25

Towards Better Accounting of Physical Design Effects in High-Level Synthesis

P. Jha † *C. Ramachandran* ‡ *F. Kurdahi* § *N. Dutt* †
†*ICS Department, University of California, Irvine, CA 92717, USA*
§*ECE Department, University of California, Irvine, CA 92717, USA*
‡*Cadence Design Systems, San Jose, CA 95134, USA*

Abstract

High level synthesis has long relied on *point models* for RT level components, where the area and delay of a component are assumed to be constant for a given style. Nowadays, many synthesis algorithms attempt to incorporate physical design information into the design process so as to better guide the synthesis tasks. In this work, we explore the combined effect of style and aspect ratio variations on the area and delay of RT level components, and on RT level designs which use such components. Our results indicate that point models are inadequate for use in the synthesis process due to the large variations in the area and delay that occur when component styles and aspect ratios are varied. We believe that our results have some deep implications with respect to the flow of the design tasks during high level synthesis.

1 INTRODUCTION

High-level synthesis (HLS) has long relied on *point models* for RT-components that assume fixed area and delay values for a given component style. However, aspect ratio variations alone can result in substantially different area-delay characteristics for a component, and thus may result in large area-delay variations for a complete design composed of the same set of components by simply changing the aspect ratios and the ensuing feasible floorplans. In this work, we attempt to realistically explore the effects of varying aspect ratios on the area and delay of complete RT-designs generated by HLS using generic components.

In particular, we are interested in finding answers to the following questions.

1. What are the 1st order factors (parameters) that determine and/or influence the size, composition, and distribution of the "real" design space of a RT-component?

2. How would the nature of the design space of individual components (from 1) influence the overall design space of a "typical" complete design composed of several RT components?

3. The main question we wish to address is: are these "point models" of RT-components adequate for characterizing the resulting space of complete design implementations using

G. Saucier et al. (eds.), *Logic and Architecture Synthesis*

these RT-components? In other words, how much penalty in accuracy is incurred by using the "point model" instead of the "comprehensive model"?

In order to answer these questions, we performed an extensive set of experiments. We present sample results showing the area-delay variations at RT component level (Section 2) and at the complete design level (Section 3). Please refer to [7] for detailed results.

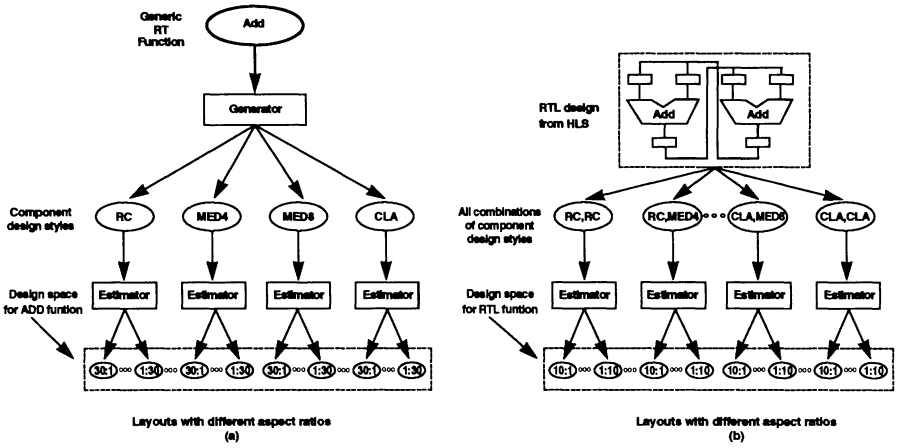

Figure 1 Experimental setup for design space exploration of an RT component

2 EFFECT OF ASPECT RATIO VARIATION ON RT-COMPONENTS

In the first set of experiments, depicted in Figure 1(a), we varied the aspect ratios and styles for an individual RT-component and observed the resulting variations in the component's area and delay attributes.

1. We generated several ADDER(a commonly used generator) components by varying the styles of designs for each bit-width, and repeated this process for different bit-widths. In particular, we generated:

   - Ripple Carry adder (RC),
   - MED4 adder, (i.e., 4-bit carry-lookahead blocks rippled)
   - MED8 adder, (i.e., 8-bit carry-lookahead blocks rippled)
   - Full carry-lookahead adder (CLA), and
   - Carry save adder (CS).

2. For each implementation, we first estimated the area and delay of the corresponding design netlist assuming a unity aspect ratio (i.e. square layout). This essentially gave us a "point model" estimate. The estimation was done by first generating a standard

cell netlist of each implementation, and then used the highly accurate predictors LAST and TELE to obtain estimates of its layout area and delay.

3. We also used the "comprehensive model" to generate area/delay values for different aspect ratios of each RT-component. This capability is built into our model because it employs the constructive-analytical technique described in [1][2] when estimating both area and delay.
4. To analyze the effects of aspect-ratio variation on each RT component, we plotted the results of the experiment in several ways and attempted to capture the overall effects with figures indicating percentage variations in area and delay for each RT component, with respect to different design styles (e.g., ripple-carry or carry-lookahead) and bit-widths.

The area-delay variations across different aspect ratios for 16 bit adder are summarized by the graph in Figure 2(a). In this graph, the point model (corresponding to designs with unity aspect ratios) are represented by big circles.

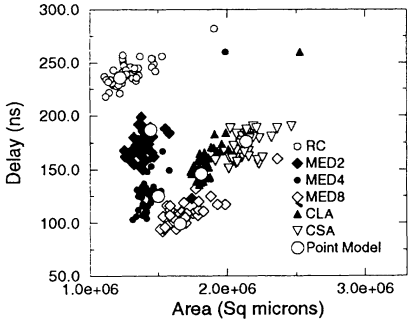

(a)

Style		Timer	Shift–add mult
Ripple carry	*area*	13.33	14.02
	delay	51.68	27.68
4–bit CLA	*area*	23.51	13.43
	delay	59.36	35.79
8–bit CLA	*area*	20.88	16.71
	delay	53.88	32.95
Full CLA	*area*	17.28	11.70
	delay	51.97	33.12
Over all designs	*area*	27.86	26.58
	delay	92.37	64.47

(b)

Figure 2 Area and delay variations with varying aspect ratios (a) 16-bit adder (b) some benchmark designs

From the area-delay graphs of Figure 2(a), we observe large variations in area and delay. In most cases, the variations exceed 30% in either area or delay by simply changing the component's aspect ratio, and hence one can easily see that the point models are not adequate for representing the RT-component attributes, and for use in High-Level Synthesis because:

1. RT components are building blocks that will be floorplanned together to form a design layout. Thus, the final aspect ratio (and hence the are-delay attributes) of each component cannot be known until physical design is started. However, point models implicitly assume a single aspect ratio configuration which yields a single area-delay

value for each component – in other words, they assume that area and delay do not vary under different aspect ratio configurations.

2. Since point models are not adequate for modeling the metrics of RT component, they cannot be used as a basis for module set selection prior to, during, or after scheduling and allocation. Instead, one must examine the full design space of these components, or alternatively, include the physical design information into the synthesis process. The latter approach means that component style *and aspect ratio* must be determined at the same time in order to obtain reliable measures of the attributes of the components.

3 EFFECT OF ASPECT RATIO VARIATIONS ON COMPLETE RT-DESIGNS OBTAINED FROM HIGH-LEVEL SYNTHESIS

In the second set of experiments, depicted in Figure 1(b), we took the output of high-level synthesis applied to some benchmark descriptions, and studied the effects of varying component styles and the design's overall aspect ratio on the area and delay of two examples: Shift-add multiplier [3] and an industrial timer example. The objective of these experiments was to study the overall RT area-delay design space as populated by different RT-component styles (e.g., CLA adder and RC adder), and by varying the overall aspect ratio of the final RT design, as shown in Figure 1(b). That is, our goal was to compare the effectiveness of the "point" model (using fixed aspect ratios) for covering the complete design space represented by the "comprehensive" model. The results are summarized using graphs that plot area vs. delay, aspect ratio vs. area, aspect ratio vs. delay, and aspect ratio vs AT, as well as tables that show percentage variations in area and delay attributes with respect to the "point" model for the complete RT designs.

The experiments were run in the following manner. For each RT design generated by a HLS tool ([4]) from a benchmark description :

1. We used the "point model" to generate the design space with a fixed aspect ratio using typical RT components, including layout and wiring which were estimated using LAST and TELE.
2. We then modified the aspect ratio of the final design and generated the RT design space for different combinations of RT component styles/implementations, and included the effects of wiring and PD. To make the comparisons realistic, we restricted the aspect ratio variations to between 1:10 through 10:1.

The results for the timer and shift-add multipler benchmarks are shown in Figure 3(a) and 3(b) respectively. Note that the "point model" (with unity aspect ratio) is represented by big circles in these figures. The percentage variation in area and delay with respect to the "point model" for the timer and the shift-add multiplier designs are shown in Figure 2(b). From these figures, we make the following observations:

- In all the cases, we observed a large size design space (as noted from the amount of variations in area and delay from Figure 2(b)). This design space tends to be densely-populated in some regions and sparse as area and delay increase. However, we note that the size of the densely populated region is large, indicating several not-so-obvious

design alternatives for similar area/delay constraints. Furthermore, a seemingly inferior design point on the traditional Area-vs-Delay curve (i.e., a design with inferior area and inferior delay) may turn out to be the *only* feasible design due to the shape constraint in the final floorplan and layout.

- In comparing the point models to the overall design space, we observe that the space defined by the set of points in the "point model" is a small subset of the overall design space(Figure 3).
- Given any target delay or area for the system within a certain range, it is possible to find *some* design point which is quite close to that target. This stems from the fact that the design space is quite dense for a wide range of area and delay values. More analysis is needed to be able to quantize this density in general.
- We observe more swing in delay than area (2-3 times or more), which indicates that in contrast to area, delay is quite sensitive to aspect ratio variations (this agrees with the results in [5]).
- The Aspect-vs-Delay graph indicates that designs using the MED* & CLA adders have close delays, while design using the RC adders are clearly slower.
- The designs with MED adders are generally the best for medium-to-high performance, while RC-based designs dominate for area-efficiency, however,
- The A*T figure of merit gives almost equal scores to all implementations, except RC which is higher in most cases. What this indicates is that it may not be worth the small savings in area to go with RC-based designs because of the drastic degradation in performance which would result from such a choice.

Based on these observations, the experimental results show conclusively that aspect ratio variations cannot be ignored during HLS, since variations are quite large and that the "point model" is indeed restricting the design space. We need to consider the "comprehensive model" that factors in aspect ratio variations in order to capture a more realistic design space.

4 CONCLUSION

In this work, we attempted to explore the effects of varying aspect ratios on the overall RT design space. Although it is intuitively obvious that aspect ratio variations of RT-components affect the area-delay attributes of complete RT-level designs, we provided conclusive data derived from a fairly extensive set of experiments to verify this fact. We presented experimental data for aspect ratio variations for complete RT-level designs generated by HLS tools from some benchmarks. The variations in area and delay were observed to be significantly large to warrant early inclusion in HLS decisions.

Indeed, we believe that the results we obtained have some deep implications on the "traditional" flow of HLS design tasks where scheduling and allocation are subject to some area and/or delay constraints. Our results raise two interesting questions related to this approach:

1. Given these large variations in the final area/delay, how can one ensure that the area/delay constraints initially imposed on the overall design are indeed satisfied when the design is laid out? This is especially relevant when the synthesized design is part

of a chip which contains other blocks, thus the shape of the final design is not exactly known in advance.

2. Does it make sense to assume "point models" during HLS? Why constrain the HLS to assume a certain (area, delay) point for each component type? Why not have HLS optimize (or select) preferred (area,delay) for each component type since we can find *some* implementation which can be very close to those specs? This argument can be extended to more recent synthesis systems (e.g. [4] and [6]) which incorporate a mix-and-match strategy of component selection, except that here only the aspect ratio is the source of variation in results.

Clearly, these and other interesting questions remain to be answered in light of the results we presented. Future work needs to address the issue of how exactly to abstract our this information for use early in the design (for example High-Level Synthesis).

REFERENCES

[1] F. Kurdahi and C. Ramachandran, "LAST: A layout area and shape function estimator for high level applications," in *European Conf. on Design Automation*, February 1991.

[2] C. Ramachandran and F. Kurdahi, "TELE: a timing evaluator using layout estimation for high level applications," in *EDAC-92*, 1992.

[3] F. Brewer and D. Gajski, "An expert system paradigm for design," in *The 23rd Design Automation Conference*, 1986.

[4] L. Ramachandran and D. Gajski, "An Algorithm for Component Selection in Performance Optimized Scheduling," *International Conference on Computer-Aided Design*, 1991.

[5] C. Ramachandran and F. Kurdahi, "A combined topological and functionality based delay estimation using a layout-driven approach for high level applications," in *The Euro-DAC*, September 1992.

[6] P. Gutberlet, J. Muller, H. Kramer, and W. Rosenstiel, "Automatic module allocation in high level synthesis," in *EuroDAC*, 1992.

[7] P. Jha, C. Ramachandran, N. Dutt, and F. Kurdahi, "The effects of variations in component styles and shapes on high-level synthesis," *Technical Report 92-115, University of California, Irvine, USA*, December 1992.

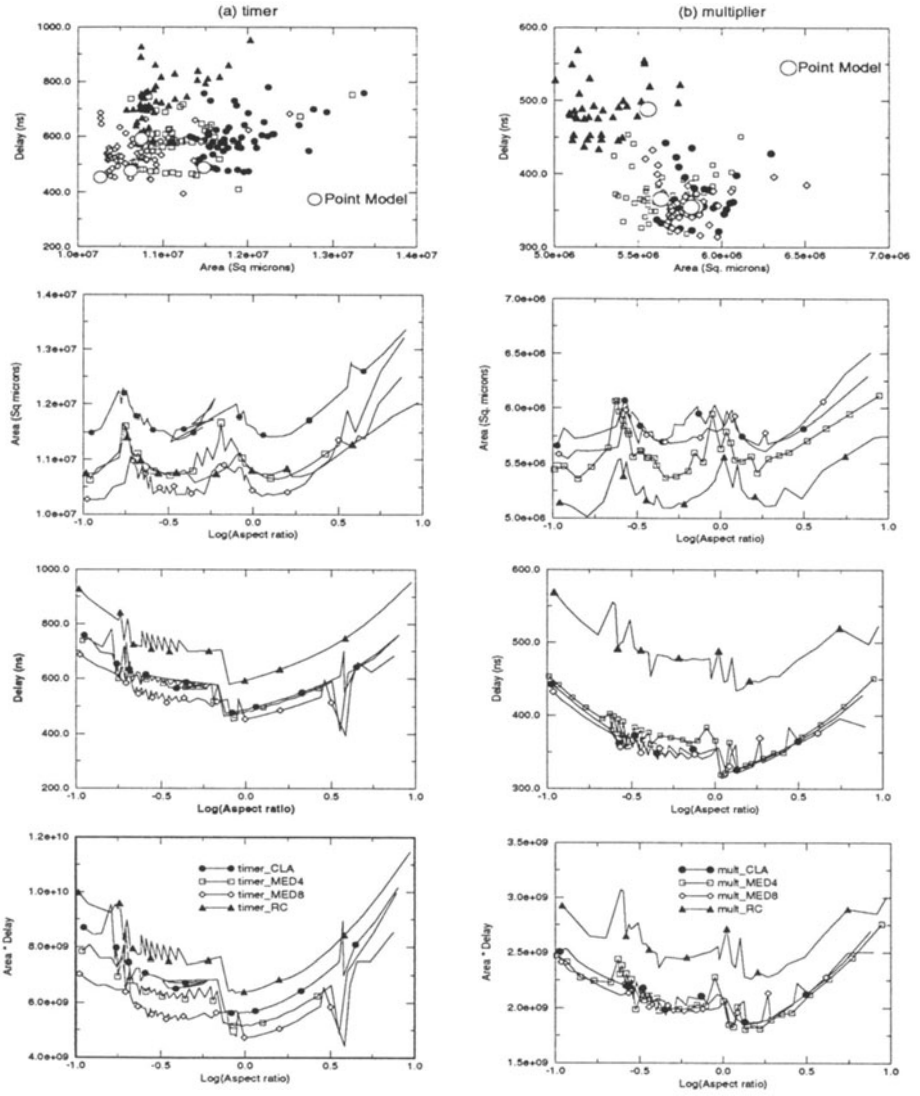

Figure 3 Area and delay variations of with varying aspect ratios (a) timer (b) multiplier. Both designs assume 16-bit words.