How Fast Will the Flip Flop? *

Mark R. Greenstreet and Peter Cahoon
Department of Computer Science, University of British Columbia
Vancouver, BC Canada

Abstract

This paper describes an experimental investigation of the application of dynamical systems theory to the verification of digital VLSI circuits. We analyze the behavior of a nine-transistor toggle element using a simple, SPICE-like model. We show how such properties as minimum and maximum clock frequency can be identified from topological features of solutions to the corresponding system of differential equations. This dynamical systems perspective also gives a clear, continuous-model interpretations of such phenomena as dynamic storage and timing hazards.

Keywords: Dynamical systems, hardware verification, hybrid models, real time systems

1 Introduction

Most verification of VLSI designs, synchronous and asynchronous, assumes discrete models for signal values and transition times. These discrete models lend themselves well to event-driven simulation [3], model checking [4], and theorem proving [17]. However, many important circuit phenomena cannot be modeled with discrete time and values, and failure to account for these phenomena can lead to faulty designs. These problems are especially apparent in the design of asynchronous circuits where computation is driven by internal events and not regulated by an external clock. This has led to many heuristic guidelines for designing such circuits referring to such things as "monotonic transitions," "isochronic forks" [12], and debates of "interleaving semantics" versus "true concurrency". Underlying these issues is a more basic question, "can discrete models of circuit behavior be based on a physically sound model of circuit behavior?"

The question is not new. About a decade ago, the "qualitative physics" community addressed issues of circuit behavior [6, 18]. These researchers performed very coarse integrations of circuit equations transformed into discrete domains. Non-determinism in the discrete domain was used to ensure that every continuous behavior had a corresponding discrete behavior. However, to make this guarantee, the discrete models tended to admit non-physical behaviors that prevent verification of typical VLSI designs. Noting this limitation, others have argued that verification of hybrid systems should be based on physics as expressed by systems of differential equations [15].

Another attempt at discretizing continuous models is described Kurshan and McMillan in [11]. As in the "qualitative physics" approach, Kurshan and McMillan partition the phase space of the circuit into fixed boxes, and map these boxes to discrete states. However, they integrate the boundaries of these boxes over fixed intervals of the *continuous* time domain to obtain a non-deterministic discrete state transition relation. Verification is performed by model checking this relation. By carefully choosing the sizes of the phase space boxes, they were able to verify an arbiter design of Seitz [16]. To perform the integration efficiently, they made several convexity assumptions that would be violated by a more detailed circuit model (e.g. one with non-linear capacitances). Their method might accommodate such models by choosing a sufficiently fine discretization, but it is not clear that this can be done without leading to an intractable state space explosion.

*This research was supported by operating grants OGPO 109688 and OGPO 138501 from the Natural Sciences Research Council of Canada, a fellowship from the Advanced Systems Institute, and an equiment donation from SUN Microsystem Canada, Ltd.

We favor an approach that bases the mapping from continuous to discrete behaviors on dynamical systems theory. Instead of quantizing the continuous system, discrete behaviors are recognized from topological features of the solutions of the system of differential equations. Examples of this approach include Hurtado [10] who examines arbiters from a dynamical systems theory perspective, Brockett [2] who describes systems of differential equations that perform simple computations such as counting, and Black [1] who proposes dynamical systems theory as a possible basis for an asynchronous design methodology.

Our application of dynamical systems theory differs from those mentioned in the previous paragraph in that our approach is analytical whereas the others view it from a perspective of synthesis. To be more specific, prior work has focused on the problem of finding a system of differential equations that performs a desired computation. We, on the other hand, start from an existing design and extract a system of differential equations from the circuit. We then seek to verify that this behavior is in fact realized by the extracted equations.

This paper describes an experimental approach to applying dynamical systems theory to VLSI design. We take an existing design for a toggle flip flop and attempt to show that, under a suitable abstraction, its discrete behavior satisfies a discrete specification. To derive the possible continuous behaviors, we employ circuit simulation. The observations presented in this paper provide explanations of the circuit behavior, but these explanations are not mathematical proofs. Our simulation based approach has revealed features of the circuit dynamics that have not, to our knowledge, been considered by more synthesis oriented approaches. We believe that the dynamical features that we have observed are essential to understanding the operation of real VLSI circuits. In particular, we show the following:

- The discrete behavior of the toggle element corresponds to a periodic attractor of the continuous dynamical system.
- Dynamic storage leads to degeneracies that preclude a one-to-one mapping between attractors and states.
- Local properties of attractors offer insight into circuit dynamics, but robust verification will require analysis of more global properties than can be established by simple simulation.

2 The toggle element

This section presents the toggle element that is used as an example in the remainder of the paper. The design was originally described by Yuan and Svensson [20]. The single-phase clocking technique proposed in that paper has received great interest because of the high clock rates that can be achieved. Using a single clock, these designs have inherent timing races that we hope to elucidate with the analysis presented here. We chose the toggle element as our example because it is a simple circuit, yet its behavior in a continuous model exemplifies many of the issues that we believe are crucial to understanding typical VLSI designs.

A discrete specification of the toggle is simple. The circuit has an input, ϕ, and an output, $\mathtt{z}$. Every time the ϕ input makes a low-to-high transition, the $\mathtt{z}$ output should make a single transition, either from low to high, or from high to low. When the ϕ input makes a high-to-low transition, the $\mathtt{z}$ output should remain unchanged. Thus, the output changes once for every other change of the input.

Figure 1 shows Yuan and Svensson's implementation of the toggle. The operation of this circuit can be understood by using a simple switch model starting from a state where the ϕ input is low. In this case, $\mathtt{y}$ will eventually become high, $\mathtt{z}$ is floating, and $\mathtt{x}$ is the logical negation of $\mathtt{z}$. If we assume that the value stored on node $\mathtt{z}$ is a well-defined logical value, then the circuit has two possible states when ϕ is low: $(\mathtt{x}, \mathtt{y}, \mathtt{z}) = (L, H, H)$, and $(\mathtt{x}, \mathtt{y}, \mathtt{z}) = (H, H, L)$. Starting from these two states, we can derive the corresponding stable successor states for when ϕ is high. If the circuit is allowed to reach a stable state before each transition of the ϕ input, then it implements a toggle as illustrated by the state transition diagram shown in figure 2.

The description of the toggle element in the previous paragraph assumed a switch-level model. Al-

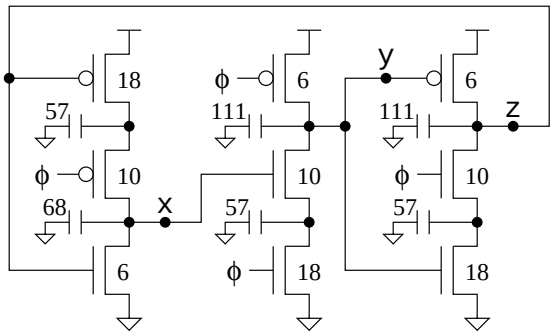

Transistors are labeled with shape factors, W/L.
Capacitances are in femptofarads.

Figure 1: The Yuan-Svensson toggle element.

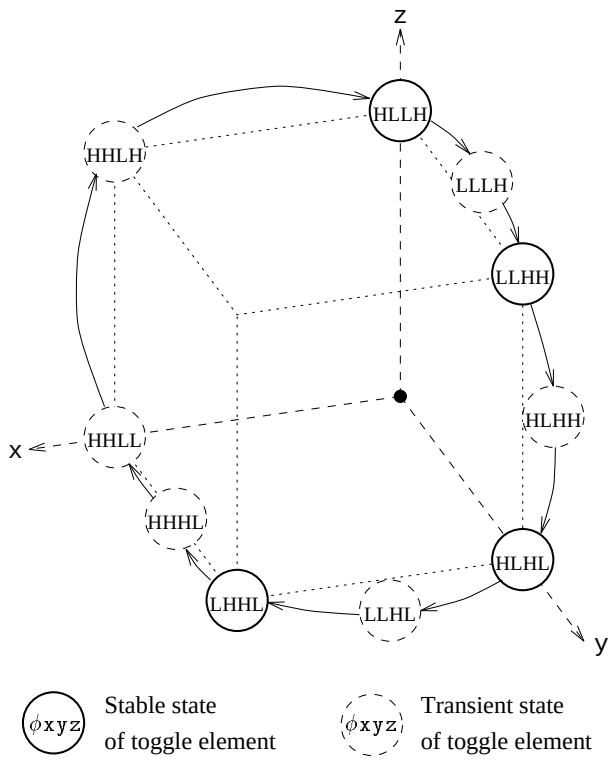

Figure 2: State transition diagram for the toggle element.

though such models are helpful when seeking an intuitive understanding of how such a circuit operates, they fail to address many analog aspects of the circuit operation that are critical to ensuring its proper behavior. Because the circuit uses a single-phase clocking scheme, there are timing hazards associated with each clock transition. For example, consider the transition from state (L, H, H) to state (L, H, L) when ϕ makes a low-to-high transition. If the high-to-low transition of `z` occurs before ϕ is completely high, then `x` may make a spurious low-to-high transition. This in turn could enable a high-to-low transition of `y` and cause `z` to return to its original high value. In section 5, we will show that such behaviors can actually arise with realistic circuit models. Other anomalous behaviors can arise as a consequence of precharged logic when spurious capacitances introduce coupling to logically unrelated signals. The Yuan-Svensson design employs techniques such as single-phase clocking and precharged logic to achieve high clock rates. Switch-level models are inadequate to verify the behavior of such designs because the behavior of the circuit *during* a clock transition is critical for the correct operation of the circuit. This precludes considering clock transitions as instantaneous events. Furthermore, traditional analog simulation cannot establish that the design is robust under a specified range of device parameters and input signal characteristics. It is for these reasons that we consider a dynamical systems approach for analyzing these circuits.

3 Dynamical systems

Dynamical systems theory provides a basis for mapping the continuous behaviors of circuit models formulated as systems of differential equations to the discrete behaviors of digital models. In this approach, we consider the phase space defined by the voltages of each node in the circuit; the state of the circuit is given by a vector of voltages, $\boldsymbol{v}$. The circuit equations give us $\dot{\boldsymbol{v}}$, the time derivative of $\boldsymbol{v}$, as a function of $\boldsymbol{v}$. Integrating these equations gives a unique trajectory for each state. Since we employ deterministic circuit models, this trajectory specifies the entire past and future behavior of the circuit. An *attractor* is

a possibly infinite set of points in phase space such that all trajectories starting in some neighborhood of this set asymptotically approach the set in the limit as time goes to $+\infty$. The neighborhood associated with each attractor is called its *basin of attraction*, and the basins of attraction of a dynamical system partition the phase space. Thus, if a dynamical system has a finite set of attractors, we can identify a finite partition of its continuous phase space. This partitioning provides a discrete interpretation of continuous behaviors. To analyze the toggle element, we consider two broad classes of attractors:

Fixed-point attractors: the trajectory asymptotically approaches a fixed point.
It is tempting to try to identify fixed-point attractors with digital states. However, as we shall describe in section 5, circuits that use dynamic storage can have extended, connected limit sets for which such a simplistic interpretation does not apply.

Periodic attractors: the trajectory enters an orbit of fixed period.
We can map this continuous orbit to a cycle of states of the finite state machine description of the circuit in a digital model. In section 6 we shall establish conditions under which trajectories for the toggle element converge to periodic orbits with periods twice that of the input clock.

In general, dynamical systems can exhibit other behaviors such as chaotic attractors and divergent flows. Although it is possible to design circuits with chaotic behavior, such designs do not seem to have widespread application. Divergent flows can be excluded on simple physical grounds (i.e. the voltages on a chip are bounded).

The system of differential equations that define a dynamical system may have parameters such as the mobility of carriers in the channel region of a transistor or the frequency of a periodic input. The value of a parameter is fixed over time as the behavior of the dynamical system evolves. By varying the value of the parameters, a family of dynamical systems is produced. At critical values of a parameter, qualitative changes may occur in the structure of the phase space as characterized by its system of attractors. For example, fixed-point attractors may be created or destroyed, or a fixed-point attractor may be replaced by a periodic attractor (the latter example is known as a Hopf bifurcation [8]). These bifurcations are discrete events that occur at precise values of the parameters.

The choice of which quantities to represent as variables and which to represent as parameters depends largely on what behaviors are of interest. Certain quantities, such as gate oxide thickness are naturally modeled as parameters, as we don't expect these quantities to vary with time for any given physical transistor. On the other hand, quantities such as the voltage or frequency of an input signal can be modeled as either a parameter or a variable depending on the objectives of the analysis.

4 A continuous circuit model

Our goal is to assess the feasibility of a dynamical systems approach to analyzing the behavior of CMOS circuits. Accordingly, we have started by using simple, intuitive circuit models. If we are successful in this framework, we hope that such success will provide a foundation for applying dynamical systems theory to more accurate circuit models.

We view a circuit as a collection of transistors and capacitors. A transistor is a device whose drain-to-source current is determined by the equations given in figure 3. This is a simple first order model that neglects such phenomena such as the body-effect, subthreshold conduction, and the dependence of I_{ds} on V_{ds} voltage in the saturation region (see [7]). In this paper, we use "typical" parameters for the MOSIS 2μ n-well process: $K_n = 2.3 * 10^{-5}\,\text{amps/volt}^2$, $V_{tn} = 1.0\text{volt}$, $K_p = -8.4 * 10^{-6}\,\text{amps/volt}^2$, $V_{tp} = -1.0\text{volt}$. S is the shape factor for the transistor.

We assume that all capacitors are of fixed value, that all capacitances are to ground, and that there is a positive capacitance from every node to ground. These capacitances are represented by a diagonal matrix, $\boldsymbol{C}$, where $\boldsymbol{C}(i,i)$ is the capacitance from node i to ground. In this paper, we ignore resistances and inductances. Again, this is a simplistic model that leads to an intuitive understanding of the circuit dynamics.

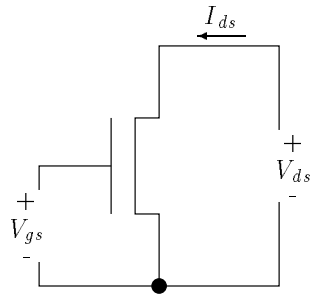

$$I_{ds} = \begin{cases} 0 & \text{if } V_{gs} - V_t \leq 0 \\ KS(V_{gs} - V_t)^2 & \text{if } 0 \leq V_{gs} - V_t \leq V_{ds} \\ \frac{KS}{2} V_{ds}(2(V_{gs} - V_t) - V_{ds}) & \\ & \text{if } V_{ds} \leq V_{gs} - vt \end{cases}$$

Figure 3: A simple transistor model.

Given a circuit state, $\boldsymbol{v}$, we use the equations given in figure 3 to determine the net current flowing out of each node due to drain-to-source currents. Let $\boldsymbol{m}(\boldsymbol{v})$ denote this current vector. By the definition of capacitance, the current flowing out of each node through the capacitors is given by the vector $\boldsymbol{C}\dot{\boldsymbol{v}}$, and by Kirchoff's current law, we have $\boldsymbol{m}(\boldsymbol{v}) = -\boldsymbol{C}\dot{\boldsymbol{v}}$. Thus, our differential-equation model for a circuit is

$$\dot{\boldsymbol{v}} = -\boldsymbol{C}^{-1}\boldsymbol{m}(\boldsymbol{v})$$

Although it is straightforward to derive these equations, finding closed-form solutions is, in most practical applications, impossible. Accordingly, the results presented in this paper are based on numerical approximations to these solutions obtained using a fourth order Runge-Kutta integrator [14].

5 Static analysis

We now consider the equilibrium behavior of the toggle element when ϕ is considered as a parameter taking on fixed values in the range of 0 to 5 volts. We call this the "static" behavior of the toggle element. In the next section, we will consider the behavior when ϕ is considered as a variable, in particular as an input driven by a sinusoidal source, and we will refer to this as the "dynamic" behavior of the circuit.

When ϕ is low, we might hope to find a pair of fixed-point attractors corresponding to the states $(\mathtt{x}, \mathtt{y}, \mathtt{z}) = (\mathtt{L}, \mathtt{H}, \mathtt{H})$ and $(\mathtt{H}, \mathtt{H}, \mathtt{L})$ and, when ϕ is high, a pair corresponding to $(\mathtt{L}, \mathtt{H}, \mathtt{L})$ and $(\mathtt{L}, \mathtt{L}, \mathtt{H})$. We might further hope that the basins of attraction for ϕ low and those for ϕ high would overlap in such a way as to give rise to the desired behavior of the toggle element. In this section, we will examine the phenomena that preclude such a simple, static characterization of the toggle's behavior.

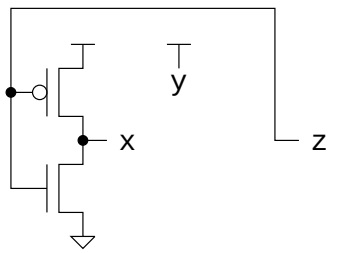

Figure 4: The toggle with $\phi \leq V_{tn}$

Consider first the behavior when $\phi < V_{tn}$. As shown in figure 4, $\mathtt{x}$ is the output of an inverter whose input is $\mathtt{z}$, $\mathtt{y}$ is driven high, and $\mathtt{z}$ is floating. There is a fixed point for any value of $\mathtt{z}$ between $-V_{tn}$ and $V_{dd} + V_{tp}$. This set of fixed points (called the ω-limit set, see [8]) forms a curve that corresponds to the transfer function of the inverter that drives $\mathtt{x}$. This curve connects the points corresponding to the states $(\phi, \mathtt{x}, \mathtt{y}, \mathtt{z}) = (\mathtt{L}, \mathtt{L}, \mathtt{H}, \mathtt{H})$ and $(\mathtt{L}, \mathtt{H}, \mathtt{H}, \mathtt{L})$. Thus, it is difficult to justify distinguishing these two states based on the static behavior of the toggle when $\phi = 0$.

The assumption that the voltage of $\mathtt{z}$ can be fixed at any value neglects the currents due to phenomena such as subthreshold conduction and the reverse-bias currents across diffusion-substrate junctions. However, a more accurate model that considers these currents will not provide the two states that we hoped to see. Instead, it will reveal that the voltage of $\mathtt{z}$ slowly drifts to some value between ground and V_{dd}, and the ω-limit set collapses to a single point, not two distinct states.

When the voltage of ϕ slightly exceeds V_{tn}, $\mathtt{z}$ is pulled to ground, which causes $\mathtt{x}$ to approach V_{dd}. This exerts a weak pull-down on $\mathtt{y}$, but for sufficiently low voltages of ϕ, the p-channel pull-up will dominate.

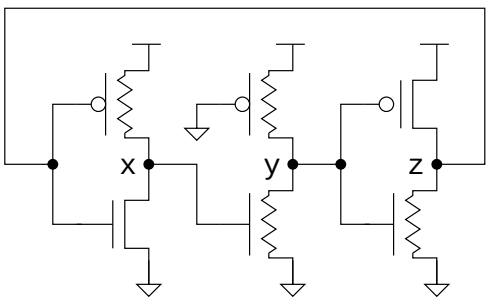

Figure 5: The toggle with $V_{tn} < \phi < V_{dd} + V_{tp}$.

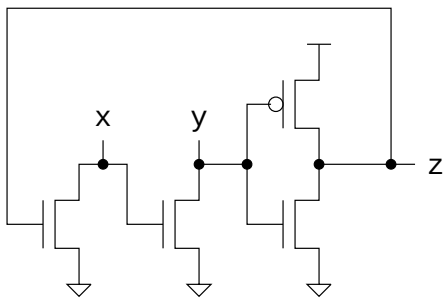

Figure 6: The toggle with $V_{dd} + V_{tp} < \phi$.

Thus, when the voltage of ϕ slightly exceeds V_{tn}, the ω-limit set collapses to a single point corresponding to the state $(\mathtt{x}, \mathtt{y}, \mathtt{z}) = (\mathtt{H}, \mathtt{H}, \mathtt{L})$.

In general, when ϕ is between V_{tn} and $V_{dd} + V_{tp}$, the circuit functions as a ring of three inverters (see figure 5), and we are not surprised to observe a Hopf-bifurcation where the single fixed-point attractor is replaced by a periodic attractor. By simulation using the model parameters given in figure 3, the bifurcation occurs at ϕ roughly equal to 1.78 volts. The toggle element continues to oscillate for any value of ϕ between this threshold and $V_{dd}+V_{tp}$. The upper threshold can be justified by considering operation with ϕ slightly less than $V_{dd} + V_{tp}$. All three stages of the toggle element continue to function as inverters with high-gain transfer functions. The two dominant poles are established by the "inverters" driving nodes $\mathtt{x}$ and $\mathtt{y}$, and these poles are at comparable frequencies. Thus, this circuit satisfies the Nyquist criterion for oscillation [9], with a period of oscillation that grows without bound as ϕ approaches $V_{dd} + V_{tp}$.

Figure 6 shows the equivalent circuit for the toggle element when the voltage of ϕ is greater than $V_{dd}+V_{tp}$. We consider points in the ω-limit set. The voltage of node $\mathtt{z}$ is determined by the voltage of node $\mathtt{y}$. If the voltage of $\mathtt{z}$ is greater than V_{tn}, then the voltage of $\mathtt{x}$ is zero. Otherwise, the $\mathtt{x}$ can take on any value between zero and V_{tn}. In a discrete interpretation, $\mathtt{x} = \mathtt{L}$ at any point in the ω-limit set, and this set includes a curve that connects the points corresponding to the states $(\phi, \mathtt{x}, \mathtt{y}, \mathtt{z}) = (\mathtt{H}, \mathtt{L}, \mathtt{L}, \mathtt{H})$ and $(\mathtt{H}, \mathtt{L}, \mathtt{H}, \mathtt{L})$. This situation arises due to the dynamic storage of the toggle state on node $\mathtt{y}$. As with the $\phi = \mathtt{L}$ scenario, we need to consider the dynamic behavior of the circuit to justify a discrete model of its behavior.

Viewing ϕ as a parameter, the toggle has four distinct regions of operation summarized below:

$\phi \leq V_{tn}$: The ω-limit set of the system is a continuous curve that connects the points corresponding to the discrete states $(\phi, \mathtt{x}, \mathtt{y}, \mathtt{z}) = (\mathtt{L}, \mathtt{L}, \mathtt{H}, \mathtt{H})$ and $(\mathtt{L}, \mathtt{H}, \mathtt{H}, \mathtt{L})$.

$V_{tn} < \phi < 1.77$ volts: The system has a single fixed point. This fixed point starts at $(\mathtt{x}, \mathtt{y}, \mathtt{z}) = (5, 5, 0)$ for ϕ slightly above V_{tn} and moves to $\sim (5.0, 2.8, 3.0)$ for $\phi = 1.77$ volts.

$1.78 < \phi < V_{dd} + V_{tp}$: The system has a single periodic attractor. The frequency of the oscillation depends on ϕ.

$V_{dd} + V_{tp} \leq \phi$: The ω-limit set of the system includes a continuous curve that connects the points corresponding to the discrete states $(\phi, \mathtt{x}, \mathtt{y}, \mathtt{z}) = (\mathtt{H}, \mathtt{L}, \mathtt{L}, \mathtt{H})$ and $(\mathtt{H}, \mathtt{L}, \mathtt{H}, \mathtt{L})$.

Thus, the discrete states identified in section 2 correspond to elements of the ω-limit set when ϕ is below the n-channel threshold or above the p-channel threshold. In both cases, the two discrete states correspond to regions of a single, connected ω-limit set. When ϕ is not between these bounds, bifurcations occur that lead to ω-limit sets that do not correspond to the desired states of the circuit. The correct operation of the toggle element requires that signals within the toggle change fast enough so that a distinct state is established after each transition of ϕ. On the other

hand, the changes must be slow enough to ensure that the intended discrete state is not lost when the continuous state approaches the "spurious" attractors for intermediate values of ϕ. These two-sided constraints on behavior of signals in the toggle reflects the timing races that are inherent in the design based on a single-phase clocking methodology. To understand the operation of the toggle element, we must consider its operation while ϕ changes.

6 Dynamic analysis

We now consider the behavior of the toggle element when ϕ is considered as a variable. For simplicity, we assume that ϕ is a periodic signal. In particular,

$$\phi \quad = \quad 2.5(1+\sin 2\pi f t)$$

where f is a parameter of the system (the clock frequency). Figure 7 shows the periodic attractor of the toggle element when $f = 100$MHz (compare with figure 2). The salient feature of the attractor is that its period is twice that of ϕ; this is the signature of a toggle element.

We found the attractor by numerically integrating the differential equations describing the system (i.e. by simulation). We let the simulation run until periodic behavior was observed. We then simulated one more traversal of the orbit, checking the eigenvalues of the Jacobian matrix at each step. These eigenvalues give the Lyapunov multipliers of the system that allow us to check for stability as described below. Each entry of the Jacobian matrix gives the partial derivative a component of the trajectory vector, $\dot{\boldsymbol{v}}$, with respect to a component of the state vector, $\boldsymbol{v}$. Intuitively, the Jacobian tells us how the trajectory would change if the state is slightly perturbed. If the eigenvalue corresponding to an eigenvector of the Jacobian is negative, then a perturbation of the state along the direction of this eigenvector will result in a change in the trajectory that tends to return the system to the original path. On the other hand, a positive eigenvalue indicates divergent behavior. Thus, if all eigenvalues are negative everywhere along the orbit, then the orbit is a periodic attractor. By computing the values of the eigenvalues at each simulation step during one traversal of the orbit, we obtain strong experimental evidence that the orbit is a periodic attractor.

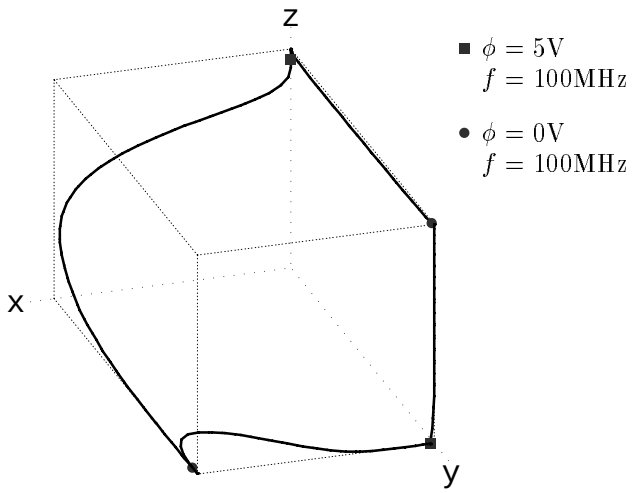

Figure 7: Phase space plot for the toggle element.

When we computed the eigenvalues, we were surprised to discover that at many points along the orbit the Jacobian had eigenvalues equal to zero. We found two causes for these degeneracies. The first is dynamic storage. If a node is floating, there is no tendency to restore its voltage after a perturbation. The second is a consequence of the simple transistor model that we used. When both the n- and p-channel devices driving a node are in saturation, the currents through these transistors are independent of the voltage on the node that they are driving. Let n be such a node. Since we assume fixed capacitances, $\dot{\boldsymbol{v}}_n$ is independent of small changes of $\boldsymbol{v}_n$.

The presence of zero eigenvalues prevents us from concluding stability based purely on the values of the eigenvalues at each point along the orbit. However, our simulations revealed no points along the proposed attractor for which any eigenvalues were positive, and there were regions where all eigenvalues were strictly negative. These two conditions are sufficient to conclude that the observed orbit is indeed a stable, periodic attractor.

When non-linear capacitances are considered, there are sections of the periodic attractor where the Jacobian has positive eigenvalues. Let n be a node, and consider a section of the trajectory where the capaci-

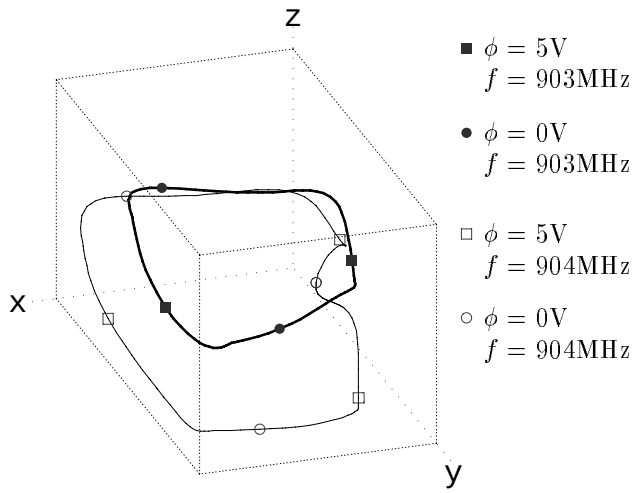

Figure 8: Bifurcation at upper frequency limit.

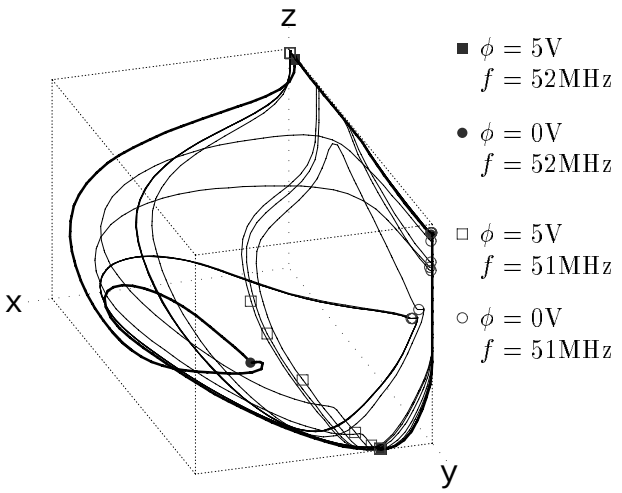

Figure 9: Bifurcation at lower frequency limit.

tance of n is decreasing. A perturbation of v_n in the direction of $\dot{\boldsymbol{v}}_n$ decreases the capacitance of node n. If the drain-source currents flowing into n are roughly unchanged by the perturbation, then $\dot{\boldsymbol{v}}_n$ will increase in magnitude. This implies that the Jacobian has an eigenvector with a positive eigenvalue. Such an orbit may still be stable. In a region where the Jacobian has a positive eigenvalue, an invariant manifold containing the periodic attractor must expand along the axis of the corresponding eigenvector. If for each such expansion there is a compensating contraction elsewhere along the orbit, then the orbit corresponds to a stable, periodic attractor. Verification of such properties requires more powerful techniques than the simple simulations described in this paper (for an example, see [5]).

Having established that the toggle element functions correctly at 100 MHz, we now address the question of the maximum and minimum frequencies of operation for the circuit. When the frequency of ϕ exceeds a critical value, there is a transition in the phase-space structure from a period two to a period three attractor. By simulation, we determined that this transition occurs at a frequency between 903 and 904 MHz. As shown in figure 8, the two attractors are quite distinct, even though the frequency difference is less than 0.2%. As the critical frequency is approached, the asymptotic behavior continues to give a clear indication of the behavior of the device; however, the time required to approach this orbit can grow very large. The attractor for 904 MHz operation also shows the cause of the failure: the transition from $(\phi, \mathtt{x}, \mathtt{y}, \mathtt{z}) = (\mathtt{L}, \mathtt{H}, \mathtt{H}, \mathtt{L})$ to $(\mathtt{H}, \mathtt{L}, \mathtt{L}, \mathtt{H})$ is not completed while ϕ is high. It is not surprising that this is the critical transition as it requires three sequential state changes, whereas each of the others only requires one. An extra clock cycle is expended for this transition, producing an attractor of period three in place of the desired one of period two.

As the frequency of ϕ is decreased, another failure mode emerges. As described in section 5, the attractor structure of the toggle collapses when ϕ is between V_{tn} and $V_{dd} + V_{tp}$, and each region of ϕ is clearly exhibited. In figure 7 the intermediate region is evidenced by the dip in the trajectory on the transition from $(\phi, \mathtt{x}, \mathtt{y}, \mathtt{z}) = (\mathtt{H}, \mathtt{L}, \mathtt{H}, \mathtt{L})$ to $(\mathtt{L}, \mathtt{H}, \mathtt{H}, \mathtt{L})$. For correct operation, ϕ crosses $V_{dd} + V_{tp}$ while the trajectory is still in the basin of attraction corresponding to state $(\mathtt{L}, \mathtt{H}, \mathtt{H}, \mathtt{L})$. However, when the clock frequency is too low, the toggle functions as a free-running oscillator for much of each clock cycle. Again, the transition is very sharp. Figure 9 shows the phase-space plots for the toggle operating at 51 and 52 MHz.

In this section we have shown how discrete behaviors can be extracted from the periodic attractors associated with a dynamical system driven by a periodic

source. In particular, the correct operation of the toggle element is characterized by a period two orbit, and the topological structure of the attractor changes dramatically at the high and low critical frequencies for the design.

We must stress that, although the period two behavior is necessary for the circuit to be considered as a toggle, in most applications, such a simple criterion will not be sufficient. In particular, as the critical frequencies are approached, the attractor is deformed, and the output may not make full excursions and transitions may not be monotonic. We are currently exploring criteria based on invariant manifolds that should allow us to verify the operation of the circuit for a wide range of inputs, not just sinusoidal signals. Furthermore, this should allow us to determine when the output of one circuit element has the required characteristics to be used as the input to another circuit.

7 Conclusions

We have described an experimental analysis of a toggle element based on dynamical systems theory. We showed that a simple interpretation that associates discrete states with fixed-point attractors of the dynamical system does not apply to this circuit because of degeneracies introduced by dynamic storage. We expect that self-timed designs that are based on precharged logic (e.g. William's divider, see [19]) would have similar properties.

We have also shown that when the input, ϕ, is viewed as a parameter, bifurcations occur for intermediate values of ϕ, resulting in a system of attractors that do not correspond to any desired discrete state. These transient phase-space configurations correspond to the race that is inherent in the single-phase clocking scheme of the toggle element. We plan to investigate the phase-space behavior of transistor level implementations of asynchronous circuits to see if similar behaviors are observed.

Because of the limitations of viewing the input as a parameter, we modeled the circuit with the input as a variable. For simplicity, we constrained ϕ to sinusoidal waveforms. With this approach, the system exhibited a periodic attractor that clearly corresponded to the desired discrete behavior. Furthermore, it gave a simple criterion for determining the range of frequencies for which toggle behavior occurs. By this approach, the verification task becomes one of identifying the gross structure of the phase space rather than one of comparing waveforms against some nebulous, intuitive criteria. We believe that this demonstrates a clear advantage of the dynamical systems approach.

To show the stability of proposed attractors, we considered the eigenvalues of the Jacobian matrix at each step of a continuous simulation. We encountered zero valued eigenvalues and noted that this precludes applying a purely local criterion for stability. We are actively exploring techniques to identify invariant manifolds of the phase space. Finding such a manifold is a prerequisite to many other verification tasks including demonstrating robustness under variations of the input signal, demonstrating robustness under process parameter variation, verifying circuits with multiple inputs, and showing that the output of one circuit is a suitable input to another.

Acknowledgments

We are grateful to Charles Molnar for his enthusiasm and suggestions that led to this work. We thank Warren Hunt, Carl Seger, and the anonymous referees for helpful comments. We are especially thankful to Robert Black who carefully read an earlier version of this paper and replicated the results described here using the Matlab™ package [13]. His patience, comments, and questions led to many improvements in this paper.

References

[1] R. D. Black. Towards a dynamical systems approach to asynchronous circuit design. Internal report, University of Waterloo, Apr. 1994.

[2] R. W. Brockett. Smooth dynamical systems which realize arithmetical and logical operations. In H. Nijmeijer and J. M. Schumacher, editors, *Three Decades of Mathematical Systems Theory: A Collection of Surveys at the Occasion of the 50th Birthday of J. C. Willems*, volume 135 of *Lecture Notes in Control and*

Information Sciences, pages 19–30. Springer-Verlag, 1989.

[3] R. Bryant, D. Beatty, et al. COSMOS: a Compiled Simulator for MOS circuits. In *Proceedings of the 27th ACM/IEEE Design Automation Conference*, pages 9–16, 1987.

[4] J. Burch, E. Clarke, et al. Symbolic model checking: 10^{20} states and beyond. In *Proceedings of 1991 International workshop on formal methods in VLSI design*, Jan. 1991.

[5] P. Cahoon. A visual workbench for analyzing the behavior of dynamical systems. In *Extracting meaning from complex data: processing, display, interaction II*, volume 1459, pages 244–253. Society of Photo-Optical Instrumentation Engineers, 1991.

[6] J. de Kleer. How circuits work. In D. G. Bobrow, editor, *Qualitative Reasoning About Physical Systems*, pages 205–280. MIT Press, Cambridge, MA, 1985.

[7] L. A. Glasser and D. W. Dobberpuhl. *The Design and Analysis of VLSI Circuits*. Addison-Wesley, 1985.

[8] M. W. Hirsch and S. Smale. *Differential Equations, Dynamical Systems, and Linear Algebra*. Academic Press, San Diego, CA, 1974.

[9] P. Horowitz and P. Winfield. *The Art of Electronics*. Cambridge University Press, second edition, 1989.

[10] M. Hurtado. *Structure and Performance of Asymptotically Bistable Dynamical Systems*. PhD thesis, Sever Institute, Washington University, Saint Louis, MO, 1975.

[11] R. Kurshan and K. McMillan. Analysis of digital circuits through symbolic reduction. *IEEE Transactions on Computer-Aided Design*, 10(11):1356–1371, Nov. 1991.

[12] A. J. Martin. Programming in VLSI: From communicating processes to delay insensitive circuits. In C. Hoare, editor, *Developments in Concurrency and Communication*, University of Texas Year of Programming Series. Addison-Wesley, 1990.

[13] The MathWorks, Natick, MA. *MATLAB User's Guide*, 1992.

[14] W. H. Press et al. *Numerical Recipes in C: The art of numerical computing*. Cambridge University Press, 1988.

[15] E. Sacks and J. Doyle. Prolegomena to any future qualitative physics. *Computational Intelligence*, 8(2):187–209, 1992.

[16] C. L. Seitz. System timing. In *Introduction to VLSI Systems*, chapter (Carver Mead and Lynn Conway) 7, pages 218–262. Addison Wesley, 1979.

[17] V. Stavridou, T. Melham, and R. Boute, editors. *Theorem Provers in Circuit Design*. North-Holland, 1992.

[18] B. C. Williams. Qualitative analysis of MOS circuits. In D. G. Bobrow, editor, *Qualitative Reasoning About Physical Systems*, pages 281–346. MIT Press, Cambridge, MA, 1985.

[19] T. E. Williams. *Self-Timed Rings and their Application to Division*. PhD thesis, Stanford University, May 1991.

[20] J. Yuan and C. Svensson. High-speed CMOS circuit technique. *IEEE Journal of Solid-State Circuits*, 24(1):62–70, Feb. 1989.