Exact Timing Analysis for Asynchronous Circuits with Multiple Periods

Rajit Manohar, Yale University

***Abstract*—The timing properties of asynchronous circuits can be summarized using cyclic graphs that capture *max*-delay constraints between signal transitions. There are many results on the timing analysis problem, but they all make various simplifying assumptions on the connectivity properties of the underlying timing graph. Most results provide approximate timing characteristics, with a few providing exact results on the circuit's timing behavior.**

In this paper we provide results that exactly characterize the timing properties for a more general class of *max*-delay constraints. We show that the circuit can be partitioned into regions with different periodicities, and provide an efficient algorithm to compute all the periods of the system.

***Index Terms*—exact timing, asynchronous systems, periodic.**

I. Introduction

Analyzing the timing properties of asynchronous circuits is a complex problem. Asynchronous circuits contain cycles of gates, and characterizing their timing behavior requires analyzing these cycles and their interactions. There are a number of results on the timing behavior of asynchronous circuits, and they use graph-based abstractions to represent timing. *Event-Rule (ER) systems*—and their finite versions *repetitive ER (RER) systems*—are commonly used to capture the timing behavior of asynchronous circuits [1], as are restricted classes of Petri nets [2], [3] such as marked graphs [4], process graphs, or timing constraint graphs. We use ER and RER systems in this paper, as the most recent results in the literature that we build on use this abstraction. RER systems also have the benefit that they are a direct generalization of the timing graph used by synchronous timing analysis, making them accessible to those unfamiliar with asynchronous timing analysis.

Many asynchronous circuits have RER systems with properties that make them more amenable to analysis. For example, RER systems corresponding to simple asynchronous controllers or coupled controllers are often a single strongly connected component [1]. Hence, many existing approaches to asynchronous timing analysis make assumptions on the connectivity properties of RER systems—in fact, most assume that the RER system is strongly connected [1], [5]. Figure 1 shows a bundled-data asynchronous pipeline, where one pipeline writes to a register and another pipeline reads from the register. In this example, the control logic for pipeline A would be in one strongly connected component in the RER system, and amenable to analysis using [5]. However, there is no path from the datapath back to the control, hence the datapath is not part of the strongly connected component. This assumption was first weakened in [6], where it was assumed that the RER system is critically connected (elaborated in Section II). In Figure 1, the datapath for pipeline A is critically connected to its control, and hence it can be included in the analysis. The same holds for pipeline B. However, the combined circuit requires a more complex RER system model, since the overall system is not critically connected, and communication between pipelines A and B occur through a shared register. Intuitively, the two pipelines may operate at different steady-state frequencies. Ideally the implementation of a CAD tool for timing analysis must be prepared to accept an arbitrary RER system as input.

In this paper we weaken the connectivity assumption further—no longer requiring that the RER system be critically connected—and provide a complete analysis of the timing properties of such RER systems. The analysis formalizes our intuitive understanding of circuits such as those from Figure 1 discussed above, showing that we can analyze various parts of the system separately and combine them to build a complete model for the entire circuit. In particular, we show that the RER system can be partitioned into a disjoint set of transitions. Each set of transitions exhibit periodic behavior, but the different sets can have different periodicity. Mathematically, there can also be a set of degenerate transitions where an infinite number of transitions occur at the same time; this is an unphysical case and should be reported as an error by the analysis engine. Based on the analysis, we show how these sets and periods can be computed in an efficient manner.

II. Background

A. Timing graphs for asynchronous circuits

In event-rule systems, a circuit is modeled by a set of *events* E, and a set of *rules* $R \subseteq E \times E \times \mathbb{R}_{\geq 0}$. An event is often

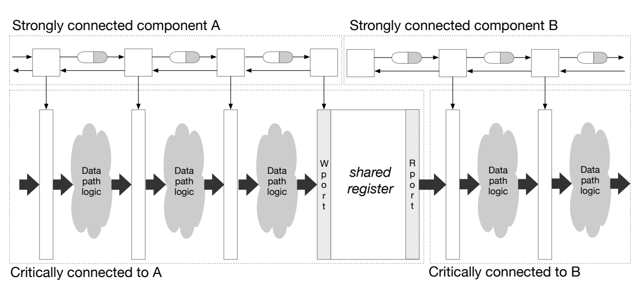

Fig. 1. Two bundled-data asynchronous pipelines interacting through a shared register. The control circuit for pipeline A is in a strongly connected component, as is the control circuit for pipeline B. The datapath logic for each pipeline is critically connected to the respective control circuits. While there are edges from the datapath of A to the datapath of B, B is not critically connected to A.

This is the author's version of an article that has been published in this journal. Changes were made to this version by the publisher prior to publication.
The final version of record is available at http://dx.doi.org/10.1109/TCAD.2019.2963271

of the form $\langle x\uparrow, i\rangle$ or $\langle x\downarrow, i\rangle$, where $i \geq 0$ is an integer that captures the *occurrence index* of the signal transition $x\uparrow$ (a zero to one transition on x) or $x\downarrow$ (a one to zero transition on x) respectively. Rules are of the form $e \overset{\alpha}{\mapsto} f$, where $e, f \in E$ are events and α is a non-negative real number that specifies a delay constraint. The rule specifies that the earliest time event f (the *target*) can occur is α time units after event e (the *source*). Multiple rules with the same target correspond to multiple constraints that must all be satisfied; therefore, ER systems capture AND-causality (a.k.a. *max*-causality). There are also approaches that include OR-causality (*min-max* causality); we restrict ourselves to AND causality in this paper.

ER systems are represented by graphs, where the vertices are events, and the edges correspond to rules. Edges are labelled with the delay constraint α corresponding to the rule. ER systems are acyclic graphs, and are typically infinite.

The *timing simulation* of an ER system $\hat{t}(\cdot)$ maps events to a real number (the time), and captures the time at which the event occurs. The timing simulation is defined recursively by

$$\hat{t}(f) = \max\{\hat{t}(e) + \alpha \,|\, e \overset{\alpha}{\mapsto} f \in R\}$$

ER systems have *initial* events ι, where the set on the right hand side of the equation above is empty; for those events, $\hat{t}(\iota)$ is defined to be zero.

Asynchronous digital circuits have oscillatory behavior, so the sets E and R for a typical circuit are infinite. However, the actual circuits themselves are finite structures. This is captured by modelling the oscillatory behavior of the circuit using a finite structure called a *repetitive* ER (RER) system [1]. An RER system captures one iteration of the circuit, as well as the timing constraints between one iteration and the next. An RER system consists of a set of *transitions* T and *rule templates* $RT \subseteq T \times T \times \mathbb{R}_{\geq 0} \times \{0, 1\}$. A transition is often of the form $x\uparrow$ or $x\downarrow$, and a rule template is of the form $t \overset{\alpha,\epsilon}{\mapsto} u$ where $t, u \in T$ are transitions, $\alpha \geq 0$ is a delay, and ϵ is either 0 or 1. The corresponding ER system is obtained by constructing the event set $E = T \times \mathbb{N}$, and rule set R as

$$R = \{\langle t, i\rangle \overset{\alpha}{\mapsto} \langle u, i+\epsilon\rangle \,|\, t \overset{\alpha,\epsilon}{\mapsto} u \in TR, i \in \mathbb{N}\}$$

Rule templates that have $\epsilon = 0$ relate transitions that have the same occurrence index (i.e. are part of the same iteration); those with $\epsilon = 1$ relate transitions from one iteration to the next. RER systems can be represented by graphs, where the vertices are transitions, and the edges correspond to rule templates. Edges are labelled with the delay α from the corresponding rule template, and also marked with a tick when $\epsilon = 1$. RER systems are cyclic structures. An example of a circuit and its corresponding RER system is shown in Figure 2.

Given a path p in the graph corresponding to an RER system, the delay along the path $\delta(p)$ is given by the sum of the values of α for all the edges in the path, and $\epsilon(p)$ is given by the sum of the values of ϵ along the path. The key performance measure for an RER system is its *cycle period*. The cycle period $p^\star$ is defined as the maximum value of the ratio $\delta(c)/\epsilon(c)$ for any cycle c in the RER system graph. A cycle that achieves the value of $p^\star$ is called a *critical cycle*. It should be clear that it suffices to examine only simple cycles of the RER system when finding a critical cycle [1].

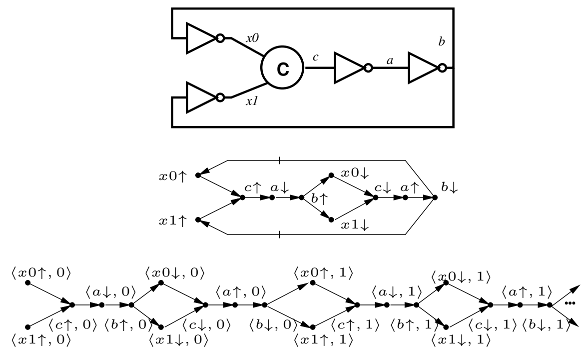

Fig. 2. Example of an asynchronous circuit and its corresponding RER system. The system is initialized so that the initial events are $\langle x0\uparrow, 0\rangle$ and $\langle x1\uparrow, 0\rangle$. Edges with $\epsilon = 1$ are denoted with a tick mark, and delays are omitted for clarity. A fragment of the ER system corresponding to the RER system is also shown.

B. Major prior results

RER systems were introduced by Burns to model asynchronous circuits with AND-causality [1]. His work assumed that RER systems are strongly connected, and showed that such RER systems are, on average, periodic with period $p^\star$. Similar results were also obtained by Murata [2] and Magott [3]. Hulgaard showed that strongly connected RER systems eventually become exactly periodic [5]. Hua generalized the exact periodicity result of Hulgaard to RER systems that were not strongly connected, but where all transitions were reachable from a critical cycle [6]. Hua also provided bounds for the time required to reach this steady-state behavior. The strongest results [6] show that there is an unfolding factor M, and an integer K such that

$$\hat{t}(\langle s, i+M\rangle) = \hat{t}(\langle s, i\rangle) + Mp^\star \qquad \forall s \in T, i \geq K \qquad (1)$$

This shows that after an initial finite period (determined by K), the gap between transitions that are M occurrence indices apart is precisely M times the cycle period $p^\star$ of the RER system.[1]

In this paper, we generalize the results of Hua to *arbitrary* RER systems. We impose no requirements on the connectivity structure of RER systems, and show that the transitions can be partitioned into sets, with each set being exactly periodic with its own $(p^\star, M)$ pair. We also provide a polynomial-time algorithm for computing both the sets as well as the cycle periods for an RER system.

C. Definitions

An obvious property of the timing simulation $\hat{t}(\cdot)$ is the following:

$$\hat{t}(e) = \max\{\delta(p) \,|\, p \text{ is a path from an initial event to } e\}$$

[1]Hulgaard's result can be viewed as the same property as equation (1), but for strongly connected RER systems.

This follows immediately from the recursive definition of the function $\hat{t}(\cdot)$. The main previous result shows that the timing simulation for critically connected RER systems is periodic.

Definition 1 (periodicity)**.** *Given an RER system $\mathcal{R} = (T, RT)$, set $S \subseteq T$, and $p \in \mathbb{R}$, we define a predicate $\mathsf{periodic}(\mathcal{R}, S, p)$ to mean that there exist integers K and M such that*

$$\forall s \in S\colon \hat{t}_{\mathcal{R}}(\langle s, i+M\rangle) = \hat{t}_{\mathcal{R}}(\langle s, i\rangle) + Mp \qquad \forall i \geq K$$

If the same condition holds for an ER system $\mathcal{E}$ that contains all the events and rules induced by the RER system $\mathcal{R}$, we use the predicate $\mathsf{periodic}(\mathcal{E}, S, p)$ to have the same meaning as above for the timing simulation of the ER system.

From prior work, we know that

Theorem 1 ([6])**.** *Consider an RER system (T, RT) where every transition is reachable from some transition on a critical cycle. Let (E, R) be the ER system obtained from the given RER system. Then $\mathsf{periodic}(\mathcal{E}, T, p^\star)$ holds, where $p^\star$ is the cycle period of the RER system (T, RT).*

We need a simple Lemma that observes that changing the times of any initial event does not modify the periodicity property of the timing simulation. Technically, we introduce fresh events that are used to modify the times of some events of the form $\langle s, 0\rangle$. Burns refers to such systems as pseudo-repetitive ER systems, because they correspond to a finite ER system followed by a repetitive ER system [1].

Lemma 1. *Consider an RER system (T, RT) where every transition is reachable from some transition on a critical cycle. Let (E, R) be the ER system obtained from the given RER system. Let EI be a (fresh) finite set of events, and let RI be a set of rules that are contained in $EI \times (T \times \{0\}) \times \mathbb{R}_{\geq 0}$. Let $\hat{t}'$ be the timing simulation of the new ER system $\mathcal{E} = (E \cup EI, R \cup RI)$. Then $\mathsf{periodic}(\mathcal{E}, T, p^\star)$ holds, where $p^\star$ is the cycle period of the original RER system (T, RT).*

Proof. (Sketch) The only difference between the two ER systems is the time of initial events—the EI events are used to potentially change the time of the initial events to some non-zero quantity. The value of the time of initial events does not affect any of the results from [6] regarding the periodicity property, since the results focus on the steady-state behavior of the ER system. □

III. General RER Systems

To analyze general RER systems, we will examine multiple distinct RER systems simultaneously. Given an RER system $\mathcal{R} = (T, RT)$, its timing simulation is denoted as $\hat{t}_{\mathcal{R}}$ to make it clear that the timing simulation is the one corresponding to the RER system $\mathcal{R}$.

A. Degenerate transitions

Consider the RER system $(\{x, y\}, \{x \overset{0.5,0}{\mapsto} y\})$. Building the ER system and computing the timing simulation, we observe that for all i,

$$\hat{t}(\langle x, i\rangle) = 0 \qquad \hat{t}(\langle y, i\rangle) = 0.5 \quad .$$

x and y are *degenerate* transitions, and we would like to eliminate transitions of these types as they have trivial timing simulations. Furthermore, such an RER system does not correspond to a physical circuit since an infinite number of events occur at a single point in time. Note that this degeneracy holds for any transition that is not reachable from a cycle in the RER system, as we establish next.

Lemma 2 (degenerate transitions)**.** *Let $S \subseteq T$ be the set of transitions that are not reachable from any cycle in an RER system (T, RT). Then there exists a K such that for all $i > K$ and $s \in S$, there is a constant c such that*

$$\hat{t}(\langle s, i\rangle) = c$$

Proof. For $i > |T|$ and $s \in S$, consider all paths from some initial event to event $\langle s, i\rangle$. All such paths must be of length less than $|T|$; otherwise, some transition t on the path must repeat, violating the assumption that s cannot be reached from any cycle in the RER system.

Each path π will be of the form $\langle a_0, \epsilon_0\rangle$, $\langle a_1, \epsilon_1\rangle$, $\ldots$, $\langle a_{n-1}, \epsilon_{n-1}\rangle$, $\langle s, i\rangle$. Observe that if there is a path from a cycle to a_i, there is also a path from the same cycle to s. Hence, all transitions a_0, $\ldots$, a_n must be members of S—in other words, paths to initial events can only refer to transitions from S. Furthermore, for any such path, $\epsilon_0 > 0$ by the selection of $i > |T|$ since $0 \leq \epsilon_k - \epsilon_{k-1} \leq 1$.

The second observation is that for $i > |T|$, path $\pi = \langle a_0, \epsilon_0\rangle$, $\langle a_1, \epsilon_1\rangle$, $\ldots$, $\langle a_{n-1}, \epsilon_{n-1}\rangle$, $\langle s, i\rangle$ is a path from an initial event to $\langle s, i\rangle$ iff $\pi^+ = \langle a_0, \epsilon_0 + 1\rangle$, $\langle a_1, \epsilon_1 + 1\rangle$, $\ldots$, $\langle a_{n-1}, \epsilon_{n-1} + 1\rangle$, $\langle s, i+1\rangle$ is a path from an initial event to $\langle s, i+1\rangle$. This follows from the fact that the ER system is generated from the RER system with rule templates, and that $\epsilon_0 > 0$.

Let $A_0 \subset S$ be the set of transitions that correspond to the zeroth event $\langle a_0, \epsilon_0\rangle$ for any path π to an event $\langle s, i\rangle$ for $s \in S$. Since $\epsilon_0 > 0$, that means that there is no rule template that has a_0 as its target; hence, $\hat{t}(\langle a_0, i\rangle) = 0$ for all $i > |T|$, $a_0 \in A_0$.

Let $A_i \subset S$ be the set of transitions that correspond to the ith event $\langle a_i, \epsilon_i\rangle$ for $i > 0$. If a transition occurs in more than one such path with different indices i, select the largest value of i to assign the transition to an A-set.

We use induction to show that if the timing simulation is constant for events $\langle s, k\rangle$ for $k > |T|$ and $s \in A_j$ $(j \leq i)$, then it is constant for the events $\langle s, k\rangle$ for $k > |T|$ and $s \in A_{i+1}$. The only rule templates that have transitions from A_{i+1} as a target have a source in A_j $(j \leq i)$—-which have a constant timing simulation by the induction hypothesis. Hence, the timing simulation for $\langle a_{i+1}, k\rangle$ will also be a constant for $k > |T|$, since the set of paths for $\langle a_{i+1}, k\rangle$ correspond to the same set of rule templates for all $k > |T|$. □

B. Sub-RER systems

In what follows, we assume that the RER system has no degenerate transitions as those can be handled by the analysis in Lemma 2. (One way to view Lemma 2 is to say that degenerate events have a cycle period that is zero.)

Construction 1 RER system decomposition

1: $i, n \leftarrow 0, 0$
2: $(T_0, RT_0) \leftarrow \mathcal{R}$
3: $\mathcal{R}_0 \leftarrow (T_0, RT_0)$
4: **while** (T_i, RT_i) has cycles **do**
5: $\quad n \leftarrow i$
6: $\quad p_i^\star \leftarrow$ critical cycle ratio of $\mathcal{R}_i$
7: $\quad C_i \leftarrow \{c | c$ is a transition on any crit. cycle of $\mathcal{R}_i\}$
8: $\quad T_i' \leftarrow \{t | t \in T_i \wedge t$ is reachable from $C_i\}$
9: $\quad T_{i+1} \leftarrow T_i - T_i'$
10: $\quad RT_{i+1} \leftarrow RT_i|_{T_{i+1}}$
11: $\quad \mathcal{R}_i' \leftarrow (T_i', RT_i|_{T_i'})$
12: $\quad \mathcal{R}_{i+1} \leftarrow (T_{i+1}, RT_{i+1})$
13: $\quad i \leftarrow i + 1$
14: **end while**

Definition 2. *Given an RER system $\mathcal{R} = (T, RT)$, the sub-RER system of $\mathcal{R}$ induced by $S \subseteq T$ is the RER system given by $(S, RT|_S)$ where $RT|_S = RT \cap S \times S \times \mathbb{R} \times \{0, 1\}$.*

Given a subset $S \subset T$ of transitions, the sub-RER system corresponds to the constraints that are only related to transitions in S. Transitions in $(T - S)$ can be also viewed as resulting in a sub-RER system. There will also be rule templates that relate transitions in S and $(T - S)$ to each other; we refer to these as *cross* rule templates, as they cross the boundary of the two sub-RER systems.

Definition 3. *Given an RER system $\mathcal{R} = (T, RT)$, a partition of $\mathcal{R}$ induced by $S \subset T$ is given by the two sub-RER systems induced by S and $T - S$, along with the set of cross rule templates $RT_\times$ given by $RT_\times = RT - (RT|_S \cup RT|_{(T-S)})$.*

A special case of a partition is when all the edges in $RT_\times$ have their source transition in S. We refer to this case as an *acyclic partition*.

Definition 4. *An acyclic partition of an RER system $\mathcal{R} = (T, RT)$ induced by $S \subset T$ is a partition whose cross rule templates satisfy $RT_\times = RT \cap (S \times T \times \mathbb{R} \times \{0, 1\})$.*

Lemma 3. *Let the partition of RER system $\mathcal{R} = (T, RT)$ induced by $S \subset T$ be acyclic, and let $\mathcal{R}_S$ denote the sub-RER system $(S, RT|_S)$. Then*

$$\hat{t}_{\mathcal{R}}(\langle s, i \rangle) = \hat{t}_{\mathcal{R}_S}(\langle s, i \rangle) \qquad \forall\ (s, i) \in S \times \mathbb{N}$$

Proof. This follows immediately from the definition of the timing simulation, since: (i) there is no path from any event in $(T - S) \times \mathbb{N}$ to any event in $S \times \mathbb{N}$ by the assumption that the partition is acyclic; and (ii) the timing simulation of an event $\langle s, i \rangle$ is equal to the maximum delay of paths to $\langle s, i \rangle$ from the set of initial events. □

C. The general case

Lemma 3 provides a way to study the timing simulation of complex RER systems through their sub-RER systems. Armed with this Lemma, we construct a sequence of RER systems using Construction 1.

It is clear that construction 1 terminates, since each loop iteration reduces the cardinality of the finite set T_i by at least one. We remark that, on termination, the set T_{n+1} is empty because degenerate transitions have been eliminated. The algorithm constructs a sequence of sub-RER systems $\mathcal{R}_0$, …, $\mathcal{R}_n$. The key idea is that we partition the original RER system into sub-RER systems, where each sub-RER system satisfies the property needed by [6]—namely, that all transitions are reachable from a critical cycle. $\mathcal{R}_0$ contains the critical cycle and reachable transitions from the original RER system. This also means that the critical cycle period $p_1^\star$ of $\mathcal{R}_1$ must be strictly smaller than $p_0^\star$. Also, any cross edges between $\mathcal{R}_0$ and (T_1, RT_1) must originate from T_1 by construction of T_0'. The net result is that construction 1 guarantees that $p_0^\star > p_1^\star > \cdots > p_n^\star$.

Our approach to analyzing the timing simulation of the original RER system is to build it incrementally from the sub-RER systems produced by construction 1. To this end, it is helpful to define a collection of auxiliary RER systems ${\mathcal{R}'}_i = (T_i, RT_i)$ from construction 1.

The RER systems constructed by construction 1 have the following relationships by the way they are constructed:

1) $\mathcal{R}_0$ is the original RER system (lines 2–3);
2) $\mathcal{R}_i$ is partitioned into sub-RER systems $\mathcal{R}_{i+1}$ and $\mathcal{R}_i'$, and $\mathsf{periodic}(\mathcal{R}_i', T_i', p_i^\star)$ holds (lines 6–12);
3) The partition is acyclic—any cross edges are from $\mathcal{R}_{i+1}$ to $\mathcal{R}_i'$, as T_i' contains all transitions reachable from the set of critical cycles of $\mathcal{R}_i$ (lines 7–8);
4) $T_i = \bigcup_{k=i}^n T_k'$, and the sets T_k' are disjoint;
5) $\mathcal{R}_{n+1}$ is empty as degenerate transitions have been eliminated.

Given these properties of the construction, we remark that

Lemma 4 (cross edges)**.** *Given the sequence of RER systems $\mathcal{R}_0'$, …, $\mathcal{R}_n'$ with transition sets T_0', …, T_n', any rule template from the original RER system is in $T_i' \times T_j' \times \mathbb{R}_{\geq 0} \times \{0, 1\}$ where $j \neq i$ satisfies $j > i$.*

Proof. This is immediate from the acyclic nature of the partitions. □

We now build the timing simulation of the original RER system by analyzing the sequence of RER systems we have constructed.

Lemma 5 (base case)**.** *Given an RER system $\mathcal{R}$ and the sequence of RER systems $\mathcal{R}_i'$ with transition sets T_i' from construction 1, then the predicate $\mathsf{periodic}(\mathcal{R}_n, T_n', p_n^\star)$ holds.*

Proof. The proof of this is immediate from Lemma 4 and Lemma 3, since $\mathcal{R}_n = \mathcal{R}_n'$ is a sub-RER system obtained by an acyclic partition of $\mathcal{R}$ induced by transition set T_n'. □

Lemma 6 (induction step)**.** *Let $\mathcal{R}$ be an RER system, and $\mathcal{R}_i$ be the the sequence of RER systems from construction 1. Furthermore, assume that for all l satisfying $k < l \leq n$, the predicate $\mathsf{periodic}(\mathcal{R}_{k+1}, T_l', p_l^\star)$ holds. Then for all l such that $k - 1 < l \leq n$, the predicate $\mathsf{periodic}(\mathcal{R}_k, T_l', p_l^\star)$ holds.*

Proof. By the induction hypothesis, we know that for l satisfying $k < l \leq n$, the predicate $\mathsf{periodic}(\mathcal{R}_{k+1}, T_l', p_l^\star)$ holds.

Consider the RER system $\mathcal{R}_k = (T_k, RT_k)$. By construction, $T_k = T_{k+1} \cup T'_k$. Furthermore, T'_k induces an acyclic partition of $\mathcal{R}_k$. Hence, by Lemma 3,

$$\forall i, s \in T_{k+1} \colon \hat{t}_{\mathcal{R}_k}(\langle s, i\rangle) = \hat{t}_{\mathcal{R}_{k+1}}(\langle s, i\rangle)$$

Hence, $\mathsf{periodic}(\mathcal{R}_k, T'_l, p^\star_l)$ holds for $k < l \le n$.

Consider the events in T'_k. From [6], we know that $\mathcal{R}'_k$ satisfies $\mathsf{periodic}(\mathcal{R}'_k, T'_k, p^\star_k)$, and $p^\star_k > \max_{k<l\le n} p^\star_l$. The ER system corresponding to $\mathcal{R}_k$ has three types of rules: those that relate transitions from T_{k+1} (contained in the ER system for $\mathcal{R}_{k+1}$); those that relate transitions from T'_k (contained in the ER system for $\mathcal{R}'_k$); and those that are from T_{k+1} to T'_k.

Consider a rule template $s \overset{\alpha;\epsilon}{\mapsto} t$ where $s \in T_{k+1}$ and $t \in T'_k$. Suppose $s \in T'_l$ for $l > k$. We know that for $i > K$, there is an M_l such that

$$\hat{t}_{\mathcal{R}_k}(\langle s, i\rangle) = \hat{t}_{\mathcal{R}_k}(\langle s, K + (i \bmod M_l)\rangle) + M_l p^\star_l \times \lfloor i/M_l \rfloor$$

Furthermore, we know that the RER system $\mathcal{R}'_k$ also satisfies $\mathsf{periodic}(\mathcal{R}'_k, T'_k, p^\star_k)$. All the constraints from $\mathcal{R}'_k$ are present in the ER system for $\mathcal{R}_k$. Hence, paths that determine the timing simulation in $\mathcal{R}_k$ include all paths from $\mathcal{R}'_k$. The latter paths imply that

$$\hat{t}_{\mathcal{R}_k}(\langle t, i\rangle) \ge \hat{t}_{\mathcal{R}_k}(\langle t, K + (i \bmod M_k)\rangle) + M_k p^\star_k \times \lfloor i/M_k \rfloor$$

Since $p^\star_k > p^\star_l$ and α is finite, for i large enough the rule $\langle s, i\rangle \overset{\alpha}{\mapsto} \langle t, i+\epsilon\rangle$ can be *deleted* from the ER system obtained from $\mathcal{R}_k$ without changing its timing simulation.

Hence, after a finite number of iterations, all cross rules from transitions in T_{k+1} to transitions in T'_k can be deleted without changing the timing simulation. Hence, after a finite number of iterations, the ER system is the same as the one obtained by taking the union of the rules from $\mathcal{R}_{k+1}$ and $\mathcal{R}'_k$ *without* including any cross rules. Hence, by Lemma 1, the timing simulation for events in T'_k satisfies $\mathsf{periodic}(\mathcal{R}_k, T'_k, p^\star_k)$. This concludes the proof. □

This construction is used to establish the main result.

Theorem 2 (multi-periodicity)**.** *Given an RER system $\mathcal{R}$ and the sequence of RER systems with transition sets T'_i from construction 1, for all j satisfying $0 \le j \le n$,* $\mathsf{periodic}(\mathcal{R}, T'_j, p^\star_j)$.

Proof. This follows by induction from Lemma 5 and Lemma 6, since $\mathcal{R}_0$ is the original RER system $\mathcal{R}$. □

D. Algorithm for computing periods

Based on construction 1, we can compute all the periods for the RER system by the algorithm below. We denote SCC as the algorithm that computes strongly connected components of a graph, and returns a partition of the vertices corresponding to the components, and YTO as the algorithm that computes the critical cycle ratio of the graph using the Young-Tarjan-Orlin algorithm [7], which is similar to the Karp-Orlin algorithm [8] but uses Fibonacci heaps to reduce computational complexity.

Algorithm 2 Cycle period computation

```
 1: 𝒯 ← SCC(T, RT)
 2: 𝒰 ← ∅
 3: while C ∈ 𝒯 has cycles do
 4:     p_C ← YTO(C, RT|_C)
 5:     𝒰 ← 𝒰 ∪ {C}
 6:     𝒯 ← 𝒯 − {C}
 7: end while
 8: while 𝒰 ≠ ∅ do
 9:     X ← arg max_{X∈𝒰} p_X
10:     𝒴 ← {X} ∪ components reachable from X in (T, RT)
11:     p_X is the cycle period for all components in 𝒴
12:     𝒰 ← 𝒰 − 𝒴
13: end while
14: Components remaining in 𝒯 are degenerate
```

The complexity of this algorithm is dominated by YTO, since the rest of the algorithm has complexity that is linear in the size of $|RT|$—the number of edges in the timing graph. (Since the graph is connected, $|RT| \ge |T| - 1$). Hence the overall algorithm has complexity $\mathcal{O}(|T||RT|\log|RT|)$, or complexity $\mathcal{O}(|RT||T| + |T|^2 \log|T|)$ when implemented with Fibonacci heaps [7]. This complexity corresponds to the case when YTO is executed on the entire graph—i.e. when the graph is a single strongly connected component. Since the time complexity of YTO increases faster than linearly with the size of the graph, partitioning the graph into components and running YTO on the individual components takes less time than the case where the entire graph is one strongly connected component. Hence, $\mathcal{O}(|RT||T| + |T|^2 \log|T|)$ is a valid complexity bound. Note that since the average fanout of a typical circuit is usually bounded by a small constant, the size of $|RT|$ is typically $\mathcal{O}(|T|)$; hence using a Fibonacci heap is unlikely to improve the asymptotic complexity of the algorithm.

IV. Discussion

We have provided a complete characterization of the periodicity properties of RER systems, and hence asynchronous circuits that can be described using AND causality. Our work extends a sequence of previous results. Early results in the concurrent systems literature showed that the asymptotic behavior of the timing simulation is periodic for strongly connected RER systems. This was followed by results that showed that strongly connected RER systems were in fact exactly periodic. Most recently, it was shown that the RER system need not be strongly connected for the exact periodicity property to hold—in sufficed that the RER system was critically connected. This paper completes the analysis, showing that the transitions in any RER system can be partitioned into (possibly empty) degenerate transitions, and sets of transitions that each are exactly periodic but where different sets have different periodicity. Our analysis also led to a simple method to compute the partitions and the cycle periods of the partitions, re-using algorithms from the literature.

This analysis can serve as the foundation for the implementation of a timing engine for asynchronous circuits. Degenerate transitions can be reported as potentially erroneous—or, more likely, an error in the input to the timing analysis engine since such transitions are unphysical. Interactions between circuit components that have different periodicity should also be flagged. Such interactions can appear in the timing graph for multiple reasons, including design errors, an error in the construction of the timing graph due to approximations in graph construction, or because the design has been specifically engineered to operate correctly in spite of such interactions. Finally, both the period and the M parameter can be reported for each component.

References

[1] S. M. Burns and A. J. Martin, "Performance analysis and optimization of asynchronous circuits," in *Conference on Advanced Research in VLSI*, 1990.

[2] T. Murata, "Petri nets: Properties, analysis and applications," *Proceedings of the IEEE*, vol. 77, no. 4, pp. 541–580, 1989.

[3] J. Magott, "Performance evaluation of concurrent systems using petri nets," *Information processing letters*, vol. 18, no. 1, pp. 7–13, 1984.

[4] F. Commoner, A. W. Holt, S. Even, and A. Pnueli, "Marked directed graphs," *Journal of Computer and System Sciences*, vol. 5, no. 5, pp. 511–523, 1971.

[5] S. Burns, H. Hulgaard, T. Amon, and G. Borriello, "An algorithm for exact bounds on the time separation of events in concurrent systems," *IEEE Transactions on Computers*, vol. 44, no. 11, pp. 1306–1317, 1995.

[6] W. Hua and R. Manohar, "Exact timing analysis for asynchronous systems," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 37, no. 1, pp. 203–216, 2018.

[7] N. E. Young, R. E. Tarjan, and J. B. Orlin, "Faster parametric shortest path and minimum-balance algorithms," *Networks*, vol. 21, no. 2, pp. 205–221, 1991.

[8] R. M. Karp and J. B. Orlin, "Parametric shortest path algorithms with an application to cyclic staffing," *Discrete Applied Mathematics*, vol. 3, no. 1, pp. 37–45, 1981.