1 Introduction

1.1 Introduction to Field-Programmable Gate Arrays

Field-Programmable Gate Arrays (FPGAs) are a revolutionary new type of user-programmable integrated circuits that provide fast, inexpensive access to customized VLSI. An FPGA consists of an array of logic cells that can be interconnected via programmable routing switches, where the routing structures are sufficiently general to allow the configuration of multiple levels of the FPGA's logic cells. FPGAs represent a combination of the features of Mask Programmable Gate Arrays (MPGAs) and Programmable Logic Devices (PLDs). From MPGAs, FPGAs have adopted a two-dimensional array of logic cells, and from PLDs the user-programmability. The research reported in this thesis is focused on FPGA routing algorithms and routing architectures.

Following their introduction in 1985, by the Xilinx Company [Cart86], FPGAs have evolved considerably as various new devices have been developed [ElGa88] [ElGa89] [Wong89] [Ahre90] [AMD90] [Gupt90] [Hsie88] [Hsie90] [Kawa90] [Marr89] [Ples89] [Plus90]. FPGAs have quickly gained widespread use, which can be attributed to the reduced manufacturing time and relatively low costs of these large-capacity user-programmable devices. As an implementation medium for customized VLSI circuits, FPGAs offer unique advantages over the alternative technologies (MPGAs, standard cells, and full custom design):

(1) FPGAs provide a reduction in the cost of manufacturing a customized VLSI circuit from tens of thousands of dollars to about one hundred dollars.

(2) FPGAs reduce the manufacturing time from months to minutes.

These advantages, which are attributable to the user-programmability of FPGAs, provide a faster time-to-market and less pressure on designers, because multiple design iterations can be done quickly and inexpensively. However, user-programmability also has drawbacks: the logic density and speed performance of FPGAs is considerably lower than those of the alternatives. While developments over the last few years have shown significant improvements in FPGAs, much research is still needed before the best FPGA designs are discovered.

1.2 Thesis Motivation

Circuits are implemented in an FPGA by interconnecting its logic cells through the user-programmable routing switches. Two distinct purposes are served by the *routing switches*: to connect the logic cells to the routing wires, and to connect one routing wire to another. One example of an FPGA routing switch is a CMOS pass-transistor controlled by a static memory bit [Cart86], but there are a number of other implementations that are used in commercial products. Regardless of the implementation, routing switches consume significant chip area and have appreciable resistance and parasitic capacitance. For these reasons, it is desirable to limit the number of routing switches in an FPGA.

All of the routing switches and wires in an FPGA, and their distribution over the surface of the chip, are collectively referred to as the FPGA's *routing architecture*. A measure of the connectivity provided by a routing architecture is its *flexibility*, which is a function of the total number of routing switches and wires. The design space for FPGA routing architectures is enormous. Choosing a good design involves a tradeoff among flexibility, logic density, and speed performance. A high flexibility yields an FPGA that is easily configured, but if the flexibility is too high then area will be wasted by unused switches, leaving less area for the logic blocks and resulting in lower logic density. Moreover, since each routing switch introduces an RC-delay, high flexibility results in reduced speed performance. Low flexibility, on the other hand, allows higher logic density and lower RC-delay, but if the flexibility is too low, then it may not be possible to interconnect the logic cells sufficiently to implement circuits. A good routing architecture is one that achieves a balance between these competing factors.

The primary focus of this thesis is the study of FPGA routing architectures with regard to the flexibility, logic density, and speed performance tradeoff. The goal of the study is to determine the minimum flexibility that is necessary to provide sufficient interconnection capability to satisfy the requirements of real circuits, and yet low enough so that routing switches are not wasted. This research is part of a large project [Brown90] [Brown91] [Fran90] [Fran91] [Rose89] [Rose90a] [Rose90b] [Rose90c] [Rose91] [Sing91] that examines many aspects of the Computer-Aided Design (CAD) and architecture of FPGAs.

1.3 Research Approach

FPGA routing architectures are studied in this thesis using both an experimental and a theoretical approach. For the experimental study, a new type of detailed routing algorithm has been developed that is able to route a wide range of FPGA routing architectures. The experiments consist of varying the routing architecture flexibility and using the router to measure the resulting effects on the routability of circuits. The results of the experiments provide insights into the amount of routing resources that is sufficient to meet the requirements of real circuits and yet low enough so that the resources are not wasted. These issues are also studied using a theoretical approach that represents both a circuit and an FPGA as simple parameters of a stochastic model.

1.4 Dissertation Organization

This dissertation is organized as follows. Chapter 2 provides background information, including a discussion of general approaches to routing problems, the definitions of routing terminology, and a short history of routing algorithms. It also describes representative examples of commercially available FPGAs, including a brief description of the routing architecture contained in each chip.

Chapter 3 presents a new detailed routing algorithm, designed specifically for FPGAs. The algorithm is unique in that it approaches FPGA routing in a general way, and is designed such that it can be used over a wide range of FPGA routing architectures. The algorithm is the main tool that is used to produce the experimental results that are shown in Chapter 4.

FPGA routing architectures are studied in Chapter 4 using an experimental approach. The algorithm in Chapter 3 is used to route a set of circuits in an FPGA based on a model that allows the routing structures in the FPGA to be changed. For each circuit, a range of flexibilities is evaluated by varying the number of routing switches and wires. The experiments measure the effect of the flexibility of the routing architecture on the percentage of interconnections that can be successfully routed for each circuit.

Chapter 5 investigates FPGA routing architectures using a theoretical approach. For this study, both the FPGA and a circuit are represented by simple parameters. A stochastic model is developed to predict the effect of routing the circuit in the FPGA. The model corroborates the experimental results from Chapter 4 and provides the foundation of a theoretical model that can be used in future studies of FPGA routing architectures.

Chapter 6 provides concluding remarks and directions for future research. References are listed at the end.