

The positive effect on IC yield of embedded Fault Tolerance for SEUs

André K. Nieuwland

Richard P. Kleihorst

Philips Research Laboratories, Prof. Holstlaan 4, 5656 AA Eindhoven, The Netherlands

Phone: +31-40-274436, fax: +32-40-2744113, e-mail: andre.nieuwlan@philips.com

Abstract

Fault tolerant design is a technique emerging in Integrated Circuits (IC's) to deal with the increasing error susceptibility (Soft Errors, or Single Event Upsets, SEU) caused by e.g. alpha particles. A side effect of these methods is that they also compensate for manufacturing defects (the Hard Errors). Currently, yield engineers focus on perfecting the manufacturing process and designers spend their effort in minimizing the area to increase the yield. In this paper, it is shown that increasing the IC area (by applying fault tolerant design techniques) leads under certain conditions to a better yield (more working dies from a wafer) and lower production cost. This is counter-intuitive for many design and yield engineers. To guide designers in deciding when the fault tolerant techniques are beneficial, break-even points between fault tolerant and regular design are presented as function of IC area, fault tolerant overhead and defect density.

1 Introduction

With decreasing device dimensions, the increasing number of devices on an Integrated Circuit (IC) and the increasing clock frequencies, the likelihood of a logic error caused by e.g. an Alpha Particle impact in silicon is increasing [1]. These so called soft errors or Single Event Upsets (SEUs) are of temporal nature, which implies that the physical integrity of the circuit is not affected. However, the temporal integrity is affected when the content of data storage (memories, registers and latches) is changed, or when a faulty logic value propagates through a circuit and is communicated with the outside world. In contrast with these SEUs, there are Hard Errors (HE) which are caused by physical imperfections in the circuit. Typical examples are “stuck-at” faults, open via's and delay faults. Production tests are used to discover this type of defects, and the silicon is discarded if these faults are detected. Especially at the start of a new IC technology generation, just after the “release for production”, there are relatively many imperfections and the yield is low [2]. In such poor processing conditions, one

way to increase yield was to reduce the IC area to minimize the possibility that a defect occurs. Alternatively, redundancy techniques are used on a large scale for regular structures such as memories (spare columns) or array processors (spare processors) [3]. At test time, it is decided if there are defective columns in the memories or processors in the array, which are then being 'replaced' by the already present spare columns/processors. Due to the regularity of these architectures, there is only a low overhead in the redundancy and the impact on IC-size is limited. Fault tolerant design is applicable for fixing hard errors in these cases, although it seems a costly method for regular designs in comparison with redundancy techniques. Fault tolerant design seems to be a better candidate for less regular structures.

For SEUs, many techniques are presented to detect and/or correct errors during operation of the device [1, 4, 5, 6]. Fixing an SEU in a non-temporal way is inherently more complex and thus more (area) costly than fixing hard errors. This area increase is often seen as a cause for lower yield, and thus an increased silicon price. Therefore, the semiconductor industry is reluctant to apply these methods, unless there is a specific request from the customer.

However, in the remainder of this article it will be shown that SEU protection provides a significant increase in yield, depending on the overhead and defect density figures. The organisation of this article is as follows. Section 2 investigates and motivates the use of Fault Tolerant IC Design for yield improvement. Section 3 does a theoretical investigation to establish the achievable results and conditions for a yield increase. Section 4 draws the conclusions and discusses some topics for future research.

2 Build In Self Repair and Yield

2.1 SEU protection for correcting hard errors

SEU protection offers protection against temporal logic errors. As long as these protection schemes are based on spatial methods (i.e. not based on retransmission, or recalculation), they are capable of dealing with logic hard errors too. E.g. an “open” fault, or a “stuck-at” fault will be

properly corrected by the SEU protection scheme. Consequently, protection for SEUs can be used as a Build In Self Repair (BISR) ability for hard errors. The drawback of using the SEU protection for repairing hard errors seems to be that it leaves the IC unprotected for SEUs, because the error correction capability has been exhausted. However, this is only true if the fault tolerant design techniques are applied at top-level. In reality, fault tolerant design techniques are applied at module level [7] and only the error protection of the defective module has been exhausted. The error protection of all the other modules remains untouched, and still offers SEU protection. With modules of reasonable small sizes (up-to 100,000 gates), the majority of the IC area remains SEU protected. Moreover, not all hard and soft error combinations are detrimental. It depends whether the correction method is able to correct multiple bit errors [6], or single bit-errors [8], and whether the SEU and the hard error are visible at the same time. E.g. in a logic tree, an SEU and a “stuck-at” could occur in mutual exclusive paths, such that in all situations only one error would appear at the correction circuitry. If in a memory one bit in some memory word is defective, standard coding techniques for memories [1] such as Hamming protect for SEUs in all other memory words. The possibility that a word with a defect bit is hit by an SEU could be acceptably small. In general, there will be only a marginal reduction in SEU protection at IC level when the BISR capability is used at module level for correcting hard errors. Therefore, with a minor reduction in SEU protection, fault tolerance can be used for hard errors.

2.2 Adding SEU protection versus silicon cost

The IC cost-price is determined by the silicon processing cost (handling of the wafers), costs for test (time and equipment) and packaging. For logic IC’s (DSP’s, ASIC’s) the cost for test and packaging are small compared to the silicon processing cost. These processing costs have to be divided by the number of *working* (read: sell-able) dies on the wafer. Since a small die has a smaller possibility to be defective than a larger die, the yield of a small die is higher. Furthermore, the number of dies that fit on a wafer is larger for small dies than for large dies, which combined with the higher yield, leads to a much larger number of working dies per wafer for small designs than for large ones. This makes small IC’s more cost-effective than large IC’s. Because adding SEU protection or BISR capabilities increases the IC area, less dies will fit on a wafer. Therefore, one would expect the silicon cost of the design to increase.

However, more crucial is the number of *working dies per wafer*. SEU protection schemes and BISR capabilities offer the possibility of fixing hard errors. Defects which normally causes the IC to be discarded, are now tolerable, which leads to a higher yield. Even though less (fault tol-

erant) dies fit on a wafer, the number of *working* dies per wafer can be much higher. This leads to the cumbersome conclusion that larger fault tolerant IC’s can be cheaper than smaller non-fault tolerant IC’s!

The potential yield improvement depends on the defect density: With zero defect density, there is no improvement possible (maximum yield). Note that fault tolerance has become part of the spec because of expected SEUs. The higher the defect density the higher the potential yield improvement. Likewise, larger chips suffer more from defects than smaller ones, hence the potential yield improvement will be larger for larger chips. It is therefore interesting to examine the trade-off between die-size, defect density and fault tolerant overhead, to identify when fault tolerant design has a cost advantage.

3 Yield improvement, theory

Because the economic benefit is determined by the number of working dies per wafer, we need to calculate the number of dies per wafer and the yield per die.

The number of dies per wafer is given by [9] and is written in Equation 1:

$$\frac{\pi(R/2)^2}{A} - \frac{\pi R}{\sqrt{2A}}, \quad (1)$$

where A is the area of the die, and R is the diameter of the wafer. For the numeric examples throughout the paper, we will assume that the wafer diameter R is 20 cm.

For the yield of larger IC’s, the negative binomial distribution is a better approximation than the Poisson Equation [10], because the negative binomial distribution takes large area clustering into account.

Disregarding full wafer destruction, using this approximation, the yield for a single die is given by [10]:

$$Y = \left(\frac{1 + D_0 A}{\alpha} \right)^{-\alpha} \quad (2)$$

Where Y indicates the die-yield, D_0 is the defect density, and α is a process-complexity factor. For the numeric examples in this paper, we consider α to be 3.

To compare the number of working dies per wafer for the protected and un-protected case, we have to equate the number of fault tolerant (FT) dies per wafer with the expected number of working dies per wafer in the non-protected case. To obtain a lower bound for the defect density or die size, we assume that fault tolerant designs are able to fix possible errors and all FT dies will be working properly. To ease the calculation, we express the area of the fault tolerant design as a combination of the area of the original design and a fraction describing the FT induced overhead: $(1 + \delta)A$, where A denotes the size of the original circuit, and δ is the FT overhead expressed as a fraction of the original size.

The break-even between FT designs and non-FT designs determines the boundary between the region where FT does improve yield. Applying FT techniques is only useful for yield if the number of FT dies is larger than the number of working regular dies. The number of fault tolerant IC's per wafer is given by the right-hand part of Equation (3). Since we assume a yield of 100% for the fault tolerant IC's, the number of working fault tolerant dies is equal to the total number of fault tolerant dies per wafer. The left-hand part of Equation (3) represents the yield of the regular design times the number of regular IC's per wafer, which provides the number of working regular dies per wafer.

$$\left(\frac{1 + D_0 A}{\alpha}\right)^{-\alpha} \times \left(\frac{\pi(R/2)^2}{A} - \frac{\pi R}{\sqrt{2A}}\right) \leq \frac{\pi(R/2)^2}{(1 + \delta)A} - \frac{\pi R}{\sqrt{2(1 + \delta)A}} \quad (3)$$

Reorganizing this equation gives us an intermediate result which provides a lower bound for the die-yield, denoted at the left-hand side of Equation (4). If this equation holds, i.e. the yield is below this threshold, it makes sense to expect a yield increase when applying the fault tolerant design. Note that the yield threshold is a function of the original die-size (A), the wafer diameter (R), and the FT overhead fraction δ . As can be expected, this equation shows that if δ increases, the expected yield of the original die has to be lower to keep the FT scheme worthwhile for yield improvement.

$$\left(1 + \frac{D_0 A}{\alpha}\right)^{-\alpha} \leq \frac{R - 2\sqrt{2(1 + \delta)A}}{(R - 2\sqrt{2A})(1 + \delta)} \quad (4)$$

Further reorganization of Equation (4) provides a lower bound for the defect density (D_0) as a function of α , die-size A , FT overhead δ and the wafer diameter (see Eq. (5)).

$$D_0 \geq \frac{\alpha \left\{ \left(\frac{R - 2\sqrt{2(1 + \delta)A}}{(R - 2\sqrt{2A})(1 + \delta)} \right)^{\frac{1}{\alpha}} - 1 \right\}}{A} \quad (5)$$

This defect density indicates whether it is beneficial in terms of yield to apply a yield improvement technique with an overhead fraction of δ for IC's of size A . If the defect density is larger than this value, the yield is higher when using error protection with a (maximum) cost of δ . If the defect density is lower than this value, then a fault tolerant design with overhead δ only reduces the yield of the wafer.

Plotting the lower-bound defect density against die-size and tolerable overhead in a three-dimensional space provides more insight in when fault tolerant design improves the yield. In Figure 1, the X-axis indicates the original die size in cm^2 . The Y-axis the fractional overhead (increase in

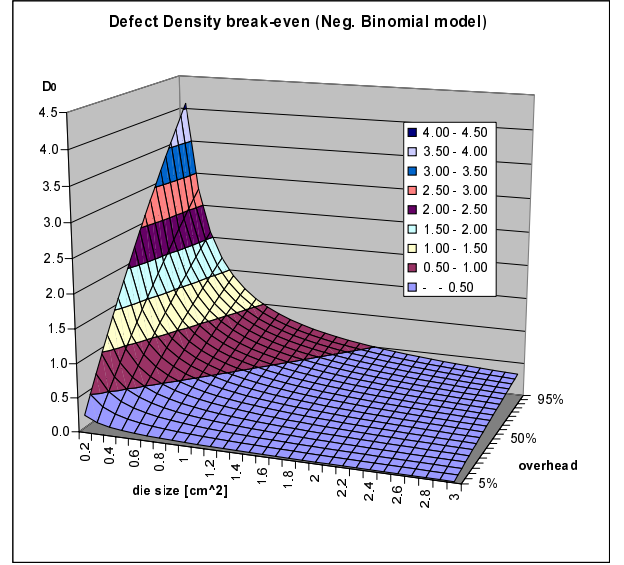


Figure 1. Break-even defect density plane as function of die size and fractional overhead

original area) as a percentage of the original IC area. The Z-axis presents the lower bound defect density. The plane drawn in this figure is the break-even defect density, D_0 as a function of A and δ for which the equality holds in Eq. (5).

For all combinations of area, overhead and defect densities above this plane, it is beneficial to spend the extra area to improve the yield. For all combinations below this plane it is counterproductive; the number of working dies per wafer is less than without the fault tolerant protection.

It is probably even more illustrative to see what can be gained in terms of production quantities. In Figure 2 we plotted the yield per wafer as a function of (un-protected) die size. We assume that fault tolerant IC consists of fault tolerant blocks of 1 mm^2 (approx. 100,000 gates), which are all capable of fixing 1 error per block. We further assume that 7% of the area is not protected (the fraction of correction logic) and that the area overhead introduced by the fault tolerant design is 30%. (These numbers are from our experiments described in [7, 8]). The defect density used for this experiment was set to '1', which is a reasonable low number. We calculated and plotted the yield of the original IC, the fault tolerant IC, and the case when there would not be a single defect on the original die due to a zero defect density (perfect processing). This last number indicates the maximum number of working dies per wafer possible.

In Figure 3, the regularly designed IC has a ratio of '1'. The 'perfect processing' line, is the maximum theoretical yield possible, that is: 100% yield without any fault tolerant overhead. The 'module based FT' line indicates the

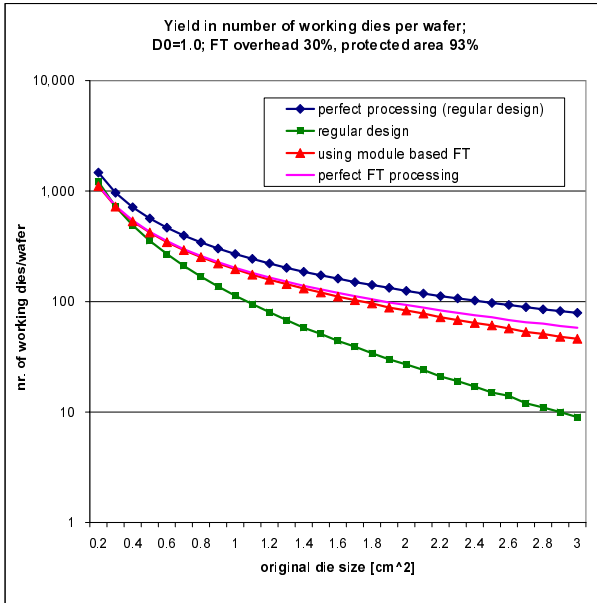


Figure 2. Yield expressed in number of working dies per wafer under perfect and normal processing conditions of the regular and the fault tolerant design. In this graph FT design delivers a higher yield for IC's larger than 0.3 cm^2

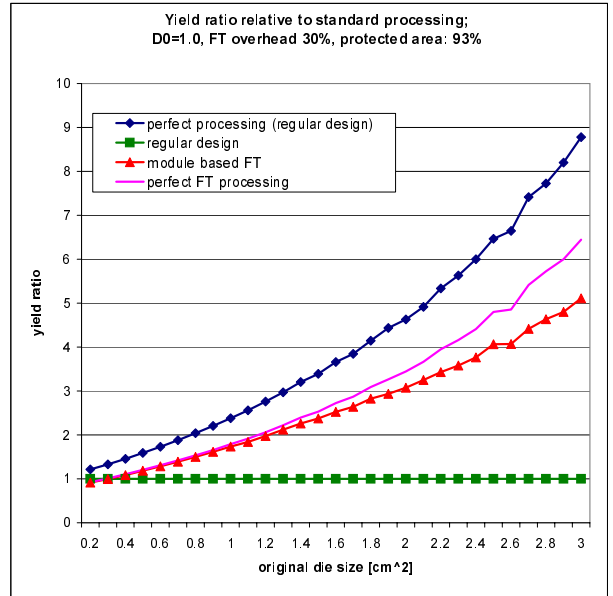


Figure 3. Yield ratio of perfect and normal processing of regular and fault tolerant design with respect to the yield of the original unprotected die. From 1.2 cm^2 on, design provides a yield which is at least a factor two higher than regular design.

yield increase when using the protection scheme described above. The 'perfect FT' line indicates the maximum FT yield possible, which is only limited by the number of (rectangular) dies per (circular) wafer. From this figure, it can be concluded that a protection of 93% of the IC area at an overhead of 30%, will increase the yield with more than a factor of two for IC's which are larger than 1.2 cm^2 ! Consequently, less wafers need to be handled to have the same production volume, which is a reduction in IC cost and frees foundry capacity for other IC productions. Also, older foundries with worse defect densities can still be economically profitable, while delaying the huge investments for new foundries and increasing the return on investment of the previous foundry. Something to look for in the current economic down-turn. Figure 3 is summarized in table 3. This table shows that a fault tolerant technique with the specifications mentioned above, is not worthwhile for an IC of 0.2 cm^2 . The fault tolerant overhead does not outweigh the increased die-yield: the total number of working dies per wafer is even 10% lower than the original yield. For IC's of twice that size, fault tolerant design leads to a production increase per wafer of 10%, which increases for the larger IC's.

4 Conclusions

Future IC technologies will require some form of fault tolerance to protect against Single Event Upsets (SEUs). This protection can be realised either in a temporal (e.g. re-transmission, result postponement [1]) or spatial (triple mode redundancy [6], circuit encoding [7]) way. It has been shown that even though fault tolerant design was aimed at fixing temporal errors (SEUs), the spatial techniques offer an additional advantage in terms of yield increase. Depending on processing conditions, fault tolerant overhead and die-size, spatial fault tolerant techniques could easily more than double the yield for realistic cases. Because these techniques are generally applied to modules on an IC rather than on a complete IC itself to handle the design complexity, the SEU protection is only affected in a limited way. Only the defective but 'repaired' modules are partly sensitive to SEUs, whereas the majority of the properly working modules are still fully SEU protected. Even with some unprotected modules, the overall IC failure rate may be well within specification. Precisely this characteristic makes fault tolerant design such a powerful way to increase the number of working dies per wafer. This could imply that the economic life-time of current clean-room facilities can be extended.

original die-size [cm^2]	nr. of reg. work. dies per wafer	nr. of work. FT dies per wafer	module based FT improvement factor
0.2	1212	1099	0.9
0.4	491	521	1.1
1.0	113	183	1.6
2.0	27	76	2.9
3.0	9	43	4.8

Table 1. Yield improvement of fault tolerant design for various chip sizes. Assumptions: $D_0 = 1.0$, FT overhead of 30%, 93% of the total IC area (including FT overhead) is protected against single errors per modules of $1 mm^2$. For a chip of $1 cm^2$, already 60% more yield is achieved, even though the FT design is 30% larger.

Future work will be directed towards the investigation of reaching a certain SEU protection limit (failure in time smaller than a certain threshold) at minimum cost, and extending this minimum scheme to increase yield in a cost effective way.

References

- [1] Eric Dupont, Michael Nicolaidis, Peter Rohr, "Embedded Robustness IPs for Transient-Error-Free IC's", *IEEE Design&Test of Computers*, May-June 2002.
- [2] Charles H. Stapper and Raymond J. Rosner, "Integrated Circuit Yield Management and Yield Analysis: Development and Implementation", *IEEE Transactions on Semiconductor Manufacturing*, vol 8, No. 2, pp. 95–102, May 1995.
- [3] D. Pradhan and J. Stiffler, "Error-correcting codes and self-checking circuits," *IEEE Computer*, pp. 27–36, Mar. 1980.
- [4] Niraj K. Jha and Jacob A. Abraham, "Totally self-checking MOS circuits under realistic physical failures" *Proc. International conference on Computer Design*, pp. 665–670, Oct 1984.
- [5] M. Nicolaidis, "Time redundancy based soft-error tolerance to rescue nanometer technologies," *Proc. VLSI Test Symposium*, pp. 86–94, 1999.
- [6] M. Alderighi *et al.*, "Achieving fault tolerance by shifted and rotated operands in TMR non-diverse ALUs," in *IEEE Int.Symp. on Defects and Fault Tolerance in VLSI systems*, (Fujiyoshida-city, Japan), Oct 2000.
- [7] G. Muurling, R. Kleihorst, N. Benschop, R. van der Vleuten, and J. Simonis, "Error correction for combinational logic circuits," in *21th Symp. on Information Theory (Benelux)*, (Wassenaar, the Netherlands), pp. 25–31, May 25-26 2000.
- [8] Richard Kleihorst, Nico Benschop, "Fault tolerant ICs by area-optimized error correcting codes," *Proceedings 7th IEEE International On-Line Testing Workshop (IOLTW01)*, Giardini Naxos, Italy, July 9-11, 2001.
- [9] John L. Hennessy, David A. Patterson, "Computer Architecture, a quantitative approach, chapter 1", 2nd edition, 1996.
- [10] Israel Koren, Zahava Koren, "Defect Tolerance in VLSI circuits: Techniques and yield analysis", *Proceedings of the IEEE*, vol 86, No. 9, Sept 1998.