Simultaneous depth and area minimization in LUT-based FPGA mapping

Jason Cong, Yean-Yow Hwang

Year of publication: *1995*
Area: *CAD*

This paper was a major step in the evolution of cut-based technology mapping algorithms.

The approach was started by the *FlowMap* algorithm, published a couple of years earlier, which was the first polynomial-time technology mapper that could claim theoretical optimality (in this case, in terms of logic depth). The key idea was to model the LUT mapping problem as the computation of a *minimum-height, K-feasible cut*, which is a *global* treatment, as opposed to the existing mapping algorithms of the time, which were based on *local* clustering.

In its original form, the *FlowMap* algorithm actually computed a *min-cut* on a transformed graph that guarantees minimum-height. While simple and efficient, this had a tendency of generating smaller LUTs, and, because each cut is computed independently, it could yield substantial amount of logic replication. Both effects contributed to higher area cost. The post-processing algorithms originally proposed to alleviate this problem, as well as the subsequent revision that considered relaxation of height requirement along non-critical paths, were both essentially local refinement, thus not an integral part of the global cut-based framework.

The *CutMap* algorithm in this paper made several key contributions. First it used a cost function to control the generation of the cuts. This not only allowed the exploration of all K-feasible cuts (not just the minimum cut), but also enabled the general modeling of secondary optimization objectives, or multiple concurrent objectives, in the cut-base mapping procedure. Second it tried to capture logic sharing globally through the pre-assignment of costs, reducing order dependency, and coupling quality with overall network structure more closely. Finally it identified *provably non-degrading* pruning criteria for its cut enumeration speed-up. While in today's light these may seem natural, at its time these strategies were not widely seen in technology mapping works; and indeed after the publication of this paper, a large number of papers in related fields adopted similar approaches. Last but not the least, *CutMap* performed well in practice; it became the base of comparison, in place of the original *FlowMap*, in many subsequent studies.

Cut-based multiple objective optimization has since become the norm of technology mapping both for and beyond FPGA world. This paper played an important role in making that happen.

Eugene Ding

DOI: http://doi.acm.org/10.1145/201310.201322