



The Effects of Physical Design Characteristics on the Area - Performance Tradeoff Curve

Alice C. Parker, Pravil Gupta and Agha Hussain
Department of Electrical Engineering
University of Southern California
Los Angeles, CA, USA

Abstract

This paper describes two experiments designed to show the effects of wiring area and delay and unused area on final chip characteristics. An example behavioral specification is used to produce a range of automatically synthesized designs with varying constraints on cost and performance, using both pipelined and non-pipelined design styles. An analysis of chip layouts is performed, and recommendations for future high-level synthesis programs are given.

1 Introduction

When high-level synthesis research began, there were no VLSI chips. Design was assumed to be done with a fixed set of available modules [1]. In at least one case, these modules were assumed to be TTL chips [2]. Such chips and modules had a fixed cost, and wiring delays between chips were minimal compared to the processing delays on chip. Power consumption could easily be computed as the sum of the power consumption of individual chips, and hot chips could be cooled with a heatsink. Now, however, we are faced with a situation where high-level synthesis programs must design datapaths and controllers to fit on one or more VLSI chips. For such chips, a large portion of the chip area is consumed by wiring. Wire delays can be important. Given this situation, high-level synthesis programs must take a number of factors into account that were by and large ignored in the past.

2 Related Research

An early study of the effects of layout on the design curve was performed by Granacki and Parker [3]. ELF

This research described in this study was funded by the Semiconductor Research Corporation under SRC Contract Number 89-DJ-075.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

[4] considers wiring costs during synthesis. BUD [5] floorplans prior to high-level synthesis, taking into account wiring space and wire delays. McFarland [6] showed that BUD could not obtain a cost-performance tradeoff curve when physical factors were taken into account. Knapp [7] floorplans in order to improve designs. Chippe [8] predicts wire delays, given the RTL design. Although many synthesis researchers have demonstrated layouts of synthesized designs, no published tradeoff studies of actual layouts based on automatically synthesized designs exist.

3 Overview of the Experiments

This paper describes a set of experiments designed to determine the impact of some physical design parameters on the high-level synthesis process. Example filter specifications were used to synthesize a number of implementations at the register-transfer level with varying cost and performance using ADAM [9]. For non-pipelined designs the cost was measured as total area of the bounding box and performance as the delay through the active area of the chip. For pipelined designs, cost was measured as the bounding box, and performance as the delay between the initiations of new data into the pipeline.

The example chosen for this study was the AR lattice filter element, a design with a clear cost-performance tradeoff curve at the register-transfer level. Designs with inner loops and conditional branches are expected to have less well-behaved tradeoff curves at the register-transfer level, and the physical impacts might well be worse. The dataflow graph for this AR filter is shown in Figure 1.

For this experiment, we produced 6 non-pipelined and 6 pipelined RTL designs. We used MAHA [9] to generate the schedule for non-pipelined designs and Sehwa [9] for pipelined designs. MABAL completed the RTL designs, which were then translated to Seattle Silicon Chipcrafter format through our netlist translation and expansion program. For our layout style, each 16-bit module is constructed with 1-bit cells arranged in a fixed order, and this constitutes a functional block. For example, a 16-bit ripple carry adder is made up of a column of 16 1-bit full adder cells connected appropriately.




Figure 1: The AR Filter Dataflow Graph

In order to optimize the layout, first local rearrangement is done by Chipcrafter at the functional block level and then overall placement of these functional blocks is performed. We ran Chipcrafter with the OKI 1.2 micron, twin-well, double-layer metal CMOS ruleset and achieved layouts ranging from 20,000 to 30,000 transistors. For the non-pipelined designs we also generated the PLA controllers and they are included in the layouts. For pipelined designs we had to restrict our study to datapaths only because the required control signal generation software was not available. We do anticipate that the controller can have a significant impact on both area and time for pipelined designs with conditional branches.

For each design, we measured individual contributions to final chip area. We then assessed whether these layouts still fit our cost-speed tradeoff curve.

4 Non-Pipelined Results

The designs varied in parallelism. Information about the non-pipelined chips is given in Tables 1-4.

Design No.	No. of Control steps	Active Area $\times 10^6 \mu m^2$	Total Delay ns	Datapath Delay ns
1.	1	44.5	171	171
2.	4	23.4	NA	NA
3.	8	19.9	836	568
4.	10	12.4	1036	670
5.	16	12.0	1693	1161
6.	18	7.4	1781	1200

Table 1: Summarized Area - Delay Statistics of the Non-Pipelined Designs

Note that the functional area (raw cells plus internal interconnect of the function blocks) is proportional to the raw cell area since the layouts of the function blocks are quite regular. A quick check of the ratio of total wiring and unused area to raw cell area shows that in

all of the designs wiring/cell ratio varies from 1.6 to 2.6. Wiring delay constitutes 15-20 % of the total delay.

Since MABAL computes the number of two-point net equivalents, an increase in global wiring like that shown in design 5 might be predictable from the MABAL statistics. However, an examination of the number of global two-point net equivalents given by MABAL shows an average number of nets for design five. Further examination of the layout will be performed to determine whether placement and routing software problems caused an anomaly in the routing area, or whether the design is truly inferior.

Multiplier area dominates the design. Hence, the two designs with two multipliers have roughly equivalent areas but quite different performance. A simple analysis excluding multiplier areas shows that a tradeoff curve still exists for the remaining logic.

A cost-performance tradeoff curve for non-pipelined datapaths is shown in Figure 3. This curve shows the register-transfer design points, and the physical parameters when layout is considered. The register-transfer design points included raw cell area and raw cell delays. Although the tradeoff curve is not a smooth convex surface when physical factors are taken into account, there are no faster designs which are cheaper, or slower designs which are also larger.

Also observe that in the case of the most parallel design (Figure 3) the actual delay is better than the predicted delay, while in all the other cases it is the opposite. The reason is that this particular design is basically a combinational circuit and the critical path in this circuit is less than the sum of critical paths in the modules, which is the predicted delay. In the other cases, due to wiring delays, actual delay is always more than the predicted delay.

Note that the multiplexer area is large and dominates the adder area. However, adders are not shared in any case where multiplexer area increases more than the decrease in the adder area. Multiplexer area does not steadily increase as sharing increases.

The most parallel non-pipelined layout is shown in Figure 2.

5 Pipelined Results

Information about the pipelined layouts is given in Tables 5-8.

The tradeoff curve in Figure 4 is for pipelined designs. The curve shows the register-transfer level design points and the actual points considering the layout. As mentioned earlier in the paper, this curve considers the datapath only and controller area and delay are not included. In these designs the basic clock period remains almost the same and therefore the delay depends on the initiation interval of the circuit. Again, the multiplier is a dominant effect. The register area drops as the design becomes more serial and less values must be stored simultaneously. Multiplexing requirements are somewhat

Design Number	Total Raw Cell Area A μm^2	Total Func. Block Area B μm^2	Total Active Area C μm^2	Internal Wiring B-A μm^2	Global Wiring C-B μm^2	Total Wiring C-A μm^2
1.	16935497	30440600	44532021	13505103	14091421	27596524
2.	7865950	13154516	23393849	5288566	10239333	15527899
3.	5578259	9152214	19902580	3573955	10750366	14324321
4.	3881830	5981386	12354541	2099556	6373155	8472711
5.	3681145	5528098	12047529	1846953	6519431	8366384
6.	2756086	3793329	7437828	1037243	3644499	4681742

Table 2: Wiring and Cell Area Statistics of the Non-Pipelined Designs

Design Number	Adder			Multiplier			Controller Area μm^2
	No. of Modules	Functional Block Area μm^2	Raw Cell Area μm^2	No. of Modules	Functional Block Area μm^2	Raw Cell Area μm^2	
1.	12	825907	715860	16	29416590	16042040	34157
2.	4	275302	238620	6	11031221	6015765	61597
3.	2	137651	119310	4	7354148	4010510	89475
4.	2	137651	119310	2	3677074	2005255	135065
5.	1	68826	59655	2	3677074	2005255	157927
6.	1	68826	59655	1	1838537	1002628	149799

Table 3: Area Statistics of the Non-Pipelined Layouts by Module Type

Design Number	Register			Multiplexer		
	No. of Modules	Functional Block Area μm^2	Raw Cell Area μm^2	No. of Modules	Functional Block Area μm^2	Raw Cell Area μm^2
1.	2	163946	143440	-*	-	-
2.	6	491839	430320	19 [†]	1294557	1119648
3.	6	491839	430320	17 [‡]	1079101	928644
4.	8	655785	573760	12 [§]	1375811	1048440
5.	7	573812	502040	8 [¶]	1050459	956268
6.	6	491839	430320	8	1244328	1113684

Table 4: Area Statistics of Non-Pipelined Layouts by Module Type

constant over the designs. Design 5 is somewhat of an anomaly, as in the non-pipelined design. The layout obtained for design 5 was initially inferior to design 6 in area. However, an examination of the design showed that the layout was not very compact, and Chipcrafter was allowed to iterate longer than for the other designs. Hence, the wiring area actually dropped. Examination of the other designs did not reveal any such obvious opportunities for optimization.

The fastest design could not take advantage of a large

combinational logic block like the fastest non-pipelined design did. Therefore the clock rate is limited by the slowest module and the wiring delay. Hence like all the other designs this is also slower than the RTL design indicated. [htb] [htb]

Design Number	No. of Control steps	Active Area $\times 10^6 \mu m^2$	Initiation Interval ns
1.	1	64.1	63
2.	4	20.6	265
3.	6	18.5	410
4.	8	12.9	529
5.	12	11.0	830
6.	16	10.0	1185

Table 5: Summarized Area - Delay Statistics of the Pipelined Designs

*Does not require any muxes.

[†]Consists of five 2to1 muxes and fourteen 3to1 muxes.

[‡]Consists of seven 2to1 muxes, one 3to1 mux and nine 4to1 muxes.

[§]Consists of four 2to1 muxes, three 3to1 muxes, one 5to1 mux, two 6to1 muxes and two 8to1 muxes.

[¶]Consists of two 2to1 muxes, one 3to1 mux, two 5to1 muxes, two 7to1 muxes and one 9to1 mux.

^{||}Consists of four 2to1 muxes, one 3to1 mux, one 5to1 mux, one 11to1 mux and one 16to1 mux.

Design Number	Total Raw Cell Area A μm^2	Total Func. Block Area B μm^2	Total Active Area C μm^2	Internal Wiring B-A μm^2	Global Wiring C-B μm^2	Total Wiring C-A μm^2
1.	22638940	36964293	64117478	14325353	27153185	41478538
2.	6648319	10391463	20646193	3743144	10254730	13997874
3.	5779969	8609260	18549425	2829291	9940165	12769456
4.	3648549	5731758	12911088	2083209	7179330	9262539
5.	3544358	5615539	11033129	2071181	5417590	7488771
6.	2803491	3864570	10041426	1061079	6176856	7237935

Table 6: Wiring and Cell Area Statistics of the Pipelined Designs

Design Number	Adder			Multiplier		
	No. of Modules	Functional Block Area μm^2	Raw Cell Area μm^2	No. of Modules	Functional Block Area μm^2	Raw Cell Area μm^2
1.	12	825907	715860	16	29416590	16042040
2.	3	206477	178965	4	7354148	4010510
3.	2	137651	119310	3	5515611	3007882
4.	2	137651	119310	2	3677074	2005255
5.	1	68826	59655	2	3677074	2005255
6.	1	68826	59655	1	1838537	1002628

Table 7: Area Statistics of the Pipelined Layouts by Module Type

Design Number	Register			Multiplexer		
	No. of Modules	Functional Block Area μm^2	Raw Cell Area μm^2	No. of Modules	Functional Block Area μm^2	Raw Cell Area μm^2
1.	82	6721796	5881040	-**	-	-
2.	19	1557489	1362680	20 ^{††}	1273350	1096164
3.	22	1803409	1577840	15 ^{††}	1152590	1074936
4.	7	573812	502040	11*	1343221	1021944
5.	7	573812	502040	9 [†]	1295828	977408
6.	6	491839	430320	8 [†]	1465369	1310888

Table 8: Area Statistics of Pipelined Layouts by Module Type

6 Conclusions

The non-pipelined layouts showed a cost-performance tradeoff between the most serial and the most parallel designs. However, the cheapest design was far slower than the fastest one, but not proportionately smaller. Buffering I/O with a RAM instead of registers subsequent to these experiments halved the area of the smallest design.

An intermediate design(5) with partial serialization

**Does not require any muxes.

††Consists of eight 2to1 muxes, three 3to1 muxes and nine 4to1 muxes.

†Consists of five 2to1 muxes, two 3to1 mux, four 4to1 muxes, two 5to1 muxes and two 6to1 muxes.

*Consists of three 2to1 muxes, three 3to1 muxes, one 5to1 mux, two 6to1 muxes and two 8to1 muxes.

†Consists of three 2to1 muxes, one 4to1 mux, one 6to1 muxes, two 7to1 muxes and two 8to1 mux.

†Consists of four 2to1 muxes, one 3to1 mux, one 6to1 mux, one 11to1 mux and one 16to1 mux.

was not significantly smaller than a more parallel design(4).

The pipelined designs showed similar results. The cheapest design was about one sixth of the area of the fastest design, but was about 20 times slower.

Twelve layouts were produced to draw these conclusions, a large effort but not comprehensive enough to allow generalization of the results. Examples with a larger variety of functions, inner loops and conditional branches must be processed. We feel that such examples are more prone to scheduling infeasibilities and allocation problems, and may not tradeoff as easily as the AR filter. We are encouraged that the tradeoff curves did exist for the physical designs, but are also aware of the effects of wiring and unused area and wiring delay on the final performance and area of each design. These effects point to the conclusion that high-level synthesis programs must take into account the effects of layout if




Figure 2: Layout of the most Serial Non-Pipelined Design




Figure 3: Cost-performance tradeoff curve for a 16-bit Non-Pipelined AR Filter Datapath Element

they are to produce designs of high quality.

References

- [1] M. Barbacci and D. Siewiorek. The CMU RT-CAD System: An Innovative Approach to Computer Aided Design. In *American Federation of Information Processing Societies Conference Proceedings, Vol. 45*, pages 643–655. Amer. Fed. of Information Processing Societies, June 1976.
- [2] L. Hafer and A. Parker. Automated Synthesis of Digital Hardware. *IEEE Transactions on Computers*, C-31(2):93–109, February 1981.

The authors would like to acknowledge Kayhan Küçükçakar for his help in running these experiments, and the feedback provided by Michael McFarland during discussions of these ideas.

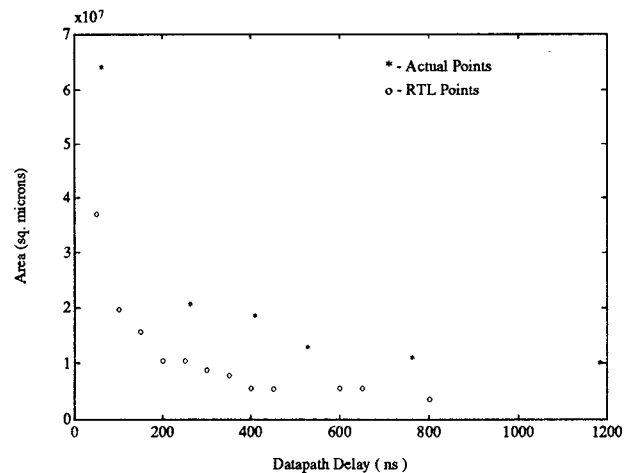


Figure 4: Overall Cost-performance tradeoff curve for a 16-bit Pipelined AR Filter Datapath

- [3] J.J. Granacki and A.C. Parker. The Effect of Register-Transfer Design Tradeoffs on Chip Area and Performance. In *Proceedings of the 20th Design Automation Conference*, June 1982.
- [4] E. Girczyc. *Automatic Generation of Microsequenced Data Paths to Realize ADA Circuit Descriptions*. PhD thesis, Carleton University, 1984.
- [5] M. McFarland. Using Bottom-Up Design Techniques in the Synthesis of Digital Hardware from Abstract Behavioral Descriptions. In *23rd Design Automation Conference*, pages 474–480. ACM, IEEE, June 1986.
- [6] Michael McFarland. Reevaluating the Design Space for Register-Transfer Hardware Synthesis. In *International Conference on Computer-Aided Design*, pages 262–265. IEEE, November 1987.
- [7] David W. Knapp. Feedback-Driven Datapath Optimization in Fasolt. In *Proceedings of the 26th Design Automation Conference*. ACM, IEEE, July 1990.
- [8] F. Brewer and D. Gajski. Chippe: A System for Constraint Driven Behavioral Synthesis. *IEEE Trans. on Computer-Aided Design*, 9(7):681–695, July 1990.
- [9] R. Jain, K. Kucukcakar, M. J. Mlinar, and A. C. Parker. Experience with the ADAM Synthesis System. In *Proceedings of the 26th Design Automation Conference*. ACM/IEEE, June 1989.