Parallel Placement for FPGAs Revisited

Cristinel Ababei

ECE Department, North Dakota State University, Fargo, ND, USA

cristinel.ababei@ndsu.edu

The runtime of classic sequential placement algorithms for FPGAs continues to represent a serious problem, aggravated by the continuous increase of FPGAs. The traditional way to parallelize the placement step is to use parallel distributed implementations run on a network of processors. This approach can suffer from significant communication and synchronization runtime overheads. To address that, we propose the use of multithreading for parallelization. The top level placement problem is decomposed into region-based placement sub-problems using four-way min-cut partitioning. These sub-problems are then processed in parallel by worker threads. The final solution, constructed using the results from all sub-problems, is further improved using a fast low-temperature annealing refinement step. Using this technique, we parallelize the simulated annealing based placement algorithm of VPR. The new parallel placement algorithm achieves an average speed-up of 2.5x using four threads, while the wirelength after placement and circuit delay after routing increased on average with 3.7% and 2.15% respectively.

ACM Categories & Descriptors: B.7.2 Placement and routing, D.1.3 Parallel programming, G.1.0 Parallel algorithms

General Terms: Algorithms, Performance, Design

Keywords: Parallel simulated annealing, FPGA placement, multithreading