The Design of Digital Machines Tolerant to Soft Errors

Yvon Savaria

Deparment of Electrical Engineering

McGill University, Montreal

July 1985

A thesis submitted to the Faculty of Graduate Studies and Research in partial fulfillment of the requirements for the degree of Doctor of Philosophy

This research was partly supported by a scholarship to the author by the National Sciences and Engineering Research Council of Canada

The Design of Digital Machines Tolerant to Soft Errors

Yvon Savaria

Deparment of Electrical Engineering

McGill University, Montreal

July 1985

A thesis submitted to the Faculty of Graduate Studies and Research in partial fulfillment of the requirements for the degree of Doctor of Philosophy

This research was partly supported by a scholarship to the author by the National Sciences and Engineering Research Council of Canada

Abstract

This thesis deals primarily with the problem of soft-error tolerance in digital machines. The possible sources of soft errors are reviewed. It is shown that the significance of ionizing radiation increases with the scaling down of MOS technologies. The characteristics of electromagnetic interference sources are also discussed. After presenting the conventional methods of dealing with soft errors, a new approach to this problem is suggested. The new approach, called Soft-Error Filtering (SEF), consists of filtering every output of the logic before latching it, in such a way that a transient injected into a machine does not change the final result of an operation. An analysis of the reduction in the error rate that is obtained by using SEF is presented. For example, this analysis demonstrates that the error rate due to alpha particles generated by the decay of radioactive elements becomes negligible. A great deal of attention is devoted to the design of filtering latches which is an essential component for implementing SEF machines. Three structures are considered and a CMOS implementation is proposed in each case. The double-filter latch is the best of the three implementations. It features a nearly optimum performance in the time domain and it is relatively insensitive to process fluctuations. An overhead analysis demonstrates that SEF usually results in a small overhead, both in area and in time simultaneously. In conclusion, SEF is the best approach to the problem of designing a machine tolerant to short transients.

Résumé

Cette thèse traite principalement du problème de la tolérance aux erreurs douces. Les sources d'erreur possibles sont passées en revue. Il est démontré ici que l'importance des radiations ionisantes augmente avec la réduction d'échelle des procédés de fabrication. Les caractéristiques de l'interférence électromagnétique comme source d'erreur sont discutées. Après la présentation des approches conventionnelles au problème des erreurs douces, une nouvelle approche est suggérée. Cette nouvelle approche appelée "Soft-Error Filtering" (SEF) consiste à filtrer toutes les sorties de la logique combinatoires avant de les mémoriser. Ceci fait qu'une transitoire courte injectée dans la machine ne peut pas changer le résultat final d'une opération. Une analyse du taux d'erreur résiduel pour une machine SEF est présentée. Par exemple, cette analyse démontre que le taux d'erreur résiduel est négligeable, pour une machine SEF affectée par la radioactivité naturelle. Une attention toute particulière est apportée à la conception de latch-filtres, qui sont des composants essentiels pour réaliser une machine SEF. Trois structures sont considérées et une réalisation CMOS est suggérée dans chaque cas. Le latch à deux filtres est le meilleur des trois. Sa performance temporelle est quasi optimale et la réalisation proposée est relativement insensible aux variations du procédé de fabrication. Une analyse du coût en temps et en matériel associé à SEF démontre que ce coût peut être faible selon les deux aspects simultanément. En conclusion, SEF est la meilleure approche pour fabriquer une machine tolérante aux erreurs douces, si ces erreurs sont causées par des transitoires courtes.

Acknowledgements

I want to thank the members of my advisory committee, Professors Jeremiah Hayes, Nicholas Rumin, and Vinod Agarwal for the support they gave me in this work. They gave me enough freedom for exploring a completely new idea, but at the same time the encouragements and tight feedback that were an invaluable help in completing this work.

Since it is too easy to forget the exact contribution of the members of a team, I will highlight some of those that are most important. In parallel with this thesis, six papers were published or have been accepted for publication. These papers were written in close collaboration with the members of my committee. Very often, they significantly improved my original texts by rewriting parts of them. This author assumes the complete responsibility for any imperfection in this thesis, but since most of the modifications suggested in the papers were subsequently embedded in the thesis, their contribution to this document is very significant. The contribution of Professor Rumin has been particularly important. He has invested a large amount of time in proofreading my derivations and suggesting better formulations. It should also be remembered that the expression Soft-Error Filtering is the fruit of one of our numerous group meetings. This expression is used in reference to the new fault-tolerance approach proposed in this thesis. An important turning point in this work was the idea of

assuming a bound on the transient duration resulting from a disturbing event. If this assumption is not made, the formalism becomes extremely difficult (a couple of months were lost there). This idea was contributed by Professor Agarwal.

A last, but not least, contribution that must be acknoledged is the one of professor Robert Dufresne of "Ecole Polytechnique de Montréal" and his wife Evelyn. There could have been no mention of cosmic radiation in this thesis if professor Dufresne had not convinced me of its significance. Cosmic radiation is largely ignored in the current literature on soft errors. Moreover, they carefully reviewed a draft of this document and suggested numerous corrections.

Preface

This thesis is, to the best of the author's knowledge, the first work specifically dedicated to soft error tolerance in logic circuits. Moreover, it is an interdisciplinary work touching on many research fields, including: the interaction of radiation with matter, the electromagnetic compatibility of electronic circuits, the theory of reliable communication systems, the design of integrated circuits and systems, and finally logic design for fault-tolerance. Therefore, in order to appreciate this work, one must not consider only one of its facets. This is the reason why the work was closely supervised by three Professors, which is fairly unusual. Professor Rumin was most capable of appreciating the optimization of the filtering latch at the transistor level. Professor Hayes, because of his background in communication theory, could review the analogy to noisy communication systems and the derivations of bounds on the error rate. The originality of Soft-Error Filtering as a new fault-tolerance technique was most appreciated by Professor Agarwal. Finally, even though Professor Dufresne is not one of my supervisors, his background in the study of cosmic radiation enabled him to review the analysis of the effects of showers of particles.

In the rest of this preface, the aspects of the thesis which are considered to be original are enumerated. The content of Chapter 2 is largely based on a review of the literature, however it is original in the sense that it unifies into one document, information from diverse sources which are scattered in the literature. Chapter 2 also contains some original work. The argument developed for ruling out electrical noise as a significant source of

soft error in static logic is new. Also, the bounds on the error rate due to the products of radioactive decay are extensions of what can be found in the literature [MAY78,SAI82]. Finally, the discussion of how the duration of a transient pulse changes with propagation are new.

Chapter 3 is a review of the conventional techniques for dealing with soft errors, however it contains some original ideas. The discussion on intrinsic tolerance to soft errors is original. Moreover, in reviewing the conventional fault-tolerance techniques, it became clear that there is a significant advantage to adapt them for soft-error-tolerance. Therefore the architectures proposed in Figs. 3.1 and 3.4 are enhancements to what can be found in the literature.

The main original contribution of this thesis is the Soft-Error Filtering approach to the problem of soft-error-tolerance. This approach is proposed in Chapter 4. The chapter includes an analysis of the reduction in error rate possible with SEF. Chapter 5 is devoted to the practical aspects of implementing SEF machines. In particular, three approches to the problem of designing filtering latches are analysed. Finally, an analysis of the overhead associated with SEF is presented in Chapter 6. This analysis demonstrates that SEF can yield an overhead significantly lower than conventional alternatives.

Table of Contents

List of Figures and Tables

Chapter 1
Introduction

Reliability may become the major obstacle in the commercialization of very large scale integrated (VLSI) circuits, when feature sizes are scaled down to submicron dimensions to satisfy the demand for increased circuit density. The increasing complexity of the systems being built creates a demand for components with a higher reliability. It forces the manufacturers to improve the quality of the integrated circuits, which results in a smaller failure rate. However, at the same time, the *soft error* rate is expected to increase, for reasons discussed in Chapter 2. Therefore the relative importance of soft errors will increase, and they should receive more attention in the future. Moreover, a low cost technique for improving the reliability of digital machines against permanent faults, such as off-line testing, does not work for soft errors. On-line tolerance is required. This thesis is devoted to soft errors and to efficient means of tolerating them.

In this thesis, a soft error is defined as a *temporary* and *non-recurrent* difference between the expected and actual behavior of a machine. In a *strict sense*, a fault is considered to be recurrent, if it is present in a given *element*, and can be correlated from cycle to cycle. An element here extends from a simple wire, the smallest logic block considered, to a complete VLSI chip. An error is called *reproducible*, and is *not* considered to be recurrent, if it occurs when a state transition of the system involves two or more separate elements. Clearly, as the fraction of the system included in a single element increases, more

reproducible errors become recurrent errors. In practice, a reasonable requirement would be that all elements used for building a system be free from recurrent error, i.e. that they work according to their specifications. However, such a system could still exhibit reproducible errors. The distinction may seem artificial but it is important, because reproducible transient faults will cause transient errors, which could be tolerated in a different way from reproducible static faults (i.e. logic errors) and recurrent faults (intermittent contacts or elements with faulty behavior).

In contrast with our definition of a soft error, it has been suggested [McC79] that, since a fault is repairable if it is testable, the testability of an error should serve as a criterion for distinguishing between intermittent and transient errors. Certainly, this is a reasonable approach for separating one from the other, but it neglects the important situation where some faults, while being testable, are not in practice repairable, since every element already functions according to its own specification. One could argue that they are repairable in a broad sense, if redesign of the system is allowed. For example, high performance systems exist where the number of drivers that are allowed to switch simultaneously is limited. If the limit is exceeded, transients larger than the noise margin are injected into the supply line. Therefore the rule may be violated when many chips are used in a system, and certain state transition of the system can result in an error. The soft-error class, as defined here, contains reproducible and thus testable errors. They are repairable in the broad sense, but usually they are never repaired.

This thesis explores the problem of soft error tolerance. The first important step is to characterize the sources of soft errors. This is covered

in Chapter 2, which is largely based on a review of the existing literature. It is demonstrated in the same chapter that electrical noise should never be significant. In contrast, an estimation of the error rate due to ionizing radiation demonstrates that it can become a major source of soft errors, with the scaling down of MOS technologies. This chapter also includes a discussion of the characteristics of electromagnetic interference, which is the most important source of soft errors in logic circuit.

Conventional methods of dealing with the soft-error problem are reviewed in Chapter 3. The knowledge required to produce state of the art VLSI circuits can be separated into a number of *levels*. These levels tend to be disjoint. Accordingly, the conventional approaches to the soft error problem are divided into two categories: some deal with the problem at the physical level only, others deal with the problem at the system level only. This chapter also suggests straightforward extensions of the system level solutions, in order to make them more appropriate for tolerating the potential causes of soft errors.

The main contribution of this thesis is a new technique for tolerating soft errors, called Soft-Error Filtering (SEF). The SEF approach is introduced in Chapter 4. This chapter presents an analogy between a noisy communication channel and a logic machine sensitive to soft errors. It also includes an analysis of the error rate for a SEF machine, which demonstrates how SEF can reduce the error rate to negligible levels.

As the name implies, Soft-Error Filtering is based on the assumption that the transients which cause soft errors can be filtered efficiently. It is demonstrated in Chapter 4 that such filtering should be done prior to

latching the result of any operation. Therefore, Chapter 5 is devoted to the important problem of designing latches capable of efficiently filtering transients in their input. Three different approaches to this problem are considered, and an efficient implementation is given in each case.

The importance of the SEF approach, in the design of soft-error-tolerant machines, follows from the fact that it can often be implemented at a lower cost than the conventional techniques [SIE82]. This claim is supported by an overhead analysis in Chapter 6. However, SEF can result in a high overhead, and the analysis also outlines the limitations of this new approach. Finally Chapter 7 suggests directions for further research, and includes the conclusion of this thesis.

Chapter 2
Characterization of Soft Error Sources

2.1 Introduction

Soft errors can result from several physical mechanisms, which fall into three categories: electromagnetic interference, electrical noise and ionizing radiation. Each of these sources is reviewed separately in Section 2.2. In this thesis, a clear distinction is made between electromagnetic interference and random noise, and for brevity the terms interference and noise will be used respectively.

A soft error is the observable consequence, on one or more output lines, of a transient injected into an internal node of a digital machine. The relevant characteristics of the transients produced by each class of physical mechanism will be described. This description is based mainly on a review of the literature. The discussion is limited to MOS VLSI systems and includes the foreseeable effects of scaling. There exists a well-established theory of scaling [DEN74,HOD83,MEA80,TOY79] and extensive experience in fabricating scaled devices [DEN79,JEC79,LIU82]. Since MOS is likely to be the dominant VLSI technology, the analysis will not be extended to the various silicon bipolar [HOD83] or gallium arsenide [MOR84] logic families. However the same problems exist in these technologies [ROO84].

A transient injected in a digital machine does not necessarily result

in an error. The first condition necessary for an injected transient to lead to an error is that it must propagate inside the combinational logic network. Assuming a synchronous system, a second condition is that the propagating transient must reach the input of a latch, during a time interval overlapping with its sensitive time slot. Consequently the transformations of the characteristics of a transient with propagation are very important, and will be analyzed in Section 2.3. A transient injected into an internal node of a latch could also result in an error if it has sufficient energy.

A discussion of the significance of soft errors and reliability trends is included in Section 2.4 and 2.5

2.2 Soft Error Sources

In this section, sources of soft errors will be presented as partitioned within three classes. It will be shown in 2.2.1 that noise will never be significant. Ionizing radiation and interference are analyzed in 2.2.2 and 2.2.3 respectively.

2.2.1 Electrical Noise

Electrical noise is the first potential source of soft error analyzed, because it is relatively easy to show that it will not be significant in the future and does not need to be considered further. The noise could cause errors in digital systems, if the energy representing a logic value is decreased to very low levels. However, fundamental limits

[MEA80] for fabricating transistors with smaller dimensions will be reached, before the voltages and currents approach values so small that electrical noise can cause errors. Therefore, as will be shown in the following paragraphs, electrical noise is a second-order consideration.

In this thesis, scaling refers to reducing the dimensions of the transistors in order to achieve a similar function but with a better performance and a lower cost. A first important observation is to note that scaling decreases rise and fall times. Since the noise equivalent bandwidth is inversely proportional to the rise and fall times, it increases the noise power. The highest error rate should be observed for devices with the smallest signal to noise ratio.

An important parameter that determines the sensitivity of a technology to soft errors is the minimum gate capacitance of a transistor, C_g. The practical limit for scaled supply voltage and gate capacitance are around V_{dd}=.5V and $C_g = 10^{-15}$ F [MEA80].

Reducing the operating voltage to a value as small as .5V is possible with CMOS, with a proper design of the process. How to choose the various parameters such as the threshold of the transistors, the doping levels, and so on, is beyond the scope of this work. After a sufficient security margin has been provided for the various possible fluctuations of parameters, one has to allow a sufficient noise margin for the various sources of interference that usually exist [MAR84]. Moreover a sufficient fraction of the noise margin must also be reserved for electrical noise.

If it can be established that only a small fraction of the supply

voltage is necessary to guarantee that the effect of electrical noise is negligible, then indeed electrical noise can be neglected. It is generally recognized that a dynamic design is not feasible with a low supply voltage such as that being considered here; consequently the machine is assumed to be static. Also, a noise event cannot cause an error if it affects a node before it is stabilized for a given clock period. This means that during the time slot of interest, all the nodes are stabilized and therefore tied to one of the supply buses through a transistor in the triode region. In this region of operation, the MOSFET transistor can be treated as an ohmic resistance, whose value is related to the transistor's transconductance [AMB82 p.196].

The rms noise voltage, E, of an ohmic resistance shunted by a capacitor, C, is given by $E=(kT/C)^{1/2}$ [MOT73 p.24], where k is the Boltzmann's constant and T is the absolute temperature. This yields E=1.4 mV rms for the more sensitive nodes, with $C=C_g$. This noise amplitude must be related to the noise margin that is necessary to make the error probability negligible. The error probability is given by the tail of a Gaussian distribution. It is easy to show, using the Chernoff bound, that the probability of a noise event with an amplitude larger than some multiple of E, N*E, is bounded by $e^{-(N*N/2)}$. A value of N=15 yields an error probability smaller than 10^{-49} per cycle-node, which is a sufficiently small value to justify neglecting this contribution to the machine error rate. With a margin of N*E= 21 mV reserved for tolerating electrical noise, its contribution can be neglected. Notice that only a few millivolts in the noise margin result in a very significant difference in the error rate. Since the supply voltage will be larger than 500 mV, increasing it by the few mV that are necessary to make the error rate negligible has an insignificant impact on the process. Therefore, noise will never be significant.

What limits the lowest voltage for which a good digital device can be designed is the so-called volt-equivalent temperature kT/q, where q is the charge of an electron. It is noteworthy that the significance of the noise would grow at low temperatures for a device with a very low supply voltage. This occurs because the noise amplitude only scales as the square root of the temperature, whereas the volt-equivalent temperature predicts a possible linear reduction of the operating voltage.

The discussion assumes that supply voltages are scaled with the dimensions of the devices. However, there are a number of good reasons why scaling at constant voltage is preferred over scaling at constant field. These include: the compatibility with existing logic families, the difficulty of controlling reliably the thresholds on a large wafer when they have to be on the order of a fraction of a volt, plus the faster switching of the high-voltage devices. Obviously, if voltages are not scaled, electrical noise is indeed an insignificant effect.

2.2.2 Ionizing Radiation

In this Section, it will be shown that ionizing radiation determines the reliability of scaled devices. Sensitivity to radiation is obviously determined by how the scaling is done, but it is argued later that at some point it will become impractical to make the MOS technology immune to the effects of ionizing radiation. The initial recognition of the significance of ionizing radiations is due to May and Wood [MAY78]. They established a direct relationship between the soft-error rate of dynamic memories and the flux of

alpha particles generated by the packaging materials.

Effects of Ionizing Radiation

In order to understand the tradeoffs involved in the sensitivity of silicon devices to ionizing radiation, the basic properties of charge injection and interaction with matter are reviewed here. For simplicity, a node in a circuit will initially be represented by its parasitic capacitance only. This will permit the development of a model where the effects of a particle injecting a charge can be represented as a current pulse with certain characteristics. It is an elementary fact in integrated circuits that in order to put more than one device on a substrate, they must be isolated. This isolation is almost always achieved by creating a PN junction which is unbiased (0 volts) or reverse-biased between the device and the substrate.

A charged particle interacts with matter by ionizing a number of atoms, leaving free electron-hole pairs along the propagation path [LAP72,WOL63]. The ionization energy in silicon is 3.6 eV per electron-hole pair [MAY78]. The amount of charge injected into an infinitesimal length of the particle trajectory, is proportional to the square of the charge on the particle for a given velocity. For example, a proton is expected to have an ionizing capacity 4 times smaller than an alpha particle. Notice that a proton is also 4 times lighter than an alpha particle and must, therefore, move at twice the velocity to carry the same energy, assuming non-relativistic velocities. So a proton and an alpha particle with the same velocity have roughly the same range, but the former generates four times less charge.

The *range*, or the depth of penetration, is determined by the particle energy. A very important factor is the *energy shedding rate*, which increases rapidly as the particle slows down due to the longer interaction time with each atom. This increase of the shedding rate translates into very intense ionization near the end of the range. Problems occur when the range is similar to the junction depth.

A consideration which completes the picture is the dynamic behavior of the collecting structure. The experimental work supporting the model used here has been performed by Hsieh et al. [HSI81]. Their model separates the actual injected current into two fractions, the *drift* and the *diffusion* components. This relates to the two basic phenomena for charge transport in semiconductor devices. The drift component is due to the electric field of the isolating junctions. This mechanism is relatively fast, occuring in less than 0.5 ns. The diffusion component, which is caused by a gradient of the carrier concentration, is a relatively slow phenomenon that can carry charges over relatively large distances but on a microsecond time scale. The diffusion component is thus capable of causing correlated errors on dynamic nodes far from the hit point, but is not significant for *static gates* which are logic functions where the output does not depend on charge storage at any moment.

The angle of incidence is another important factor in determining the sensitivity of a particular structure. It has an impact on the amount of *charge sharing*, which is the injection of the same total charge but on more than one node in a region. When a hit occurs at a high angle of incidence (less than 45 degrees from the normal), charge sharing is

negligible unless there is more than one collecting junction within a radius of 2 microns from the propagation track. The majority of the hits are in this category when the dominant source is the packaging material. This figure of 2 microns is subject to some controversy, because a node as far as 10 microns from the hit point collects charge in the range of 1 femto Coulomb [SAI82]. Nevertheless for a hit at normal incidence, the region where charge injection is intense is roughly delimited by a circle of 2 microns radius. Sai-Halasz et al. [SAI82] analyzed the problem in a dynamic RAM context, without differentiating between diffusion and drift. Presumably the small charge collected at distances larger than 2 microns is collected by diffusion over many nanoseconds and cannot upset a static gate.

The partition of the total charge injected into drift and diffusion components is not obvious. It is a function of the doping levels and the geometry of the structure. This problem is beyond the scope of the present work but an analysis is possible and has been performed for some simple structures by Hsieh et al. [HSI81]. The conclusion of their analysis, which was confirmed by experimentation, is that the maximum charge collected by the drift mechanism is larger than 60 fC, if the substrate's resistivity is larger than 2 ohm-cm. A typical value for the minimum gate capacitance C_g of a 5 microns process is 10^{-14}F, which implies pulses as large as 3 Volts, for a maximum injected charge of 60 fC on a node of 2 C_g. The maximum collection occurs for a particle energy in the range of 3 MeV. Notice that particles with a higher energy are possible, but result in a smaller drift charge, because the intense ionization region occurs too deep in the substrate. This range of substrate resistivity is consistent with a one micron technology [DEN79,LIU82].

Obviously, for submicron technology charge sharing is significant. If a hit is modelled by a current source, then the load is composed of the capacitance of more than one node. Notice that the logic value of a node has a strong influence on its collection efficiency. Obviously if the node's potential is the same as the substrate's potential, then injecting free carriers in the isolating junction will not cause a current to flow. This makes the problem difficult to analyze in a charge sharing situation, since the charge not collected by one node will be collected by its neighbor, and the sensitivity of an internal node of a logic function is dependent on the input value.

The paper by Sai-Halasz discussed earlier [SAI82] addresses the serious question of hits at low angle of incidence. One may attempt to reduce the flux of particles by coating the chip with a material of very little radioactivity, with the hope that this layer will absorb all the particles that would otherwise hit the surface. There will be a remaining flux for which the majority of the particles originate in the top layers of the chip itself. They also contain significant traces of radioactive impurity and, therefore, will often produce hits at a low angle of incidence. With a range as long as 60 microns in silicon, a large number of nodes could be affected by a single hit. An alpha particle loses more than 2.5 MeV in the last 10 microns of its range, which corresponds to 110 fC of injected charge. At very low angles of incidence, on a dense device, the majority of this charge would be collected by the drift mechanism. Assuming a 1 micron pitch, the charge could be anywhere between 15 fC on 7 adjacent nodes and 110 fC on a single wire (see Fig. 2.1), depending on the angle between the layout orientation and the track direction.

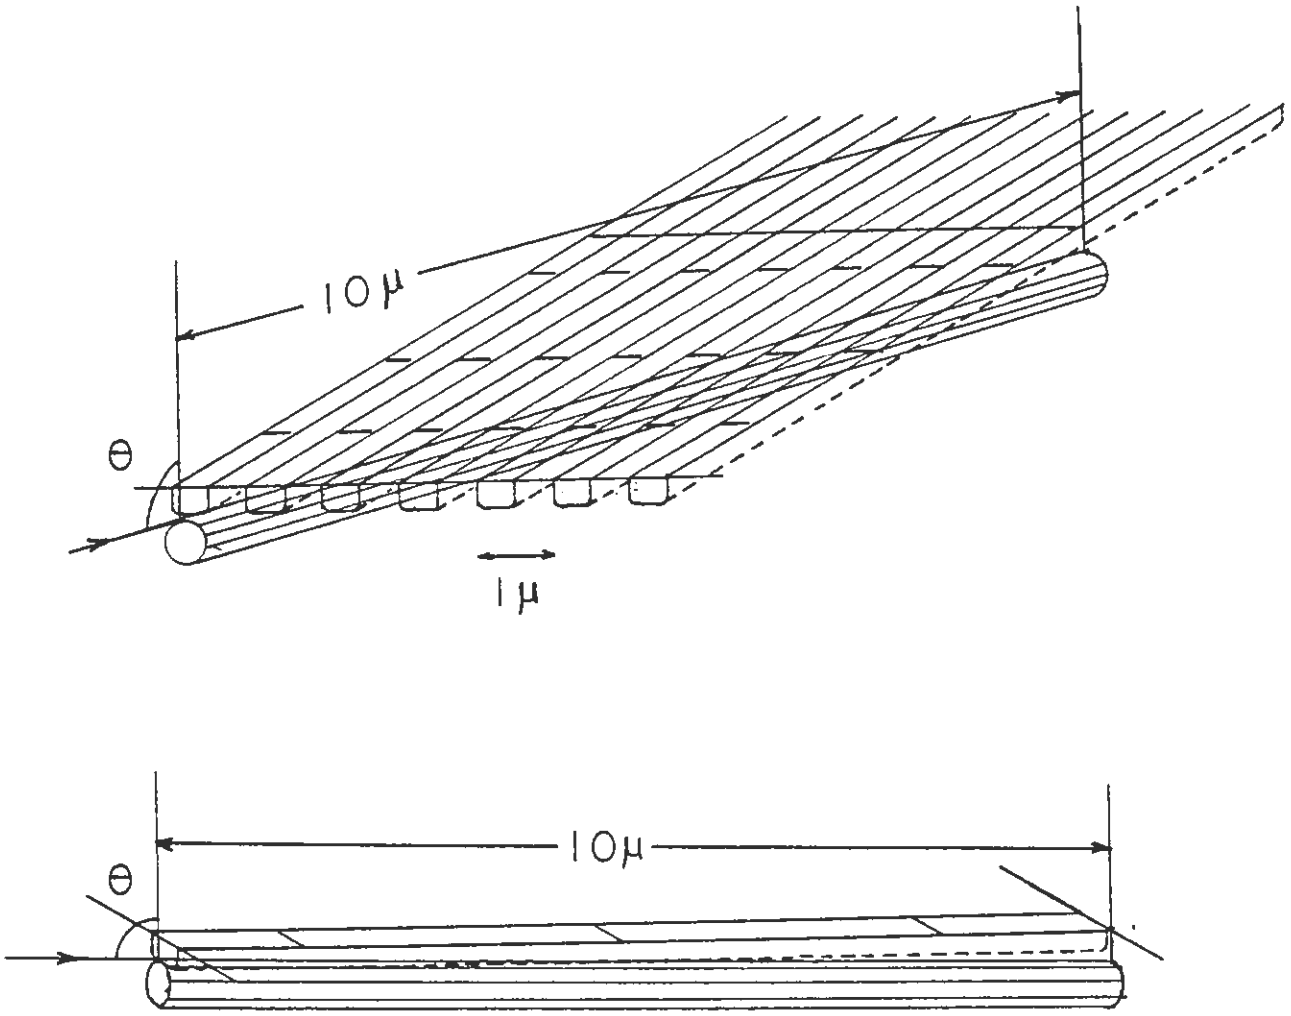

Figure 2.1 Hits at a low angle of incidence ($\theta \approx 90^{\circ}$) on diffusion lines with a 1μm pitch. If the total charge injected in the last 10 μm is 110 fC it may inject (a) 15 fC per node into 7 adjacent nodes, (b) or the total charge into a single wire.

The discussion of the effects of low incidence hits, shows that a large fraction of the total area of a VLSI chip is to some extent sensitive. An accurate quantitative estimate of this sensitivity is only possible by a detailed simulation. If the chip is not coated, the events described are still possible but occur with a relatively low probability, and they are not included in our analysis of the error rate since they are only second order effects. Nevertheless, for estimating the probability of multiple simultaneous events, which have a serious impact on the effectiveness of tolerance schemes, the contribution of events at a low angle of incidence is very important.

A convenient layout-independent unit for comparing the sensitivity of technologies is the Volt-C_g. Thus the preceeding example in this Section with $C_g=10^{-14}F$ yields a sensitivity of 6 V-C_g. The possible amplitude of a pulse on a given node can then be estimated from the layout, by dividing the sensitivity by the capacitance of this node. It should be clear that an injected pulse cannot have an amplitude larger than the supply voltage, since a difference of potential is required in order to collect a charge. Consequently if the result of calculating the amplitude of the transient is larger than the supply voltage, it simply means that this node would be shorted to the substrate (or well) potential for a time sufficient to either: diffuse the charge in the substrate (may be as long as a microsecond), or compensate it with a pull-up (in NMOS).

It should be clear from this discussion that, at the injection point, the transients due to ionizing particles are unipolar in NMOS: a high level can be driven to a low level but not vice-versa. A similar argument in CMOS shows that for each sensitive node there is a state which is not affected by

an ionizing radiation hit.

A technology scaled to the limit with $C_g = 10^{-15}F$ yields a transient amplitude of 60 V-C_g. There is a significant uncertainty in this limiting value for C_g and the suggested figure corresponds to a scaling down of 10 from a typical 5μm technology. If the scaling could be pushed to 0.25μm or lower by the use of novel approaches to overcome the fundamental limitations of small geometry devices, and considering the difference between the constant field scaling theory and what is done in practice [JEC79], a further reduction of C_g by a factor of 5 is possible. The corresponding sensitivity value of 300 V-C_g is too speculative to serve as a basis for this work. It can now be argued that the scaled technology will become sensitive at some point, even if one manages to scale at constant voltage. Moreover, ionizing radiation is not the only disturbance that may be encountered and, if more than half of the noise margin must be reserved for it, the design of a reliable system will be difficult [MAR84].

Injected Voltage Transient Characteristics

A model was developed in the preceding Section to represent the effects of physical interaction between an ionizing particle and an integrated circuit. This model can now be used to pursue the analysis with a circuit-level representation of the device. In particular the properties of the transient and the sensitivity of static nodes will be analyzed, since they exibit a much better intrinsic tolerance than dynamic nodes. In fact, the first step in designing a machine capable of tolerating soft errors is to restrict the designer to static logic. The better tolerance of static logic follows from its capacity to recover from injected transients, which results from the existence of a low impedance path to one of the supply voltages. This low impedance path also makes each node harder to drive into the wrong state. For reliable machines capable of tolerating soft errors, dynamic design would be confined to pure memory structures, where the problem can be considered to be solved with the use of the proper error-correcting codes [SAR84].

In a static circuit, injecting a sufficient charge will not cause an error, unless it is done over a short enough period to override the pull-up or pull-down device. The interaction time between a particle and a junction is between 0.1 ns and 0.5 ns [HSI81]. The effect of the particle can be approximated by a current pulse having a peak amplitude which is a function of the substrate resistivity, and is larger than 0.25 ma [HSI81] (see Fig. 2.2(a)). Another condition necessary for errors to result from ionizing radiation is that the injected transient must propagate, or in other words, the technology must be fast enough. There are two cases of interest depending on whether the supply voltage is scaled or not.

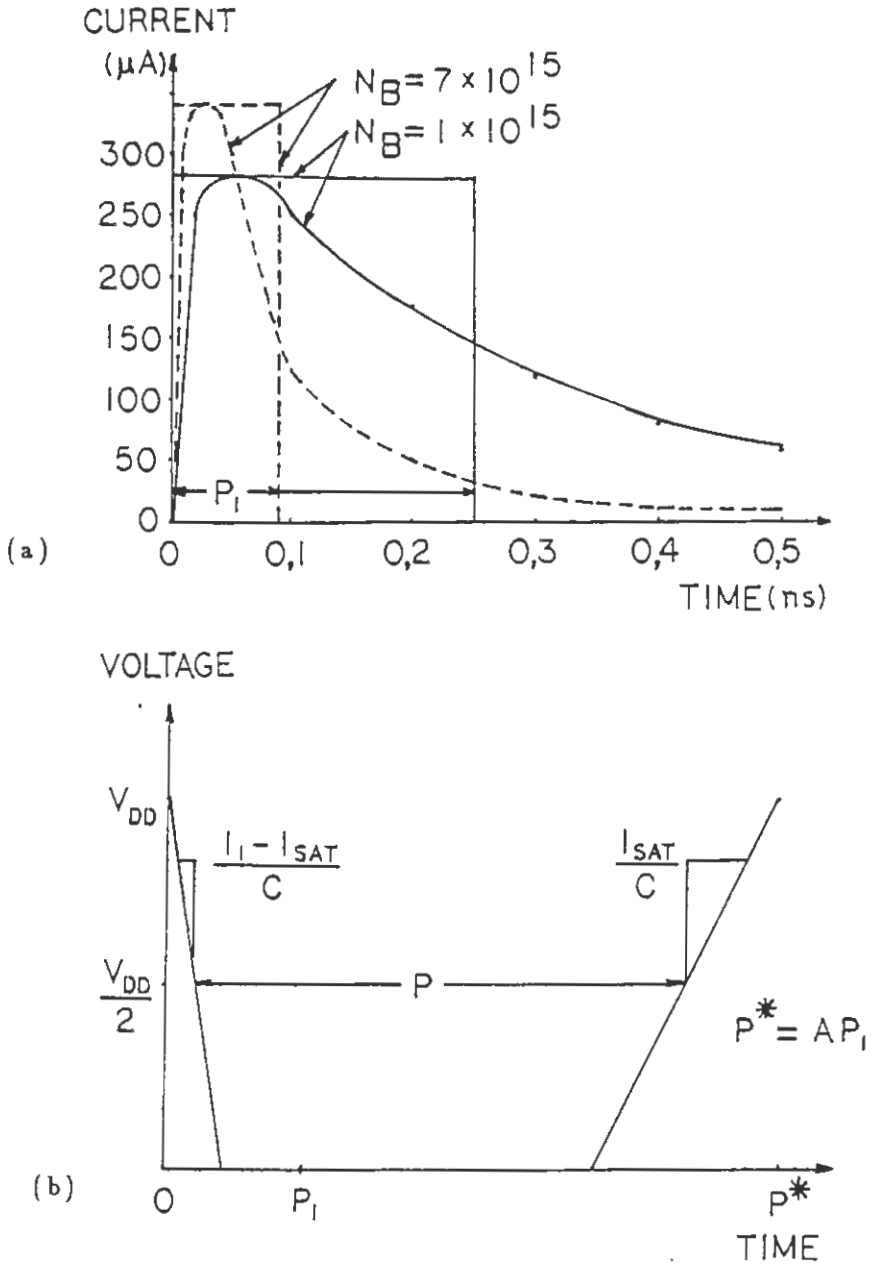

Figure 2.2 (a) Calculated injected current pulses (reproduced from [HSI81]), and a rectangular approximation, for a 4.8 MeV alpha particle at normal incidence with a polarization of 8V (Injected charges of 70 and 30 fC). N_B is the substrate doping.
(b) A first order approximation of the resulting voltage transient.

If the supply voltage is scaled, the maximum saturation current of the transistors, I_{sat}, also decreases with scaling. Since the peak current injected by the particle, I_i, does not decrease until the device reaches feature size where charge sharing is significant, it leads to $I_i=AI_{sat}$, where the parameter A can become much greater than 1. It is assumed that the injected charge is sufficient to drive a node from 5 to 0 volts, the substrate potential of a N device. The injected current is approximated by a rectangular pulse of amplitude equal to the peak value of the real current pulse, and having a duration P_1 that yields the same total charge as shown in Fig. 2.2(a). When the sensitive node reaches the substrate potential, the charge is assumed to stay available at the junction until the pull-up can compensate for it. The justification of this assumption is that the diffusion mechanism that could carry the free carriers far away is a relatively slow phenomenon and should not be significant on a time scale of a few nanoseconds.

The current from the pull-up is approximated by its maximum saturation current I_{sat} as soon as the event begins. The corresponding shape of the resulting voltage pulse is illustrated in Fig. 2.2(b). It is possible to estimate the duration of the transient pulse, P, resulting from the initial event of duration P_1. The total transient duration, P_1A, corresponds to the time necessary for the pull-up to compensate for the injected charge, assuming that the pull-up current is equal to the maximum saturation current for the whole interval. The duration of P is obtained by subtracting the time where the transient amplitude is less than $V_{dd}/2$.

$$P = P_1 A - \left(\frac{\text{rise-time}}{2} \quad \frac{A}{A-1}\right) \qquad (2.1)$$

This estimate is reasonably accurate if A is larger than 5. The point here is not to estimate P accurately but to demonstrate that the injection time and the transient duration may be significantly different. Even though the injection time is roughly limited to .2 ns, the voltage transient observed can be longer than 1 ns. This calculation corresponds fairly well to the NMOS 1 micron technology proposed by Dennard et al. [DEN79], where the pull-up of a minimum size inverter has a saturation current of 50 μA and the propagation delay of an unloaded gate with a 2.5V supply is 230 ps.

The second situation corresponds to scaling without decreasing the supply voltage, which results in a transistor saturation current that increases with scaling. This approach to scaling can lead to a technology where the pull-up saturation current is larger than or the same as the maximum injected current. Notice that it does not necessarily mean that the technology is immune to ionizing radiation, even though that is possible. It is important to remember that transistors are far from ideal current sources and, in fact, only a fraction of the maximum saturation current actually 'combats' the injected current, when the amplitude of the transient is small. At the same time, device switching times are also decreased by scaling, which makes it possible for a pulse with a duration as short as 200 ps to propagate. For example, a technology like HMOSIII [LIU82] would be marginally sensitive on its minimum-size logic gates, and the equivalent CMOS technology, CHMOSIII [SER84], would be tolerant with P devices having a length-to-width ratio L/W=1/2 that yields a 330 μA saturation current. Notice that it is not likely that this constant supply voltage approach to scaling will be possible for submicron devices, since, for example, the punchthrough

voltage of 1.5 μm channel-length devices is only 10 V [SER84].

According to the simple model presented, one may be tempted to say that because of the lower impedance of "on" transistors, CMOS will always have a better tolerance to ionizing radiation than NMOS. Unfortunately the exact effects of charge injection are much more complex in CMOS. This is due to the existence in CMOS circuits of the parasitic bipolar structures shown in Fig. 2.3. For example, if the end of the track of a particle is just inside what happens to be the base of a parasitic NPN or PNP device, the charge may be amplified in a fashion similar to the action of a saturated transistor after the base current has been decreased to 0. An alpha particle loses .4 MeV in the last 2 microns of its range [LAP72], which corresponds to 17 fC, so even with a modest amplification of 10, the injected current becomes much larger than expected. Unfortunately, with the decreasing well depth of high performance CMOS, the gain of these parasitic devices tends to be large, since the current gain of a bipolar transistor increases as W^{-2}, where W is the base width [GIB66,p.342]. Moreover, to make matters even more complex, the high ionization density produced by the alpha particle is known to change the shape of the depletion region, and thus the effective base width, in a dynamic fashion [HSI81].

Even more troublesome in CMOS is the parasitic PNPN structure that introduces the possibility of radiation-induced latch-up. This presents a serious reliability problem. If one supposes that a region of the chip is operated near the point where it will latch, the injected charge could behave as a triggering current if it occurs in the right spot. These considerations are certainly strong motivations to develop a silicon-on-insulator technology, where the parasitic structures do not exist [DAV83]. Analyzing

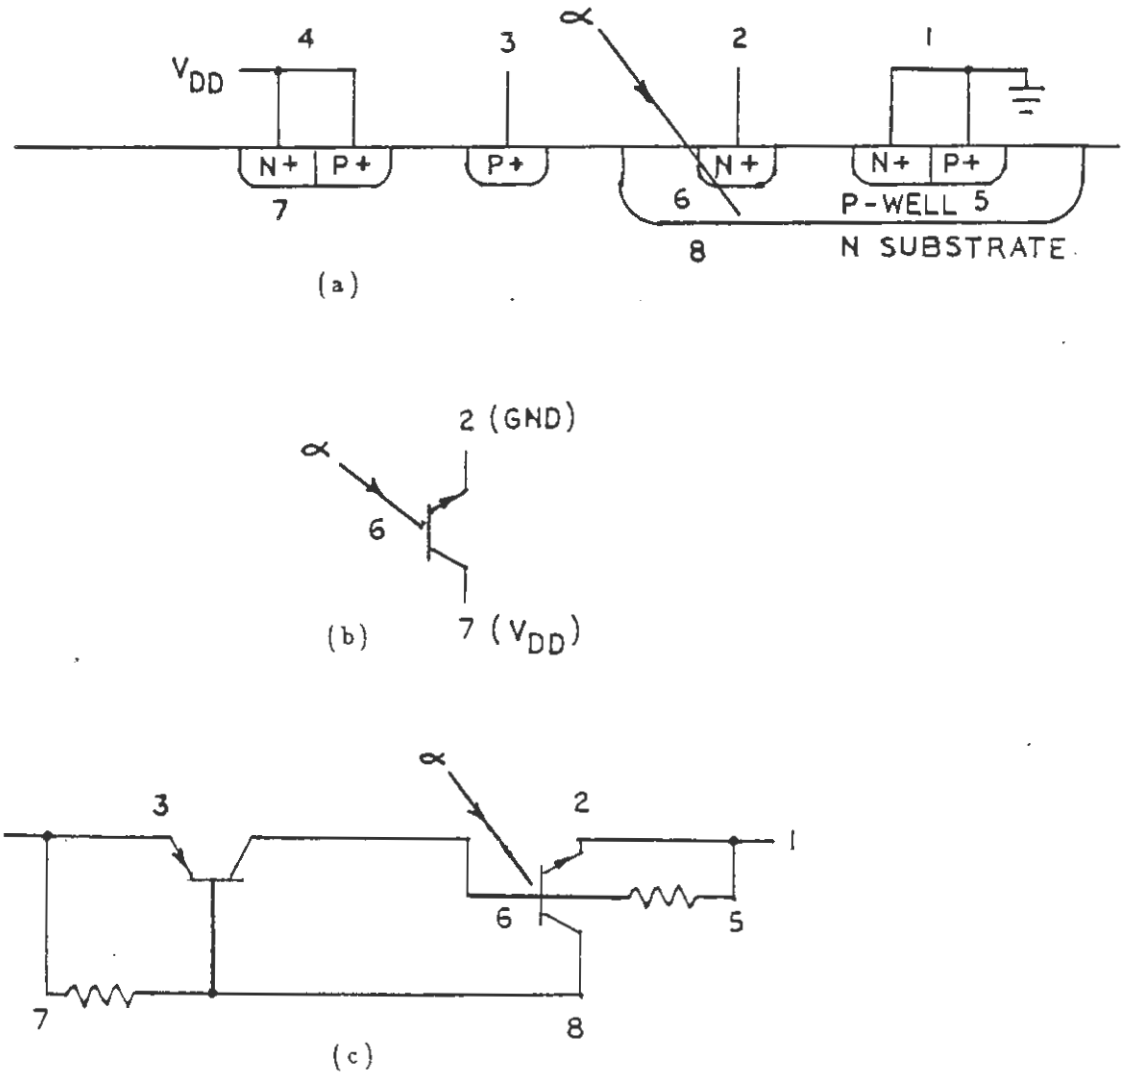

Figure 2.3 Parasitic bipolar structures in CMOS.
(a) A p-well device
(b) A parasitic NPN transitor
(c) A parasitic thyristor.

the effects of intense ionization on multilayer structures is a complex problem and is outside the scope of this work.

In conclusion, a particular technology can be immune to the effects of ionizing radiation, but this is not true in general for submicron devices. Obtaining intrinsic immunity to alpha particles does not come free, and it is certainly worthwhile to analyze alternatives to solutions in which the process or the design rules are modified, in such a way that each gate is individually immune.

Sources of Ionizing Radiation

There are two distinct sources of ionizing radiation. The most often considered and most important for a system operated at sea level is alpha particles from the decay of trace levels of Uranium and Thorium in packaging materials [MEI79]. This decay also produces a significant flux of beta particles, but because the mass of an electron is much lower than that of an alpha particle, a beta of sufficient energy to ionize a significant number of atoms has a harmless range (over 1000 μm) and, for a range similar to the feature size of an integrated circuit, the energy is less than .05 MeV [LAP72 p.271] (.015 MeV in the last 2 μm). The beta particles would start to be significant for a technology scaled to the limit at constant field. Notice that an alpha particle from natural decay can only come from a thin layer inside the package, since the energy spectrum is limited to 9 MeV, which corresponds to a range of less than 60 μm in silicon, as mentioned earlier.

The second important source of ionizing radiation at sea level is the

flux of secondary particles produced by high energy cosmic rays. It is composed mainly of mesons and electrons [WOL63]. Since mesons are singly-charged particles, and their mass is approximately 25 times smaller than that of an alpha particle, the injected charge should be one order of magnitude smaller than for alpha particles. This means that mesons should not be neglected for a technology scaled at constant field.

The flux of alphas from packaging materials is expected to be approximately 0.1 part./cm^2-hr [MEI79]. By comparison, the basic flux of mesons is much larger, with 80 part./cm^2-hr considering only the soft component [WOL63]. The hard component can be neglected because the interaction with matter is not intense for high energy particles. Moreover, only a very small fraction of the mesons, exhausting their energy in a layer of about 20 μm from the chip surface, can cause errors due to the high energy-shedding rate in the low end of the energy spectrum. Consequently the effective remaining flux is expected to be well below that resulting from the packaging material.

An important feature of the meson flux is the existence of a correlation between the events, which does not exist for alpha particles. This obviously has a significant impact on any tolerance technique that is based on independence of events. Moreover, the meson energy spectrum is continuous up to very high energies, and their interaction with matter is less intense than for alpha particles. Thus they are much harder to eliminate with a shield. Note in passing that, at sea level, having no shield is probably better than an imperfect shield for reducing the error rate due to high energy particles, since an imperfect shield plays the role of a target, increasing the local flux of secondary particles and the correlation between

events. Consequently cosmic rays introduce a background level of radiation that cannot be eliminated. The effect of cascades of particles is analyzed in Chapter 4.

It is also important to realize that the impact of cosmic rays could become dominant for space applications since, again, the energy spectrum of the flux of particles is not limited as it is for radioactive decay. This makes them difficult to eliminate by shielding. Moreover without the shielding effect of the atmosphere, there will be a significant flux of low energy protons, alphas and heavier nuclei.

Ionizing-Radiation-Induced Error Rate

In this subsection, it is shown that errors occur with an observable probability. The earlier discussions show clearly that a large number of significant parameters must be considered in order to obtain an accurate estimate of the error rate. It should also be clear that every node inside an integrated circuit is a special case, since the sensitivity is determined by: the geometry of the layout, the hit rate, the distribution of the energies and angles of incidence, and the logic function as well as the state and dynamic behavior of the integrated circuit. The complexity and cost of the simulation needed to achieve an accurate estimate of error rate are simply prohibitive for VLSI logic chips, and can only be performed on simple structures [HSI81,SAI82].

The approach followed here consists of obtaining an approximate first-order estimate of the error rate, which gives a much better insight

into the tradeoffs involved. The error rate for a chip can be expressed as a summation of the error rates for the individual nodes, neglecting the probability of counting the same error twice because of the dependence on neighboring nodes, which follows from charge sharing. The error rate is certainly proportional to the effective hit rate, which is composed of the events that can inject a charge above a certain threshold. It is also determined by the composition of the basic particle flux and the node's area. Many nodes are insensitive for a significant fraction of the clock period, and therefore their sensitivity must be multiplied by the fraction of the time where a hit on the node does in fact result in an error. Finally one must consider the *visibility* of an event, which is the probability that an injected transient will propagate to a primary output, and which reflects the structure of the machine and the time spent in each state. The visibility is clearly a first-order-parameter, considering for example a Triple Modular Redundant [SIE82] machine for which most of the nodes have no visibility. For a machine which is not designed for masking errors this parameter is difficult to estimate.

The simplicity of the first-order expression for error rate which is presented below results from making some important approximations, hence it is more realistic to determine upper and lower bounds. This is consistent with our goal of estimating the significance of soft errors for a class of VLSI chips fabricated with a given sensitive technology. The bounds are obtained by either underestimating or overestimating the parameters in an expression describing the error rate. The tightness of the bounds is determined by the amount of resources that one is ready to invest in calculating them.

A pair of values can be estimated for each node, giving the range where the true value for each parameter should be. Taking the product of a particle flux times a sensitive area yields the effective hit rate. The error rate by node is calculated by multiplying the effective hit rate by the visibility and by the fraction of the clock period where the node is sensitive. Consequently the error rate can be estimated by the following expression:

$$\frac{\sum_i f_i a_i t_{1i} v_i}{T} < \text{soft error rate} < \frac{\sum_i F_i A_i t_{2i} V_i}{T} \qquad (2.2)$$

where

f_i, F_i = effective particle flux

lower case: low estimate

upper case: high estimate

a_i, A_i = sensitive area

v_i, V_i = visibility

t_{1i} = latch set-up time for static nodes,

active period for dynamic nodes

t_{2i} = latch set-up time plus pulse duration

after propagation for static nodes,

active period for dynamic nodes

T = clock period

and subscript i refers to node i

In the following, each of the parameters in (2.2) will be discussed. This permits one to understand the effects of scaling and the approximations involved. It is relatively easy to understand the effects of scaling on each

of the important factors determining the error rate.

The effective particle flux is composed of those particles with an energy and angle of incidence that can result in an error for the particular node considered. It increases when the parasitic capacitances decrease. This is particularly important when the technology is at the level where the maximum injected charge becomes just sufficient to inject transients that can propagate. The fraction of the total flux that can cause an error is different for each node, and depends on the size of the node and the impedance of the driver for a static node.

In the case of large-dimension devices, the targets correspond to the diffusion regions and, to some extent, to the channels of the transistors. The nodes can be treated as lumped elements if they do not include long polysilicon lines. When the minimum device dimensions approach 1 μm, the radius of the ionization region [WOL63] cannot be neglected in calculating the sensitive area. This means that, for submicron devices, the sensitive area is much larger than the area of the diffusion regions, and the chances of simultaneous transient injection are high. Therefore, in the upper bound, the areas of some regions are counted more than once, which makes sense, since it corresponds to adding the visibilities of adjacent nodes when charge sharing is significant. This is a union bound which may be fairly tight when the visibilities are small and the delays to the primary outputs are different. This shows that the sensitive area of a submicron chip can be significantly larger than the sum of the areas of the diffusion regions.

For the lower bound on the error rate, no area on the chip should be counted in the sensitive region of more than one node. Notice that

pathological situations can exist involving the use of reconvergent fanout where two adjacent nodes are individually visible, but a hit on the joint portion of their sensitive area is not visible because of a cancellation, as in Fig. 2.4. In such a situation the lower bound is not correct. However, even though it is easy to imagine circuits with this property, because of the high symmetry needed, it is also clear that in a real circuit the occurrence of such events should be relatively low. Therefore it should not significantly affect the accuracy of the lower bound.

The set-up time of the latches in (2.2) has the conventional meaning of the time interval during which the input must not change. The active period for a dynamic node is the time interval during which a logic value stored on it can affect a primary output. It can obviously be different from node to node. This is also true for the maximum expected transient pulse duration from a hit. This last variable is difficult to estimate, since an accurate determination involves a circuit simulation of all the existing paths from every node to all primary outputs. Notice that the state of the machine, by making different sets of reconverging paths active from cycle to cycle, determines the maximum value of the transient duration after propagation. The computational effort required can be reduced by assuming that all reconvergent paths are simultaneously sensitized, which gives a looser upper bound. The amount of computation required for an exact estimation of this variable for a VLSI chip is prohibitive. Experimentation with a prototype of the logic system seems to be the most practical way of achieving an accurate estimation.

Visibility, which was defined earlier, will now be discussed further. This variable reflects the chances of finding a given node on a sensitized

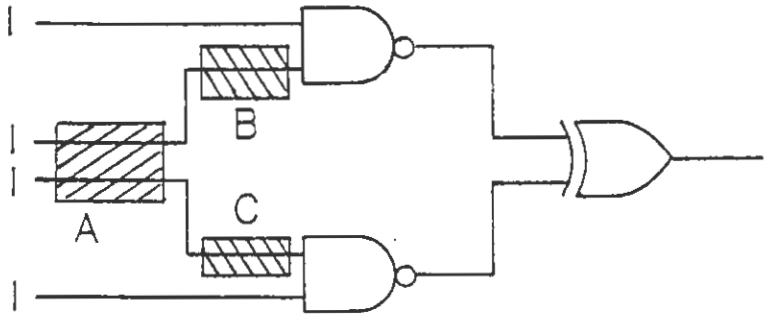

Figure 2.4 A circuit for which the regions "B" and "C" are both individually visible, but a hit on "A" is not visible if the injection is almost the same for both nodes.

path, and it is obviously dependent on the existing redundancy and on the fraction of the time spent in a particular state. This parameter also includes the reduction of the error probability that results from the cases where there is no generated transient, because the injection polarity is the same as the level on the node. In order to estimate the visibility accurately, except for certain structures like a TMR machine, a large amount of logic level simulation would be required. The TMR machine, with a 0 visibility on all internal nodes, proves that this parameter cannot be ignored, even though it is unlikely that one would really try to calculate it for each node of a VLSI chip.

The last important factor for the determination of the error rate is the clock period duration, which becomes shorter as devices become faster with scaling. It is remarkable that, for ionizing radiation, simply running a machine faster amplifies the error rate without any consideration of noise margin or switching energy. A corollary is that the error rate per unit of time decreases when the clock period increases. Nevertheless, it must be stressed that reducing the error rate by simply increasing the clock period is unwise, because it negates the speed benefit of scaling and also, since each cycle is longer, the error probability per cycle is unchanged. There exist much better ways of exploiting time, as will be demonstrated in the next chapters.

The main purpose of this Section is the calculation of a reasonable estimate of the soft error probability. A simple way to make this estimate is to put reasonable estimates for the parameters into (2.2), assuming that all the nodes behave similarly. Suppose a large VLSI chip of 1 cm^2 area, with 20% of its area sensitive to an alpha particle hit. The figure of 20% is

largely influenced by the size of the nodes, which determines the parasitic capacitance of the diffusion regions. Also, a consequence of scaling is that it tends to make sensitive, all the diffusion regions not directly tied to the supply. Moreover, as mentioned earlier, for submicron devices the sensitive area is much larger than the nodes themselves.

Assume a reasonably 'cold' package which yields a particle flux of .1 part./cm^2-hr [MEI79]. The charge injection time is approximately 200 ps [HSI81], and, according to the earlier discussions on transient injection and propagation, the resulting transient is expected to be significantly longer than the injection time. The typical transient is assumed to last 1 ns after propagation, with a register set-up time that is negligibly short in comparison, which will be the case for submicron MOS VLSI. With a 40 ns clock period, which is expected to be typical for a 32 bit microprocessor on a chip based on 1 μm CMOS [GHE84], and an average visibility of 20%, the estimated error rate is 10^{-4}/hour. Obviously the error rate for a particular chip could be very different from this simple estimation, which is believed to be typical for a large chip fabricated with a fast and relatively low power sensitive technology.

2.2.3 Electromagnetic Interference

Electromagnetic interference is another important source of soft errors. Since ionizing radiation also produces transient errors it is not immediately obvious whether an observed error rate is due to radiation or interference. It is argued here that the known data on transient errors are in fact measurements of the effects of interference, because the technologies which were used in these experiments are intrinsically tolerant to the effects of ionizing radiation. How a technology can be intrinsically tolerant to ionizing radiation is discussed later in Section 3.1.1 .

The best available results on the measurement and characterisation of the transient error rate for real computers is the one by McConnel et al [McC79,McC81]. An important result of this work is that the interarrival time for transient errors is better described by the Weibull distribution than by the Poisson distribution. Note in passing that, from the earlier discussion, alpha particles from radioactive decay should result in a Poisson distribution because the individual hits are really independent. The results of McConnel suggest that interference causes a crash rate that is 10 to 50 times larger than the failure rate. This result is even more significant when one considers that not all errors are detected by McConnel's experiment (erronous results may not cause a crash). Moreover if some sections of a machine are overstressed, as is usually the case, they will enter the wearout period much earlier than the rest. This means that the failure rate observed for a mature system is typically higher than the random failure rate.

The methodology used in McConnel's work is not a sufficient characterization for the purpose of designing machines tolerant to the soft

errors generated by interference. The main reason for this is that a single event can corrupt a great deal of data before it is detected and, due to the latency of some errors, it is not practical to separate almost simultaneous transients. Therefore a time threshold must be defined as the minimum delay between two detected errors in order to count them as two distinct events. A time threshold of five minutes was used in [McC79]. This macroscopic information is useful for estimating the probability of various events, but a characterization on a microscopic time scale is also needed.

To obtain such a microscopic characterization of the errors due to interference, each possible source of errors must be considered separately. Some interference sources are tolerable by means similar to those used to tolerate transients due to ionizing radiation, but others are not. The error sources can be separated into two classes: external interference and self-interference. The time of occurrence of events from external sources is independent of the state of the machine, and the expected events may be long pulses with sharp transients (lightning [NEW74], transients in the power distribution line [HAG74]) or a continous high frequency wave (RF sources [HAG74]). On the other hand, for internal sources, the time of occurrence is determined by the machine's state transition, and the expected events are usually short since the duration is a function of the switching time for a given technology [MAR84,RAM84].

The long events will usually last for a large number of machine cycles. For example, the time scale for lightning is measured in milliseconds [NEW74]. For power distribution, there are basically three types of disturbances: direct coupling from power distribution lines which are stable, under-voltage and over-voltage lasting seconds, and relatively sharp

transients lasting a few microseconds, with amplitudes as high as many hundreds of volts [MAR84]. The longest events are usually easy to deal with by good design, since electromagnetic shielding is very effective at low frequency, and a wire need not be considered as a transmission line. The sharp rise of the short power transients and lightning implies a very significant harmonic content. Moreover, the power in the original transients is so high that significant energy may remain in the 100 MHz region, where shielding is difficult and even a wire of modest size, say on a printed circuit board, makes a good antenna and is best represented as a transmission line. No significant pulse can be induced inside a chip, but the connections to the outside world and especially the power lines may experience transients of significant amplitude, thus reducing the available noise margin for internally-generated transients.

The interaction between RF sources and digital circuits involves a completely different type of effect. Clearly, an RF source, like a radar transmitter for example, will not be blocked efficiently by a shield designed for lower frequencies. The signal can easily couple to wires on PCBs and will reach the gates, superimposed on the logic levels. Since latent diodes exist everywhere in an integrated circuit, one should not be surprised if rectification takes place, forming a peak amplitude detector, and therefore shifting the logic levels [WHA79]. At a high RF power, this may cause a gate to behave like a "stuck at" as long as the RF signal is present, but at lower power, it simply reduces the available noise margin for internally-generated transients inside the system.

The internally-generated transients can be separated into two categories: those injected into the supply and those injected into the

neighboring signal paths. For a synchronous system which is not built from current-mode logic, there is always a large current pulse injected into the supply lines due to the quasi-simultaneous switching of a large number of gates. The supply lines exhibit series resistance and inductance. Inside the chip the resistance is important, outside the chip the inductance is usually the problem. This phenomenon is so important that it has to be considered from the very beginning, at the stage where a technology is designed. Means of dealing with the problem at board level are well known [MAR84], but with the scaling of technology, the remaining transients injected inside a chip can be sufficient to cause errors [RAM84]. Technological solutions to this problem exist [RAM84,SON84], but a significant fraction of the noise margin must be reserved to deal with it economically.

The second type of internally-generated transients involves parasitic coupling between adjacent propagation paths. Capacitive and inductive coupling are both important at the board level, whereas at the chip level only the capacitive coupling is important. Reflection on non-terminated signal paths may also cause significant transients at the board level. Since the effect is only observable on the affected line, this type of transient may go undetected more easily than those in the supply and, therefore, is more likely to remain as a reproducible error.

The remaining consideration in interference-induced errors is the effect of scaling. Scaling makes the technology more sensitive to self interference because the switching time decreases. As the number of state transitions per unit time increases, the chances of producing one of those that results in a reproducible error also increase. Another consequence of scaling is that the series resistance per square on a chip scales in the same

way as the coupling capacitance between adjacent wires, but the impedance of the coupling capacitance decreases for higher frequencies. Moreover, the resistance of a wire of a given length increases with scaling, which reflects on the impedance of the supply rails. For VLSI, this problem is amplified further because the relative length of the wires increases with complexity, and so does their series resistance. Reducing the operating voltage also contributes to the problem if it results in a smaller relative noise margin, which follows from the higher relative variance of the transistor thresholds.

When signals go off-chip the problem is associated with the faster transients and their effect on parasitic coupling. This severe problem can usually be solved by modifications to the packaging technology. It is also clear that if the external interference sources are not scaled with the machine's supply voltage, the occurrence of transients of sufficient amplitude to exceed the noise margin can only increase.

In conclusion, interference is significant and the phenomenon may involve complex interrelations between the various sources. It should be clear that a given VLSI chip may have a zero error rate when taken separately but, when used inside a system, once in a while a transient will exceed the noise margin and cause an error. It should also be obvious that, by being conservative, the designer has a direct impact on the error rate. Clearly there is a tradeoff between the cost and the intrinsic (or non-redundant) reliability of a technology. If an efficient technique can be devised for tolerating rarely occurring transients, it may be possible to reduce the cost and keep the same reliability.

2.3 Pulse Propagation

The following discussion applies to pulses injected from all possible sources in a combinational logic network. There are three obvious conditions for pulse propagation in a logic network: the pulse must have a sufficient amplitude, a sufficient duration and must be on a sensitized path.

What is less obvious is how the duration of a pulse is modified by propagation. In particular, assuming a single sensitized path, the model simplifies to a cascade of inverters. There is a large difference between the maximum duration of a pulse that will not propagate to the next stage, and the minimum duration of a pulse that would propagate in a cascade of arbitrarily large depth. This is demonstrated by the simulation results in Fig. 2.5, which show that a 2,5ns pulse does propagate through one inverter, whereas a 8ns pulse easily propagates through 8 inverters, but would not propagate to depth much larger than 8 since the pulse is decreasing in amplitude and duration with propagation. This means that the minimum duration for an event to propagate is a function of the logical depth to the primary outputs.

Moreover, if a rising edge does not propagate at the same speed as a falling edge along a given path, the duration of a pulse may increase or decrease with propagation. This will happen if the rise and fall times are different for a cascade of gates which is not evenly loaded, as demonstrated in Fig. 2.6. The load imbalance in that simulation is an area of 20 squares (5μ by 5μ) of diffusion on the outputs of the even inverters. It is clear that, for a positive pulse on the input, the duration decreases with propagation, whereas for a negative pulse it increases. This phenomenon is

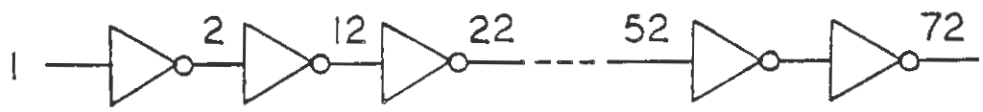

(a)

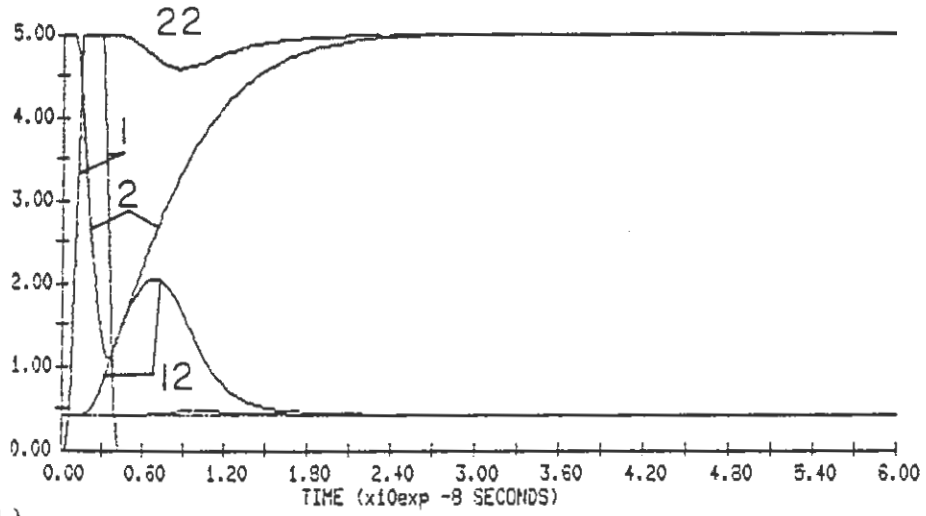

(b)

CHAIN OF INVERTERS, NMOS, K=4, PULSE OF 2.5NS 10-JUN-85

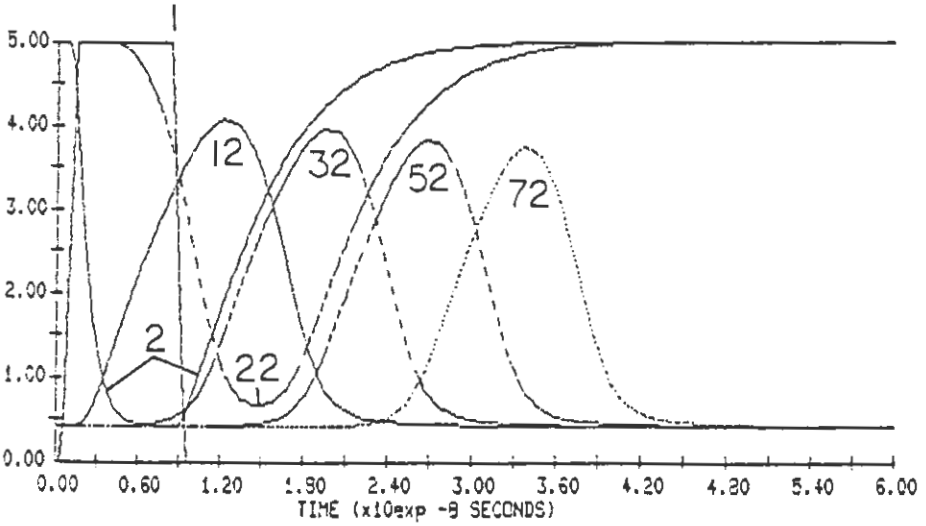

(c)

CHAIN OF INVERTERS, NMOS, K=4, PULSE OF 8NS 10-JUN-85

Figure 2.5 (a) A chain of 8 inverters (5μm NMOS not loaded),
(b) Response to a pulse of 2.5 ns,
(c) Response to a pulse of 8 ns.

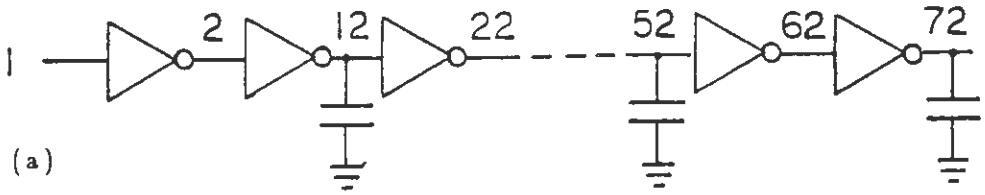

(a)

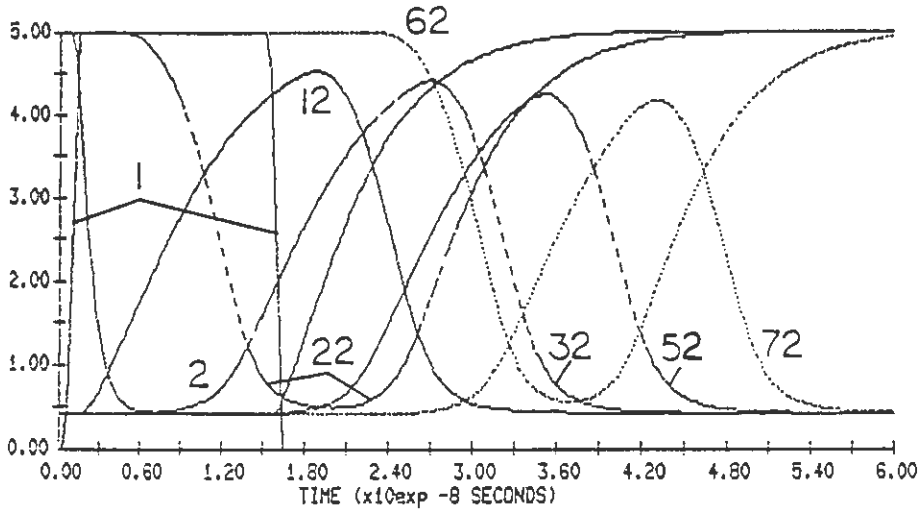

(b) CHAIN OF UNEVENLY LOADED INVERTERS, POSITIVE PULSE, 15NS 10-JUN-85

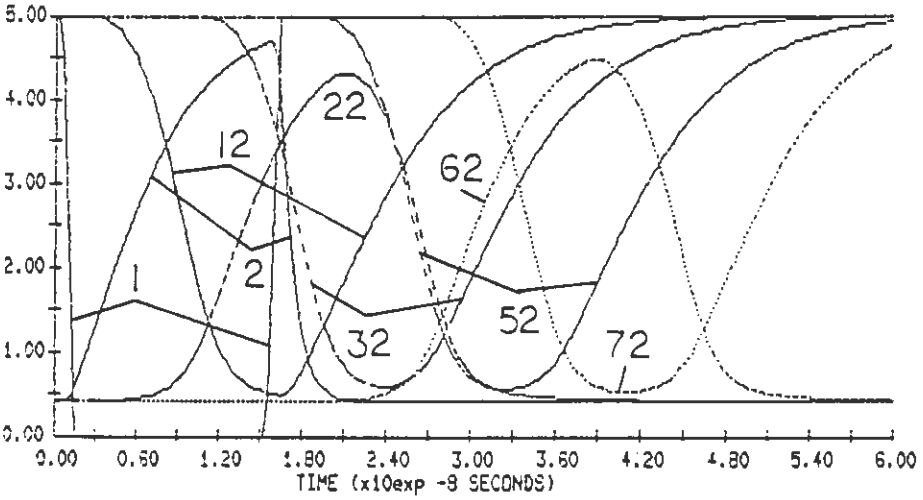

(c) CHAIN OF UNEVENLY LOADED INVERTERS, NEGATIVE PULSE, 15NS 10-JUN-85

Figure 2.6 (a) The same chain of inverters shown in Fig. 2.5, with a load of 20 squares of diffusion on nodes 12, 32, 52, and 72,
(b) Response to a positive pulse,
(c) Response to a negative pulse.

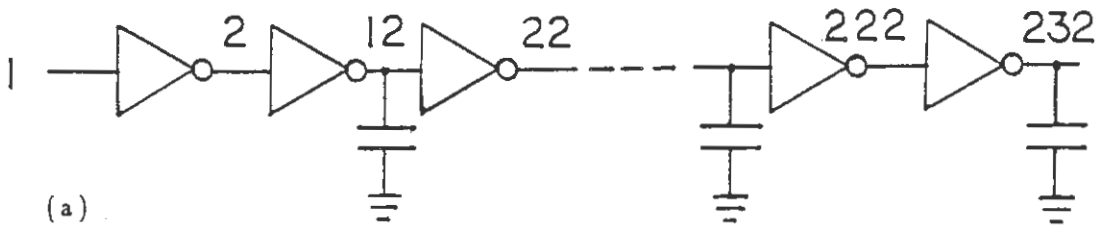

(a)

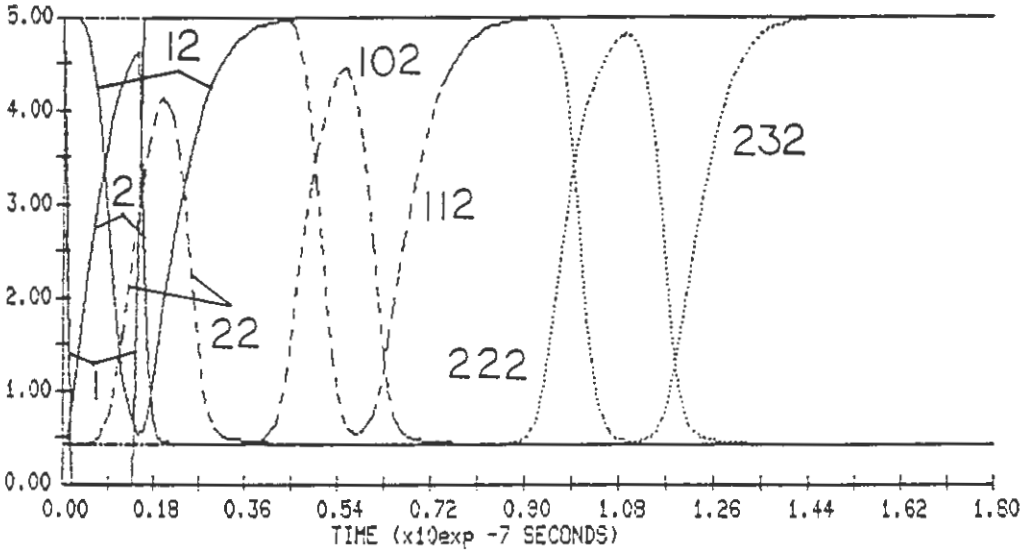

(b) CHAIN OF UNEVENLY LOADED INVERTERS, NEGATIVE PULSE, 14NS 10-JUN-85

Figure 2.7 (a) A chain of 24 NMOS inverters unevenly loaded with 20 squares of diffusion.
(b) Response to a negative pulse of 14 ns.

significant because the duration gets increased or decreased by a fraction of the difference between the rise and fall times, which is independent of the initial pulse duration. In other words a relatively short pulse just sufficient to propagate can become arbitrarily long, provided the existence of a sufficient logic depth, as demonstrated in Fig. 2.7 with a cascade of 24 inverters.

A second phenomenon which is even more important in practice is the effect of reconvergent fanout. When two or more paths are simultaneously sensitized from an affected node to a primary output, the individual pulses may add up cumulatively to increase the transient duration. The significance of reconvergent fanout is demonstrated by the example in Fig. 2.8. Here, a number of paths with slightly different delays transform a pulse, of duration just sufficient to propagate, into a pulse as long as the maximum propagation delay in the logic minus one gate delay. This approaches the duration of the clock period if the maximum operating frequency of a machine is determined by the delay in the combinational logic.

The examples given are possible, but certainly not typical. In order to increase significantly the duration of a pulse by propagation along a single path, a large logic depth is necessary. Otherwise, if the delay is lumped in a single gate, the transient does not propagate and there is no problem. The practical logic depth of a network will rarely exceed 20, and 10 is more typical [GHE84].

The case of a large number of reconverging paths shown in Fig. 2.8 can be interpreted differently, if one realizes that a large load on the IN node is implied, thus making the node intrinsically tolerant. It is

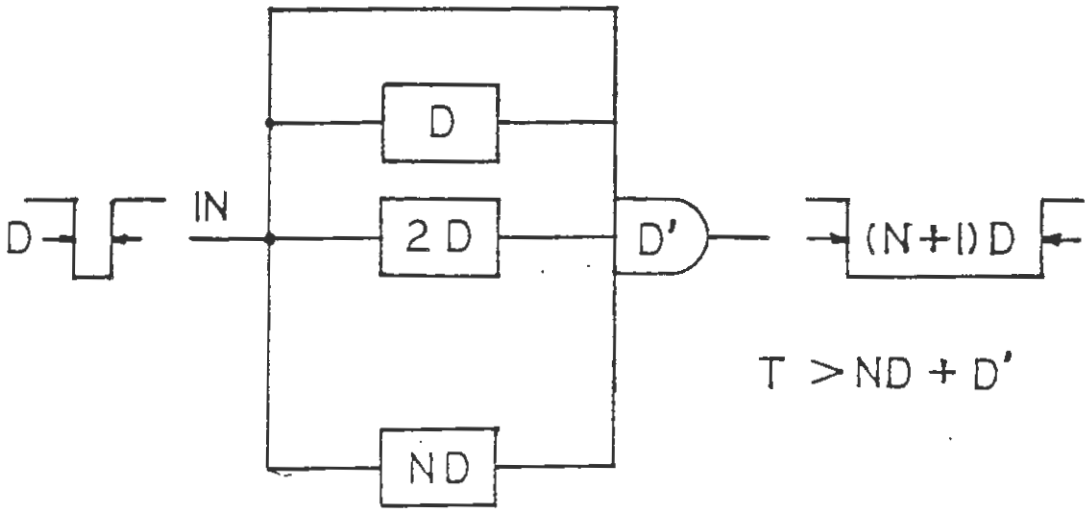

Figure 2.8 Pulse spreading due to reconvergent fanout.

noteworthy that increasing the capacitance of one node on the propagation path is not necessarily sufficient to stop a propagating pulse, particularly if the system has been optimized for speed with the insertion of a suitable buffer to speed up propagation. A buffer can act as power amplifier for propagating a transient pulse.

There is a good reason to believe that reconvergent fanout has a stronger effect than pulse spreading on a single path. The effective number of paths with different delays grows as the product of the fanout of the reconvergence points in series on a path, as shown in Fig. 2.9. If all these paths have effectively different delays to the output node, a short transient injected on the input of this structure could easily be transformed into a transient of a duration approaching the longest delay in the structure.

In conclusion of this Section, in the worst case the spreading of a pulse in a combinational network can be very important. However, we believe that in real circuits it is usually limited. A quantitative characterization of pulse spreading for combinational logic function would be needed, to support the design methodology presented later in Chapters 4 to 6. This analysis is not included in the thesis and is left for further work.

2.4 Significance of Soft Errors

It is appropriate here to discuss the significance of soft errors, since the error rate is sufficiently low to be ignored in many situations. A user of a digital machine usually assumes that his machine is error-free until it develops a permanent failure, which is reasonable for small systems.

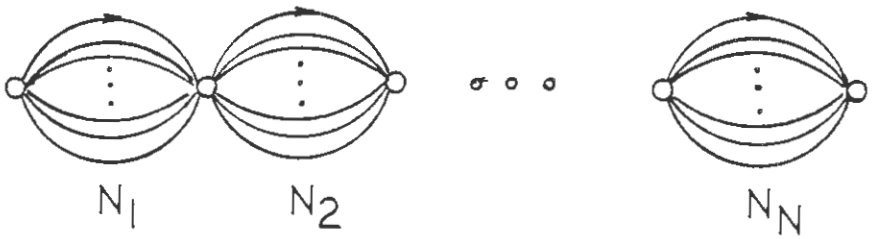

$$N = N_1 N_2 N_3 \cdots N_N$$

Figure 2.9 For the simple linear structure shown, the total number of reconverging paths in a logic network, with possibly different delays, is given by the product of the internal reconverging fanouts.

The earlier discussions show that this is not true in general. The error rate is obviously a function of the system's complexity. This function is a straight proportionality relation if the system has no built-in tolerance to soft errors. On the other hand, because of wearout of components or obsolescence, even a small system has a limited life. Consequently, there exists a minimum complexity below which the error rate is not significant. When the effect of scaling is considered, it is clear that this minimum complexity, measured in area of silicon, will permit one to fabricate fairly complex devices.

For systems larger than the minimum complexity, in the ideal situation, the error rate of the machine should be determined by the chip failure rate. It is well known that for very complex systems, even the low random failure rate, after burn-in, yields a significant failure rate. If the soft error rate per chip is larger than the random failure rate, then it will limit reliability and is therefore significant.

The assumption that error-free computation is necessary may be too pessimistic in certain situations. There exists a special case where the error rate may not be significant, even though it dominates the failure rate. Consider a complex machine that has a fairly simple controller, where the data on which the machine operates has no effect on the state of the machine. Examples of machines with this structure are hardwired digital filters and decoders. If the controller can be hardened in such a way that it is immune to soft errors, there remain only the errors in the data manipulating sections. If this machine operates on data with an error rate per bit in the range of 10^{-6}, the error rate contributed by the machine itself is probably negligible.

2.5 Reliability Trends

It is clear from the previous discussion that a knowledge of the error rate is not sufficient to determine its significance. The error rate due to radiation and interference tends to increase with the scaling of the technology, as shown earlier. On the other hand, the evolution of the random failure rate with scaling also affects the significance of the error rate.

The failure rate is a strong function of how the technology is scaled. A number of reliability problems exist [WOO81] including: electromigration [GHA82,HO82,NAG79], hot electrons, dielectric breakdown [ANO79], radiation exposure [DAV82] and, in general, the effects of heating due to power dissipation. Methods of dealing with these problems have been developed as they became significant [GHA82,SON84,MOR84,PEA83, WOO81]. The failure rate varies as a function of time, and is usually represented by the lognormal distribution and a refinement of it, assuming that a fraction of the population are freak devices [ANO79,GHA82,WOO81].

After each failure mechanism in a given process is sufficiently well understood, design rules and process parameters can be chosen to adjust the stress to a level that yields the required reliability.

An important question is the relationship between the complexity and the failure rate. In the MSI to LSI range of complexity the failure rate grows as the square root of the number of gates [SIE82]. For commercial microprocessor chips, the measured failure rate was around 0.03 failure /

1000 hours [PEA81]. Assuming that the failure rate for VLSI continues to grow as the square root of the complexity, it places the failure rate of a VLSI chip around 0.1 % per 1000 hours or 10^{-6}/hr (temperature and package quality can change this figure by 2 orders of magnitude in both directions [WOO81]).

It is interesting to examine the reliability trends and goals for components in Table 2.1. Notice that for a device of a given complexity, the trend clearly goes in the direction of a reducing failure rate with time. Peattie [PEA81] mentions that goals for failure rate as low as 1 per billion hours are envisioned.

Table 2.1 Reliability trends and goals

Failure rate, % per 1000 hours

Year	Automotive electronic engine control [FLI81]	Digital logic circuits [PEA81]
79	0.12	0.0005
81	0.035	0.0004
83-85 (goal)	0.0025	0.0003
88 (goal)	0.00025	-

The prediction that the failure rate will increase as the square root of the number of gates is pessimistic. Firstly, the smaller geometries imply smaller chances of incorporating a weakness in each device. Also, if the same yields are to be achieved from VLSI devices as those achieved a few years ago

by LSI devices, then the quality of the fabrication process must improve. Moreover, VLSI systems are likely to be built with as many chips as the LSI systems of the earlier generation. At the same time, these complex systems have to be more dependable because they are often performing critical functions. This will translate into a demand for high reliability components. Therefore, it seems plausible to have a 1 cm^2 chip with a failure rate of 10^{-8}/hour.

The soft-error rate due to ionizing radiation calculated earlier for a large VLSI sensitive chip was 10^{-4}/hr. Depending upon whether the reliability of VLSI chips will correspond to the higher or lower estimate, the error rate dominates by two to four orders of magnitude. This demonstrates that a system built from very reliable components is more affected by soft errors.

An important factor in determining the significance of soft errors is the gamma ray exposure. This yields a predictable gradual shift in the transistor's threshold [DAV82], which will eventually result in a *stuck at* behavior. However, since the phenomenon is gradual with exposure, the system exposed to gamma rays will operate with a decreasing noise margin. Therefore the error rate due to interference may become unacceptable, long before a permanent failure can be observed.

To summarize, it has been demonstrated that the soft error rate will have a greater significance as the technology scales down. Both interference and ionizing radiation induce short transients that may propagate inside a circuit, and cause errors if they are memorized. Therefore a machine designed to be immune to short transients would have a significantly lower error rate.

Ideally the error rate should become smaller than the failure rate, which would then remain as the only important factor determining reliability. A design methodology aimed at tolerating the expected transients would be very useful, enabling the exploitation of the full potential of VLSI devices for building complex, reliable systems.

Chapter 3
Conventional Methods for Decreasing the Soft Error Rate

The important sources of soft errors were identified and their salient characteristics were analyzed in Chapter 2. A new general method of tolerating soft errors will be presented in the following chapters of this thesis. However, before proposing the new approach, it is appropriate to review the conventional methods of dealing with soft errors.

Methods found in the literature for decreasing the soft error rate can be separated into two classes. The first will be called physical approaches, because they are typical of how physicists attack this kind of problem. An interaction mechanism is identified between a source and a receiver, for example, an ionizing particle and an electronic circuit. If the interaction jeopardizes the normal operation of the device, a means is proposed to reduce the said interaction to an insignificant level. The second class of methods will be called system approaches. Typically the system designer works with a very abstract model of how the device operates. This model may lose important features of the mechanism by which an error occurs. For example, the fault model may be as simple as assuming that a given output line takes the wrong logical value. The proposed solution will generally permit the masking of this incorrect logic value by adding redundancy into the circuit.

It will be seen that the new approach proposed in this thesis does

not fit into either category, however it borrows from both; the physical level by using a refined interaction model, and the system level by introducing a form of redundancy which does not involve reconverging signal paths.

The known solutions at the physical level are presented in Section 3.1 and those at the system level are presented in Section 3.2. The system-level techniques are not usually specialized for soft error tolerance. Straightforward extensions of the standard system level techniques are presented. They result in a better tolerance to transient errors. In each case, reasons are given to support our opinion that there is room for a new, more general and efficient method of tolerating soft errors.

3.1 Physical Level Solutions

Three means of decreasing the soft error rate due to ionizing radiation are presented first. Then the conventional methods of decreasing the error rate due to interference are discussed. Finally this section concludes with an evaluation of the effectiveness of these techniques in solving the problem of soft errors.

3.1.1 Ionizing-Radiation Induced Soft Errors

A first category of solutions to the problem of reducing the error rate due to ionizing radiation, consists of decreasing the particle flux that reaches the sensitive regions. The first and more drastic solution, proposed

by May and Woods [MAY78], consists of refining all the materials that compose an integrated circuit, in order to decrease the concentration of radioactive impurities. Since these impurities already are present as traces in the material, this solution is obviously a very expensive one, if not economically unfeasible. The error rate reduction achieved by such further material refinement is of the order of one or two orders of magnitude, but at a substantial cost [MEI79].

The depth of penetration of alpha particles in materials of average density is less than 100 μm, thus a fair compromise is achieved by coating the chip with a layer of very pure material, in order to absorb the particles emitted by the package. This solution was originally proposed by May and Woods [MAY78]. This is much less costly, but is not as efficient as the former proposal of extreme purity of materials. It was mentioned in Chapter 2 that, when this solution is applied, radiation mainly originates in the top layers of the chip itself [SAI82]. Therefore the hit rate does not decrease to zero, furthermore the fraction of the hits with a low angle of incidence increases significantly. Therefore, reducing the flux of particles is an effective method for decreasing the error rate, but it does not eliminate the problem and it entails significant costs.

A second solution consists of developing an intrinsically tolerant technology. In the case of dynamic RAMs, it corresponds to increasing the capacitance of the storage nodes in such a way that the maximum injected charge becomes insufficient to cause an error. For the case of logic machines built with static circuits, the saturation current of conducting transistors can also be increased. By increasing saturation current and parasitic capacitance simultaneously, a point can be reached where the amplitude of the

injected voltage transient is not sufficient to cross the threshold of the gates. The saturation current needed to achieve intrinsic tolerance is of the order of 0.3 mA. This assumes that there is no current amplification by parasitic bipolar structures such as described in Chapter 2.

The value of the saturation current of a transistor is determined by both the aspect ratio and the details of process scaling. Therefore intrinsic tolerance can be obtained by increasing the minimum width of the transistors, to a value larger than the minimum permitted by the resolution of the fabrication process, which obviously would imply a penalty in circuit density. It may be necessary to increase the width of transistors by a factor of more than 5, in order to obtain intrinsic tolerance with a submicron process. But in doing so, one would readily double the area necessary for implementing a given system. This factor would have to be increased further, if only a fraction of the total noise margin must be reserved for ionizing radiation, or if the parasitic bipolar structures in CMOS turn out to amplify the peak injected current. Therefore this approach can become very costly. Moreover, increasing the saturation current of transistors by increasing their width also increases the dissipated power proportionately. This, eventually, would limit the complexity of chips due to cooling problems.

Scaling a process at constant voltage can make it intrinsically tolerant, because it tends to increase the saturation current of the minimum size devices, whereas scaling at constant field tends to decrease it [HOD83 p.114]. Therefore, by scaling at a constant voltage, it may be possible to increase the saturation current of a minimum size device, to a level that makes a scaled technology intrinsically tolerant.

For the NMOS technology, assuming a 5 V supply, the minimum dissipated power per ON gate is set to 1.5 mW (0.3 mA of saturation current), or 75 W for a 100 kilogate chip. This is beyond the capabilities of economic air-cooled packages. Therefore this approach, which defines a minimum power per gate, limits the maximum complexity of a chip to somewhere around 3 kilogates per chip, which is a serious limitation. It is important to stress that the figure of 1.5 mW per chip assumes MOS simple gates, and is not valid for TTL or ECL gates which have internal nodes with an impedance much higher than their output impedance. However, the same argument holds with a different value of minimum dissipated power per gate.

CMOS needs to be considered separately due to its smaller dissipation per gate. Again, scaling theory [HOD83] predicts that scaling CMOS at constant voltage leads to an increase of the switching power per gate. Unavoidably, scaling leads to a dissipation problem one or two generations later, which is similar to that encountered with NMOS. This problem will either limit the maximum complexity, or the maximum frequency at which a chip can be operated.

There are clear indications that the semiconductor industry evolves toward a reduction of the dissipated power per gate for NMOS, as indicated by the statistics on the VHSIC program published by Fischetti [FIS82]. Power per gate of 37μW and 100μW are reported for 1.25μm NMOS at Texas Instruments and IBM respectively [FIS82]. Also, at the speed and complexity now attainable with CMOS, a reduction of the switching power per gate appears unavoidable. This leads to a reduction of the saturation current of minimum sizes devices. State of the art cooling techniques [PEA83] could be used to continue with the constant power per gate approach for a few generations. However, these

techniques are costly and therefore they are usually confined to the high performance mainframe or supercomputer market niche, which represents a very small fraction of the fabricated integrated circuits. Thus, an intrinsically tolerant machine can be designed with a small error rate, but it will require either an area significantly larger than that required by a non-tolerant machine, or will be limited by the dissipated power per chip.

A third solution for dealing with ionizing radiation consists of modifying the fabrication process, in such a way that the collection efficiency is reduced. For example, Sai-Halasz et al. [SAI82] propose the fabrication of integrated circuits, with a layer of inverted dopant polarity buried under the active devices. This effectively reduces the collection efficiency and thus the error rate. However, it does not eliminate the problem. Also, it is not clear at all that such a technique is scalable, since from generation to generation the collection efficiency must be progressively reduced. It seems unlikely that the ionizing radiation problem will be solved by this means in the future.

3.1.2 Interference

Interference control is a classic problem in electronic design. It is generally possible to identify a source, a coupling mechanism, and a receiver [PAU81,SPI81]. The interference problem usually follows from characteristics not considered by a designer in the model of a system. A more accurate model of device operation will permit the inclusion of the coupling between the source and the receiver. The solution of an observed problem consists of modifying the system in order to reduce the coupling, the amplitude of the

source, or the sensitivity of the receiver. The book by Mardiguian [MAR84] is a good survey of the known techniques for dealing with interference in the context of computer design.

In theory, there is no reason why a machine could not be designed with a zero error rate due to interference, but in practice this could only be achieved at a high cost or at a loss of performance. For example, if a machine is an efficient design, with a noise margin just sufficient for proper operation, it is probably sensitive to abnormal electromagnetic events in its vicinity. Also, it is well known that a complex design is generally used without being completely tested. This means that certain untested state transitions can result in an error. Consequently, in theory, there is no problem in making the interference-induced error rate negligible, but in fact interference determines the machine error rate, as described in Section 2.2.3.

3.1.3 Efficiency of the Physical Level Techniques

Each of the physical level techniques described earlier is specific for dealing with a particular type of error-generation mechanism. When many sources of soft error affect a given machine, a number of physical level techniques must be used simultaneously to make the error rate negligible. However, the cost of soft error tolerance is the sum of the costs of each individual technique. Even though some of these techniques can decrease to 0 the error rate due to the source for which they are designed, they usually imply high costs or restrictive constraints that make them inapplicable in general.

3.2 System Level Solutions

The usual underlying model adopted for system level solutions is relatively simple. A fault that changes the output of a gate, may change the state or the output of a machine, thereby resulting in an error. A fault may be transient, but is usually treated as a permanent fault for a given number of machine cycles. The theory of fault tolerant systems [SIE82] permits two types of solution to the problem of soft errors. The first type consists of detecting the occurrence of a transient error and retrying as necessary. The second type of solution consists of masking an error when it occurs. Both types of solution are discussed in the following sections. It is assumed here that tolerance to a single transient error is sufficient for neglecting the error rate.

3.2.1 Detection and Retry

Detection and retry yields the lowest overhead for tolerating transient errors. Depending upon how much tolerance is required, the overhead can be anywhere from very small to more than 100%. If a machine is sufficiently versatile, detection may be provided by software which systematically checks the consistency of the result. Hardware overhead remains small, but usually at the expense of a high time overhead. Low overhead detection techniques exist, but they usually result in a reduced error coverage.

A problem might limit the reliability of a machine where the

tolerance is based on redundant software: the existence of a *hard core*. The hard core is composed of the logic circuits that can affect the critical part of the state. If, for example, a soft error results in an arbitrary jump, the atomicity [AND81] of actions cannot be guaranteed, and the error cannot be confined. Consequently the machine could still crash as a consequence of a transient fault which corrupts the hard core of its state.

Moreover, with the current trends in hardware and software costs, the shifting of complexity from hardware to software is a questionable choice in low volume applications. The advantage of a lower cost for the hardware could be outweighed by the cost of the software, if non-stop operation is to be achieved with some level of confidence in a large machine.

In conclusion, if only the cost of hardware is considered, detect and retry based on software is a possible low-cost alternative, however it does not achieve the same level of tolerance to transient errors as hardware techniques, and it also assumes that the machine has a computer-like architecture, which is not always the case. Thus the detection and retry solution is often not satifactory. Therefore the rest of this chapter deals with hardware techniques only.

The use of arithmetic codes is a relatively low overhead detection technique, if one is prepared to sacrifice coverage. If an error is not caught when it occurs, no further testing of that transient fault can be done to determine the source of the error. Moreover, this technique applies only to selected portions of a machine and, in particular, is not applicable to the control section.

Assuming full duplication, it becomes possible to design a machine that is very robust with respect to transient errors. In a duplicated system, two machines operate in parallel, and a comparator on the output detects any single error as it occurs. After an error has been detected, a recovery mechanism is initiated. In many situations, where the machine can be stopped for brief periods, software retry will be sufficient. If a system cannot be stopped, a hardware retry mechanism is more appropriate.

Duplication does not provide a sufficient amount of redundancy for resolving a conflict between the states of two machines. Therefore, the state must be unique and the memory elements forming it are not duplicated, which decreases the overhead. The key to building robust machines is to guarantee the integrity of the state. This can be achieved by coding the state or by making each bit intrinsically tolerant. The best solution depends on the source of error that is to be neutralized, and on the number of memory words among which a decoder could be shared. One way to ensure the integrity of the next output and state is to latch them only when they match.

If a machine is designed for soft error tolerance from the beginning, the cost of hardware retry is negligible, but yields a much more robust system, especially when bursts of errors are expected. However, to minimize hardware overhead, a sufficiently precise fault model must be adopted. In particular, it may be necessary to consider the duration of the transient events. For example, a transient could affect an output line in such a way that the output bit is changed, but the output of the comparator does not reflect this change at the sampling time, because both comparing and latching require a finite amount of time. Consequently a tolerance technique developed for permanent faults may fail with transient faults. The tightly coupled

Double Modular Redundant (DMR) machine, shown in Fig. 3.1, is the lowest overhead general solution obtained by modifying a conventional technique. An alternative to the solution in Fig. 3.1 is to duplicate the Φ_2 register, and to compare only after the registers. This increases the overhead, but the duration of the transients no longer needs to be considered.

For the machine in Fig. 3.1, if a single transient fault in the logic is assumed, it must occur either in the functional part or in the comparator. In any case, the comparator must flag the event and invalidate the output of the logic. If the transient is short, there may be no overlap between the mismatch on the output lines and its detection pulse by the comparator, which defeats the purpose of duplication.

A possible solution is shown in Fig. 3.2, where the output of the comparator is monitored by a Set-Reset latch. The latch must be reset before the data valid period, and the detection of a mismatch sets this latch. The time window, during which the mismatch line must be false, should extend for at least one comparator delay after the clock of the output latch. A careful design at the circuit level can guarantee that a glitch sufficiently long to upset an output latch will propagate through the comparison logic, and set the S-R latch. This can be done by slowing down the output latch. The tradeoffs involved in this kind of design are discussed at length in Chap. 5.

Such a machine is a very robust one with respect to the bursts of errors that would typically result from a very intense electromagnetic event or a radiation flux, provided that the state bits are intrinsically tolerant. In such a case many output bits would be corrupted and only a perfect match

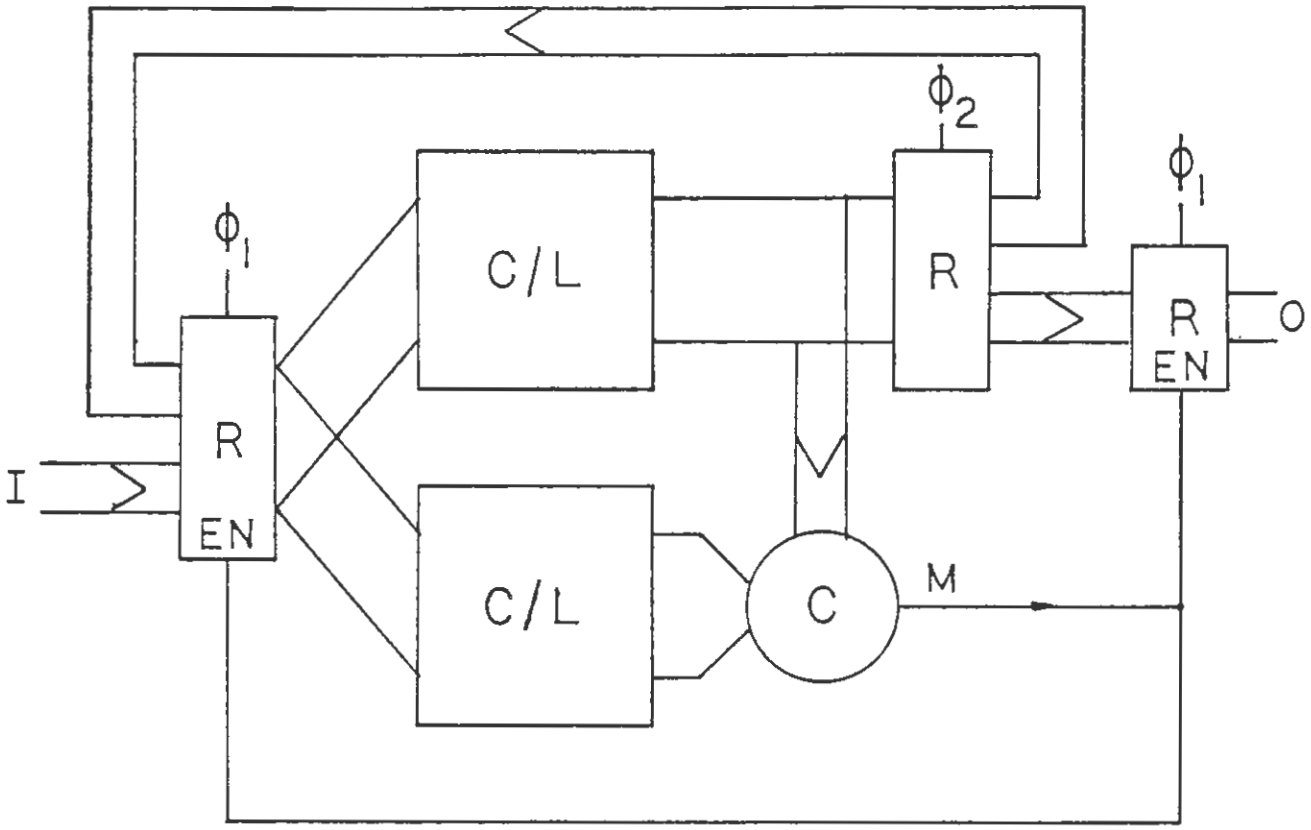

Figure 3.1 Tightly coupled Double Modular Redundancy
R: register C/L: combinational logic
C: comparator EN: enable

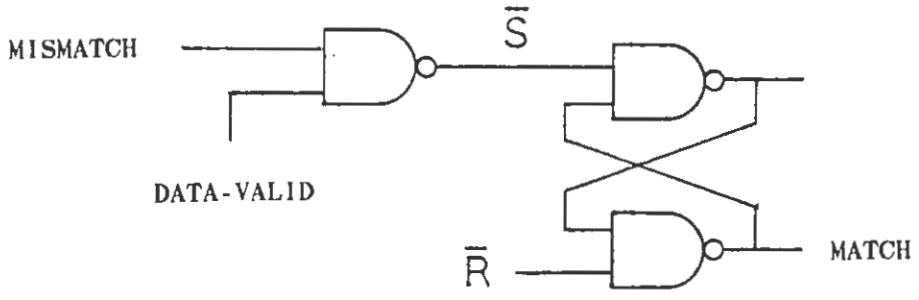

Figure 3.2 A circuit for validating the output of a DMR machine. This circuit is connected on the output of the comparator.

of all the outputs, for the total duration of the sampling window, would result in an error. This machine would be frozen for as many cycles as is necessary for the output to become noise-free again.

When performance and overhead are considered, tightly coupled DMR is a powerful general technique for tolerating transient errors. An efficient detection mechanism generally requires duplication, and regardless of the details of implementation, the overhead is at least on the order of 100%.

3.2.2 Masking Redundancy

Masking redundancy, as the name implies, consists of providing redundancy in such a way that the consequences of a fault are not visible on the output. A general form of masking redundancy, which is always used when a memory needs to be protected, is error correcting codes such as Hamming codes. In this thesis, using error correcting codes for protecting memories from transient errors is considered to be a solved problem [SAR84]. As long as sufficient precautions are taken for limiting the effects of correlated events, memory should be implemented with the densest possible dynamic RAM, and protected by a code. It is already possible to make dynamic RAMs significantly denser by not trying to make them intrinsically tolerant, and the ratio of the area of a tolerant RAM compared to that of a non-tolerant one can only grow with scaling.

The most common form of masking redundancy for logic is modular redundancy. In the general form, N modules are performing the same computation in parallel, thus the name NMR. The result is derived by taking

the majority of the outputs for the N modules. A well known particular case is the Triple Modular Redundancy or TMR with N=3. Other forms of masking redundancy exist, including arithmetic codes, interwoven logic, and the coded state machine [REE70].

Arithmetic codes can correct errors, but require a substantial overhead. Moreover, the technique lacks generality since it only applies to logic performing selected arithmetic operations. Interwoven logic, by providing tolerance at the gate level, results in a high overhead. A coded state machine could be an interesting alternative in selected applications, but only when it requires less overhead than TMR. However, in general, it would result in a much higher overhead, and the only means of determining the overhead is by a detailed design.

From the above considerations, one concludes that the most efficient general method for masking an error is TMR. The reliability of a TMR machine, where the modules are subject to transient errors, is a function of the exact implementation structure. This is made clear by comparing a loosely coupled TMR machine as in Fig. 3.3(a), with a tightly coupled one as in Fig. 3.3(b). An error affecting the state in one of the modules of the loosely coupled machine, may result in a loss of synchronism. If the states are not systematically compared, this error may have a long latency period, where an error in one of the remaining modules can cause a crash of the TMR system. Therefore, all the state bits need to be regularly compared to remove such a discrepancy. This cannot happen with the tightly coupled version.

The reason for considering the loosely coupled version is a practical one: such systems can be built from off-the-shelf modules, not specifically

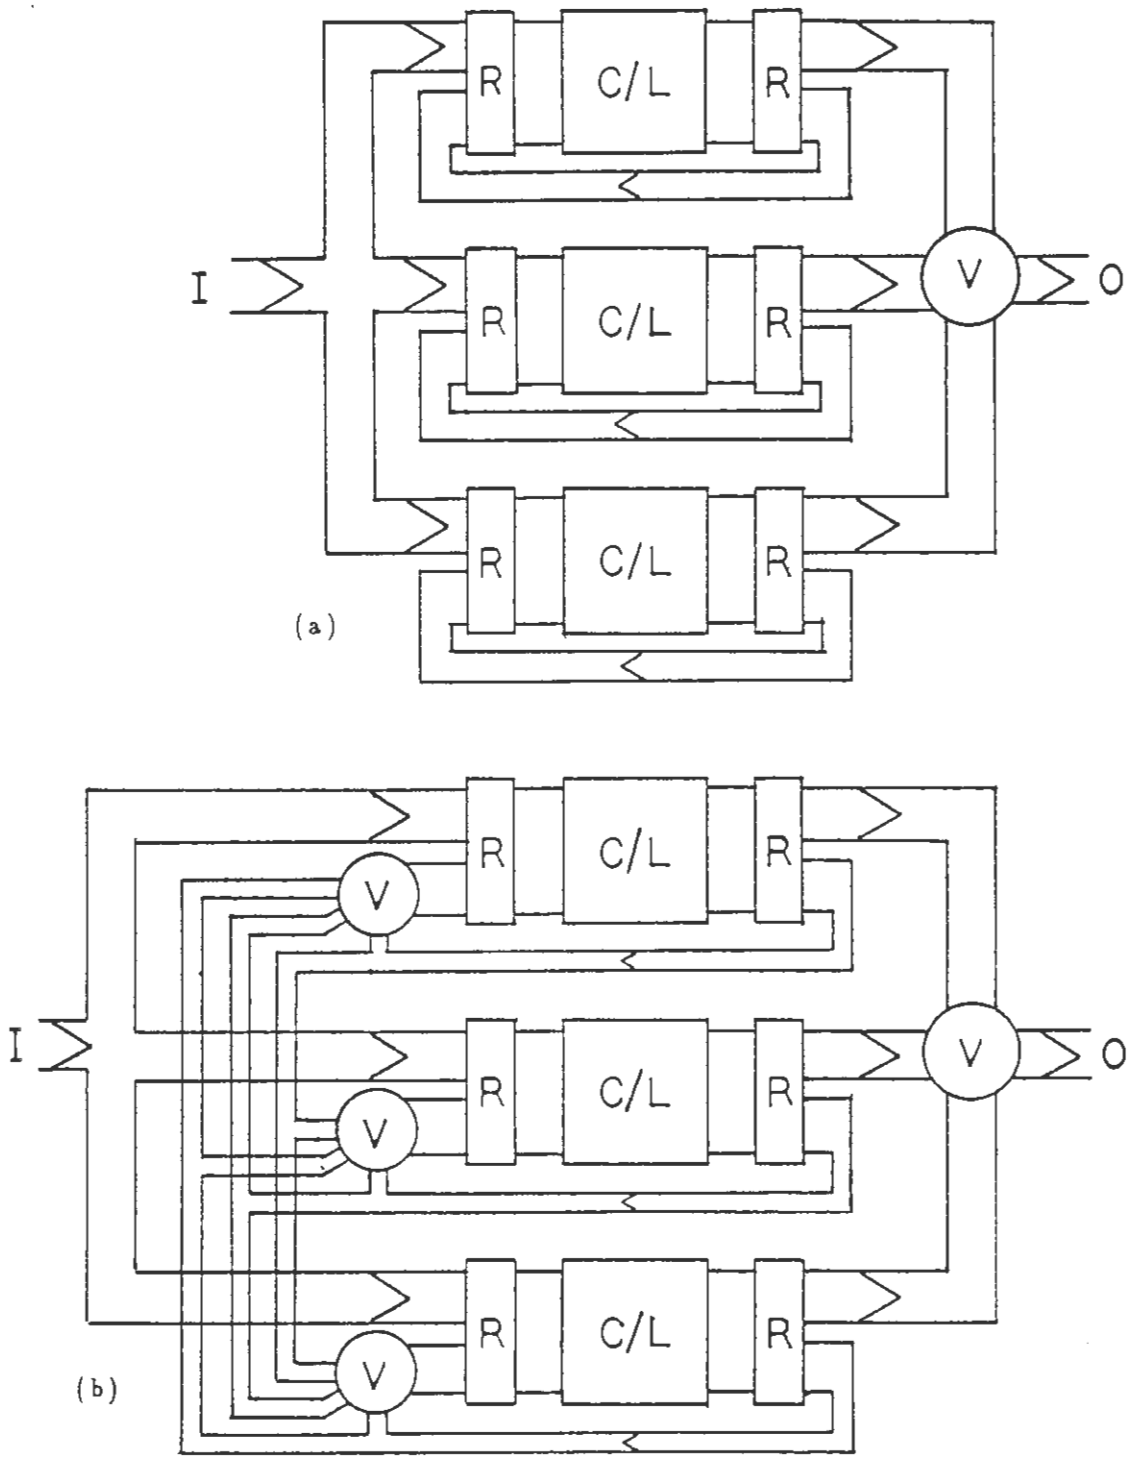

Figure 3.3 (a) A loosely coupled TMR machine
(b) A tightly coupled TMR machine

designed for being part of a TMR system. The main advantage of the TMR system, over the tightly coupled DMR already presented, is its capacity to tolerate at least one permanent fault. However, the pure TMR is less robust with respect to bursts of transients than the tightly coupled DMR. This is particularly true if the outputs are voted on a bit by bit basis, as is usually the case in practice. A burst of transients violates the basic independence assumption that gives TMR its ability to improve reliability. It is very likely that a burst of transients would corrupt all three modules, resulting in an error or a crash. This weakness of TMR can be solved by the scheme proposed in Fig. 3.4. It is noteworthy that the 3 combinations of 2 machines out of 3, form 3 tightly coupled DMR machines. These machines would simply ignore a noisy output, leaving as many cycles as necessary for the burst of transients to disappear. After the occurrence of a permanent fault, the circuit would continue to operate as a tightly coupled DMR machine, keeping the attribute of tolerance to at least one transient fault. The reliability gain for the machine shown in Fig. 3.4 could be significant, considering the relatively higher frequency of soft errors.

The best solution with masking redundancy implies at least 200% overhead. Therefore, among the conventional techniques, the detect and retry approach that leads to tightly coupled DMR, is the most efficient system approach for dealing with transient errors. The Soft-Error Filtering technique, proposed in the next chapters of this thesis, is a new masking technique that can break the 100% hardware overhead barrier while keeping a small time overhead. It is demonstrated later that the overhead can be much smaller than 100%.

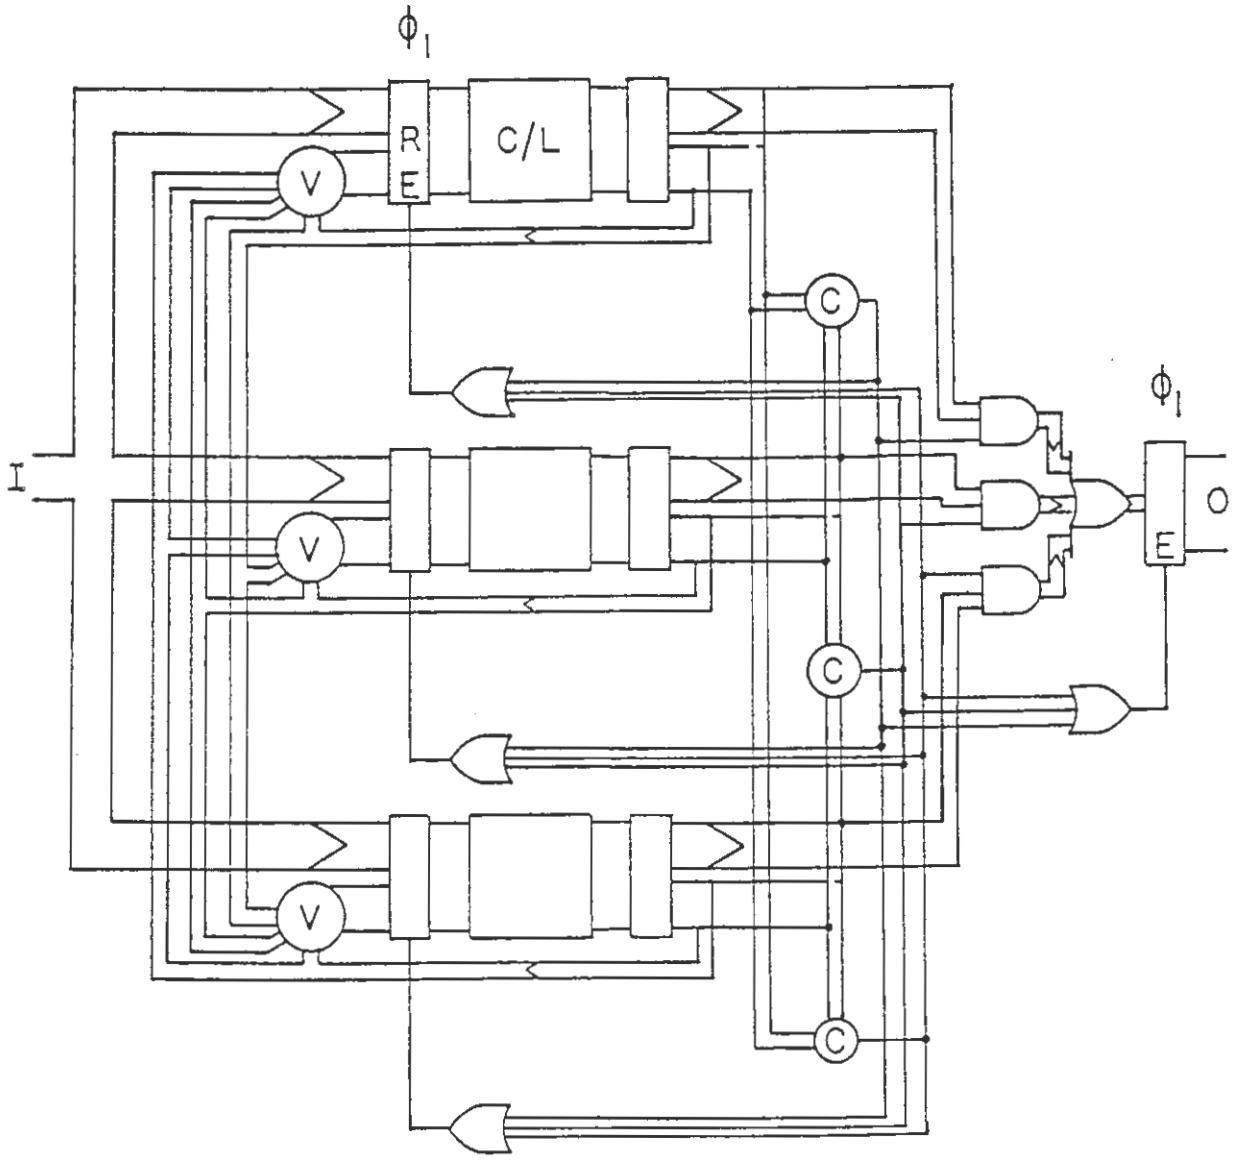

Figure 3.4 A tightly coupled TMR machine, hardened for tolerating bursts of transients. The comparators work on all outputs simultaneously. V: voter (bit-by-bit)

Chapter 4
Soft-Error Filtering

This chapter presents the Soft-Error Filtering (SEF) approach aimed at decreasing the soft error rate. SEF is a general design methodology intended to make machines tolerant to soft errors. The basic idea and the choice of a model are discussed in Section 4.1. A discussion of how the SEF approach is rooted in the fundamentals of communication theory follows in Section 4.2. An analogy is drawn between a digital machine subject to transient errors and a digital communication channel corrupted by noise. This serves as a useful guide to finding means of improving the reliability of a digital system at a modest cost.

An analysis of the error rate due to radioactive decay for a SEF machine is developed in Section 4.3. This analysis demonstrates that a SEF machine can have a negligibly small soft error rate when bombarded by alpha particles due to radioactive decay. Section 4.4 shows how a variable hit rate could increase the error rate of a tolerant machine by orders of magnitude for a given average hit rate. However, even though cosmic rays produce a variable hit rate, it is shown in Section 4.5 that the error rate for a SEF machine can be neglected. A notable exception is the case where a dense but imperfect shield is used in close proximity to the machine. The error rate of a SEF machine due to interference is discussed in Section 4.6. Section 4.7 discusses the applicability of SEF for solving the soft error problem in general.

4.1 Basic Model

A widely applicable model for digital machines is needed in order to develop a general method of tolerating soft errors. The finite-state-machine, shown in Fig. 4.1(a), is such a simple model which generalizes easily to a wide variety of digital machines. A two-phase clock is used and the first clock, CK1, stores the present state of the machine in the left-hand register R. The combinational network, C/L, computes the output, OUT, and the next state, ST, on the basis of the contents of this register. During the second phase CK2 transfers the output of C/L, θ, to the right-hand output register R. The state outputs, ST, of this register, as well as the primary inputs, IN, are stored in the left-hand register by CK1. Figure 4.1(c) illustrates the clocking scheme. The results derived for this model can be easily extended to more complex register-transfer machines, such as a pipelined computer or a multiphase machine.

In reference to Fig. 4.1(a), a soft error is a non-recurrent and temporary difference between the actual behavior and the specification, as observed on one of the output OUT or state ST lines. Since a transient can also be injected directly into the registers, the latches composing them must be intrinsically tolerant in order to mask the transient. In this Chapter, it will be assumed that such tolerant latches can be fabricated. The design of these latches is discussed later in Chapter 5.

Assuming intrinsic tolerance of the latches, a soft error can only result from a transient injected into the combinational logic section, as

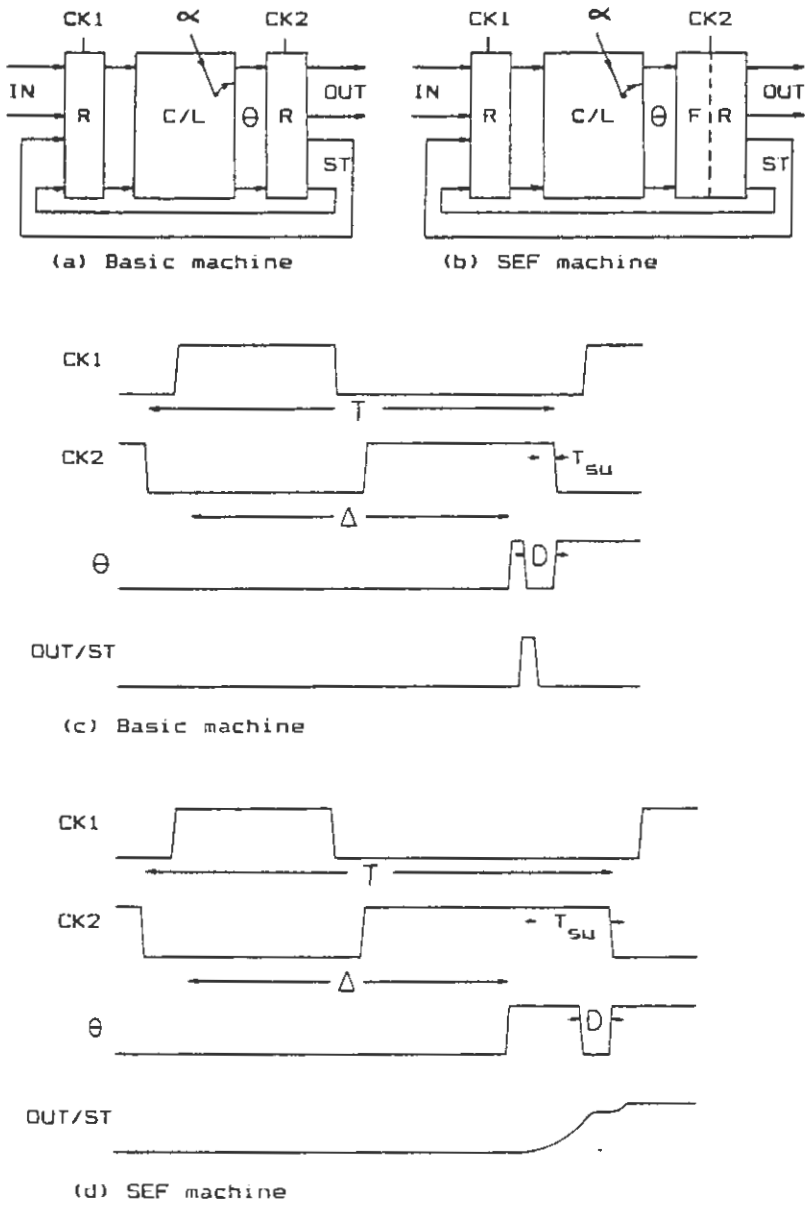

Figure 4.1 A conventional finite-state machine based on a two-phase clock, CK1, CK2, and consisting of two registers, R, and a combinational logic block, C/L. Only C/L is assumed to be sensitive to a hit by an alpha particle, α. (b) A functionally-equivalent SEF machine in which the outputs, θ, of C/L pass thru filters, F, before being latched in the output register. (c) Timing diagram for the machine in (a), showing the effect of a θ=1 being corrupted by an alpha-induced transient of duration D during the register set-up time T_{su}. (d) Same as (c) for the SEF machine in (b). Note that T_{su} is longer than in (c) necessitating a longer duration of CK2.

illustrated in Fig. 4.1(a) and (c). The transient must appear on one of the output lines of the combinational logic during the interval when the latch is sensitive, and its duration must exceed a certain minimum value related to the set-up time of the latches, T_{su}.

If most of the soft errors are due to short transients, an important reduction of the soft error rate would be achieved by filtering these transients. One way that this can be achieved is by making all the nodes inside the machine slower. However, considerations of efficiency, both in terms of speed and area, suggest that the number of nodes which have to behave like filters should be minimized.

As long as a logic network is combinational, an injected transient remains a transient after propagation. If a boundary encloses only combinational functions, the effect of all nodes behaving as filters can be achieved by filtering all the lines fanning out of this boundary. For any machine, if a complete system is to be filtered, the places that yield the smallest number of filters are at the output of the combinational function. Such a SEF machine with filters (F) between every output of the combinational logic and the output register is shown in Fig. 4.1(b). A timing diagram for this SEF machine is shown in Fig. 4.1(d). Notice that the timing diagrams in Figs. 4.1(c) and (d) are very similar. The propagation time from the input of the first register to the output of the logic, Δ, is the same for both machines. Only the set-up time T_{su} of the SEF machine is longer. The disturbing transient on one of the θ lines has the same duration D in both cases. However, for the SEF machine, the inertia accumulated in the filtering register is sufficient to tolerate the effect of the injected transient. The filters impose a slightly longer duration for the phase 2 clock.

A simple extension of this idea permits an improvement of the machine, when a region is known to be relatively noisier. If this region can be enclosed by a boundary, with a small number of outgoing lines, the hardware overhead could be reduced by adding filters inside the combinational network. This extension of the design methodology is reasonably obvious and will not be treated explicitly.

In this thesis the expression set-up time is not used with its conventional meaning. The set-up time T_{su} of a register reflects the time interval during which data is latched. Hence T_{su} is normally defined as the time interval during which input data must not change, and is usually measured with respect to one of the clock edges. In the present case, this definition is relaxed to permit the momentary corruption of input data by a short duration pulse. Therefore T_{su} can be used as a measure of the register's tolerance to soft errors. For convenience of analysis, the registers are assumed to be level-sensitive (as against edge-triggered) without any loss of generality. It turns out that all types of latches or flip-flops have a minimum time interval during which data should be stable in order to function properly.

To summarize, SEF consists of transforming a basic machine by replacing its memory elements with filtering latches. The combinational logic network is not replicated, which reduces the hardware overhead. Since SEF uses logic elements as fast as in the basic machine, the performance of the machine is minimally affected by the inclusion of filters at the input of the latches. The difference between SEF and using a slow technology is that, in the former a single slow node is included in every propagation path, whereas

in a slow technology every node is slow.

4.2 A Parallel With Communication Systems

There are strong similarities between a digital circuit sensitive to soft errors and a communication system. The input and output registers, designed to be noise-free, are analogous to the transmitter and receiver. The combinational logic circuit disturbed by injected transients plays the same role as the communication channel disturbed by noise. This similarity suggests that the techniques developed for optimizing communication systems may furnish guidance for techniques that may be used to provide reliable logic circuits. In particular, an approach which relies on filtering, in the output register, will be considered. The additional overhead that is required is analogous to increasing the signal power in a communication channel.

There have been previous efforts to apply ideas explicitly from communication systems to computational systems. In general the idea is to add redundancy to combat the effects of failed components or wiring defects in a computational system. The overhead in combinational circuits, required by an error-correcting code, is analogous to the increased bandwidth required for transmitting a fixed amount of information at a given rate in a communication system.

Earlier work along these lines was unified and extended by Winograd and Cowan [WIN63]. In their work, as in much of the work done around that period (1963), the focus was upon the channel capacity concept of the

classical information theory, and upon using error-correcting codes to try to achieve that capacity. With the same mathematical ideas as in classical information theory, namely entropy and equivocation, a concept analog to channel capacity, called computation capacity, is defined. These authors have shown that, as long as the automata are composed of modules with positive computation capacity, they can be constructed with arbitrarily high reliability, apart from errors in the output circuits. A systematic way of using error-correcting codes in order to add redundancy was also developed.

Unfortunately, the above theory is not useful for solving the problem considered in this thesis, since it assumes very unreliable elements and, therefore, results in very high overhead for achieving a reasonable system error rate. As was discussed in Chapter 2, the basic elements are highly reliable, therefore very simple error-correcting codes such as majority voting, presented in Chapter 3, are generally sufficient for the reliability improvement required. It is of interest that all the standard techniques for masking errors are applications of coding.

The new approach proposed here is based on another analogy to communication theory. In contrast to the coding approach, SEF is analogous to optimum filtering. A fundamental difference between the two is that SEF leaves the machine with the same combinational logic network, thus avoiding a significant fraction of the hardware redundancy.

The filtering technique is based on the consideration that the time overhead and the effective energy in the signal are both proportional to the set-up time of the latches composing the output register. The minimum clock period is determined by the sum of the worst case propagation delay in the

combinational logic circuit plus the set-up time of the registers. If the delay in the logic is much larger than the minimum set-up time of the registers, which is normally the case, then there is a possibility of a significant improvement in reliability without a large increase in time overhead. The relative weight of these factors is, of course, technology dependent.

The similarity between the problems of designing a SEF machine and a reliable communication link is interesting, because it will be shown in the sequel that the error rate for a SEF machine decreases exponentially with the energy in the signal, in a manner similar to a communication link with a properly designed receiver. The design of the filtering register is similar to the design of the said proper receiver and is covered in Chapter 5. It is of interest that the proposed filtering register, which produces a binary output, can be viewed as a device that takes the majority of M samples in the analog domain. Since many bits are manipulated by the same physical device in the analog domain, such a filtering register gives a relatively compact realization of the majority function.

The idea that a signal corrupted by a noise event can carry information in a reliable way is well established in the communication field. Nevertheless, as yet, no one seems to have exploited the idea that the output of an uncoded digital machine can be recovered reliably, even though the machine has been corrupted by a noise event. This idea is the fundamental reason why it is possible to mask errors with less than a 100% overhead, both in hardware and in time, simultaneously.

4.3 Products of Radioactive Decay; Error Rate Improvement With SEF

The calculation presented in this section is based on several facts and a single key assumption. Firstly, filtering registers are feasible, this will be established in Chapter 5. Secondly, radioactive decay produces a flux of particles with a Poisson distribution. Therefore, intervals between hits are independent and exponentially distributed. Another important fact concerns the form of the transient induced by a hit. A hit may produce a transient composed of more than one pulse, because it may affect more than one node, and more than one sensitized path may exist. However, it is assumed that the sum of the durations of the individual pulses is bounded by some value, P, as shown in Fig. 4.2. This assumption permits a very important simplification of the formulation. The worst case occurs when the transient consists of a single pulse of duration P.

As will be shown in Chapter 5, a single hit will not cause an error, if the set-up time T_{su} of a latch is sufficiently longer than P. Accordingly, in order to provide immunity to soft errors, registers are constructed with a set-up time longer than the minimum possible for a given technology. The SEF machines built from such registers could equivalently be called Set-Up-Time-Redundant (SUTR) [SAV84a].

The necessary condition for an error to occur is, therefore, that two or more independent hits happen in the same clock period and with a proper timing relationship. In order to evaluate the tolerance of a SEF machine, bounds on the error probability are calculated for a machine with and without

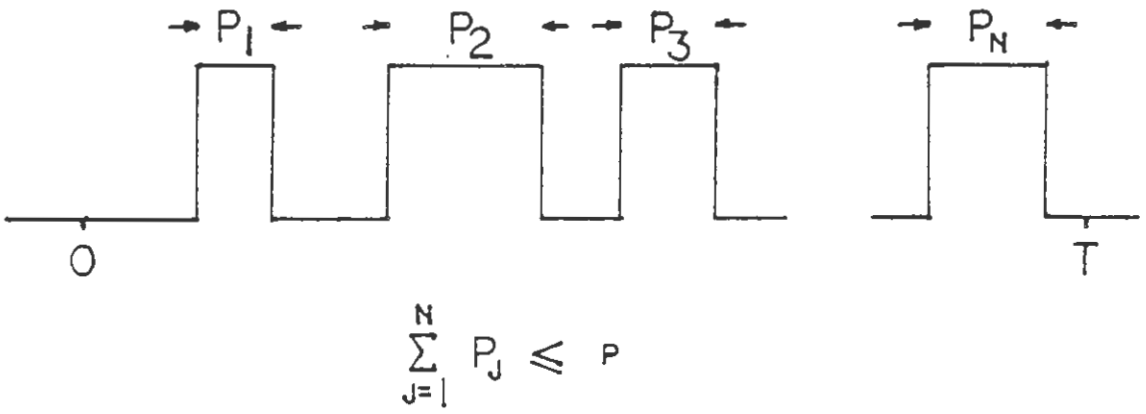

Figure 4.2 A transient composed of N pulses

SEF. The bounds to be derived are loose, but, nevertheless, sufficient to demonstrate that SEF does effectively reduce the error rate to insignificant levels. How to calculate tighter bounds for a sensitive machine has been discussed in Chapter 2. The ideas used in Chapter 2 could to some extent be applied to the error rate calculation for tolerant machines, but the conclusion obtained with the simple bounds would not change, and therefore the computational effort required would not be justified.

4.3.1 Error Rate Analysis

The analysis begins by considering the error rate for a single output line. Let D be the maximum duration of a tolerable pulse, i.e., the duration of a pulse that is guaranteed not to be latched at the output of the logic. The minimum number of hits that is required to cause an error is given by $\eta=[D/P]+$, where $[X]+$ is the smallest integer larger than X. For an error to occur, the pulses must fall in the sensitive time interval of the output register, in such a way that the line is at the wrong value for a duration longer than D. Consequently, Pr(error and n hits) < Pr(n hits), because not all hits propagate to a primary output. Also, the duration of the transients is generally smaller than the bound P, and the pulses may overlap, thereby resulting in a composite pulse shorter than the sum of the durations of the individual pulses. From these considerations, the probability of error per clock cycle per output line can be bounded by

$$\Pr(\text{error}) < \sum_{I=\eta}^{\infty} \Pr(I \text{ hits}) \qquad (4.1)$$

Since radioactive decay has a Poisson hit rate, the probability of exactly I hits as a function of the average effective hit rate N is given by

$$\Pr(I \text{ hits}) = \frac{e^{-N} N^{I}}{I!} \qquad (4.2)$$

where N can be computed in a way similar to the error rate in Chapter 2.

The following discussion explains how (4.3), which is an expression for N as a function of the basic parameters of a machine, is derived. In particular, this discussion emphasizes the assumptions and approximations involved.

First of all, only the hits that can potentially disturb the machine are counted. F is the effective hit rate in hits/cm^2-hr. The sensitive area in the cone of the considered output is designated by A_{sr} in cm^2. A cone is the set of all the nodes for which a path exists to a given output.

If T is the clock period in seconds, then the average number of hits per clock period which have the potential to cause problems is $FA_{sr}T/3600$. Moreover, at a given moment, only a portion of these hits will generate a pulse that falls within a time slot corresponding exactly to one propagation delay prior to the sensitive time slot of the register. It has been assumed that the maximum pulse duration is P.

The set-up time of the register should be related to D, the pulse duration which is tolerated. A constant S, called the *security margin* is introduced here. This constant reflects the ability of a register to filter out spurious events. By definition of S and D, their product gives the set-up time of the register. Therefore, the set-up time is not sufficient by itself to model the tolerance of various SEF or conventional machines, because the

value of S to be optimized for SEF machines, can be significantly smaller than that of a standard latch.

When a single event affects the machine, the worst case is in general a single transient pulse on the output of the combinational logic with a duration P. It is a worst case because it is the most difficult situation to filter. With the pessimistic assumption that a transient is always composed of a single pulse of duration P, the fraction of the hits that can potentially disturb the machine is given by (P+SD)/T.

The numerator of this factor overestimates the sensitive period of the SEF machine in a conservative manner. This corresponds to the assumption that, if the transient composed of a single pulse overlaps the set-up time interval, then the register is affected as if the pulse lies completely inside the set-up time.

Also it implicitly takes into account the case where P becomes significantly longer than SD, which is important if this development is to hold for a machine that does not have filtering registers. When a machine has registers which are not filtering registers, their set-up time may become extremely short, but the sensitive period does not decrease to zero. In this case the sensitive period is determined by the duration of the disturbing transient. The difference is particularly significant if the expected pulses are long.

And finally, it is also necessary to add P to the sensitive period in the numerator of that ratio, when there is more than one hit during a given cycle. In particular, if a first pulse of duration $P=D-\in$, where $\in$ is greater

than 0 and small, falls completely within the set-up time, any overlap between this set-up time and a second pulse results in an error.

Notice that increasing D has two opposite effects. The first and desired effect is a reduction of the fraction of the single hit events which can cause an error by themselves. Unfortunately, it also has the unwanted effect of increasing the fraction of events that are potentially harmful by overlapping with the set-up time interval. The first effect dominates because it increases the minimum value of η in (4.1), which becomes the exponent in (4.2), whereas the second effect only results in a linear increase of the effective hit rate, which is fairly low in any case.

A simple expression for the effective hit rate is obtained by multiplying the expected number of hits per clock period, with the fraction of the time where each node is sensitive, yielding

$$N = \frac{FA_{sr}(P+SD)}{3600} \qquad (4.3)$$

There is an approximation in (4.3), which is associated with assuming that a single pulse of duration P is always the worst case. Consider, for example, the situation illustrated in Fig. 4.3(a), where a first hit results in a single pulse of duration P that falls completely inside the set-up time, and $P=D-\in$. In this case, the worst situation for a second hit would be a transient composed of $M=P/\in'$ pulses evenly distributed with $\in'>\in$ as shown in Fig. 4.3(b). For this transient, the sensitive slot is the entire clock period if $\in$ is small enough. The error rate contributed by this sequence of events is significantly underestimated.

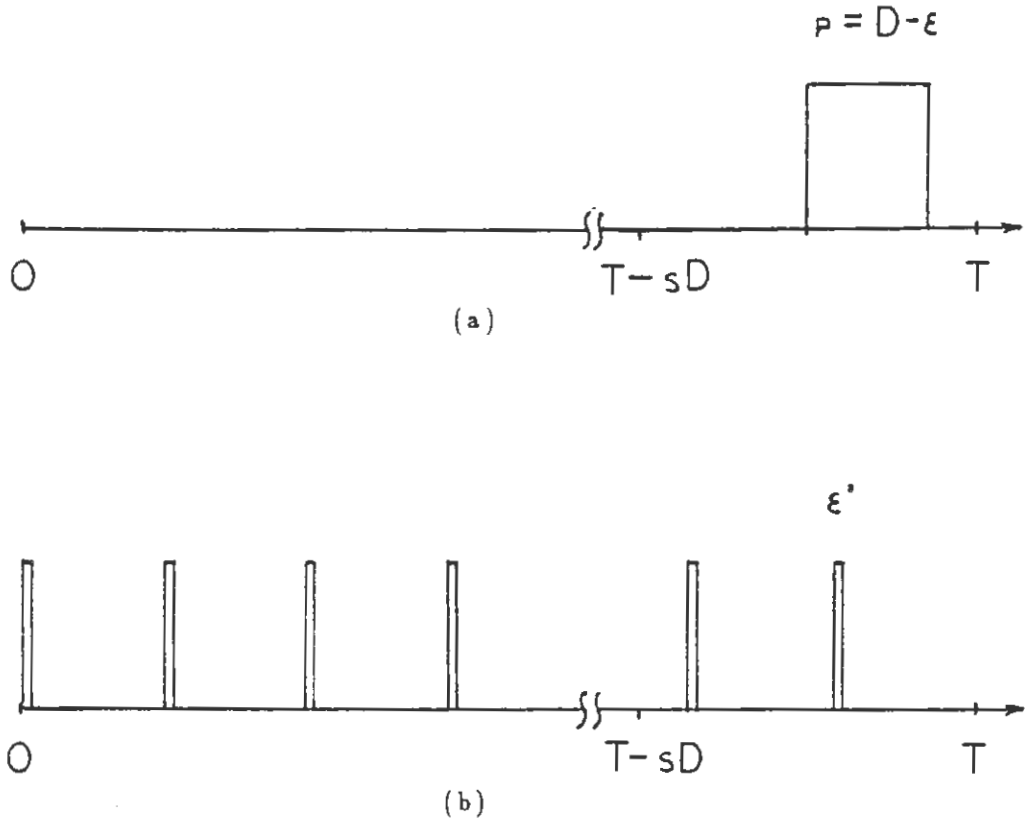

Figure 4.3 (a) A transient of duration D-∈ in the set-up time interval of the machine. (b) A transient formed by m pulses of duration ∈' (∈'>∈). The sensitive period for this second transient is the whole clock period.

The equation could be modified accordingly, but, since for most of the events, the contribution to the error rate is grossly overestimated, (4.2) is a loose upper bound except for some exceptional situations. Therefore, a simpler expression is prefered to a more complex and looser bound. There are two reasons why, in general, (4.2) overestimates significantly the error rate of a tolerant machine. Firstly, P is an upper bound on the duration of a transient that results from a single hit, and the total duration of the joint transient that results from two hits (defined as in Fig. 4.2), may not be sufficient to cause an error. Secondly, if the two individual transients are fragmented, the duration of the joint transient may be larger than D, but no window of duration SD includes a transient of duration D.

In the following discussion, the emphasis will be on simplifying the expressions. However the calculations could also be done with the more complex expressions. The simpler expressions give a better intuitive feeling for the tradeoffs involved with negligible effects on accuracy.

The effective hit rate N is of the same order of magnitude as the error rate calculated for a non-redundant machine in Chapter 2. This will become clear later when the general expression, valid for both redundant and non-redundant machines, will be simplified for the latter. The error rate calculated in Chapter 2 was very small; therefore, the contribution to the error rate of more than i hits can be neglected by comparison to the probability of i hits. Consequently the error rate is determined by the first term of the summation in (4.1). This term corresponds to the minimum number of events sufficient to cause an error. Notice also that the factor e^{-N} in (4.2) can be replaced by 1, without loosening the upper bound significantly.

The bound for the error probability per line can be rewritten as

$$\Pr(\text{error per line}) < \frac{1}{\eta!}\left(\frac{F\,A_{sr}\,(P+SD)}{3600}\right)^{\eta} \qquad (4.4)$$

If there are R outputs to a chip, then the error rate can be bounded as follows:

$$\text{Soft error prob. per cycle for the chip} < \sum_{j=1}^{R}\left\{\frac{1}{\eta!}\left(\frac{FA_{sr,j}\,(P+SD)}{3600}\right)^{\eta}\right\} \qquad (4.5)$$

where the subscript j refers to the region in the cone of the j^{th} output.

This is a union bound since the occurrence of errors on different outputs is not disjoint. Notice that the noisiest output line tends to determine the error probability when $\eta>1$, because in this case, the exponent amplifies the relative differences in the average hit rate. In general, a first approximation of the bound on the error rate can be computed by assuming that all the lines are as noisy as the noisiest one, in which case the summation is replaced by a multiplication by R.

When the machine is sensitive to a single hit ($\eta=1$) the expression can be simplified as follows:

$$\text{Soft error prob.} < \left(\frac{F\,A_{s}\,(P+SD)}{3600}\right) \qquad (4.6)$$

where A_s = total sensitive area of the chip.

4.3.2 Discussion

The simplified expression for a non-tolerant machine in (4.6) is very similar to the upper bound of (2.2) derived in Chapter 2. Therefore, the expression for tolerant machines is consistent with more accurate bounds developed for non-tolerant machines. A first important difference is the visibility, which is more difficult to introduce for tolerant machines. Also the formulas derived in this chapter do not distinguish the different sensitivities of the nodes. This difference follows from the fact that the error rate is derived with respect to each individual node in Chapter 2, with a potentially better accuracy, whereas, for simplicity, in the case of a tolerant machine, it is derived as a function of the sensitivity of each output line.

Notice that a slow technology could have D>P, even with no explicit utilization of SEF. In such a case (4.4) would apply for calculating the error rate of a non-modified machine, and this technology would be intrinsically tolerant.

At this point, it is appropriate to comment on the implications of the fact that the error probability per cycle is independent of the clock period. By comparison, the average number of failures per cycle is proportional to T/MTBF, where MTBF is the mean time between failures, and this ratio decreases when the device is operated faster. This contributes to making soft errors relatively more important in a scaled technology. Another implication is that a computation performed on a pipelined machine suffers from a relative increase in error probability, which is directly proportional to the improvement in performance due to pipelining. The reason is that the

error probability increases with the number of times a signal must be sampled, which is a corollary of the independence of the error probability with respect to the clock period for a two-phase machine.

The error probability of a SEF machine in (4.4) can be rewritten in the form

$$Pr(error) < \frac{1}{\eta!} \; e^{\eta \ln N} \qquad (4.7)$$

Where ln(N) is typically smaller than -30 and η=[D/P]+ should be limited to a small value such as 2 or 3. Remember that η is equivalent to a signal to noise ratio, since both the signal and the transient have the same amplitude. Therefore the division by $\eta!$ can be neglected, because it is not the dominant term, and the remaining expression for the bound on the error rate is an exponential function of the energy in a bit. Unlike the case of Gaussian noise in a communication system, this exponent has a staircase behavior. This is due to the quantum nature of the transients injected by ionizing radiation.

The approximations in the development of the error rate usually hold as long as the hit rate is sufficiently low. Also, when the hit rate is sufficiently low, the error rate can be made as low as required by increasing the latch set-up time. When high radiation levels are expected, it could become necessary to tolerate two or more events in order to obtain the required reliability. Usually it is sufficient to tolerate one event, which corresponds to η=2. This can be realized by making the longest tolerated event D slightly longer than P, the longest expected event. This is a situation, where an incremental change in D, gives a large improvement in the

error rate.

4.3.3 A Numerical Example

The same numerical example developed in chapter 2 can be used here to illustrate the reduction in error rate for a SEF machine in the case of radioactive decay. Three other parameters must be assumed to calculate the error probability per cycle with (4.5). Let R, the number of output lines be 20, and assume that 20% of the total sensitive area of the chip is included in each region, and also let SD the set-up time be 2ns. With P=1ns and A_s=0.2cm^2, (4.5) yields an error probability of $1.1*10^{-28}$/cycle. The calculated error probability was $1.1*10^{-15}$/cycle in Chapter 2, for the equivalent sensitive machine, taking the visibility into account. With the same 40ns clock period, the error rate for the tolerant machine is 10^{-17}/hour, and for the sensitive machine it is 10^{-4}/hour. Remembering that the failure rate is in the range of 10^{-6}/hour to 10^{-8}/hour, it is clear that the error rate can be neglected for a machine with η=2. Therefore, this SEF machine is tolerant to transients induced by radioactive decay. On the basis of these calculations, the error rate becomes so small by comparison to the significance threshold, that it is justified to trade accuracy for simplicity of the expressions.

It may seem surprising that some regions are counted in the sensitive area of more than one output line, however it is correct, since the transients are tolerated independently on each output line, and a node can be part of the cone of more than one output line. As mentioned earlier, the events "error on line X" and "error on line Y" are not disjoint, which is the reason why (4.5) is a union bound on the error rate.

4.4 Effect on the Error Rate of a Variable Hit Rate

The analysis of error rate presented earlier is generalized here, under the assumption that a particular source can be modeled as a Poisson process with a variable hit rate [PAP65 p.286]. It is shown that two sources with the same average hit rate may have very different error rates.

It was shown earlier in (4.2) and in the subsequent discussion that, for a given hit rate per cycle N, the error probability per output line and per cycle is bounded by $N^{\eta}/\eta!$. Here, η is the minimum number of events required to cause an error. The average error rate is bounded by the time average of this quantity, and if ergodicity is assumed it can also be expressed by

$$\text{Average error rate} < E\left[\frac{N^{\eta}}{\eta!}\right] \qquad (4.7)$$

With (4.7), one can study the effects of very noisy periods occuring with a low probability. It is shown here that some hit rate distributions have a minor effect on the average hit rate, but increase the error rate significantly. For example, assuming a machine that can tolerate all single events, if the distribution is discrete with probabilities p_i:

$p_1 = 0.999 \qquad N_1 = 1$

$p_2 = 0.001 \qquad N_2 = 1000$

then

$E[N] = 1.999$, and Error rate < 500.5

whereas if

$N = 2$, then Error rate < 2

Thus, if for a given average hit rate, there is a very noisy period with a low probability, then significant differences exist for the error rate. On the other hand, if the hit rate variations occur in the form of small deviations from the average, then only a minor variation of the error rate results from neglecting the deviations in the error rate calculation, as long as the proper average hit rate is used. The magnitude of the variation has to be compared with the precision with which the error rate is estimated. This is demonstrated by the following situation, where a relatively small deviation of the hit rate from its average, yields a small difference in the error rate by comparison to the earlier example with N=2:

$f(N) = 1/4$ for $0 < N < 4$

$E[N] = 2$ Error rate < 8

4.5 Significance of the Correlated Events Due to Cosmic Rays

The charge injection resulting from cosmic rays was discussed in Chapter 2. In particular, it was argued that the mesons are expected to inject a charge, which is one order of magnitude smaller than alpha particles. Cosmic rays also contain a small quantity of heavier nuclei, but, their average flux would be too small to result in a significant error rate

if their distribution had a constant hit rate. However, there is a property of the particle flux associated with cosmic rays that could be very detrimental to machines designed to tolerate a fixed number of hits in each cycle. This is the time and space correlation that exists between the secondary particles, generated by a single high-energy primary particle. Based on the assumption that a particular technology is sensitive to hits by these secondary particles, the significance of such a correlation is discussed here.

The correlation follows from the fact that a high-energy particle loses its energy in a cascade involving a large number of collisions. In the atmosphere, a number of these collisions generate particles with a significant range. Initially, the particles do not diverge much from the trajectory of the parent particle, but after subsequent collisions the secondary particles gradually scatter. A large number of secondary particles may reach the ground and almost simultaneously, however these particles are spread over some area.

The preceding corresponds to the situation analyzed in Section 4.4, where a short period with a hit rate much higher than the average could cause a significantly higher error rate. This is an important issue, since the amplification of the peak hit rate affects only the machines with built-in tolerance.

In Section 4.3.3, where the expected hit rate was constant, there was a very important difference between the error rate of a SEF machine and the failure rate. Therefore, the amplification of the error rate for a given average hit rate, may turn out to be insufficient to cause problems.

Consequently, further analysis is required in order to determine the significance of such a peaking phenomenon.

Before the significance of cascades of secondary particles can be discussed, their basic properties must be reviewed. This information is extracted from a book by Wolfendale [WOL63]. The cascade can be separated into three fractions: the electrons, the mu-mesons, and the nuclear component. At sea level, the electrons constitute the majority of the particles in the cascade, and the size of a shower is generally expressed in terms of this electron flux. A dense shower may contain 10^9 electrons spread over an area of approximately 10^4 m^2, which yields an average of 10^5 electrons/m^2. However, near the axis of a cascade, the density can be as much as 100 times higher, or 10^7 electrons/m^2. For reasons discussed in Chapter 2, digital electronic circuits are generally insensitive to incident electrons, and this dense flux of charged particles can safely be ignored.

The type of secondary particle affects both its range and its scattering angle, therefore, the composition of the cascade varies with the distance from its center. The mu-mesons form the majority of the particles at large distances from the axis of the cascade. Consequently, the mu-mesons can also be safely ignored, because they are scattered over such a large area, that the resulting low density of the particle flux results in a negligible amplification of the error rate due to correlation. This is demonstrated by calculating that the nuclear component, which is denser, does not result in a significant amplification of the error rate.

Near the center, where the cascade is denser, the heavier particles

represent a serious hazard even if they constitute less than 2% of the total particle flux. In a first approximation, the total particle flux and the electron flux can be equated. Therefore, based on the electron flux mentioned earlier, the average number of hits per square cm at the center of the cascade is 20.

These particles are penetrating, in the sense that their spectrum of energy is wide, and a fraction can penetrate a shield as thick as 20 cm of lead. Of these penetrating particles, at any point along the cascade trajectory, only the small fraction that terminates its range in an active region can cause an error. A rough estimate of that fraction is $2*10^{-5}$, assuming that the particle must stop in a layer of 20 μm from the chip surface to cause an error. Therefore the effective nuclear flux in a dense cascade is on the order of $4*10^{-4}$ particles/cm^2. Assuming that the flux of particle in the densest region is uniformly distributed, the probability of finding more than one hit in a small region is given by the Poisson distribution with this effective hit rate.

In the worst case, the difference between the arrival time of these events is small, and it is assumed that all the hits occur in the same machine cycle. If the sensitive area in the cone of every output line is smaller than 0.1 cm^2, then the maximum expected number of hits in each region is $4*10^{-5}$. If all single hits are tolerated, at least two hits are required to observe an error, and the conditional probability of an error, given a dense shower, is then $8*10^{-10}$. This assumes that two hits always result in an error, which is pessimistic. It is of interest that the conditional probability of an error, given a dense cascade, can be made as small as desired by decreasing the maximum sensitive area in the cone of any

output line. For a machine that tolerates all single hits, reducing the sensitive area by a factor K decreases the error probability by a factor K^2.

If only mesons and nuclear particles with just the proper energy to stop in an active region are considered to be harmful, then it can be shown that the probability of a harmful hit from cosmic radiation is smaller than that of an alpha particle hit generated by the products of radioactive decay. Since cosmic radiation includes all cascades, the probability of a hit form a large cascade is even smaller. The error probability due to cascades is given by the product of the probability of observing a cascade times the conditional probability of an error given a cascade, and the conditional probability is small as demonstrated earlier. Therefore, in conclusion, the correlated particle flux in the cascades should not make the error probability significant, unless the sensitive area for a given line is very large.

However an important case exists where this result does not hold: when a shield of dense material is used, and this shield is not thick enough to completely absorb the nuclear cascade. When a particle travels in a dense material, the distance required for this particle to experience a certain number of collision is very small when compared with that required in the atmosphere. For the same number of collisions in a solid, a shower of secondary particles is generated with similar scattering angles as in the atmosphere, but a very small propagation distance is available to spread the particles.

For example, if a shield of 200 g/cm^2 with a thickness of 18 cm

(lead) is used, a primary particle with an energy of 10^{15}eV yields 5000 secondary particles (protons or heavier) [WOL63 p.193], which are distributed over a small region. The area of this region can be estimated from the distribution of the scattering angles and the radiation length (see [WOL63] pp.24-29). The majority of the secondary particles will fall in an area of 1 cm^2. Again, in this case, the density of the particle flux in the cascade may be as much as 100 times higher near its center, and if this flux is multiplied by the fraction that will stop in the active region calculated earlier, the effective hit rate at the center of the cascade is on the order of 10 hit/cm^2. If the cone area is .04 cm^2, as for the example in Section 4.3, then a cascade generates two hits or more in a given region with a probability of 0.062. The numerical example chosen is particularly significant, since for a shield in the vicinity of 200g/cm^2, the flux of secondaries contains many more protons and heavier particles than mesons, and the heavier particles are more efficient than mesons at causing errors.

It should be clear from Section 4.3.1 that two hits in the same cone, during the same machine cycle, do not necessarily result in an error. However, this calculation demonstrates that for a machine capable of tolerating a limited number of hits, a dense shield can significantly reduce the reliability improvement obtained with any fault-tolerance approach based on independence of the events. If, in theory, the formalism of (4.7) could be applied to (4.5), it will not be useful in practice, unless distributions of primaries and secondaries are characterized in detail. This topic is left for further work.

As a final comment with respect to cosmic radiation, if a machine were to be used in space, the shielding and scattering effects of the

atmosphere are absent. The primary flux of particles is known to contain many protons and heavier nuclei. This flux is significantly higher than the flux of alpha particles generated by the packaging material of a chip. If a technology is not intrinsically tolerant to ionizing radiation, the use of fault-tolerance is easy to justify. Moreover, in such a situation, tolerance to multiple hits may be necessary.

4.6 Effectiveness of SEF to Combat Interference

For machines in operation today, interference is the major source of soft errors (in the logic). There are many different sources of interference with very different characteristics. Moreover, the underlying sources of interference are often deterministic. If the model of the system were sufficiently complete, it would often be possible to predict the occurrence of errors. Therefore, it is very difficult to calculate the improvement in reliability obtained by using an approach like SEF. However, the significance of the problem remains, and the difficulty of quantifying the reduction in error rate does not mean that such an improvement is not possible.

The most important work on the characterization of the error rate due to interference has been done by McConnel [McC79,McC81]. The crash rate of several systems was measured, and a distribution for the interarrival time was obtained in this work. It is argued here that the crashes which were not due to a permanent fault were caused by interference. Since the machines used for this experiment were designed with bipolar or MOS technology available around 1976, this assumption is reasonable on the basis of the earlier

discussions on intrinsic tolerance to ionizing radiation. If not all the transient errors observed for these machines were due to interference, one can certainly argue that most of them were.

Any divergence from a Poisson error rate in the distribution of crashes, in McConnel's work, must be due to a different distribution for the underlying physical sources of errors. An important result of his work is that a significant divergence from Poisson distribution does exist. The observed crash rate was best described by a decreasing hazard rate distribution, and a Weibull distribution was fitted on the observed interarrival-time data. The Weibull distribution can be seen as a generalization of the exponential distribution, and the latter describes the interarrival time of a Poisson process. Since the Weibull distribution has a shape parameter that permits the adjustment of the mean and the variance independently, it is not surprising that a better fit can be achieved to the observed data. McConnel does not attempt to justify this distribution on the basis of the intrinsic properties of the physical sources of transient error. Therefore it is not necessarily the only distribution that can be used to fit the data.

The discussion in Chapter 2 shows that a large number of interference sources can contribute to the transient error problem. It is reasonable to postulate that these sources are independent. Therefore, if the sources were Poisson, the resulting error rate should be Poisson. It is clear from McConnel's results that the Poisson assumption is not correct. On the basis of the physical properties of the interference sources discussed in Chapter 2, one necessary property of the Poisson process is violated; namely, disjoint time intervals are not independent. In other words, the fact that an

error is observed is an indication that an unusual electromagnetic activity is taking place in the vicinity. Therefore, it is likely that another error will be observed within a delay shorter than the average. This differs from the variable Poisson hit rate because disjoint intervals are correlated.

For many random processes, a correlation exist for short adjacent time intervals, however, if the delay between intervals of similar duration increases, the correlation tends to decrease. This is a reason to believe that the distribution of interarrival times is probably very different for short time intervals, even though no data is available to confirm it. Therefore, the data collected by McConnel [McC79,McC81] are of limited interest for characterizing the error rate on a short time scale, because for practical reasons they have been truncated to interarrivals larger than 5 minutes. The daily variations of the error rate observed in [McC81] should have an insignificant effect, because the differences from the average are not very large, as discussed in Section 4.4.

Assuming that the non-Poisson behavior of the distribution is due to a correlation between events from a given source during its active periods, and observing that the resulting error rate per cycle is fairly low, one observes that the probability that two independent interference sources are active at the same clock period is very low. The lack of basic data and the deterministic nature of interference sources precludes a quantitative estimation of the error rate reduction with SEF. However, if all the interference sources can be tolerated when they are taken individually, an error will only occur if two independent sources are active during the same cycle. This event is possible but not probable; therefore, the resulting transient error rate can be neglected. The tolerance scheme will be defeated

easily if interference sources are perfectly synchronized.

Notice that SEF is not affected by a source with arrivals correlated from cycle to cycle. In fact, a high rate of transient events from an interference source can be tolerated more easily with SEF than a similar rate from a Poisson source, because the behavior of the first is more predictable. Similarly, when two sources are active in the same time window, they are again easily tolerated if they are slightly skewed by at least one machine cycle.

4.7 Discussion

In the case of ionizing radiation, SEF usually works well as long as the required operating frequency is sufficiently low. If the duration of the expected transients approaches the duration of the clock period then the overhead becomes excessive, as will be demonstrated later. Recall also that the determination of the maximum possible duration of the transients on the outputs of a VLSI chip is a difficult problem, even if a perfect knowledge of the maximum injection time is assumed. Therefore, because of the uncertainty on the maximum transient duration, a SEF machine can be built that is still sensitive to a single hit on a small fraction of its nodes. Notice that, even though in such a case the formulas derived for the error rate no longer hold, the machine can still be tolerant. The error rate would be determined by the remaining small sensitivity to single hits. An optimum design would be the one that makes the random failure rate slightly dominant. This would be achieved by tolerating events long enough to decrease the fraction of the area of the machine which is sensitive to a single hit, and also the

sensitive time period of each node.

SEF is efficient at dealing with interference sources injecting short transients. However, it is clear that SEF is not practical for interference sources which inject transients as long as or longer than the clock period. The question of time overhead is discussed quantitatively in Chapter 6, but it is already clear that tolerating an injected pulse of one microsecond or more generally leads to prohibitively high overhead, for a high speed machine. SEF would have to be used in conjunction with the standard techniques developed for dealing with electromagnetic interference, which are very efficient for reducing the coupling at low frequencies [MAR84].

Chapter 5
The Design of a Filtering Register

The SEF approach to the design of soft-error-tolerant digital systems is based on the assumption that it is possible to build a register which can filter out transients at its input. These transients can be significantly longer than the basic switching time of the technology used. Moreover, the space and time overheads associated with the registers must be acceptably low. Also, it is implicit in the SEF technique that the registers must be intrinsically tolerant to sources of soft errors such as direct alpha particle hits.

There exist a number of different approaches to the problem of designing an efficient filtering latch. To determine the most efficient, every one must be analyzed and compared. A simple but inefficient means of making a filtering latch, is to slow down a standard one. This approach is considered in Section 5.1. Another design which permits a significant reduction of the time overhead is considered in Section 5.2. This second latch is based on a single filter. It is first optimized at the functional level and then a transistor implementation is proposed. The steps necessary for redesigning a version of this latch with a different set of constraints are also outlined. Finally a double-filter configuration is analyzed and a practical realization is proposed in Section 5.3. The implementation of the double-filter configuration turns out to be simpler than the single-filter latch, and the design is more efficient than either of the other two.

In order to avoid assumptions about a future hypothetical sensitive technology, it was decided to use the Nothern Telecom 5μ CMOS1B process as a benchmark [CMC85]. This choice follows from the fact that it is easier to scale up injected transients and their effects, than to predict all the parameters describing accurately a scaled down technology. Moreover, the chosen process is available to universities, and therefore it becomes possible to fabricate these devices and experimentally confirm the results derived here.

The reference from which the scaling factors are derived is a hypothetical 0.5μ technology with a 1.5V supply. This represents a sensitive technology for which the appropriateness of the SEF approach is reasonably evident. Such a technology corresponds to first scaling CMOS1B at constant voltage by a factor of 3, which is representative of the state of the art, followed by a subsequent scaling at constant field by another factor of 3. Consequently, the basic time constant of the 5μ process is approximately 27 times longer than that of the 0.5μ process. The scaling by 3 at constant voltage contributes a factor of 9, and a factor of 3 results from the scaling at constant field. The saturation currents for transistors with equivalent length/width ratios is roughly the same for both technologies. The capacitances on the 5μ process are 9 times higher than those of the 0.5μ process, and the supply voltage is 3 times higher.

The duration of the injected current pulse is usually shorter than 0.25 ns as shown in Fig. 2.2. It was also demonstrated in Chapter 2, that the duration of the voltage transient after propagation, may be significantly longer than the injection time. Consequently, a reasonable objective for the filtering register in the 0.5μ process is that it must filter out all pulses

up to a duration of 1 ns. The 1 ns transient serves as a reference for testing and comparing the different approaches to the design of filtering latches. Therefore, after scaling up of this transient, the 5μ implementation must tolerate a transient of 27ns.

The second objective is intrinsic tolerance to a direct hit. The current pulse to be tolerated is shown in Fig. 2.2. After scaling up and some approximation for making it easier to simulate with SPICE, it gives the pulse in Fig. 5.1. The basic pulse in Fig. 2.2 corresponds to an injected charge of 82 fC, and the simulated pulse corresponds to a charge of 2.2 pC. Thus, the scaling factor for charge is 27, it can be interpreted as the same current with a duration 27 times longer, or as an equivalent charge on a capacitor 9 times larger when the supply is 3 times higher.

The above events do not necessarily represent the worst possible case but are relatively large disturbances. If the design is to be conservative but efficient, it is essential to know accurately these worst possible events for a given process. This chapter demonstrates how a latch can be designed for a given set of expected worst-case events. An efficient filtering latch must be optimized for the level of tolerance required.

5.1 Slow Latch

It is well known that a minimum amount of energy must be imparted to the input of a latch, for it to switch regeneratively to the opposite stable state. Thus, if constant amplitude pulses are injected into the latch's input, there will be a threshold to the pulse duration below which the

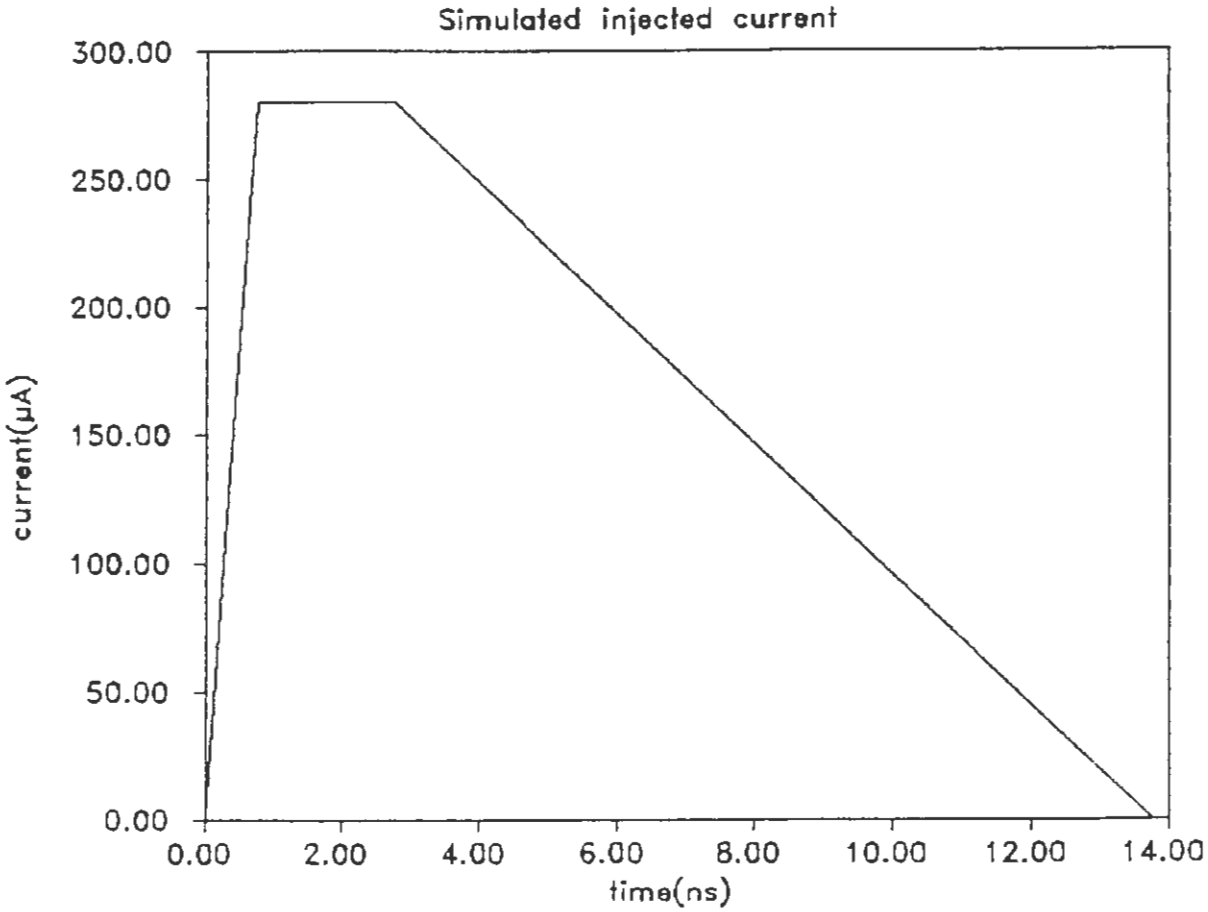

Figure 5.1 Approximation of the scaled up version of a current pulse injected by a direct alpha hit on a node.

content of the latch will remain unaffected. This observation leads to the conclusion that a slower latch can tolerate pulses that are proportionately longer.

Consider the standard D latch shown in Fig. 5.2(a), where all the transistors are of minimum size. One can modify it so that its state cannot change rapidly, this yields a slow latch. This can be achieved by loading nodes 7 and 8 with a capacitor C as shown in Fig. 5.2(b). The value of C is determined from the amplitude of the injected charge, when the latch experiences a direct hit on nodes 7 and 8.

An iterative procedure to determine the optimum value of C is to simulate it once with a capacitor which is known to be too large. Comparing the amplitude of the injected transient with the noise margin, it is possible to calculate a second value for C which is nearer to the minimum. One or two iterations should be sufficient, because the tolerance on the value of a capacitor is usually on the order of 20%, and thus in practice nothing can be improved.

The capacitor C does not have to be linear. This is important because many processes do not permit the fabrication of a linear capacitor. In the simulations discussed below, C was realized by the gates of two large MOS transistors in parallel. One of them is an N device, whereas the other is a P device. Both are 36 times larger than the corresponding minimum size transistor which were used in the rest of the latch. This gives some measure of the area overhead which is required for intrinsic tolerance.

Figure 5.3 shows some results of a simulated experiment with the slow

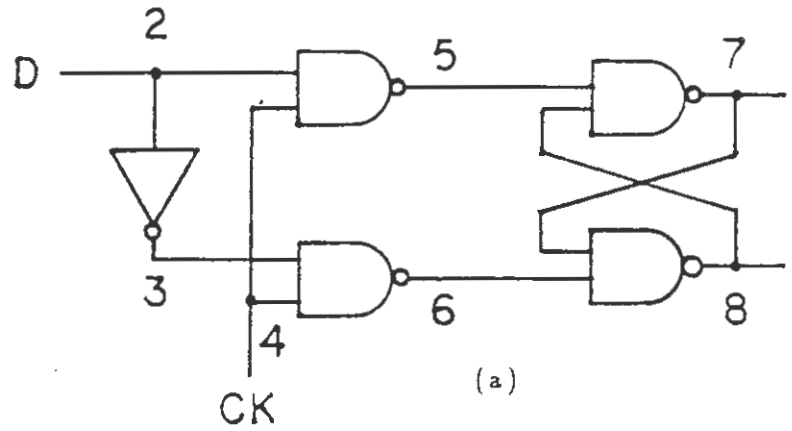

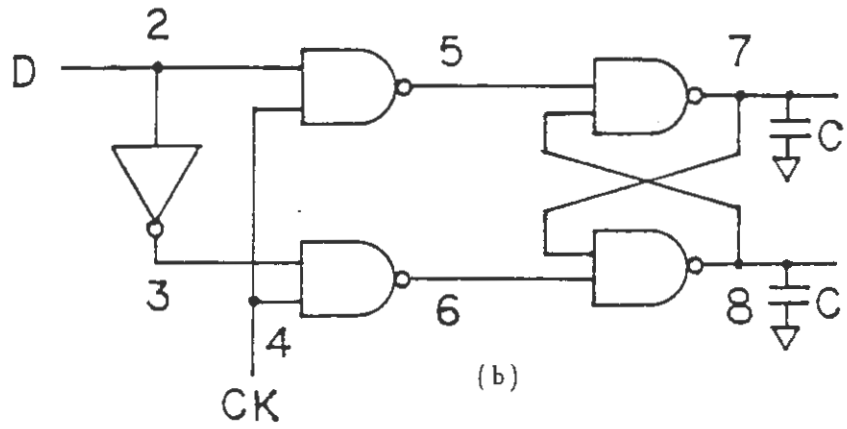

Figure 5.2 (a) A level-sensitive D latch (b) The same latch modified to have a slower response

latch, in which a 0 signal on its input is corrupted by a transient 1 having a duration of 27 ns. In Fig. 5.3(a) the clock pulse width is 91 ns. The figure shows the voltage on node 7 falling after the clock input is returned to zero and, indeed, a longer simulation confirms that node 7 settles to a 0 state. Hence, the slow latch recovers properly the input signal with a clock pulse of 91 ns. In Fig. 5.3(b), the clock pulse is only 2 ns shorter, but node 7 eventually reaches 5 V. Therefore the data is not recovered successfully with a clock pulse of 89 ns.

The filters in this paper are designed and compared using a performance measure called the *security margin*, S. This was defined in Section 4.3 through the relationship T_{su}=SD. In view of the relaxed definition of T_{su} introduced in Section 4.1, S represents the factor by which T_{su} must exceed D for a corrupting pulse of width D not to cause the latch to malfunction. Since the latch failed with a clock pulse duration of less than 91 ns, therefore, for the conditions depicted in Fig. 5.3, S=91/27= 3.37.

A small value of S is associated with a more efficient filtering latch. However, it should be clear from Fig. 5.3 that if the corrupting transient had occured earlier with respect to the clock's falling edge, a value of S less than 3.37 would have been obtained. This is indeed what simulation shows. Finding the worst time of occurence for a pulse of constant width D is an expensive iterative process. The conditions shown in Fig. 5.3 in fact are not far from the worst case.

The security margin does not reflect one important property of a particular filtering latch design, namely, the time it takes the output to

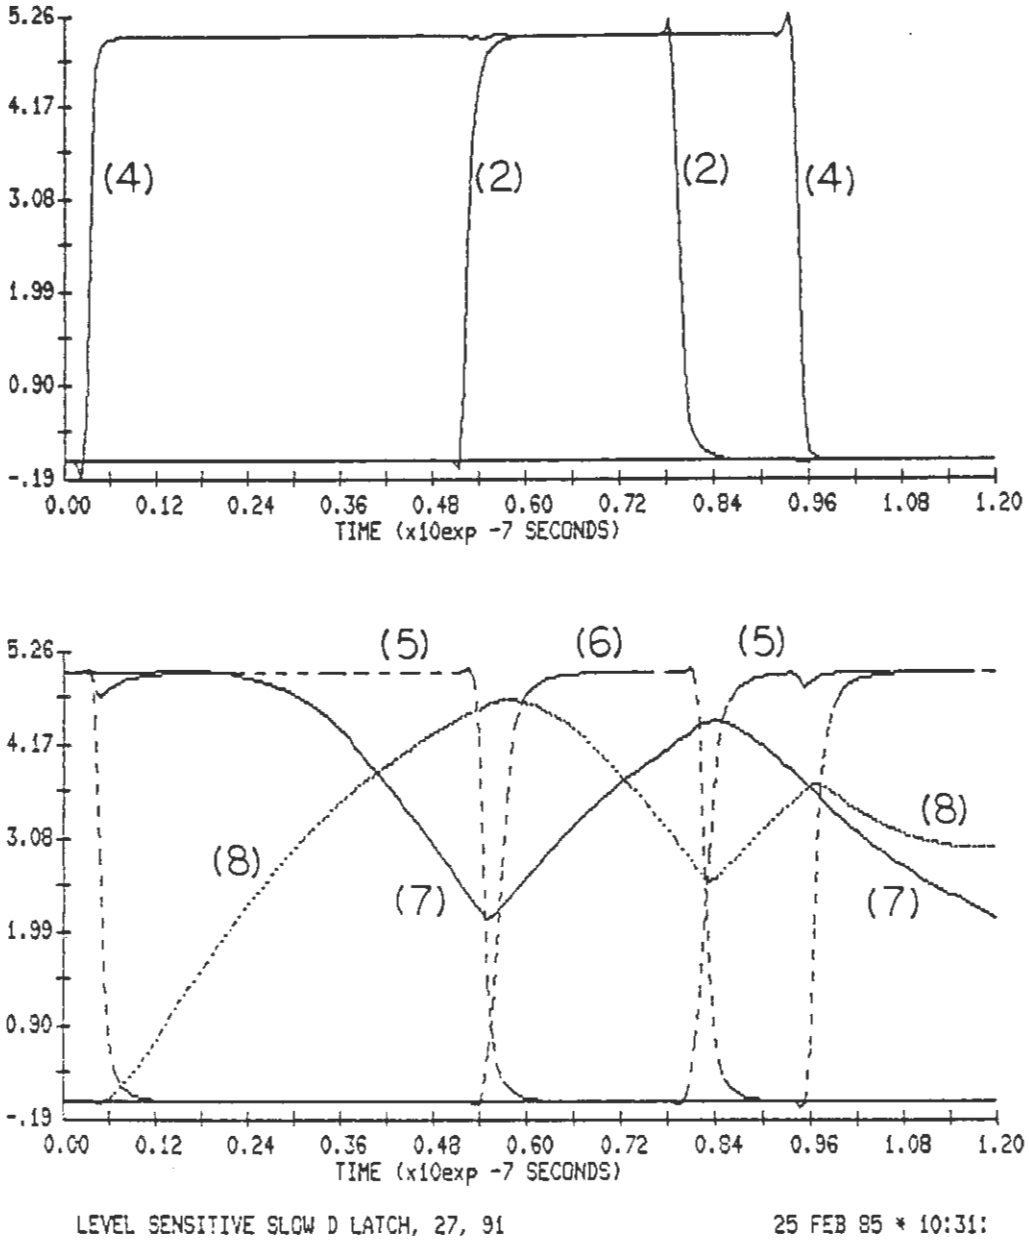

(a)

Figure 5.3 (a) Response of a slow latch (Fig. 5.2(b)) to a transient pulse of 27 ns for a clock signal lasting 91 ns. The input signal is a 0 corrupted by 1. The latch eventually recovers to the right output value. (b) Same as (a) with a clock pulse of 89 ns. The latch does not recover.

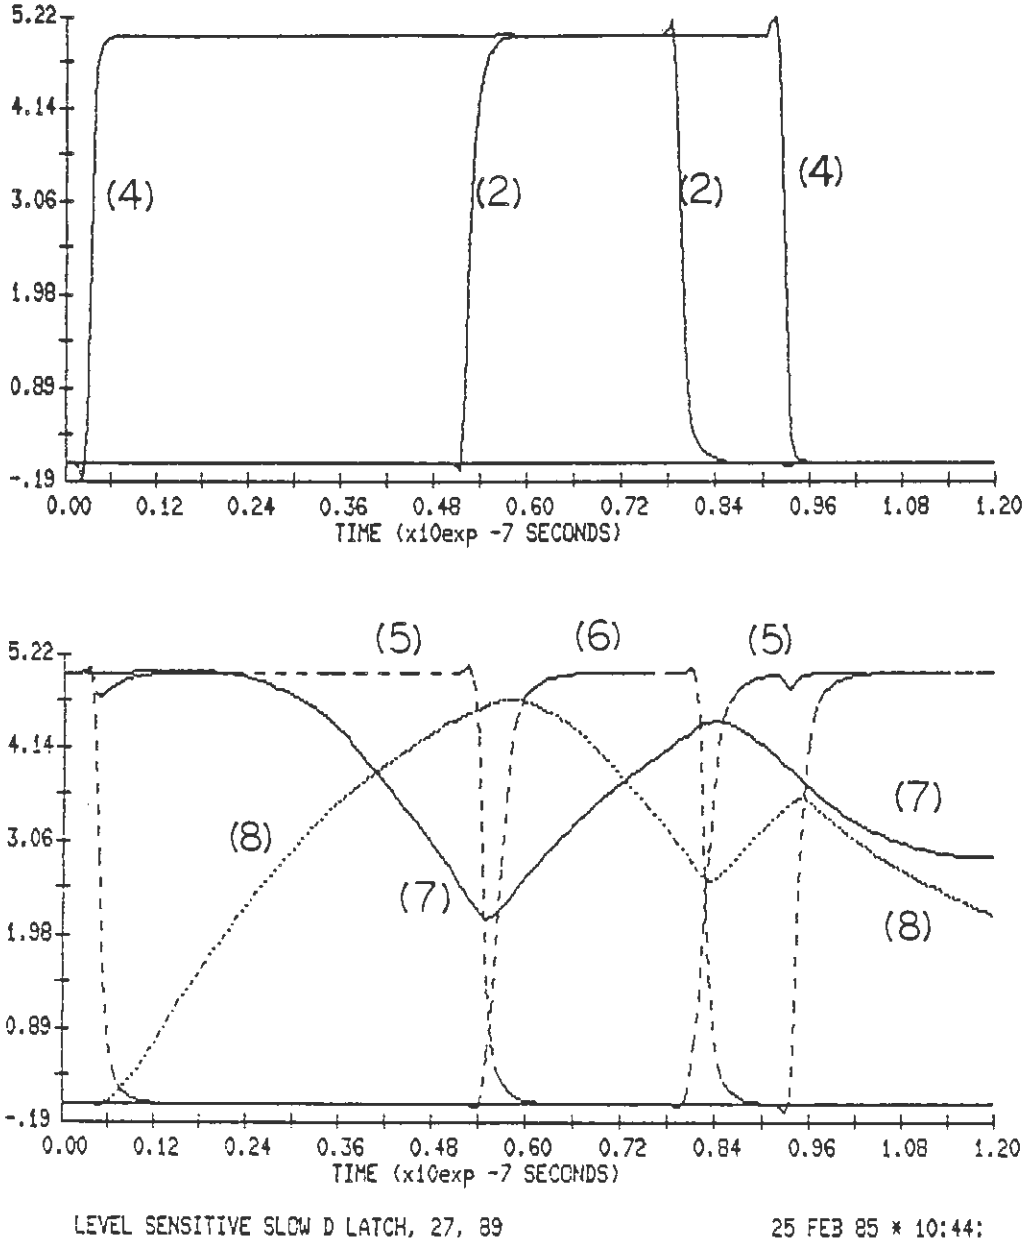
5.22
4.14
3.06
1.98
0.89
-.19
(4)
(2)
(2)
(4)
0.00 0.12 0.24 0.36 0.48 0.60 0.72 0.84 0.96 1.08 1.20
TIME (x10exp -7 SECONDS)
(5)
(6)
(5)
(8)
(7)
(7)
(8)
LEVEL SENSITIVE SLOW D LATCH, 27, 89
25 FEB 85 * 10:44:

Figure 5.3 (b)

recover to a valid state. Thus the slow latch used to obtain the data in Fig. 5.3 has a worst case of $S \approx 3.4$, which is better than some alternative designs will be seen to be capable of yielding. However, it can be noted in Fig. 5.3(a), for example, that the output voltage on node 7 is still far from a valid 0, 120 ns after the rising edge of the clock pulse, a time which is more than 4 times the duration D of the disturbing transient.

Although it is possible to improve the S of a slow latch by fine tuning its time constants, a much more significant improvement is obtained in the next sections by modifying the structure of the circuit. Also, the latch in Fig. 5.2(b) has another important weakness, namely, the direct dependence of the recovery time on the time constant of the latch. In other words, by making the latch slower not only are longer transient filtered, but the time necessary for restoring a valid level after the clock pulse is finished is also increased. This difficulty will be overcome with the circuits proposed in the following sections.

5.2 Single-Filter Latch

It is assumed here that tolerance to transients is achieved by filtering the corrupted signals. Clearly, the filter will be a source of overhead. Therefore if a single-filter latch can approach the optimum performance, there should not be much room left for improvement unless the latch becomes significantly bigger. This approach was first explored.

5.2.1 Functional Design

A model of a latch is needed to demonstrate the feasibility of designing filtering registers. A model which includes only logic elements such as gates and switches is not sufficient for this purpose. Figure 5.4 shows a functional model of a latch which retains the necessary and sufficient features for a functional optimization. The low-pass filter accounts for the property which determines the shortest event that can be latched. Another property which a latch possesses is logic level restoration, which is accounted for in the model by the quantizer. This memory element is, of course, the essential part of the latch. Because the registers are assumed to be immune to effects which cause soft errors, the memory cannot be based on charge storage as in a dynamic MOS register, and must be implemented with a bistable latch. Notice that the latch model in Fig. 5.4 is very similar to that of a matched filter receiver for a noisy communication channel, shown in Fig. 5.5. This observation reinforces the analogy discussed in Chapter 4, between a noisy communication channel and a machine sensitive to soft-errors. Moreover, the existing knowledge on matched filter design [PRO83] can serve as a guide to the design of filtering latches.

If the only non-ideal component in the latch is the filter, then the problem can be simplified to optimizing it. Since the filter receives its input signal from a logic gate, it can be assumed that all signals and transients have an amplitude equal to the supply voltage which is normalized to 1. Consequently there will be only two situations to analyze: a valid 1 disturbed by a transient 0, or a valid 0 disturbed by a 1, as shown in Fig. 5.6. The origin of time has been taken, without loss of generality, as the moment when the data is known to be valid.

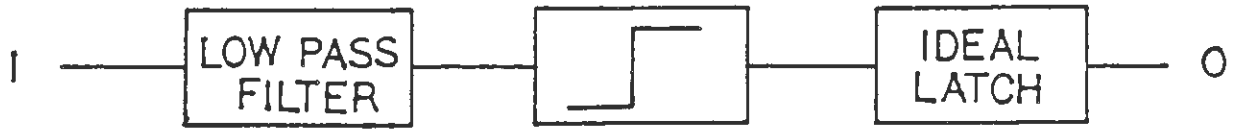

Figure 5.4 Functional model of a latch

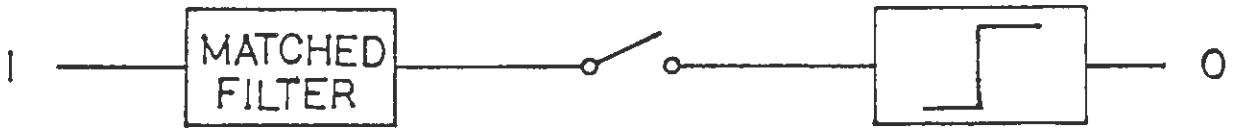

Figure 5.5 Matched filter receiver

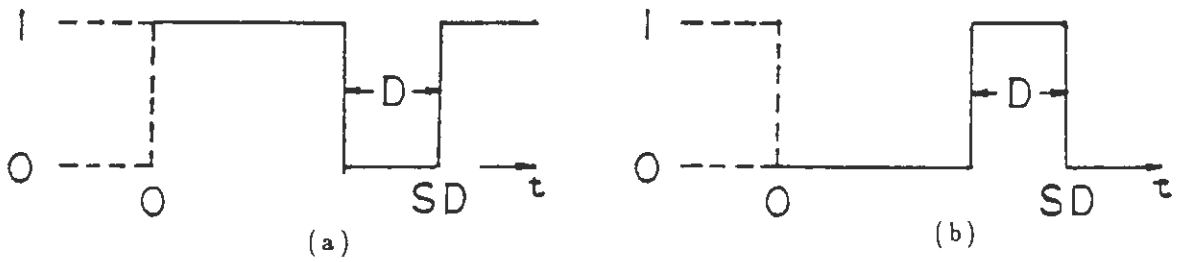

Figure 5.6 Signals disturbed by transient pulses during the sensitive time of a latch (a) A valid 1 disturbed by 0 (b) A valid 0 disturbed by 1

A lower bound on S can be estimated by using the following simple argument. It is assumed that the probability of a 1 or a 0 occurring at the filter input is the same. Furthermore, the distributions of expected transients of either polarity are equivalent and symmetrical. In such a case, if a signal has one polarity during more than half of some time interval, the probability of making an error is minimized by choosing that polarity as the probable signal. In other words S must be greater than 2.

A bias in the distribution of transient polarities does exist at the generation point. Nevertheless the assumption of symmetrical distribution of the disturbing transients is justified by the fact that, for a large machine, there is no a *priori* bias for an odd or even number of inversions from the generation point to the filter input.

A filter which has a potentially good performance is the ideal integrator. This is suggested by analogy to the matched filter for square pulses, which is optimum for additive white Gaussian noise.

Assuming that the output of the integrator is reset to 0 at $t = 0$, the signals shown in Fig. 5.6 will be recovered properly if the following three conditions, discussed below, are satisfied:

$$KD(S-1) \leq 1 \tag{5.1}$$

$$KD(S-1) > Th \tag{5.2}$$

$$KD < Th \tag{5.3}$$

where K is the gain of the integrator, D is the maximum duration of a transient that the latch is designed to filter, as defined in Chapter 4, and Th is the quantizer threshold.

Expression (5.1) ensures that the output remains within the range of the supply voltage when the signal is corrupted and, therefore, that clipping does not occur. Clipping is allowed to occur if the input signal is not corrupted, in which case the input signal is declared to be 1. Inequalities (5.2) and (5.3) correspond to the requirements for recovering valid data in the two possible situations depicted in Figs. 5.6(a) and (b), respectively. The solution of (5.2) and (5.3) yields $S > 2$ in agreement with the preceding argument.

It is convenient to compare filters in the present context on the basis of their *discrimation*, defined as the difference between the output values corresponding to the two conditions illustrated in Fig. 5.6, evaluated at the sampling time SD. This quantity is an important figure of merit because, in a practical implementation, the quantizer cannot be assumed to be perfect and, therefore, the larger the discrimination, the easier it is to design a quantizer for signal recovery.

For the integrator, the discrimination is given by the difference between the left-hand sides of (5.2) and (5.3), i.e. $KD(S-2)$. Hence the maximum discrimination, Δ_1, which, from (5.1), corresponds to $KD(S-1)=1$, is given by

$$\Delta_1 = \frac{S - 2}{S - 1} \qquad (5.4)$$

Even though the integrator achieves the lower bound on S, as will be shown below, it appears to be an impractical solution from the point of view of both speed and chip-area overhead. A logical alternative is the simple RC filter. Fig. 5.7 shows the response of this filter in the two situations depicted in Fig. 5.6, assuming that the initial voltage on the capacitor corresponds to the complement of the valid signal. It is important to note that the results obtained above for the integrator are independent of where the corrupting transient occurs in the time interval SD. On the other hand, the situation depicted in Fig. 5.6, namely corruption of data just before sampling, is the worst case for the RC filter because of its exponential response.

Because of the symmetry of the two responses in Fig. 5.7, it follows that Th = 0.5 is the optimum. Also because of the symmetry, it suffices to analyze only one case.

In the situation where a valid 1 is corrupted by a 0 transient (Fig. 5.7(a)), proper data recovery requires that

$$\left(1 - e^{-(S-1)\frac{D}{RC}}\right) e^{-\frac{D}{RC}} > 1/2 \qquad (5.5)$$

For a given ratio of D/(R C), the maximum output occurs when

$$\frac{D}{RC} = \frac{\ln S}{S - 1} \qquad (5.6)$$

Substitution of (5.6) into (5.5) yields a lower bound on S

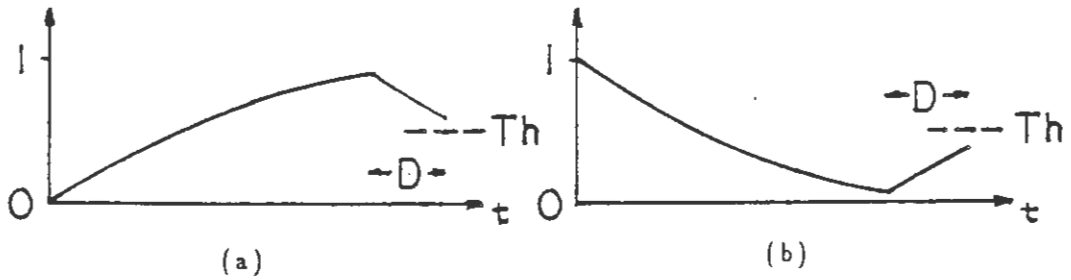

Figure 5.7 Response of the RC filter to the input signal shown in Fig. 5.6, assuming an initial charge which is the complement of the valid signal

$$S^{-\frac{1}{S-1}} - S^{-\frac{S}{S-1}} > 1/2 \qquad (5.7)$$

which has a numerical solution $S > 4.4035$. The limiting case $S=4.4035$ yields $D/RC = 0.4355$.

The maximum discrimination for the RC filter, Δ_2, is easily shown to be given by

$$\Delta_2 = 2\left(S^{-\frac{1}{S-1}} - S^{-\frac{S}{S-1}} - \frac{1}{2}\right) \qquad (5.8)$$

Figure 5.8, which shows the dependence of Δ_2 on S, clearly illustrates the superiority of the integrator. Note that each point of the maximum discrimination curve represents the best performance of an RC filter with a different time constant. This comment also applies to the other curves derived later for RC filters. Therefore an RC filter is optimum for only a single value of S, and the discrimination for this particular filter is smaller than the value given by (5.8) for all other values of S.

The two principal reasons for the relatively poor discrimination of the RC filter are the exponential nature of the response and the absence of an initialization of the output voltage. Whereas the former is an intrinsic property of the circuit and can only be compensated for by the use of nonlinear elements, the latter can be corrected quite simply.

Figure 5.9 shows the simplified circuit of an RC filter with precharge. Since the input signals of both polarities must be tolerated with the same performance, it can be shown that the optimun initial voltage on the

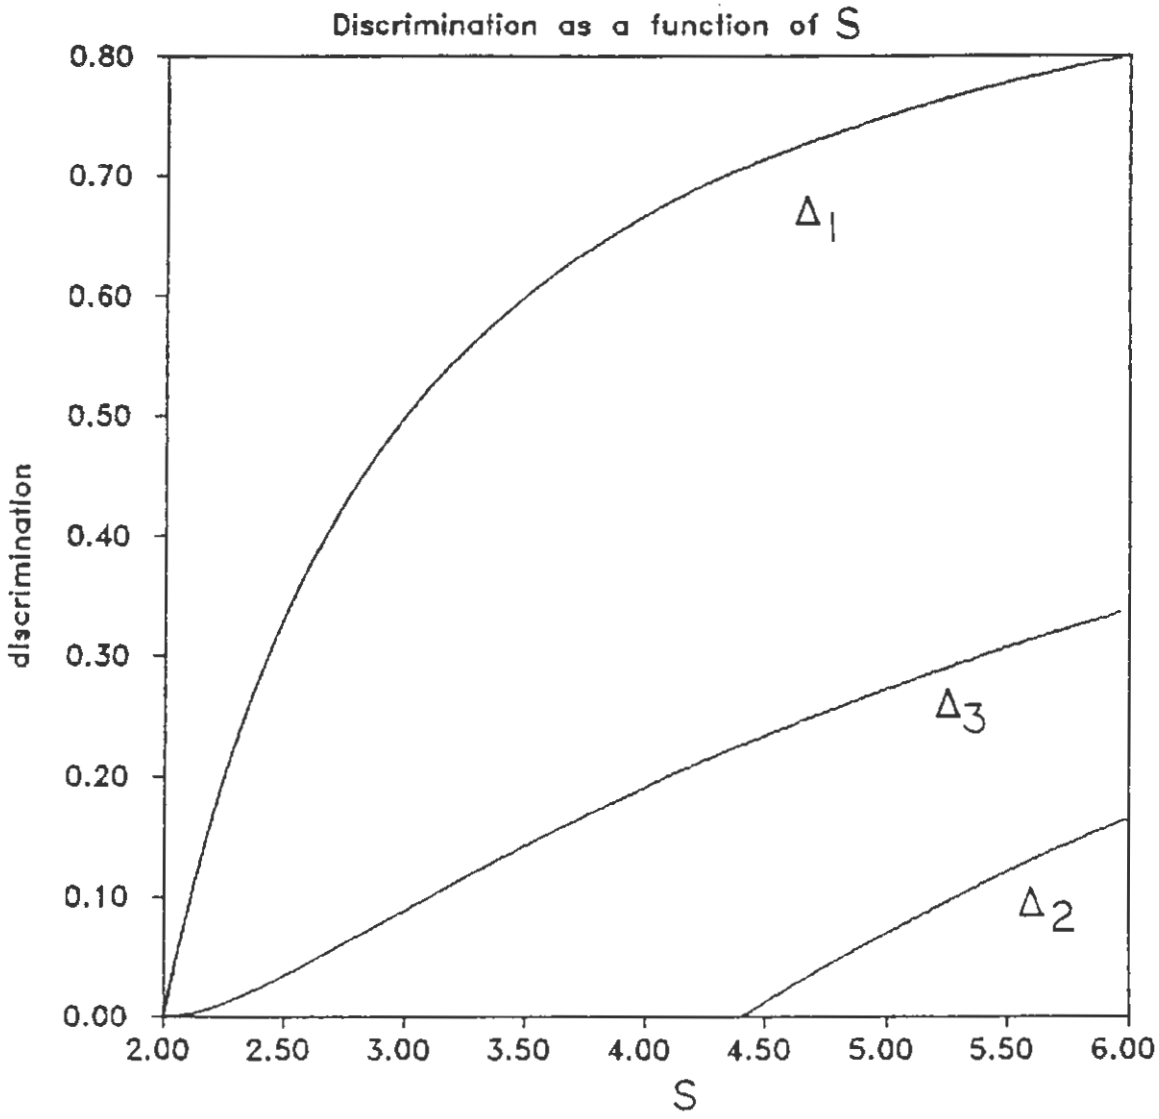

Figure 5.8 Plots of the computed maximum discrimination as a function of S for the integrator, Δ_1, the pure RC filter, Δ_2, and the RC filter with precharge, Δ_3.

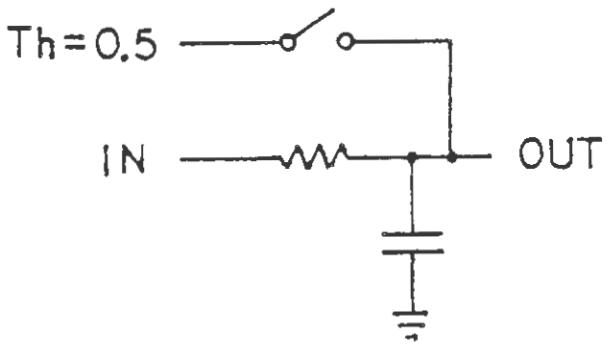

Figure 5.9 RC filter with precharge.

capacitor should be equal to the threshold of the quantizer Th = 0.5. An analysis which is completely analogous to that for the simple RC filter leads to the results that, for a given D/RC, the discrimination is maximized when

$$\frac{D}{R\,C} = \frac{1}{S-1} \ln\left(\frac{S}{2}\right) \tag{5.9}$$

and the lower bound on S is given by

$$\left(\frac{S}{2}\right)^{-\frac{1}{S-1}} - \frac{1}{2}\left(\frac{S}{2}\right)^{-\frac{1}{S-1}} = \frac{1}{2} \tag{5.10}$$

The solution of (5.10) is S = 2, which means that this filter has the same lower bound on performance as the integrator. Not surprisingly (5.9) shows that, to achieve this performance, the condition RC >> D must be satisfied, i.e. the filter should behave like an integrator.

Figure 5.8 shows a comparison of the discrimination of three filters, where that of the RC filter with precharge is given by

$$\Delta_3 = 2\left(\frac{S}{2}\right)^{-\frac{1}{S-1}} - \left(\frac{S}{2}\right)^{-\frac{S}{S-1}} - 1 \tag{5.11}$$

Clearly the use of the precharge with the RC filter leads to a very significant improvement in performance.

In conclusion, the RC filter with precharge is the best choice, taking into account both circuit complexity and performance. In a monolithic IC design, the resistor is synthesized using transistors, and in the MOS technology in particular, it is replaced by a FET operated in the triode

mode, which helps to minimize the degradation in discrimination associated with an exponential response, as will be shown in the next section.

5.2.2 Circuit Implementation

An implementation is now considered for a single-filter latch. The latch must have intrinsic tolerance to ionizing radiation, in addition to the capability of efficiently filtering all pulses on its input up to a prescribed duration.

The filtering latch presented here is derived from a standard dynamic RAM sense amplifier configuration [MAV83 p.134] shown in Fig. 5.10. This circuit can be viewed as a pair of inverters with tight feedback. These inverters form a static bistable element. In the RAM context, the recovery of the content of a memory cell can be achieved by polarizing and modifying the feedback loop of this circuit, in such a way that the final state of the bistable element reflects the initial charge on a storage node. The data recovery process involves three different modes. First the circuit must be precharged at threshold with M5, then a fraction of the charge on the storage node is injected into the gate of M3, which is floating when Φ_2=0, and finally the bistable element has a regeneration phase, where the level must start from threshold plus or minus a small voltage, and reach a valid logic 1 or 0 respectively. Therefore, this circuit is capable of regenerating a valid logic signal from a small difference in initial voltage. This property is particularly interesting if a low time overhead is desired, since, as demonstrated earlier, it means that the filter operates with a small discrimination.

After a number of refinements, the resulting circuit is shown in Fig. 5.11. It can be noted that there is not always a one to one correspondance between the models in Figs. 5.4, 5.9, and the implementation in Fig. 5.11.

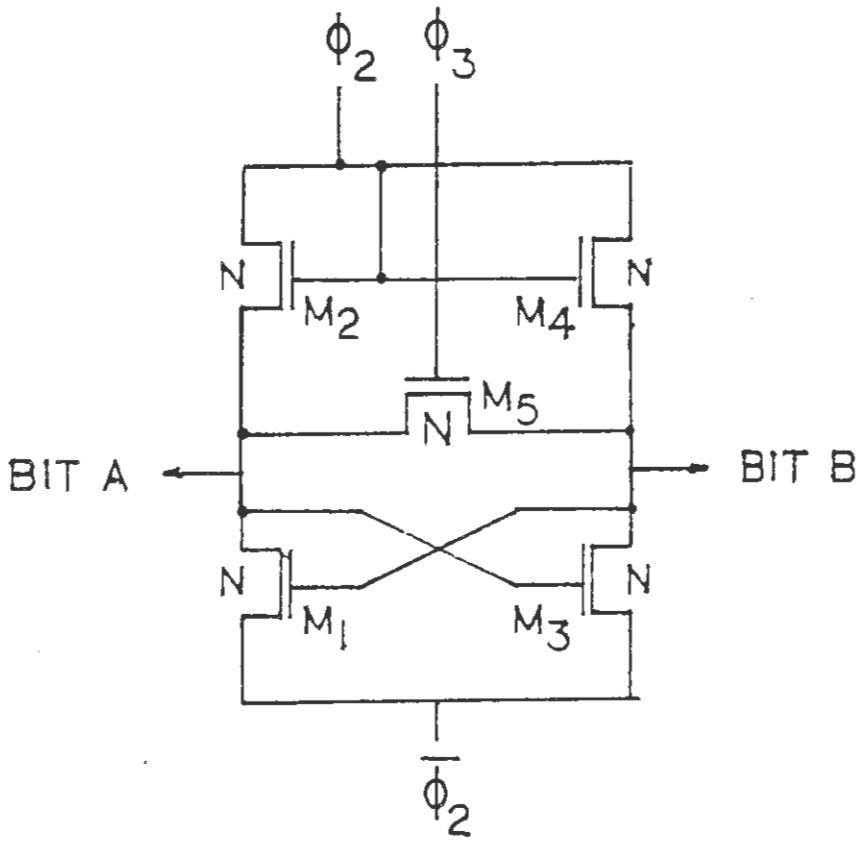

Figure 5.10 A standard sense amplifier configuration. Reproduced from [MAV83 p.134].

For example, M1-M4, are used to implement the quantizer and the latch of the idealized model in Fig. 5.4. The filter is formed by a transmission gate, M12, M13, used as a saturable resistor, feeding a nearly linear capacitor, implemented by transistors M8 and M9. The switch in Fig. 5.9 is implemented by M5-M7, but M5 and M6 are also necessary for the operation of the latch. The transmission gate formed by M10, and M11, is necessary for defining the time interval during which data is supposed to be valid, but it is not shown in the model of Fig. 5.4.

A normal sequence of data recovery begins with the precharging of the circuit at the threshold of the inverters. This is achieved by forcing Φ_2=0V with Φ_3=0V and Φ_1=0V. A simplified equivalent circuit for this mode is shown in Fig. 5.12(a), and the sequence of timing pulses is shown in Fig 5.13. After the circuit is established at threshold, it is ready to filter an input signal with Φ_2=5V, Φ_3=5V, and Φ_1=5V which corresponds to the equivalent configuration shown in Fig. 5.12(b). If the input signal is limited to 0 or V_{dd}, with a relatively short transition between the two, this circuit emulates the behavior of a true integrator with fairly good accuracy. Finally, the third mode consists of having Φ_1=0V, Φ_2=5V, and Φ_3=0V, which yields the equivalent circuit shown in Fig. 5.12(c).

5.2.3 Choice of Dimensions for the Transistors

The optimization of this circuit has to be performed with a circuit simulator such as SPICE, but a good initial guess based on its basic properties will facilitate convergence. The first requirement is that a direct alpha particle hit on the latch must not affect the stored value.

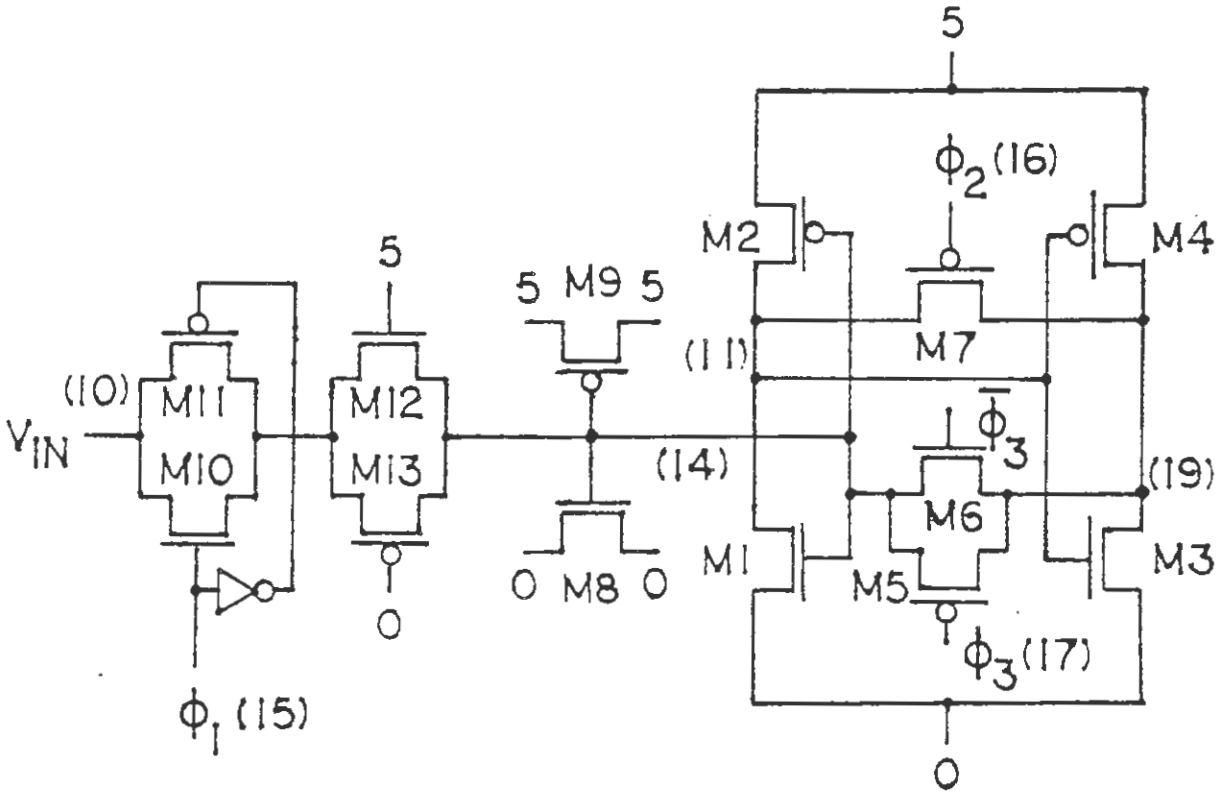

Figure 5.11 Single-filter implementation of the filtering latch. Numbers in paratheses correspond to nodes in simulations. (L,W) in microns, M1=M3=(5,10), M2=M4=(5,26), M5=(5,30), M6=(5,15), M7=(5,150), M8=M9=(20,75), M10=(5,5), M11=(5,13), M12=(16,5), M13=(6,5).

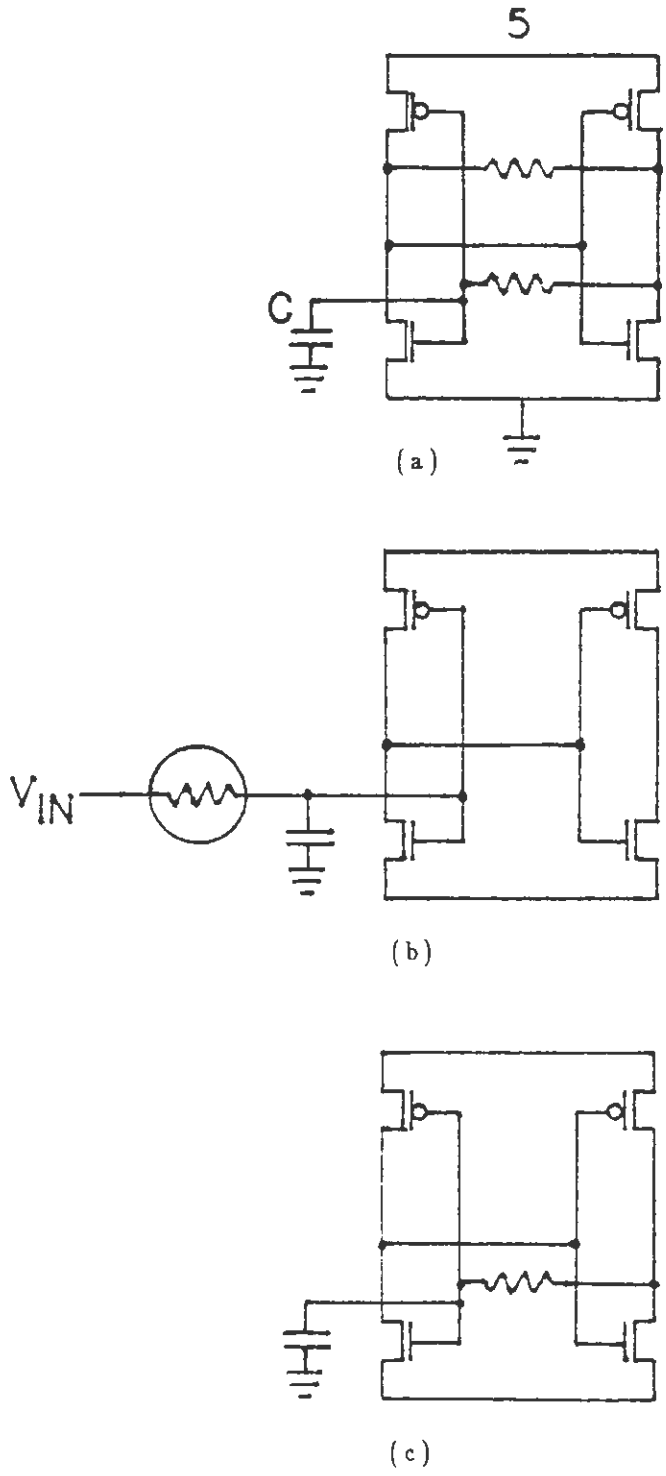

Figure 5.12 Equivalent circuit for Fig. 5.11 with different combinations of clocks applied
(a) $\Phi_1=\Phi_2=\Phi_3=$ 0V, precharging state
(b) $\Phi_1=\Phi_2=\Phi_3=$ 5V, filtering state
(c) $\Phi_1=\Phi_3=$ 0V, Φ_2=5V, restoring state

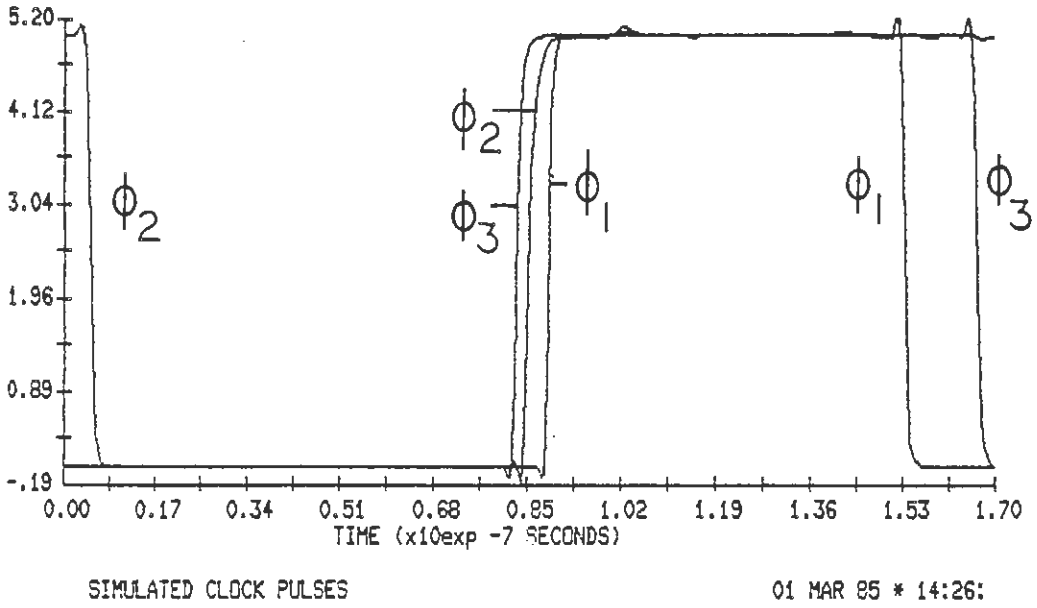

Figure 5.13 Clock pulses as simulated

The number of transistors in Fig. 5.11 is modest: however, there exist a large number of possibilities for the choice of their dimensions. Moreover, this circuit has three modes of operation with different sensitivities. Therefore the relative sensitivities of the various nodes and modes of operation must be discussed, in order to facilitate the convergence of the design process. It is demonstrated below that a direct hit on node 19 need not be considered in the following analysis, if nodes 11 and 14 can tolerate a direct hit. Moreover, node 14 is most sensitive during the filtering period, and node 11 is most sensitive just before the end of the filtering period.

When the feedback loop of the latch is closed, M5 and M6 are on, and the logic levels are restored to either 0 or V_{dd}. In this case, the transient injected by a direct hit on the latch is partly neutralized by a low impedance path to one of the supply busses, therefore the latch is more tolerant.

During the filtering phase, when M5 and M6 are off, node 14 is not protected by a low impedance path to the supply, moreover node 14 is polarized near $V_{dd}/2$, and the amplitude of the transient required to change the state of the latch is reduced. Therefore the tolerance of node 14 is reduced. In this situation, the tolerance of node 14 is determined by the inertia associated with the capacitor formed by the gates of M8 and M9.

If the voltage on node 14 is near the threshold of the latch, which is the case during filtering, the ability of M1 and M2 to combat an injected transient is reduced. Moreover if the hit affects node 11 just before the

feedback loop is closed, the time allowed for M1 and M2 to recover from this hit is minimized. Therefore, the tolerance of node 11 is at a minimum if a direct hit happens just before the end of the filtering period.

If node 11 and 14 are tolerant to a direct hit, then node 19 is tolerant. It is easy to show that if equivalent transistors are used, the ability of M3 and M4 to neutralize an injected current pulse on node 19, is always better than, or equivalent to, that of M1 and M2 to neutralize the same pulse on node 11. Since the signal on the gates of M3 and M4 is an amplified version of the one on the gates of M1 and M2, it is greater than or equal to the signal on the gates of M1 and M2. Moreover when M5 and M6 are on, node 19 is protected by the large parasitic capacitance of node 14, and if the transient is injected just before the end of the filtering period, charge charing between nodes 14 and 19, when M5 and M6 are on again, would help to neutralize any injected transient.

It follows from this discussion that a reasonable starting point for this design is to choose the value of the capacitor C on node 14, assuming that the parasitic capacitance must itself be sufficient to guarantee intrinsic tolerance. The injected current pulse due to a direct alpha hit is shown in Fig. 5.1. The voltage transient on node 14, resulting from this current pulse, must have an amplitude smaller than $V_{dd}/2$, because in any case, it must be smaller than the signal from which the state of the latch is restored, which is smaller than $V_{dd}/2$. If the performance of this latch is to approach that of an ideal latch with a true integrator input filter, the amplitude of this transient must in fact be significantly smaller. The maximum amplitude resulting from a direct injection on node 14 was thus chosen to be $V_{dd}/5$. The total injected charge is $2.2*10^{-12}$C and, since

C_g is on the order of $2*10^{-2}$ pF for CMOS1B, this translates into a pulse amplitude of 110 V-C_g. As mentioned earlier, the capacitor C has been formed by the gate to channel capacitance of an N and a P device of identical area (M8 and M9). An area 55 times that of a minimum size transistor for M8 and M9yields a nominal injected transient amplitude of $V_{dd}/5$. In fact, the simulations discussed later have been done with M8 and M9 having an area that is 60 times larger than that of a minimum-size transistor channel. Transient injection on nodes 11 and 14 (of the complete design) has been simulated, and in both cases the transients were insufficient to change the final state of the latch.

Limiting the amplitude of the injected transient to $V_{dd}/5$ is an arbitrary choice. If larger transients were allowed, the area occupied by M8 and M9 could be decreased. On the other hand, in order to approach S=2 the amplitude of the signal on node 14 must be relatively small. If large transients were allowed, a larger signal would be necessary, and the filter would follow the exponential response associated with a RC filter, which increases S for a given discrimination as demonstrated earlier. Therefore the area reserved for C is a compromise between the area and the time overheads introduced by the latch.

The inverters formed by M1, M2, M3, and M4 could in principle be minimum size devices. However the width of the P devices was set to 2.6 times that of the N. It was observed by simulating variants of this circuit that even though the performance of the circuit is almost insensitive to an imbalance in the capacitive loading of nodes 11 and 19, it is sensitive to the relative conductance of the N and P devices. An imbalance of the conductance of the N and P devices changes the threshold of the inverters,

and it is important to have a threshold at the middle of the supply voltage. Moreover, the conductance between nodes 10 and 14 of M10, M11, M12, and M13 must be as equal as possible for the two input signal polarities. Any imbalance in threshold or in conductance results in one polarity of input data corrupted by a transient, for which the response on node 14 is faster for the transient polarity than for the signal polarity. Consequently the latch would require a greater filtering time in comparison with a latch where the slope of the response is symmetrical.

The width of M3 and M4 has been increased by a factor of 2 to decrease the initial precharging time. Even though, strictly speaking, the precharging phase does not generate time overhead, if this period is too long it becomes a limiting factor. Also, a similar scaling of 2 for M1 and M2 helped to reduce the lag between nodes 14 and 19. This lag would have introduced more delay in the feedback loop, a limiting factor on the performance to be discussed later.

The transistor M7 is responsible for precharging the circuit at the inverter threshold. The decision to use only a P device follows from the important body effect that affects the N device in a P well process. When the precharging voltage approaches the threshold of the inverter, the conductance of a N device becomes negligible when compared with that of a P device of similar size. The exclusive use of a P device is most significant because of its large size, which follows from the fact that if its ON impedance were too high, a steady state precharging offset would remain between nodes 11 and 19. This phenomenon is similar to what happens in NMOS, a ratioed impedance logic family, where the steady state voltage representing a logic 0 is not 0V, but depends on the ratio of impedances. If the width of M7 is more than 5 times

the larger of M2 and M4, the steady state offset becomes negligibly small (this effect is non-linear). Therefore, a significant saving result from not completing the transmission gate with a N device.

Finally the slope of the response from V_{in} to 14 is controlled by transistors M10 to M13. Logically it can be seen as a single transmission gate. However, a significant reduction of the feedthrough from Φ_1 to 14 is obtained when two transmission gates are used, with the second permanently on. This can make a significant difference if the saturation current of the device is reduced by increasing its length, in order to filter very long transients.

5.2.4 Simulation Results

Figure 5.13 shows the clock pulses necessary to operate the latch shown in Fig. 5.11. The first period with Φ_2 and Φ_3 low, polarizes the latch at its threshold. After the latch has stabilized to its threshold, the feedback is turned off by making Φ_3 high, and then Φ_2 is switched to a high level leaving the storage node 14 in a floating state. The filtering operation begins by turning on Φ_1. As long as Φ_1 remains on, the input signal is integrated on top of the threshold voltage with a fair accuracy, for a limited period of time, because the circuit emulates an integrator when the input signal is at one of the supply voltages. Then Φ_1 is turned off and the feedback is turned on again by making Φ_3 low. The signal on node 14 drives the latch to 1 if it is larger than the threshold of the latch and to 0 otherwise. The order in which the three clock signals switch from 0 to 1 before the filtering period is important. But leaving a small delay between

the transitions does not increase the time overhead, and it ensures the best performance for the latch. If the order is violated, or if the delays become too small, the precharging is imperfect, and the performance of the latch degrades gradually. However, after the filtering period, the delay between the 1 to 0 transitions of Φ_1 and Φ_3 should not be 0, if the latch is to be operated with almost no margin (minimum value of S possible), as discussed later.

Figure 5.14 shows the response of this latch corrupted by a transient pulse of 27 ns in four situations. The duration of Φ_1 is 60 ns, which is the shortest one that permits recovery in the 4 situations simulated. It is important, as can be noted in Fig. 5.14, that the simulation always begins with a polarity on node 14 which is the opposite of the appropriate final value. This conservatively takes into account the small offset that remains on node 14 after precharging. With the settling time allowed, and for this set of transistor sizes, a 38 mV difference remains at the end of the precharging period, for the various simulations with a logical one or a zero as the initial value on node 14. The two polarities of the transient at the edge were simulated, since it was the worst situation for a filter with an exponential response. For large amplitude signals, the slope of the response does decrease with time and, therefore, a pulse at the edge yields the smallest discrimination.

Even though the discrimination is larger when a pulse arises some time before the falling edge of Φ_1, the internal delays of the feedback loop make the situation depicted in Fig. 5.14 (a) and (c) more difficult to tolerate. Trial and error were used to determine what is a *bad delay* between the transient on input and the trailing edge of Φ_1. The transient

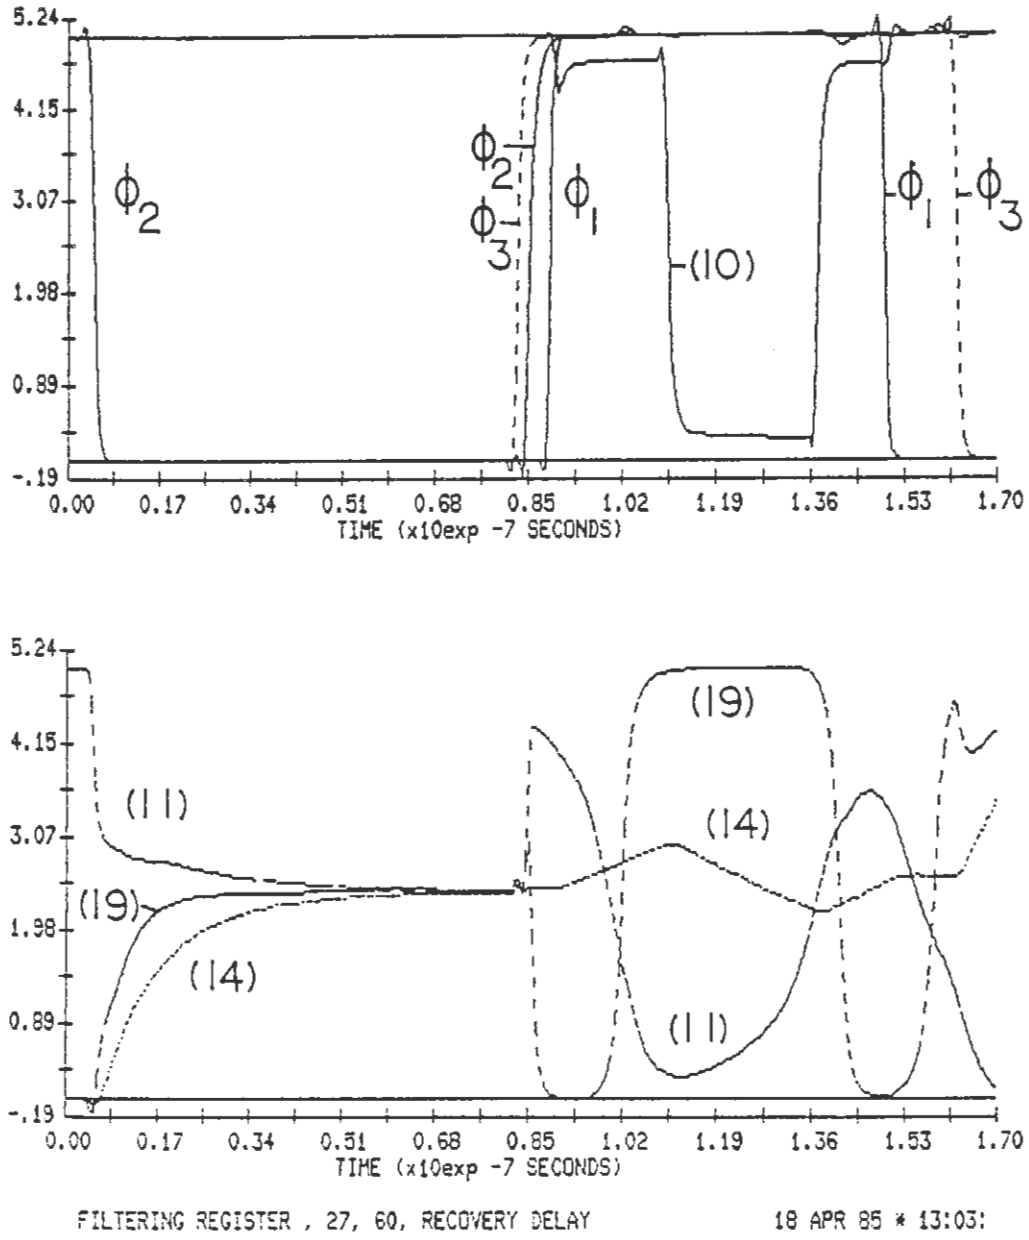

Figure 5.14(a) Simulated response of the circuit shown in Fig. 5.11. The transient pulse lasts 27 ns and the filtering time is 60 ns. The input signal is a 1 corrupted by 0 (node 10), the transient pulse occurs 14 ns before the trailing edge of the clock Φ_1.

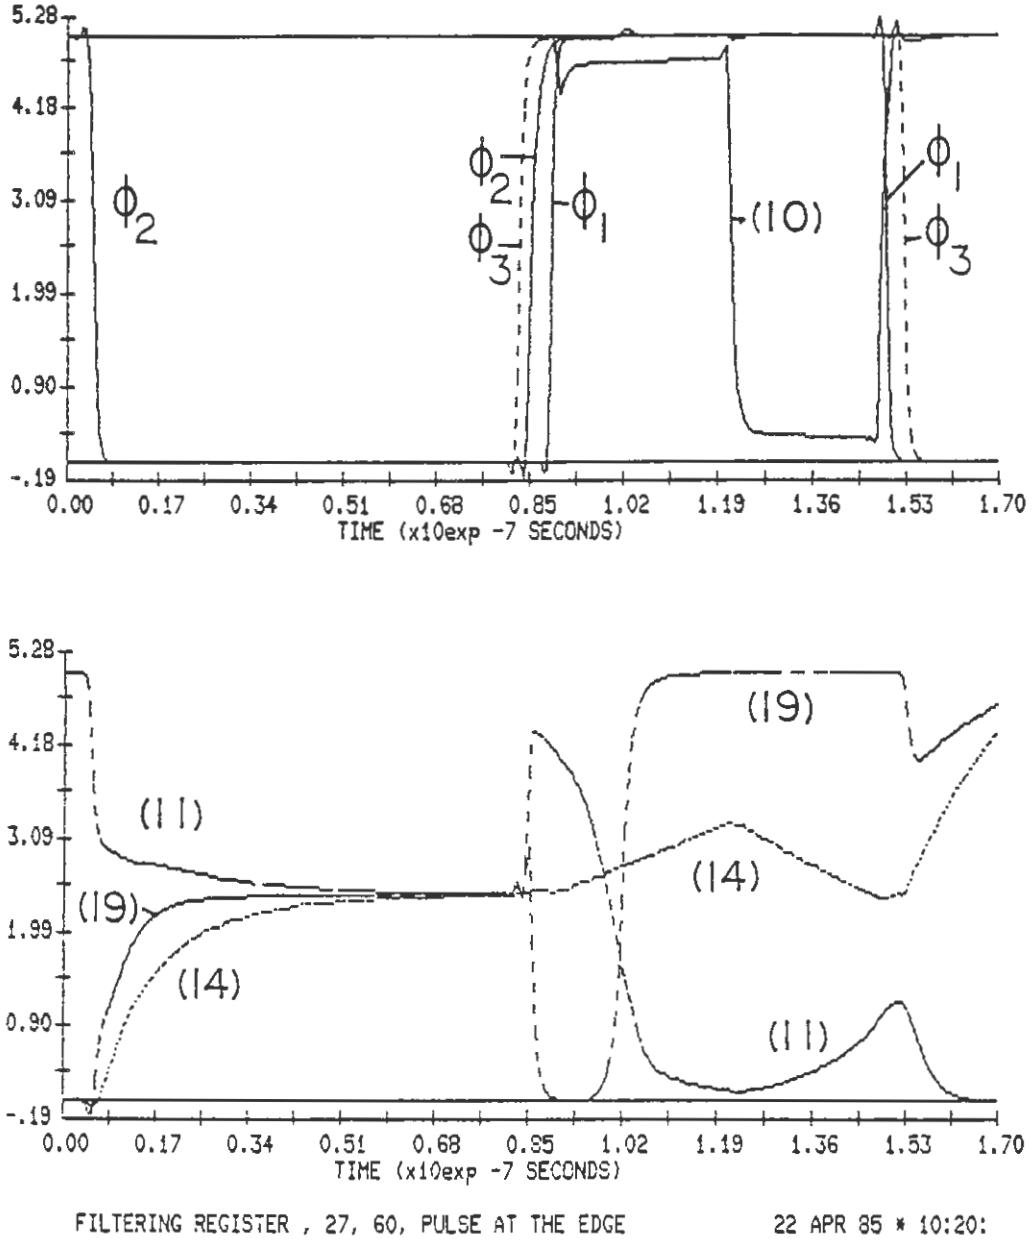

Figure 5.14(b) The input signal is a 1 corrupted by 0 at the trailing edge of Φ_1.

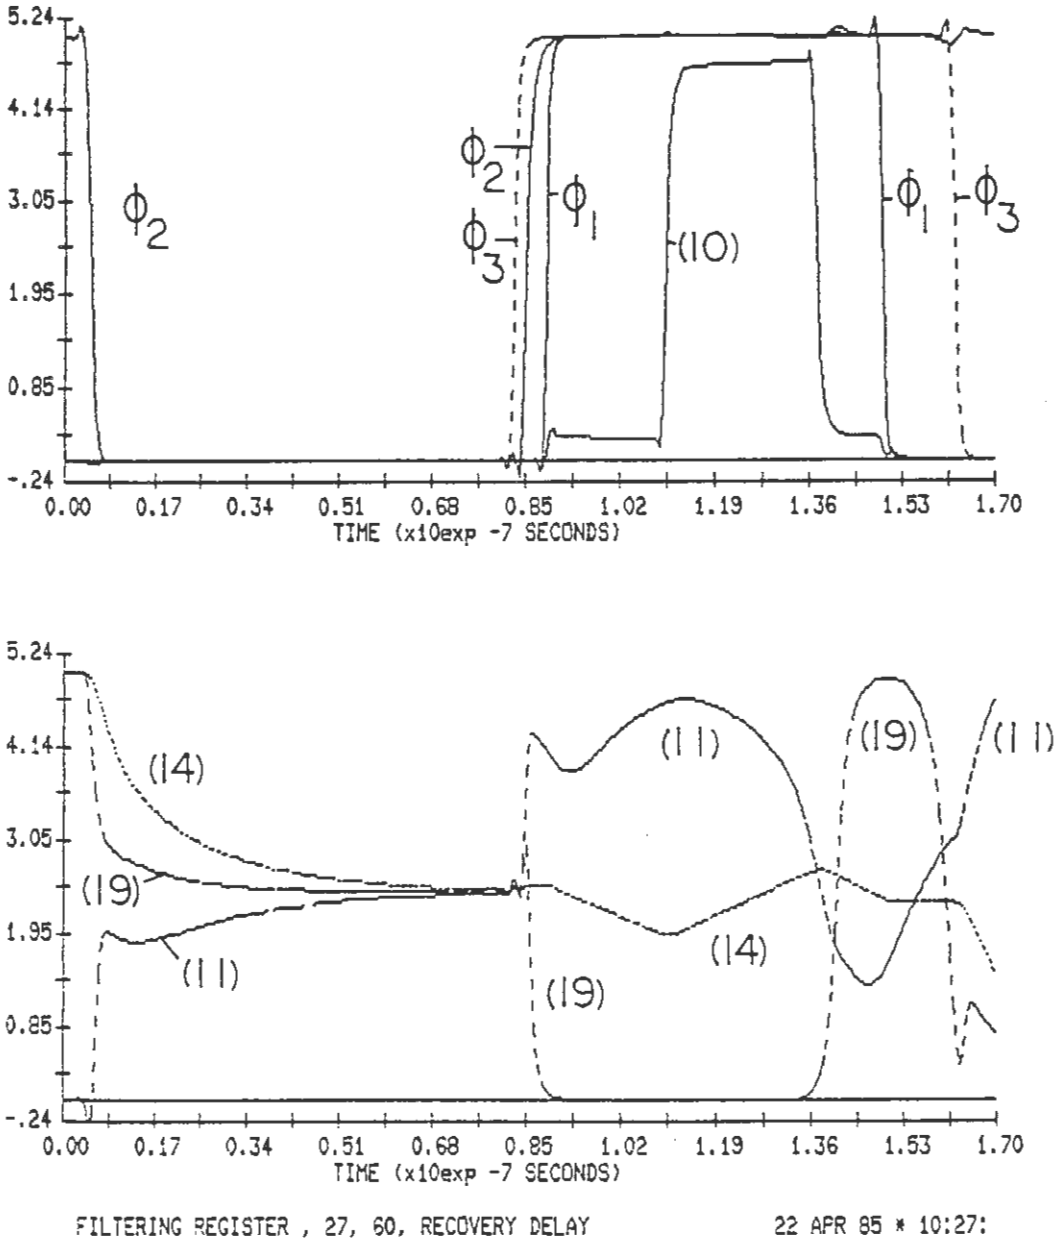

Figure 5.14(c) The input signal is a 0 corrupted by 1, the transient pulse occurs 14 ns before the trailing edge of Φ_1.

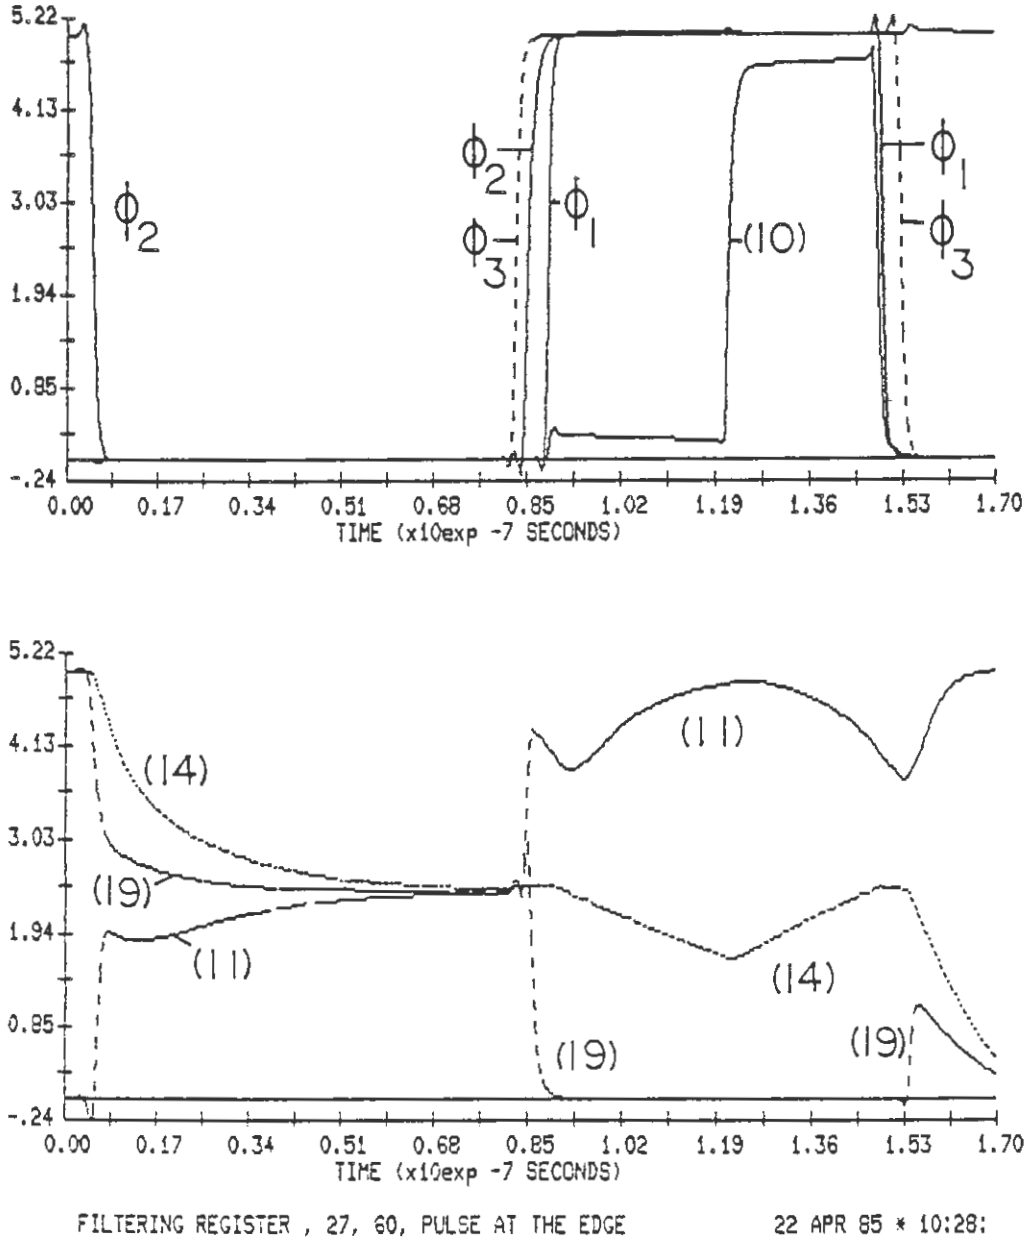

Figure 5.14(d) The input signal is a 0 corrupted by 1 at the trailing edge of Φ_1.

Table 5.1
Simulation Results for the Single-Filter Design in Fig. 5.11

Transient pulse at the edge Φ1

Φ1 (ns)	S	Voltage on node 11 at trailing edge of Φ1		Discrimination (V)	(normalized)
		1 by 0 (V)	0 by 1 (V)		
57	2.11	2.337	2.528	-.191	-.0382
57.5	2.13	2.340	2.525	-.185	-.0370
58	2.15	2.349	2.511	-.162	-.0324
60	2.22	2.384	2.473	-.089	-.0178
62	2.30	2.418	2.439	.021	.0042
64	2.37	2.456	2.402	.054	.0108

Transient pulse injected 14ns before the trailing edge of Φ1

Φ1 (ns)	S	stable output (node 11) before restoration		Discrimination (V)	(normalized)
		1 by 0 (V)	0 by 1 (V)		
57	2.11	2.551	2.353	.198	.0396
57.5	2.13	2.564	2.340	.224	.0448
58	2.15	2.568	2.336	.232	.0464
60	2.22	2.600	2.302	.298	.0596
62	2.30	2.630	2.271	.359	.0718
64	2.37	2.659	2.248	.411	.0822

Characterization of recovery delays (pulse before the edge of Φ1)

Φ1 (ns)	Crossing delays		Delays to 1 *		Delays to 0 *	
	1 by 0 (ns)	0 by 1 (ns)	1 by 0 (ns)	0 by 1 (ns)	1 by 0 (ns)	0 by 1 (ns)
57	10.2	-	16.2	-	20.1	-
58	9.1	-	14.5	-	16.9	-
60	4.8	5.4	9.3	10.2	12.3	12.6
62	0	0	4.5	3.8	6.9	5.2
64	0	0	0	0	3	0

* The thresholds for a valid 1 and 0 are taken to be respectively, 2.85V and 1.35V. These values are the input voltages corresponding to the minimum and maximum unit-slope points, on the voltage transfer characteristics of CMOS NOR and NAND gates built with minimum size transistors.

The precharging values are 2.411V and 2.449V when the starting voltage on the storage node are 0V and 5V respectively.

is difficult to tolerate when it produces a large excursion of the wrong polarity on node 11, and the peak of this excursion is reached near the trailing edge of Φ_1. As can be noted in the graphs labeled RECOVERY DELAY in Fig. 5.14 (a) and (c), when the pulse finishes, the input signal has been to the wrong level longer than to the correct level. The signal on node 11 reflects this, but with a significant delay. If the feedback loop is closed immediately after the end of sampling, the delay effectively neutralizes the discrimination. It is necessary to delay the falling edge of Φ_3 with respect to that of Φ_1, for at least long enough to ensure that the value on node 11 will cross that on node 14. This suggests the label *recovery delay* for the simulations (it should not be confused with the expression *recovery delay to a valid level*, which is the exact time necessary for reaching a valid level; the meaning is clear by the context).

A number of simulations were done to characterize the behavior of this latch. The results of these simulations are summarized in Table 5.1. The discrimination is the difference between the output signals on node 14, for the two polarities of the input signal shown in Fig. 5.6. The output signals are measured at the falling edge of Φ_1, for the simulations done with a corrupting pulse at the edge. Whereas the measurements were made at a later time, when the output were stabilized, for the simulations with a pulse 14ns before the trailing edge of Φ_1. In the second case, a recovery time was allowed for the latch before turning on the feedback and, therefore, it is more realistic to use the stabilized difference on node 14 after the trailing edge of Φ_1. The *crossing delay* is the time necessary, after the trailing edge of Φ_1, for the voltage on node 11 to cross that on node 14. The crossing delay is important because it determines the earliest time, after the trailing edge of Φ_1, when the feedback can be turned on and the

latch will still recover properly. Notice also that the recovery delays in Table 5.1 are measured as the time needed for the output signals to become valid logic levels. For the logic circuits fed by the output of the filtering latch, valid logic levels are defined to be V_{il}=1.35V and V_{ih}=2.85V. These values are the input voltages corresponding to the minimum and maximum unit-slope points, on the voltage transfer characteristics of 2 inputs CMOS NOR and NAND gates built with minimum size transistors.

As can be noted in Table 5.1, a pulse at the edge effectively yields a smaller discrimination but, due to the internal delays of the feedback loop, the signal can be recovered even with a negative discrimination of 191 mV, with a clock pulse of 57 ns. The results of the simulations with a pulse 14 ns before the edge demonstrates an imbalance in the tolerance of the two transient polarities (this is due to a small remaining difference between the conductance of N and P devices). Also, the delay required to reach a valid level is very important, and this delay is minimized by using a clock pulse slightly longer than the shortest required for recovering to the correct level. Figure 5.15 compares the simulated values of discrimination as a function of S with the theoretical results obtained earlier. It is clear that this realization can yield a better discrimination than the optimum RC filter with precharge. For sufficiently large S, the performance is between those of the integrator and the optimum RC filter with precharge.

The fact that a smaller discrimination can permit an easier recovery demonstrates that it cannot be the only performance criterion. For example, the recovery delay to a valid level is also very important and has been characterized. Other considerations are more difficult to quantify in general but may be even more significant. A very important consideration in the case

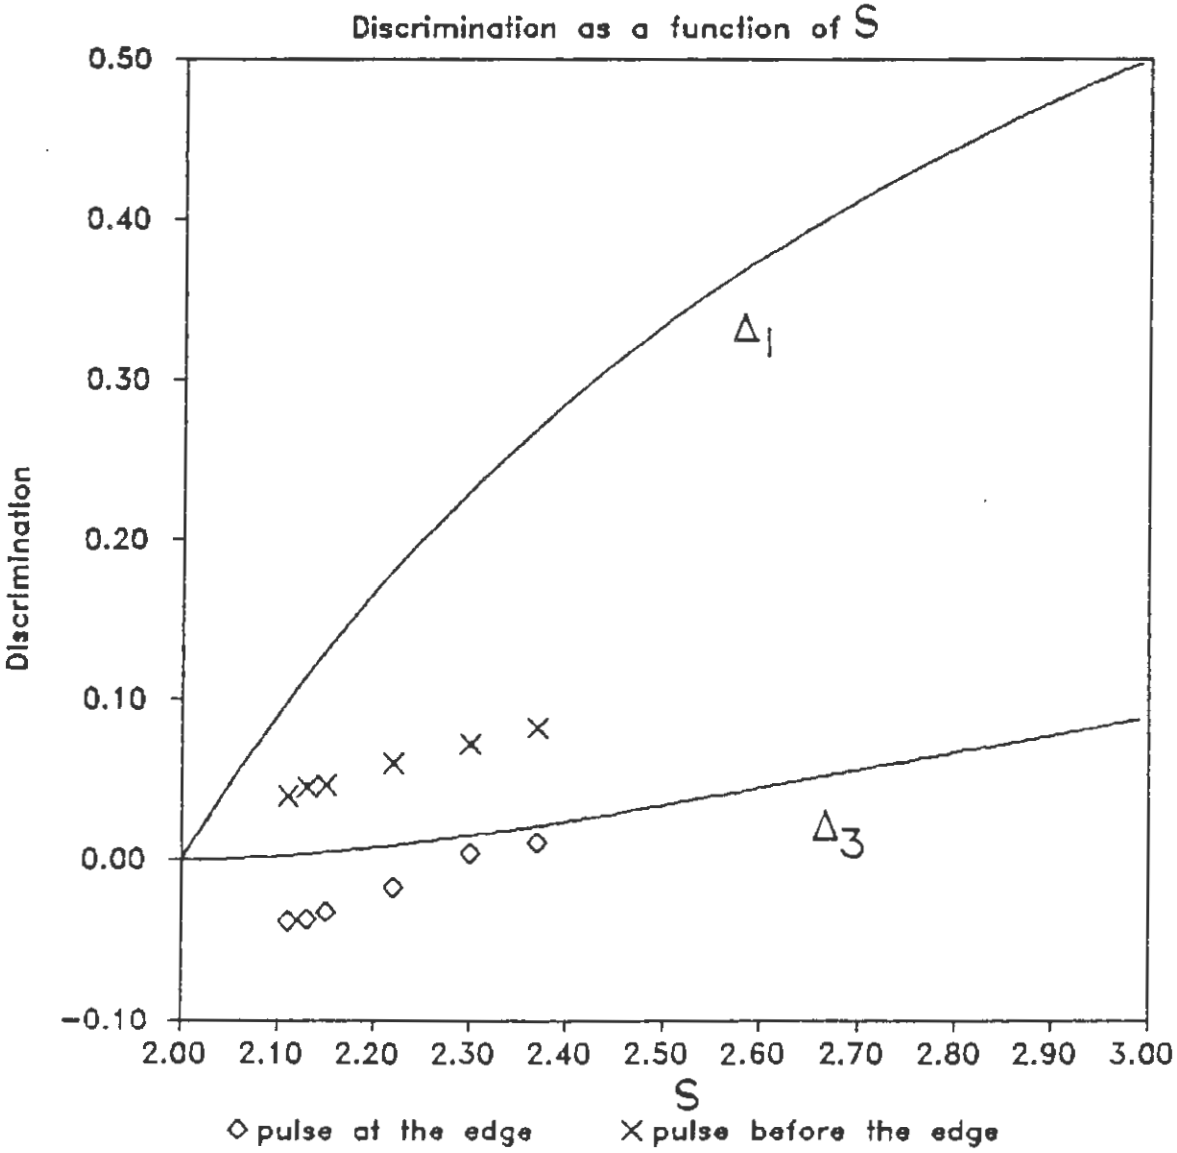

Figure 5.15 Comparison of simulated discriminations with theoretical results. If S is sufficiently large, the circuit in Fig. 5.11 has a performance between those of the best RC filter with precharging and the pure integrator.

of an integrated realization is the sensitivity to process fluctuations. To achieve a performance similar to the simulation results reported here would require an accurate match of the conductances of the P and N devices. The ratio of these conductances is precisely a characteristic which is not well controlled. This problem is complicated further by the fact that the absolute value of the load capacitance, which determines the time constant, is not accurately controlled. It is clear that the performance of the latch is sensitive to these fluctuations, but this analysis is left for further work.

Another consideration which is difficult to quantify is the constraint put on the designer by the series of clock pulses to be supplied to the latch. Generating and distributing them is not a trivial task, if good performance is the objective. Also, the dead time necessary for precharging is not counted as overhead. This is correct only if the precharging time is shorter than the delay in the logic; otherwise the fraction of the precharging time exceeding the delay in the logic becomes overhead, and the simple slow latch described earlier could be preferable. Moreover, from the beginning of precharging to the end of recovery, the signal on the output of the latch is not a valid logic level. This is a serious limitation if useful data processing takes place between every pair of registers of a general register transfer machine.

In conclusion, even though the realization proposed in this section approaches the theoretical optimum performance, a different structure is required to overcome these practical limitations. The double-filter realization overcomes most of these limitations and, contrary to what the simple argument in the preamble of this section suggests, this solution does *not* require a larger area.

5.3 Double-Filter Latch

5.3.1 Functional Design

A different realization for a latch capable of filtering long transients is considered here. The structure of this latch is shown in Fig. 5.16. It will be demonstrated that this filter has the same lower bound for S as the one based on a single integrator, with the further advantage of a larger discrimination for all values of S.

If the set-up time of the circuit in Fig. 5.16 is SD, a corrupting transient pulse of duration D as shown in Fig.5.6(a) results in a value on the node α at the sampling time given by:

$$\alpha = K(S-1)D - KD = K(S-2)D \qquad (5.12)$$

Therefore if $S > 2$, the signal can be recovered if the threshold of the quantizer is 0. The problem is completely symmetrical with respect to the polarity of the signal and therefore only one polarity needs to be analyzed. If these results are to be compared with those of the earlier section, the output amplitude of the integrators must be limited to the supply voltage. There is no difficulty in having a negative value at α, since it represents the difference of the two inputs of a differential amplifier. The quantizer and the differential amplifier are realized by the same physical device. Limiting the output of the integrators inside the supply region yields:

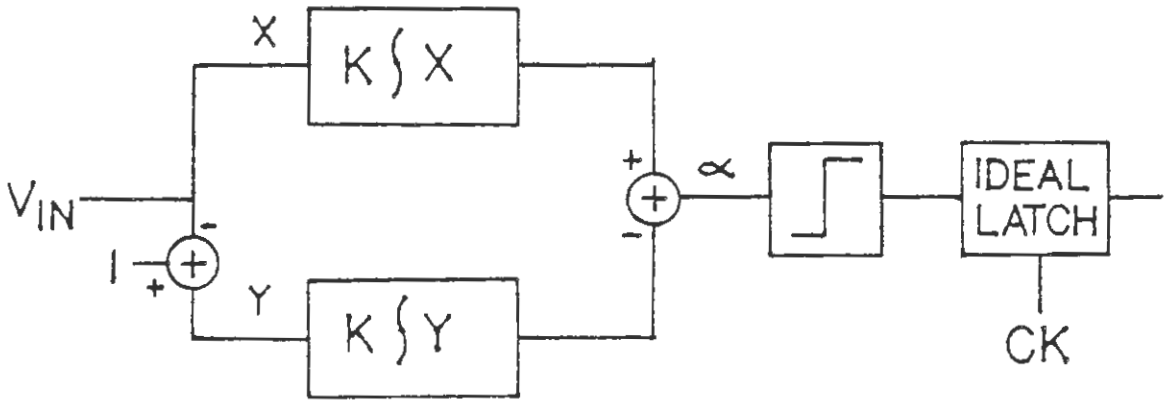

Figure 5.16 Filtering latch with a double-integrator structure.

$$K\,(S-1)\,D < 1 \tag{5.13}$$

$$K\,D < 1 \tag{5.14}$$

If $S > 2$, only (5.13) needs to be considered. Therefore, the maximum value of K is $1/((S-1)D)$, which yields the maximum discrimination for this configuration Δ_4:

$$\Delta_4 = 2\,\frac{S-2}{S-1} \tag{5.15}$$

The factor of two follows from the fact that the input signal with the opposite polarity yields a negative signal of the same amplitude. Comparing (5.15) with (5.4) shows that the double-filter realization yields a discrimination that is twice as large as that of the single-filter realization.

Consider the idealized realization of the filtering section shown in Fig. 5.17, where the integrators are replaced by RC networks. Obviously, in practice, the resistors are implemented with transistors. A transistor in its saturation region can be used to emulate accurately the behavior of an integrator, but in its triode region its behavior resembles that of a RC network. Therefore the performance attainable with this realization based on RC filters can be thought of as a lower bound on the discrimination. Not a bound in the sense that all implementations are better than the one based on an RC filter, but in the sense that a good one should perform at least as well as the best RC realization.

This circuit is nonlinear because of the switches. Both capacitors are first discharged by the signal clear, CL, then the switches CL are left

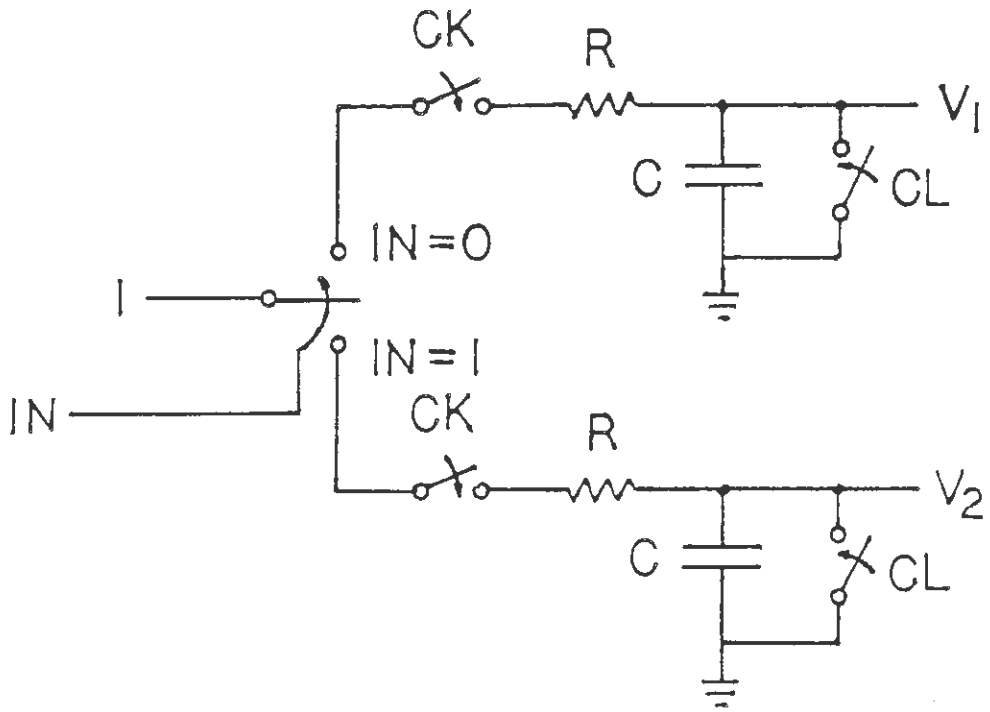

Figure 5.17 A realization of the integrator section of Fig. 5.16 based on switched RC networks.

open. For the total duration of the set-up time, the switches CK are closed. The input signal controls the remaining switch. If the input value is low, the supply is connected to the upper filter, and the reverse is obtained if the signal is high. The branch which is not connected to the supply is left floating. When the input signal is a 1 corrupted by 0, the final values of V_1 and V_2 are:

$$V_1 = 1 - e^{-\frac{D}{R\,C}} \tag{5.16}$$

$$V_2 = 1 - e^{-\frac{(S-1)\,D}{R\,C}} \tag{5.17}$$

The useful signal is the difference between these two quantities, which is given by the following equation:

$$V_2 - V_1 = e^{-\frac{D}{R\,C}} - e^{-\frac{(S-1)\,D}{R\,C}} \tag{5.18}$$

The important design parameter is the ratio U=D/(RC). If U is too small, the deflection and the output signal are small. If the ratio is too large, both signals are far in the exponential response and their difference, which composes the output signal, is also small. This shows the existence of an optimum for U. The procedure for determining it is very similar to the one followed in Section 5.2. Assuming that S is fixed, and taking the derivative of (5.18) with respect to U, after elementary manipulations, yields the following equation:

$$\frac{U}{S-1} = e^{(2-S)\,U} \tag{5.19}$$

This equation has no explicit solution but can easily be solved using Newton's algorithm, yielding the optimum U for a given S. The maximum discrimination Δ_5 of the realization based on a switched RC filter is obtained by multiplying by 2 the result calculated from (5.18), where U is replaced by its optimum value for a given S. Again the factor of 2 follows from the input signal with opposite polarity yielding exactly the same amplitude, but with the opposite sign. The result of this calculation has been plotted in Fig. 5.18, together with the discriminations of all the other configurations. Note the rapid increase of Δ_5 near S=2 which is to be compared with Δ_3 for the single RC filter with precharge. It means that the performance of the double-filter realization is much less sensitive to the exponential response of the RC filter and, therefore, the performance of the transistor realization should suffer much less from operation in the triode region which permits a larger discrimination.

Interest in the double filter realization is further increased because it can be implemented very efficiently. In fact the circuit was discovered by the author before the theory was developed. Consider the conventional level-sensitive D latch shown in Fig. 5.19(a), and the same circuit with two capacitors on the $\overline{S}$ and $\overline{R}$ lines, as shown in Fig. 5.19(b). With CK=0, both capacitors are precharged to the supply voltage. With CK=1, either the $\overline{S}$ or $\overline{R}$ line is ramped to a low value depending on IN. When a transient pulse corrupts the input, the wrong line ($\overline{R}$ for 1 corrupted by 0) starts to ramp exactly like the idealized network in Fig. 5.17, while the 'good' line ramps back toward its precharged value. It is easy to show that this circuit cannot do better than S=3. Removing the second P transistor in the input NAND gates, as in Fig. 5.19(c), neutralizes the ramping back during the pulse, leaving the 'good' line in a floating state. It is easy to see

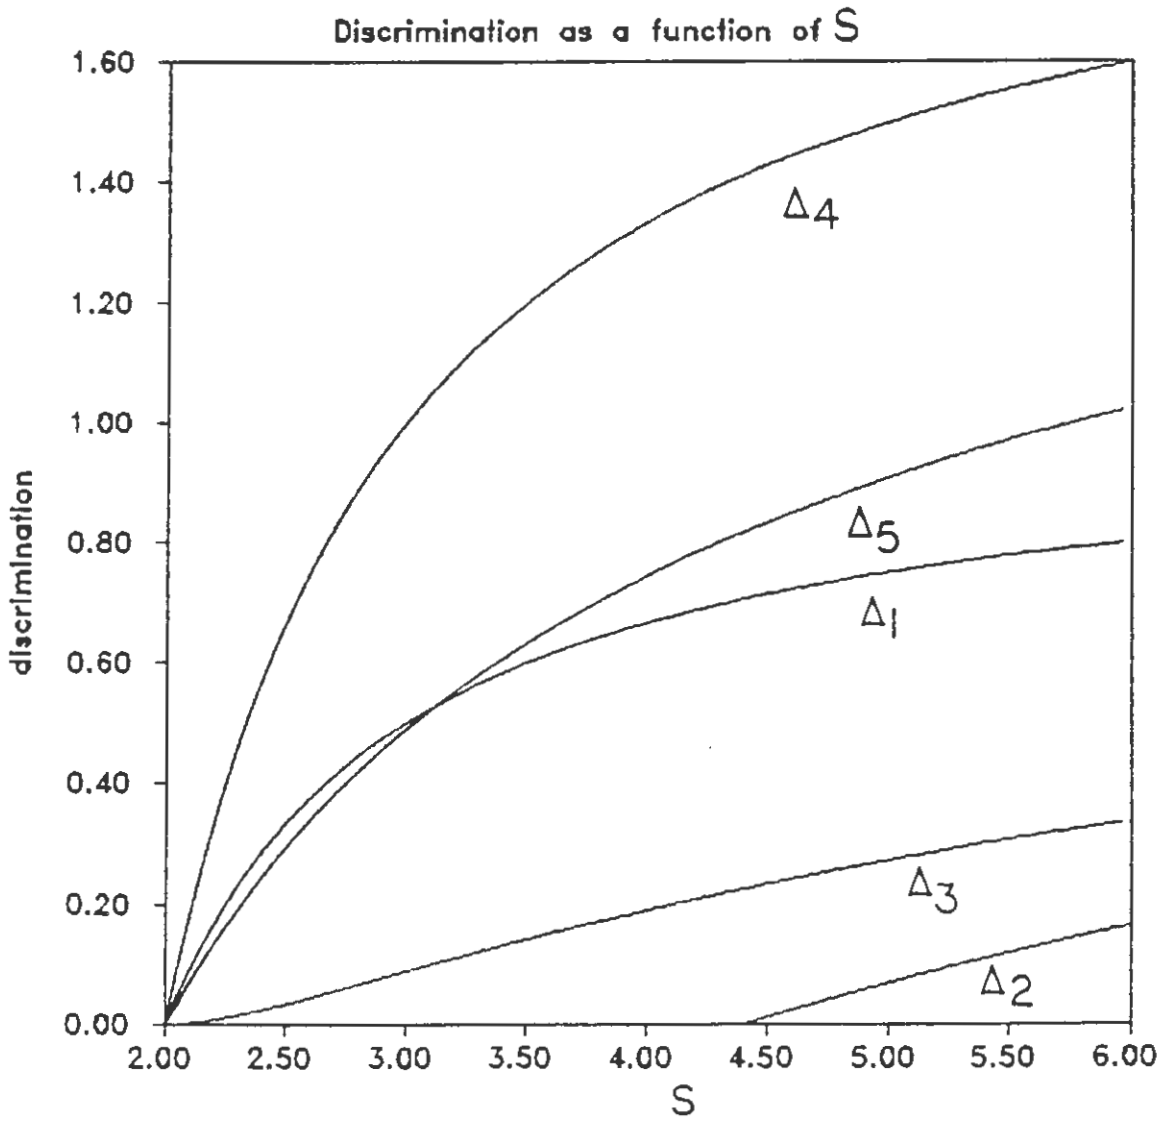

Figure 5.18 Plots of the computed maximum discrimination as a function of the security margin, S, for the single integrator, Δ_1, the single RC filter, Δ_2, the RC filter with precharge (Fig. 5.9), Δ_3, the double integrator (Fig. 5.16), Δ_4, and the double switched-RC filter (Fig. 5.17), Δ_5.

that, except for a polarity inversion, this circuit is functionally equivalent to that of Fig. 5.17.

5.3.2 Implementation of the Double-Filter Latch

The complete double-filter latch design is shown in Fig. 5.20. Two variations have been simulated, with the channel length of transistors 2, 3, 5, and 6 equal to 5μ in one case and 6μ in the other. The best design is the one with 5μ transistors, recovering the signal correctly with an S as low as 2.07 and with a differential signal of only 42 mV. The estimation of S is derived from the two simulation results in Fig. 5.21, where it is shown that a 27ns pulse is recovered with a clock pulse of 56 ns, whereas it is not recovered with a clock pulse of 55 ns. The input signal is a 0 corrupted by one. It is noteworthy that only one polarity needs to be simulated, since the response to a 'set' is completely symmetric to that of a 'reset'; in other words the definitions of 'set' and 'reset' can be exchanged if the output definitions are exchanged.

The results of the simulations are summarized in Table 5.2. The circuit was simulated in three slightly different configurations. For each configuration, the duration of the corrupting pulse on the input was kept constant, and the duration of the clock pulse was varied. The discrimination and the delays to valid levels after the trailing edge of the clock pulse are listed in the table. Table 5.2 includes the performance of a loaded version of the circuit in Fig. 5.20. The load is a capacitor on each output equal to those on the $\overline{S}$ and $\overline{R}$ lines. The loaded version, being slower, requires a larger discrimination in order to recover the signal.

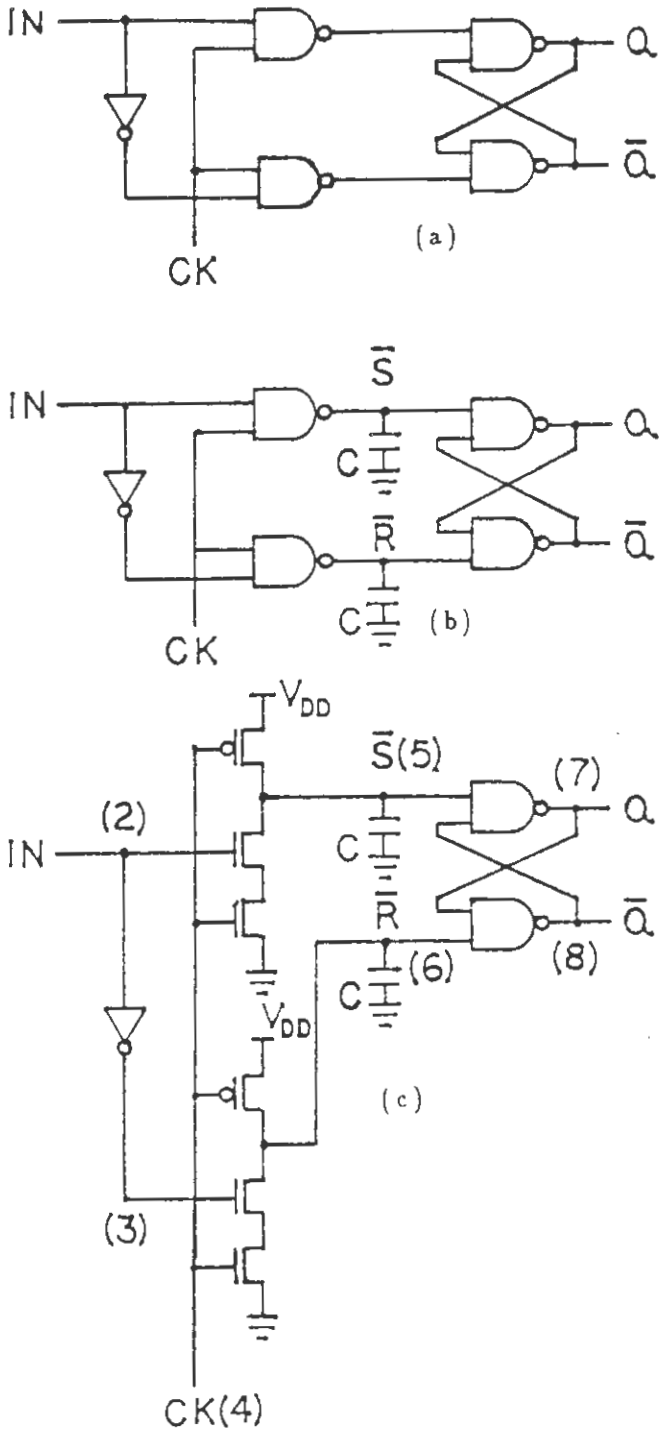

Figure 5.19 Evolution of the double-filter latch. (a) A standard level-sensitive D latch (b) The set and reset lines are used as filters. (c) The input NAND gates are converted to dynamic inverters.

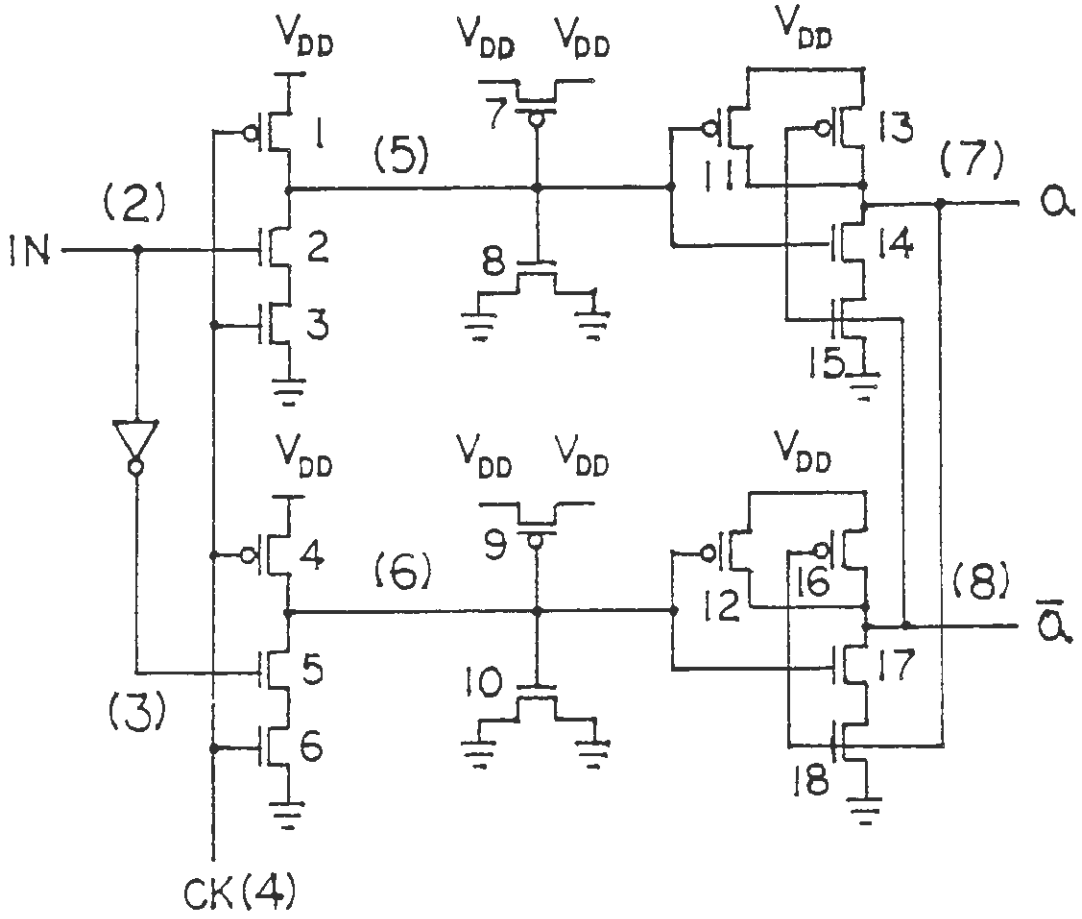

Figure 5.20 Circuit of a practical CMOS double-filter latch. Transistor dimensions (length,width) in microns:
1=4=(5,5), 7=8=9=10=(30,30), 11=12=14=17=(5,60),
3=15=16=18=(5,30), 2=3=5=6=(5,5) for 5 μm version and (6,5) for 6 μm version.

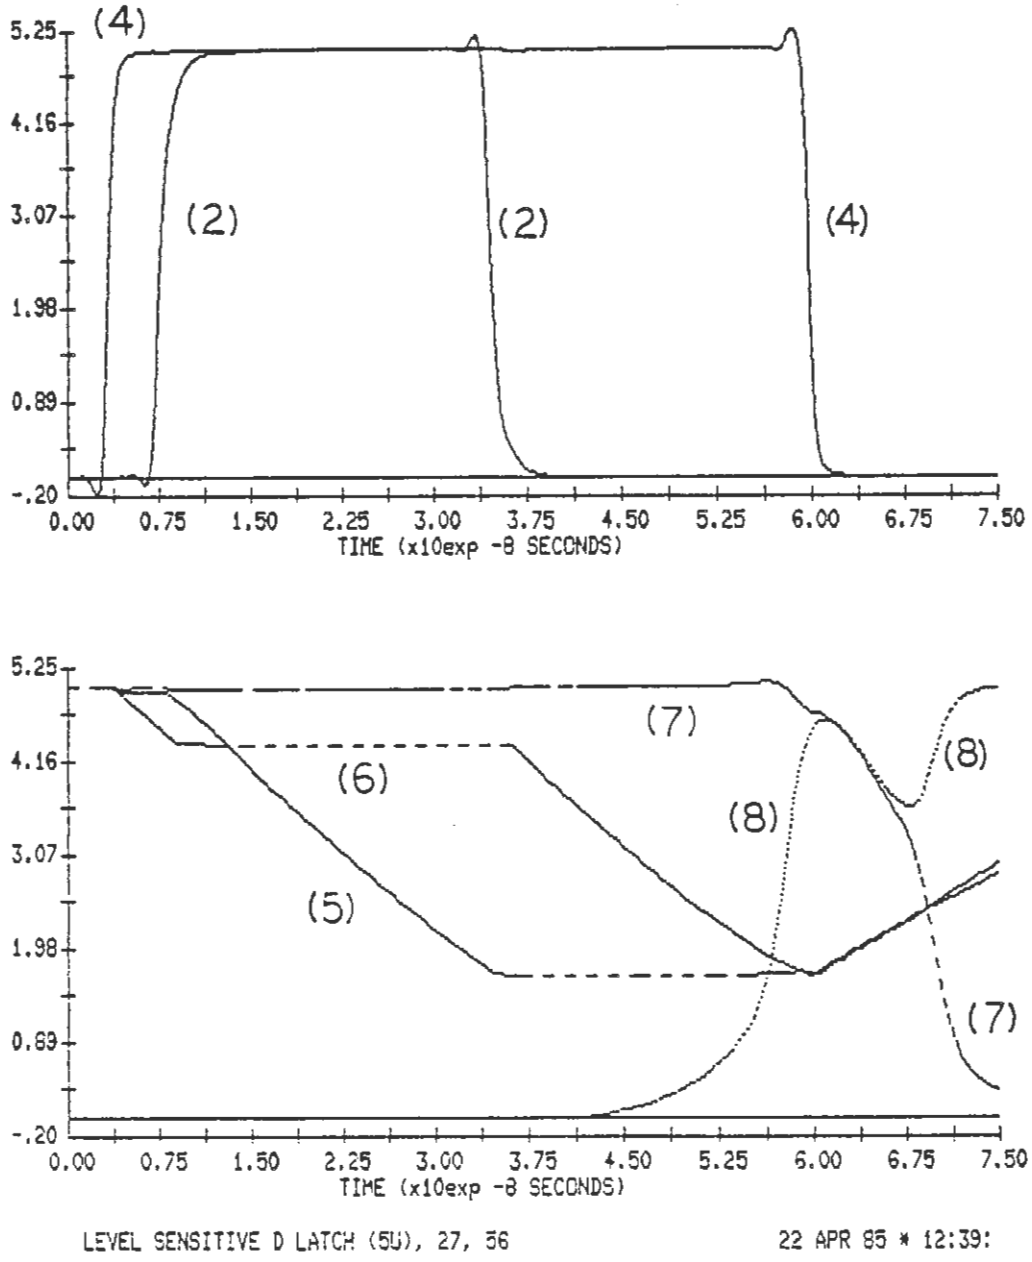

Figure 5.21(a) Simulation results for the 5 μ version. Input signal is 0 corrupted by 1. A clock pulse of 56 ns is sufficient to recover.

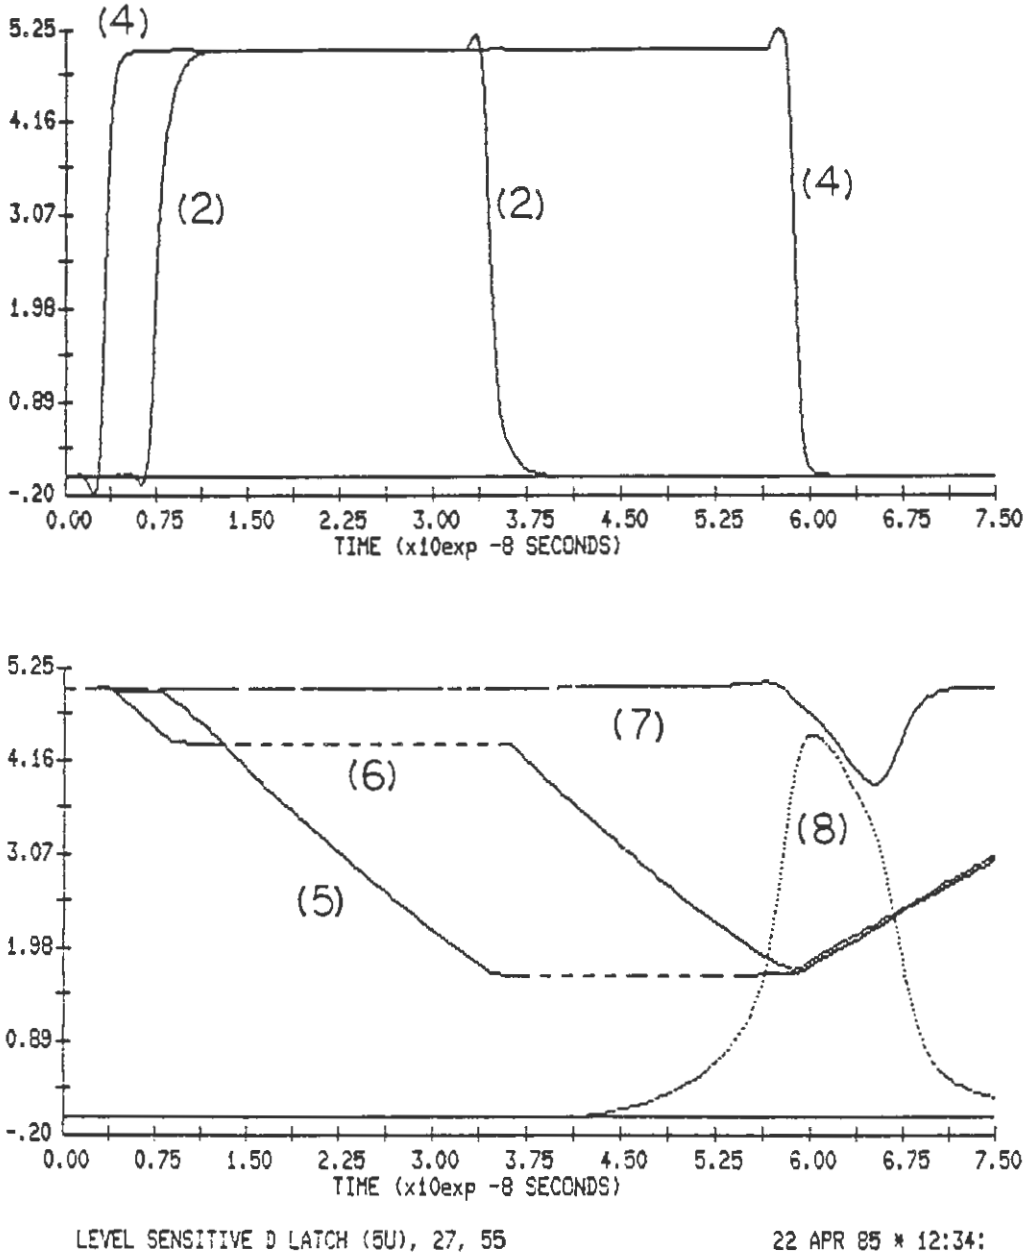

Figure 5.21(b) A clock pulse of 55 ns is not sufficient to recover

Table 5.2
Simulation Results for the Double-filter Latch in Fig. 5.20

CK (ns)	S	Output (V)	Normalized discrim.	Delays* to 1 (ns)	to 0 (ns)
(5 microns, not loaded)					
55	2.04	-.044	-.0176	-	-
56	2.07	.042	.0168	0	9.2
57	2.11	.128	.0512	0	8.7
57.5	2.13	.172	.0696	0	8.4
58	2.15	.193	.0772	0	8.3
(6 microns, not loaded)					
58	2.15	.260	.104	-	-
59	2.19	.328	.131	-	-
60	2.22	.387	.155	0	4.2
(6 microns, loaded)					
62	2.30	.574	.230	-	-
63	2.33	.640	.256	-	-
64	2.37	.707	.283	1.2	14.4
65	2.41	.767	.307	0.2	11.4

* The thresholds for a valid 1 and 0 are taken to be respectively, 2.85V and 1.35V. These values are the input voltages corresponding to the minimum and maximum unit-slope points, on the voltage transfer characteristics of CMOS NOR and NAND gates built with minimum size transistors.

These results are compared in Fig. 5.22 to the theoretical results derived earlier. For a large enough S, the discrimination obtained is a compromise between the integrator and the RC filter with the best performance. Note in Table 5.2 that for S=2.15, the 6μ version has a discrimination of .104, whereas the 5μ version has only .077. Therefore the 6μ version yields a better discrimination; however, it requires S=2.22, whereas S=2.07 is sufficient for the 5μ version, which means that the 5μ version is more efficient.

The smaller discrimination of the 5μ version is a consequence of the larger deflection on nodes 5 and 6, thus the transistors 2, 3, 5 and 6 operate more in the triode region. Consequently, the exponential response begins to play a significant role. However, what makes the 5μ version better are the characteristics of the quantizer. For the given choice of transistor dimensions, the threshold of the latch is such that the gain is higher for the 5μ version. This means that the design could be refined further by shifting the threshold of the latch with 6μ transistors to a higher value. This was not done for two reasons: first, the point of diminishing return on investment has clearly been reached, and second, in doing so, one would neglect the effect of process fluctuations. The observed difference in performance is equivalent to the effect of a 20% difference in the time constant of the filters, which is typical of what could be expected in an integrated circuit.

The 6μ version is more conservative since a real implementation is likely to use S$\approx$2.3 in order to tolerate process fluctuations. The slower version can always be used successfully by lengthening the clock pulse.

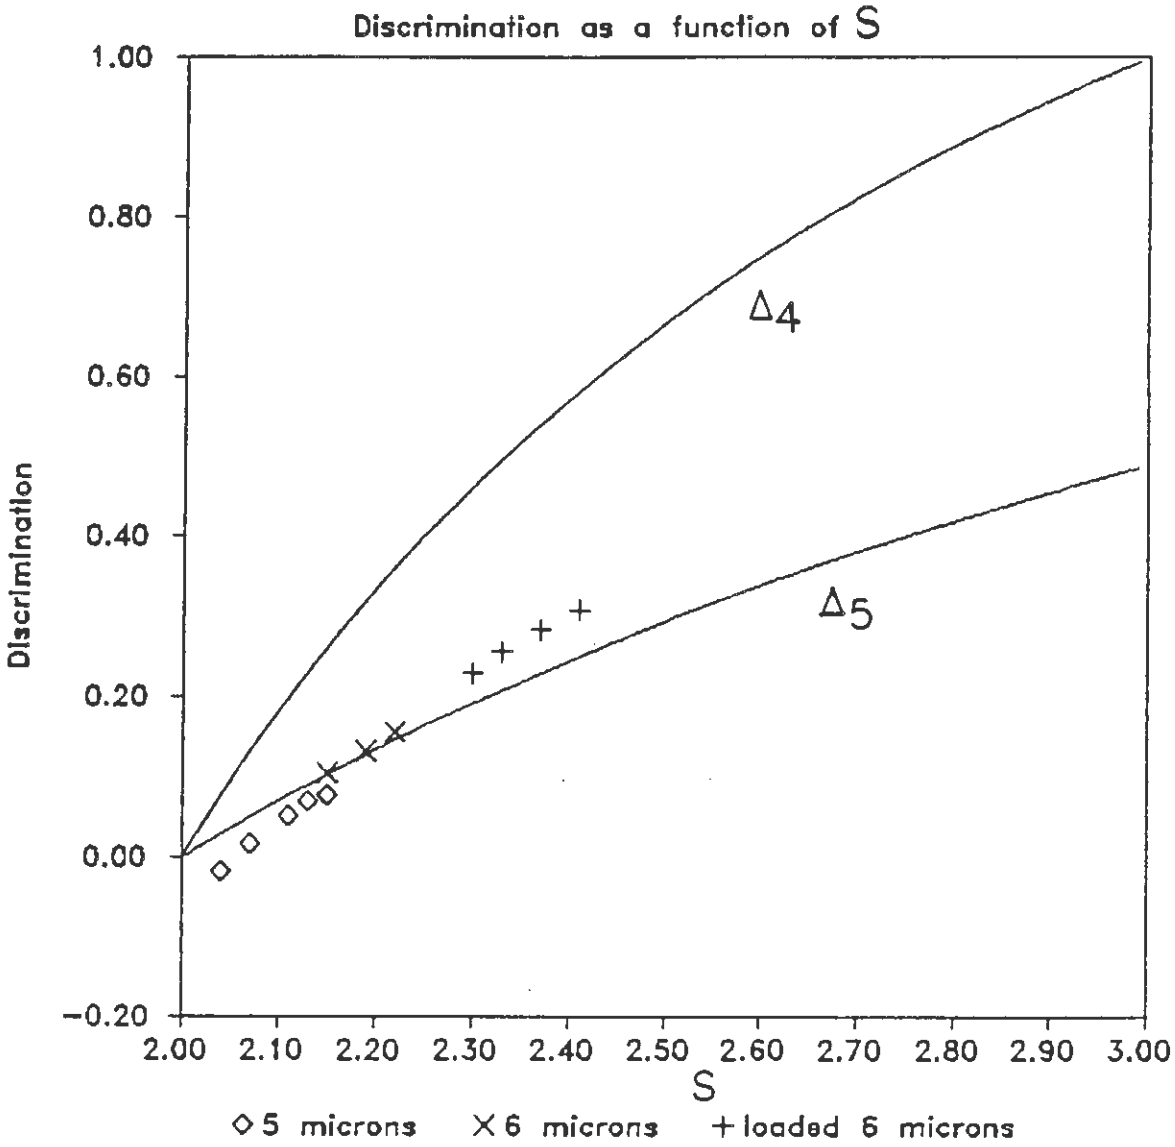

Figure 5.22 Comparison of simulation results for the circuit in Fig. 5.20 with theoretical variation of discrimination with S for the circuits in Figs. 5.16 and 5.17 (Δ_4 and Δ_5, respectively)

However, the faster version can reach the point where the gain of the quantizer falls back again, and making the clock pulse longer does not work.

It is noteworthy that unlike the idealized RC realization in Fig. 5.17, there are not two, but a single clock signal. Consequently the signal recovery task is more difficult, since the precharging of the filters begins immediately at the end of the filtering period, when the signal is ready. However, having a single clock is a definite advantage of this configuration that is desirable in the real implementation.

If the common mode bias on the signal puts the circuit in a low gain region, the quantizer loses its efficiency. For these reasons, the common mode component on the outputs of the filter must not change too rapidly. Slowing down the ramping back also increases tolerance to a transient on the clock line itself, at the expense of a dead time between the successive clock pulses. The reset slope was chosen to be nominally equal to the filter slope.

The slow latch described in Section 3.1 would be very inefficient if it were loaded in the manner described above for the double-filter design. Both the recovery time and S would increase substantially. The essential weakness of the slow latch vis-a-vis the double-filter one is that, in the former, the load capacitance also determines the time constant of the latch and its immunity to direct hits. In contrast, the capacitors at the filter outputs in the latter design set the time constant, while the sizes of the latch transistors are chosen to harden the latch to direct hits and to handle larger capacitive loads.

Chapter6
Overhead Analysis

This chapter is devoted to an analysis of the overhead associated with SEF as it affects area, time and energy. The analysis presented in Section 6.1 demonstrates that SEF is attractive in practice. However, the overhead is very dependent on the function to be implemented. It is also shown in Section 6.2 that SEF generally implies less overhead than other techniques for tolerating transient errors. Section 6.3 concludes this chapter with a number of practical considerations.

To obtain an accurate overhead estimation for a design methodology such as SEF requires a detailed design of a number of systems. However, as will be shown later, even when a detailed implementation is available, it is not always obvious to determine what is overhead. Moreover, SEF can be regarded as a design style, which means that the sources of overhead can be identified at an early stage in the design process, and the details of implementation modified to decrease the overhead significantly.

6.1 Overhead With SEF

There are three important aspects to the analysis of overhead. These are:

1) Area overhead, which is the most obvious, and is discussed in Section 6.1.1. It is obtained by comparing the area occupied by

the redundant digital machine with that occupied by an equivalent non-redundant one.

2) The time overhead, since tolerant machines are often slower. This is discussed in Section 6.1.2.

3) The energy overhead, discussed in Section 6.1.3.

6.1.1 Area overhead

The area overhead is a measure of the amount of hardware redundancy. In an integrated circuit context, it is more realistic to measure overhead in area than in transistor count or gate count. An overhead analysis based on a gate count as in DasGupta et al. [Das82] neglects the fact that more than 50% of the area of a chip can be reserved for interconnections (bus and pads). Moreover, the only meaningful basis for comparison is area, when various types of logic structures are used in the same machine, such as an ALU, some PLAs, random logic, and the special registers proposed in Chapter 5. A consequence of measuring overhead in terms of area is that two independent realizations of exactly the same machine could result in significantly different overheads.

The area overhead for SEF is given by the following equation:

$$O_A = \frac{A_1 - A_2}{A} \qquad (6.1)$$

where A_2 is the area occupied by the standard registers, A_1 is the area occupied by the filtering registers, and A is the total area of the machine with standard registers. There are a few difficulties here, the first one resulting from the fact that if the standard machine could be dynamic (Domino

for example), except for the use of SEF, the overhead would be larger, and an accurate estimation would require a detailed design of the two machines. Another difficulty arises because, even though a filtering register requires no more global interconnection than a standard static D flip-flop, the bigger size of the latches causes the SEF machine to be larger, therefore longer interconnections are required. Again there are no simple means of estimating the impact of the bigger size of the registers on interconnection area, except by a detailed design of two versions of a machine. Since reoptimizing two versions of a machine just for the sake of estimating overhead is too expensive with current design tools, (6.1) is used.

In the following, a refinement of (6.1) is obtained. The derivation is similar to the overhead analysis for Level Sensitive Scan Design [Das82], and assumes that a function is implemented as a network of gates. The expression to be derived depends on four parameters:

Q: The ratio of the area occupied by one bit of filtering register to that of a 2-input gate.

K: The ratio of the number of gates needed to realize the combinational logic part of a machine to the number of memory bits needed.

C: The fraction of the area reserved for communication of data. It includes the area reserved for global routing plus that reserved for input/output pads.

R: The ratio of the area for one bit of a SEF register to that of one bit of a standard register.

From these definitions, the area of the non-SEF machine, A, measured in equivalent gates and normalized for one bit of memory, is given by

$$A = \frac{K + \frac{Q}{R}}{1 - C} \tag{6.2}$$

The difference between the area of one bit of filtering register and that of one bit of a standard register is the overhead per bit in the machine. This quantity is expressed in equivalent gates as

$$\text{overhead} = Q - \frac{Q}{R} \tag{6.3}$$

Dividing (6.3) by (6.2) gives the area overhead, O_A

$$O_A = \frac{(1 - C)\, Q\, (1 - \frac{1}{R})}{K + \frac{Q}{R}} \tag{6.4}$$

Equation 6.4 becomes interesting if upper and lower bounds on each of the parameters are known. The complexity of the filtering latch shown in Fig. 5.19 is at least that of 4 equivalent gates, and it should be possible to realize a layout smaller than the area of 8 gates, therefore $4 < Q < 8$. From DasGupta [Das82], the number of gates in the combinational logic per memory bit is usually in the interval $5 < K < 25$. It is well known that the fraction of a chip reserved for communication can be more than 50%, but it can also be as low as 20 % for very regular structures, therefore $0.2 < C < 0.5$. Finally, considering Fig. 5.19 again, the area occupied by a filtering latch should be from twice the size of a standard level-sensitive D flip-flop, to somewhere around 5 times the complexity of a C^2MOS latch, consequently $2 < R < 5$.

Using the upper and lower bounds for Q, K, C, and R, in (6.4), yields

upper and lower bounds for the area overhead. A typical value can also be calculated by using the medians of these bounds in (6.4). These calculations yield a typical value of 17% and an interval of $3.7\% < O_A < 78\%$. The interval for the area overhead is wide and obviously depends on the type of function being implemented. When a machine has relatively few memory elements, the area overhead is small.

A category of machines exists where the overhead estimation obtained above is not valid. Consider machines such as RISCs [FIT81], based on large arrays of registers and very simple control logic. To implement large arrays of memory with SEF latches is not practical. Therefore a straight application of SEF as described in Chapters 4 and 5 is not realistic. Machines of this category would require a modification of the architecture before implementation. A possiblity, mentioned earlier, is to use a coded register array.

6.1.2 Time overhead

Conventional methods of tolerating errors exist, whereby area overhead is traded for execution time. The equivalent tradeoff exists with SEF, as will be discussed later in Section 6.3. However, SEF can be implemented with a low overhead in area and in time simultaneously. The time overhead is defined to be the ratio of the difference between the clock periods of a SEF and a non-SEF machine, to the clock period of the non-SEF machine. Since only the set-up times of the registers are different in a SEF and a non-SEF machine, the time overhead O_T is given by the following equation.

$$O_T = \frac{S_1 D - S_2 \delta}{T} \qquad (6.5)$$

Where the set-up time of the SEF machine is S_1D, and the set-up time of the non-SEF machine is $S_2\delta$ (here δ is the longest event that a conventional latch will not memorize). It is clear from Chapter 5 that S_1 and S_2 are not equal in general, since the value of S depends on the structure of the latch. Also, from Chapter 4, D must be greater than P, the longest expected transient.

The necessary condition for SEF to yield a low time overhead is now known; the duration of the longest expected transient must be smaller than the clock period of the machine. For example, if the longest expected transient is 4 ns, and filtering latches are used with S slightly larger than 2, then transforming a standard machine with a 50 ns clock period into an SEF machine, results in a time overhead of approximately 20%.

The time overhead of a SEF machine varies enormously. A time overhead of a few percent is possible if the clock period T is greater than 100 ns, and P is on the order of 1 ns. At the other extreme, a transient at the output of the logic could become longer than the clock period, if either the injected transient is long or the pulse spreading is important, as discussed in Section 2.3. Using SEF for combatting such long transients could result in a time overhead of more than 200%.

The preceding discussion demonstrates that although the time overhead could be small, it is still highly dependent on the machine to be hardened, and on the transient source to be combatted. However, as will be discussed in

Section 6.3, the general SEF approach can be specialized to reduce the time overhead further.

6.1.3 Energy overhead

The last aspect of the overhead analysis is the energy per computation. This is important in any situation where the power supply is limited, like in space applications. It is clear that the design style (dynamic or static, optimized for low power consumption or optimized for speed, etc.) has a strong impact on the power consumption of a machine. Therefore, the following discussion is only meaningful if similar design styles are adopted for the SEF and non-SEF machine.

If the SEF machine is built with a CMOS technology, the energy per computation in the combinational logic is nominally unchanged. However, intrinsic tolerance of the latches is achieved by keeping their switching energy above a critical level. As a first step for calculating the energy overhead, O_E, the fraction of the total energy dissipated in the latches of a conventional machine must be evaluated. This fraction is given by

$$\frac{L}{KG + P + L} \tag{6.6}$$

where K: The number of gates in the combinational logic section divided by the number of memory bits in the machine

G: The average energy dissipated by a gate

P: The energy dissipated in the I/O pads, divided by the number of memory bits in the machine.

L: The energy dissipated by a conventional latch.

If E is the ratio of the energy dissipated by a SEF latch to that of a standard latch, and since only the energy dissipated in the latches can be counted as overhead, then O_E is given by

$$O_E = \frac{(E - 1)\, L}{KG + P + L} \tag{6.7}$$

In principle, it is easy to estimate ranges for E, L, G, and P, but in practice they are very dependent on the details of implementation. Therefore, it would be necessary to derive these estimates from data collected on real designs, and since such data is not available, no attempt is made here to obtain a numerical range for O_E. Despite the lack of quantitative knowledge for the various parameters, (6.7) is interesting because it demonstrates that again, in this case, the overhead can be small when the main limiting factor is not the latches but something else, for example, the energy dissipated in the logic network or in the I/O connections.

If a SEF machine is built with a technology that dissipates a significant amount of DC power, like NMOS or pseudo NMOS implemented in a CMOS technology, the energy overhead can be expressed as follows:

$$O_E = \frac{(E - 1)\, L}{KG + P + L} + \frac{E_{DC} T\, O_T}{E_{DC} T + P_{SW}} \tag{6.8}$$

Here E_{DC} and P_{SW} are, respectively, the standby dc energy and the switching power of the conventional machine. The second term takes into account the situation where SEF increases the DC consumption of the machine

by increasing the clock period. Therefore an SEF machine that dissipates a significant amount of DC power has a high energy overhead if its time overhead is high.

6.2 Comparison With Alternatives

This section considers the overhead implied by the alternative solutions described in Chapter 2. These alternatives include: intrisic tolerance, tightly coupled DMR, loosely coupled TMR, tightly coupled TMR, and tightly coupled TMR hardened for bursts of transients. All these alternatives allow one to build a machine tolerant to independent transients injected at an exponential interval. The error rates are not exactly the same in each case, but the differences are not significant, considering the failure rate. However, when bursts of transients are expected, the two basic TMR schemes are not really appropriate, therefore overhead considerations are of secondary importance.

To evaluate the cost of intrinsic tolerance to transient errors in general is not possible, because it depends too much on the particular situation. After all known inexpensive techniques of decreasing the error rate have been applied, if a significant error rate still remains, as is usually the case, there exists a point beyond which it is less expensive to use system solutions like TMR, than to achieve the required reliability level by increasing intrinsic tolerance.

The solution of increasing the power per gate, for example, requires expensive cooling techniques, and these may well introduce reliability

hazards. Moreover, the system will have to be implemented with many more chips, with a direct consequence on the cost and speed of the machine. Also, this solution is only efficient for combatting ionizing radiation, and ignores the effects of interference.

It is easier to compare SEF with the system solutions, because overhead can be quantified more easily. The easiest is loosely coupled TMR as in Fig. 3.3(a). A natural means of implementing a loosely coupled TMR machine is to use three off-the-shelf modules in parallel, and vote only on the final results. It implies a 200% area overhead for the two redundant machines, and the voter also contributes a small overhead in terms of gate count.

The overhead in the voter remains small because only the final outputs are compared. For example a 10000 gate system may have 40 output lines. A voter can be implemented with an equivalent complexity of 4.5 gates (The carry line of a 1 bit full adder is a voter, and the carry can be generated with 18 transistors [MAV83 p.92], and 4 transistors are counted as 1 equivalent gate). If only the overhead contributed by the gates were counted, voting on 40 lines would require a 2% overhead.

However, the real overhead, in merging 3 times 40 outputs to obtain the final result, is not that of the gates themselves. Three input pads and one output pad are required for each bit of the voter. With a 5μ technology one pad occupies an area larger than 10 gates. Moreover, when the technology is scaled down, the discrepancy between the size of a pad and that of a gate grows as the square of the scaling factor. Therefore the communication cost of voting is at least 10 times that of the gates which perform the voting operation.

Another practical consideration too easily overlooked is that, for a chip with a low gate-to-pin ratio, such as a voter chip, the package is usually more expensive than the silicon die. Therefore voting may end up costing 50% or more of the cost of the original non-redundant chip, which is higher than the typical area overhead contributed by SEF.

The loosely coupled TMR presents the advantage of not causing any time overhead, because voting is not in the feedback path. Finally, the energy overhead is larger than 200%.

Estimating the overhead for tightly coupled machines becomes more difficult. Apart from the communication costs, tightly coupled DMR and TMR respectively require at least 100% and 200% overhead respectively. However, it was clearly demonstrated earlier that in an integrated realization of a fault-tolerant machine, the communication cost of comparing and voting is not negligible. While the number of lines to be compared was small for loosely coupled TMR, the same is not true for tightly coupled machines. Therefore the overhead of tightly coupled machines could become much higher than it appears from Figs. 3.1, 3.3(b), and 3.4.

It is difficult to derive a generally valid estimate of overhead, because the communication cost can become so high as to force a redesign of the system in directions which cannot be quantified accurately. For example, a possible alternative that gives obvious benefits, is to partition the machine in order to minimize the number of times a signal must exit the chip for comparison and voting. The implications of partitionning in terms of overhead could be estimated with Rent's rule [MUR82], which relates the

number of gates to the average number of I/O required.

Despite the difficulty mentioned earlier, it is noteworthy that when K, the number of gates per bit, decreases, the relative importance of the overhead introduced for comparing and voting grows for tightly coupled machines. Therefore, when the area overhead of SEF is high, comparators and voters also contribute a high overhead in tightly coupled machines.

The time overhead of a tightly coupled machine is not zero. A sufficient time for comparing and voting must be allowed, since it takes place on the feedback lines of the machine. Moreover, when the machine has to be implemented on more than one chip, the communication delays contribute further to the time overhead. Nevertheless, in contrast with SEF, this time overhead is not a function of the expected transient duration. In a first approximation, the energy overhead for the tightly coupled alternatives should be proportional to the amount of hardware redundancy.

In conclusion, the area overhead for SEF is comparable or smaller than the communication costs of system alternatives. Therefore not duplicating nor tripling the logic function is an advantage for SEF. However, the system solutions prevail over SEF if the machine must tolerate long transients. But if the transients are short, SEF keeps the time overhead comparable or smaller. If CMOS is used, or if the time overhead is small, the energy overhead should always be smaller for SEF. Therefore, considering simultaneously the time, area, and energy overheads, SEF is the best approach for building a machine tolerant to short transients.

6.3 Practical Considerations

SEF as presented earlier is a general approach that may be improved when adapted to a particular situation. The context may easily change the relative importance of the different types of overhead. For example, time overhead may be more significant than area overhead in some situations, or conversely. Moreover this can be true even for different sections of the same machine. For example, the time overhead is not important for an output which is not on the critical propagation path. The importance of this observation can be amplified if pulse spreading is very different from one output line to the other. Therefore, a technique for exchanging time overhead for area overhead would be useful.

Conventional techniques exist for transforming a critical propagation path into a non-critical one, for example, buffering may be provided on the slow nodes, or pipelining may be introduced in order to allow more than one clock cycle for data propagation along the critical path. In addition to the standard techniques, *delay equalization* is a technique that permits one to exchange time overhead for area overhead. The author demonstrated in an earlier paper [SAV84b] that, for a technology with equal rise and fall times, an arbitrary logic function can always be redesigned with the same worst case propagation delay, but nominally with no pulse spreading. This is achieved by adding delays on the faster propagation paths, in such a way that all path delays are equalized.

One aspect of SEF requires special attention: the duration of transients generated by interference. These transients can be so long that SEF becomes impractical because of excessive time overhead. However, if SEF

is used as a complement to standard electromagnetic shielding, it will handle efficiently the remaining short transients that may exist due to imperfections of the shield, or that may be generated inside the shield. Therefore SEF should be supported with appropriate shielding.

Another important practical consideration with SEF is that transients injected by interference outside the chips, are usually much longer than those injected by ionizing radiation inside the chip. Therefore, the set-up time of SEF latches should reflect this in order to minimize the time overhead. The limitations imposed by the relatively long delays required for off-chip connections are usually recognized in VLSI systems, where a single bit can pass through a pin at each clock cycle. That cycle could be separated into a number of internal microcycles. For example, the clock period could be 50 ns with a set-up time of 10 ns for latches that include in their fan-in a connection from outside the chip. The same machine could have an internal microcycle of 10 ns with internal set-up times of 2 ns, resulting in a time overhead of only 20%, even though some latches have a set-up time as long as the microcycle. This architecture would allow to filter out transients as long as 4 ns at the board level.

Finally, overhead can be wasted if it is required that the SEF latches are to be capable of tolerating every transient resulting from a single hit. If the combinational logic block has a relatively small number of outputs which terminate paths with long propagation delays, and if the fan-in of these outputs includes nodes that are the sources of transients much longer than those appearing on the rest of the outputs, a large time overhead results. The time overhead is large because the clock period must be longer than the sum of long set-up times plus long propagation delays. If this

situation occurs with a sufficiently small probability, using SEF latches optimized for the shorter transients can result in a soft-error rate less than the failure rate. Thus, to ignore the long transient would be justified in such a case and would, of course, lead to a lower area and time overhead.

Chapter 7
Conclusions and Further Work

7.1 Conclusions

The first contribution of this thesis is to unify in one document the literature on the characterization of soft-error sources in digital machines. The characteristics of electrical noise, ionizing radiation, and electromagnetic interference were reviewed. It has been demonstrated that electrical noise should never be significant. It is noteworthy that SEF would be very efficient to combat electrical noise as a potential source of soft errors, if a technology sensitive to its effects is ever developed. The direct relationship that exists between the duration of a transient and the bandwidth required to propagate it makes the probabilty of longer events much smaller (decreasing exponential relationship). A comparison between (permanent) failure rates and error rates due to ionizing radiation was developed to demonstrate the significance of the latter as a source of soft errors.

The second important contribution is the recognition of the fact that conventional fault-tolerance techniques may not be the most efficient way to tolerate soft errors. This led to the proposal of tightly coupled DMR and TMR machines in Figs. 3.1 and 3.4. But it also led to the main contribution of this thesis, which is the Soft-Error Filtering technique. The error rate reduction achieved with this technique has been analyzed for different soft error sources, in order to demonstrate how SEF can make the error rate negligible. A great deal of attention has been devoted to the design of

filtering latches, because they are essential components of SEF machines, and also because they are the main source of overhead. The present work on filtering latches resulted in the proposal of a double-filter latch, which can be implemented efficiently in CMOS. The simulated performance of the proposed implementation is nearly optimum and, the implementation is relatively insensitive to fluctuations of the fabrication process. Finally, an overhead analysis supports the significance of SEF, by demonstrating that it is feasible with less overhead than conventional fault-tolerance techniques. It should be stressed that SEF permits a low overhead in hardware and in time simultaneously.

When short transients are expected at a sufficient rate to cause a significant error rate, SEF is the most appropriate tolerance technique applicable. Another potential and very interesting application of SEF is to enhance the reliability of machines with a degraded noise margin. As mentioned in Chapter 2, aging and gamma ray exposure both reduce the noise margin.

7.2 Suggestions for Further Research

This thesis is, to the best of the author's knowledge, the first work specifically dedicated to soft error tolerance in logic circuits. Moreover, it is an interdisciplinary work, touching on many research fields, including: the interaction of radiation with matter, the electromagnetic compatibility of electronic circuits, the theory of reliable communication systems, the design of integrated circuits and systems, and finally logic design for

fault-tolerance. Therefore, several extensions of this work are possible.

A first domain where further research could confirm the basis of the theory proposed in this thesis is the interaction of ionizing particles with logic circuits. This work is based on the measured characteristics of the injected current pulses on simple PN junctions by alpha particle hits. Extrapolating these measurements to complex multilayer structures with submicron feature sizes in different technologies is not obvious. Therefore, more experimental data are required. For example, there remains the question of the extent to which a bipolar structure will amplify an injected charge, as a function of the polarization and the device geometry. Other fundamental data that are necessary to quantify the effect of cosmic rays are the characteristics and the distribution of nuclear cascades, very near their propagation axis, as a function of the shield used.

The characterization of pulse spreading is another area of future study. If a transient is injected on any internal node of a machine, what are its characteristics at the input of the latches after propagation? Such statistics were not needed previously; therefore, these fundamental data are not available. Moreover, special techniques or tools for estimating the pulse spreading in a circuit being designed are needed, in order to achieve a reliable implementation of SEF with as low an overhead as possible.

Furthermore, as mentioned in Chapter 6, it is possible to exchange time and area overhead with *delay equalization*, in order to reduce the time overhead. However, for this technique to be practical, either a strict design methodolgy that yields circuits with low pulse spreading, or a silicon compiler capable of equalizing the delays in a circuit are needed. This is

another possible research direction.

Much more work is needed on the filtering latch, because only CMOS realizations were considered here, whereas SEF could be implemented with other technologies. Also, even though the performance achieved by the double-filter latch presented in Chapter 5 is almost optimum, a better CMOS implementation of the filtering latch is still possible. Needless to say that all these designs must be implemented in silicon and tested.

Finally, in practice, the most important complementary work is in the domain of overhead analysis. At this stage, it is necessary to prove that SEF is a practical technique by implementing integrated circuits of reasonable complexity with built-in SEF. With that goal in mind, this author has already undertaken the redesign of an existing microprocessor. This experiment should demonstrate clearly that SEF is indeed a practical technique.

References

[AMB82] A. Ambrozy, *Electrical noise*, McGraw Hill 1982

[AND81] T. Anderson, and P. A. Lee, *Fault tolerance, principles and practice*, Prentice Hall 1981.

[ANO79] E. S. Anolick, and G. R. Nelson, "Low field time dependent dielectric integrity", *IEEE Proc. of the International Reliability Physics Symposium*, 1979, pp. 8-12.

[CMC85] Canadian Microelectronics Corporation, *Guide to the integrated circuit implementation services of the Canadian Microelectronics Corporation*, version 1.0 march 1985.

[Das82] S. DasGupta, P. Goel, R. G. Walther, and T. W. Williams, "A variation of LSSD and its implication on design and test pattern generation", *IEEE International Test Conference* 1982 pp. 63 - 66.

[DAV82] R. T. Davis, M. H. Woods, W. E. Will, and P. R. Measel, "High-performance MOS resists radiation", *Electronics*, Nov. 17,1982, pp. 137-139.

[DAV83] R. D. Davies, "The case for CMOS", *IEEE Spectrum*, vol.20, pp.26-32, Oct. 1983.

[DEN74] R. H. Dennard, F. H. Gaensslen, H-N Yu, V. L. Rideout, E. Bassous, and A. R. LeBlanc, "Design of ion-implanted MOSFETs with very small physical dimensions," *IEEE Journal of Solid State Circuits* vol. SC-9, pp. 256-268, October 1974

[DEN79] R. H. Dennard, F. H. Gaenslen, E. J. Walker, and P. W. Cook, "1um MOSFET VLSI technology: PartII-Device designs and characteristics for high-performance logic applications", *IEEE Journal of Solid-State Circuits*, vol. SC-14, pp. 247-255, April 1979.

[FIS82] M. A. Fischetti, "VHSIC contractors tell their story", *IEEE Spectrum*, vol. 19, no. 12, Dec. 1982, pp.36-38.

[FIT81] D. T. Fitzpatrick et al., "VLSI implementation of a reduced instruction set computer", in H. T. Kung, B. Sproul, and G. Steele, *VLSI systems and computations*, Computer Science Press, 1981, pp. 327 - 336.

[FLI81] S. J. Flint, and L. L. Kent, "Electronic engine control: Auto makers contend with one of the harshest environments", *IEEE Spectrum*, vol. 18, no. 10 Oct. 1981, pp. 61 - 62.

[GHA82] P. B. Ghate, "Electromigration-induced failures in VLSI interconnects", *IEEE International Reliability Physics Symposium* 1982 pp. 292-299.

[GHE84] T. R. Gheewala, "System level comparison of high speed technologies", *IEEE International Conference on Computer Design*, 1984, pp. 245-250.

[GIB66] J. F. Gibbons, *Semiconductor Electronics*, McGraw-Hill 1966.

[HAG74] G. H. Hagn, and R. A. Shepherd, "Man-made electromagnetic noise from unintentional radiators: a summary", *AGARD Conference Proceedins No.159*, pp. 3-1 3-23, 1974.

[HO82] P. S. Ho, "Basic problems for electromigration in VLSI applications", *IEEE Proc. of the International Reliability Physics Symposium*, 1982, pp.288-291.

[HOD83] D. A. Hodges, and H. G. Jackson, *Analysis and design of digital integrated circuits*, McGraw-Hill, 1983.

[HSI81] C. M. Hsieh, P. C. Murley, and R. R. O'Brien, "Dynamics of charge collection from alpha-particle tracks in integrated circuits", *Proc. of the International Reliability Physics Symposium* 1981, pp.38-42.

[JEC79] R. M. Jecmen, C. H. Hui, A. V. Ebel, V. Kynett, and R. J. Smith, "HMOSII static RAMs overtake bipolar competition" *Electronics*, vol. 52, pp. 124-128, Sept.13,1979.

[LAP72] R. E. Lapp, and H. L. Andrews, *Nuclear radiation physics, Fourth edition*, Prentice-Hall 1972.

[LIU82] S.-M. S. Liu, C.-H. Fu, G. Atwood, H. Dun, J. Langston, E. Hazani, E. Y. So, S. Sachdev, and K. Fuchs, "HMOSIII technology", *IEEE Journal of Solid-State Circuits*, vol. SC-17, pp.810-815, Oct. 1982.

[MAR84] M. Mardiguian, *Interference control in computers and microprocessor-based equipment*, Don White consultants, 1984.

[MAV83] J. Mavor, M. A. Jack, and P. B. Denyer, *Introduction to MOS LSI design*, Addison-Wesley 1983.

[MAY78] T. C. May, and M. H. Woods, "A new physical mechanism for soft errors in dynamic memories", *Proc. of International Reliabilty Physics Symposium* 1978, pp. 33-40.

[McC79] S. R. McConnel, D. P. Siewiorek, and M. M. Tsao, "The measurement and analysis of transient errors in digital computer systems", *Proc. of the International conference on Fault Tolerant Computer Systems*, 1979, pp. 67-70.

[McC81] S. R. McConnel, *Analysis and modeling of transient errors in digital computers*, PhD dissertation, Carnegie Mellon University 1981.

[MEA80] C. Mead, and L. Conway, *Introduction to VLSI systems*, Addison Wesley, 1980

[MEI79] E. S. Meieran, P. R. Engel, and T. C. May, "Measurement of alpha particle radioactivity in IC device packages", *Proc. of the International Reliability Physics Symposium*, 1979, pp.13-22.

[MOR84] H. Morkoc, and P. M. Solomon, "The HEMT a superfast transistor", *IEEE Spectrum*, vol. 21, Feb. 1984, pp. 28-35.

[MOT73] C. D. Motchenbacher, and F. C. Fitchen, *Low-noise electronic design*, John Wiley, 1973.

[MUR82] S. Muroga, *VLSI system design* Wiley 1982.

[NAG79] E. Nagasawa, H. Okabayashi, T. Nozaki, and K. Nikawa, "Electromigration of sputtered Al-Si alloy films", *IEEE Proc. of the International Reliability Physics Symposium*, 1979, pp. 64-71.

[NEW74] M. M. Newman, and J. D. Robb, "Atmospheric discharges and noise (and communication systems interference reduction)", *AGARD Conference Proceedings No. 159*, pp. 2-1 2-21, 1974.

[PAP65] A. Papoulis, *Probability, random variables, and stochastic processes*, McGraw Hill 1965.

[PEA81] G. Peattie, "Quality control for ICs", *IEEE Spectrum* vol. 18, no. 10, pp. 93 - 97.

[PEA83] R. F. W. Pease, "Fabrication issue for next generation circuits", *IEEE Spectrum*, vol. 20, Nov. 1983, pp. 102-105.

[PRO83] J. G. Proakis, *Digital communications*, McGraw Hill 1983.

[RAM84] S. Ramaswamy, L. Nguyen, T. Brooks, and A. Gokhale, "Simultaneous switching noise analysis in VLSI", *IEEE International Symposium on Circuits and Systems vol.2*, 1984, pp. 706-709.

[REE70] I. S. Reed, "Error tolerant sequential circuits", US Pat. no. 3529141 Sept. 1970 (filed in Sept. 1967)

[ROO84] S. A. Roosild, "DARPA GaAs plans and pilot production line project", *Proc. of the International Conference on Computer Design*, pp.251-257, 1984

[SAI82] G. A. Sai-Halasz, M. R. Wordeman, and R. H. Dennard, "Alpha-particle-induced soft error rate in VLSI circuits", *IEEE Journal of Solid-State Circuits*, vol. SC-17, April 1982, pp.355-361.

[SAR84] D. B. Sarrazin, and M. Malek, "Fault-tolerant semiconductor memories", *IEEE Computer*, vol.17, number 8, august 1984, pp. 49 - 56.

[SAV84a] Y. Savaria, V. K. Agarwal, N. Rumin, and J. F. Hayes, "A design for machines with built-in tolerance to soft errors", *IEEE International Test Conference*, 1984, pp. 649-659.

[SAV84b] Y. Savaria, V. K. Agarwal, N. C. Rumin, and J. F. Hayes, "Delay equalisation for soft error tolerance of VLSI logic circuits", *1984 Canadian Conference on Very Large Scale Intergration*, pp. 1.2 - 1.5.

[SER84] G. Sery, K. Kokkoken, P. Dishaw, B. Mantha, J. McCollum, J. Orton, J. Smudsky, and R. J. Smith, "CHMOSIII technology for VLSI applications", *Proc. of the International Conference on Computer Design*, 1984, pp.551-554.

[SIE82] D. P. Siewiorek, R. S. Swarz, *The theory and practice of reliable system design*, Digital Press 1982.

[SON84] W. S. Song, L. A. Glasser, "Power distribution techniques for VLSI circuits", *Conference on advanced research in VLSI, MIT, 1984*, pp.45-52.

[TOY79] T. Toyabe, and S. Asai, "Analytical model of threshold voltage and breakdown voltage of short-channel MOSFETs derived from two-dimensional analysis", *IEEE Journal of Solid-State Circuits*, vol. SC-14, pp. 375-383, April 1979.

[WHA79] J. J. Whalen, J. G. Tront, C. E. Larson, and J. M. Roe, "Computer-aided analysis of RFI effects in digital integrated circuits", *IEEE Trans. on Electromagnetic Compatibility*, vol. EMC-21, pp.291-297, Nov.1979.

[WIN63] S. Winograd, and J.D. Cowan, *Reliable computation in the presence of noise*, The MIT press 1963.

[WOL63] A. W. Wolfendale, *Cosmic rays*, George Newnes 1963.

[WOO81] J. Wood, "Reliability and degradation of silicon devices and integrated circuits", in *Reliability and degradation, semiconductor devices and circuits*, M. J. Howes, and D. V. Morgan ed., John Wiley 1981, ch. 4, pp. 191-236.