

A Faithful Binary Circuit Model with Adversarial Noise

Matthias Függer*, Jürgen Maier†, Robert Najvirt†, Thomas Nowak‡, Ulrich Schmid†

*CNRS & LSV, ENS Paris-Saclay

†Technische Universität Wien

‡Université Paris-Sud

Abstract—Accurate delay models are important for static and dynamic timing analysis of digital circuits, and mandatory for formal verification. However, Függer et al. [IEEE TC 2016] proved that pure and inertial delays, which are employed for dynamic timing analysis in state-of-the-art tools like ModelSim, NC-Sim and VCS, do not yield faithful digital circuit models. Involution delays, which are based on delay functions that are mathematical involutions depending on the previous-output-to-input time offset, were introduced by Függer et al. [DATE'15] as a faithful alternative (that can easily be used with existing tools). Although involution delays were shown to predict real signal traces reasonably accurately, any model with a deterministic delay function is naturally limited in its modeling power.

In this paper, we thus extend the involution model, by adding non-deterministic delay variations (random or even adversarial), and prove analytically that faithfulness is not impaired by this generalization. Albeit the amount of non-determinism must be considerably restricted to ensure this property, the result is surprising: the involution model differs from non-faithful models mainly in handling fast glitch trains, where small delay shifts have large effects. This originally suggested that adding even small variations should break the faithfulness of the model, which turned out not to be the case. Moreover, the results of our simulations also confirm that this generalized involution model has larger modeling power and, hence, applicability.

I. INTRODUCTION

Modern digital circuit design relies heavily on fast functional simulation tools like Cadence NC-Sim, Mentor Graphics ModelSim or Synopsis VCS, which also allow dynamic timing validation using suitable delay models. In fact, for modern VLSI technologies with their switching times in the picosecond range, static timing analysis may not be sufficient for critical parts of a circuit, where e.g. the presence of glitch trains may severely affect correctness and power consumption. Fully-fledged analog simulations, on the other hand, are often too costly in terms of simulation time.

Delay models like CCSM [9] and ECSM [13] used in gate-level timing analysis tools make use of elaborate characterization techniques, which incorporate technology-dependent information like driving strengths of a gate for a wide range of voltages and load capacitances. Based on these data, dynamic timing analysis tools compute the delay for each gate and wire in a specific circuit, which is then used to parametrize pure and/or inertial delay *channels* (i.e., model components

representing delays). Recall that pure delay channels model a constant transport delay, whereas inertial delay channels [14] allow an input transition to proceed to its output only if there is no subsequent (opposite) input transition within some time window $\Delta > 0$. Subsequent simulation and dynamic timing analysis runs use these pre-computed delays as *constants*, i.e., they are not reevaluated at every point in time.

More accurate simulation and dynamic timing analysis results can be achieved by the *Degradation Delay Model (DDM)*, introduced by Bellido-Díaz et al. [2], [3], which allows channel delays to vary and covers *gradual* pulse cancellation effects.

Függer et al. [7] investigated the *faithfulness* of digital circuit models, i.e., whether a problem solvable in the model can be solved with a real physical circuit and vice versa. Unfortunately, however, they proved that none of the existing models is faithful: for the simple *Short-Pulse Filtration (SPF)* problem, which resembles a one-shot variant of an inertial delay channel, they showed that every model based on *bounded single-history channels* (see below for the definition) either contradicts the unsolvability of SPF in bounded time or the solvability of SPF in unbounded time by physical circuits [11].

Single-history channels allow the input-to-output delay for a given input transition to depend on the time of the *previous* output transition. Formally, a single-history channel is defined by a *delay function* $\delta : \mathbb{R} \rightarrow \mathbb{R}$, where $\delta(T)$ determines the delay of an input transition at time t , given that the previous output transition occurred at time $t - T$. Fig. 1 depicts the involved parameters. Note that T and $\delta(T)$ are potentially negative in the case of a short input pulse, where a new input transition occurs earlier than the just scheduled previous output transition. Together with the rule that non-FIFO transitions cancel each other, this allows to model attenuation and even suppression of glitches. Fig. 2 shows an example input/output-trace generated by a single-history channel. Note that, for *bounded* single-history channels, $\delta(T)$ cannot point arbitrarily far back into the past.

In [6], Függer et al. introduced an *unbounded* single-history channel model based on *involution channels*, which use a delay function $\delta(T)$ whose negative is self-inverse, i.e., fulfills the involution property $-\delta(-\delta(T)) = T$. They proved that, in sharp contrast to bounded single-history channels, SPF cannot be solved in bounded time with involution channels, whereas it is easy to provide an unbounded SPF implementation, which is in accordance with real physical circuits [11]. Hence, binary-

This research was supported by the FATAL (grant P21694) and SIC project (grant P26436-N30) of the Austrian Science Fund (FWF).




Fig. 1: Input/output signal of single-history channel, involving the previous-output-to-input delay T and input-to-output delay $\delta(T)$.




Fig. 2: Single-history channels allow to model pulse attenuation: The delay $\delta(T)$ becomes smaller with smaller previous-output-to-input time T . Observe the cancellation of the second pulse due to non-FIFO-scheduled output transitions.

valued circuit models based on involution channels are faithful with respect to the SPF problem. We note that this actually implies faithfulness also w.r.t. other, practically more relevant problems: analogous to [1], it is possible to implement a one-shot version of a latch (that allows a single up- and a single down-transition of the enable input) using a circuit solving SPF, and vice versa. Consequently, the involution model is also faithful for one-shot latches. Moreover, in [12], Najvirt et al. used both measurements and Spice simulations to show that the involution model can also be made reasonably accurate by suitable parametrization, in the sense that it nicely (though not perfectly) predicts the actual glitch propagation behavior of a real circuit, namely, an inverter chain.

As it is easy to replace the standard pure or inertial delays currently used in VITAL or Verilog models by involution delays, the model is not only a promising starting point for sound formal verification, but also allows to seamlessly improve existing dynamic timing analysis tools.

Main contributions: Notwithstanding its superiority with respect to faithfulness, like every *deterministic* delay model, the involution model has limited modeling power: many different effects in physical circuits cause various types of noise in signal waveforms and, hence, *jitter* in the digital abstraction [4]. No deterministic delay function can properly capture the resulting variability in the signal traces.

In this paper, we relax the involution model introduced in [6] by adding limited *non-determinism* $\eta = [-\eta^-, \eta^+]$, for some fixed $\eta^-, \eta^+ \geq 0$, on top of the (deterministic) involution delay function $\delta(T)$. We prove that this can be done without sacrificing faithfulness: both the original SPF impossibility result and, in particular, a novel SPF possibility hold for this generalized model. We need to stress, however, that adding non-determinism is merely a convenient way of securing maximum generality of our results: no practically observable bounded jitter phenomenon, neither bounded random noise, from white to slowly varying flicker noise [4], nor even *adversarially* chosen transition time variations can invalidate the faithfulness of the resulting η -involution model. Deterministic effects, like slightly different thresholds due to process variations, are of course also covered.

Note carefully that, albeit the non-determinism (η^+ and η^-) must be restricted to ensure faithfulness, the mere fact that we can afford some non-determinism here *at all* is very surprising: comparing the faithful original involution model and the non-faithful DDM model reveals that they primarily differ in handling fast glitch trains, where small delay shifts have large effects. We thus conjectured originally that adding even small non-determinism would break the border between both models, which we now know is not the case.

Our generalization also results in an improved *principal*² modeling accuracy of the η -involution model: thanks to the additional freedom for choosing transition times provided by η , it is obviously easier to match the real behavior of a circuit with some feasible behavior of the circuit in the model. We provide some simulation results (in a similar setting as used in [12]), which demonstrate that it is indeed possible to match the behavior of a real inverter chain with the η -involution model if the variations of operating conditions resp. process variations are small. Whereas this does not hold for larger variations, we observed that excessive deviations occur for relatively large values of T only, which are essentially irrelevant for faithfulness. We are of course aware that more validation experiments, with more complex circuits, will be needed to actually claim good accuracy of the η -involution model, nevertheless, our preliminary results are encouraging.

Regarding applicability, we consider the η -involution model interesting for primarily two reasons: First, it facilitates accurate modeling and analysis of circuits under (restricted) noise, varying operating conditions and parameter variations. Second, to the best of our knowledge, it is the first model that appears to be a suitable basis for the sound formal verification of a circuit, which aims at proving that the circuit meets its specification in *every* feasible trace. We thus believe that our η -involution model might eventually turn out to be an interesting ingredient for a novel verification tool.

Paper organization: In Section II, we provide some indispensable basics of standard involution channels taken from [6]. Section III defines our η -involution model, Section IV provides the proofs for faithfulness. Our simulation results are presented in Section V, and some conclusions and directions of our current/future work are appended in Section VI.

II. THE INVOLUTION MODEL WITHOUT CHOICE

Before we can present the generalized η -involution model with non-deterministic delay variations, we recall the basics from the circuit model introduced in [6].

Signals. A *falling transition* at time t is the pair $(t, 0)$, a *rising transition* at time t is the pair $(t, 1)$. A *signal* is a list of alternating transitions where only finitely many alternating transitions may occur in a finite time interval.

Circuits. Circuits are obtained by interconnecting the external interface, i.e., a set of input and output ports, and a set

²We stress that we do not aim at resolving the non-determinism of the η -involution model to build an accurate simulator in this paper, but rather at providing a model that makes this possible.

of combinational gates via channels. The valid connections are constrained by demanding that gates and channels must alternate on every path in the circuit and that any gate input and output port is attached to only one channel output. Formally we describe a circuit by a directed graph with potentially multiple edges between nodes. Its nodes are in/out ports and gates, and edges are channels. A channel has a channel function, which maps input signals to output signals, whereas a gate is characterized by a (zero-time) Boolean function and an initial Boolean value that defines its output until time 0. Channels connecting input and output ports are assumed to have zero delay, in order to facilitate the composition of circuits.

Executions. An execution of circuit C is an assignment of signals to the vertices and edges of C that respects channel functions, Boolean gate functions, and initial values of gates. Signals on input ports are unrestricted. For an edge c representing a channel with channel function f_c from vertex v in C , we require that the signal s_c assigned to c fulfills $s_c = f_c(s_v)$.

Involution Channels. An involution channel propagates each transition at time t of the input signal to a transition at the output happening after some input-to-output delay $\delta(T)$, which depends on the previous-output-to-input delay T (cf. Fig. 1).

An involution channel function is characterized by two strictly increasing concave delay functions $\delta_\uparrow : (-\delta_\infty^+, \infty) \rightarrow (-\infty, \delta_\infty^+)$ and $\delta_\downarrow : (-\delta_\infty^+, \infty) \rightarrow (-\infty, \delta_\infty^+)$ such that both $\delta_\infty^+ = \lim_{T \rightarrow \infty} \delta_\uparrow(T)$ and $\delta_\infty^+ = \lim_{T \rightarrow \infty} \delta_\downarrow(T)$ are finite and

$$-\delta_\uparrow(-\delta_\downarrow(T)) = T \text{ and } -\delta_\downarrow(-\delta_\uparrow(T)) = T \quad (1)$$

for all T . All such functions are necessarily continuous. For simplicity, we will also assume them to be differentiable; δ being concave thus implies that its derivative δ' is monotonically decreasing. In this paper, we assume all involution channels to be *strictly causal*, i.e., $\delta_\uparrow(0) > 0$ and $\delta_\downarrow(0) > 0$.

A particular and important special case are the so-called *exp-channels*: They occur when gates drive RC-loads and generate digital transitions when reaching a certain threshold voltage V_{th} (typically $V_{th} = 1/2$ of the maximum voltage V_{DD}). We obtain

$$\delta_\uparrow(T) = \tau \ln(1 - e^{-(T+T_p - \tau \ln(\overline{V_{th}}))/\tau}) + T_p - \tau \ln(1 - \overline{V_{th}})$$

$$\delta_\downarrow(T) = \tau \ln(1 - e^{-(T+T_p - \tau \ln(1 - \overline{V_{th}}))/\tau}) + T_p - \tau \ln(\overline{V_{th}}),$$

where τ is the RC constant, T_p the pure delay component and $\overline{V_{th}} = V_{th}/V_{DD}$.

The channel function f_c mapping input signal s to output signal $f_c(s)$ (cp. Fig. 2) is defined via the following algorithm. It can easily be implemented in e.g. VHDL to be used by existing simulators like ModelSim, as these simulators automatically drop transitions on signals violating FIFO order.

Output transition generation algorithm: Let t_1, t_2, \dots be the transitions times of s , set $t_0 = -\infty$ and $\delta_0 = 0$.

- *Initialization:* Copy the initial transition at time $-\infty$ from the input signal to the output signal.
- *Iteration:* Iteratively determine the tentative list of pending output transitions: Determine the input-to-output delay δ_n for the input transition at time t_n by setting




Fig. 3: The η -involution channel: Non-deterministic choice of the tentative output transition after applying $\delta(T)$.

$\delta_n = \delta_\uparrow(t_n - t_{n-1} - \delta_{n-1})$ if t_n is a rising transition and $\delta_n = \delta_\downarrow(t_n - t_{n-1} - \delta_{n-1})$ if it is falling. The n^{th} and m^{th} pending output transitions *cancel* if $n < m$ but $t_n + \delta_n \geq t_m + \delta_m$. In this case, we mark both as canceled.

- *Return:* The channel output signal $f_c(s)$ has the same initial value as the input signal, and contains every pending transition at time $t_n + \delta_n$ that has not been marked as canceled.

III. INTRODUCING ADVERSARIAL CHOICE

We now generalize the circuit model from the previous section to allow a non-deterministic perturbation of the output transition times after the application of the delay functions δ_\uparrow and δ_\downarrow . Note that the resulting output shifts need *not* be the same for all applications of the delay functions; they can vary arbitrarily from one transition to the next. However, each perturbation needs to be within some pre-determined interval $\eta = [-\eta^-, \eta^+]$. These non-deterministic choices can be used to model various effects in digital circuits that cannot be captured by single-history delay functions, ranging from arbitrary types of noise [4] to unknown variations of process parameters and operating conditions. Fig. 3 shows the possible variation of the output transition time caused by the non-deterministic choice.

Formally, we change the notion of the *channel function* to accept an additional parameter: A channel has a channel function, which maps each pair (s, H) to an output signal, where s is the channel's input signal and H is a parameter taken from some suitable set of admissible parameters (see below). We also adapt the definition of an *execution* to allow an adversarial choice of H : For an edge c from v in C , we require that there exists some admissible parameter H such that the signal s_c fulfills $s_c = f_c(s_v, H)$.

For η -involution channels, we let the admissible parameters H be any sequence of choices $\eta_n \in \eta$. The output transition generation algorithm's *Iteration* step for the n^{th} transition of the input signal is adapted as follows: $\delta_n = \delta_\uparrow(t_n - t_{n-1} - \delta_{n-1}) + \eta_n$ if t_n is a rising transition and $\delta_n = \delta_\downarrow(t_n - t_{n-1} - \delta_{n-1}) + \eta_n$ if it is falling.

Fig. 4 depicts two example signal traces, out_1 and out_2 , obtained by an η -involution channel with the same underlying δ as the one in Fig. 2. Observe that the adversary has the freedom to “de-cancel” pulses that would have canceled according to the delay function (second pulse in out_2), extend pulses (first pulse in out_1), and shift pulses (first pulse in out_2).




Fig. 4: The η -involution channel covers pulse attenuation under (bounded) adversarial noise, varying operating conditions, parameter variations and other modeling inaccuracies.

IV. FAITHFULNESS OF INVOLUTION CHANNELS WITH ADVERSARIAL CHOICE

In this section, we will prove that η -involution channels are faithful with respect to *Short-Pulse Filtration (SPF)*.

A pulse of length Δ at time T has initial value 0, one rising transition at time T , and one falling transition at time $T + \Delta$.

Definition 1 (Short-Pulse Filtration). A circuit with a single input and a single output port solves Short-Pulse Filtration (SPF), if it fulfills the following conditions for all admissible channel function parameters H :

- F1) The circuit has exactly one input and one output port. (*Well-formedness*)
- F2) A zero input signal produces a zero output signal. (*No generation*)
- F3) There exists an input pulse such that the output signal is not the zero signal. (*Nontriviality*)
- F4) There exists an $\varepsilon > 0$ such that for every input pulse the output signal never contains a pulse of length less than ε . (*No short pulses*)

Note that we allow the SPF circuit to behave arbitrarily if the input signal is not a (single) pulse.

To show faithfulness of the η -involution model, we start with the trivial direction: we prove that no circuit with η -involution channels can solve the bounded-time variant of SPF (where the output must stabilize to constant 0 or 1 within bounded time). Note that this matches the well-known impossibility [10] of building such a circuit in reality. Indeed, the result immediately follows from the fact that the adversary is free to always choose $\eta_n = 0$, i.e., make the η -involution channels behave like involution channels. In [6], [5], it has been shown that no circuit with involution channels can solve bounded-time SPF, which completes the proof.

What hence remains to be shown is the existence of a circuit that solves SPF (with unbounded stabilization time) with η -involution channels. We can prove that the circuit shown in Fig. 5, which consists of a fed back OR-gate forming the storage loop and a subsequent buffer with a suitably chosen (high) threshold voltage (modeled as an exp-channel), does the job. As a consequence, a circuit model based on




Fig. 5: A circuit solving unbounded SPF, consisting of an OR-gate, with initial value 0, fed back by channel c , and a high-threshold buffer HT.

η -involution channels enjoys the same faithfulness as the involution channels of [6], even though its set of allowed behaviors is considerably larger.

Informally, we consider a pulse of length Δ_0 at time 0 at the input and reason about the behavior of the feed-back loop, i.e., the output of the OR gate. There are 3 cases: If Δ_0 is small, then the pulse is filtered by the channel in the feed-back loop. If it is big, the pulse is captured by the storage loop, leading to a stable output 1. For a certain range of Δ_0 , the storage loop may be oscillating, possibly forever. In any case, however, it turns out that a properly chosen exp-channel can translate this behavior to a legitimate SPF output.

Theorem 2. Consider the circuit in Fig. 5 subject to constraint $\eta^+ + \eta^- < \delta_{\downarrow}(-\eta^+) - \delta_{\min}$. The fed-back OR gate with a strictly causal η -involution channel has the following output when the input pulse has length Δ_0 :

- If $\Delta_0 \geq \delta_{\infty}^+ + \eta^+$, then the output has a single rising transition at time 0.
- If $\Delta_0 \leq \delta_{\infty}^+ - \delta_{\min} - \eta^+ - \eta^-$, then the output only contains the input pulse.
- If $\delta_{\infty}^+ - \delta_{\min} - \eta^+ - \eta^- < \Delta_0 < \delta_{\infty}^+ + \eta^+$, then the output may resolve to constant 0 or 1, or may be an (infinite) pulse train, with $\Delta_n \leq \Delta$ for some $0 < \Delta < \delta_{\min}$ and duty cycle $\gamma_n \leq \gamma = \frac{\Delta}{\delta_{\uparrow}(-\Delta) + \eta^+} < 1$ for $n \geq 1$. If $\Delta_n > \Delta$ for some n , bounded time later, the output resolves to 1 and $\Delta_m > \Delta_{m-1}$ for all $m > n$.

Finally, a high-threshold buffer with arbitrary threshold can be modeled by an exp-channel with properly chosen V_{th} :

Lemma 3 ([5, Lem. 14]). Let $\Theta > 0$ and $0 \leq \Gamma < 1$. Then, there exists an exp-channel C such that every finite or infinite pulse train with pulse lengths $\Theta_n \leq \Theta$, $n \geq 0$, and duty cycles $\Gamma_n \leq \Gamma$, $n \geq 1$, is mapped to the zero signal by C .

By choosing $\Gamma = \gamma(1 + \varepsilon) < 1$ for some $\varepsilon > 0$ sufficiently small and Θ so large that the feed-back loop in Figure 5 has already locked to constant 1 at time $T + \Theta$, where T is the time when some pulse Δ_n , $n \geq 1$, of the feed-back loop with duty cycle $\gamma(1 + \varepsilon)$ has started, we get the following: If SPF input pulse lengths Δ_0 and adversarial choices are such that no Δ_n reaches duty cycle $\gamma(1 + \varepsilon)$, the output of the exp-channel is constant zero; otherwise, there is a single up-transition (occurring only after $T + \Theta$) at the output. Therefore:

Theorem 4. There is a circuit that solves unbounded SPF.

V. SIMULATIONS

In this section, we complement the proof of faithfulness provided in the previous section with simulation experiments




Fig. 6: Schematics of the ASIC used for validation measurements. It combines an inverter chain with analog high-speed sense amplifiers.




Fig. 7: Measured δ_{\downarrow} for UMC-90 inverter chain for $V_{DD} \in \{0.3, 0.4, 0.6, 0.7, 0.8, 1\}$ V and simulated (dashed brown) δ_{\downarrow} for $V_{DD} = 0.6$ V, taken from [12, Fig. 7].

and measurement results, which confirm that our η -involvement model indeed captures reality better than the original involvement model [12]. Whereas more experiments, with different technologies and more complex circuits (including multi-input gates), would be needed to actually claim improved model coverage, our results are nevertheless encouraging.

We employ the same experimental setup as in [12], which uses UMC-90 nm and UMC-65 nm bulk CMOS 7-stage inverter chains as the primary targets. For UMC-65, we resorted to Spice simulations of a standard cell library implementation, for UMC-90, we relied on a custom ASIC [8]. The latter provides a 7-stage inverter chain built from 700 nm x 80 nm (W x L) pMOS and 360 nm x 80 nm nMOS transistors, with threshold voltages 0.29 V and 0.26 V, respectively, and a nominal supply voltage of $V_{DD} = 1$ V. As all inverter outputs are connected to on-chip low-intrusive high-speed analog sense amplifiers (gain 0.15, -3 dB cutoff frequency 8.5 GHz, input load equivalent to 3 inverter inputs), see Fig. 6, which can directly drive the 50 Ω input of a high-speed real-time oscilloscope, the ASIC facilitates the faithful analog recording of all signal waveforms. Independent power supplies and grounds for inverters and amplifiers also facilitate measurements with different digital supply voltages V_{DD} .

For convenience, we provide the delay functions determined in [12] in Fig. 7 (δ_{\downarrow} for UMC-90, measurements).

In order to validate the η -involvement model, we use the following general approach: Given simulated/measured output waveforms of a single inverter excited by input pulses of different width, we compare (i) the digital output obtained from the simulated/measured waveforms with (ii) the pre-

dictions for some given delay function. The differences of the transition times of predicted and real digital output is a measure of modeling inaccuracy of the original involvement model. If these differences can be compensated by suitable output shifts within $[\eta^-, \eta^+]$, however, we can claim that the η -involvement model matches the real behavior of the circuit for the given waveforms. Since faithfulness puts the severe constraint $\eta^+ + \eta^- < \delta_{\downarrow}(-\eta^+) - \delta_{min}$ on η^+, η^- , it is not clear under which conditions this claim indeed holds. In our evaluation, η^+ was first set to a suitable value ($\eta^+ > 0$) and afterwards η^- was calculated according to $\eta^- = \delta_{\downarrow}(-\eta^+) - \delta_{min} - \eta^+$. Clearly, this results in different η bounds in each of the figures below.

The particular questions addressed in our experiments are the following: Is the allowed range for η^+ and η^- sufficient for the η -involvement model to capture: (a) The circuit behavior under variations of certain operating conditions. After all, circuit delays change with varying supply voltage and temperature, so the question remains to what extent the resulting fluctuations are covered by the η -involvement model. (b) The circuit behavior under process variations. In general, circuit delays vary from manufactured chip to chip, so the question arises whether the η -involvement model based on a “typical” delay function covers typical process variations. (c) The real behavior of our inverter chain with a (suitably parametrized) standard involvement function, in particular for exp-channels. This would simplify model calibration, as it is typically easier to determine the exp-channel model parameters for a given circuit [2], rather than its entire delay function.

To investigate question (a), i.e., the robustness against voltage variations, we added a sine wave to the voltage supply source (nominally 1.2 V = V_{DD}) with a period similar to the full range switching time of the inverter and a magnitude of 0.012 V (1 % of V_{DD}). We applied pulses with differing width to the input of the inverter and recorded the output, whereat the phase of the sine wave was set for each pulse randomly between 0 and 360 degrees. In Fig. 8a, the deviation D between the prediction and the actual crossing over the previous-output-to-input delay T is shown. Despite the stringent bounds on η , it is possible to fully cover the resulting delay variations for low T , for higher values however, the η -involvement model does no longer apply. Please note that the huge difference between δ_{\downarrow} and δ_{\uparrow} can be easily explained by the fact that δ_{\uparrow} results in a falling transition at the output of the inverter. For this transition, the transistor connecting the output to the power supply gets closed more and more, reducing also the impact of the voltage variations. (When varying the ground level, the reverse case can be observed.)

To answer question (b), we chose to vary the transistor width, which increases/decreases the maximum current and allows us to model variations of resistance and capacitance as well. The simulations themselves were carried out in the same fashion as described in the last paragraph, except that $V_{DD} = 1.2$ V was constant. Fig. 8b shows the results for 10 % wider transistors, where the η -bound is even bigger than required. In contrast, the deviations for 10 % narrower ones

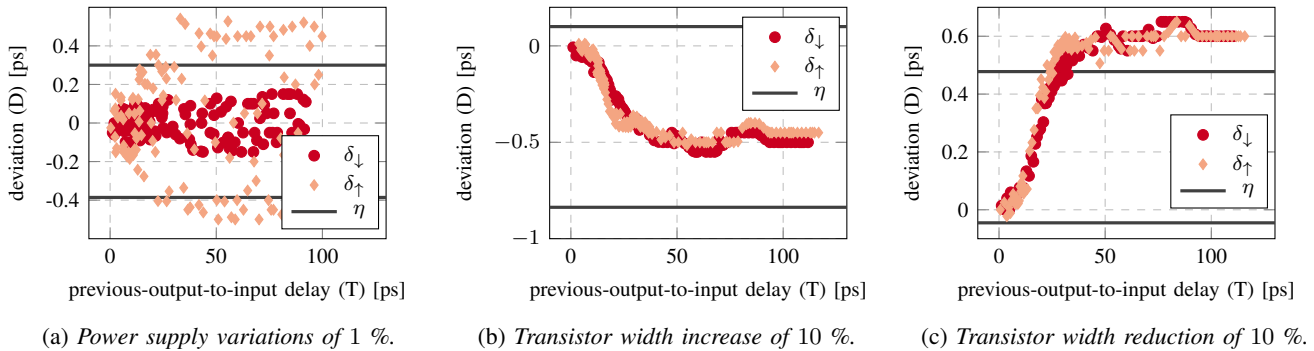


Fig. 8: Deviation between predicted and actual V_{TH} crossings for different variations.




Fig. 9: Fitting an exp-channel involution to measured data.

(Fig. 8c) exceed the η -bound with increasing values of T . Unlike V_{DD} variations, varying transistor sizes, as expected, either increases or decreases the delay. This can be seen very clearly in the figures, as one trace is well below and one well above $D = 0$.

For question (c), we tried to fit the parameters of the involution function (2) for exp-channels w.r.t. the measurement data published in [12] and evaluated the deviations D between the resulting model predictions and the real digital output. Whereas the deviations over the whole range of T exceed the feasible η -bounds, one can observe that even this very simple exp-channel only results in minor mispredictions near $T = 0$. As shown in Fig. 9, it again turns out that, when using the resulting involution function, excessive deviations occur (quite naturally) for large values of T only.

We hence conclude that the η -involution model indeed improves the modeling accuracy of the original involution model, despite the fact that the allowed non-determinism, i.e., η , is quite restricted. Moreover, our simulation experiments indicate that the absolute deviations $|D|$ between model predictions and real traces is increasing with increasing previous-output-to-input delay T , making it possible to fully compensate D via η near $T = 0$. This is crucial, as our η -bounds result from proving faithfulness, which involves the range $T \in [-\delta_{min}, 0]$ only. For larger T , D grows bigger, but in this region, it might be feasible to also increase the allowed non-determinism as these values are almost irrelevant w.r.t. faithfulness.

VI. CONCLUSIONS AND FUTURE WORK

We proved the surprising fact that adding non-determinism to the delays of involution channels, the only delay model

known so far that is faithful for the SPF problem, does not invalidate faithfulness. As confirmed by some simulation experiments and even measurements, noise, varying operating conditions and process parameter variations hence do not a priori rule out faithful continuous-time, binary value models. Part of our future work will be devoted to further increase the level of non-determinism sustained by our model, the handling of more complex circuits, and the first steps for incorporating the η -involution model in a suitable formal verification tool.

REFERENCES

- [1] José C. Barros and Brian W. Johnson. Equivalence of the arbiter, the synchronizer, the latch, and the inertial delay. *IEEE ToC*, 32(7):603–614, 1983.
- [2] M. J. Bellido-Díaz, J. Juan-Chico, A. J. Acosta, M. Valencia, and J. L. Huertas. Logical modelling of delay degradation effect in static CMOS gates. *IEE Proceedings – Circuits, Devices, and Systems*, 147(2):107–117, 2000.
- [3] Manuel J. Bellido-Díaz, Jorge Juan-Chico, and Manuel Valencia. *Logic-Timing Simulation and the Degradation Delay Model*. Imperial College Press, London, 2006.
- [4] C. E. Calosso and E. Rubiola. Phase noise and jitter in digital electronics. *arXiv:1701.00094*, 2016.
- [5] Matthias Függer, Robert Najvirt, Thomas Nowak, and Ulrich Schmid. Faithful glitch propagation in binary circuit models. *arXiv:1406.2544*, 2014.
- [6] Matthias Függer, Robert Najvirt, Thomas Nowak, and Ulrich Schmid. Towards binary circuit models that faithfully capture physical solvability. In *Proceedings of the 2015 Design, Automation & Test in Europe Conference & Exhibition, DATE '15*, pages 1455–1460, San Jose, CA, USA, 2015. EDA Consortium.
- [7] Matthias Függer, Thomas Nowak, and Ulrich Schmid. Unfaithful glitch propagation in existing binary circuit models. *IEEE Transactions on Computers*, 65(3):964–978, March 2016.
- [8] Michael Hofbauer, Kurt Schweiger, Horst Dietrich, Horst Zimmermann, Kay-Obbe Voss, Bruno Merk, Ulrich Schmid, and Andreas Steininger. Pulse shape measurements by on-chip sense amplifiers of single event transients propagating through a 90 nm bulk CMOS inverter chain. *IEEE Transactions on Nuclear Science*, 59(6):2778–2784, December 2012.
- [9] Synopsis Inc. CCS timing library characterization guidelines, October 2016. Version 3.4.
- [10] Leonard R. Marino. The effect of asynchronous inputs on sequential network reliability. *IEEE ToC*, 26(11):1082–1090, 1977.
- [11] Leonard R. Marino. General theory of metastable operation. *IEEE ToC*, 30(2):107–115, 1981.
- [12] Robert Najvirt, Ulrich Schmid, Michael Hofbauer, Matthias Függer, Thomas Nowak, and Kurt Schweiger. Experimental validation of a faithful binary circuit model. In *Proceedings of the 25th Edition on Great Lakes Symposium on VLSI, GLSVLSI '15*, pages 355–360, New York, NY, USA, 2015. ACM.
- [13] Cadence Design Systems. Effective current source model (ECSM) timing and power specification, January 2015. Version 2.1.2.
- [14] Stephen H. Unger. Asynchronous sequential switching circuits with unrestricted input changes. *IEEE ToC*, 20(12):1437–1444, 1971.