A Hybrid Complete-Graph Partial-Crossbar Routing Architecture for Multi-FPGA Systems

Mohammed A. S. Khalid and Jonathan Rose

Department of Electrical and Computer Engineering, University of Toronto
10 King's College Road, Toronto, Ontario, CANADA M5S 3G4
{khalid, jayar}@eecg.toronto.edu

Abstract

Multi-FPGA systems (MFSs) are used as custom computing machines, logic emulators and rapid prototyping vehicles. A key aspect of these systems is their programmable routing architecture; the manner in which wires, FPGAs and Field-Programmable Interconnect Devices (FPIDs) are connected. Several routing architectures for MFSs have been proposed [Arno92] [Butt92] [Hauc94] [Apti96] [Vuil96] and previous research has shown that the partial crossbar is one of the best existing architectures [Kim96] [Khal97]. In this paper we propose a new routing architecture, called the **H**ybrid **C**omplete-**G**raph and **P**artial-Crossbar (HCGP) which has superior speed and cost compared to a partial crossbar. The new architecture uses both hard-wired and programmable connections between the FPGAs.

We compare the performance and cost of the HCGP and partial crossbar architectures experimentally, by mapping a set of 15 large benchmark circuits into each architecture. A customized set of partitioning and inter-chip routing tools were developed, with particular attention paid to architecture-appropriate inter-chip routing algorithms. We show that the cost of the partial crossbar (as measured by the number of pins on all FPGAs and FPIDs required to fit a design), is on average 20% more than the new HCGP architecture and as much as 35% more. Furthermore, the critical path delay for designs implemented on the partial crossbar increased, and were on average 9% more than the HCGP architecture and up to 26% more.

Using our experimental approach, we also explore a key architecture parameter associated with the HCGP architecture: the proportion of hard-wired connections versus programmable connections, to determine its best value.

1 Introduction

Field-Programmable Gate Arrays (FPGAs) are widely used for implementing digital circuits because they offer moderately high levels of integration and rapid turnaround time [Brow92]. Multi-FPGA systems (MFSs), which are collections of FPGAs and memory joined by programmable interconnection network as illustrated in Figure 1, are used when the logic capacity of a single FPGA is insufficient, and when a quickly re-programmable system is desired. The typical uses are for logic emulation [Apti96] [Quic96], rapid prototyping [Van92] [Alte94] [Gall94] [Lewi97] and reconfigurable custom computing machines [Arno92] [Cass93] [Dray95] [Vuil96] [Lewi97].

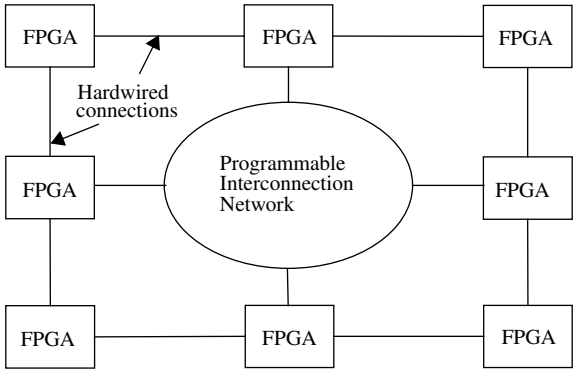

Figure 1 - A Generic Multi-FPGA System

The routing architecture of an MFS is the way in which the FPGAs, the fixed wires, and the programmable interconnect chips are connected. The routing architecture has a strong effect on the speed, cost and routability of the system. Many architectures have been proposed and built [FCCM] [Butt92] [Van92] [Apti96] [Lewi97] and some research work has been done to empirically evaluate and compare different architectures [Kim96] [Khal97].

These studies have shown that the partial crossbar is one of the best existing MFS architectures. In this paper we present HCGP, a routing architecture for MFSs that uses both hardwired and programmable connections to reduce cost and increase speed. We evaluate and compare the HCGP architecture and the partial crossbar architecture using an empirical approach. In particular we compare architectures on the basis of pin cost and speed.

The speed comparisons are based on post inter-chip routing critical path delay of real benchmark circuits, which, to our knowledge, is the first time such detailed timing information has been used in the study of board-level MFS architectures.

Previous work has been done evaluating mesh [Hauc94] and other architectures [Chan93]. Although this work provides some theoretical insight into these architectures, empirical studies that evaluate the implementation of real circuits on different architectures provide a more clear picture of the 'goodness' of each architecture relative to the others [Kim96] [Khal97]. Our own previous research has shown that partial crossbar is vastly superior to the best mesh architecture

[Khal97]. In [Kim96], several MCNC circuits were mapped to seven different architectures, including the partial crossbar architecture. Each circuit was mapped to a fixed size MFS (containing 30 FPGAs). The size of the FPGA was varied depending upon the circuit size. Each architecture was evaluated on the basis of total number of CLBs needed across all circuits (where fewer CLBs used implies better architecture), the type of FPGA chips used (smallest FPGAs implies better architecture), and maximum number of hops needed across all inter-FPGA nets (as a metric for speed). A *hop* is defined as a chip-to-chip connection, i.e. a wire segment that connects two different chips on a board. It was shown that one of the proposed architectures, FPGAs connected together as a tri-partite graph, gave the best results (slightly better than partial crossbar). In this work, relatively few large circuits were used that would have really 'stressed' the architectures, as only three reasonably large circuits (>2000 CLBs) were employed. Also, for the speed estimate only the worst case net delay in terms of the number of hops was considered; which is not as representative of the true delay as post-routing critical path delay.

This paper is organized as follows: In Section 2 we describe the experimental evaluation procedure and the evaluation metrics used, and give details on the suite of large benchmark circuits used in this experimental work. In Section 3 we cover the architectural issues and assumptions that arise when mapping real circuits to the HCGP and partial crossbar architectures. We also briefly describe architecture-specific inter-chip routing algorithms employed. Experimental results and their analysis are presented in Section 4, and we conclude in Section 5.

2 Experimental Overview

To evaluate the two routing architectures considered in this paper, we used the experimental procedure illustrated in Figure 2. Each benchmark circuit was partitioned and routed into each architecture. We also assume that each MFS architecture will be implemented on a single board. Section 2.1 describes the general toolset used in this flow. The cost and delay metrics that we use to evaluate architectures are described in Section 2.2. A description of the 15 benchmark circuits is given in Section 2.3.

2.1 General CAD Flow

As illustrated in Figure 2, we start with a (technology mapped) netlist of 4-LUTs and flip flops of the circuit. The circuit is partitioned into a minimum number of sub-circuits using a multi-way partitioning tool which accepts as constraints the specific FPGA logic capacity and pin count. For all the experiments presented in this paper we used a Xilinx 4013E-1 FPGA, which consists of 1152 4-LUTs, 1152 flip flops, and 192 usable I/O pins [Xili97]. Multi-way partitioning is accomplished using a recursive bi-partitioning procedure. The partitioning tool used is called 'part' and was originally developed for the Transmogrifier-1 rapid prototyping system [Gall94]. It is based on the Fiduccia and Mattheyses partitioning algorithm [Fidu82] with an extension for timing-driven pre-clustering [Shih92]. The output of the partitioning step is a netlist of connections between the FPGAs that contain the circuit.

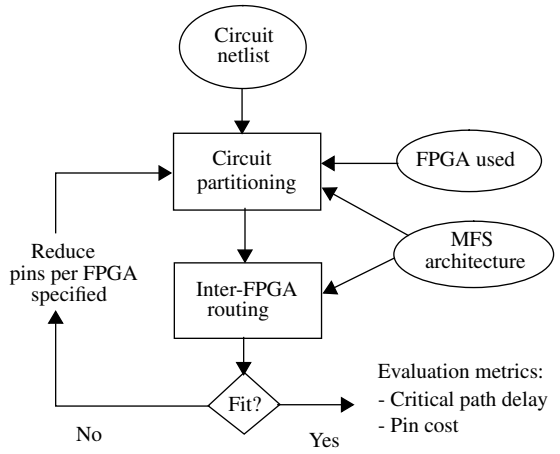

Figure 2 - Experimental Evaluation Procedure for Multi-FPGA Systems

Given the chip-level interconnection netlist, the next step is to route each inter-FPGA net using the most suitable routing path. The routing path chosen should be the shortest path (use the minimum number of hops) and it should cause the least possible congestion for subsequent nets to be routed. Depending on the architecture, the routing resources available in an MFS could be wires that are direct connections between FPGAs, or wires that connect FPGAs and FPIDs.

If the routing attempt fails, the partitioning step is repeated after reducing the number of I/O pins per FPGA specified to the partitioner. This usually increases the number of FPGAs needed, and helps routability by decreasing the pin demand from each FPGA, and providing more "route-through" pins in the FPGAs which facilitate routing.

Note that in an actual MFS, the inter-FPGA routing step is followed by pin assignment, placement and routing within individual FPGAs. We need not perform these tasks because we are only interested in knowing the MFS size needed to fit the circuit. Our previous research has shown that we can afford to assign pins randomly for each FPGA without jeopardizing routability and speed [Khal95]. During recursive bi-partitioning, we restrict the logic utilization of each FPGA to be at most 70% to avoid placement and routability problems within individual FPGAs. Thus we ensure that if inter-FPGA routing attempt succeeds, it is almost guaranteed that the subsequent pin assignment, placement, and routing steps will be successful for each FPGA in the MFS.

We developed a specific router for each of the architectures compared. (We had attempted to create a generic router but found that it had major problems with different aspects of each architecture [Khal98].)

2.2 Evaluation metrics

To compare the two routing architectures we implement benchmark circuits on each and contrast the pin cost and post-routing critical path delay, as described below.

2.2.1 Pin cost

The cost of an MFS is likely a direct function of the number of FPGAs and FPIDs: if the routing architecture is inefficient, it will require more FPGAs and FPIDs to implement the same amount of logic as a more efficient MFS. While it is difficult to calculate the price of specific FPIDs and FPGAs, we assume that the total cost is proportional to the total number of pins on all of these devices. Since the exact number of FPGAs and FPIDs varies for each circuit implementation (in our procedure above, we allow the MFS to grow until routing is successful), we calculate, for each architecture, the total number of pins required to implement each circuit. We refer to this as the *pin cost* metric for the architecture.

2.2.2 Post-Routing Critical Path Delay

The speed of an MFS, for a given circuit, is determined by the critical path delay obtained after a circuit has been placed and routed at the inter-chip level. We call this the *post-routing critical path delay*. We have developed an MFS static timing analysis tool (MTA) for calculating the post routing critical path delay for a given circuit and MFS architecture.

Item	Delay (ns)
Intra-FPGA CLB-to-CLB routing delay	2.5
FPGA input pad delay	1.4
FPGA output pad delay	3.2
CLB delay (without using H-LUT)	1.3
CLB delay (via H-LUT)	2.2
FPID crossing delay (including pad delays)	10
PCB trace delay	3
FPGA Route Through Delay	10

Table 1 - Delays Used in Timing Analyzer Model

The operation and modeling used in the MTA are described briefly as follows: It first calculates the critical path delay of the un-partitioned design using a widely used method called the *block oriented technique* [Joup87]. It then reads the inter-FPGA netlist and the routing path for each inter-FPGA net, as provided by the inter-chip router, and the MFS architecture description. From this information the circuit is annotated with the inter-chip delays, from which the critical path delay can be calculated.

In the delay annotation step, the delay values given in Table 1 (obtained from data sheets [Xili97] and [Icub97] and some design experience) are used.

Note that since we do not perform individual FPGA place and route, we approximate the CLB-to-CLB delay as a constant. The value of 2.5 ns for CLB-to-CLB routing delay is roughly half the delay on a long line for XC4013E-1 FPGA. This is a pessimistic estimate. Although using a single delay value is somewhat inaccurate, it still gives us a good estimate of the post-routing critical path delay of an MFS because it is dominated by off-chip delay values.

2.3 Benchmark Circuits

A total of fifteen large benchmark circuits were used in our experimental work. An extensive effort was expended to collect this suite of large benchmark circuits. The details of each benchmark circuit are shown in Table 2 which provides the circuit name, size (in 4-LUTs, D flip flops, and I/O count), rough description of the functionality, the source of the circuit and the manner in which it was synthesized. Four circuits were obtained from MCNC [Yang91], two from FPGA synthesis benchmarks [Prep96], and the remaining nine were developed at the University of Toronto (UofT). The circuits from MCNC were available in the XNF [Xili97] gate-level netlist format required by our front end tools. All the circuits from [Prep96] and UofT were originally available as VHDL or Verilog HDL models and were synthesized into XNF netlists using Exemplar [Exem94] and Synopsys Behavioral Compiler [Knap96] and/or Design Compiler [Syno97] synthesis tools. We show these details of the benchmark circuits because we feel that the MCNC circuits that have been used so far in MFS architecture studies are insufficient in terms of size and variety to 'stress' different architectures and the mapping tools used. Specifically, we found that they are easier to partition and map compared to the other real circuits that we use in this work.

3 Routing Architecture Description and Algorithms

In this Section we describe the partial crossbar and the HCGP architectures. For each architecture, we briefly describe an architecture-specific inter-chip router.

3.1 Architectural Description and Routing for the Partial Crossbar

The partial crossbar architecture [Butt92] [Varg93] is used in logic emulators produced by Quickturn Design Systems [Quic96]. An example partial crossbar using four FPGAs and three FPIDs is shown in Figure 3. The pins in each FPGA are

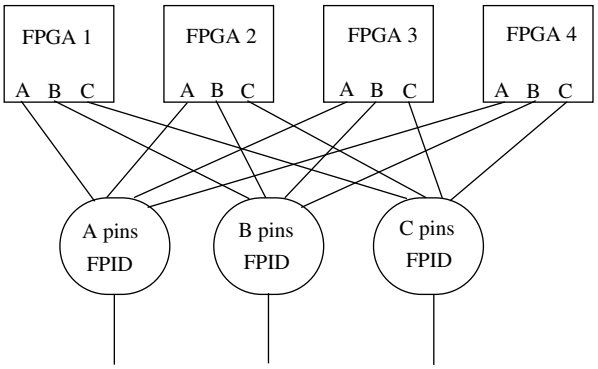

Figure 3 - The Partial Crossbar Architecture

Circuit	Size	Function	Source, Synthesis tool used (if applicable)
s35932	4374 LUTs, 1728 FFs, 357 I/Os	Sequential circuit	MCNC
s38417	6097 LUTs, 1463 FFs, 134 I/Os	Sequential circuit	MCNC
s38584	4396 LUTs 1451 FFs, 292 I/Os	Sequential circuit	MCNC
mips64	2900 LUTs 440 FFs, 260 I/Os	Scaled down version of MIPS R4000	[Prep96], Verilog model synthesized using Exemplar
spla	3423 LUTs 0 FFs, 62 I/Os	Combinational Circuit	MCNC
cspla	2039 LUTs 0 FFs, 62 I/Os	Clone of spla	UofT, Generated using GEN[Hutt96]
mac64	2560 LUTs 64 FFs, 133 I/Os	64-bit multiply-accumulate ckt.	UofT, Verilog model synthesized using Synopsys
sort8	1540 LUTs 200 FFs, 20 I/Os	8-bit HW sort engine	UofT, Verilog model synthesized using Synopsys
fir16	5366 LUTs 1040 FFs, 60 I/Os	16-bit, 8-stage FIR filter	UofT, Verilog model synthesized using Synopsys
gra	2494 LUTs 1156 FFs, 144 I/Os	Graphics acceleration circuit	UofT, circuit generated using tmcc[Gall95]
fpsdes	3484 LUTs 1008 FFs, 69 I/Os	Fastest pseudo DES circuit	UofT, Verilog model synthesized using Synopsys
spsdes	2452 LUTs 982 FFs, 69 I/Os	Smallest pseudo DES circuit	UofT, Verilog model synthesized using Synopsys
ochip64	3617 LUTs 5810 FFs, 84 I/Os	Output chip for ATM switching chip set	UofT, VHDL model synthesized using Exemplar
ralu32	2553 LUTs 584 FFs, 98 I/Os	32-bit register file, ALU, and control logic	[Prep96], VHDL model synthesized using Synopsys
iir16	3149 LUTs 522 FFs, 52 I/Os	16-bit IIR filter	UofT, VHDL model synthesized using Synopsys

Table 2 - Benchmark Circuits

divided into N subsets, where N is the number of FPIDs in the architecture. All the pins belonging to the same subset in different FPGAs are connected to a single FPID. Note that any circuit I/Os will have to go through FPIDs to reach FPGA pins. Thus, a certain number of pins per FPID are reserved for circuit I/Os.

The number of pins per subset (P_t) is a key architectural parameter that determines the number of FPIDs needed and the pin count of each FPID. The extremes of the partial crossbar architecture can be illustrated by considering an architecture with four FPGAs (assuming 192 usable I/O pins per FPGA). A P_t value of 192 will require a single 768-pin FPID that acts as a full crossbar. A P_t value of 1 will require 192 4-pin FPIDs. Both of these cases are impractical.

A good value of P_t should require low cost, low pin count FPIDs. For the above example, a P_t value of 12 will require 16 48-pin FPIDs. Taking into account the extra FPID pins required for circuit I/Os we will need to use 64 or 96-pin FPIDs, which are commercially available [ICub97]. When choosing a value of P_t, we must ensure that the number of usable I/Os per FPGA is evenly divisible by P_t or at least the remainder should be a very small number so that we can use such pins for routing high fanout inter-FPGA nets. Our previous research [Khal97] has shown that, for real circuits, the routability and speed of the partial crossbar is not affected by the value of P_t used. This result is contingent upon using an

intelligent inter-chip router that understands the architecture and routes each inter-FPGA net using only two hops to minimize the routing delay.

3.1.1 Routing Algorithm for the Partial Crossbar

For any MFS architecture in general and for the partial crossbar in particular, it is important to use a routing algorithm that exploits architecture-specific features in order to obtain good results.

We have developed a routing tool, PCROUTE, for the partial crossbar architecture that gives excellent results for all the circuits. Irrespective of the value of P_t, it achieves 100% routing completion and produces two-hop routing for all the nets in almost all circuits. For only two circuits, for the specific case of P_t= 4, it produced multi-hop routing paths for a negligible number of nets (1 out of 991 nets for the first circuit and 3 out of 645 nets for the second). In practical terms, this means it gives almost optimal results for all of our benchmark circuits.

The PCROUTE algorithm works as follows: for each net (irrespective of fanout), it evaluates paths through all available FPIDs. It uses a suitable cost function to choose an FPID that will guarantee balanced usage of FPIDs and will preserve the most options for two-hop routing of subsequent nets to be routed. We show in [Khal98] that PCROUTE is equivalent in quality to other partial crossbar routers that have been proposed so far [Kim96] [Mak97a] [Lin97]. PCROUTE is better than [Mak97b] in terms of both speed and routability because that algorithm splits each multi-terminal into a set of two-terminal nets and routes them independently, leading to multiple hops and even possible routing failures.

3.2 Architectural Description and Routing for HCGP

The new HCGP architecture is shown in Figure 4 for four FPGAs and three FPIDs. The I/O pins in each FPGA are divided into two groups: hardwired connections and programmable connections. The pins in the first group connect to other FPGAs and the pins in the second group connect to FPIDs. The FPGAs are directly connected to each other using a complete graph topology, i.e. each FPGA is connected to every other FPGA. The connections between FPGAs are evenly distributed, i.e. the number of wires between every pair of FPGAs is the same. The FPGAs and FPIDs are connected in exactly the same manner as in a partial crossbar. As in the partial crossbar, circuit I/Os have to go through FPIDs to reach FPGA pins. Again, a certain number of pins per FPID are reserved for circuit I/Os.

The direct connections between FPGAs can be exploited to obtain reduced cost and better speed. For example, consider a net that connects FPGA 1 to FPGA 3 in Figure 4. If there were no direct connections as in the partial crossbar, we would have used an FPID to connect the two FPGAs. This will cost extra delay and two extra FPID pins. A natural question to ask is: why not dispense with FPIDs and just use FPGAs connected as a completely connected graph as investigated in [Kim96]? The answer is that routing multi-terminal nets in an FPGA-only architecture is expensive in terms of routability because there may not be enough extra FPGA pins for routing multi-terminal nets, as illustrated in Figure 5. In Figure 5(a) two extra FPGA pins are used for routing a fanout 3 multi-terminal net. If we use too many FPGA pins for routing, not enough pins remain for accessing the logic in each FPGA. If we use an FPID for routing the same multi-terminal net, we do not need even a single extra FPGA pin, other than the FPGA pins needed to access the source and sinks of the net as shown in Figure 5(b).

A key architectural parameter in the HCGP architecture is the percentage of programmable connections, P_p. It is defined as the percentage of each FPGA's pins that are connected to FPIDs (the remainder are connected to other FPGAs). If P_p is too high it will lead to increased pin cost, if it is too low it will adversely affect routability. If P_p is 0% the HCGP architecture degrades to a completely connected graph of FPGAs with no FPIDs used. If P_p is 100% the HCGP architecture degrades to a standard partial crossbar. A key issue we address later is the best value of P_p for obtaining minimum cost and good routability.

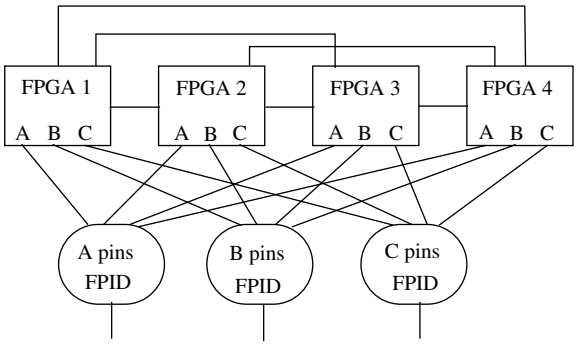

Figure 4 - The HCGP Architecture

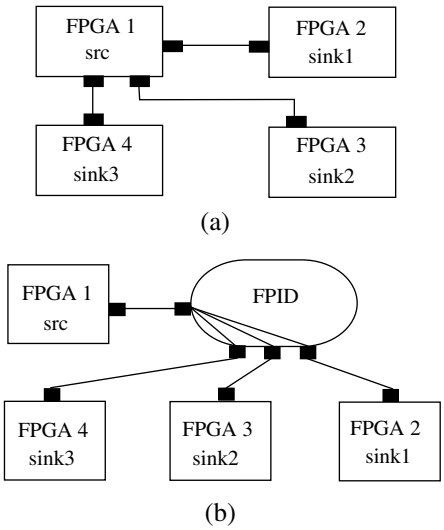

Figure 5 - Multi-terminal routing (a) without an FPID (b) with and FPID

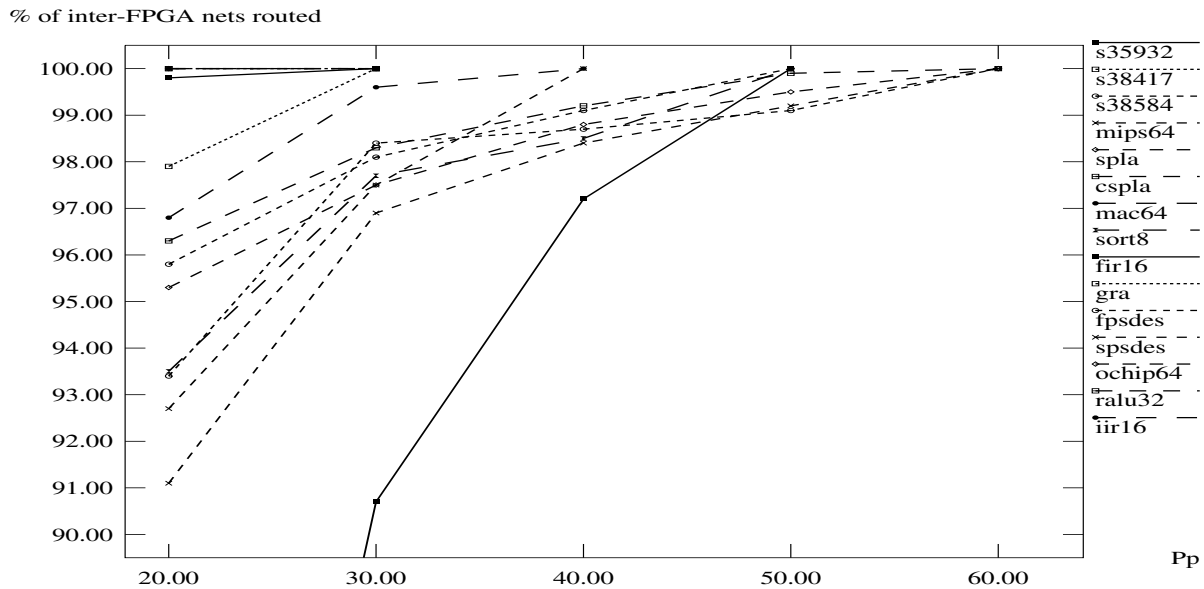

Figure 6 - The Effect of P_p on Routability of the HCGP Architecture

3.2.1 Routing Algorithm for HCGP

The inter-chip routing algorithm for HCGP is similar to the partial crossbar routing algorithm in the sense that the same algorithm is used when routing nets through FPIDs. However, the difference here is that the router should also exploit the direct connections between FPGAs to minimize the number of FPGA and FPID pins used for routing and to minimize the number of hops for routing each inter-FPGA net.

We have developed a routing tool, called HROUTE, that understands the HCGP architecture and gives excellent routability and speed results for all the benchmark circuits.

The main objective of HROUTE is to route all nets using no more than two hops for each source-sink path. Wherever possible, we try to use direct connections to minimize source-sink net delay when routing both two-terminal and multi-terminal nets. We first try to route all possible two-terminal nets using the direct connections between FPGAs to minimize usage of pins and net delay. Next, we route all multi-terminal nets through FPIDs using a routing algorithm similar to that used in PCROUTE, described above. Finally, the remaining two terminal nets are routed using FPGAs or FPIDs. Any nets that remain unrouted are processed by a maze router. Our experience has shown that net ordering is crucial for obtaining good routability and speed results in HROUTE. A detailed description of HROUTE is given in [Khal98].

4 Experimental Results

In this section we determine the effect of varying the value of P_p on the routability of the HCGP architecture and compare the partial crossbar and HCGP architectures.

4.1 HCGP Architecture: Analysis of P_p

Recall the definition of P_p, the percentage of pins used for programmable connections, given in Section 3.2. P_p is important because it affects the cost, routability and speed of the HCGP architecture. Here we explore the effect of P_p on the routability of the HCGP architecture. We mapped the fifteen benchmark circuits to the HCGP architecture using five different values of P_p (20, 30, 40, 50 and 60). The results are shown in Figure 6. The Y-axis represents the percentage of inter-FPGA nets routed and the X-axis represents the P_p values. The first clear conclusion is P_p = 60% gives 100% routability for all the benchmark circuits. Notice that about half of the circuits routed at P_p <= 40%, and for the remaining half, more than 97% of the nets routed. This implies that there is a potential for obtaining 100% routabilty for all circuits at P_p = 40% if we use a routability driven partitioner like the one used in [Kim96]. This should lead to further reduced pin cost for HCGP compared to the partial crossbar.

We conjecture that the P_p value required for routing completion of a given circuit on HCGP depends upon how well the circuit structure ‘matches’ the topology of the architecture.

Circuit	Number of FPGAs		Normalized pin cost		Normalized post-routing critical path delay	
	Partial crossbar	HCGP	Partial crossbar	HCGP	Partial crossbar	HCGP
s35932	8	8	1.30	1.0	1.0	1.0
s38417	9	9	1.34	1.0	1.0	1.0
s38584	9	9	1.34	1.0	1.23	1.0
mips64	14	16	1.16	1.0	0.99	1.0
spla	18	25	0.91	1.0	0.96	1.0
cspla	18	21	1.13	1.0	1.01	1.0
mac64	6	8	0.98	1.0	1.11	1.0
sort8	12	14	1.11	1.0	0.99	1.0
fir16	10	10	1.30	1.0	1.24	1.0
gra	4	4	1.35	1.0	1.20	1.0
fpsdes	9	9	1.34	1.0	1.16	1.0
spsdes	8	8	1.30	1.0	1.15	1.0
ochip64	8	8	1.30	1.0	1.26	1.0
ralu32	9	15	0.76	1.0	1.06	1.0
iir16	6	6	1.32	1.0	1.05	1.0
Average			**1.20**	**1.0**	**1.09**	**1.0**

Table 3 - Comparison of the Partial Crossbar and HCGP Architectures

4.2 Comparison of HCGP and Partial Crossbar

The 15 benchmark circuits described in Table 2 were mapped to the partial crossbar and HCGP architectures using the experimental procedure described in Section 2. The results obtained are given in Table 3 and Table 4. In Table 3, the first column shows the circuit name. The second column shows the number of FPGAs needed to implement the circuit on each architecture (recall that we increase the MFS size until routing is successful). The third column shows the normalized pin cost (where the number of pins used by the HCGP architecture is set as 1) and the fourth column shows the normalized critical path delay obtained for each architecture. Table 4 is similar to Table 3 except that it shows actual (un-normalized) pin cost and delay values.

The number of FPIDs used is not shown because it is constant for each architecture. All the results for partial crossbar use $P_t = 17$. The parameter P_t determines the number of FPIDs required and the number of FPGAs in the architecture determine the pin count of each FPID. We have shown that the value of P_t used has no effect on the routability and speed of the partial crossbar [Khal97]. Therefore any arbitrary value of P_t can be used. However, for practical reasons, the value chosen should require FPIDs that have reasonable pin counts (about 400 pins or less, which are commercially available) for the largest partial crossbar required in our experiments. A reasonable choice in this respect is $P_t = 17$.

The value of P_p for the HCGP architecture was set to 60% to obtain good routability across all circuits, as discussed in Section 4.1. Notice that the parameter P_t also applies to the programmable connections in the HCGP. For the same reasons as in the partial crossbar (given in the previous paragraph), we chose $P_t = 14$ for the HCGP architecture.

In reviewing Table 3, consider the circuit mips64. The first partitioning attempt resulted in 14 FPGAs required to implement the circuit on partial crossbar. However, the circuit was not routable on HCGP and the partitioning was repeated after reducing the number of pins per FPGA specified to the partitioner by 10%. This resulted in 16 FPGAs required to implement the circuit. The second partitioning attempt was routable on the HCGP architecture because more ‘free pins’ were available in each FPGA for routing purposes. The pin

Circuit	Number of FPGAs		Pin cost		Post-routing critical path delay (in ns)	
	Partial crossbar	HCGP	Partial crossbar	HCGP	Partial crossbar	HCGP
s35932	8	8	2992	2296	57.4	57.4
s38417	9	9	3366	2520	94.6	94.6
s38584	9	9	3366	2520	139.4	113.4
mips64	14	16	5236	4528	461.9	467.5
spla	18	25	6732	7400	196.3	203.9
cspla	18	21	6732	5964	192.5	191.2
mac64	6	8	2244	2296	622.9	563
sort8	12	14	4488	4046	532.8	538.3
fir16	10	10	3740	2870	238	192.7
gra	4	4	1496	1112	70	58.5
fpsdes	9	9	3366	2520	226.5	195.4
spsdes	8	8	2992	2296	248.8	216.2
ochip64	8	8	2992	2296	63.2	50.1
ralu32	9	15	3366	4410	316.8	298
iir16	6	6	2244	1704	160.2	152.8
			Total: 55352	**Total: 48778**	**Avg.: 241.42**	**Avg.: 226.2**

Table 4 - Actual Pin Cost and Delay Values for the Partial Crossbar and HCGP Architectures

cost for the partial crossbar was still more than that for HCGP because it uses many more programmable connections, and hence more FPID pins. A partial crossbar always requires one FPID pin for every FPGA pin; the HCGP architecture requires a lower ratio.

Inspecting Table 3, we can make several observations. First, the partial crossbar needs 20% more pins on average, and as much as 35% more pins compared to the HCGP architecture. Clearly, the HCGP architecture is superior to the partial crossbar architecture in terms of the pin cost metric. This is because the HCGP exploits direct connections between FPGAs to save FPID pins that would have been needed to route certain nets in partial crossbar. However, for routability purposes, the HCGP needs some free pins in each FPGA and may require repeated partitioning attempts for some circuits.

Table 3 also shows that the typical circuit delay is lower with the HCGP architecture: the HCGP gives significantly less delay for six circuits compared to the partial crossbar and about the same delay for the rest of the circuits. The reason is that the HCGP utilizes fast and direct connections between FPGAs, whenever possible. From the delay values in Table 1, we can show that the interconnection delay is much smaller (12.6 ns) if we use direct connections between FPGAs compared to the delay value (25.6 ns) when connecting two FPGAs through an FPID. Another interesting observation is that even for the circuits where the HCGP needs more FPGAs compared to the partial crossbar, it still gives comparable delay value. This clearly demonstrates that the HCGP architecture is inherently faster due to the nature of its topology. It gives significant speed up even though we have not yet employed timing driven inter-FPGA routing.

Table 4 shows the actual pin cost and delay values obtained for the partial crossbar and HCGP architectures. It is interesting that the estimated clock speeds for the partial crossbar architecture range from 17.4 MHz for the *s35932* circuit to 1.61 MHz the *mac64* circuit. This range is representative of the clock rates expected in MFSs [Quic96].

5 Conclusions and Future Work

In this paper we presented the Hybrid Complete-Graph and Partial-Crossbar (HCGP), a new routing architecture for

multi-FPGA systems. Using an experimental approach, we evaluated and compared this architecture to the partial crossbar architecture and showed that it is superior in terms of pin cost and speed. To our knowledge, this is the first architectural study of board-level MFSs that considers post-routing critical path delay when evaluating the speed performance of different architectures.

We explored a key parameter, P_p, associated with the HCGP architecture and determined its best value (60%) for obtaining good routability for a variety of circuits.

We believe that the HCGP architecture would give even better results if we use better mapping (CAD) tools for partitioning and inter-FPGA routing. First, a timing driven router that routed all or most of the nets on the critical paths using fast and direct connections would lead to larger reductions in the critical path delay. Second, a routabilty driven partitioner, similar to the one used in [Kim96], would result in further reduced pin cost by making circuits routable for even lower values of P_p (say 40%).

The HCGP architecture is suitable for single board MFSs using a maximum of about 20 FPGAs. As FPGA logic and pin capacities continue to rise, it makes sense to use single board systems using a few high capacity FPGAs to avoid the problems associated with using high pin count connectors for multi-board systems [Lewi97]. For applications where hundreds of FPGAs are needed, like logic emulation, we could use 'clusters' of HCGPs interconnected using a hierarchical partial crossbar scheme [Butt92]. The hardwired connections, within each cluster and between different clusters, would still help in reducing the overall pin cost. Determining the P_p value suitable for such hierarchical architectures is an open research problem. We will need extremely large benchmark circuits and appropriate CAD tools to explore hierarchical architectures.

Acknowledgments

The authors would like to thank Dave Galloway for his help with the partitioning tool and Jason Anderson for his help in synthesizing the benchmark circuits. This research was supported by the Information Technology Research Center (ITRC) of Ontario and MICRONET.

References

[Alte94] Altera Corporation, Reconfigurable Interconnect Peripheral Processor (RIPP10) Users Manual, Version 1.0, 1994.

[Apti96] Aptix Corporation, Product brief: The System Explorer MP4, 1996. Available on Aptix Web site: http://www.aptix.com.

[Arno92] J. M. Arnold, D. A. Buell, and E. G. Davis, "Splash 2," Proceedings of 4th Annual ACM Symposium on Parallel Algorithms and Architectures, pp. 316-322, 1992.

[Brow92] S. Brown, R. Francis, J. Rose, and Z. Vranesic, Field Programmable Gate Arrays, Kluwer Academic Publishers, 1992.

[Butt92] M. Butts, J. Batcheller, and J. Varghese, "An Efficient Logic Emulation System," Proceedings of IEEE International Conference on Computer Design, pp. 138-141, 1992.

[Cass93] S. Casselman, "Virtual Computing and The Virtual Computer," Proceedings of IEEE Workshop on FPGAs for Custom Computing Machines, pp. 43-48, 1993.

[Chan93] P. K. Chan, M. D. F. Schlag, "Architectural Tradeoffs in Field-Programmable-Device-Based Computing Systems," Proceedings of IEEE Workshop on FPGAs for Custom Computing Machines, pp. 152-161, 1993.

[Dray95] T. H. Drayer, W. E. King, J. G. Tront, and R. W. Conners, "MORRPH: A Modular and Reprogrammable Real-time Processing Hardware," Proceedings of IEEE Symposium on FPGAs for Custom Computing Machines, pp. 11-19, 1995.

[Exem94] Exemplar Logic, VHDL Synthesis Reference Manual, 1994.

[Fidu82] C. M. Fiduccia, and R. M. Mattheyses, "A Linear-Time Heuristic for Improved Network Partitions", Proc. of 19th ACM/IEEE Design Automation Conference, pp. 241-247, 1982.

[FCCM] Proceedings of IEEE Workshops/Symposia on FPGAs for Custom Computing Machines, 1992 to 1996.

[Gall94] D. Galloway, D. Karchmer, P. Chow, D. Lewis, and J. Rose, "The Transmogrifier: The University of Toronto Field-Programmable System", CSRI Technical Report (CSRI-306), CSRI, University of Toronto, 1994.

[Gall95] D. Galloway, "The Transmogrifier C Hardware Description Language and Compiler for FPGAs," Proceedings of IEEE Symposium on FPGAs for Custom Computing Machines, pp. 136-144, 1995.

[Hauc94] S. Hauck, G. Boriello, C. Ebeling, "Mesh Routing Topologies for FPGA Arrays", Proceedings of FPGA'94, 1994.

[Hutt96] M. Hutton, J.P. Grossman, J. Rose and D. Corneil, "Characterization and Parameterized Random Generation of Digital Circuits," Proc.of the Design Automation Conference, pp. 94-99, 1996.

[Icub97] I-Cube, Inc., The IQX Family Data Sheet, May 1997. Available at: www.icube.com.

[Icub94] I-Cube, Inc., "Using FPID Devices in FPGA-based Prototyping," Application note, Part number:D-22-002, February 1994.

[Joup87] Norman P. Jouppi, "Timing Analysis and Performance Improvement of MOS VLSI Designs," IEEE Trans. on CAD, vol. CAD-6, no. 4, pp. 650-665, July 1987.

[Khal95] M. A. S. Khalid and J. Rose, "The Effect of Fixed I/O Pin Positioning on The Routability and Speed of FPGAs," Proceedings of The Third Canadian Workshop on Field-Programmable Devices (FPD'95), pp. 92-104, 1995.

[Khal97] M. A. S. Khalid and J. Rose, "Experimental Evaluation of Mesh and Partial Crossbar Routing Architectures for Multi-FPGA Systems," Proceedings of the Sixth IFIP International Workshop on Logic and Architecture Synthesis (IWLAS'97), 1997.

[Khal98] M. A. S. Khalid, Routing Architecture and Layout Synthesis for Multi-FPGA Systems, Ph.D. Thesis, University of Toronto, in progress.

[Kim96] C. Kim, H. Shin, "A Performance-Driven Logic Emulation System: FPGA Network Design and Performance-Driven Partitioning," IEEE Trans. on CAD, vol. 15, no. 5, pp. 560-568, May 1996.

[Knap96] D. W. Knapp, Behavioral Synthesis: Digital System Design Using the Synopsys Behavioral Compiler, Prentice Hall PTR, 1996.

[Lewi97] D. M. Lewis, D. R. Galloway, M. Van Ierssel, J. Rose, and P. Chow, "The Transmogrifier-2: A 1 Million Gate Rapid Prototyping System," Proceedings of FPGA'97, pp. 53-61, 1997.

[Lin97] S. Lin, Y. Lin, and T. Hwang, "Net Assignment for the FPGA-Based Logic Emulation System in the Folded-Clos Network Structure," IEEE Trans. on CAD, vol. 16, no. 3, pp. 316-320, March 1997.

[Mak97a] Wai-Kei Mak, D. F. Wong, "On Optimal Board-Level Routing for FPGA-based Logic Emulation," IEEE Trans. on CAD, vol. 16, no. 3, pp. 282-289, March 1997.

[Mak97b] Wai-Kei Mak, D. F. Wong, "Board-Level Multi-Terminal Net Routing for FPGA-based Logic Emulation," ACM Trans. on Design Automation of Electronic Systems, vol. 2, no. 2, pp. 151-167, April 1997.

[Prep96] Programmable Electronics Performance Corporation, HDL models for different circuits (synthesis benchmarks) are available on their Web site: http// www.prep.org.

[Quic96] Quickturn Design Systems, Inc., System Realizer Data Sheet, 1996. Available on Quickturn Web site:http://www.quickturn.com.

[Shih92] M. Shih, E. S. Kuh, "Performance-Driven System Partitioning on Multi-Chip Modules," Proc. of the Design Automation Conference, pp. 53-56, 1992.

[Syno97] Synopsys, Inc., Design Compiler(Version 3.4a), Behavioral Compiler (Version 3.4a), and Library Compiler (Version 3.4a), Reference Manuals. Documents available on-line.

[Van92] D. E. Van Den Bout, et al, "Anyboard: An FPGA-Based Reconfigurable System," IEEE Design and Test of Computers, pp. 21-30, June 1992.

[Varg93] J. Vargese, M. Butts, and Jon Batcheller, "An Efficient Logic Emulation System", IEEE Trans. on VLSI Systems, vol. 1, no. 2, pp. 171-174, June 1993.

[Vuil96] J. E. Vuillemin, P. Bertin, D. Roncin, M. Shand, H. Touati, and P. Boucard, "Programmable Active Memories: Reconfigurable Systems Come of Age," IEEE Transactions on VLSI, Vol 4, No. 1, pp. 56-69, March 1996.

[Xili97] Xilinx, Inc., Product Specification: XC4000E and XC4000X Series FPGAs, Version 1.2, June 16, 1997. Available on Xilinx Web site: www.xilinx.com.

[Yang91] S. Yang, Logic Synthesis and Optimization Benchmarks User Guide, Version 3.0, Microelectronics Center of North Carolina, January 1991.