University of Pennsylvania
Department of Electrical and Systems Engineering
Electronic Design Automation

ESE535, Spring 2015 — Assignment # 4 Exercise — Wednesday, February 4

Due: Assign 4: Thursday, February 12, 10PM

This is just the exercise supplement to assignment 4. See the Assignment 3–6 for the project portion.

Exercise:

Our primary assumption for the heterogeneous multicontext computing array is that we will make one pass through the interconnect for every LUT evaluation (every level). As we will be time-multiplexing the upper levels of the interconnect, it will take multiple primitive clock cycles per LUT evaluation. Assuming it take N_{route} cycles to route one wave of signals through the interconnect, it will take: $N_{total} = levels \times N_{route}$ to evaluate the entire design through all $levels$ levels.

This certainly suggests that delay-oriented LUT mapping that minimizes the number of LUTs in the critical path is beneficial for delay minimization when mapping to this substrate.

One thing that is unfortunate about $N_{route} > 1$ is that sequentialized execution of a pair of LUTs at a PE can be performed on back-to-back cycles. That is, the second LUT can be evaluated on clock cycle after the first, not N_{route} cycles later. So, perhaps the design could be accelerated by clustering some LUT-to-LUT connections in the same PE and evaluating them without an intervening signal propagation wave across the interconnect. This might not help much, if there were isolated cases, but if entire level-to-level connections could be captured in PEs, perhaps we could reduce the number of network-crossing routing waves?

Specifically, you could image clustering together a number of LUTs that should be evaluated on a single PE and keeping those together during partitioning. So, for example, as long as the PE capacity is at least 2, you could cluster together a LUT and its successor to guarantee they were assigned to the same PE.

1. Building on the techniques we have seen (specifically Day 3 and 5), how could you formulate an optimization that would create clusters to reduce the number of network crossing routing waves?

   (a) Describe your revised evaluation model that allows a PE to evaluate multiple LUTs sequentially between network crossing routing waves and give an equation for the number of cycles required for netlist evaluation.

   (b) Sketch your mapping algorithm.

   (c) Estimate the complexity of your algorithm.

   (d) Describe the kind of delay reduction you expect your algorithm to achieve.

   - This is deliberately open-ended. We doubt there is a single optimal answer. We think there are some good answers—at least, answers that are much better than doing nothing.
   - For the sake of motivation, you might specifically think about PEs that hold four 4-LUTs and $N_{route} = 10$, but you should ultimately consider a solution that works for any size PE and N_{route}.
   - Is there a reason the four 4-LUTs is simpler (less computationally complex?) than the case where the PE capacity is larger?
   - To simplify this problem, you may want to initially assume that any number of LUTs may be evaluated in a single time-step within a PE.
   - How does it change things when you must consider serialization of the LUTs assigned to a single PE?
   - This will be forth 15 of 100 points on assignment 4. Suggestion is to not spend more than 1.5 hours on this exercise.