide by scientific experiment within the simulated annealing research program if one annealing temperature schedule is superior to another. This requires the ability to control every possible variable and is possible in the bench mark experimental system. Of course, in an area dealing with human design, a possible experimental result might be that one method is better for one design or substrate type while another is better for another. The bench mark system may provide the capability to systematically characterize such differences. This last sort of puzzle solving is known as problem shifting or problem splitting [9].

Another condition first identified by Professor Popper is the ability to falsify hypotheses (see [17] [9]). As a trivial example, without a controlled and widely available layout environment it is impossible to falsify random placement Imagine a claim of discovery that random placement with no evaluation function is superior to all over techniques. Without a controlled experimental system, the advocate of random placement could reasonably claim any falsification based on another implementation of random placement simply uncovered flaws in the implementation.

6. Proposed Bench Mark Experimental Opportunities

I believe the following experimental questions merit further study.

a. Controlled Algorithm Comparison.

It seems to me too much effort has been put into trying to produce improved layout systems without putting a corresponding amount of effort into attempting to understand what makes the system "good". The algorithm is usually named after the search strategy it uses, but it has not yet been shown that search strategy is a significant determinant of layout quality. Even the acknowledged best openly described layout systems still are effectively black boxes. For example, the success of Timberwolf [22] could as easily be related its combination of placement and global routing into one process, as to its search strategy. It is not clear what aspect of quadrisection makes it work [23].

I believe many people would claim simulated annealing is the search strategy of choice, but this has never been proven. There is even considerable negative evidence against its utility from theoretical studies. See [16] for a discussion of the weak mathematical power of simulated annealing. Also, in all the work that sought good searching heuristics, mostly for solving the traveling salesman problem, simulated annealing was not even considered (see for example [11]). Finally, many physical design workers still believe in continuous systems theory type approaches such as simulated annealing while those approaches are obsolete in nearly every other area of computer science.

A controlled experiment that keeps all aspects of a layout system constant except for search strategy would be interesting. Of course, considerable ingenuity may be required to eliminate problems caused by algorithm aspects whose implementation is related to properties of simulated annealing. Hopefully, the experiments will be described in sufficient detail so that an industrial implementer can gain insight that allows the algorithm discoveries to be applied to a particular industrial problem. The industrial problem may be totally new due to an integrated circuit break through. It seems to me this experimentation will have much wider applicability to heuristic search questions in general.

b. Substrate Organization Comparison.

It is possible to evaluate new or competing substrate organizations by holding the layout program and bench mark circuit as constant as possible, and then varying the substrate organization upon which the circuit is layed out. Finally the resulting layouts can be compared. The comparison may involve considerable ingenuity and intuition, but if the various details are published, people will be able to decide for themselves. I believe too many product line substrate organization decisions in semiconductor companies are made according to how the advertising copy will read.

c. Testing and Improvement of the Bench Mark Layout Environment.

It is probable that once the bench marks are used in controlled experiments, changes will be necessary in the bench mark environment and circuits. This process will be more efficient if controlled experiments are used to assist in deciding which changes to make.

d. Evaluation of Heuristics and Approximations in a Controlled Environment.

There is considerable disagreement on which approximations can safely be made during especially placement and which can not. For example, it is still open whether it is sufficient to treat all nets as if they contained only two pins with suitable adjustments. Many developers believe using the net half perimeter is better. A few even believe only exact Steiner tree net measurements are acceptable. Experiments that keep everything constant but change the various low level calculations in the evaluation function would be interesting.

e. *Experiments to systematically classify circuits types.*

This area could be of particular value to commercial developers for products aimed at particular circuit types. It would be interesting to see the results of experiments presented by circuit type that measure in a controlled manner various layout parameters. If it is possible to improve layout results given a narrow range of allowed circuits types, this could have large economic value.

7. Example Third Metal Layer Substrate Evaluation Experiment

In order to illustrate the experimental approach, an experiment measuring the possible utility of including a third layer of metal as part of the substrate organization. The experiment holds all factors fi xed and then rewires a placement (at the global routing level) using three metal layers. It is interesting because the improvement from a third layer metal is seemingly not large. This experiment can, of course, be criticized since remapping a placement optimized for two layer metal may not be valid. The purpose of this example is to encourage more accurate experimentation.

A number of experimental master slice ASIC circuits use a third metal layer (metal3). The most obvious use dedicates metal3 to long vertical cell row feed through wires (actually feed over) that connect pins on rows separated by an intervening cell rows. The connection pattern illustrated in fi gure 3 would use the third metal layer but the patterns shown in fi gures 1 and 2 would not. The advantage of dedicating metal3 to long vertical connections is that such wire cause no congestion in the regions they cross. The alternative of dedicating metal3 to horizontal wires suffers from the problem that the vias connecting metal1 and metal3 horizontal wires in, for example, channel jogs block scarce vertical metal2 feed through grids. There may be some better mixed direction metal3 use not considered here.

Known good bench mark circuit placements can be used to measure the possible benefit from dedicating metal3 to vertical wires. The results discussed here use placements made for two layer metal by the Timberwolf system [22] to see how much track and feed through reduction is possible with simple rewiring. The Timberwolf 17 row primary1 and 23 row primary2 placements from the 1988 workshop are used The next step would possibly be to repeat this experiment using placements from a placer modified to maximize vertical wiring. But since maximizing vertical wiring causes an increase in real wire length, the rewiring scheme discussed here may actually show the limit of potential size reduction from the addition of a third metal layer.

The following data measures YAL coded Timberwolf placements and uses a placement level routing estimate approach described in [14]. The routing estimate decomposes each net into pin pairs using the minimum spanning tree and then assumes the connection can be made with at most one via. Steiner MST decomposition does not materially change the results [14]. This measurement scheme gives an optimistic estimate of the gain from a third metal layer since if extra vias are needed, the long metal3 wires will cause additional congestion.

Tables 1 and 2 compare the rewired vertical feed through numbers for the primary1 and primary2 circuits to the original two metal layer wire requirement where the spanning tree pin pair decomposition used exact physical wire length rather than weighing to maximize vertical wires. For exact physical wire length, a wire connecting rows separated by two intervening wiring channels and one cell row (see fi gure 3) along one vertical grid is as distant as a pin connecting to another pin 60 grids (20 gates) distant along the same row (see figure I). For the three metal layer vertical wire maximizing measurement, the spanning tree distance metric assumes the same vertical connection is equivalent to a pin only six grids (two gates) distant. Further decrease in vertical wire cost leads to no increase in horizontal wiring and no channel height decrease at least for the benchmark circuits. Column two contains the required vertical feed through wire number for each row for the two layer metal case. Column three gives the percentage of available feed through used. Available means non blocked metal2 grids and assumes no vertical feed through would be lost to congestion problems. To understand the meaning of column three, consider the most congested primary2 circuit row 17. The required 1039 feed through wires is 23.7 percent more than available because the 23 row placement row length is 1005 grids of which 165 are blocked by vertical intracell wiring or unconnected pins (only 14). The percentage used is 1039/840 or 123.7. A value of more than 100 percent means a possible feed through shortage that requires either additional feed through cells which would lengthen every row or a better global routing approach. Column four gives the number of long vertical wires possibly movable to metal3. Column five gives the percentage of metal3 feed through that would be used above each row if every possible wire were moved to metal3.

Since the primary1 circuit has available feed through, adding a third metal layer does not lead to much area reduction. Nearly two thirds of the available metal3 feed through grids go unused over most internal rows. Global routing that makes maximum use of metal3 would, at a minimum, reduce the number of required track from 226 to 209 or 7.5 percent. The trade off is a 10.4 percent increase in total wire length. This assumes the bench mark system uses the same pitch for all metal layers. A better global router could reduce the required maximum channel density by moving wires into horizontal channel sections with unused grids below the maximum channel requirement thereby reducing channel density peaks, but then the same router could probably also reduce

peaks in the two layer metal case [10].

Table 3 shows the saving for the primary2 circuit The advantage of a third metal layer for the primary2 circuit is potentially larger since there is a shortage of metal2 vertical feed through tracks (see table 3 column 3). It is possible that a global router that changes connecting pin pairs to reduce vertical feed through requirements by using wiring channel sections below maximum channel density could possibly route the rows with a shortage of feed through. If the improvement is not possible, a global routing that maximizes use of metal3 would, at a minimum, reduce horizontal routing channel tracks from 544 to 486 or 10.7 percent. The trade off is a 15.1 percent increase in total wire length. Column fi ve shows that more than half of the available metal3 tracks are used over the central rows and also that around half metal2 feed through tracks now are unused. If all the master slice row lengths really need to be increased by the 199 grids required by row 17 in the two metal layer case, addition of a third metal layer can reduce total area by 25 percent (1005/1204*486/544). Of course, the actual decision to add a third metal layer will be determined by manufacturing and electrical considerations.

8. Impossibility of Cell Generation

Cell generation (sometimes called cell layout synthesis) differs from oilier layout problems that assign macro cells and wires to substrate locations in the sense that it is not an optimization problem solvable in principle if computers were fast enough to overcome the limitations of NP completeness. Rather, it is an unstructured intuitive problem that is no easier than the general human intelligence problem. Examples of this general problem are natural language understanding, scientific discovery, and legal decision making, but not chess. Since it is not possible to abstract out a formal problem and be sure solving the formal problem also solves the intuitive cell synthesis problem, usable cell generation computer programs are no more likely than usable natural language understanding programs. Fortunately, human cell design is not particularly diffi cult since cells are small and at most a few hundred are needed for any given product line.

There are a number of ways to view the general intelligence problem. One view characterizes it as an intuitive problem for which all human scientific and cultural background knowledge can be brought to bear. Another characterization is as a problem that requires balancing of conflicting requirements such that any algorithm for balancing can be beaten by intuitive knowledge. Another viewpoint defines the problem as one of meaning rather than syntactic understanding ([21], p. 31-32). According to this view, is necessary to understand in an intuitive, gestalt, unconscious, subconscious, and even formalist sense all at once.

The cell synthesis problem is equivalent to the general intelligence problem since it requires the balancing of a number of conflicting requirements, is amenable to scientific progress since completely new circuits designs can be discovered using the background knowledge embedded in solid state physics, requires subtle intuitive judgment, and fi nally is a problem whose solution is improved by human trial and error experience. The following conflicting requirements must be balanced and 'synthesized' (used in the informal sense here) into a circuit represented fi nally as a mask set These conflicting criteria, which are embedded in the conceptual matrix of circuit electrical parameters (resistance, capacitance, stability, etc.) and manufacturing process technology, must be combined and balanced. Cells need porosity for feed through wires yet can not be too porous or area and speed will be inefficient. Cell interfaces must be well behaved since future electrical environment is only predictable in an intuitive sense, yet use of too many buffers will again result in inefficient cells. The advantages of the small area and delays of pass transistors must be balanced against unpredictable delays and a tendency to become bidirectional (an expert would run spice experiments to determine if pass transistors are acceptable in a given place).

Cell power distribution simplicity (resulting in small area) must be balanced against substrate area, electrical field, and magnetic field characteristics of an entire chip's power bus distribution system. Geometric design rules that require large area must be balanced against smaller but potentially unreliable intuitive circuit feature patterns in an environment of changing process manufacturing methods and parameters. Second, the recent rapid growth in ASIC technology switching delays minimization and area reduction (outside of process feature size shrinkage) has been made possible by scientific discoveries [1]. Finally, experienced cell designers produce both smaller and faster cells than experts with similar background but no experience.

9. References

- 1. Dell'oca, C. Gate array technology. *Proceedings IEEE ICCD*, 1988, 296-299.
- 2. Dunlop, A. E., Preas, B., and Roberts, K. Cell-based layout benchmarks. *Proceedings 24th Design Automation Conference*, 1987, 318.
- 3. Garey, M. R., and Johnson, D. S. *Computers and Intractability. A Guide to the theory of NP Completeness.* W.H. Freeman, San Francisco, 1979.
- 4. Hanan, M., and Kurtzberg, J. Placement Techniques. In M. Breuer (Ed). *Design Automation of Digital Systems. Vol. 1,* Prentice Hall, 1972, 213-282.
- 5. Hanan, M., Wolff, P., and Agule, J. A study of Placement Techniques. J. Des. Automat. Fault Tolerant Comput. 2, 2 (May 1978) 28-61.
- 6. Hartoog, M. R. Analysis of placement procedures for VLSI standard cell layout. *Proceedings 23rd Design Automation Conference*, 1986, 314-319.
- 7. Igusa, I., Beardslee, M., and Sangiovani-Vincentelli, A. ORCA a sea-of-gates place and route system. *Proceedings 26th Design Automation Conference*, 1989, 122-127.
- 8. Johannsen, D. L. Silicon Compilation, in Seitz, C. L. (ed.), *Advanced Research in VLSI, Tenth Cal Tech Conference*, MIT Press, Cambridge, 17-36.
- 9. Lakatos, I. The Methodology of Scientific Research Programs. In Lakatos, I. and Musgrave, A. (Eds). *Criticism and the Growth of Knowledge*. Cambridge, 1970, 91-196.
- 10. Lee, K., and Sechen, C. A New Global Router for Row-Based Layout. *Proceedings IEEE ICCAD-88*, 1988, 180-183.

- 11. Lin, S. Heuristic programming as an aid to network design. *Networks*, 5(1975), 33-43.
- 12. Kuhn, T. S. *The Structure of Scientific Revolutions*. Princeton University Press, 1962.
- 13. Meyer, S. J. A constructive placement algorithm for logic arrays. *Proceedings IEEE ICCD*, 1983, 58-61.
- 14. Meyer, S. J. A new placement level wirability estimate. In preparation, 1990.
- 15. Meyer, S. J. CAD tool interchangeability through net list translation. In preparation, 1990.
- 16. Sasaski, G. H., and Hajek, B. The time complexity of maximum matching by simulated annealing. *J. ACM 35*, 2(*ApriI1988*), 387-403.
- 17. Popper, K. R. The Logic of Scientific Discovery. London, Hutchinson, 1959.
- 18. Preas, B. Benchmarks for cell-based layout systems. *Proceedings 24th Design Automation Conference*, 1987, 319-320.
- 19. Roberts, K. A. Automatic Layout in the Highland System. *Proceedings of ICCAD*, 1984, 224-226.
- 20. Roberts, K. A. Notes and results from 1987 UTMC bench mark routings. unpublished, 1987.
- 21. Searle, J. Minds, Brains, and Science. Harvard University Press, Cambridge, 1984.
- 22. Sechen, C. VLSI Placement and Global Routing Using Simulated Annealing. *Kluwer Academic Publishers*, Boston, 1988.
- 23. Suaris, P. R., and Kedem, G. Quadrisection: A new approach to standard cell layout Proceedings. *IEEE ICCAD-87*, 1987, 474-477.
- 24. YAL Reference Manual. On bench mark distribution tape from Micro Electronics Center of North Carolina, 1987, 1988.
- 25. Zhang, X., Pillage, L., and Rohrs, R. Efficient final placement based on nets-aspoints. *Proceedings 26th Design Automation Conference*, 1989, 578-581.

Tables and Figures

	Real Wire Distance		Two Rows Equal Two Gates	
Cell	Total	% of Metal2	Possible	%s of
Row	Crossing	if no	Metal3	Metal3
	Wires	Metal3	Wires	
1	188	57.7	33	6.9
2	243	73.3	48	10.1
3	223	70.8	68	14.1
4	286	85.5	120	25.1
5	279	84.3	144	29.9
6	283	83.4	177	36.9
7	291	88.2	195	40.5
8	324	94.1	208	43.2
9	334	95.7	192	40.0
10	307	89.7	176	36.6
11	324	94.2	153	31.8
12	284	86.3	137	28.5
13	263	82.1	119	24.8
14	252	77.2	102	21.3
15	246	74.7	78	16.2
16	227	68.8	50	10.4
17	220	67.9	24	5.0

 TABLE 1. Primary1 Metal3 Vertical Feed Through Wires

	Real Wire Distance		Two Rows Equal Two Gates	
Cell	Total	% of Metal2	Possible	%s of
Row	Crossing	if no	Metal3	Metal3
	Wires	Metal3	Wires	
1	478	63.9	128	12.8
2	633	79.4	221	22.0
3	715	90.2	297	29.5
4	700	88.0	346	34.4
5	802	97.4	408	40.6
6	833	100.7	430	42.7
7	938	110.5	448	44.5
8	896	107.2	500	49.8
9	923	91.7	520	51.7
10	890	107.4	535	53.2
11	858	105.0	583	58.0
12	839	104.6	560	55.7
13	953	117.0	574	57.1
14	901	108.6	558	55.5
15	944	113.7	537	53.4
16	983	116.5	553	55.0
17	1039	123.7	566	56.3
18	1008	117.2	576	57.3
19	929	109.5	535	53.2
20	847	102.1	445	44.2
21	821	97.5	359	35.7
22	755	91.8	208	20.7
23	549	74.8	94	9.4

 TABLE 2. Primary2 Metal3 Vertical Feed Through Wires



Figure 1. Along Channel Possible Connections



Figure 2. Cross Channel Connection



Figure 3. Cross Row One Via Possible Connections