

Yield Enhancements of Design-Specific FPGAs

Nicola Campregher, Peter Y.K. Cheung, George A. Constantinides, Milan Vasilko*
Electrical & Electronic Engineering, Imperial College London, UK

* School of Design, Engineering and Computing, Bournemouth University, UK
{nc300,p.cheung,g.constantinides}@imperial.ac.uk, m.vasilko@computer.org

ABSTRACT

The high unit cost of FPGA devices often deters their use beyond the prototyping stage. Efforts have been made to reduce the part-cost of FPGA devices, resulting in the development of Design-Specific FPGAs. These parts offer cost reductions by limiting manufacturing tests and improving the number of working devices in a wafer. This paper addresses the issue of yield enhancement in Design-Specific FPGAs. In this paper, an analytical model predicting the probability of mapping a specific design onto potentially defective FPGAs is developed. When combined with existing yield modelling techniques, a quantitative measure of the potential yield improvements of the Design-Specific FPGA approach is reported for current and future technology nodes. It is found that this approach, while beneficial with current manufacturing technology, may not be suitable for 22nm technology or beyond.

Categories and Subject Descriptors

B.8.1 [Hardware]: Performance and Reliability—*Reliability, Testing, and Fault-Tolerance*; B.7.3 [Hardware]: Integrated Circuits—*Reliability and Testing*

General Terms

Reliability

Keywords

yield prediction, FPGA interconnect, yield enhancement, interconnect faults, Structured ASIC, Design-Specific FPGA, interconnect utilization.

1. INTRODUCTION

With ever increasing ASIC manufacturing costs, companies are looking for alternatives to ensure that their small to medium volume devices can be manufactured at reasonable costs. Although FPGAs have long been used for prototyping

designs, their higher unit cost when compared with ASIC often discourages their use in production. Various attempts have been made to reduce the part-cost of designs that originate from FPGA prototypes [1]. One such effort is based on the idea of restricting an FPGA part to one specific design, an approach designated **Design-Specific FPGA** in this paper. Xilinx's Easypath [2] is one example. The cost saving of this approach is mostly due to two factors: 1) It reduces the cost of manufacturing test because only limited functionality appropriate to the specific design is required; 2) Devices may contain defects as long as faulty resources are not used by the design, thus increasing the number of working devices that can be obtained from a wafer. The device therefore cannot be guaranteed to function under all circumstances, and is expected to be used with only one bitstream, losing the possibility to reprogram the FPGA. This paper investigates the potential yield enhancement of the Design-Specific FPGA approach for current and future technology nodes. Only catastrophic interconnect defects are being considered in this work.

The impact of interconnect defects on yield of current and future FPGA devices was first reported in [3]. It was shown that the yield expected for devices manufactured using 22nm technology can be very low for large devices, suggesting that some form of fault tolerant scheme is likely to be required. In a later study the yield enhancement offered by three different fault tolerant schemes was reported [4]. The issue of yield enhancement gained by switch-box with redundancy has also been studied in a recent work [5], suggesting that the problem is being addressed by manufacturers and researchers. As far as the authors are aware, no published work so far addresses the issue of yield enhancement in Design-Specific FPGAs. The novel contributions of this paper are: 1) To develop a probabilistic model that a given design would map to a FPGA device using the Design-Specific approach without the need for redundancy or rerouting; 2) To combine this probabilistic model with the interconnect yield model developed in [4] so that the yield enhancement of Design-Specific FPGAs can be quantified for current and future technology nodes; 3) To investigate the probability that more than one design may be mapped to a given device.

This paper is structured as follows. Section 2 introduces the concept of Design-Specific FPGAs and explains some of the issues related to this approach. Section 3 provides a brief account of the yield models presented in [3, 4]. In Section 4, the problem of yield prediction for Design-Specific FPGAs is formulated. Section 5 presents the results of our analysis

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

FPGA'06, February 22–24, 2006, Monterey, California, USA.
Copyright 2006 ACM 1-59593-292-5/06/0002 ...\$5.00.




Figure 1: Split of costs of IC design and manufacturing [2]

and their implications, and finally Section 6 concludes the paper and suggests areas for future research on the subject.

2. DESIGN-SPECIFIC FPGAS

With ever increasing development and manufacturing costs, as shown in Figure 1[1], small to medium companies have begun to look for alternatives to ASICs. While for small sized designs FPGAs and CPLDs can still provide good returns, for big designs FPGAs are often too expensive. The introduction of Structured ASICs has filled a gap in the market, allowing the fast and cheap development of small to medium volume devices.

Structured ASICs are usually made of an array of standard logic cells, connected by custom interconnect layers [6]. By using a common logic layer and only a limited number of custom metal layers, non-recurring engineering (NRE) costs are greatly reduced, allowing small volumes to be produced at much lower costs. Furthermore, the use of prefabricated common logic layers allows for faster turnaround time, which is becoming an important factor as the lifespan of devices reduces [7].

FPGAs manufacturers have recently introduced their own solutions to compete with Structured ASICs, which threaten a significant portion of the FPGA market, as shown in Figure 2 [2]. In particular, FPGA manufacturers have concentrated on efficient migration from the prototyping stage, which is often carried out using FPGAs, to volume manufacturing. While some manufacturers have implemented their own Structured ASIC architecture coupled with a CAD flow to migrate FPGA bitstreams to the Structured ASIC [1], others have identified advantages in using partially tested FPGA devices as non-reprogrammable parts and only guaranteed to work with given bitstreams [2]. It is the latter approach which is the subject of this work.

When customers begin their prototyping stage, they are offered a defect-free device to implement their design. Once the design is finalized, the bitstream is sent to the FPGA manufacturer, who prepares custom test vectors to test the devices. The chips which pass the tests are then guaranteed to work only with the provided bitstream, and only minor tweaks are allowed [2]. This method assures full compatibility with the prototype FPGA, and the timing requirements are guaranteed to be met, at lower overall device cost.

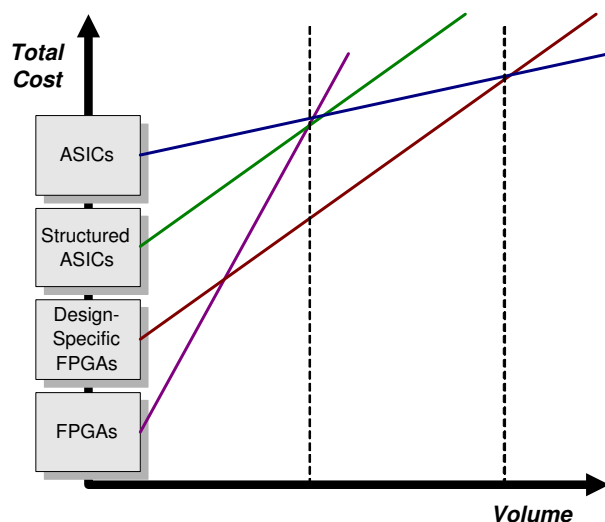


Figure 2: Cost vs Volume for different ASIC and FPGA Methodologies[13]




Figure 3: Catastrophic faults relative to size. Similar sized defects may only cause a fault if a pattern is broken or two patterns joined

3. YIELD MODELING

This section explains the concept of critical area and how it is used to predict device yield. Defects are assumed to be square, with side dimensions x . The critical area analysis applies to metal patterns drawn on a single fabricated metallization layer. Modern chips are constructed with multiple metallization layers, meaning that the critical area analysis has to be carried out for each layer with different parameters. Unless specified, the yield in this paper refers to a single metal layer yield.

3.1 Critical area

The critical area of a lithographic pattern is defined as the portion of the total chip area within which the occurrence of a defect results in a fault [8]. In more general terms, a defect of size x will only cause a fault if its center falls in a particular section of the chip, as shown in Figure 3. Defects of equal size may or may not cause a catastrophic fault, depending on where their centers fall.

The critical area is defined in (1):

$$A_C = A_{Total} \int_0^{\infty} K(x)S(x)dx \quad (1)$$

where A_{Total} is the total die area, x is the defect size, $K(x)$ is the *fault probability kernel*, and $S(x)$ is the defect




Figure 4: (a) Parameters in the fault probability kernel $K(x)$. L , w and d are architectural parameters, while h is dependent on x , the size of the defect. (b) Fault probability kernel $K(x)$, defect size distribution $S(x)$, and fault probability $\phi(x)$. Note that the majority of defects have size similar to the minimum feature size x_0 .

size distribution. The integral term is sometimes referred to as $\phi(x)$, the *fault probability*. Figure 4(b) shows a graphical representation of these functions.

The fault probability kernel shows how the portion of the defect-sensitive chip area varies as the defect size varies. For the sake of clarity, only the case for a single metal line susceptible to an open fault is shown.

The critical area for open defects for a single, metal interconnect is depicted in Figure 4(a) as the area enclosed within the two dotted lines. w is the width of the conducting paths, and d is defined as the minimum strip of metal needed in order to guarantee conduction. The total critical area is thus $L(w + 2h)$, and $h = x/2 - (w - d)$. Details on how to calculate a value for d are given in [4]. The fault probability kernel $K(x)$ is shown in Figure 4(b).

There has been much discussion regarding the defect size distribution [9]. It is now widely accepted that $S(x)$ should increase linearly until the defect size reaches the minimum feature size x_0 , which is known as the critical defect size (see Figure 4(b)), and fall away from a maximum value as a function that is inversely proportional to the cube of the defect size [8]. Defects of size smaller than x_0 , are not considered as they will not result in a catastrophic fault.

3.2 Yield equations

For a non-constant defect density, the probability of finding m defects in a chip of critical area A_C , assuming a defect density D , is given by (2), where $f(D)$, α and B are defined by (3). α is known as the *clustering parameter*, whereas D_0 is known as the *average defect density* [8], and $\Gamma(\cdot)$ is the gamma function.

$$p(m, A_C, D) = \int f(D) \frac{(A_C D)^m e^{-A_C D}}{m!} dD \quad (2)$$

$$f(D) = \frac{1}{\Gamma(\alpha) B^\alpha} D^{\alpha-1} e^{-\frac{D}{B}}, \alpha = \frac{D_0^2}{\text{var}(D)}, B = \frac{\text{var}(D)}{D_0} \quad (3)$$

Combining (2) and (3) results in (4).

$$p(m, A_C, D) = \frac{\Gamma(\alpha + m)}{m! \Gamma(\alpha)} \frac{(A_C D_0 / \alpha)^m}{(1 + A_C D_0 / \alpha)^{m+\alpha}} \quad (4)$$

As the yield is the probability of obtaining defect-free




Figure 5: Interconnect metal layer. The gaps between the metal lines account for vias connections between layers

chips, the yield prediction is made using (5). Defect free yield is referred to as the functional yield.

$$Y = p(0, A_C, D) = \frac{1}{(1 + A_C D_0 / \alpha)^\alpha} \quad (5)$$

3.3 Metal Layer model

FPGAs have, by nature, a regular, repeating structure. Their logic architecture is formed by an array of identical logic blocks and switch matrices. As a result, all the metal connections between logic blocks are also regularly shaped and distanced.

FPGAs offer lines of specific length to connect one logic block to another. An interconnect metal layer can therefore be modelled as a collection of lines of similar length, grouped in channels, leading from one logic block to another. A model of a possible metal layer design is shown in Figure 5.

The parameters used to define the model are defined below:

- M - width and height of CLB array in the FPGA. The device is assumed to be a square array of $M \times M$ CLBs.
- $lines$ - Number of interconnects in a wiring channel.
- L - Length of line. This measure differs depending on the metal layer.
- w - width of the conduction path.
- s - space between conducting paths.

For simplicity, it is assumed that the width and the space between paths have identical size. The size of each parameter can then be found by halving the wire pitch value. Short lines are manufactured on lower layers, whereas the higher layers host the longer, global lines.

All inter-layer patterns (vias, contacts) are assumed to be contained in the areas above the switch matrices. It is therefore possible to model all lines as straight, parallel patterns equally spaced between each other. The patterns are only broken over the switch matrices, which are regularly arranged in the FPGA logic array.

The predicted array size is calculated assuming that M is approximately inversely proportional to the minimum feature size. It is further assumed that halving the minimum feature size will result in doubling the parameter M .

$$K(x) = \frac{1}{A_{Total}} \begin{cases} 0 & x < w - d \\ M^2 * lines * L' * (x - (w - 2d)) & w - d \leq x < s + 2(w - d) \\ H * M * L' & s + 2(w - d) \leq x < \frac{W}{M} - L \\ H * W & x \geq \frac{W}{M} - L \end{cases} \quad (6)$$

$$\text{where } L' = L + \sqrt{x^2 - (w - d)^2}$$

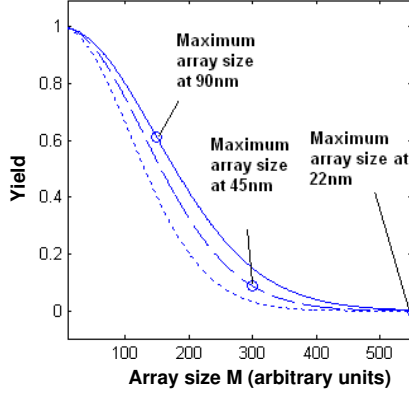


Figure 6: Yield losses due to interconnect defects versus FPGA array size M (for $M \times M$ array of CLBs) [8].

With regards to the metal layers, it is assumed that the silicon space is used to a maximum, i.e. there are no free areas on the silicon. Area not occupied by the metal lines may be occupied by vias and contacts.

Calculating the fault probability kernel for such a structure is a relatively trivial task. For open defects, the fault probability kernel is shown in (6).

Using the data from the SIA roadmap [10] in Equation 5 it has been possible to analyze the yields of FPGA devices as array size increases. The main results have been summarized in Figure 6. The graph also highlights our predictions of maximum array size for devices manufactured using more advanced technology nodes [3, 4].

4. PROBLEM FORMULATION

This section describes the analysis carried out to examine the yield of Design-Specific FPGAs. Section 4.1 describes the main assumptions used and the analysis of the probability that a given design would map to a device exhibiting catastrophic interconnect faults, while Section 4.2 describes how the work is extended to provide an analysis of multiple design mappings.

4.1 Probability of Successful Mapping

Considering a design which needs to map n signals on a layer made of k interconnects, we aim to find the probability that all signals are assigned to faultless interconnects, given that faulty interconnects exist on the layer. k , the total number of interconnects on any given layer, is a parameter of our interconnect model, calculated assuming minimal line spacing and no free silicon left on the device.

The probability that a signal is assigned to an interconnect j at location x, y on the layer is modelled as a bivariate

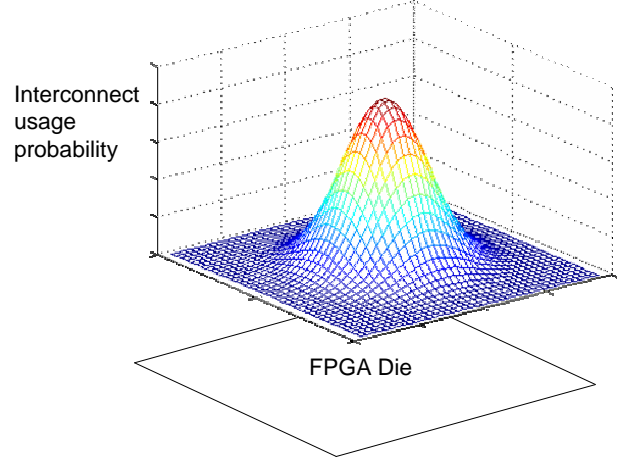


Figure 7: $f_j(x, y)$, a bivariate normal distribution, indicates higher congestions towards the middle of the device.

normal distribution $f_j(x, y)$, as shown in Figure 7.

$$f_j(x, y) = \frac{1}{2\pi\sigma_x\sigma_y} e^{-\frac{1}{2}\left(\frac{x^2}{\sigma_x^2} + \frac{y^2}{\sigma_y^2}\right)} \quad (7)$$

The function reflects the higher usage of interconnects in the middle of the device, a behavior observed in the operation of the Place and Route tool used for the analysis, VPR [11], which tends to create congestions towards the middle of the device [12].

The probability that an interconnect at location x, y exhibits a catastrophic fault, $g_j(x, y)$, is modelled as a bivariate gamma distribution, expressed in Equation 8, where D is the wafer defect density and α is the clustering factor, as introduced in Section 2.2 [8]. The bivariate gamma distribution models the spatial distribution of defects, as explained in Section 2.2.

$$g_j(x, y) = \frac{1}{\Gamma(x, y) \left[\frac{\text{var}(D)}{D_0}\right]^\alpha} e^{-\frac{DD_0}{\text{var}(D)}} \quad (8)$$

For any set S of interconnects the probability that S is a set of usable lines is given by the function $G(S)$ in (9) [13].

$$G(S) = \prod_{j \in S} g_j(x, y) \prod_{j \notin S} (1 - g_j(x, y)) \quad (9)$$

The probability that all signals are assigned into set S is given by (10).

$$F(S) = \left[\sum_{j \in S} f_j(x, y) \right]^k \quad (10)$$

Line length	Segment Frequency	Fc	Fs
1	0.33	1	1
4	0.25	1	1
8	0.25	0.33	1
long	0.17	0.5	1

channel width - 24

Table 1: Architecture used for Place and Route analysis. The parameters Fc and Fs indicate the connection block and switch block population respectively

And the overall probability that a successful signal assignment is achieved over a single metal layer is given by (11).

$$P(\text{success}) = \sum_S G(S)F(S) \quad (11)$$

The probability that a successful signal assignment is made over all layers is the product of the individual layer probabilities.

The final part of the analysis is concerned with interconnect utilization. Interconnect utilization is independent of logic utilization, and is often dependent on the nature of the design as opposed to its size. For example, the requirements for memory and buses in a design is likely to affect the routing utilization independent of the logic design underneath.

The aim of this part of the analysis is to quantify the number of active metal lines on any given metal layer. Interconnect of different lengths are likely to be manufactured on different metal layers, using different geometries. Larger geometries are less susceptible to defects, meaning that higher layer yields will be achieved [3]. The scope of this analysis is to quantify the extent to which lines built on layers with higher yields are used with respect to interconnects built on other layers. For the analysis, we assume that longer lines will be built on higher layers and larger geometries.

To quantify interconnect utilization, an FPGA architecture which resembles that of commercial FPGAs has been prepared for VPR, a tool which provides full reports on interconnect usage for each design. The architecture is modelled in VPR using the parameters shown in Table 1. The full suite of MCNC benchmarks [14] has then been placed and routed, to acquire data regarding usage of each interconnect type. The benchmark circuits were placed and routed using the smallest possible array that each design would fit so that utilization figures across all designs are normalized. The interconnect utilization results is summarized in Table 2.

The interconnect utilization ratio is, by definition, the ratio n/k , thereby indicating how many routing resources are being used by a given device. Assuming that each line type is manufactured on the FPGA using two metal layers, one for the vertical interconnects and one for the horizontal ones, the probability that all signals are mapped on the layer using the faultless interconnects can be found using Equation (11). The right most column in Table 2 indicates the average interconnect utilization of each design, which has been used to categorize designs into high and low utilization.

Finally, the last row in Table 2 summarizes the relative interconnect usage for each interconnect type. This has been used to calculate which layers are more likely to be more populated, and allow analysis to be carried out using over-

all interconnect utilization as opposed to individual layer utilization.

The successful mapping probability indicates the proportion of devices that could be used with a specific design even when they may exhibit interconnect faults. The overall yield of a Design-Specific FPGA for a given design is then calculated as the sum of the functional (i.e. fault-free) yield and the fraction of usable devices exhibiting functional faults.

4.2 Multiple bitstreams

If multiple designs are to be mapped to a device, the correlation of the designs will affect the probability of successful mappings. If the designs are identical, the overall success probability is obviously the same as the individual design success probability. If, on the other hand, the designs are very different, the overall success probability is likely to be significantly reduced. It is therefore important to quantify the design correlation when analyzing the overall success probability.

In order to determine the correlation of two designs, we calculate the Degree of Similarity, a measure of how many resources are shared between the designs. A value of 1 for the Degree of Similarity means that the designs are identical, while a value of 0 means that no common resources exist between the designs.

For the case of two designs being mapped to a device, the design netlists are placed and routed using VPR. A comparison of the interconnect resources used from the routing files then yields the Degree of Similarity between the designs. The technique can easily be extended to analyze multiple bistream mappings.

5. RESULTS

The placed and routed results from VPR (Table 2) was used to compute the probability that a given design can be mapped onto a FPGA that may have interconnect defects. The MCNC benchmark circuits were divided into three categories: those with high interconnect utilization (such as *spla*), those with low interconnect utilization (such as *des*) and those with average utilization. Using Equations 4 and 11 we predicted the interconnect yield for conventional and design-specific FPGAs as a function of the complexity of the device as shown in Figure 8. The solid line shows the predicted functional yield (i.e. yield of defect-free devices) due to interconnect defects for devices built using a 90nm process. The dotted lines show the expected yield with design-specific FPGAs for different interconnect utilizations. Assuming that the largest array size built using this technology to be 160 x 160 CLBs [3], the functional yield is around 58%. According to these results, limiting FPGAs to a specific design can improve the yield by 10% to 30% depending on the utilization of interconnect resources. Since Design-Specific FPGAs can tolerate device defects as long as the faulty interconnects are not used, such improvement in yield is expected. Furthermore, the probability of a defective interconnect remains unused for a low utilization design is higher and therefore yield gain is higher as expected.

As manufacturing technology moves to new technology nodes and smaller geometries, it is expected that larger devices will exhibit multiple functional faults. Figure 9 shows the proportion of devices with different number of faults as

<i>Design</i>	<i>Line length percentage utilization</i>				<i>Average</i>
	1	4	8	long	
ex5p	0.337	0.806	0.692	0.827	0.666
tseng	0.069	0.669	0.462	0.676	0.469
apex4	0.283	0.793	0.682	0.78	0.46
misex3	0.25	0.797	0.661	0.833	0.635
alu4	0.155	0.761	0.69	0.72	0.581
diffeq	0.0827	0.715	0.57	0.794	0.54
dsip	0.0382	0.506	0.395	0.257	0.299
seq	0.266	0.804	0.737	0.785	0.648
apex2	0.297	0.813	0.738	0.847	0.674
des	0.0673	0.454	0.374	0.133	0.257
bigkey	0.0543	0.57	0.436	0.214	0.319
s298	0.0987	0.771	0.681	0.756	0.577
spla	0.435	0.826	0.771	0.938	0.743
frsic	0.305	0.794	0.74	0.945	0.696
elliptic	0.178	0.762	0.701	0.907	0.637
ex1010	0.255	0.834	0.765	0.625	0.62
s38417	0.0898	0.737	0.631	0.355	0.453
clma	0.392	0.833	0.799	0.903	0.732
Average	0.169	0.736	0.640	0.633	0.545
Relative usage	7.8%	33.8%	29.3%	29.1%	

Table 2: Interconnect utilization of MCNC benchmarks placed and routed using a FPGA architecture which resembles commercial FPGAs

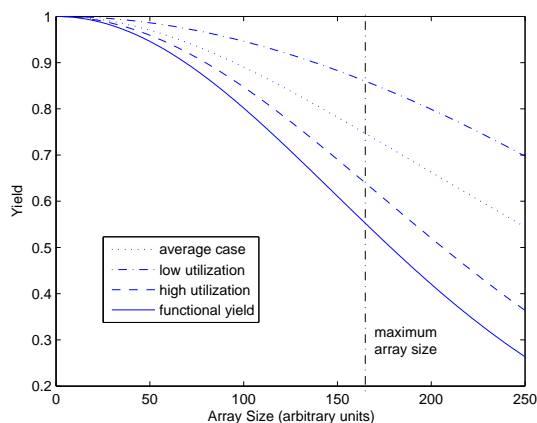


Figure 8: Yield of Structured FPGAs vs Array size (M x M) for devices built using a 90nm process. An average yield gain of 20% is expected over the functional yield, with peaks of 30%.

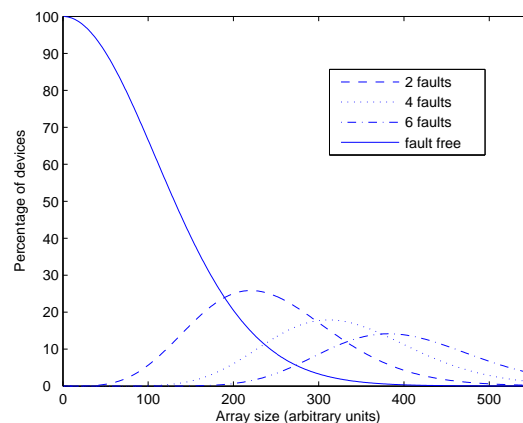


Figure 9: Device yields for 22nm process. Functional yields for larger devices are expected to be close to 0%. The overall yield is likely to consist of devices exhibiting multiple faults

a function of device size for the 22nm process. The solid line indicates the expected fraction of devices which are defect-free. This drops to near zero for devices larger than 400 x 400 CLBs. The dotted lines show the expected percentage of devices with different number of defects. Figure 10 shows the probability of successfully finding a match between a design of different interconnect utilization and an FPGA with one to four defects. As expected, the probability of a match decreases with increasing interconnect utilization. It is also harder to find a match if the device exhibits higher number of defects. In fact for devices with four or more defects, the probability for a match falls below 10% for interconnect utilization as low as 0.3.

Figure 11 shows the maximum achievable interconnect utilization for different target yields in Design-Specific FPGA manufactured in 90nm, 45nm and 22nm processes. For this result, we assume that the largest arrays are used for all technologies (160x160 for 90nm, 300x300 for 45nm and 550x550 for 22nm). This result suggests that for the 90nm process, if the target yield is 50%, the maximum interconnect utilization can be almost 0.6. It also suggests that the Design-Specific FPGA approach does not really help the large 22nm devices. Even for low interconnect utilization (say 0.3), the achievable yield is as low as 10%.

If multiple designs are to be mapped to the same device, the yield also depends on the Degree of Similarity between




Figure 10: Probability of finding a design match for Design-Specific FPGA vs interconnect utilization.




Figure 12: Mapping two bitstreams to a large device manufactured on a 90nm process




Figure 11: Maximum utilization vs target yield in Design-Specific FPGA for 90nm and 22nm processes




Figure 13: An example of multiple bistream mappings on 90nm devices

the designs. As shown in Figure 12, yields of large devices built on a 90nm process rapidly decrease as the number of common interconnect resources between two designs decreases. There exists a region where no yield is expected if two very dissimilar designs, both with high utilization, are to be used. This is due to the high number of faultless resources needed in order to map the two designs.

Figure 13 shows how yields would compare for two designs to be mapped on an FPGA built using a 90nm process. The solid line represents the reference functional yield. The graph shows the yields of two designs exhibiting interconnect utilization characteristics similar to two designs from the MCNC benchmark suite *apex4* (low utilization) and *misex3* (high utilization). These two designs were chosen so that the Degree of Similarity between them could be calculated and used in this analysis. The Degree of Similarity between the *apex4* and *misex3* design was found to be 0.63. This value was kept constant for varying array sizes in order to carry out the analysis. Figure 13 shows how the need to combine the two designs in the same device would affect the yield; for large arrays, the yield loss compared with the high

utilization design yield used on its own is of the order of 5%.

It is relevant to note that, while mapping individual designs offers significant yield advantages even when high utilization is needed, the yield enhancement of mapping two designs to the same device are at most in the region of 10%, which may not justify the use of Design-Specific FPGAs as a yield enhancement scheme.

6. CONCLUSIONS

In this paper we have introduced an analysis of the yield advantages of using FPGAs as one-time programmable devices for small to medium volumes. By using fixed design bitstreams, devices which might exhibit function faults in other conditions can be used to map designs, therefore driving up manufacturing yields. An analysis of interconnect layer usage has been presented, which allowed the yield of a specific design to be calculated.

Using existing yield models it has been possible to analyze the improvements which can be derived by losing the reprogrammability of FPGAs. By using the SIA roadmap predictions it has been possible to study the effects of such a

scheme as more advanced technology nodes are introduced.

The effect of higher interconnect utilization on the yields has been studied, to gauge future requirements for the scheme. Finally, we have examined the success probability of this technique if multiple bitstreams are to be used on a device.

We have shown that significant yield advantages can be achieved using the Design-Specific FPGA approach, however the requirements for successful design mappings denature FPGAs, whose main quality lies in the infinite reprogrammability of devices. As the results in Section 5 show, the potential yield improvement of this approach for 22nm technology and beyond is questionable.

This study has a number of limitations. The interconnect model used may not match the existing heterogeneous FPGAs and could therefore be improved. Verification of the results against manufacturer's data could be made if such data can be obtained. Other assumptions made (such as the bivariate normal distribution in (7)) could be refined. Notwithstanding this, the work provides interesting insights into the benefits and limitations of Design-Specific FPGAs.

7. ACKNOWLEDGMENTS

The authors would like to thank the following people for their helpful discussion relating to this work: Steve Wilton and Robert Isreal of University of British Columbia; Steve Trimberger and Kamal Chaudhary of Xilinx Research Labs; Vaughn Betz, Andrew Leaver, Mike Hutton and Robert Blake of Altera.

8. REFERENCES

- [1] Altera Corp. *Hardcopy 2 Datasheet*. 2005.
- [2] Xilinx Inc. *EasyPath Devices Datasheet*. 2005.
- [3] N. Campregher, P.Y.K. Cheung, G.A. Constantinides and M. Vasilko. Analysis of Yield Loss due to Random Photolithographic Defects in the Interconnect Structure of FPGAs. In *Thirteenth ACM International Symposium on Field-Programmable Gate Arrays*, Monterey, CA, 2005.
- [4] N. Campregher, P.Y.K. Cheung, G.A. Constantinides and M. Vasilko. Yield Modelling and Yield Enhancement for FPGAs Using Fault Tolerance Schemes. In *Proceedings of FPL 2005. 15th International Conference on Field Programmable Logic and Applications*, pages 409-414, Tampere, Finland, 2005.
- [5] A.J. Yu and G.G.F. Lemieux. Defect Tolerant FPGA Switch Block and Connection Block with Fine-grain Redundancy for Yield Enhancement. In *Proceedings of FPL 2005. 15th International Conference on Field Programmable Logic and Applications*, pages 255-262, Tampere, Finland, 2005.
- [6] K.-C. Wu and Y.-W. Tsai. Structured ASIC, Evolution or Revolution? *Proceedings of the International Symposium on Physical Design*, pages 103 - 106, 2004.
- [7] T. Makimoto. The Hot Decade of Field Programmable Technologies. *2002 IEEE International Conference on Field-Programmable Technology (FPT)*, pages 3 - 6, 2002.
- [8] A. V. Ferris-Prabhu. *Introduction to Semiconductor Device Yield Modeling*. Artech House materials science library. Artech House, Boston, 1992.
- [9] H. Sato, M. Ikota, A. Sugimoto, and H. Masuda. A New Defect Distribution Metrology with a Consistent Discrete Exponential Formula and its Applications. *IEEE Transactions on Semiconductor Manufacturing*, 12(4):409-418, 1999.
- [10] Semiconductor Industry Association. *The international roadmap for semiconductor*, 2004.
- [11] V. Betz and J. Rose. Vpr: a New Packing, Placement and Routing Tool for FPGA Research. *Field-programmable Logic and Applications. 7th International Workshop, FPL '97*. Proceedings, pages 213 - 22, 1997.
- [12] Dr. Steve Wilton, University of British Columbia. Private Communication. July 2005.
- [13] Professor Robert Isreal, University of British Columbia. Private Communication. June 2005.
- [14] S. Yang. *Logic Synthesis and Optimization Benchmarks*, version 3.0. Microelectronics Centre of North Carolina, 1991.