

Oscillator Verification with Probability One

Chao Yan
Intel

Mark Greenstreet
University of British Columbia

Abstract—This paper presents the formal verification of start-up for a differential ring-oscillator circuit used in industrial designs. Dynamical systems theory shows that any oscillator must have a non-empty failure; however, it is possible to show that these failures only occur with zero probability. To do so, this paper generalizes the “cone argument” initially presented in [1] and proves the soundness of this generalization. This paper also shows how concepts from analog design such as differential operation can be soundly incorporated into the verification to produce simpler models and reduce the complexity of the verification task.

I. INTRODUCTION

System-on-Chip (SoC) and analog-mixed-signal (AMS) designs have created new challenges for analog circuit designers. Typical analog design relies heavily upon simulation tools such as HSPICE and Spectre. Long simulation times along with the continuous nature of device parameters, operating conditions and input waveforms mean that simulation tools can only provide partial verification of analog designs. In practice, designers typically focus their simulation efforts on parametric and small-signal sensitivity analysis when the circuit is in or near its intended operating mode. Such analysis can be used to determine the gain and bandwidth of an amplifier, the jitter transfer function of a phase-locked loop, along with finding transistor sizes to optimize a given circuit topology for an objective function formulated in terms of steady-state properties of the circuit. However, simulations cannot show that the circuit will eventually reach its intended operating condition from all possible starting conditions.

This paper presents a rigorous, formal verification that a commonly used differential ring-oscillator circuit correctly starts oscillation with probability 1. As shown in Figure 1, the oscillator consists of two stages, where each stage has a pair of “forward” inverters (labeled fwd in the figure) and a pair of “cross-coupling” inverters (labeled cc). If the forward inverters are much larger than the cross-coupling inverters, then the circuit acts like a ring of four inverters settles to one of two states:

- State 1: X_1 and X_3 , are low; and X_2 and X_4 are high.
 - State 2: X_1 and X_3 , are high; and X_2 and X_4 are low.
- (1)

Conversely, if the cross-coupling inverters are much larger than the forward ones, then the circuit acts like two separate static latches and has four stable states. If the forward and cross-coupling inverters have comparable strength, then the circuit should oscillate in a stable fashion.

In 2008, researcher from Rambus posed the problem of showing that the oscillator circuit shown in Figure 1 starts from all initial conditions for a particular choice of transistor




Fig. 1. Ring-Oscillator Example from Rambus

sizes [2]. They described this as a “real-world” problem noting that oscillators of this type had been observed to fail in the test-lab. They posed a further problem of determining the range of transistor sizes for which proper start-up is guaranteed. This paper presents solutions to these problems.

A. Prior Work

Oscillator circuits have been a popular example for applying formal methods to analog circuit verification [3]–[6]. These early papers focused on simple oscillators, such as a tunnel-diode based design, that are not representative of the oscillator circuits used in real VLSI designs. More recently, several groups have reported results for the Rambus oscillator problem described above.

The earliest attempted verification of the oscillator that we have seen [7] predates the formulation by [2] and considers the behaviours of a 128 stage oscillator for a pulse-width modulated voltage regulator. Their “proof” of correct operation assumes differential and periodic operation, and does not consider weak coupling between stages (e.g. due to power supply noise), that could stabilize undesired, higher harmonic modes of oscillation.

A more rigorous approach was taken in [8] which used monotonicity properties of the i_{ds} functions of MOSFETs to reduce the search for DC-equilibria in a Rambus ring oscillator with an arbitrary, even number of stages to a one-dimensional search, regardless of the number of stages in the oscillator. They then used standard, small-signal analysis techniques to determine if any of these equilibria are stable. If an oscillator circuit had no stable DC equilibria, it was deemed free from DC lock-up. The authors noted that their proof did not rule out other behaviours such as higher harmonic oscillations or chaotic behaviours.

Several subsequent papers have also treated the verification problem as one of ruling out the existence of DC equilibria. For example, Tiwari *et al.* [9], [10] used a SAT solver to identify DC equilibria. To find *stable* equilibria, they added

constraints that at least one node of the circuit must be within 0.2 volts of power or ground. They did not state how they had arrived at these extra constraints or whether or not they can be shown to be sound. Steinhorst *et al.* [11] presented a particle filtering approach and compared it with a model-checking method. The correctness condition for the model-checking was lack of stable DC-equilibria. While higher-harmonic oscillations or chaotic behaviours were not considered, they presumably would show up in the visualizations if suitable particles were included in the state-space sampling. Zaki *et al.* [12] presented an approach where the “pencil-and-paper” analysis from [8] was automated using HySAT [13] and Matlab toolboxes for interval arithmetic and matrix pseudo-spectrum calculations.

Little *et al.* [14] showed that trajectories in a neighbourhood of the nominal periodic trajectory for the oscillator remain close to that nominal trajectory. This replaces the small-signal analysis of traditional analog design with linear hybrid Petri net (LHPN) model checking and confirms the stability of the desired oscillating behaviour. As the analysis only considers a portion on the state space near the desired trajectory, it does not verify proper start-up for all initial conditions.

B. Contributions

This paper combines analytical techniques based on dynamical systems theory with reachability tools to present the first verification of the Rambus oscillator problem that actually addresses the question posed by Jones *et al.*: “Will the oscillator start up from all initial conditions?” In Section II we consider the dynamics of *any* oscillator that is modeled by non-linear differential equations and show that it must have some set of initial conditions for which the circuit fails to oscillate. However, this failure set can be *negligible*, i.e. have zero probability. We present a generalization of the cone argument from [1] to verify that the failure set has zero probability, and thus that the oscillator starts with a probability of one. We also introduce a symmetry reduction method that allows us to exploit the differential operation of the oscillator in a formal verification context. Section IV describes our implementation of the verification method using Matlab and Coho [15]. Section V presents the results of verifying the oscillator circuit with these methods.

II. NO PERFECT OSCILLATOR

This section shows that no physically plausible oscillator starts from all initial conditions.

A. Dynamical Systems and Oscillators

We assume that a circuit, such as an oscillator, is modeled by a system of ordinary differential equations. If the model has d variables, states of the circuit correspond to points in \mathbb{R}^d . The model includes a function, $f: \mathbb{R}^d \rightarrow \mathbb{R}^d$ that is the *time derivative* of the system: for state $\mathbf{x} \in \mathbb{R}^d$, $\dot{\mathbf{x}} = f(\mathbf{x})$ is the time-derivative of the system in state \mathbf{x} . For $\mathbf{x}_0 \in \mathbb{R}^d$, the *initial value problem* is to find a function $\mathbf{x}: \mathbb{R}^+ \rightarrow \mathbb{R}^d$ such that for all $t \geq 0$, $\frac{d}{dt}\mathbf{x}(t) = f(\mathbf{x}(t))$ and $\mathbf{x}(0) = \mathbf{x}_0$. Let $Q \subseteq \mathbb{R}^d$ be a

closed set. Q is invariant with respect to f if all trajectories that start in Q remain in Q forever. To show that Q is invariant with respect to f , it is sufficient to show that for every $\mathbf{x} \in Q$ there is an $\varepsilon > 0$ such that $\mathbf{x} + \varepsilon f(\mathbf{x}) \in Q$. We impose two restrictions on f :

- R1:** There is a $Q \subseteq \mathbb{R}^d$ and some $K \in \mathbb{R}$ such that Q is invariant with respect to f and for every $\mathbf{x} \in Q$, $\|f(\mathbf{x})\| < K$.
- R2:** f is \mathcal{C}^1 in Q . This means that $f(\mathbf{x})$ is differentiable with respect to the components of \mathbf{x} , and these derivatives are continuous.

These two conditions guarantee the existence and uniqueness of solutions to the initial value problem for f and any $\mathbf{x}_0 \in Q$ (see [16, chap. 8.3]). We can define a function $\Phi_f(\mathbf{x}_0, t)$ such that if \mathbf{x} is the solution to the initial value problem for f with $\mathbf{x}(0) = \mathbf{x}_0$, then $\mathbf{x}(t) = \Phi_f(\mathbf{x}_0, t)$. Given restrictions R1 and R2, $\Phi_f(\mathbf{x}_0, t)$ is a \mathcal{C}^1 function with respect to \mathbf{x}_0 and t for any $\mathbf{x}_0 \in Q$ and $t \geq 0$ (see [16, chap. 8.4]). We extend Φ_f to sets in the natural way: if $X \subseteq \mathbb{R}^d$, then $\Phi_f(X, t) = \{\mathbf{x}_2 | \exists \mathbf{x}_1 \in X. \mathbf{x}_2 = \Phi_f(\mathbf{x}_1, t)\}$.

We assume that any physically plausible oscillator can be modeled by an ODE with f and Q satisfying restrictions R1 and R2. The requirement that f is \mathcal{C}^1 follows from the smoothness of the underlying physical models for electric fields, charge distributions, etc. The requirement of the existence of the set Q is satisfied because VLSI circuits generally have node voltages that are bounded by the voltages of ground and the power supply or that have limited excursions beyond these power supply voltages.

We now define “oscillation.” If there is a $\mathbf{x}_0 \in Q$ and a $P > 0$ such that $\Phi_f(\mathbf{x}_0, P) = \mathbf{x}_0$, and for all $0 < t < P$, $\mathbf{x}(t) \neq \mathbf{x}_0$, then f has a solution with period P . In this case, we write $\Gamma_{f, \mathbf{x}_0} = \{\mathbf{x} | \exists t \in [0, P]. \mathbf{x} = \Phi_f(\mathbf{x}_0, t)\}$ to denote the set of points in this periodic orbit. It is straightforward to show $\forall t > 0. \Phi_f(\Gamma, t) = \Gamma$. Let $J = \text{Jac} \Phi_f(\mathbf{x}_0, P)$, i.e., J is the matrix of partial derivatives of $\Phi_f(\mathbf{x}_0, P)$ with respect to \mathbf{x}_0 . If J has $d - 1$ eigenvalues with magnitude less than 1, then the periodic solution for \mathbf{x}_0 is a *periodic attractor* [16, Theorem 13.2]. We say that a system is an oscillator with period P if it has a periodic attractor with period P .

B. Oscillator Start-Up

First consider the set of possible initial states. Labeling one terminal of the power supply as “ground” and the other as “ V_{dd} ” is simply a designer convention. Depending on circuit details, the node voltages on power-up may be arbitrary values. Rather than trying to analyse the circuit in detail, we simply assume that each node has an arbitrary initial voltage in $[V_{lo}, V_{hi}]$; typically V_{lo} is ground or close to ground, and V_{hi} is close to V_{dd} . Let $X_0 = [V_{lo}, V_{hi}]^d$ denote the set of initial node voltages. Because X_0 contains all reachable states of the circuit, we assume $\gamma \subseteq X_0 \subseteq Q$, where γ is the desired periodic attractor of the oscillator.

We can now describe an ideal oscillator.

A d -dimensional dynamical system with time-derivative function f is an *ideal oscillator* iff

The system is physically plausible: There is a set $Q \subseteq \mathbb{R}^d$ such that f and Q satisfy conditions R1 and R2.

Periodic behavior: The system has a periodic attractor. Let Γ be the orbit associated with this attractor.

Start up: There is a convex set $X_0 \subseteq \mathbb{R}^d$ of initial states such that $\Gamma \subseteq X_0 \subseteq Q$ and for every point $\mathbf{x}_0 \in X_0$ and every $\varepsilon > 0$, there is a $t > 0$ and a point $\mathbf{x}_1 \in \Gamma$ such that $\|\mathbf{x}_1 - \Phi(\mathbf{x}_0, t)\| < \varepsilon$.

The first two conditions were described in the previous section. The last condition states that the set of initial states must contain the periodic orbit as described above, and that for any initial state, the trajectories emanating from that state must eventually be arbitrarily close to the periodic orbit. The requirement that this initial set be convex reflects the topological properties of sets such as $[V_{lo}, V_{hi}]^d$ described above. We believe that this definition of an ideal oscillator captures the notion of the oscillator starting from all initial conditions requested in [2].

Theorem 1. *There is no ideal oscillator.*

Proof: This follows directly from the property that solutions of ODEs that satisfy properties R1 and R2 are continuous in their initial conditions. Thus, the topology of the initial set, X_0 , is preserved by $\Phi_f(X_0, t)$. However, any small neighborhood of a periodic attractor must have genus 1 (be “torus-like”) whereas the set of initial states has genus 0 (i.e. it is “sphere-like”). Thus, it is not that case that all initial conditions lead to trajectories that are arbitrarily close to the desired attractor. This establishes the claim. ■

III. VERIFICATION OUTLINE

Our verification proceeds in three main phases:

Differential Operation The oscillator shown in Figure 1 is a differential design: nodes X1 and X3 form a “differential pair” and likewise for nodes X2 and X4. The first phase of the verification shows that each of these differential pairs can be treated as a single signal.

Escape from the Failure Set As shown in Section II, for any oscillator, there must be initial conditions from which it does not properly start. The second phase of the verification shows that this occurs with probability zero.

Proper Oscillation The first two phases show that most initial conditions lead to a fairly small subset of the full state space. In the final phase, we use existing reachability methods to show that the oscillator starts up properly from the region.

This section describes the dynamical systems issues associated with each of these phases. Section IV describes our verification method based on these observations.

We model the oscillator circuit from Figure 1 using non-linear ordinary differential equations (ODEs) of the form:

$$\dot{\mathbf{x}} = f(\mathbf{x}) \quad (2)$$

where \mathbf{x} is a vector of node voltages, $\dot{\mathbf{x}}$ is the vector of time derivatives for these voltages, and f is the function

modeling the non-linear dynamics of this circuit. Let d be the dimensionality of \mathbf{x} . We assume that f is C^1 which guarantees that Equation 2 has a unique solution for any initial state, $\mathbf{x}(0)$. For simplicity, we model the system as being autonomous (no inputs or outputs). Inputs (e.g. to model VCO control inputs, power supply noise), can be modeled by giving f additional parameters, i.e. $f(\mathbf{x}, \mathbf{in})$.

A. Differential Behaviour

Nodes X1 and X3 in the oscillator from Figure 1 form a “differential pair” and likewise for nodes X2 and X4. Let x_i denote the voltage on node X_i . The *differential component* of the differential pair is $x_1 - x_3$, and $x_1 + x_3$ is the *common mode component*. When the oscillator is operating properly, the common mode components are roughly constant and the oscillation is manifested in the differential components. Let V_0^+ be the nominal value for the common mode components. We show that for a relatively small V_{err} if $|x_1 + x_3 - V_0^+| > V_{err}$, then $\frac{d}{dt}(x_1 + x_3)$ and $(x_1 + x_3 - V_0^+)$ have opposite signs. This shows that that the common mode component for nodes X1 and X3 converges to within V_{err} of the nominal value. Likewise for nodes X2 and X4.

B. Escape from the Failure Set

Theorem 1 shows that there is no perfect oscillator. For the Rambus ring-oscillator, there is an equilibrium point, \mathbf{x}_{fail} , i.e. a point where $\dot{\mathbf{x}} = 0$, and there is a manifold, X_{fail} such that

$$\forall \mathbf{x} \in X_{fail}. \lim_{t \rightarrow \infty} \|\Phi_f(\mathbf{x}, t) - \mathbf{x}_{fail}\| = 0.$$

Thus, direct application of continuous state-space model checkers (e.g. [3], [17]) to the oscillator start-up problem will identify regions where trajectories might stay forever. Because we cannot show that the set of failure states is empty, we must settle for showing that it is *negligible* (i.e. occurs with probability zero). This is sufficient in practice, as designers are not worried about a design that fails with probability zero.

For intuition, consider an oscillator where all inverters are identical. We define V_{eq} as the voltage that can be applied to the input of the inverter such that the output settles to the same voltage. When all of the inverters are identical, \mathbf{x}_{fail} is the point at which all node voltages are V_{eq} . Furthermore, any trajectory starting at a point where $x_1 = x_3$ and $x_2 = x_4$ converges to \mathbf{x}_{fail} ; thus, such points are in X_{fail} .

Using existing reachability methods, we can find a small region, U_{fail} , that contains the point \mathbf{x}_{fail} . Furthermore, we can show that if an oscillator starts any point where each node has a voltage in the interval $[0, V_{dd}]$, then within bounded time, the oscillator state will either be in U_{fail} , or it will be in a region where we can show convergence to the desired periodic orbit.

We will show that the set of failing trajectories is sufficiently small as to ensure that the oscillator fails to start with a probability of zero. As in the previous section, we write \mathbb{R}^d to denote the phase space. We will avoid a detailed treatment of measure theory (see [18]) by noting that when we say that $B \subseteq \mathbb{R}^d$ is measurable, we mean that it has a well-defined d -dimensional “volume” (i.e. it is Lebesgue measurable),

and we write $|B|$ to denote this volume (i.e. measure). We write $\mu(B)$ to denote the probability that the initial state of the oscillator is in B . Our assumption that μ is smooth (i.e. absolutely continuous) means that if $|B|$ is zero, then $\mu(B)$ is zero as well. For example, let

$$B = \{(x_1, x_2, x_3, x_4) \mid (x_1 = x_3) \wedge (x_2 = x_4)\}$$

i.e. the plane described above. Because this plane has zero volume, $|B| = 0$, and by our smoothness assumption, $\mu(B) = 0$ as well.

Let U be a bounded, measurable subset of \mathbb{R}^d . We define

$$\begin{aligned} \text{escape}_f(\mathbf{x}, U) &= \exists t \in \mathbb{R}^+. \Phi_f(\mathbf{x}, t) \notin U \\ \text{trapped}_f(U) &= \{\mathbf{x} \in U \mid \neg \text{escape}_f(\mathbf{x}, U)\} \end{aligned}$$

For any $U \subseteq \mathbb{R}^+$, and any $t \in \mathbb{R}$, $|U| = 0 \Leftrightarrow |\Phi_f(U, t)| = 0$. Thus, it suffices to show that $|\text{trapped}_f(U_{\text{fail}})| = 0$. The next theorem presents conditions that ensure $\mu(\text{trapped}_f(U)) = 0$.

Theorem 2. *Let μ be a smooth probability measure over \mathbb{R}^d . Let U be a bounded, measurable subset of \mathbb{R}^d , and $f: \mathbb{R}^d \rightarrow \mathbb{R}^d$ be bounded and C^1 in U . If there is a matrix $H \in \mathbb{R}^{d \times d}$ such that at least one eigenvalue of H has a positive real part, and $k > 0$ such that for all $\mathbf{x}_1, \mathbf{x}_2 \in U$:*

$$\begin{aligned} (\mathbf{x}_2 - \mathbf{x}_1)^T H (\mathbf{x}_2 - \mathbf{x}_1) &> 0 \\ \Rightarrow (\mathbf{x}_2 - \mathbf{x}_1)^T H (f(\mathbf{x}_2) - f(\mathbf{x}_1)) &> k(\mathbf{x}_2 - \mathbf{x}_1)^T H (\mathbf{x}_2 - \mathbf{x}_1), \end{aligned}$$

then $\mu(\text{trapped}_f(U)) = 0$.

Proof: Assume that $\text{trapped}_f(U) \neq \emptyset$ as the other case is trivial. Let ρ_{\max} be the maximum real part of any eigenvector of H . Let \mathbf{u} be a unit vector such that $\mathbf{u}^T H \mathbf{u} = \rho_{\max}$. Let \mathbf{x}_0 be any point in $\text{trapped}_f(U)$, and $\alpha \in \mathbb{R}$ such that $\alpha > 0$ and $\mathbf{x}_0 + \alpha \mathbf{u} \in U$. We'll define $\mathbf{x}_1 = \mathbf{x}_0 + \alpha \mathbf{u}$.

We now show $\mathbf{x}_1 \notin \text{trapped}_f(U)$. Consider two trajectories,

$$\begin{aligned} \eta_0(t) &= \Phi_f(\mathbf{x}_0, t), \quad \text{the trajectory that starts at } \mathbf{x}_0 \\ \eta_1(t) &= \Phi_f(\mathbf{x}_1, t), \quad \text{the trajectory that starts at } \mathbf{x}_1 \end{aligned}$$

Note that both trajectories start in U . We'll show that these two trajectories diverge, and thus at most one of them can remain in U . Let

$$w(t) = (\eta_1(t) - \eta_0(t))^T H (\eta_1(t) - \eta_0(t))$$

We claim that for $t \geq 0$, $w(t) \geq \alpha^2 \rho_{\max} e^{kt} > 0$. First note that $w(0) = \alpha^2 \rho_{\max}$ which satisfies the claim (at $t = 0$). Both $w(t)$ and $\alpha^2 \rho_{\max} e^{kt}$ are continuous functions of t . Thus, if the claim were ever to be violated, there would have to be a value of t for which $w(t) = \alpha^2 \rho_{\max} e^{kt}$ and $\frac{d}{dt} w(t) < \frac{d}{dt} \alpha^2 \rho_{\max} e^{kt}$. For the sake of contradiction, let t be such a time. Then

$$\begin{aligned} \frac{d}{dt} w(t) &= (\eta_1(t) - \eta_0(t))^T H (f(\eta_1(t)) - f(\eta_0(t))) \\ &> k(\eta_1(t) - \eta_0(t))^T H (\eta_1(t) - \eta_0(t)) \\ &= kw = k\alpha^2 \rho_{\max} e^{kt} = \frac{d}{dt} \alpha^2 \rho_{\max} e^{kt} \end{aligned}$$

But this shows that $\frac{d}{dt} w(t) > \frac{d}{dt} \alpha^2 \rho_{\max} e^{kt}$, a contradiction. Thus, $w(t) \geq \alpha^2 \rho_{\max} e^{kt}$ as claimed.

Because, $w(t) \geq \alpha^2 \rho_{\max} e^{kt}$, $\|\eta_1(t) - \eta_0(t)\|$ must diverge as $t \rightarrow \infty$. By assumption, $\eta_0(t)$ stays in U , and U is bounded. Therefore, $\eta_1(t)$ must exit U .

We have shown that for any point $\mathbf{x}_0 \in \text{trapped}_f(U)$, all points in the cone defined by H whose apex is at \mathbf{x}_0 must escape from U . This shows that $\text{trapped}_f(U)$ must have lower dimension than the full space. Thus, $|\text{trapped}_f(U)| = 0$, and therefore $\mu(\text{trapped}_f(U)) = 0$ as claimed. ■

Note: Theorem 2 was based on the cone argument from [1]. The present theorem generalizes the result from [1] to systems of arbitrary dimensions and whose Jacobian matrices have complex eigenvalues. The conditions for Theorem 2 are slightly stronger than those from [1] (for the systems where the latter applies) – this is mainly for simplicity.

C. Proper Oscillation

For the trajectories under consideration after the first two steps, the common mode components of both differential signal pairs are within V_{err} of V_0^+ . This allows the differential equation model from Equation 2 to be rewritten as a *differential inclusion* [19]:

$$\dot{\mathbf{u}} \in F(\mathbf{u}) \quad (3)$$

where \mathbf{u} is the vector $(\sqrt{2}/2)[x_1 - x_3, x_2 - x_4]$. By using an inclusion, F accounts for all values of the common mode components in $[V_0^+ - V_{\text{err}}, V_0^+ + V_{\text{err}}]$. Reducing the four-dimensional state space of the original problem to a two-dimensional space makes the exploration of trajectories from all remaining start conditions straightforward.

By showing that all such trajectories lead to an oscillation in the fundamental mode, we solve the first part of the challenge problem from [2]: we show that for a particular choice of transistor sizes, the circuit will start oscillation from all initial conditions except for a set of zero measure. Section V provides a brief description of how these methods can be extended to establish a range of transistor sizes for which the oscillator will start with probability one.

IV. IMPLEMENTATION

This section describes our implementation of the verification techniques described in the previous section. We construct an ODE model for the ring oscillator circuit using standard, modified nodal analysis. We obtain drain-to-source current data by tabulating HSPICE outputs and fitting piece-wise quadratic functions to this tabulated data. The resulting errors are less than 1%; thus, our transistor models closely match those used by practicing circuit designers in industry.

A. Differential Operation

This verification phase starts by changing the coordinate system to one based on the differential and common mode representation of signals.

Let \mathbf{u} be the circuit state in “differential” coordinates:

$$\begin{aligned} \mathbf{u} &= M^{-1} \mathbf{x} \\ M &= \frac{\sqrt{2}}{2} \begin{bmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ -1 & 0 & 1 & 0 \\ 0 & -1 & 0 & 1 \end{bmatrix} \end{aligned} \quad (4)$$

We assume each of nodes X1, X2, X3 and X4 can independently have initial voltages anywhere in $[0, 1.8]V$. Thus, the differential components, u_1 and u_2 , are initially in $[-0.9\sqrt{2}, +0.9\sqrt{2}]$, and the common mode components, u_3 and u_4 , are initially in $[0, 1.8\sqrt{2}]$.

To establish differential operation, we divide the range of each component u_i of \mathbf{u} into m intervals, creating m^4 cubes. We construct a graph, $G = (V, E)$ to represent the reachability relationship between these cubes. Let $v_{i,j,k,\ell}$ be a vertex corresponding to the i^{th} interval for u_1 , the j^{th} interval for u_2 and so on. There is an edge from v to w if f allows a flow out of the cube for v directly into the cube for w , and there is a self-loop for v if each component of f is zero somewhere in v . If a vertex of G has no incoming edges, then any trajectory that starts in the corresponding cube will eventually leave that cube, and no trajectories will ever enter the cube. Such a cube can be eliminated from further consideration. Thus, we only need to consider cubes whose vertices are members of cycles. These vertices can be identified in $O(V + E) = O(m^4)$ time. With a direct implementation of this computation, constructing G dominates the entire time for verifying the oscillator.

To obtain a more efficient computation, we first note that the goal is to establish differential operation. It is sufficient to project the vertices of V onto the common-mode components of the differential signals and show that most of this projection can be eliminated from further consideration. Let $G' = (V', E')$ where $v'_{k,\ell}$ corresponds to the k^{th} interval of u_3 and the ℓ^{th} interval of u_4 . There is an edge in E' from v'_{k_1,ℓ_1} to v'_{k_2,ℓ_2} iff there exist i and j such that $(v_{i,j,k_1,\ell_1}, v_{i,j,k_2,\ell_2}) \in E$. Clearly, G' over approximates reachability. Thus, if a vertex of G' has no incoming edges, then all of the corresponding vertices in G must have no incoming edges as well. Computing the edges in E' requires examining all of the edges of E , but subsequent operations on the graph G' are much faster than those on G .

To reduce the time required to find edges of E , we start with a small value of m and thus a coarse grid. Many large blocks can be eliminated from G' even with a coarse grid. We then double m (i.e. divide each vertex of G' into four) and recompute reachability using the finer grid for finding edges in E as well. In practice this adaptive gridding approach eliminates blocks quickly while achieving enough precision to allow the rest of the verification to proceed without difficulties.

B. Escape from the Failure Set

At the end of establishing differential operation, there are a few cubes with self-loops – there is more than one such cube because of the over approximations described above. These cubes contain the point \mathbf{x}_{fail} . We now construct a larger cube that contains all of these and make a change of variables so that this cube is centered at the origin. We'll write \mathbf{x} for vectors in the original coordinate system and \mathbf{u} for vectors in the coordinates where the center of a cube with a self-loop is at the origin. Let r be the maximum ℓ_2 distance of any point in this cube from the origin.

As described at the beginning of this section, we use piecewise quadratic models for transistor currents and model

node capacitances as constants. Thus, the derivative function, f , is piecewise quadratic. Our repeated subdivision of cubes when establishing differential operation ensures that the cube containing \mathbf{x}_{fail} is modeled by a simple quadratic (i.e. a single ‘‘piece’’). We can write this model as:

$$\dot{\mathbf{u}} = A_0 + A_1\mathbf{u} + \sum_{j=1}^d (\mathbf{u}^T A_{2,j} \mathbf{u}) \mathbf{b}_j \quad (5)$$

where \mathbf{b}_j is a unit vector corresponding to the j^{th} component of \mathbf{u} . We will assume wlog that the $A_{2,j}$ matrices are symmetric throughout paper.

To establish the hypotheses of Theorem 2, we again exploit the differential operation of the oscillator and choose $H = \text{diag}([+1, +1, -1, -1])$. The two $+1$ elements of H anticipate a growing, differential component of the state, and the two -1 elements are for a diminishing common-mode component. Consider $(\mathbf{u}_2 - \mathbf{u}_1)^T H (f(\mathbf{u}_2) - f(\mathbf{u}_1))$:

$$\begin{aligned} & (\mathbf{u}_2 - \mathbf{u}_1)^T H (f(\mathbf{u}_2) - f(\mathbf{u}_1)) \\ &= (\mathbf{u}_2 - \mathbf{u}_1)^T H A_1 (\mathbf{u}_2 - \mathbf{u}_1) \\ & \quad + (\mathbf{u}_2 - \mathbf{u}_1)^T H \sum_{j=1}^d ((\mathbf{u}_2 - \mathbf{u}_1)^T A_{2,j} (\mathbf{u}_2 + \mathbf{u}_1)) \mathbf{b}_j \end{aligned} \quad (6)$$

We now derive a lower bound for

$$\frac{(\mathbf{u}_2 - \mathbf{u}_1)^T H A_1 (\mathbf{u}_2 - \mathbf{u}_1)}{(\mathbf{u}_2 - \mathbf{u}_1)^T H (\mathbf{u}_2 - \mathbf{u}_1)} \quad (7)$$

and an upper bound for

$$\left| \frac{(\mathbf{u}_2 - \mathbf{u}_1)^T H \sum_{j=1}^d ((\mathbf{u}_2 - \mathbf{u}_1)^T A_{2,j} (\mathbf{u}_2 + \mathbf{u}_1)) \mathbf{b}_j}{(\mathbf{u}_2 - \mathbf{u}_1)^T H (\mathbf{u}_2 - \mathbf{u}_1)} \right| \quad (8)$$

when $(\mathbf{u}_2 - \mathbf{u}_1)^T H (\mathbf{u}_2 - \mathbf{u}_1) > 0$.

Equation 7 is a convex conic program and can be solved by standard techniques (see [20, chap. 4.4]); let lin_{\min} be the minimum value for Equation 7. To bound the magnitude of the quadratic term, let σ_{\max} denote the largest singular value of any of the $A_{2,j}$ matrices. Then, for all $j \in 1 \dots d$,

$$(\mathbf{u}_2 - \mathbf{u}_1)^T A_{2,j} (\mathbf{u}_2 + \mathbf{u}_1) \leq \sigma_{\max} (\mathbf{u}_2 - \mathbf{u}_1)^T (\mathbf{u}_2 + \mathbf{u}_1)$$

Therefore,

$$\begin{aligned} & \left\| \sum_{j=1}^d ((\mathbf{u}_2 - \mathbf{u}_1)^T A_{2,j} (\mathbf{u}_2 + \mathbf{u}_1)) \mathbf{b}_j \right\| \\ & \leq \sqrt{d} \sigma_{\max} (\mathbf{u}_2 - \mathbf{u}_1)^T (\mathbf{u}_2 + \mathbf{u}_1) \end{aligned}$$

Noting that the largest singular value of H is 1, and $\|\mathbf{u}_2 + \mathbf{u}_1\| \leq 2r$, we get:

$$\begin{aligned} & (\mathbf{u}_2 - \mathbf{u}_1)^T H \sum_{j=1}^d ((\mathbf{u}_2 - \mathbf{u}_1)^T A_{2,j} (\mathbf{u}_2 + \mathbf{u}_1)) \mathbf{b}_j \\ & \leq 2r\sqrt{d}\sigma_{\max} \|\mathbf{u}_2 - \mathbf{u}_1\|^2 \end{aligned} \quad (9)$$

By our choice of H ,

$$(\mathbf{u}_2 - \mathbf{u}_1)^T H (\mathbf{u}_2 - \mathbf{u}_1) \leq \|\mathbf{u}_2 - \mathbf{u}_1\|^2 \quad (10)$$

Now, let $k = \text{lin}_{\min} - 2r\sqrt{d}\sigma_{\max}$. Combining the results from Equations 6 through 10, we get

$$(\mathbf{u}_2 - \mathbf{u}_1)^T H (f(\mathbf{u}_2) - f(\mathbf{u}_1)) \geq k (\mathbf{u}_2 - \mathbf{u}_1)^T H (\mathbf{u}_2 - \mathbf{u}_1)$$

If $k > 0$, then we can satisfy the conditions of Theorem 2. In practice, the conditions of Theorem 2 can be satisfied by choosing r to be sufficiently small.

C. Proper Oscillation

As described in Section III-C, we reduce the state space from four dimensions to two by replacing the differential equation model for the circuit with a differential inclusion. The space to be considered forms a ring: the outer boundary is determined by the assumption that all signals have voltages between ground and V_{dd} , and the inner boundary is established by eliminating trajectories in a neighborhood near x_{fail} . Figure 3 shows the remaining region. We use a collection of “spokes” as shown in Figure 4, and show that all trajectories in these wedges converge to a unique, periodic attractor. The computation has three parts:

- 1) Starting from each “spoke”, show that all trajectories starting at that spoke eventually cross the next spoke.
- 2) Show that all trajectories starting from the inner or outer boundary eventually cross the next spoke.
- 3) Starting from one spoke, compute the reachable set until it converges to a limit set.

V. RESULTS

We generated transistor models using HSPICE to determine drain-to-source currents for 0.18μ long and 1μ wide nMOS and pMOS devices with the gate and drain voltages swept from 0 to 1.8V in 0.01V steps. For the nMOS transistors, we assume that the source and body are at 0V, and for the pMOS devices, we assume that they are at 1.8V. We assume that all transistors have a length of 0.18μ , and obtain current for other widths by linear scaling from the 1μ data. For all inverters, we use pMOS devices that are twice as wide as the nMOS devices. All forward inverters have transistors of the same size, and likewise for the cross-coupled inverters. In the following, s denotes ratio of the cross-coupled inverter size to the forward inverter size. This section first presents the verification of an oscillator with $s = 1$. Then, the oscillator is verified for $0.673 \leq s \leq 2.0$.

The verification routines were implemented using Matlab with Coho used for the final reachability computation. All times were obtained running on a dual Xeon E5520 (quad core) 2.27GHz machine with 32GB of memory. The computations described here are all performed using a single core.

A. Verification with equal-size inverters

The first phase of the verification establishes differential operation. Initially, the computation partitions the space for each of the u_i variables into 8 regions, creating a total of $8^4 = 4096$ cubes to explore. After eliminating cubes that have no incoming or self-circulating flows, the remaining cubes are subdivided and rechecked until there are 64 intervals for each variable. Figure 2 shows the remaining cubes projected onto the common-mode variables, u_3 and u_4 at the end of this phase.

With 8 intervals per region, there are 752 cubes under consideration (18% of the total space). With each subdivision, the number of cubes remaining increases by a factor of roughly 4.6, and thus the volume of the space under consideration drops by about a factor of roughly 0.29. With 64 intervals per region, 74676 cubes remain (0.45% of the total space).




Fig. 2. Common-mode convergence to $V_{dd}\sqrt{2}/2$

The decrease in the volume is steady, suggesting that further reductions would be possible with more iterations. However, the time per iteration increases with the number of cubes under consideration, and the time for this phase dominates the total verification time. Thus, for verifying this circuit, there is no incentive to further refine the region bounding the common-mode signal.

The second phase of the verification eliminates the unstable equilibrium. The equilibrium is near the point where all node voltages are 0.867V. We chose U to be the hyper-rectangle with sides of length 0.1V whose center is at this point. The region U contains all cubes that correspond to graph-vertices with self-loops from phase 1. There is more than one such cube due to the use of interval arithmetic in computing the adjacency graph to ensure soundness. Using the least-squares best-fit quadratic model for points in U yields:

$$\begin{aligned} \text{lin}_{\min} &> 5 \times 10^{10} \text{sec}^{-1}, \\ \sigma_{\max} &< 2 \times 10^9 \text{sec}^{-1} \text{V}^{-1}, \text{ and} \\ r &= 0.1 \text{V} \end{aligned}$$

from which we get that the conditions of Theorem 2 are satisfied for any k with $0 < k < 4.92 \times 10^{10} \text{sec}^{-1}$. Thus, we can safely remove the cubes in U .

We can now repeat the procedure from phase 1 to remove all cubes that transitively have no incoming flows. This phase eliminates roughly half of the remaining cubes, leaving 38384 cubes for analysis by the final phase.

The final phase starts with the 38384 cubes from the second phase. As described in Section IV-C, we divide these cubes into 16 wedges divided by “spokes” in the $u_1 \times u_2$ projection. As described in Section IV-C, it is sufficient to show trajectories starting on the boundary of the wedge lead to points inside the next wedge in the clockwise direction. With 16 wedges, we perform 48 reachability computation runs. At this point, the oscillator is verified.

We also ran a longer reachability computation starting from a spoke and completing two complete cycles of the oscillation. The second cycle starts from a smaller region than the first and establishes tighter bounds on the limit cycle. The blue polygons in Figure 4 indicate this limit cycle. The remaining




Fig. 3. Eliminating the unstable equilibrium

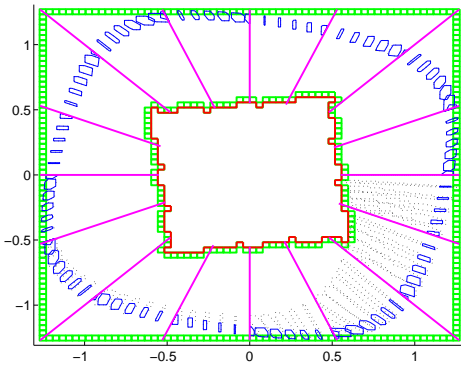


Fig. 4. Computing the invariant set

width of the limit cycle is mainly due to approximating the four-dimensional differential equation with a differential inclusion.

The computation is very efficient. The run-time of the first phase is about 720 seconds, and the run-time of reachability computation is less than 470 seconds. Reducing state space helps to improve performance significantly. It takes several hours to complete the reachability computation in full-dimensional space from a single cube. In contrast, the run-time in reduced-space is less than 10 minutes as shown above. Reducing the space also introduces over-approximations to the reachable regions. However, this did not lead to false-negative results, as the circuit converges to the oscillation orbit rapidly.

B. Verification for a range of sizes

Phases 1 and 3 of our verification method use conservative over-approximations to guarantee soundness of the results. These approximations make it straightforward to model s as being in an interval rather than having a precise value. We have verified escape from the failure set for values of s from 0.673 to 2.0 by testing values of s in steps of 0.01 for $0.6 \leq s \leq 2$ and in steps of 0.001 for $0.67 \leq s \leq 0.7$. The lower-bound for s is slightly higher than the one reported in [8]. We conjecture that our transistor current tables are slightly different than those used in [8] perhaps due to an updating of the SPICE

models provided by the foundry. For $s > 2$, the third phase of the verification fails to show that trajectories leave the “corners” of the $u_1 \times u_2$ space. These correspond to lock-up of the cross-coupled inverters. The DC analysis method shows that these lock-up states become stable for $s > 2.25$. The gap between the reachability computation and the DC analysis is presumably due to conservative over-approximations used in the reachability method.

VI. CONCLUSIONS

This paper has presented the first, formal verification that the differential oscillator circuit presented in [2] properly starts from almost all initial conditions. In particular:

- no “physically plausible” oscillator starts from all initial conditions (Theorem 1, Section II);
- we presented a generalization of the “cone-argument” from [1] to show that the failures occur with probability zero and thus the oscillator starts with probability one (Theorem 2, Section III);
- our approach shows how reachability analysis can be combined effectively with dynamical systems analysis;
- we showed how differential-operation, a common feature of analog designs, can be exploited for model reduction.

We elaborate on some of these below.

First, metastable behaviors is unavoidable for most mode-switching circuits. While metastability is most often associated with synchronizer circuits [21], [22], it arises anytime the state of a continuous system can evolve to two or more distinct states. For example, when a phase-locked loop (PLL) locks, the VCO phase may advance to match the phase of the reference, or the VCO may drop back depending on the initial conditions. Thus, there are conditions where any physically realizable PLL takes an arbitrarily long time to lock. On the other hand, there are published verifications of bounded lock time for phase-locked loops (e.g. [23]). The discrepancy is resolved by noting that [23] uses an abstract model for the phase-comparator that makes a discontinuous step as the phase-difference passes through 180° . For many designs, this is a reasonable abstraction; yet, we note that a PLL can fail to lock if there is a dead-spot in the response of the phase-comparator at the wrap-around point. We see our work as complementary to that of [23] – they provide powerful abstractions that enable the verification of larger designs, and we provide methods of ensuring that those abstractions are sound.

Second, our verification combined analytical methods from dynamical systems theory with reachability methods that are more typical of the formal methods community. Neither alone is sufficient to verify the oscillator. Reachability techniques are inadequate because they cannot show escape from a failure set of zero measure. Such “failures” are not of concern to practical designers as they are unobservable in the physical system. On the other hand, the dynamical systems methods that allow us to establish probability-one results are arguments about local dynamics. The reachability computations are needed to go from these local results to proving global properties.

The notion of probability that we used, a smooth distribution over initial states, was simplistic. A more physical model would use stochastic integration techniques to determine the evolution of this distribution under the circuit dynamics as perturbed by noise processes such as thermal noise. While this might be more satisfying, it would mainly serve to make the mathematics more complicated, and quantitative results would be hard to obtain due to the highly non-linear dynamics of the circuits. However, the basic topological observations on which we base our results would be preserved. Thus, we believe that our probability one results would continue to hold in a more detailed, stochastic model.

Proving that something happens “eventually” can be unsatisfying, as such proofs often don’t give an indication of how long one needs to wait. Our proof for Theorem 2 shows that the divergence is at least as fast as an exponential with time-constant k . For the oscillator considered, $k \approx 1/(20\text{ps})$. Thus, we can make a quantitative conclusion that in a few nanoseconds, the probability that the oscillator has not started is extremely small. This should satisfy practicing designers.

Of course, there are many areas of future work. Most immediately, we claimed escape from the failure set for a wide-range of inverter sizes by verifying the property for a large number of closely spaced choices of the sizes. We would like to use interval-arithmetic methods to show that these intervals are completely covered. To do so, we are making a few extensions to the intlab package [24]. Likewise, we plan to show that the method can be applied to a design in a more state-of-the-art process (e.g. using PTM models [25]). We expect to include results for interval arithmetic and other processes in the final version of this paper.

We would like to verify larger circuits. For example, a ring oscillator with six or more stages can have stable higher harmonic modes if small inter-stage couplings are included in the model. We would like to verify (and refute) such designs. We expect that the first two phases of our verification could readily be generalized to a oscillators with an arbitrary number of stages with straightforward inductive formulations. We don’t see induction working directly to extend the reachability analysis to larger designs. Instead, we are looking further into dynamical systems approaches to rule out entire classes of failure modes. Then we use reachability analysis techniques like those presented in the paper to complete the verification. The Rambus oscillator circuit is a good example for detailed analysis of how reachability computation complexity scales with circuit size.

REFERENCES

[1] I. Mitchell and M. Greenstreet, “Proving Newtonian arbiters correct, almost surely,” in *Proceedings of the Third Workshop on Designing Correct Circuits*, Båstad, Sweden, Sep. 1996.

[2] K. D. Jones, J. Kim, and V. Konrad, “Some “real world” problems in the analog and mixed-signal domains,” in *Proc. Workshop on Designing Correct Circuits*, Apr. 2008.

[3] W. Hartong, L. Heidrich, and E. Barke, “Model checking algorithms for analog verification,” in *Proceedings of the 39th ACM/IEEE Design Automation Conference*, Jun. 2002, pp. 542–547.

[4] S. Gupta, B. H. Krogh, and R. A. Rutenbar, “Towards formal verification of analog designs,” in *Proceedings of 2004 IEEE/ACM International Conference on Computer Aided Design*, Nov. 2004, pp. 210–217.

[5] G. Frehse, B. H. Krogh, and R. A. Rutenbar, “Verifying analog oscillator circuits using forward/backward abstraction refinement,” in *Proceedings of Design Automation and Test Europe*, Mar. 2006, pp. 257–262.

[6] S. Little, N. Seegmiller, D. Walter, C. Myers, and T. Yoneda, “Verification of analog/mixed-signal circuits using labeled hybrid petri nets,” in *Proceedings of the International Conference on Computer Aided Design*, Nov. 2006, pp. 275–282.

[7] J. Xiao, A. V. Peterchev, and S. R. Sanders, “Architecture and IC implementation of a digital VRM controller,” in *IEEE 32nd Annual Power Electronics Specialists Conference (PESC’01)*, vol. 1, Jun. 2001, pp. 38–47.

[8] M. R. Greenstreet and S. Yang, “Verifying start-up conditions for a ring oscillator,” in *Proceedings of the 18th Great Lakes Symposium on VLSI (GLSVLSI’08)*, May 2008, pp. 201–206.

[9] S. K. Tiwari, A. Gupta *et al.*, “fSpice: a boolean satisfiability based approach for formally verifying analog circuits,” presented at the 2008 Workshop on Formal Verification for Analog Circuits (FAC’08), Princeton, NJ, Jul. 2008.

[10] S. Tiwari, A. Gupta *et al.*, “First steps towards SAT-based formal analog verification,” in *Proceedings of the 2010 International Conference on Computer Aided Design*, Nov. 2010.

[11] S. Steinhilber, M. Peter, and L. Heidrich, “State space exploration of analog circuits by visualized multi-parallel particle simulation,” in *International Conference on Signal Processing Systems (ICSPS’09)*, May 2009, pp. 858–862.

[12] M. H. Zaki, I. Mitchell, and M. R. Greenstreet, “Towards a formal analysis of DC equilibria of analog designs,” Presented at the 2009 Workshop on Formal Verification for Analog Circuits (FAC’09), Grenoble, France, Jun. 2009.

[13] M. Fränzle, “HySAT: An efficient proof engine for bounded model checking of hybrid systems,” *Formal Methods in System Design*, vol. 30, no. 3, pp. 179–198, 2007.

[14] S. Little and C. Myers, “Abstract modeling and simulation aided verification of analog/mixed-signal circuits,” presented at the 2008 Workshop on Formal Verification for Analog Circuits (FAC’08), Princeton, NJ, Jul. 2008.

[15] C. Yan and M. R. Greenstreet, “Faster projection based methods for circuit-level verification,” in *Proceedings of the 2008 Asia and South Pacific design automation conference (ASPDAC’08)*, Jan. 2008, pp. 410–415.

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

[17] G. Frehse, “PHAVer: Algorithmic verification of hybrid systems past HyTech,” in *Proceedings of the Fifth International Workshop on Hybrid Systems: Computation and Control*. Springer-Verlag, 2005, pp. 258–273, LNCS 3414.

[18] D. Pollard, *A User’s Guide to Measure Theoretic Probability*. Cambridge University Press, 2001.

[19] T. A. Henzinger, P.-H. Ho, and H. Wong-Toi, “Algorithmic analysis of nonlinear hybrid systems,” *IEEE Transactions on Automatic Control*, vol. 43, no. 4, pp. 540–554, Apr. 1998.

[20] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.

[21] T. Chaney and C. Molnar, “Anomalous behavior of synchronizer and arbiter circuits,” *IEEE Transactions on Computers*, vol. C-22, no. 4, pp. 421–422, Apr. 1973.

[22] D. J. Kinniment, C. Dike *et al.*, “Measuring deep metastability and its effect on synchronizer performance,” *IEEE Transactions on VLSI Systems*, vol. 15, pp. 1028–1039, Sep. 2007.

[23] M. Althoff, A. Rajhans *et al.*, “Formal verification of phase-locked loops using reachability analysis and continuization,” in *Proceedings of the 2011 International Conference on Computer Aided Design*, Nov. 2011, pp. 659–666.

[24] S. Rump, “INTLAB - INTerval LABoratory,” in *Developments in Reliable Computing*, T. Csendes, Ed. Dordrecht: Kluwer Academic Publishers, 1999, pp. 77–104, <http://www.ti3.tu-harburg.de/rump/>.

[25] Y. Cao, “PTM: predictive technology model,” <http://ptm.asu.edu>, 2008.