Improving Heterogeneous FPGAs by Maximizing Use of Hard Circuits

Peter Jamieson — University of Toronto

Supervisor: Jonathan Rose

Introduction

There is a dramatic logic density gap between Field Programmable Gate Arrays (FPGAs) and Application Specific Integrated Circuits (ASICs), and this gap is the main reason FPGAs are not cost-effective in high volume applications. Modern FPGAs (like that pictured in Figure 1) narrow this gap by including “hard" circuits such as memories and multipliers, which are very efficient *when* they are used. However, if these hard circuits are not used, they go wasted and have a negative impact on logic density. The largest impact is that 70% of the area of the soft logic of an FPGA is dedicated to the routing (Figure 2)

The general theme of this research concerns using hard circuits as efficiently as possible either by better architectures or better CAD mapping tools. The main contributions of my thesis are:

1. A new architectural concept, called shadow clusters, (described below) which significantly improve the area-efficiency of FPGAs.
2. The creation of a front-end RTL synthesis tool called Odin that comprehends heterogeneous FPGAs. Its quality is almost equal to an industrial front-end tool.
3. A new technique for using un-used multipliers as multiplexers, which results in a 7% improvement in logic density.

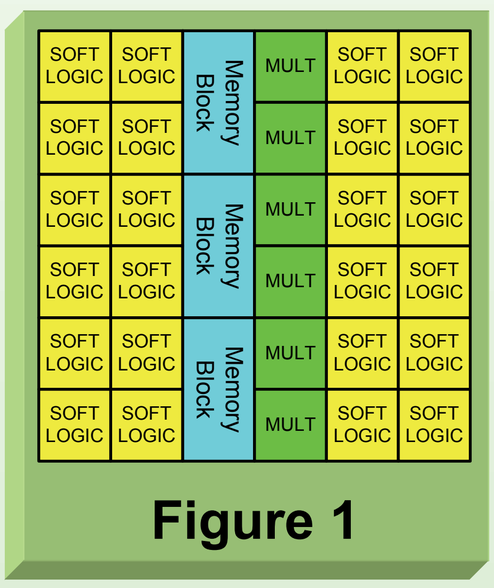

Figure 1

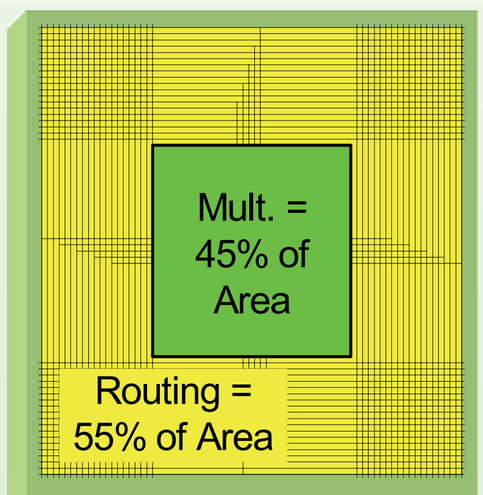

Figure 2

Shadow Cluster Design

A shadow cluster (Figure 3) is a standard FPGA logic “cluster” (typically consisting of a group of LUTs and flip-flops) that exists within the same logical tile as a hard circuit such as a multiplier. In the event that the hard circuit cannot be used by an application circuit, simple fabric-programmable multiplexers “swap” in the shadow cluster which can be used just like the regular soft logic fabric.

The key design decisions are architecting the hard circuit and the size of the shadow cluster. We design hard circuit’s routing to match the routing provided by the FPGA. This means that the routing architecture is experimentally designed to match the soft logic’s demand for input and output pins (pin demand). We match the hard circuits pin demand to that of the soft logic, meaning that the shadow cluster size has the same number of LUTs (N) as the soft logic cluster tile (Figure 3).

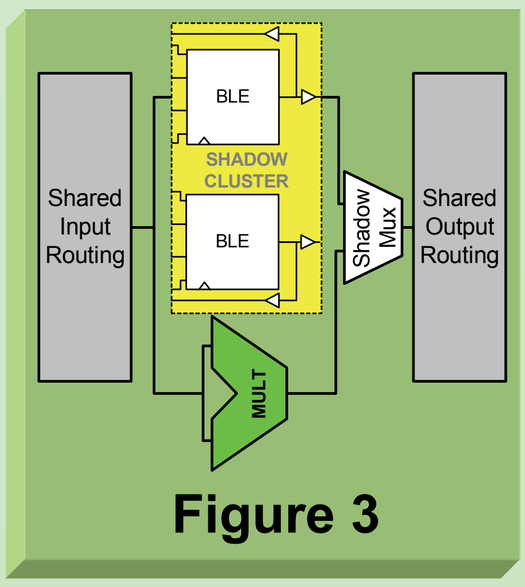

Figure 3

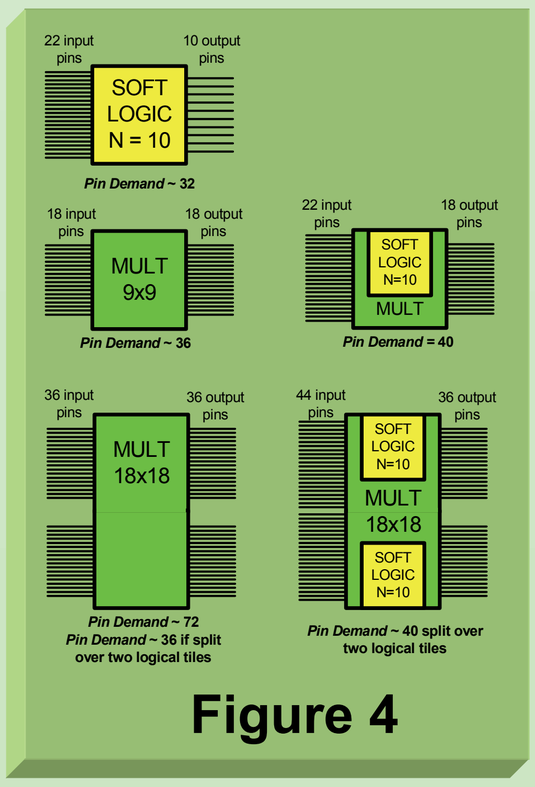

Figure 4

Measurement Methodology

Is to map benchmarks modeled as multipliers and clusters to FPGA sized by supply ratio (Example 1). FPGA area equals tiles used multiplied by tile area (Important to get proper area estimation of tiles).

Example 1 –Shadow cluster benefit

FPGA has a supply ratio of 1:4

Design 1 – Non Shadow FPGA

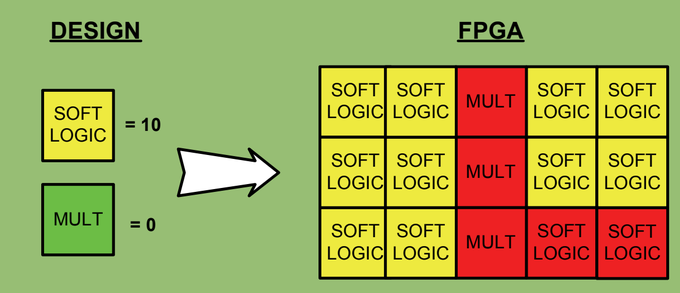

Design 1 – Shadow Cluster FPGA

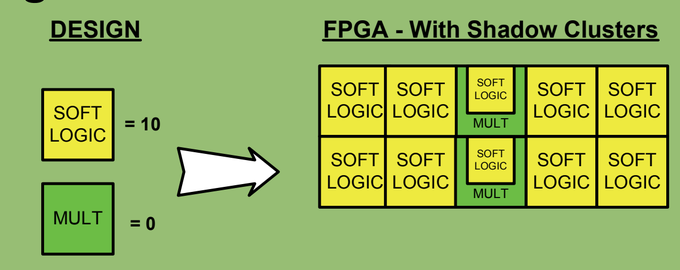

Shadow architecture is ~ 1/3 smaller

Results

We measure the effectiveness of adding shadow clusters to heterogeneous FPGAs with multiplier tiles.

Table 1 shows results for 10 benchmarks with average demand of 1 multiplier for every 15 soft logic tiles on an architecture with a supply of 1 multiplier for every 15 soft logic tiles. It achieves a gain of 4.6%.

# Soft Logic Clusters	# Multipliers (18x18)	Demand Ratio	A - Area – no shadow clusters (10e5 u^2)	B- Area – shadow clusters (10e5 u^2)	Area-Efficiency Ratio (A/B)
1849	0	0	252	226	**1.115**
1638	0	0	223	201	**1.111**
1420	0	0	193	174	**1.109**
1042	0	0	142	128	**1.106**
1904	0	0	258	234	**1.101**
1924	89	1:21.6	262	257	**1.019**
1523	141	1:10.8	268	273	**0.980**
1309	121	1:10.8	229	234	**0.980**
1528	284	1:5.4	292	297	**0.980**
1502	349	1:4.3	315	321	**0.980**
		Average: 1:15			Average: **1.046**

Acknowledgments

This work was done under the supervision of Jonathan Rose. This work was funded by an NSERC scholarship and NSERC grant.

For further information

Please contact *jamieson@eecg.toronto.edu*