4 The Flexibility of Field-Programmable Gate Array Routing Architectures

4.1 Introduction

This chapter uses an experimental approach to investigate the effect of the flexibility of an FPGA's routing architecture on its routability. *Flexibility* is a measure of the connectivity provided by a routing architecture and is a function of the total number of routing switches and wires. A loose definition of *routability* is the percentage of the connections in a given circuit that can be successfully routed. The experiments consist of routing a set of circuits in an FPGA that is modelled in such a way that the number of routing switches and wires can be changed. Routability is measured over a range of flexibility, using the routing algorithm described in Chapter 3.

Recall that a discussion in Chapter 1 showed that the choice of a good routing architecture involves a tradeoff among flexibility, logic density, and speed performance. High flexibility results in an FPGA that has good routability, but routing switches will be wasted if the flexibility is higher than is necessary for 100 percent routability. Also, high flexibility results in lower logic density and speed performance because each routing switch consumes significant area and has appreciable resistance and capacitance. Low flexibility, on the other hand, permits higher logic density and speed performance, but if flexibility is too low then 100 percent routability of circuits may not be possible.

The routing architectures that will studied in this chapter fit the FPGA model that was introduced in Chapter 3. Figure 4.1 reproduces the model, for ease of reference. The FPGA shown in the figure has three tracks per routing channel and two pins on each side of a logic cell (L). Recall that the connection (C) blocks contain the routing switches that are used to connect the pins of the logic cells to the routing channels, and the switch (S) blocks at the intersections of horizontal and vertical routing channels provide the routing switches that can connect wiring segments in one channel segment to those in another. The flexibility of the routing architecture can be altered by changing the connectivity in the C blocks, the S blocks or the number of tracks in each routing channel [Rose90b] [Rose91].

The specific questions concerning FPGA routing architectures that are answered in this chapter include:

> *What is the effect of the flexibility of the C blocks on routability?* Section 4.5.1 shows that a high flexibility in the C blocks is desirable for good routability.

> *What is the effect of the flexibility of the S blocks on routability?* Section 4.5.2 shows that good routability can be achieved with a relatively low flexibility in the S blocks.

> *How do the S block and C block flexibilities interact?* Section 4.5.3 shows that a high flexibility in either of the S or C blocks can compensate to some extent for a relatively low flexibility in the other.

> *What is the effect of the flexibilities of the C and S blocks on the number of tracks per routing channel required to achieve 100 percent routability?* Section 4.5.4 shows that very close to the theoretical minimum number of tracks per channel (this is defined by the maximum channel density over all channels in the FPGA) can be achieved for surprisingly low flexibilities.

> *What is the effect of the flexibilities of the C and S blocks on the total number of routing switches required in an FPGA to achieve 100 percent routability?* Section 4.5.5 shows that there is a range of C and S block flexibilities that results in a minimum number of switches.

This chapter is organized as follows: Section 4.2 presents assumptions that are made about the architecture of the FPGA, Section 4.3 describes the experimental procedure, the limitations of this work are discussed in Section 4.4, Section 4.5 presents the experimental results and explanations, and concluding remarks appear in Section 4.6.

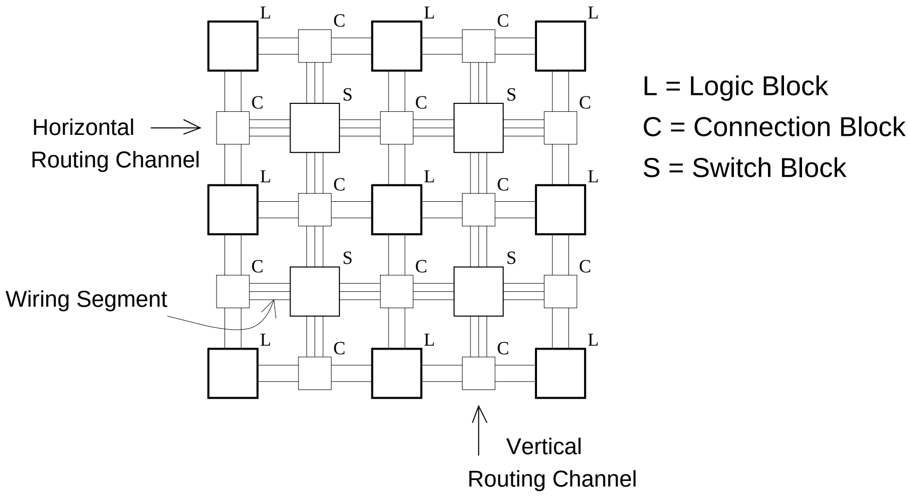

Figure 4.1 - *The Model of the FPGA Routing Architecture*

4.2 FPGA Architectural Assumptions

The architecture of an FPGA consists of its I/O cells, logic cells, and routing architecture. The I/O cells are assumed, as in Chapter 3, to be logic cells on the periphery of the chip. The other assumptions that are made in this chapter concerning the logic cells and the routing architecture are given in the following subsections.

4.2.1 The Logic Cell

The logic cell that is used here is the result of a previous study [Rose89][Rose90c] and is illustrated in Figure 4.2. It has a four input look-up table, a D flip-flop, and a tri-state output. In the previous study, this cell achieved the minimum total area over a set of circuits when compared to other cells that had differing numbers of inputs, including and excluding a D flip-flop. The logic cell has a total of 7 logical pins (4 inputs, 1 output, 1 clock, 1 tri-state) but they may not physically appear on all sides of the cell. The number of physical occurrences of each pin is an important architectural parameter, as explained below.

The number of logic cell sides on which each logical pin physically appears is called T. To illustrate this concept, the cases $T = 4$ and $T = 1$ are shown in Figure 4.3, where the logic cell pins are numbered from 0 to 6. As an example, the figure also shows the connection of pin 0 on one logic cell to pin 6 on another. The particular choice of T affects the routing problem in a number of ways. Selecting a low value of T implies that there will be fewer routing switches, which means the switches will use less area and add less capacitance to the tracks, but as shown in Figure 4.3, connections may be longer since it may be necessary to route to a certain side. This increases the channel densities and causes the connections to pass through more routing switches. Conversely, choosing a higher value of T allows shorter connections and minimizes the channel densities, but if T is higher than necessary, switches will be wasted.

For the experiments shown here, the value used for T is 2. This was chosen for area considerations only and was determined by performing the global routing of several circuits for each value of T and measuring the number of tracks per channel (maximum channel density) required in each case. The results are shown in Table 4.1, which gives the average maximum channel density of the circuits, for each value of T. The table shows a significant decrease in track count from the $T = 1$ to $T = 2$ case but diminishing

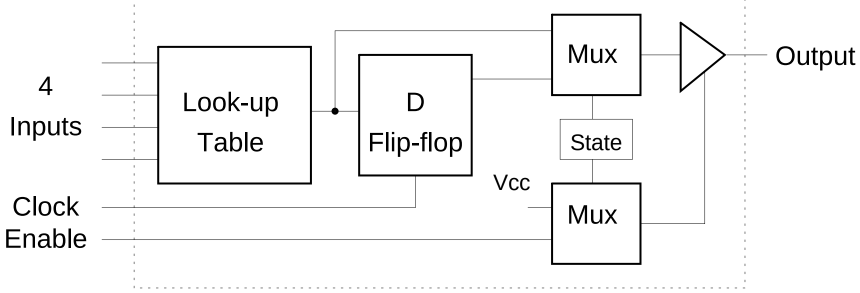

Figure 4.2 - *The Logic Cell*

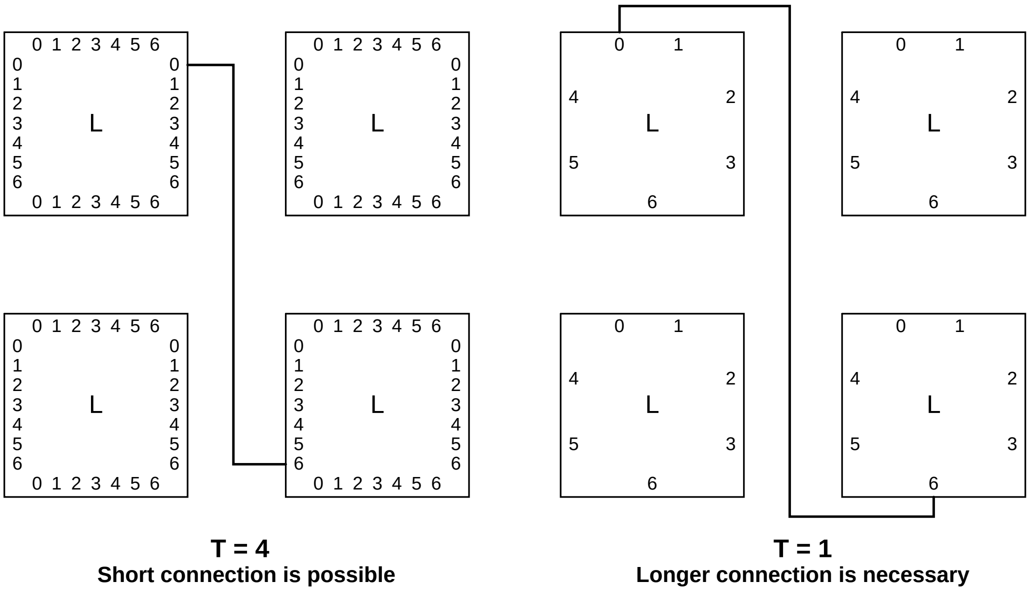

Figure 4.3 - *Example, and Effect on Connection Length, of T*

returns for higher values of *T*. Note that the routing tools used for these experiments did not make use of the functional equivalence of the logic cell inputs (the inputs to a look-up table are functionally equivalent), and if they had it may have been possible to choose a value of $T = 1$ without an increase in the number of tracks.

T	**Avg. Maximum Channel Density**
1	15
2	12
3	11
4	11

Table 4.1 - *The Effect of T on Channel Density*

The following two sections provide a detailed discussion of the C and S blocks used in these experiments. Some of this information also appeared in Chapter 3, but is repeated here for continuity.

4.2.2 The Connection Block

The connection block used is illustrated in Figure 4.4, where a routing switch is indicated by an X. The channel wires (drawn vertically in the figure) pass uninterrupted through the C block and have the option of connecting to the logic cell pins through the switches. The flexibility of the C block is represented by the variable F_c, which defines the number of tracks that each logic cell pin can connect to. For the example shown in the figure, each logic cell pin can connect to 2 vertical tracks, and so F_c is 2. In this chapter, no assumption is necessary about the implementation of routing switches, except that the switches are assumed to be bi-directional.

4.2.2.1 Connection Block Topology

The *topology* of the connection block (the pattern of the switches) can have a significant effect on routability, particularly when F_c is low. To illustrate this consider Figure 4.5 which shows two different C block topologies and one connection (from pin A to pin B) that must be routed. In this example each logic cell has three pins on a side and there are four tracks per routing channel. By examining the locations of the routing switches, it is clear that it is not possible to route the connection with Topology 1, while

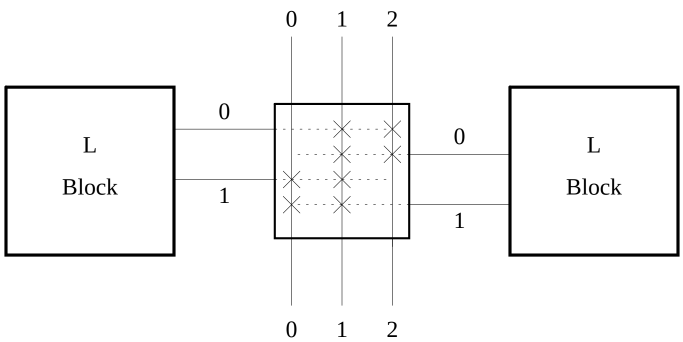

Figure 4.4 - *The Connection Block*

it is possible to do so using Topology 2. Topology 2 works because it has a common wire that can be reached by both pin A and pin B. This example illustrates the fact that a C block topology must provide common wires for every pair of pins that may need to be connected. At the same time, however, it is easy to recognize that it is desirable for the routing switches in the C block to be spread evenly among the tracks, so that there is a reasonable opportunity for each track to be used. A good C block topology should achieve a balance of these tradeoffs. Given W tracks per channel, the design of a C block is straight-forward if F_c is close to W, but for lower values of F_c the C block should be carefully designed. The issue is most acute if $F_c \leq 0.5W$ because at this point some pairs of pins may not have any common wires if the C block is poorly designed.

The topology of the C block that is used for the results presented in this chapter is illustrated by Figure 4.6. In the figure, there are 10 tracks per routing channel, seven pins per logic cell, and F_c is 6. The design of this topology is based on statistics, from a set of

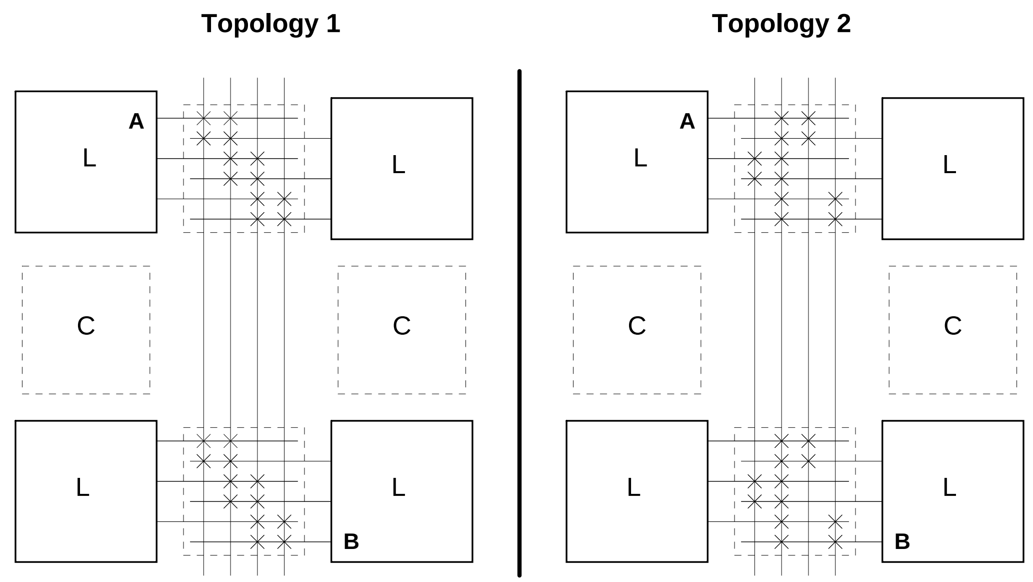

Figure 4.5 - *Two Connection Block Topologies*

circuits, that show how frequently each pair of pins is connected. For pins that are often connected the topology tends to provide common tracks, whereas for pins that are seldom connected different tracks are used. For example, the statistics say that pin 0 (a logic cell input) is often connected to pin 6 (an output), so these two pins share six tracks, whereas pin 0 is seldom connected to pin 5 (an input), so this pair shares only three tracks. This type of analysis is possible because logic cell inputs tend to be connected to outputs, and vice-versa. In this way, the topology provides as much overlap as practical for each pair of logic cell pins, while also balancing the distribution of the switches among the channel wires.

4.2.3 The Switch Block

The general nature of the switch block used is illustrated in Figure 4.7. Its flexibility, F_s, defines the number of other wiring segments that each wiring segment entering an S block can connect to. For the example shown in the figure, the wiring segment at the top left of the S block can be switched to six other wiring segments, and so F_s is 6. Although not shown, the other wiring segments are similarly connected.

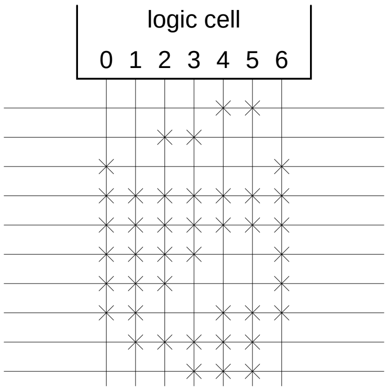

Figure 4.6 - *The C Block Topology Used for this Research*

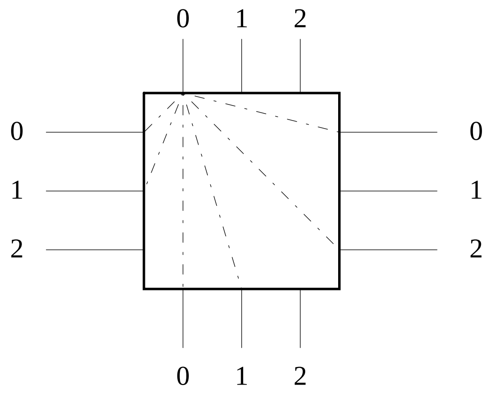

Figure 4.7 - *The Switch Block*

4.2.3.1 Switch Block Topology

The topology of the S blocks can be very important since it is possible to choose two different topologies with the same flexibility measure (F_S) that result in very different routabilities. This is particularly important if the flexibility is low. As an illustration, consider the two different topologies shown in Figure 4.8. In both topologies, the switch block has the same flexibility measure, $F_S = 2$. Assume that a global router has specified that a wiring segment at A must be connected to another at B by traveling through the two switch blocks shown. By examining the routing switches, it is easy to see that it is not possible to reach B from A with topology 1, while it is with topology 2. The reason that topology 2 is successful can be explained as follows. Consider the two vertical wires in topology 2 that connect from A to the two horizontal wires on the right side of the S block. At the next S block, one of the horizontal wires can connect to the top of the block (to B) and one to the bottom. The key is that any turn that is taken at one S block does not prohibit any other turn at the next S block, and this is true for all possible sequences of turns. For the results that are presented in this chapter, topology 2 is used. For higher flexibilities, switches are added such that the basic pattern is preserved.

4.3 Experimental Procedure

This section describes the experimental procedure that is used to investigate FPGA routing architectures. Given a functional description of a circuit, the following steps implement the circuit in an FPGA.

(1) Perform the *technology mapping* [Keut87] of the original network into the FPGA logic cells. This step transforms the functional description of the network into a circuit that interconnects only logic cells of the type shown in Figure 4.2. The technology mapping was done using an early version of the algorithm described in [Fran90].

(2) Perform the placement of the netlist of logic cells. The logic cells are placed by the Altor placement program [Rose85], which is based on the min-cut placement algorithm [Breu77]. Altor makes the resulting two-dimensional array of logic cells as square as possible.

(3) Perform the global routing of the logic cell interconnections. This step finds a path through the routing channels for each pair of logic cell pins that are to be connected.

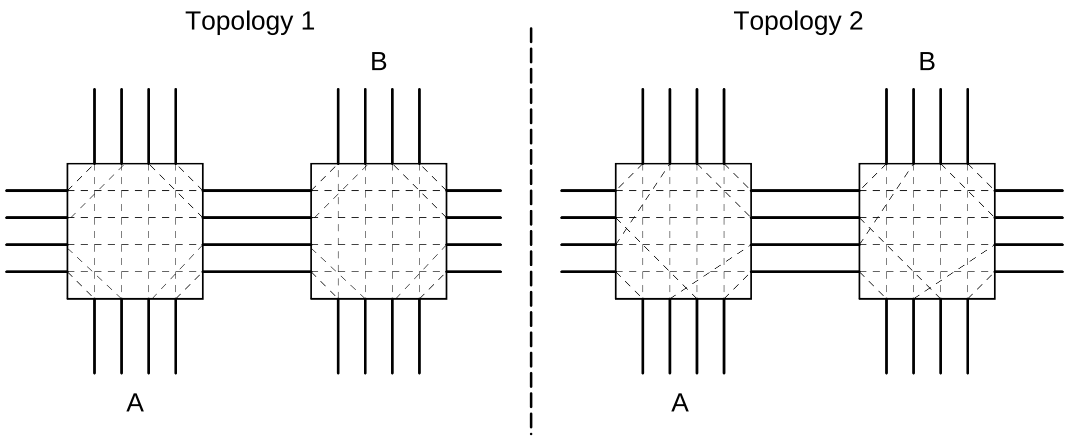

Figure 4.8 - *Two Switch Block Topologies*

Since each connection is assigned to specific channels this determines the *maximum channel density* of the circuit, which is defined as the maximum number of connections that pass through any channel segment. This sets the theoretical minimum number of tracks per channel (for the particular global router used) that is needed to route the circuit. The global router that is used here is based on the LocusRoute standard cell global routing algorithm [Rose90a], that is described in Chapter 2.

(4) Perform the detailed routing of each connection, using the path assigned by the global router. The CGE detailed router, described in Chapter 3, is used for this purpose, and yields two kinds of results. If a specific W (number of tracks per channel) is given as input, CGE determines the percentage of connections that can be successfully routed for specific values of F_s and F_c. Alternatively, if the desired output is the number of tracks per routing channel required to route 100% of connections for a specific F_s and F_c, then CGE is invoked repeatedly, with an increasing number of tracks, until complete routing is achieved.

The salient point in this procedure is that the global router is used only once for each circuit, and this determines the densities of all of the routing channels. The number of tracks required per channel to route each circuit then depends on the flexibility of the routing architecture. Thus, to investigate the effect of flexibility on routability, step (4) is performed over a range of values of F_c, F_s, and W.

4.4 Limitations of this Work

This section discusses the effects of the architectural assumptions and the experimental procedure on the accuracy of the results that are presented later in this chapter.

The models that have been used for the C and S blocks are based on balanced topologies, in that each L block pin can be connected to exactly F_c tracks and each wiring segment that enters an S block can connect to exactly F_s others. Also, every wiring track must use a routing switch to pass through an S block - i.e. all the tracks comprise short wiring segments only. Although it is also interesting to consider other classes of architectures, the assumptions made here allow interesting and useful results to be generated with experiments that have simple parameters.

The experimental procedure described in Section 4.3 limits each connection to a single global route. A better approach would be one that provides a feedback mechanism that allows the detailed router to request a different global route for connections that fail. Finally, the accuracy of the routability results that are presented in this chapter depends on the quality of the routing CAD tools, which includes both the global and detailed routers.

4.5 Experimental Results

The experimental results that are presented here are based on the five circuits that were described in Table 3.1. This section first investigates the effect of the flexibilities of the C and S blocks on the routability of these circuits and shows the tradeoffs that exist between these two blocks. Following this, the effect of different values of F_c and F_s on the number of tracks required per channel is shown. Finally, the effect of the C and S block flexibilities on the total number of switches required in an FPGA is measured.

4.5.1 Effect of Connection Block Flexibility on Routability

Figure 4.9 is a plot of the percentage of successfully routed connections versus connection block flexibility, F_c, for the circuit BNRE. Each curve in the figure corresponds to a different value of switch block flexibility, F_s. The lowest curve in the figure corresponds to the case $F_s = 2$ and the highest curve to $F_s = 10$. The number of tracks, W, is set to 14, which is two greater than W_g, the minimum possible number of tracks as indicated by the global router. The value of $W = W_g + 2$ was chosen to give the detailed router a reasonable chance of success. Using a higher or lower value of W would shift the curves slightly upward or downward, respectively. Figure 4.9 indicates that the routing completion rate is very low for small values of F_c and only achieves 100% when F_c is at least one-half of W. The figure also shows that increasing the switch block flexibility improves the completion rate at a given F_c, but to get near 100% the value of F_c must always be high (above 7 for this circuit).

Table 4.2 summarizes the results for the other circuits. It gives the minimum values of F_c and $\frac{F_c}{W}$ required to achieve 100% routing completion for each circuit, for nine

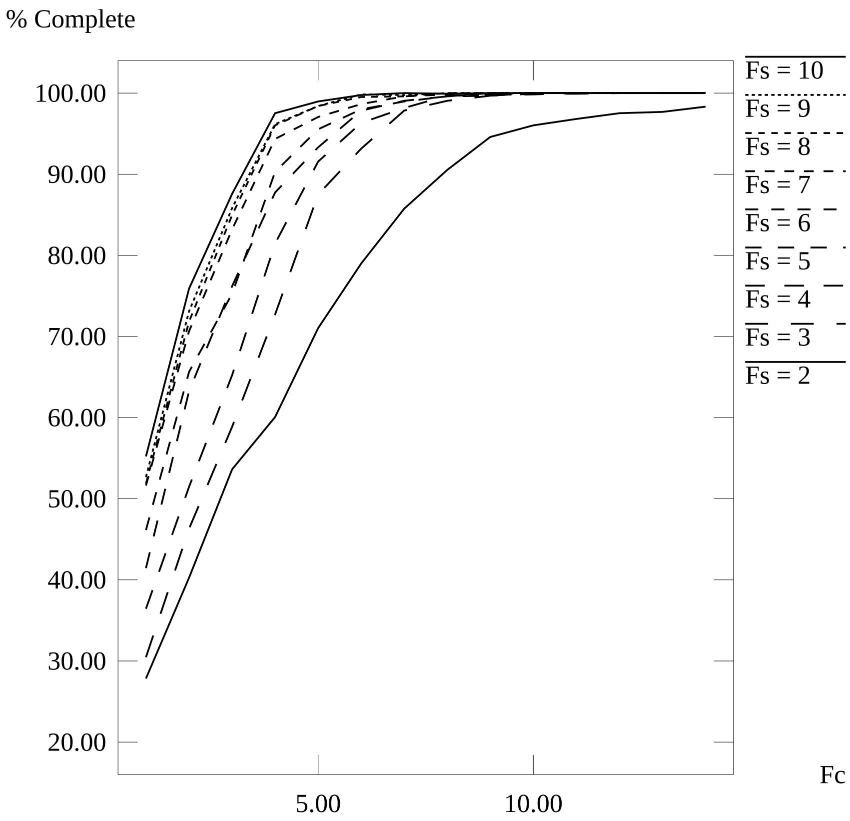

Figure 4.9 - *Percent Routing Completion vs. Fc, Circuit BNRE*

values of F_s. W is fixed at $W_g + 2$, in all cases, to give a reasonable chance for success. The value "nr" in the table indicates that 100% routing was not achieved.

The key observation from the data of Table 4.2 is that there appear to be minimum values of F_c and $\frac{F_c}{W}$ below which circuits are not routable. However, since this data is based on a fixed value of $W = W_g + 2$, it is interesting to investigate whether F_c or $\frac{F_c}{W}$ can be reduced if W is not fixed. To study this, a similar experiment was conducted in which W was allowed to vary to a maximum of $3 \times W_g$. Again, the experiments measure the minimum possible values of F_c and $\frac{F_c}{W}$ for which 100 percent routing can be achieved, for a range of values of F_s. The results are shown in Figure 4.10, which for conciseness gives the average results for the five circuits. The left curve in the figure shows that $\frac{F_c}{W}$ can be substantially reduced by allowing W to vary, but the curve to the right shows that F_c still reaches about the same minimum value.

To see why there exists a minimum value of F_c below which circuits are not routable, consider the following discussion concerning C block topology. Assume that a C block must connect n logic cell pins to a set of tracks, and that some pin, p_i, must be able to connect to all of pins p_j, $1 \le j \le n$. Some connections between these pin pairs will occur within one C block and others will involve two different C blocks. To simplify the analysis, assume that $F_s \le 3$, so that no jogging is allowed among the tracks. As the discussion in Section 4.2.2.1 showed, the C block topology must provide at least one common track that connects to both p_i and each p_j. To accomplish this, the design of the topology may:

Circuit	W	Fs	100% Fc	Fc/W
BUSC	11	2	nr	nr
BUSC	11	3	9	0.82
BUSC	11	4	7	0.64
BUSC	11	5	7	0.64
BUSC	11	6	6	0.54
BUSC	11	7	6	0.54
BUSC	11	8	5	0.45
BUSC	11	9	6	0.54
BUSC	11	10	5	0.45
DMA	12	2	nr	nr
DMA	12	3	8	0.67
DMA	12	4	7	0.58
DMA	12	5	7	0.58
DMA	12	6	7	0.58
DMA	12	7	7	0.58
DMA	12	8	5	0.42
DMA	12	9	5	0.42
DMA	12	10	5	0.42
BNRE	14	2	nr	nr
BNRE	14	3	12	0.86
BNRE	14	4	11	0.79
BNRE	14	5	10	0.71
BNRE	14	6	9	0.64
BNRE	14	7	10	0.71
BNRE	14	8	8	0.57
BNRE	14	9	8	0.57
BNRE	14	10	7	0.50
DFSM	13	2	nr	nr
DFSM	13	3	9	0.69
DFSM	13	4	9	0.69
DFSM	13	5	9	0.69
DFSM	13	6	8	0.62
DFSM	13	7	8	0.62
DFSM	13	8	7	0.54
DFSM	13	9	7	0.54
DFSM	13	10	7	0.54
Z03	13	2	nr	nr
Z03	13	3	10	0.77
Z03	13	4	9	0.69
Z03	13	5	9	0.69
Z03	13	6	9	0.69
Z03	13	7	7	0.54
Z03	13	8	7	0.54
Z03	13	9	7	0.54
Z03	13	10	7	0.54

Table 4.2 - *Minimum F_c Required for 100% Completion*

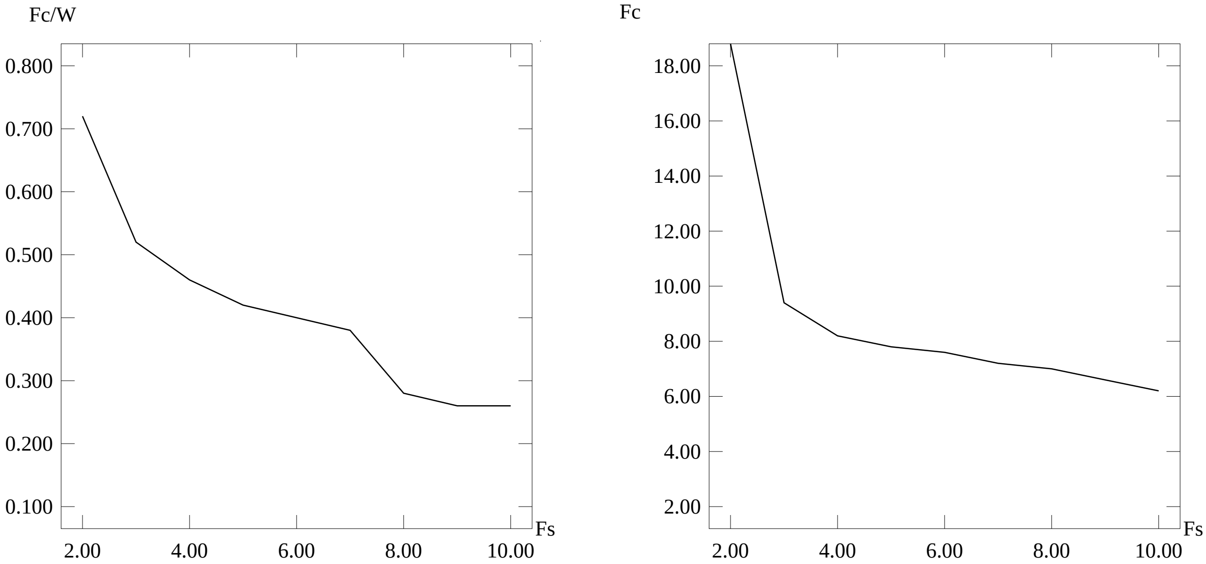

Figure 4.10 - $\frac{F_c}{W}$ *vs.* F_s *and* F_c *vs.* F_s*, With a Variable W*

(1) Attach two switches to each of n different tracks, such that each track connects one p_j to p_i. In terms of Section 4.2.2.1, this corresponds to spreading the switches evenly across the tracks.

(2) Attach n switches to any one track, such that p_i can connect to any p_j on that track.

(3) Use a combination of options (1) and (2).

Option (1) leads directly to a minimum value for F_c because it entails attaching n switches to p_i. The effect of option (2) is more subtle, as discussed below. Consider Figure 4.11, which shows a C block topology in which each pin connects to exactly the same F_c tracks as every other pin. The figure shows that, with this topology, when one pin is connected to a track, one choice of track is eliminated for every other pin. In this scenario, it follows that the minimum possible value of F_c is determined by the maximum number of pins that are connected at any C block.

A more realistic C block, such as the one that was shown in Figure 4.6, is based on option (3). This means that a combination of the effects of options (1) and (2) determines a minimum value for F_c. The key to this discussion is that any realistic C block must provide connections between a number of different pairs of pins and this leads directly to a minimum possible value for F_c.

Note that the minimum value of F_c can be reduced slightly by increasing F_s to be above three, because this increases the connectivity between pairs of pins by allowing jogging from one track to another. However, this only affects connections that involve two different C blocks, and since some connection's pins are both within one C block, an absolute minimum value exists for F_c.

4.5.2 Effect of Switch Block Flexibility on Routability

Figure 4.12 is a plot of the percentage routing completion versus switch block flexibility, F_s. Each curve in the figure corresponds to a different value of F_c, with the lowest

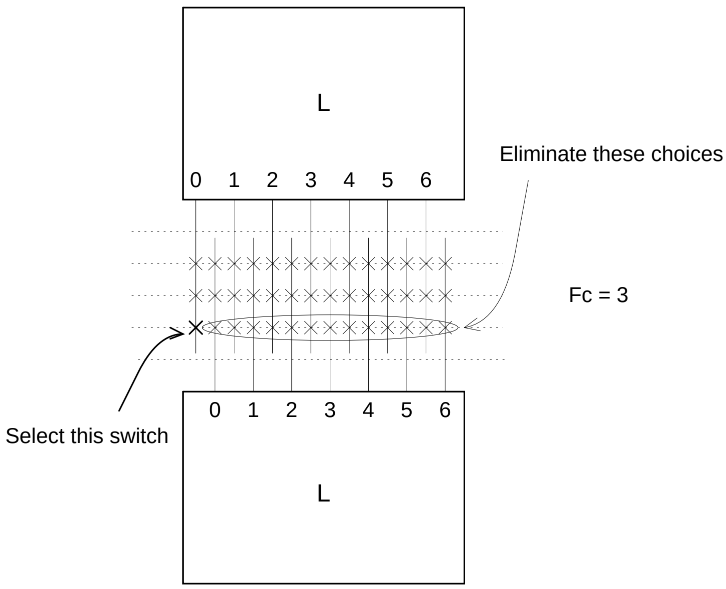

Figure 4.11 - *Connecting One Pin Eliminates One Choice for Every Other*

curve representing $F_c = 1$ and the highest curve corresponding to $F_c = W$. This plot is for the circuit BNRE, with W set to 14. The figure shows that if F_c is high enough, then very low values of F_s can achieve 100% routability. These F_s values are low in comparison with the maximum possible value of F_s, which is $3 \times W$. For the results in Figure 4.12 this maximum is 42, whereas 100% routing completion is often achieved for $F_s = 3$. This makes intuitive sense because even for $F_s = 3$ every track that passes through a particular C block is guaranteed to connect to one other track at every other C block. To further quantify the effect of F_s on routability, consider the connection of two logic cell pins that are separated by n S blocks. The number of tracks connectable at the first logic cell pin is F_c and the number of paths available to reach the connection block adjacent to the second logic cell is

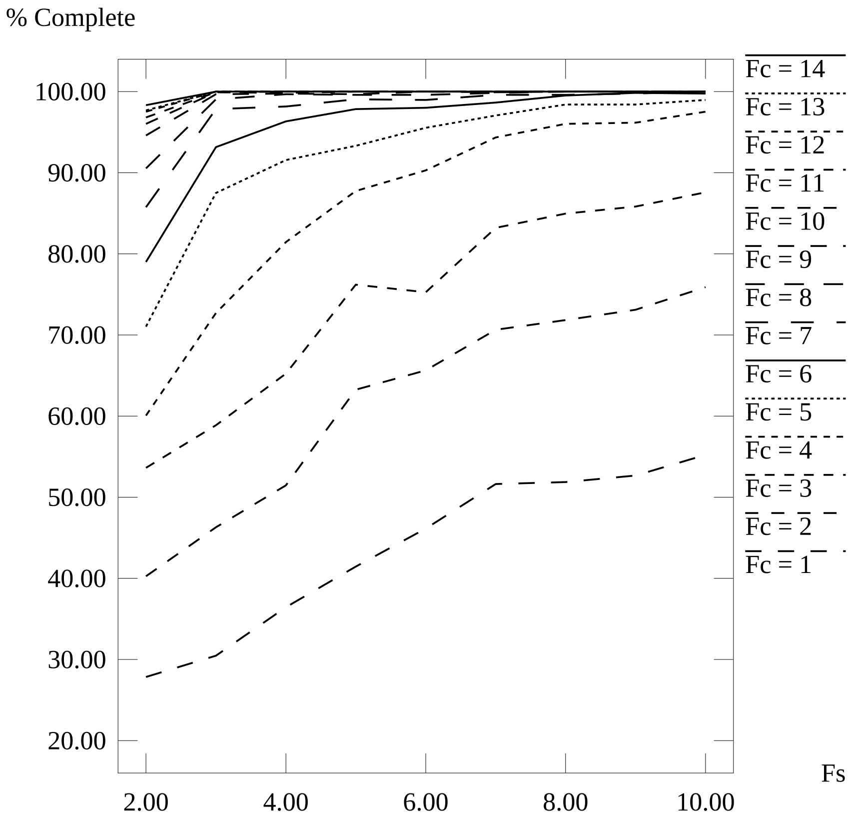

Figure 4.12 - *Percent Routing Completion vs. Fs, Circuit BNRE*

$$\# \, Paths = F_c \times \left(\frac{F_s}{3} \right)^n$$

Using the average value of n of about 3 for typical circuits, if $F_s = 3$ and $F_c = 10$, then there are 10 paths available. If F_s is increased to 6, there are 80 paths available. Thus a small increase in F_s greatly increases the number of paths, and hence the routability.

4.5.3 Tradeoffs in the Flexibilities of the S and C Blocks

Figures 4.9 and 4.12 can be combined in three-dimensions to show that a tradeoff exists between the flexibilities of the S blocks and the C blocks. This is illustrated by the three-dimensional surface plot in Figure 4.13. The plot shows, for example, that a decrease along the F_c axis can be compensated for by an increase along the F_s axis, and vise-versa. This can also be seen in Figure 4.14, which is a plot of the minimum value of F_c (averaged over the five circuits) for which 100 percent routing can be attained for a range of values of F_s. In this plot, W is constant for each circuit, at two higher than W_g. Note that there is no data point for the case $F_s = 2$ in the figure because the circuits are not routable for that S block flexibility with $W = W_g + 2$. The slope of the curve in Figure 4.14 will flatten for higher values of F_s since there exists a minimum value of F_c for each circuit, as was discussed in Section 4.5.1.

Note that Figure 4.14 alone does not provide enough information to choose the best values of F_c and F_s because it is based on the fixed value of W. The required value of W for different values of F_c and F_s is the subject of the following section.

4.5.4 Track Count Requirements

This section investigates the effect of the S and C block flexibilities on routability by measuring the number of tracks per channel, W, required to route 100 percent of a circuit's connections for specific values of F_s and F_c. For these experiments, the detailed

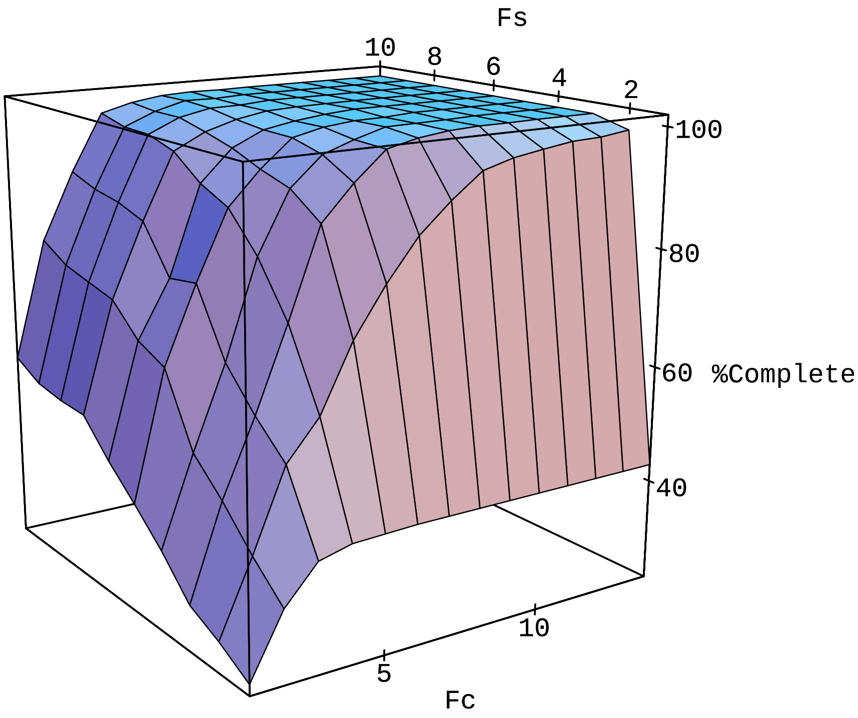

Figure 4.13 - *C Block and S Block Tradeoff*

router was invoked repeatedly over a range of values of W, from W_g (the maximum channel density of the circuit) to a maximum of $3 \times W_g$, until 100 percent routing was achieved. Each routing architecture is assessed by comparing how close the required W is to W_g. For these experiments the flexibility of the C block is expressed as the ratio of F_c over W - as W is changed, the ratio is kept constant. This provides a convenient way to coordinate changes in W with the number of routing switches in a C block. This is automatic for an S block, since F_s applies to each track.

Table 4.3 shows the number of tracks required to achieve 100 percent routing completion for circuit BNRE over a set of FPGA routing architectures. Each entry in the table corresponds to a different pair of $(F_s, \frac{F_c}{W})$ values. The value "nr" in the table

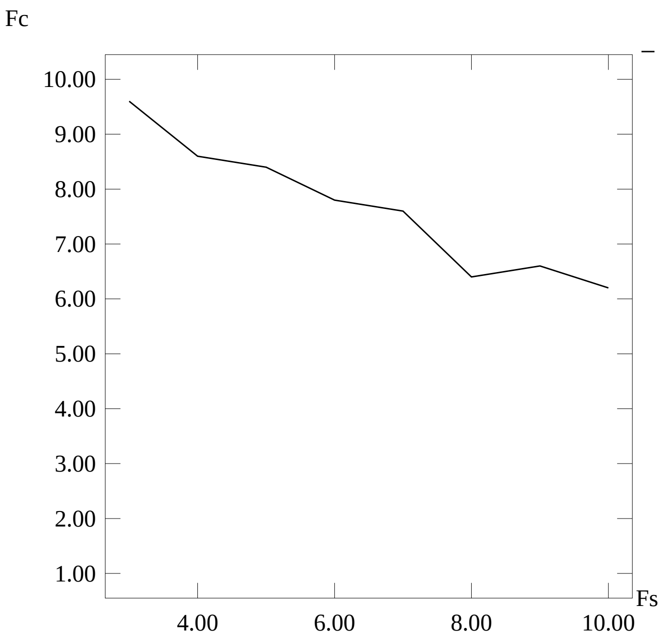

Figure 4.14 - *F_c for 100% Routing vs. F_s, With $W = W_g + 2$*

means that 100 percent routing was not achieved. Since W_g for this circuit is 12, the table shows that even with low flexibilities it is possible to achieve 100 percent routing using very near to the minimum possible number of tracks.

Table 4.4 summarizes the results for all the circuits. Each entry in this table is the average over all the circuits of the number of tracks that are required *in excess* of the minimum (W_g) to route 100 percent of the connections. These results show that with very low flexibilities it is possible to achieve a number of tracks only slightly greater than the minimum. In particular, for $F_s \geq 3$ and $\frac{F_c}{W} \geq 0.6$, excluding the case ($F_s = 3$, $\frac{F_c}{W} = 0.6$), the number of tracks required in excess of the minimum is no greater than three.

Fs	Fc/W 0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
2	nr	nr	nr	nr	nr	23	23	17
3	nr	nr	nr	18	14	13	13	13
4	nr	nr	18	14	13	13	13	13
5	nr	nr	18	14	13	12	12	12
6	nr	21	16	14	14	12	13	13
7	nr	19	17	14	12	12	12	12
8	nr	19	15	13	12	12	12	12
9	23	17	13	13	12	12	12	12
10	19	16	13	13	12	12	12	12

Table 4.3 - *Track Count Requirements for BNRE (Minimum = 12)*

4.5.5 Architectural Choices

The choice of a particular FPGA routing architecture must be based on the cost of its implementation in terms of area and delay. Although these depend on the technology used for the routing switches, it is possible to make general comments because, regardless of their implementation, routing switches consume area and cause time delays. It is instructive to examine the total number of routing switches required by different routing architectures. The number of switches in a C block and an S block depends on the number of logic cell pins (P), the number of sides that each pin appears on (T), and on the parameters F_c, F_s, and W, according to the following equations:

$$\#\ Switches\ in\ Connection\ Block = \frac{1}{2} \times T \times P \times F_c$$

$$\#\ Switches\ in\ Switch\ Block = 2 \times F_s \times W$$

For the results presented here P is 7, T is 2, and the last three parameters are read from the equivalent of Table 4.3 for each circuit. As the flexibility of the routing structures

Fs	Fc/W 0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
2	nr	nr	nr	nr	nr	11.2	10.8	9.0
3	nr	nr	nr	4.6	2.4	1.2	0.8	0.8
4	nr	nr	6.2	3.0	1.6	0.6	0.6	0.6
5	nr	nr	4.2	2.4	1.0	0.4	0.2	0.2
6	nr	7.6	3.8	1.8	0.8	0.4	0.4	0.4
7	nr	5.2	3.4	1.4	0.2	0.2	0.2	0.2
8	nr	4.4	1.4	0.6	0.2	0.0	0.4	0.2
9	10.0	4.2	1.0	0.4	0.2	0.2	0.0	0.0
10	8.0	3.2	1.4	0.6	0.2	0.0	0.0	0.0

Table 4.4 - *Average Excess Track Count Requirements over all Circuits*

increases (F_s and F_c) the number of switches per track increases, but the number of tracks per channel may decrease, as shown in Table 4.3. Hence there should be an architecture that exhibits a minimum total number of switches. Table 4.5 gives the number of switches *per tile* required for circuit BNRE for each FPGA routing architecture. A tile is the section of the FPGA that would be replicated across the entire chip, and includes the logic cell, two connection blocks and one switch block.

As Table 4.5 shows, flexibilities of $F_s = 3$ and $\frac{F_c}{W} = 0.7$ achieve a minimum number of switches for this circuit, at 221. Note that several neighboring architectures have similar switch counts. For all of the test circuits the minimum number of switches was between 172 and 223, and occurred when the architecture's parameters were in the range $3 \le F_s \le 4$ and $0.7 \le \frac{F_c}{W} \le 0.8$.

Fs	Fc/W 0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
2	nr	nr	nr	nr	nr	349	381	306
3	nr	nr	nr	259	221*	223	241	260
4	nr	nr	270	229	231	249	267	286
5	nr	nr	306	257	257	254	271	288
6	nr	369	304	285	305	278	319	338
7	nr	372	357	313	285	302	319	336
8	nr	410	345	317	309	326	343	360
9	510	401	325	343	333	350	367	384
10	459	409	351	369	357	374	391	408

Table 4.5 - *Average Number of Switches per Tile for Each Architecture*

4.6 Conclusions

This chapter has explored the relationships between the flexibility of routing architectures and routability in FPGAs. The principal conclusions are that connection blocks should have high flexibility to achieve high percentage routing completion, but that a relatively low flexibility is sufficient in the switch blocks. Furthermore, with low flexibilities the number of tracks per channel required is very close to the minimum. Finally, it has been shown that routing architectures with these properties yield the lowest total number of routing switches.