

Analysis of glitch power dissipation in CMOS ICs

M. Favalli and L. Benini*

DEIS - University of Bologna - Viale Risorgimento, 2, 40136 Bologna, Italy

* CIS - Stanford University - Stanford CA-94305 USA

Abstract

Since a consistent fraction of the total power dissipated in CMOS IC's is due to glitches, power estimation tools should correctly account for their presence. This can be done at the electrical level but only for medium size circuits. The accuracy achievable at the logic level, instead, has not yet been analyzed in details. In this paper, the phenomenon of glitch power dissipation is analyzed at the electrical level showing the main sources of error in gate-level simulation (with different delay models) and providing basic guidelines for the development of adequate logic level models.

1 Introduction

With the spreading use of portable equipments, power consumption is joining area and speed as the more relevant parameters in VLSI design [1]. As a consequence, the problem of estimating this quantity, with particular regard to CMOS circuits, has been recently addressed in several works.

Tools performing this task are widely based on existing simulation approaches and can be distinguished between those: a) performing a gate level analysis and evaluating power consumption as a sum of contributions from switching gates [2, 3, 4, 5]; b) based on electrical level simulation and reconstructing time domain power supply current waveforms [6, 7, 8, 9].

Such approaches differently trade off accuracy for efficiency. In particular, fast logic-level simulators suitable for the interactive design of large circuits are attracting increasing interest as power estimators, but their accuracy is far to be validated in detail and, as it will be seen, major problems are still unsolved.

A distinctive feature of static CMOS circuits is that the total power dissipation is mainly caused by signal switching. Hence, gate-level logic simulation algorithms estimate the average power dissipated by monitoring the activity (i.e. the number of transitions) of gate outputs and by using:

$$P_{avg} = f \frac{V_{DD}^2}{2} \sum_i^n C_{L_i} a_i \quad (1)$$

where: f is the inverse of the clock period, n is the number of

gates, while C_{L_i} and a_i are the output capacitance and the activity (a_i is equal to the number of gate output transitions in the considered period) of the gate i , respectively.

Expression 1 does not take into account the power dissipated because of gate internal capacitances and short circuit currents. These quantities, however, can be taken into account with library characterization procedures that collect data on power dissipation in look-up tables or similar data structures to be used at simulation time (see for example [5]). As a result, an accurate logic-level estimate of power dissipation is possible in hazard free circuits, but the presence of hazards gives rise to major accuracy problems because the power dissipated in such transitions may be lower than that needed for a complete transition.

This paper shows that results achieved at the logic level critically depends on the adopted simulation model, in fact, the transport delay model may give rise to an unrealistic high number of signal transitions (thus overestimating gate activity); while the results achieved by the inertial delay model are sensitive to the distribution of signal timing.

These problems cannot be overlooked because, as shown by analyses performed at the electrical level [10], the power dissipated because of hazards is a consistent fraction (15%-20%, or more) of the global power and it is strongly dependent on both the circuit topology and the applied test vector.

We have performed a study at the electrical level of glitch power dissipation in static CMOS gates that is specifically oriented to characterize this phenomenon from the point of view of logic level simulation.

We first consider the case of hazard generation (at a gate output because of hazard free transitions at gate inputs) analyzing the power estimate error of gate level simulation and identifying all parameters that are relevant to power estimate. The problem of hazard propagation through a gate is then considered.

The analysis of data collected from electrical-level simulation suggests that the estimation of power dissipated in glitch generation can be performed locally at the gate level. Unfortunately, this is not true for glitch propagation. In particular, the power estimate error of logic level simulators has been found to depend on the length of the paths through which the hazard propagates.

2 Glitch power dissipation

The basic contributes of hazards to the power dissipated in a digital circuit (see Fig. 1) are given by:

1. the generation of a static hazard at the output of a gate;
2. the propagation of a hazard through a gate.




Figure 1: Schematic of the basic ways in which a hazard may contribute to power dissipation.

The propagation of multiple static hazards as well as the generation of dynamic hazards are here considered as a second order effect of the above mentioned possibilities.

As in the case of glitch free transitions, in cases 1 and 2, the energy dissipated (here, we will refer in the same way to energy and power, since the performed analyses are independent from the circuit operating frequency) can be expressed such as the sum of three components, namely:

- a) the energy (E_{out}) used so as to perform the (possibly partial) output switch (E_{out} is proportional to $C_L V_{DD} \Delta V$, where C_L is the load capacitance and ΔV is the gate output swing);
- b) that (E_{par}) used to charge the parasitic capacitances;
- c) that (E_{cc}) dissipated because of penetration currents.

At the gate level no information on the shape of the waveform or the electrical characteristics of the devices and parasitics is specified. In the large majority of cases, the input waveforms are described as sequence of transitions, and the environment of a gate is specified using macroscopic “black box” parameters such as input and output loading and gate propagation delays. We will therefore employ these quantities as parameters for our electrical-level simulations.

As an example of the electrical level behavior in the presence of glitches, the case of a CMOS static NAND gate will be discussed throughout the paper.

3 Hazard generation

In this case, the input signals of a gate switches between configurations producing the same gate output value, but, depending on signal switching instants, a temporary gate input configuration may be present that produces a complementary output value thus resulting in a static hazard.

3.1 Simulation set-up

The simulation set-up used to characterize glitch power dissipation in case (1) is constituted the by circuit of Fig. 2 that is composed of NAND gates designed by means of a standard $1.2\mu m$, 5V power supply, CMOS process. Gates have been taken from a MOSIS library of digital cells. The gate where the hazard is generated is G and it is driven by two other gates (G_A and G_B) so as to provide realistic input waveforms. Both the analyzed gate and the driving gates have a variable fan-out so as to account for changes in input signals slope and variations in the output load capacitance (in particular, the fan-out of G , G_A and G_B are denoted by N , N_A and N_B , respectively).

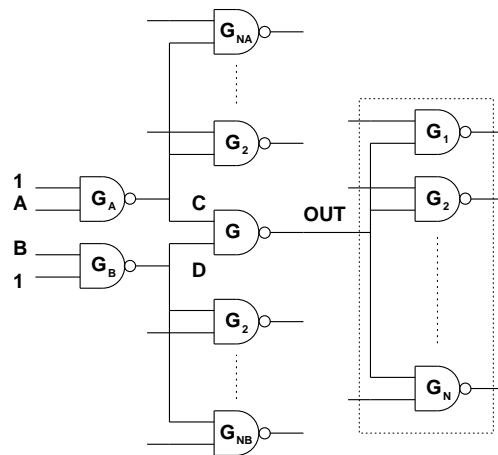


Figure 2: Schematic of the circuit used so as to characterize hazard generation from the point of view of power dissipation. G is the gate where the hazard is generated because of skewed transitions at the outputs of G_A and G_B .

The electric schematic of gate G is shown in Fig. 3, this scheme makes use of lumped capacitors to be referred in the following discussions. The SPICE simulations, however, have been performed by using level 3 transistors models and layout extracted parasitic capacitance parameters.




Figure 3: Electrical schematic of the considered NAND gate.

As for the input waveforms of G (signals C and D), the two cases (a) $CD = 01 \rightarrow 10$, b) $CD = 10 \rightarrow 01$ giving rise to hazard generation are schematically illustrated in Fig. 4 where symmetric waveforms have been considered.

The characterization of input signal waveforms is here simply made by means of the skew (τ) between signals C and D that is measured at half of the logic swing (see Fig. 4) and of the slope of the signal waveforms. To simplify, the slopes of signals C and D are assumed in the discussion to be inversely proportional to the load of the gates G_A and G_B . These choices are consistent with the kind of informations available in event driven logic simulation.

As for skew measurement, it should be noticed that it is measured as the difference between the instants in which the rising and the falling signal reach half the logic swing. Then, if the falling signal anticipates the skew is negative, but the glitch power is non null because the pull-down transistors are briefly turned ON.




Figure 4: Possible input waveforms of the gate G of Fig. 4 giving rise to an output static hazard.

Depending on the skew, the gate output voltage may perform a full swing or not; the main parameter at this regard is the minimum voltage (V_{min}) that provides an approximation of the energy drawn during the hazard as:

$$E_{out} = \frac{1}{2}C_L V_{DD}(V_{DD} - V_{min}), \quad (2)$$

that, as already noticed, does not takes into account neither parasitic capacitances nor penetration current.

All glitch power/energy results presented in this work will be normalized to the power/energy dissipated in case of a rising glitch-free output transition ($CD = 11 \rightarrow 10$) in the case of unitary fan-out for G_A , G_B and G .

We will analyze the dependence of glitch power dissipation from the macroscopic parameters that characterize the environment of the gate. Using the input skew τ as independent variable to modulate the amplitude of the glitch, we will study the variation of glitch power for varying fan-out, applied input pattern and fan-in loading.

3.2 Output load dependency

The circuit of Fig. 2 has been simulated with different values of the fan-out of the gate G and with the driving gates with unitary fan-out. The normalized power dissipated at gate G has been evaluated for different values of the skew τ between its inputs and it is shown in Fig. 5.

In this set of simulations, the slope of the input signals are of the same order or larger than the output slope. Glitch power dissipation is therefore dominated by E_{out} which is proportional to $V_{DD} - V_{min}$. Notice that when $V_{out} = V_{min}$ the current provided to the load is null ($dV_{out}/dt = 0$) and the circuit working point lies on the static voltage characteristic of the considered gate.

In Fig. 5, three regions can be recognized:

1. for low values of τ (approximately up to 0.2ns) the glitch power does not depend on the load capacitance and exhibits a parabolic course;
2. for intermediate values of τ (between 0.2ns and the propagation delay) glitch power starts to become slightly sensitive to the load capacitance and it depends almost linearly on the skew;
3. for values of the skew larger than the propagation delay, the glitch power saturates to the value corresponding to a hazard free low to high transition.




Figure 5: Normalized power dissipated for hazard generation at the output of gate G in the circuit of Fig. 4. The fan-out of G is varied, while the fan-out of driving gates is kept at 1. Data are plotted for different values of the skew. Vertical lines are placed in correspondence of the gate propagation delay for each loading condition.

A first-order approximation of this behavior could employ a piecewise linear fitting, with coefficient tabulated for various load conditions in a look-up table. Notice however that, glitch power dissipation is non negligible even for $\tau < 0$, this means that even if the two events at the inputs do not propagate any event at the output (either in transport and inertial simulation), the gate dissipates a non negligible power. This suggests that, at the logic level, power estimate by gate output waveforms analysis presents intrinsic accuracy problems that can be overcome by analyzing gate input waveforms.

We will refer to our circuit example to briefly explain the observed behavior. It is immediate to recognize the correspondence between the above mentioned regions and those of the NAND gate static voltage characteristic. In region 1), V_{min} is slightly below V_{DD} so that the transistor (M_{NC}) that is turned ON during the hazard remains in the saturation region when conducting, the pull-up transistors are in linear region. In region 2, both M_{NC} and the pull-up transistors remain in the saturation region for the more relevant part of the falling transient of V_{out} . In both cases, the current flowing in the NAND pull-down can be roughly assumed to depend only on input waveforms and not on the load so that ΔV results inversely proportional to the load capacitance and E_{out} is independent of the load capacitance.

Finally, in region 3 the transient is exhausted with M_{NC} in linear region and consequently ΔV is no longer proportional to $1/C_L$ and it exponentially approximates V_{DD} .

3.3 Pattern dependency

As for pattern sensitivity, the curves in case (b) of Fig. 4 have shapes similar to those of Fig. 5 although presenting lower values of glitch power. This can be seen in Fig. 6 that shows a comparison between the results achieved in case (a) and (b) for the specific case of $N = 2$.

As can be seen, for low input skews the glitch power in case (b) is much lower than in (a). This because at the beginning of the glitch $V_N = V_{DD} - V_{TN}$ (instead of $V_N = 0$) and the transistor turned ON (M_{ND}) has to discharge C_N to allow for M_{NC} to conduct a non negligible current. In this way the discharging of the gate output capacitance is less efficient.

For higher values of the skew, the ratio between the power dissipated in the cases (a) and (b) becomes proportional to the ratio between the capacitances to be loaded during the rising transient of V_O (that is $C_L + C_N$ in case (a) and C_L in case (b)) such as in the case of glitch free transitions ($CD = 11 \rightarrow 10$ and $CD = 11 \rightarrow 01$).




Figure 6: Glitch power dissipated for two different sequences of input stimuli ((a) and (b)) in a CMOS NAND gate.

This behavior has been found in different CMOS static gates and, in the considered NAND case, it is instantiated in Fig. 7, where the ratio between the glitch power dissipated in cases (a) and (b) is plotted as a function of τ and for different values of load. A first-approximation model of pattern-dependent power dissipation should then keep track of the charge status of the internal capacitances.




Figure 7: Ratio between the glitch power dissipated in cases (a) and (b).

3.4 Input slope dependency

The influence of the driving conditions on glitch power can be expressed in term of the input slopes.

It should be noticed that, with the exception of low values of the skew (to be discussed later), faster input signals correspond to faster output transients and to decreases in the short circuit current. These conflicting effects lead, for the same gate output load and skew, to opposite variations of E_{out} and E_{cc} when the input slopes are varied.

The shape of the glitch power curves is still determined by the charging of the load, but, depending on the loading

conditions, the contribution of short circuit current may be non negligible possibly prevailing on E_{out} .

In the NAND case, Fig. 8 shows the glitch power dissipated as a function of the skew for different values of the fan-out of driving gates under symmetric conditions (that is $N_A = N_B$). The difference between the curves depends on both E_{out} and E_{cc} (that cannot be neglected for $N_A, N_B > 2$).




Figure 8: Glitch power (in case (a)) evaluated by varying the loading conditions of the driving gates (G_A and G_B) and by keeping fixed the fan-out of the considered gate ($N = 2$).

Notice however that the relative importance of short-circuit power decreases when the ratio between the input and the output slopes is increased [11]. In particular, in the case of Fig. 8 ($N = 2$) for the same value of τ the power increases with N_A, N_B . This is not true when the output load is increased as in Fig. 9 showing a detail of the curves obtained with $N = 8$.




Figure 9: Glitch power evaluated by varying the loading conditions of the driving gates (G_A and G_B) and by keeping fixed the fan-out of the considered gate ($N = 8$).

To complete the analysis of the effects of input slopes, let us consider the case of asymmetric loading of gate fan-in.

At this regard, we have considered $N = 2$ and N_A and N_B that may assume the values 1 or 8. The glitch power dissipated in these cases is shown in Fig. 10 together with that dissipated in the two symmetric cases. The gate inputs transition corresponds to case (a) ($CD = 01 \rightarrow 10$).

As can be seen, the dissipated power is almost the same




Figure 10: Glitch power dissipated in the presence of asymmetric driving slopes.

for low values of the skew and for values larger than the gate propagation delay.

From the point of view of logic level simulation, it should be noticed that the curves evaluated for asymmetric input driving conditions remains between the two limiting curves obtained under symmetric conditions. This suggests the possibility of using suitable averages or worst case conditions without the need of considering all the possible combinations of input loading.

In developing an approximate logic level model of glitch power dissipation, we need to verify that the effect of input slopes can be approximated as a local property of the considered gate and it is not necessary to propagate informations about signal slopes through successive levels of logic.

To this purpose, we have varied the input slopes to driving gates G_A and G_B starting from the case of single fan-out up to fan-out=16. The obtained values of glitch power at gate G are plotted in Fig. 11 as a function of the skew. The data have been evaluated for different values of N_A and N_B and for $N = 2$ as in Fig. 11.

As can be seen, for the same skew the glitch power is rather insensitive on the slope of signals two levels upstream. Therefore, only local information on slopes has to be considered for a correct delay and power modeling.




Figure 11: Glitch power as a function of the skew for asymmetric input loading and fan-out = 2.

3.5 Logic level modeling

Gate level event-driven simulation considers only full swing transitions, and consequently, it implies some approximation in evaluating the power dissipated during hazard generation.

In the transport delay model, no filtering is in order so that glitches are generated whatever the skew between input signals is. In practice, the transport delay model estimates the energy dissipated in hazard generation with the scheme illustrated in Fig. 12, in comparison with a typical result achieved at the electrical level.




Figure 12: Glitch power dissipation evaluated by varying the loading conditions of the gate driving G and by keeping fixed its fan-out.

In the case of the inertial delay model, filtering effects are accounted for by generating only glitches whose amplitude exceeds a given threshold that is typically equal to the gate propagation delay. The corresponding model for energy dissipation estimate is shown in Fig. 12. In this case, differently from the transport delay model that always overestimate power, the error may present different sign depending on whether the actual skew between input waveforms is larger or not than the gate propagation delay.

In both cases, the error strongly depends on the distribution of the skew at the inputs of the considered gate.

These errors can be avoided by modeling the glitch power using input waveforms. At this regard, the previous analysis has established that glitch power dissipation can be studied with sufficient approximation employing local information and it is univoquely determined by the knowledge of:

- gate output loading and parasitic capacitances;
- driving gates loading and driving capabilities (from which the gate input slopes are known);
- input signals values and timing (that allows to determine the skew between input vectors).

All these informations are available in logic simulation (with accurate delays).

4 Hazard propagation

Once generated, a hazard may be propagated through gates whose output is sensitive to the value of the considered line. In this case, some filtering effect is in order: for instance, if the input glitch voltage remains below the conductance threshold of transistors, the gate through which the

glitch attempts to propagate will not dissipate power. The filtering is typically an analog phenomenon since it changes with continuity both the peak voltage and the glitch amplitude, therefore it is hard to model at the logic level.

Glitch power propagation has been studied in a way similar to that used for the case of hazard generation, in practice, with respect to the circuit of Fig. 2, a load has been added to one (referred as G_P) of the gates in the fan-out of the gate G at which the hazard is generated (such load is represented by N_P NAND gates).

The glitch power results will be studied as a function of the same parameter (the skew τ between signals A and B) used to characterize the hazard generation at gate G . In Fig. 13, the glitch power dissipated at gate G_P is plotted as a function of the skew. Results have been achieved by keeping $N = 2$ and by varying N_P . In the same figure, thin lines represent the glitch power dissipated in hazard generation for the same input and output loading conditions (that is $N_{A,B} = 2$ and N varying in the same way of N_P).




Figure 13: Glitch power dissipated at the gate (G_P) where a glitch propagates compared with that dissipated at the gate where the glitch is generated (with the same loading).

As it can be seen, the shape of the curves is almost the same in the two cases, but that achieved for glitch propagation is, in practice, translated. It is worth mentioning that, even if the skew is larger than the propagation delay of G and G_P , the power dissipated is lower than that corresponding to a full gate output transition.

As discussed in the previous section, the inertial delay model attempts to account for the filtering effects of gates, but presents an on/off characteristic that leads to error during hazard generation.

In case of propagation, the inertial delay model produces inaccuracies in the power dissipated at the gate through which the glitch propagates, but the choice of propagating or not the hazard may provoke additional errors. In particular, when the skew between the edges is larger than the gate propagation delay, the hazard is propagated without any attenuation even if the results achieved at the electrical level shows that this is not realistic.

In this way, the logic simulation of the propagation of a hazard along a path may give rise to a remarkable overestimate in evaluating the glitch power.

This can be seen in Fig. 14, showing the power dissipated by gates G and G_P (in case $N = N_P = 2$) as a function of the skew as compared to that evaluated by logic simulation with the inertial delay model. In this latter case, since both gates have the same propagation delay, any skew larger than such parameter gives rise to a hazard that is propagated as a full transition of both gates thus resulting in a consistent error at gate G_P .




Figure 14: Glitch power dissipated at gate G , G_P and their sum as a function of the skew as compared with that evaluated in logic simulation.

It has been verified that such error increases if other gates are added to the path.

5 Conclusions

Glitch power has been analyzed in detail showing the inaccuracies of gate-level event-driven simulation. In the case of hazard generation, results show the feasibility of an accurate modeling since all relevant parameters are available at the logic level. In the case of hazard propagation, instead, a major problem has been detected that is intrinsic to event-driven simulation that cannot track the continuous (analog) filtering of a hazard that propagates through a path. The error in this case may be remarkable since it depends on the length of the considered path.

References

- [1] A. Chandrakasan et al., *J. of Solid State Circuits*, vol. 27, no. 4, pp. 473 – 484, 1992.
- [2] F. Najm et al., in *Proc. of ICCAD*, pp. 204 – 207, 1988.
- [3] M. A. Cirit, in *Proc. of ICCAD*, pp. 534 – 537, 1987.
- [4] S. Devadas et al., *Trans. on CAD*, vol. 11, no. 3, pp. 373 – 383, 1992.
- [5] J. Lin et al., in *Proc. of DAC*, pp. 304 – 309, 1994.
- [6] A. Tyagi, *Proc. of ICCAD*, pp. 530 – 537, 1987.
- [7] A. Deng et al., in *Proc. of ICCAD*, pp. 208 – 211, 1988.
- [8] S. Chowdury et al., *Trans. on CAD*, vol. 9, no. 6, pp. 642 – 654, 1990.
- [9] L. Benini et al., in *Proc. of DAC*, pp. 354 – 360, 1993.
- [10] L. Benini et al., in *Int. Symp. on Low Power Design*, pp. 27 – 32, 1994.
- [11] H. J. M. Veendrick, *J. of Solid State Circuits*, vol. SC-19, pp. 468 – 473, August 1984.