

Timing Vulnerability Factors of Sequentials

Norbert Seifert, *Senior Member, IEEE*, and Nelson Tam, *Member, IEEE*

Abstract—Single-event upsets (SEUs) from particle strikes have become a key challenge in microprocessor design. Modern superpipelined microprocessors typically contain many thousands of sequentials whose soft-error rate (SER) cannot be neglected anymore. An accurate assessment of the SER of sequentials is therefore crucial. This paper describes a method for computing timing vulnerability factors (TVFs) of sequentials. Our methodology captures the impact of the circuit environment which sequentials are typically placed in. Further, upsets occurring in local clock nodes have been accounted for. Results are presented for master–slave type flip flops and for flow-through latches of a high-performance microprocessor. Our investigations demonstrate that TVFs are a strong function of the propagation delay of the combinational logic and typically vary between $\sim 0\%$ and 50% . For high-performance microprocessors, we predict average TVF values of the order of 20% – 30% . Further, we expect TVFs to be largely technology independent for the same design.

Index Terms—Jitter, radiation effects, sequential logic circuits, SEE, SER, SEU, soft error.

I. INTRODUCTION

TECHNOLOGY scaling has driven the computer industry for several decades now. Scaling has not only resulted in cheaper and more powerful microprocessors, it has also resulted in microprocessors that contain many millions of devices with ever-decreasing node charges. This trend has the potential to increase the radiation-induced failure rate of future microprocessors substantially.

Single event upsets (SEUs) arise from the interaction of energetic particles—such as neutrons and alpha particles—with the semiconductor. This interaction generates electron–hole pairs, which are separated by the electric fields near reverse-biased junctions. If sufficient charge is collected at a device’s diffusion, the state of the logic device (SRAM, sequential, etc.) can be inverted, thereby introducing a logical fault. For a comprehensive introduction to the subject of soft errors, the reader is referred to [1]. In the remainder of this paper, soft errors (SE) denotes exclusively radiation-induced failures.

The radiation-induced soft error rate (SER) of modern microprocessors, whose caches and large memory arrays are protected by parity or ECC, is dominated by the failure rate of sequentials [2]–[4]. An accurate assessment of the SER of sequentials is therefore very important. An excellent introduction into how the SER at the chip level can be estimated has been given by

Nguyen and Yagil [5]. The SER of a circuit can be described by the following equation [3], [5], [6]:

$$\text{SER}^{\text{circuit}} = \sum_{i \text{ over all nodes}} \left(\text{SER}_i^{\text{nominal}} * \text{TVF}_i * \text{AVF}_i \right) \quad (1)$$

where the nominal SER ($\text{SER}^{\text{nominal}}$) refers to the underated SER which is independent of the circuit environment. The impact that the circuit environment and the architecture have on SER of the circuit are accounted for by the TVF and AVF factors, respectively. TVF denotes the timing vulnerability factor and is defined as the fraction of time a node or device is susceptible to upsets. AVF is the architectural vulnerability factor and equals the probability that a fault in device i will be observed by the system or user. Please note that AVF and TVF are called logic derating (LD) and timing derating (TD) factors in Nguyen’s and Yagil’s paper¹ [5].

A methodology to compute the architectural vulnerability factors has recently been laid out by Mukherjee *et al.* [6]. Since TVFs and AVFs are independent of each other, AVFs are set equal to one for the remainder of this paper. Please note that this study focuses on the calculation of timing vulnerability factors only.

This work introduces a methodology to compute the timing vulnerability factors of sequentials at the example of master–slave (MS)-type flip-flops (FFs) and flow-through latches used in a high-performance microprocessor. What is new in this study is that key aspects of the circuit environment, in which sequentials are typically placed, are taken into account. Our work demonstrates the impact that the propagation delay of combinational logic has on TVF and therefore on the SER of the sequential. Further, this work addresses and quantifies for the first time radiation-induced jitter and its impact on the failure rate of sequentials under use conditions.

Since most sequentials are placed in data or control paths, the impact that the logic depth of the path has on the timing derating of sequentials has to be accounted for. In the past, this effect has been neglected, which resulted in very conservative timing derating factors and correspondingly conservative SER estimates. Therefore, a direct and immediate benefit of using the proposed methodology is that overly conservative timing derating values, such as a TVF of 50% for master–slave type latches [5], are avoided. It is worth mentioning that the methodology described here is applicable to any type of sequential.

The rest of the paper is organized as follows. Section II discusses the details of the simulation methodology used to compute the timing vulnerability factors. Results are presented and discussed in Section III.

¹We chose to use the term vulnerability instead of derating, since a larger derating corresponds to a higher failure rate, which is counter-intuitive. A larger vulnerability, in contrast, is consistent with a higher SER

Manuscript received February 23, 2004; revised May 4, 2004.

N. Seifert is with Intel Corporation, Logic Technology Development Q&R, Hillsboro, OR 97124 USA (e-mail: Norbert.Seifert@ieee.com).

N. Tam is with Intel Corporation, Enterprise Processor Division, Santa Clara, CA 95052 USA (e-mail: nelson.tam@intel.com).

Digital Object Identifier 10.1109/TDMR.2004.831993

II. METHODOLOGY

In the following, the methodology to compute TVF values is described in detail. First, the difference between the nominal and derated SER is explained, then the simulation procedure used in this study to compute the derated and nominal SER values is described.

A. Nominal Versus Derated SER

As mentioned in the previous section, the timing vulnerability factor is a measure of the fraction of time a circuit is sensitive to upsets. The actual, derated SER of a circuit is—according to (1)—the product of the nominal SER times the vulnerability factors.

The nominal SER is defined as the soft failure rate of a circuit/node under *static* conditions, assuming all the inputs and outputs are driven by a constant voltage. This is typically how the SER of latches is measured experimentally, and the failure rate can be characterized by a constant Q_{crit} . However, when the circuit is placed in its natural environment, i.e., a datapath of a microprocessor, its inputs typically vary as a function of time, resulting in dynamic biasing conditions. As a result, the Q_{crit} of the circuit becomes a function of time. Since Q_{crit} varies with time, the SER rate of a circuit under actual operation conditions also varies in time. The nominal SER usually represents the worst case failure rate of a circuit, and the timing tends to reduce its SER. This is reflected by the fact that the vulnerability factors in (1), which take into account the timing and logical masking, are smaller or equal to one. Once the nominal and derated SER values of a circuit have been determined, the TVF of a circuit is computed by the following equation:

$$\text{TVF}^{\text{circuit}} = \frac{\text{SER}^{\text{circuit}}}{\text{SER}^{\text{nominal}}}. \quad (2)$$

Similarly, one wants to factor out the architectural vulnerability from the SER of a device or circuit, which is not a subject of this paper, however. The AVFs are assumed to equal one in this work. The impact of the circuit environment on the vulnerability factors has been modeled using circuits and miniature data paths with a fanout of one. This, and the fact that every latched fault is considered a soft error, justifies the usage of an AVF of one [5].

B. Simulation Setup

To compute the SER of a circuit (nominal or derated), one needs: 1) the waveform of the particle-induced injected current; 2) the critical charges (Q_{crit}) of each node in the circuit as a function of time when the charge is injected; and 3) calibrated models that estimate the SER as a function of Q_{crit} , the charge collection area, and the particle fluxes. Q_{crit} denotes the minimum collected charge that corrupts the state of a node. Q_{crit} s are computed as explained in [7] by inserting current sources into the netlist that account for the injected charge. In all the SPICE-level simulations, particle strikes were modeled with piecewise-linear current waveforms to account for funneling and diffusion charge collection. Alpha-particle hits and neutron hits have been modeled using two different waveforms, whose exact shapes have been extracted from a combination of experimental and device simulation data. The critical charge of

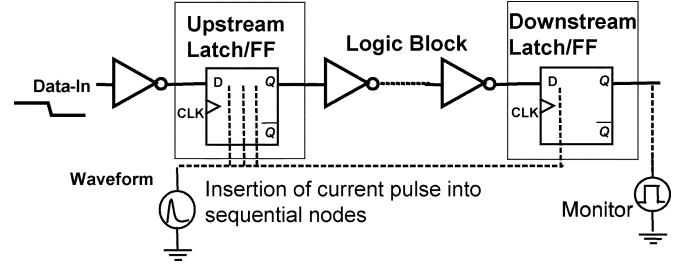


Fig. 1. Simulation setup for computation of the derated SER of sequentials. The miniature data path is fully clocked and Q_{crit} as well as the SER contributions of each node are time dependent. The number of inverters is changed to vary the propagation delay in the combinatorial logic.

a node is then found iteratively by increasing the magnitude of the current pulse until an upset condition is observed.

In this paper, the nominal SER is defined as the sum of the failure rates of the state nodes of the investigated sequential under static conditions. Further, it is assumed that all input vectors have the same probability of occurrence. We therefore get

$$\text{SER}_{\alpha,n}^{\text{nominal}} = \sum_{\text{State Node } i} \text{SER}_{\alpha,n}^{\text{nominal}}(Q_{\text{crit},\alpha,n}^i) \quad (3)$$

where α, n refer to the alpha-particle and neutron-induced SER and Q_{crit} values and i enumerates all state nodes of the studied sequential. Please note that the weighting of the input vectors is not reflected in (3) for reasons of simplicity. Other definitions of the nominal SER are possible in principle and would yield corresponding TVF values. The general trend and key observations do not depend on the exact nature of the definition as long as the nominal SER does not account for the circuit environment in which the sequentials are placed. The above definition is consistent with experiments and other simulation efforts conducted at Intel Corporation. As mentioned previously, all Q_{crit} and SER values are independent of time for the nominal case.

The computation of the *derated* SER is more involved and is based on the following equation, which is derived in the Appendix:

$$\langle \text{SER} \rangle_{\alpha,n} = \sum_{\text{Node } i} \sum_{\text{Time } j} \text{SER}_{\alpha,n}(Q_{\text{crit},\alpha,n}^{i;j}) \frac{\Delta t}{T_{\text{cycle}}} \quad (4)$$

where T_{cycle} denotes the clock cycle and Δt the time-step (see below and Fig. 2 for more details).

Fig. 1 depicts the simulation setup used to implement (4). This methodology naturally accounts for the timing vulnerability factors, since it computes the SER of a circuit taking into account the functionality and environment of the circuit as a function of time when the charge is injected. Since the computational methodology is SPICE based, glitch propagation, electrical and logical masking, and latching-window masking are automatically included in this approach [2]. If a node, for instance, cannot be upset during a certain time interval, then Q_{crit} will be large and the computed SER contribution very small. One can see in Fig. 2, depicting the critical charge values of the state node of a slave latch as a function of time for two different path lengths, that Q_{crit} is now time dependent. The implications of Fig. 2 will be discussed in the next section.

To compute the average derated SER of a circuit, a clock cycle is divided up into N time-steps of length Δt (Fig. 3). At each

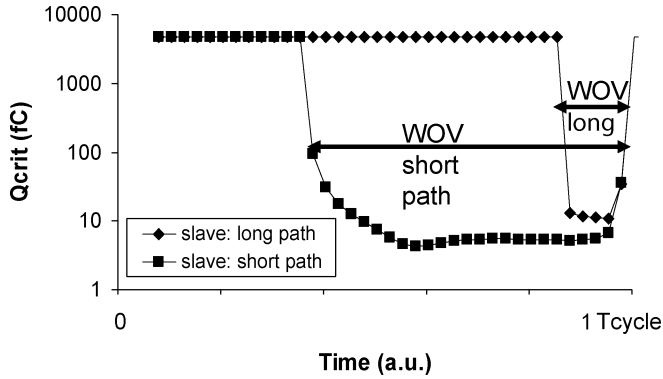


Fig. 2. Computed critical charge values (Q_{crit}) as a function of time when charge is being injected into the slave node. One can clearly see that the window of vulnerability (WOV), i.e., the time when the node is susceptible to upsets and the Q_{crit} is low, decreases with increasing logic path lengths.




Fig. 3. To compute the derated SER, charge is injected into sequential nodes during one clock cycle, which is divided up into N time-steps. The downstream sequential output node is checked for latched faults when the sequential is in hold mode.

time-step Q_{crit} is determined and the nominal SER is computed. N typically equals 40–80 time-steps. The monitor node for the derated SER computations is the output node of the last (downstream) latch or FF. It is important to emphasize that this approach implicitly assumes that every clock cycle is equivalent in terms of SER. If that is not the case, then the simulations have to run and average the SER over the corresponding number of clock cycles.

To compute the derated SER the sequentials are placed along a data or control path. As mentioned above, the monitor node is at the end of the miniature data path, as shown in Fig. 1. We are only interested in faults that manage to propagate to the next downstream sequential. If the fault generated in the upstream sequential is masked by the timing constraints of the path, then it is a “don’t care” and should not contribute to the actual SER. This way, the impact of the propagation delay on the SER is modeled properly. To vary the propagation length, the number of inverters between the sequentials is varied. Inverters instead of NANDs, NORs, etc., have been chosen to be consistent with an AVF equal to one, i.e., to make sure no logical masking is occurring.

This simulation setup implicitly assumes periodic boundary conditions, i.e., that all sequentials are equivalent in terms of TVFs. This is strictly valid only for very long paths where the SER contribution of the last downstream sequential is negligible.

In this study, solely nonclock nodes of the *upstream* sequential are hit during one clock cycle. In contrast, to estimate the impact of upsetting local clock nodes on the SER, charge is injected into the *downstream* clock nodes. Please note that *local* here refers to clock nodes that belong to the schematic of the studied sequential. The main reason for injecting charge into the downstream clock nodes is that any jitter introduced into a clock node has the potential to cause a setup violation, i.e., that the downstream sequential cannot latch the correct data, which is also a SER contribution. Hitting clock nodes in the *upstream* sequential does not cause any setup violations, since no data is propagating toward that sequential in the setup used here, since the miniature data path used here comprises only one stage (without any loss of generality).

III. RESULTS AND DISCUSSION

The dependence of TVF on the propagation delay between two sequentials is of major interest, since most sequentials are located along data and control paths in the case of modern high-speed microprocessors. The nominal SER of a sequential, which is independent of the circuit environment and therefore independent of the propagation delay of a particular data path, is usually readily available. To determine the total chip-level SER, however, one needs to factor in the impact of different logical depths of different pipeline stages and paths. This dependence is implicitly contained in TVF. The assessment of the total SER contribution of sequentials on a chip therefore involves two steps (neglecting AVFs):

- 1) modeling of the dependence of TVF on the propagation delay in the combinational logic at use conditions (i.e., for given V_{cc} , temperature, clock speed, etc.);
- 2) extracting the chip-level distribution of propagation delays.

The chip-level SER contribution of sequentials is then computed by integrating the TVF dependence on the propagation delay over the minimum path delay distribution. Minimum path statistics have been chosen, in order to stay on the conservative side. In this study, a proprietary static timing tool has been used to collect the delay distribution of a couple of sections on a modern high-performance microprocessor.

A. TVF Results

Fig. 4 depicts TVF as a function of the ratio of propagation delay to cycle time for a master–slave type FF and for a flow-through latch at the target clock speed of the studied microprocessor. One can see that TVF decreases as a function of increasing propagation delays, mainly because the fault has to arrive at the *downstream* sequential at least a setup time before the clock asserts (assuming that skew and jitter are negligible). If the upset in the *upstream* sequential occurs late in the cycle, it might not make it to the next sequential in time, and therefore will be masked (see Fig. 5 and Fig. 2). In Fig. 2 the impact of the path length on Q_{crit} of a slave state node is demonstrated. Consistent with the above reasoning, the window of vulnerability (WOV) strongly depends on the number of combinational gates placed between two sequentials. A larger WOVS corresponds to a larger derated SER and consequently a larger TVF as defined in




Fig. 4. Calculated TVF for master-slave FFs and flow-through latches as a function of the relative propagation delay. TVF decreases with increasing propagation delay. The offset at $t_{prop} = 0$ is due to the intrinsic delay in the sequentials.




Fig. 5. Any upset occurring outside the sensitive time window WOVS will not propagate in time to the next downstream sequential and therefore will not contribute to the SER. The propagation delay impacts the WOVS of the master only for $T_{prop} \geq T_{phase}$.

(2). In Fig. 2 the width of the WOVS appears wider than 50% of the cycle time for the shorter path, suggesting a TVF larger than 50%. Please note that this is in Q_{crit} space only. Since SER decreases dramatically with Q_{crit} , the corresponding TVF is still below 50%.

For flow-through latches as well as MS-type FFs, TVF can be modeled as

$$TVF_{sequential} \propto \frac{(T_{cycle} - (\Delta t_{tot}^{prop} + \Delta t_{setup} + \Delta t_{clk} \pm \Delta t_{skew}))}{T_{cycle}} \quad (5)$$

where Δt_{tot}^{prop} denotes the sum of the propagation delay through the combinational logic and the intrinsic delay within the sequential, Δt_{setup} the setup time, Δt_{clk} the clock rise and fall times and also accounts for clock jitter, and finally Δt_{skew} the clock skew.

The simulation results depicted in Fig. 4 do not exactly extend to T_{cycle} for FFs and T_{phase} for latches mainly because of the finite delay of the combinational gates and because of the fact that the fault has to arrive about Δt_{setup} before the clock asserts.

At very slow clock speeds (i.e., $\Delta t_{tot}^{prop}/T_{cycle} \ll 1$) the TVF of a master-slave FF is expected to equal 50% for a 50% clock duty cycle. In this case the master as well as the slave is each driven for 50% of the time and therefore practically not vulnerable to upsets during this time. This “zero-delay” initial value of TVF is not reflected in (5). Equation (5) solely shows the trend of TVF as a function of T_{cycle} and various delay and clock parameters. The initial TVF value might be different for different types of sequentials, however, the dependency on T_{cycle} etc should be valid for a wide range of sequentials.

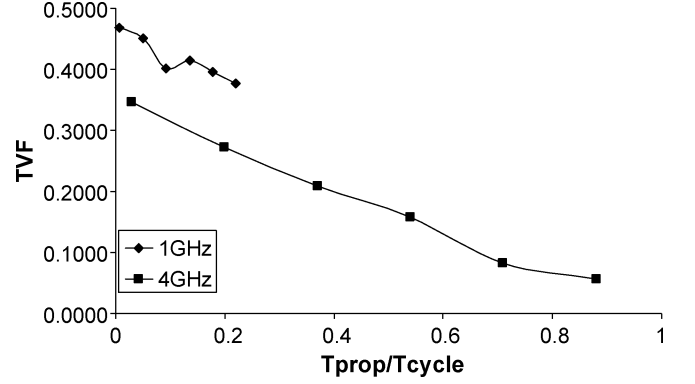


Fig. 6. Computed TVF are depicted for two different cycle times. TVF clearly increases for decreasing clock frequency.

For latches, T_{cycle} needs to be replaced with T_{phase} . It is well known that setup and skew do not impact the cycle time of latches in contrast to FFs, because of the transparency of latches [8]. For latches, data arrive usually more than a setup time before the closing edge of the clock. However, here we are focusing on *faults*, as opposed to data, being latched and probe the monitor node during the hold phase (i.e., when the latch is closed). This is why skew and setup time do impact the TVF of a latch as described by (5). A fault generated in the upstream latch still has to arrive at least a setup time before the clock edge at the downstream latch in order to be latched and become a SE. This, however, is not the case when clock nodes are being upset, as we will discuss further below.

Based on this, one therefore expects that the SER contribution of sequentials decreases with increasing clock frequency. In Fig. 6, the TVF for different clock frequencies is plotted for FFs. The TVFs have been calculated at the same power supply voltages (V_{cc}). Clearly, the TVF increases with decreasing clock speed. This makes sense, since for larger cycle times the relative contribution of the propagation delay becomes smaller and the window during which an upset in the upstream sequential can be latched by the next sequential increases. This has indeed been observed experimentally [9]. Please note that for clock speeds approaching zero, the maximum TVF equals 50% for the shortest paths. The slope of the TVF versus (T_{prop}/T_{cycle}) curve is independent of the clock speed, as one can see in (5) and Fig. 6, if radiation-induced jitter is neglected.

For an MS-type FF, TVF is actually the average of the contributions of both the master and slave part. Fig. 7 depicts both individual contributions. The TVF of the slave latch decreases faster than that of the FF (compare with Fig. 4). The main reason for this is the almost constant TVF of the master latch at small propagation delays. It takes propagation delays longer than about T_{phase} to impact the time during which the master is susceptible to upsets (i.e., when the master is in hold). This is different for the slave latch, which is in hold and therefore susceptible to upsets during the clock phase when the data and faults have to arrive at the downstream latch. In that case, any propagation delay decreases the vulnerability time window of the slave latch.

Interestingly, the master TVF *increases* for very long propagation times (see Fig. 7). This increase is due to faults caused

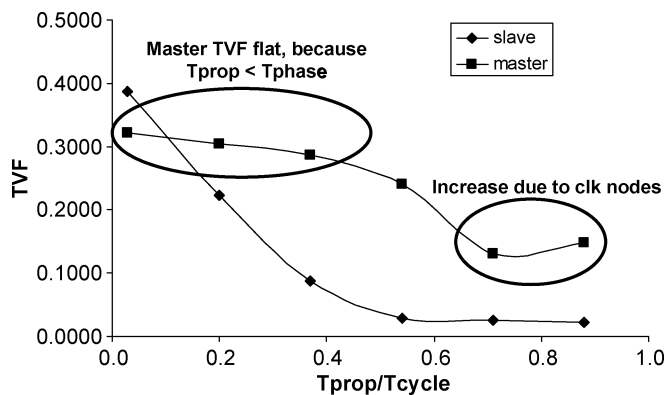


Fig. 7. Calculated TVF of a master–slave FF. Both contributions, the master TVF and the TVF of the slave, are depicted as a function of the relative propagation delay.

by hits in the local clock nodes. For very long propagation times (i.e., critical paths), the data barely make it to the downstream FF. Any radiation-induced jitter in the clock tree can cause the data to be incorrectly latched. Therefore, clock node hits contribute only for critical paths in the case of FF's. In extreme cases, jitter-induced TVF could dominate the SER of the sequential and could reach values beyond 100%, since upsets occurring in local clock nodes do not contribute to the *nominal* SER. One example would be radiation-hardened sequentials, where the clock nodes have not been hardened.

For paths that have plenty of delay margin, jitter does not result in any faults. Note that this radiation-induced jitter is about *data* not being latched, whereas when sequential nonclock nodes are hit, TVF is determined by irradiation-induced *faults* being latched. Faults *not* being latched because of jitter introduced into the clock tree do not contribute significantly to reducing the SER, since that case would require two simultaneous hits in the same data-path segment, which is very unlikely.

In the case of latches, TVF behavior for clock node hits is not as simple and straightforward as in the case of FFs. The main reason for this is the transparency of the latches, as discussed previously. Even in the case of long data paths, data could arrive at the latch sufficiently far away from the clock edge, because of the time borrowing potential of latches [8]. So, even if the path delay is the maximum possible (i.e., $\sim T_{\text{phase}}$), data could be traversing the latch in the middle of the transparent period. Only in those unfortunate cases where the data happen to arrive at the clock edge, jitter introduced by radiation will increase TVF. This is the case when the circuit designer exploits the time borrowing potential of latches and reduces the path length of the previous path, which allows for paths with propagation times larger than T_{phase} in the subsequent segment. But this is rather unlikely and the increase of TVF for long propagation times is therefore neglected here, as can be seen in Fig. 4.

Another issue with latches in contrast to FFs is that a fault induced in the upstream sequential could propagate through the downstream latch while it is open and be latched in a different cycle further downstream. This has not been considered in our simulations, where we accounted for two sequentials only. Please note that the probability of the fault/glitch being latched at another sequential downstream is relatively small, since it

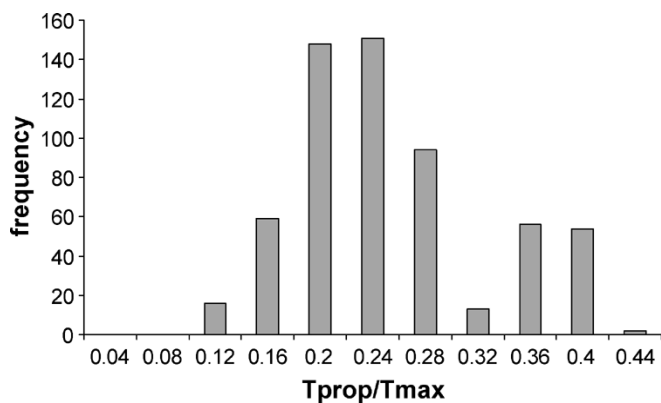


Fig. 8. Path delay histogram for a few sections of a high-performance microprocessor. The average delay in this case is about 25% of the target cycle time.

has to arrive at the receiver within the setup and hold time window [5].

Please note that for the shortest path lengths, upsets occurring in the clock tree could result in race, i.e., data propagating through two pipeline stages in one clock cycle or clock phase. However, in our simulations we have assumed zero skew and the potential increase of TVF is therefore negligible for the cases studied. Further, since race cannot be fixed by lowering the cycle time [8], designers usually apply sufficient margin to their design, which makes radiation-induced hold-time violations rather unlikely, even in the presence of skew.

Although the present study focuses on sequentials, it can easily be expanded to dynamic logic. Depending on the propagation delay to the next synchronization point, which could either be another dynamic gate or a sequential, TVF of dynamic logic is expected to decrease with increasing propagation delay as well.

We have collected delay distribution statistics for a few sections of one high-performance microprocessor (Fig. 8). Please note that the results depicted in Fig. 8 are *not* representative of the delay distribution of the microprocessor studied. The average FF–FF path delay of the chip sections depicted in Fig. 8 is about 25% of the cycle time. This translates into a TVF of $\sim 25\%$ for FFs. Chip-wide statistics suggest average TVF values for FF and latches of the order of 20%–30% for high-performance microprocessors.

We have addressed the impact of the length of data and control paths on the TVF and therefore on the SER of latches and MS-type FFs. Our results predict that the SER of the studied sequentials decreases as a function of clock frequency because of this path length dependence. We would like to emphasize that this dependence on the clock speed has been observed experimentally in [9] for the 21164 Alpha microprocessor.

The question remains how applicable the reported results are for other types of sequentials. Reference [9] also reports the observation of an *increasing* SER for 21264 Alpha microprocessors, which seems to contradict the results presented here. In the case of the 21264 sense amplifier (SA)-type FFs have been used. SA-type FFs are only sensitive during a very short and most importantly constant time window. Therefore, TVFs for SA-type FFs are proportional to a constant divided by the cycle

time and increase with increasing clock frequency. This observation underlines the fact that the functional dependence of TVF of sequentials depends strongly on the type of sequential being used. However, the methodology described here is valid and applicable to any type of sequential. Further, the TVF trends presented here should be valid for all types of flow-through latches and MS FFs.

B. Scaling Trend

Within the same technology and design, the TVFs of the investigated sequentials decrease for increasing clock speed, since the propagation time, setup, and clock skew are independent of the cycle time. Of greater interest, however, is the extrapolation of TVF to different, scaled processes. For the same type of sequential, one expects that all quantities in brackets in (5) should to first order scale roughly the same as the cycle time. This translates into roughly technology-independent TVF for the same type of sequentials. TVF for the same product built in two different technologies therefore will be similar, although not perfectly aligned, since designers usually only fix the critical paths to make the circuits work, but the average of the delay distribution might shift depending on the scaling of the interconnect RC and transistor speeds. No accurate prediction can be made for new designs built in a different process.

IV. CONCLUSION

Timing vulnerability factors (TVFs) of sequentials have been investigated in this work. The main importance of TVFs lies in the fact that they account for the circuit environment in which the sequentials are typically placed. The nominal SER does not capture this dependence and therefore says little about the actual SER of sequentials on real chips. This work underlines the impact that the propagation delay of combinational logic has on the actual SER of the sequentials.

Master-slave-type flip-flops and transmission-type flow-through latches have been studied. The SPICE-level simulation results suggest that TVFs vary between 50% (valid at very slow clock speeds relative to the combinational and intrinsic sequential propagation delays) down to almost 0%. Particularly for critical paths, TVF is close to the minimum. However, TVFs do not become exactly zero because of the increasing importance of clock node upsets for the longest paths. Clock jitter introduced by ionizing radiation increases the SER of sequentials in the case of critical paths and in extreme cases can result in a TVF of larger than 100%.

The key observations of this work are that the timing vulnerability factor and consequently the SER decreases for the studied type of sequentials with increasing clock frequency, increasing propagation delay, and for decreasing supply voltage (data not shown here). The main reason for this observed trend is the relative increase of the propagation delay with respect to the cycle time. For high-performance microprocessors, TVFs are expected to be of the order of 20%–30% for the sequentials studied.

Finally, TVF for the same type of sequentials are believed to be largely technology independent.

APPENDIX

The particle-induced soft error rate of a device is a statistical quantity and its average $\langle \text{SER} \rangle$ can be defined as

$$\langle \text{SER} \rangle = \frac{1}{T_{\text{cycle}}} \sum_n^{\text{nodes}} \sum_i^Q \sum_{j=t_{\text{inj}}}^{T_{\text{cycle}}} \text{upset}(Q_i, t_{\text{inj}}, n) \times \text{PROB}(Q_i, t_{\text{inj}}, n) \quad (6)$$

where upset as a function of the collected charge Q_i injected at time t_{inj} into node n denotes whether an upset has occurred and equals either one (upset detected) or zero (no upset detected). PROB denotes the probability that charge Q_i is collected, and the charge was injected at time t_{inj} into node n . All the charge collection properties, as well as the particle flux, are contained in $\text{prob}(Q_i)$. Assuming that the probabilities are uncorrelated, i.e.,

$$\begin{aligned} \text{PROB}(Q_i, t_{\text{inj}}, n) &= \text{prob}(Q_i) * \text{prob}(t_{\text{inj}}) * \text{prob}(\text{node}_n) \\ &= (\text{prob}(Q_i) \Delta q) \left(\frac{\Delta t}{T_{\text{cycle}}} \right) \left(\frac{A_n}{A_{\text{tot}}} \right) \end{aligned} \quad (7)$$

we get

$$\langle \text{SER}_{\text{circuit}}^{\text{total}} \rangle = \sum_{\text{Node } n} \sum_{\text{Time } j} \text{SER}(n; j) \frac{\Delta t}{T_{\text{cycle}}} \quad (8)$$

where $\text{SER}(n, j)$ is the nominal SER of node n at time-step j defined by [3], [10]

$$\text{SER}(n, j) = A_n \sum_i^Q \text{prob}(Q_{i,n}) \Delta q \frac{\text{upset}_{j,i,n}}{T_{\text{cycle}}} \quad (9)$$

ACKNOWLEDGMENT

The authors would like to thank B. Griesbach for running most of the TIDEST simulations, D. Grundmann for writing some of the code used, S. Walstra for all his support with respect to the SER models and algorithms, M. Pant for help and advice regarding sequentials, and S. Mukherjee and J. Maiz for many insightful discussions.

REFERENCES

- [1] J. F. Ziegler *et al.*, "IBM experiments in soft fails in computer electronics (1978–1994)," *IBM J. Res. Devel.*, vol. 40, no. 1, pp. 3–18, Jan. 1996.
- [2] P. Shivakumar, M. Kistler, S. W. Keckler, D. Burger, and L. Alvisi, "Modeling the effect of technology trends on the soft error rate of combinatorial logic," in *Proc. Int. Conf. Dependable Systems and Networks*, 2002, pp. 389–398.
- [3] N. Seifert, X. Zhu, and L. W. Massengill, "Impact of scaling on soft-error rates in commercial microprocessors," *IEEE Trans. Nuclear Sci.*, vol. 49, pp. 3100–3106, Dec. 2002.
- [4] R. Baumann, "The impact of technology scaling on soft error rate performance and limits to the efficacy of error correction," in *Int. Electron Devices Meeting (IEDM) Dig.*, Dec. 2002, pp. 329–332.
- [5] H. T. Nguyen and Y. Yagil, "A systematic approach to SER estimation and solutions," in *Proc. IEEE Int. Reliability Physics Symp.*, Mar.–Apr. 2003, pp. 60–70.
- [6] S. S. Mukherjee, C. T. Weaver, J. Emer, S. K. Reinhardt, and T. Austin, "A systematic methodology to compute the architectural vulnerability factors for a high-performance microprocessor," in *Proc. 36th Annu. IEEE/ACM Int. Symp. Microarchitecture (MICRO-36)*, Dec. 2003, pp. 29–40.

- [7] R. J. McPartland, "Circuit simulations of alpha-particle-induced soft errors in MOS dynamic RAMs," *IEEE J. Solid-State Circuits*, vol. SC-16, pp. 31–34, Feb. 1981.
- [8] D. Harris, *Skew-Tolerant Circuit Design*, 1st ed. San Mateo, CA: Morgan Kaufmann, 2000, p. 42.
- [9] N. Seifert, X. Zhu, D. Moyer, R. Mueller, R. Hokinson, N. Leland, M. Shade, and L. Massengill, "Frequency dependence of soft error rates for sub-micron CMOS technologies," in *Int. Electron Devices Meeting (IEDM) Dig.*, Dec. 2001, pp. 14.4.1–14.4.4.
- [10] T. V. Rajeevakumar, N. C. C. Lu, W. H. Henkels, W. Hwang, and R. Franch, "A new failure mode of radiation-induced soft errors in dynamic memories," *IEEE Electron Device Lett.*, vol. 9, pp. 644–646, Dec. 1988.



Norbert Seifert (M'99–SM'04) received the M.S. degree in physics from Vanderbilt University, Nashville, TN, and the Ph.D. degree in physics from the Technical University of Vienna, Austria, in 1993. His Ph.D. thesis focused on radiation-induced defect formation and diffusion in wide bandgap ionic crystals.

He has conducted research in a wide range of physics topics, from charge-transfer processes in atomic collisions as a postdoctoral associate at North Carolina State University from 1993 to 1994, to computational fluid dynamics of high-power laser material processing as a postdoctoral associate at the Technical University of Vienna from 1994 to 1997. In 1997, he joined the Alpha Development Group (DEC/Compaq/HP) where he worked in the fields of device physics, device reliability, and digital design. He is currently a Design and Reliability Engineer with Intel Corporation, Hillsboro, OR. His research interests include the interdependence of design and reliability.



Nelson Tam (M'03) received the B.S. degree in chemical engineering and the M.S. and Ph.D. degrees in electrical engineering and computer science from the University of California, Berkeley, in 1984, 1989, and 1991, respectively.

At UC Berkeley, his research was on the characterizing and modeling of optical resists under electron-beam lithography. In 1991, he joined Intel Corporation, Santa Clara, CA, where he worked on the development of phase shifting mask (PSM). In 1997, he joined the Enterprise Processor Division as a Quality and Reliability Engineer focusing on IPF processor pre-Si reliability verification. His research interests include simulation and experimental techniques for determining radiation effects on microprocessors.