

# Bounded Cycle Synthesis<sup>\*</sup>

Bernd Finkbeiner and Felix Klein

Reactive Systems Group, Saarland University, Germany  
{finkbeiner,fklein}@cs.uni-saarland.de

**Abstract.** We introduce a new approach for the synthesis of Mealy machines from specifications in linear-time temporal logic (LTL), where the number of cycles in the state graph of the implementation is limited by a given bound. Bounding the number of cycles leads to implementations that are structurally simpler and easier to understand. We solve the synthesis problem via an extension of SAT-based bounded synthesis, where we additionally construct a witness structure that limits the number of cycles. We also establish a triple-exponential upper and lower bound for the potential blow-up between the length of the LTL formula and the number of cycles in the state graph.

## 1 Introduction


There has been a lot of recent progress in the automatic synthesis of reactive systems from specifications in temporal logic [4, 6, 7, 9, 12]. From a theoretical point of view, the appeal of synthesis is obvious: the synthesized implementation is guaranteed to satisfy the specification. No separate verification is needed.

From a practical point of view, the value proposition is not so clear. Instead of writing programs, the user of a synthesis procedure now writes specifications. But many people find it much easier to understand the precise meaning of a program than to understand the precise meaning of a temporal formula. Is it really justified to place higher trust into a program that was synthesized automatically, albeit from a possibly ill-understood specification, than in a manually written, but well-understood program? A straightforward solution would be for the programmer to *inspect* the synthesized program and confirm that the implementation is indeed as intended. However, current synthesis tools fail miserably at producing readable code.

Most research on the synthesis problem has focused on the problem of finding *some* implementation, not necessarily a high-quality implementation. Since specification languages like LTL restrict the behavior of a system, but not its structure, it is no surprise that the synthesized implementations are often much larger and much more complex than a manual implementation. There has been some progress on improving other quality measures, such as the runtime performance [4], but very little has been done to optimize the *structural quality* of

---

<sup>\*</sup> Partially supported by the DFG project “AVACS” (SFB/TR 14). The second author was supported by an IMPRS-CS PhD Scholarship.



**Fig. 1.** Three implementations of the TBURST4 component of the AMBA bus controller. Standard synthesis with Acacia+ produces the state graph on the left with 14 states and 61 cycles. Bounded synthesis produces the graph in the middle with 7 states and 19 cycles. The graph on the right, produced by our tool, has 7 states and 7 cycles, which is the minimum.


the synthesized implementations (cf. [14]). Can we develop synthesis algorithms that produce implementations that are small, structurally simple, and therefore easy to understand?

A first step into this direction is *Bounded Synthesis* [9]. Here, we bound the number of states of the implementation and can therefore, by incrementally increasing the bound, ensure that the synthesized solution has minimal size.

In this paper, we go one step further by synthesizing implementations where, additionally, the number of (*simple*) cycles in the state graph is limited by a given bound. Reducing the number of cycles makes an implementation much easier to understand. Compare the three implementations of the TBURST4 component of the AMBA bus controller shown in Figure 1: standard synthesis with Acacia+ produces the state graph on the left with 14 states and 61 cycles. Bounded Synthesis produces the middle one with 7 states and 19 cycles. The graph on the right, produced by our tool, has 7 states and 7 cycles, which is the minimum.

An interesting aspect of the number of cycles as a parameter of the implementations is that the number of cycles that is potentially needed to satisfy an LTL specification explodes in the size of the specification: we show that there is a triple exponential lower and upper bound on the number of cycles that can be enforced by an LTL specification. The impact of the size of the specification on the number of cycles is thus even more dramatic than on the number of states, where the blow-up is double exponential.

Our synthesis algorithm is inspired by Tiernan’s cycle counting algorithm from 1970 [17]. Tiernan’s algorithm is based on exhaustive search. From some arbitrary vertex  $v$ , the graph is unfolded into a tree such that no vertices repeat on any branch. The number of vertices in the tree that are connected to  $v$  then corresponds to the number of cycles through  $v$  in the graph. Subsequently,  $v$  is removed from the graph, and the algorithm continues with one of the remaining vertices until the graph becomes empty. We integrate Tiernan’s algorithm into



**Fig. 2.** Witness for an example state graph with three cycles. The state graph is shown on the left. The first graph on the right proves that vertex 1 is on two cycles (via vertex 2 and vertices 2 and 3). The second graph proves that vertex 2 is on a further cycle not, containing vertex 1, namely via vertex 3. There are no further cycles through vertex 3.

the Bounded Synthesis approach. Bounded Synthesis uses a SAT-solver to simultaneously construct an implementation and a *witness* for the correctness of the implementation [9]. For the standard synthesis from an LTL specification  $\varphi$ , the witness is a finite graph which describes an accepting run of the universal tree automaton corresponding to  $\varphi$ . To extend the idea to Bounded Cycle Synthesis, we define a second witness that proves the number of cycles, as computed by Tiernan’s algorithm, to be equal to or less than the given bound. An example state graph with three cycles is shown on the left in Figure 2. The witness consists of the three graphs shown on the right in Figure 2. The first graph proves that vertex 1 is on two cycles (via vertex 2 and vertices 2 and 3). The second graph proves that vertex 2 is on a further cycle, not containing vertex 1, namely via vertex 3. There are no further cycles through vertex 3.

Our experiments show that Bounded Cycle Synthesis is comparable in performance to standard Bounded Synthesis. The specifications that can be handled by Bounded Cycle Synthesis are smaller than what can be handled by tools like Acacia+, but the quality of the synthesized implementations is much better. Bounded Cycle Synthesis could be used in a development process where the programmer decomposes the system into modules that are small enough so that the implementation can still be inspected comfortably by the programmer (and synthesized reasonably fast by using the Bounded Cycle Synthesis approach). Instead of manually writing the code for such a module, the programmer has the option of writing a specification, which is then automatically replaced by the best possible implementation.

## 2 Preliminaries

The non-negative integers are denoted by  $\mathbb{N}$ . An alphabet  $\Sigma$  is a non-empty finite set.  $\Sigma^\omega$  denotes the set of infinite words over  $\Sigma$ . If  $\alpha \in \Sigma^\omega$ , then  $\alpha_n$  accesses the  $n$ -th letter of  $\alpha$ , starting at  $\alpha_0$ . For the rest of the paper we assume  $\Sigma = 2^{\mathcal{I} \cup \mathcal{O}}$  to be partitioned into sets of input signals  $\mathcal{I}$  and output signals  $\mathcal{O}$ .

A *Mealy machine*  $\mathcal{M}$  is a tuple  $(\mathcal{I}, \mathcal{O}, T, t_I, \delta, \lambda)$  over input signals  $\mathcal{I}$  and output signals  $\mathcal{O}$ , where  $T$  is a finite set of states,  $t_I \in T$  is the initial state,

$\delta: T \times 2^{\mathcal{I}} \rightarrow T$  is the transition function, and  $\lambda: T \times 2^{\mathcal{I}} \rightarrow 2^{\mathcal{O}}$  is the output function. Thereby, the output only depends on the current state and the last input letter. The size of  $\mathcal{M}$ , denoted by  $|\mathcal{M}|$ , is defined as  $|T|$ . A path  $p$  of a Mealy machine  $\mathcal{M}$  is an infinite sequence  $p = (t_0, \sigma_0)(t_1, \sigma_1)(t_2, \sigma_2) \dots \in (T \times \Sigma)^\omega$  such that  $t_0 = t_I$ ,  $\delta(t_n, \mathcal{I} \cap \sigma_n) = t_{n+1}$  and  $\lambda(t_n, \mathcal{I} \cap \sigma_n) = \mathcal{O} \cap \sigma_n$  for all  $n \in \mathbb{N}$ . We use  $\pi_1(p) = \sigma_0\sigma_1\sigma_2 \dots \in \Sigma^\omega$ , to denote the projection of  $p$  to its second component.  $\mathcal{P}(\mathcal{M})$  denotes the set of all paths of a Mealy machine  $\mathcal{M}$ .

Specifications are given in *Linear-time Temporal Logic* (LTL). The atomic propositions of the logic consist of the signals  $\mathcal{I} \cup \mathcal{O}$ , resulting in the alphabet  $\Sigma = 2^{\mathcal{I} \cup \mathcal{O}}$ . The syntax of an LTL specification  $\varphi$  is defined as follows:

$$\varphi := \text{true} \mid a \in \mathcal{I} \cup \mathcal{O} \mid \neg\varphi \mid \varphi \vee \varphi \mid \bigcirc\varphi \mid \varphi \mathcal{U} \varphi$$

The size of a specification  $\varphi$  is denoted by  $|\varphi|$  and is defined to be the number of sub-formulas of  $\varphi$ . The semantics of LTL are defined over infinite words  $\alpha \in \Sigma^\omega$ . We define the satisfaction of a word  $\alpha$  at a position  $n \in \mathbb{N}$  and a specification  $\varphi$ , denoted by  $\alpha, n \models \varphi$ , for the different choices of  $\varphi$ , respectively, as follows:

- $\alpha, n \models \text{true}$
- $\alpha, n \models a$  iff  $a \in \alpha_i$
- $\alpha, n \models \neg\varphi$  iff  $\alpha, n \not\models \varphi$
- $\alpha, n \models \varphi_1 \vee \varphi_2$  iff  $\alpha, n \models \varphi_1$  or  $\alpha, n \models \varphi_2$
- $\alpha, n \models \bigcirc\varphi$  iff  $\alpha, n+1 \models \varphi$
- $\alpha, n \models \varphi_1 \mathcal{U} \varphi_2$  iff  $\exists m \geq n. \alpha, m \models \varphi_2$  and  $\forall n \leq i < m. \alpha, i \models \varphi_1$

An infinite word  $\alpha$  satisfies  $\varphi$ , denoted by  $\alpha \models \varphi$ , iff  $\alpha, 0 \models \varphi$ . The language  $\mathcal{L}(\varphi)$  is the set of all words that satisfy  $\varphi$ , i.e.,  $\mathcal{L}(\varphi) = \{\alpha \in \Sigma^\omega \mid \alpha \models \varphi\}$ . Beside the standard operators, we have the standard derivatives of the boolean operators, as well as  $\diamond\varphi \equiv \text{true} \mathcal{U} \varphi$  and  $\square\varphi \equiv \neg\diamond\neg\varphi$ . A Mealy machine  $\mathcal{M}$  is an implementation of  $\varphi$  iff  $\pi_1(\mathcal{P}(\mathcal{M})) \subseteq \mathcal{L}(\varphi)$ .

Let  $G = (V, E)$  be a directed graph. A (*simple*) *cycle*  $c$  of  $G$  is a tuple  $(C, \eta)$ , consisting of a non-empty set  $C \subseteq V$  and a bijection  $\eta: C \mapsto C$  such that

- $\forall v \in C. (v, \eta(v)) \in E$  and
- $\forall v \in C. n \in \mathbb{N}. \eta^n(v) = v \Leftrightarrow n \bmod |C| = 0$ ,

where  $\eta^n$  denotes  $n$  times the application of  $\eta$ . In other words, a cycle of  $G$  is a path through  $G$  that starts and ends at the same vertex and visits every vertex of  $V$  at most once. We say that a cycle  $c = (C, \eta)$  has length  $n$  iff  $|C| = n$ .

We extend the notion of a cycle of a graph  $G$  to Mealy machines  $\mathcal{M} = (\mathcal{I}, \mathcal{O}, T, t_I, \delta, \lambda)$ , such that  $c$  is a cycle of  $\mathcal{M}$  iff  $c$  is a cycle of the graph  $(T, E)$  for  $E = \{(t, t') \mid \exists \nu \in 2^{\mathcal{I}}. \delta(t, \nu) = t'\}$ . Thus, we ignore the input labels of the edges of  $\mathcal{M}$ . The set of all cycles of a Mealy machine  $\mathcal{M}$  is denoted by  $\mathcal{C}(\mathcal{M})$ .

A *universal co-Büchi automaton*  $\mathcal{A}$  is a tuple  $(\Sigma, Q, q_I, \Delta, R)$ , where  $\Sigma$  is the alphabet,  $Q$  is a finite set of states,  $q_0 \in Q$  is the initial state,  $\Delta \subseteq Q \times \Sigma \times Q$  is the transition relation and  $R \subseteq Q$  is the set of rejecting states. A run  $r = (q_0, \sigma_0)(q_1, \sigma_1)(q_2, \sigma_2) \dots \in (Q \times \Sigma)^\omega$  of  $\mathcal{A}$  is an infinite sequence such that  $q_0 = q_I$  and  $(q_n, \sigma_n, q_{n+1}) \in \Delta$  for all  $n \in \mathbb{N}$ . A run  $r$  is accepting if it has a suffix

$q_n q_{n+1} q_{n+2} \dots \in (Q \setminus R)^\omega$  for some  $n \in \mathbb{N}$ . An infinite word  $\alpha \in \Sigma^\omega$  is accepted by  $\mathcal{A}$  if all corresponding runs, i.e., all runs  $r = (q_0, \sigma_0)(q_1, \sigma_1)(q_2, \sigma_2) \dots$  with  $\alpha = \sigma_0 \sigma_1 \sigma_2 \dots$ , are accepting. The language  $\mathcal{L}(\mathcal{A})$  of  $\mathcal{A}$  is the set of all  $\alpha \in \Sigma^\omega$ , accepted by  $\mathcal{A}$ .

The *run graph*  $G$  of a universal co-Büchi automaton  $\mathcal{A} = (2^{\mathcal{I} \cup \mathcal{O}}, Q, q_I, \Delta, R)$  and a Mealy machine  $\mathcal{M} = (\mathcal{I}, \mathcal{O}, T, t_I, \delta, \lambda)$  is a directed graph  $G = (T \times Q, E)$ , with  $E = \{((t, q), (t', q')) \mid \exists \sigma. \delta(t, \mathcal{I} \cap \sigma) = t', \lambda(t, \mathcal{I} \cap \sigma) = \mathcal{O} \cap \sigma, (q, \sigma, q') \in \Delta\}$ . A vertex  $(t, q)$  of  $G$  is rejecting iff  $q \in R$ . A run graph is accepting iff there is no cycle of  $G$ , which contains a rejecting vertex. If the run graph is accepting, we say,  $\mathcal{M}$  is accepted by  $\mathcal{A}$ .

### 3 Bounds on the number of cycles

Our goal is to synthesize systems that have a simple structure. System quality most certainly has other dimensions as well, but structural simplicity is a property of interest for most applications.

The purpose of this section is to give theoretical arguments why the number of cycles is a good measure: we show that the number of cycles may explode even in cases where the number of states is small, and even if the specification enforces a large implementation, there may be a further explosion in the number of cycles. This indicates that bounding the number of cycles is important, if one wishes to have a structurally simple implementation. On the other hand, we observe that bounding the number of states alone is not sufficient in order to obtain a simple structure.

Similar observations apply to modern programming languages, which tend to be much better readable than transition systems, because their control constructs enforce a simple cycle structure. Standard synthesis techniques construct transition systems, not programs, and therefore lose this advantage. With our approach, we get closer to the control structure of a program, without being restricted to a specific programming language.

#### 3.1 Upper bounds

First, we show that the number of cycles of a Mealy machine  $\mathcal{M}$ , implementing an LTL specification  $\varphi$ , is bounded triply exponential in the size of  $\varphi$ . To this end, we first bound the number of cycles of an arbitrary graph  $G$  with bounded outdegree.

On graphs with arbitrary outdegree, the maximal number of cycles is given by a fully connected graph, where each cycle describes a permutation of states, and vice versa. Hence, using standard math we obtain an upper bound of  $2^{n \log n}$  cycles for a graph with  $n$  states. However, our proof uses a more involved argument to improve the bound even further down to  $2^{n \log(m+1)}$  for graphs with bounded outdegree  $m$ . Such an improvement is desirable, as for LTL the state graph explodes in the number of states, while the outdegree is constant in the number of input and output signals.

**Lemma 1.** *Let  $G = (V, E)$  be a directed graph with  $|V| = n$  and with maximal outdegree  $m$ . Then  $G$  has at most  $2^{n \log(m+1)}$  cycles.*

*Proof.* We show the result by induction over  $n$ . The base case is trivial, so let  $n > 1$  and let  $v \in V$  be some arbitrary vertex of  $G$ . By induction hypothesis, the subgraph  $G'$ , obtained from  $G$  by removing  $v$ , has at most  $2^{(n-1) \log(m+1)}$  cycles. Each of these cycles is also a cycle in  $G$ , thus it remains to consider the cycles of  $G$  containing  $v$ . In each of these remaining cycles,  $v$  has one of  $m$  possible successors in  $G'$  and from each such successor  $v'$  we have again  $2^{(n-1) \log(m+1)}$  possible cycles in  $G'$  returning to  $v'$ . Hence, if we *redirect* these cycles to  $v$  instead of  $v'$ , i.e., we insert  $v$  before  $v'$  in the cycle, then we cover all possible cycles of  $G$  containing  $v$ .<sup>1</sup> All together, we obtain an upper bound of  $2^{(n-1) \log(m+1)} + m \cdot 2^{(n-1) \log(m+1)} = 2^{n \log(m+1)}$  cycles in  $G$ .  $\square$

We obtain an upper bound on the number of cycles of a Mealy machine  $\mathcal{M}$ .

**Lemma 2.** *Let  $\mathcal{M}$  be a Mealy machine. Then  $|\mathcal{C}(\mathcal{M})| \in O(2^{|\mathcal{M}| \cdot |Z|})$ .*

*Proof.* The Mealy machine  $\mathcal{M}$  has an outdegree of  $2^{|Z|}$  and, thus, by Lemma 1, the number of cycles is bounded by  $2^{|\mathcal{M}| \log(2^{|Z|} + 1)} \in O(2^{|\mathcal{M}| \cdot |Z|})$ .  $\square$

Finally, we are able to derive an upper bound on the implementations realizing a LTL specification  $\varphi$ .

**Theorem 1.** *Let  $\varphi$  be a realizable LTL specification. Then there is a Mealy machine  $\mathcal{M}$ , realizing  $\varphi$ , with at most triply exponential many cycles in  $|\varphi|$ .*

*Proof.* From [15, 16, 9] we obtain a doubly exponential upper bound in  $|\varphi|$  on the size of  $\mathcal{M}$ . With that, applying Lemma 2 yields the desired result.  $\square$

### 3.2 Lower bounds

It remains to prove that the bound of Theorem 1 is tight. To this end, we show that for each  $n \in \mathbb{N}$  there is a realizable LTL specification  $\varphi$  with  $|\varphi| \in \Theta(n)$ , such that every implementation of  $\varphi$  has at least triply exponential many cycles in  $n$ . The presented proof is inspired by [1], where a similar argument is used to prove a lower bound on the distance of the longest path through a synthesized implementation  $\mathcal{M}$ . We start with a gadget, which we use to increase the number of cycles exponentially in the length of the longest cycle of  $\mathcal{M}$ .

**Lemma 3.** *Let  $\varphi$  be a realizable LTL specification, for which every implementation  $\mathcal{M}$  has a cycle of length  $n$ . Then there is a realizable specification  $\psi$ , such that every Mealy machine  $\mathcal{M}'$  implementing  $\psi$  contains at least  $2^n$  many cycles.*

<sup>1</sup> Note that not every such edge needs to exist for a concrete given graph. However, in our worst-case analysis, every possible cycle is accounted for.

*Proof.* Let  $a$  and  $b$  be a fresh input and output signals, respectively, which do not appear in  $\varphi$ , and let  $\mathcal{M} = (\mathcal{I}, \mathcal{O}, T, t_I, \delta, \lambda)$  be an arbitrary implementation of  $\varphi$ . We define  $\psi := \varphi \wedge \Box(a \leftrightarrow \bigcirc b)$  and construct the implementation  $\mathcal{M}'$  as

$$\mathcal{M}' = (\mathcal{I} \cup \{a\}, \mathcal{O} \cup \{b\}, T \times 2^{\{b\}}, (t_I, \emptyset), \delta', \lambda'),$$

where  $\lambda'((t, s), \nu) = \lambda(t, \mathcal{I} \cap \nu) \cup s$  and

$$\delta'((t, s), \nu) = \begin{cases} (\delta(t, \mathcal{I} \cap \nu), \emptyset) & \text{if } a \in \nu \\ (\delta(t, \mathcal{I} \cap \nu), \{b\}) & \text{if } a \notin \nu \end{cases}$$

We obtain that  $\mathcal{M}'$  is an implementation of  $\psi$ . The implementation remembers each input  $a$  for one time step and then outputs the stored value. Thus, it satisfies  $\Box(a \leftrightarrow \bigcirc b)$ . Furthermore,  $\mathcal{M}'$  still satisfies  $\varphi$ . Hence,  $\psi$  must be realizable, too.

Next, we pick an arbitrary implementation  $\mathcal{M}''$  of  $\psi$ , which must exist according to our previous observations. Then, after projecting away the fresh signals  $a$  and  $b$  from  $\mathcal{M}''$ , we obtain again an implementation for  $\varphi$ , which contains a cycle  $(C, \eta)$  of length  $n$ , i.e.,  $C = \{t_1, t_2, \dots, t_n\}$ . We obtain that  $\mathcal{M}''$  contains at least the cycles

$$\mathbb{C} = \{(\{(t_i, f(t_i)) \mid i \in \{1, 2, \dots, n\}\}, (t, s) \mapsto (\eta(t), f(\eta(t)))) \mid f: C \rightarrow 2^{\{b\}}\},$$

which concludes the proof, since  $|\mathbb{C}| = 2^n$ .  $\square$

Now, with Lemma 3 at hand, we are ready to show that the aforementioned lower bounds are tight. The final specification only needs the temporal operators  $\bigcirc$ ,  $\Box$  and  $\Diamond$ , i.e., the bound already holds for a restricted fragment of LTL.

**Theorem 2.** *For every  $n > 1$ , there is a realizable specification  $\varphi_n$  with  $|\varphi_n| \in \Theta(n)$ , for which every implementation  $\mathcal{M}_n$  has at least triply exponential many cycles in  $n$ .*

*Proof.* According to Lemma 3, it suffices to find a realizable  $\varphi_n$ , such that  $\varphi_n$  contains at least one cycle of length doubly exponential in  $n$ . We choose

$$\varphi_n ::= \Box \left( \underbrace{\bigcirc \bigwedge_{i=1}^n (a_i \rightarrow \bigcirc b_i)}_{\varphi_n^{prem}} \rightarrow \underbrace{\bigcirc \bigwedge_{i=1}^n (c_i \rightarrow \bigcirc d_i)}_{\varphi_n^{con}} \right) \leftrightarrow \Box \Diamond s$$

with  $\mathcal{I} = \mathcal{I}_a \cup \mathcal{I}_b \cup \mathcal{I}_c \cup \mathcal{I}_d$  and  $\mathcal{O} = \{s\}$ , where  $\mathcal{I}_x = \{x_1, x_2, \dots, x_n\}$ . The specification describes a monitor, which checks whether the invariant  $\bigcirc \varphi_n^{prem} \rightarrow \bigcirc \varphi_n^{con}$  over the input signals  $\mathcal{I}$  is satisfied or not. Thereby, satisfaction is signaled by the output  $s$ , which needs to be triggered infinitely often, as long as the invariant stays satisfied.

In the following, we denote a subset  $x \subseteq \mathcal{I}_x$  by the  $n$ -ary vector  $\vec{x}$  over  $\{0, 1\}$ , where the  $i$ -th entry of  $\vec{x}$  is set to 1 if and only if  $x_i \in x$ .

The specification  $\varphi_n$  is realizable. First, consider that to check the fulfillment of  $\varphi_n^{prem}$  ( $\varphi_n^{con}$ ), an implementation  $\mathcal{M}$  needs to store the set of all requests  $\vec{a}$  ( $\vec{c}$ ), whose 1-positions have not yet been released by a corresponding response  $\vec{b}$  ( $\vec{d}$ ). Furthermore, to monitor the complete invariant  $\diamond \varphi_n^{prem} \rightarrow \diamond \varphi_n^{con}$ ,  $\mathcal{M}$  has to guess at each point in time, whether  $\varphi_n^{prem}$  will be satisfied in the future (under the current request  $\vec{a}$ ), or not. To realize this guess,  $\mathcal{M}$  needs to store a mapping  $f$ , which maps each open request  $\vec{a}$  to the corresponding set of requests  $\vec{c}$ <sup>2</sup>. This way,  $\mathcal{M}$  can look up the set of requests  $\vec{c}$ , tracked since the last occurrence of  $\vec{a}$ , whenever  $\vec{a}$  gets released by a corresponding vector  $\vec{b}$ . If this is the case, it continues to monitor the satisfaction of  $\varphi_n^{con}$  (if not already satisfied) and finally adjusts the output signal  $s$ , correspondingly. Note that  $\mathcal{M}$  still has to continuously update and store the mapping  $f$ , since the next satisfaction of  $\varphi_n^{prem}$  may already start while the satisfaction of current  $\varphi_n^{con}$  is still checked. There are double exponential many such mappings  $f$ , hence,  $\mathcal{M}$  needs to be at least doubly exponential in  $n$ .

It remains to show that every such implementation  $\mathcal{M}$  contains a cycle of at least doubly exponential length. By the aforementioned observations, we can assign each state of  $\mathcal{M}$  a mapping  $f$ , that maps vectors  $\vec{a}$  to sets of vectors  $\vec{c}$ . By interpreting the vectors as numbers, encoded in binary, we obtain that  $f: \{1, 2, \dots, 2^n\} \mapsto 2^{\{1, 2, \dots, 2^n\}}$ . Next, we again map each such mapping  $f$  to a binary sequence  $b_f = b_0 b_1 \dots b_m \in \{0, 1\}^m$  with  $m = 2^n$ . Thereby, a bit  $b_i$  of  $b_f$  is set to 1 if and only if  $i \in f(i)$ . It is easy to observe, that if two binary sequences are different, then their related states have to be different as well.

To conclude the proof, we show that the environment has a strategy to manipulate the bits of associated sequences  $b_f$  via the inputs  $\mathcal{I}$ .

To set bit  $b_i$ , the environment chooses the requests  $\vec{a}$  and  $\vec{c}$  such that they represent  $i$  in binary. The remaining inputs are fixed to  $\vec{b} = \vec{d} = \vec{0}$ . Hence, all other bits are not affected, as possible requests of previous  $\vec{a}$  and  $\vec{c}$  remain open.


To reset bit  $b_i$ , the environment needs multiple steps. First, it picks  $\vec{a} = \vec{c} = \vec{d} = \vec{0}$  and  $\vec{b} = \vec{1}$ . This does not affect any bit of the sequence  $b_f$ , since all requests introduced through vectors  $\vec{c}$  are still open. Next, the environment executes the aforementioned procedure to set bit  $b_j$  for every bit currently set to 1, except for the bit  $b_i$ , it wants to reset. This refreshes the requests introduced by previous vectors  $\vec{a}$  for every bit, except for  $b_i$ . Furthermore, it does not affect the sequence  $b_f$ . Finally, the environment picks  $\vec{a} = \vec{b} = \vec{c} = \vec{0}$  and picks  $\vec{d}$  such that it represents  $i$  in binary. This removes  $i$  from every entry in  $f$ , but only resets  $b_i$ , since all other bits are still open due to the previous updates.

With these two operations, the environment can enforce any sequences of sequences  $b_f$ , including a binary counter counting up to  $2^{2^n}$ . As different states are induced by the different sequences, we obtain a cycle of doubly exponential length in  $n$  by resetting the counter at every overflow.  $\square$

---

<sup>2</sup> Our representation is open for many optimizations. However, they will not affect the overall complexity result. Thus, we ignore them for the sake of readability here.





**Fig. 3.** The Mealy automata  $\mathcal{M}_n$  (red/dotted) and  $\mathcal{M}'_n$  (blue/dashed). Solid edges are shared between both automata.

### 3.3 The trade-off between states and cycles

We conclude this section with some observations regarding tradeoffs between the problem of synthesizing implementations, which are minimal in the number of states, versus the problem of synthesizing implementations, which are minimal in the number of cycles. The main question we answer, is whether we can achieve both: minimality in the number of states and minimality in the number of cycles. Unfortunately, this is not possible, as shown by Theorem 3.

**Theorem 3.** *For every  $n > 1$ , there is a realizable LTL specification  $\varphi_n$  with  $|\varphi| \in \Theta(n)$ , such that*

- *there is an implementation of  $\varphi$  consisting of  $n$  states and*
- *there is an implementation of  $\varphi$  containing  $m$  cycles,*
- *but there is no implementation of  $\varphi$  with  $n$  states and  $m$  cycles.*

*Proof.* Consider the specification

$$\varphi_n = (\neg b \wedge c) \wedge \bigcirc^{k+2} (\neg b \wedge c) \wedge \bigwedge_{i=1}^k \bigcirc^i (\neg c \wedge \bigcirc \neg c \wedge (a \leftrightarrow \bigcirc b))$$

over  $\mathcal{I} = \{a\}$  and  $\mathcal{O} = \{b, c\}$ , where  $\bigcirc^i$  denotes  $i$  times the application of  $\bigcirc$ . The specification  $\varphi_n$  is realizable with at least  $n = 2k + 1$  states. The corresponding Mealy machine  $\mathcal{M}_n$  is depicted in Figure 3. However,  $\mathcal{M}_n$  has  $m = 2^k$  many cycles. This blowup can be avoided by spending the implementation at least one more state, which reduces the number of cycles to  $m = 1$ . The result  $\mathcal{M}'_n$  is also depicted in Figure 3.  $\square$

Our results show that the number of cycles can explode (even more so than the number of states), and that sometimes this explosion is unavoidable. However, the results also show that there are cases, where the cycle count can be improved by choosing a better structured solution. Hence, it is desirable to have better control over the number of cycles that appear in an implementation. In the remainder of the paper, we show how to achieve this control.

## 4 Bounding the Cycles

In this section, we show how to synthesize an implementation  $\mathcal{M}$  from a given LTL specification  $\varphi$ , while giving a guarantee on the size and the number of cycles of  $\mathcal{M}$ . We first show how to guarantee a bound on the number of states of  $\mathcal{M}$ , by reviewing the classical Bounded Synthesis approach. Our encoding uses Mealy machines as implementations, and Boolean Satisfiability (SAT) as the underlying constraint system.

We then review the classical algorithm to count the cycles of  $\mathcal{M}$  and show how this algorithm gets embedded into a constraint system, such that we obtain a guarantee on the number of cycles of  $\mathcal{M}$ .

### 4.1 Bounded Synthesis

In the bounded synthesis approach [9], we first translate a given LTL specification  $\varphi$  into an equivalent universal co-Büchi automaton  $\mathcal{A}$ , such that  $\mathcal{L}(\mathcal{A}) = \mathcal{L}(\varphi)$ . Thus, we reduce the problem to finding an implementation  $\mathcal{M}$  that is accepted by  $\mathcal{A}$ , i.e., we look for an implementation  $\mathcal{M}$  such that the run graph of  $\mathcal{M}$  and  $\mathcal{A}$  contains no cycle with a rejecting vertex. This property is witnessed by a ranking function, which annotates each vertex of  $G$  by a natural number that bounds the number of possible visits to rejecting states. The annotation itself is bounded by  $n \cdot k$ , where  $n$  is the size bound on  $\mathcal{M}$  and  $k$  denotes the number of rejecting states of  $\mathcal{A}$ .

Fix some set of states  $T$  with  $|T| = n$  and let  $\mathcal{A} = (2^{\mathcal{I} \cup \mathcal{O}}, Q, q_I, \Delta, R)$ . Then, to guess a solution within SAT, we introduce the following variables:

- $\text{TRANS}(t, \nu, t')$  for all  $t, t' \in T$  and  $\nu \in 2^{\mathcal{I}}$ , for the transition relation of  $\mathcal{M}$ .
- $\text{LABEL}(t, \nu, x)$  for all  $t \in T, \nu \in 2^{\mathcal{I}}$  and  $x \in \mathcal{O}$ , for the labels of each transition.
- $\text{RGSTATE}(t, q)$  for all  $t \in T$  and  $q \in Q$ , to denote the reachable states of the run graph  $G$  of  $\mathcal{M}$  and  $\mathcal{A}$ . Only reachable states have to be annotated.
- $\text{ANNOTATION}(t, q, i)$  for all  $t \in T, q \in Q$  and  $0 < i \leq \log(n \cdot k)$ , denoting the annotation of a state  $(t, q)$  of  $G$ . Thereby, we use a logarithmic number of bits to encode the annotated value in binary. We use  $\text{ANNOTATION}(t, q) \circ m$  for  $\circ \in \{<, \leq, =, \geq, >\}$ , to denote an appropriate encoding of the relation of the annotation to some value  $m$  or other annotations  $\text{ANNOTATION}(t', q')$ .

Given a universal co-Büchi automaton  $\mathcal{A}$  and a bound  $n$  on the states of the resulting implementation, we encode the Bounded Synthesis problem via the SAT formula  $\mathcal{F}_{BS}(\mathcal{A}, n)$ , consisting of the following constraints:

- The target of every transition is unambiguous:

$$\bigwedge_{t \in T, \nu \in 2^{\mathcal{I}}} \text{exactlyOne}(\{\text{TRANS}(t, \nu, t') \mid t' \in T\})$$

where  $\text{exactlyOne}: X \mapsto \mathbb{B}(X)$  returns a SAT query, which ensures that among all variables of the set  $X$  exactly one is *true* and all others are *false*.

- The initial state  $(t_I, q_I)$  of the run graph for some arbitrary, but fix,  $t_I \in T$  is reachable and annotated by one. Furthermore, all annotations are bounded by  $n \cdot k$ :

$$\text{RGSTATE}(1, 1) \wedge \text{ANNOTATION}(1, 1) = 1 \wedge \bigwedge_{t \in T, q \in Q} \text{ANNOTATION}(t, q) \leq n \cdot k$$

- Each annotation of a vertex of the run graph bounds the number of visited accepting states, not counting the current vertex itself:

$$\bigwedge_{t \in T, q \in Q} \text{RGSTATE}(t, q) \rightarrow \bigwedge_{\sigma \in 2^X} \text{label}(t, \sigma) \rightarrow \bigwedge_{t' \in T} \text{TRANS}(t, \mathcal{I} \cap \sigma, t') \rightarrow$$

$$\bigwedge_{q' \in \Delta(q, \sigma)} \text{RGSTATE}(t', q') \wedge \text{ANNOTATION}(t, q) \prec_q \text{ANNOTATION}(t', q')$$

where  $\prec_q$  equals  $<$  if  $q \in R$  and equals  $\leq$  otherwise. Furthermore, we use the function  $\text{label}(t, \sigma)$  to fix the labeling of each transition, i.e.,  $\text{label}(t, \sigma) = \bigwedge_{x \in \mathcal{O} \cap \sigma} \text{LABEL}(t, \mathcal{I} \cap \sigma, x) \wedge \bigwedge_{x \in \mathcal{O} \setminus \sigma} \neg \text{LABEL}(t, \mathcal{I} \cap \sigma, x)$ .

**Theorem 4 (Bounded Synthesis [9]).** *For each bound  $n \in \mathbb{N}$  and each universal co-Büchi automaton  $\mathcal{A}$ , the SAT formula  $\mathcal{F}_{BS}(\mathcal{A}, n)$  is satisfiable if and only if there is a Mealy machine  $\mathcal{M}$  with  $|\mathcal{M}| = n$ , which is accepted by  $\mathcal{A}$ .*

## 4.2 Counting Cycles

Before we bound the number of cycles of a Mealy machine  $\mathcal{M}$ , we review Tienan’s classical algorithm [17] to count the number of cycles of a directed graph  $G$ . The algorithm not only gives insights into the complexity of the problem, but also contains many inspirations for our latter approach.

*Algorithm 1.* Given a directed graph  $G = (V, E)$ , we count the cycles of  $G$  using the following algorithm:

- (1) Initialize the cycle counter  $c$  to  $c := 0$  and some set  $P$  to  $P := \emptyset$ .
- (2) Pick some arbitrary vertex  $v_r$  of  $G$ , set  $v := v_r$  and  $P := \{v_r\}$ .
- (3) For all edges  $(v, v') \in E$ , with  $v' \notin P \setminus \{v_r\}$ :
  - (3a) If  $v' = v_r$ , increase  $c$  by one.
  - (3b) Otherwise, add  $v'$  to  $P$  and recursively execute (3). Afterwards, reset  $P$  to its value before the recursive call.
- (4) Obtain the sub-graph  $G'$ , by removing  $v_r$  from  $G$ :
  - (4a) If  $G'$  is empty, return  $c$ .
  - (4b) Otherwise, continue from (2) with  $G'$ .

The algorithm starts by counting all cycles that contain the first picked vertex  $v_r$ . This is done by an unfolding of the graph into a tree, rooted in  $v_r$ , such that there is no repetition of a vertex on any path from the root to a leaf. The number of vertices that are connected to the root by an edge of  $E$  then represents the corresponding number of cycles through  $v_r$ . The remaining cycles of  $G$  do not contain  $v_r$  and, thus, are cycles of the sub-graph  $G'$  without  $v_r$ , as well. Hence, we count the remaining cycles by recursively counting the cycles of  $G'$ . The algorithm terminates as soon as  $G'$  gets empty.

The algorithm is correct [17], but has the drawback, that the unfolded trees, may become exponential in the size of the graph, even if none of their vertices is connected to the root, i.e., even if there is no cycle to be counted. For an example consider the induced graph of  $\mathcal{M}'_n$ , as depicted in Figure 3. However, this drawback can be avoided by first reducing the graph to all its strongly connected components (SCCs) and then by counting the cycles of each SCC separately [18, 13]. A cycle never leaves an SCC of the graph.

As a result, we obtain an improved algorithm, which is exponential in the size of  $G$ , but linear in the number of cycles  $m$ . Furthermore, the time between two detections of a cycle, during the execution, is bounded linear in the size of the graph  $G$ .

### 4.3 Bounded Cycle Synthesis

We combine the insights of the previous sections to obtain a synthesis algorithm, which not only bounds the number of states of the resulting implementation  $\mathcal{M}$  but also bounds the number of cycles of  $\mathcal{M}$ . We use the unfolded trees from the previous section as our witnesses.

We call a tree that witnesses  $m$  cycles in  $G$ , all containing the root  $r$  of the tree, a witness-tree  $\mathcal{T}_{r,m}$  of  $G$ . Formally, a *witness-tree*  $\mathcal{T}_{r,m}$  of  $G = (V, E)$  is a labeled graph  $\mathcal{T}_{r,m} = ((W, B \cup R), \tau)$ , consisting of a graph  $(W, B \cup R)$  with  $m = |R|$  and a labeling function  $\tau: W \rightarrow V$ , such that:

1. The edges are partitioned into blue edges  $B$  and red edges  $R$ .
2. All red edges lead back to the root:

$$R \subseteq W \times \{r\}$$

3. No blue edges lead back to the root:

$$B \cap W \times \{r\} = \emptyset$$

4. Each non-root has at least one blue incoming edge:

$$\forall w' \in W \setminus \{r\}. \exists w \in W. (w, w') \in B$$

5. Each vertex has at most one blue incoming edge:

$$\forall w_1, w_2, w \in W. (w_1, w) \in B \wedge (w_2, w) \in B \Rightarrow w_1 = w_2$$

6. The graph is labeled by an unfolding of  $G$ :

$$\forall w, w' \in B \cup R. (\tau(w), \tau(w')) \in E,$$

7. The unfolding is complete:

$$\forall w \in W. \forall v' \in V. (\tau(w), v') \in E \Rightarrow \exists w' \in W. (w, w') \in B \cup R \wedge \tau(w') = v'$$

8. Let  $w_i, w_j \in W$  be two different vertices that appear on a path from the root to a leaf in the  $r$ -rooted tree  $(W, B)$ <sup>3</sup>. Then the labeling of  $w_i$  and  $w_j$  differs, i.e.,  $\tau(v_i) \neq \tau(v_j)$ .

9. The root of the tree is the same as the corresponding vertex of  $G$ , i.e.,  $\tau(r) = r$ .

**Lemma 4.** *Let  $G = (V, E)$  be a graph consisting of a single SCC,  $r \in V$  be some vertex of  $G$  and  $m$  be the number of cycles of  $G$  containing  $r$ . Then there is a witness-tree  $\mathcal{T}_{r,m} = ((W, B \cup R), \tau)$  of  $G$  with  $|W| \leq m \cdot |V|$ .*

*Proof.* We construct  $\mathcal{T}_{r,m}$  according to the strategy induced by *Algorithm 1*, where an edge is colored red if and only if it leads back to the root. The constructed tree satisfies all conditions 1 – 9. By correctness of *Algorithm 1*, we have that  $|R| = m$ .

Now, for the sake of contradiction, assume  $|W| > m \cdot |V|$ . First we observe, that the depth of the tree  $(W, B)$  must be bounded by  $|V|$  to satisfy Condition 8. Hence, as there are at most  $m$  red edges in  $\mathcal{T}_{r,m}$ , there must be a vertex  $w \in W$  without any outgoing edges. However, since  $G$  is a single SCC, this contradicts the completeness of  $\mathcal{T}_{r,m}$  (Condition 7).  $\square$

**Lemma 5.** *Let  $G = (V, E)$  be a graph consisting of a single SCC and let  $\mathcal{T}_{r,m}$  be a witness-tree of  $G$ . Then there are at most  $m$  cycles in  $G$  that contain  $r$ .*

*Proof.* Let  $\mathcal{T}_{r,m} = ((W, R \cup B), \tau)$ . Assume for the sake of contradiction that  $G$  has more than  $m$  cycles and let  $c = (C, \eta)$  be an arbitrary such cycle. By the completeness of  $\mathcal{T}_{r,m}$ , there is path  $w_0 w_1 \dots w_{|C|-1}$  with  $w_0 = r$  and  $\tau(w_i) = \eta^i(r)$  for all  $0 \leq i < |C|$ . From  $w_i \neq r$  and Condition 2, it follows  $(w_{i-1}, w_i) \in B$  for all  $0 < i < |C|$ . Further,  $\eta^{|C|}(r) = r$  and thus  $(w_{|C|-1}, w_0) \in R$ . Hence, by the tree shape of  $(W, B)$ , we get  $|R| > m$ , yielding the desired contradiction.  $\square$

From Lemma 4 and 5 we derive that  $\mathcal{T}_{r,m}$  is a suitable witness to bound the number of cycles of an implementation  $\mathcal{M}$ . Furthermore, from Lemma 4 we also obtain an upper bound on the size of  $\mathcal{T}_{r,m}$ .

We proceed with our final encoding. Therefore, we first construct a simple directed graph  $G$  out of the implementation  $\mathcal{M}$ . Then, we guess all the sub-graphs, obtained from  $G$  via iteratively removing vertices, and split them into their corresponding SCCs. Finally, we guess the witness-tree for each such SCC.

To keep the final SAT encoding compact, we even introduce some further optimizations. First, we do not need to introduce a fresh copy for each SCC, since the SCC of a vertex is always unique. Thus, it suffices to guess an annotation for each vertex, being unique for each SCC. Second, we have to guess  $n$  trees  $\mathcal{T}_{i,r_i}$ , each one consisting of at most  $i \cdot n$  vertices, such that the sum of all  $i$  is equal to the

<sup>3</sup> Note that the tree property is enforced by Conditions 3 – 5.

overall number of cycles  $m$ . One possible solution would be to overestimate each  $i$  by  $m$ . Another possibility would be to guess the exact distribution of the cycles over the different witness-trees  $\mathcal{T}_{i,r_i}$ . However, there is a smarter solution: we guess all trees together in a single graph bounded by  $m \cdot n$ . Additionally, to avoid possible interleavings, we add an annotation of each vertex by its corresponding witness-tree  $\mathcal{T}_{i,r_i}$ . Hence, instead of bounding the number of each  $\mathcal{T}_{i,r_i}$  separately by  $i$ , we just bound the number of all red edges in the whole forest by  $m$ . This way, we not only reduce the size of the encoding, but also skip the additional constraints, which would be necessary to sum the different witness-tree bounds  $i$  to  $m$ , otherwise.

Let  $T$  be some ordered set with  $|T| = n$  and  $S = T \times \{1, 2, \dots, m\}$ . We use  $T$  to denote the vertices of  $G$  and  $S$  to denote the vertices of the forest of  $\mathcal{T}_{i,r_i}$  s. Further, we use  $M = T \times \{1\}$  to denote the roots and  $N = S \setminus M$  to denote the non-roots of the corresponding trees. We introduce the following variables:

- $\text{EDGE}(t, t')$  for all  $t, t' \in T$ , denoting the edges of the abstraction of  $\mathcal{M}$  to  $G$ .
- $\text{BEDGE}(s, s')$  for all  $s \in S$  and  $s' \in N$ , denoting a blue edge.
- $\text{REDGE}(s, s')$  for all  $s \in S$  and  $s' \in M$ , denoting a red edge.
- $\text{WTREE}(s, i)$  for all  $s \in S$ ,  $0 < i \leq \log n$ , denoting the witness-tree of each  $s$ . As before, we use  $\text{WTREE}(s) \circ x$  to relate values with the underlying encoding.
- $\text{VISITED}(s, t)$  for all  $s \in S$  and  $t \in T$ , denoting the set of all vertices  $t$ , already visited at  $s$ , since leaving the root of the corresponding witness-tree.
- $\text{RBOUND}(c, i)$  for all  $0 < c \leq m$ ,  $0 < i \leq \log(n \cdot m)$ , denoting an ordered list of all red edges, bounding the red edges of the forest.
- $\text{SCC}(k, t, i)$  for all  $0 < k \leq n$ ,  $t \in T$ , and  $0 \leq i < \log n$ , denoting the SCC of  $t$  in the  $k$ -th sub-graph of  $G$ . The sub-graphs are obtained by iteratively removing vertices of  $T$ , according to the pre-defined order. This way, each sub-graph contains exactly all vertices that are larger than the root.

Note that by the definition of  $S$  we introduce  $m$  explicit copies for each vertex of  $G$ . This is sufficient, since each cycle contains each vertex at most once. Thus, the labeling  $\tau$  of a vertex  $s$  can be directly derived from the first component of  $s$ .

Given a universal co-Büchi automaton  $\mathcal{A}$ , a bound  $n$  on the states of the resulting implementation  $\mathcal{M}$ , and a bound  $m$  on the number of cycles of  $\mathcal{M}$ , we encode the Bounded Cycle Synthesis problem via the SAT formula  $\mathcal{F}_{BS}(\mathcal{A}, n) \wedge \mathcal{F}_{CS}(\mathcal{A}, n, m) \wedge \mathcal{F}_{SCC}(n)$ . The constraints of  $\mathcal{F}_{CS}(\mathcal{A}, n, m)$ , bounding the cycles of the system, are given by Table 1. The constraints of  $\mathcal{F}_{SCC}(n)$ , enforcing that each vertex is labeled by a unique SCC, can be found in the technical report [8].

**Theorem 5.** *For each pair of bounds  $n, m \in \mathbb{N}$  and each universal co-Büchi automaton  $\mathcal{A}$  with  $|\mathcal{A}| = k$ , the formula  $\mathcal{F} = \mathcal{F}_{BS}(\mathcal{A}, n) \wedge \mathcal{F}_{CS}(\mathcal{A}, n, m) \wedge \mathcal{F}_{SCC}$  is satisfiable if and only if there is a Mealy machine  $\mathcal{M}$  with  $|\mathcal{M}| = n$  and  $|\mathcal{C}(\mathcal{M})| = m$ , accepted by  $\mathcal{A}$ . Furthermore,  $\mathcal{F}$  consists of  $x$  variables with  $x \in O(n^3 + n^2(m^2 + 2^{|\Sigma|}) + n|\mathcal{O}| + nk \log(nk))$  and  $|\mathcal{F}| \in O(n^3 + n^2(m^2 + k|\Sigma|))$ .*

**Table 1.** Constraints of the SAT formula  $\mathcal{F}_{CS}(\mathcal{A}, n, m)$ .

$\bigwedge_{t,t' \in T, \nu \in 2^I}$	$\text{TRANS}(t, \nu, t') \rightarrow \text{EDGE}(t, t')$	Construction of $G$ from $\mathcal{M}$ .
$\bigwedge_{t,t' \in T}$	$\text{EDGE}(t, t') \rightarrow \bigvee_{\nu \in 2^I} \text{TRANS}(t, \nu, t')$	
$\bigwedge_{r \in T}$	$\text{WTREE}((r, 1)) = r$	Roots indicate the witness-tree.
$\bigwedge_{s \in S, (r,1) \in M}$	$\text{REDGE}(s, (r, 1)) \rightarrow \text{WTREE}(s) = r$	Red edges only connect vertices of the current $\mathcal{T}_{i,r_i}$ .
$\bigwedge_{s \in S, s' \in N}$	$\text{BEDGE}(s, s') \rightarrow \text{WTREE}(s) = \text{WTREE}(s')$	Blue edges only connect vertices of the current $\mathcal{T}_{i,r_i}$ .
$\bigwedge_{s' \in N}$	$\text{exactlyOne}(\{ \text{BEDGE}(s, s') \mid s \in S \})$	Every non-root has exactly one blue incoming edge.
$\bigwedge_{(t,c) \in S, r \in T}$	$\text{REDGE}((t, c), (r, 1)) \rightarrow \text{EDGE}(t, r)$	Red edges are related to the edges of the graph $G$ .
$\bigwedge_{(t,c) \in S, (t',c') \in N}$	$\text{BEDGE}((t, c), (t', c')) \rightarrow \text{EDGE}(t, t')$	Blue edges are related to the edges of the graph $G$ .
$\bigwedge_{\substack{(t,c) \in S, r \in T, \\ t \geq r}}$	$\text{EDGE}(t, r) \wedge \text{SCC}(r, t) = \text{SCC}(r, r) \wedge \text{WTREE}((t, c)) = r \rightarrow \text{REDGE}((t, c), (r, 1))$	Every possible red edge must be taken.
$\bigwedge_{\substack{(t,c) \in S, r, t' \in T, \\ t \geq t'}}$	$\text{EDGE}(t, t') \wedge \text{SCC}(r, t) = \text{SCC}(r, t') \wedge \text{WTREE}((t, c)) = r \wedge \text{VISITED}((t, c), t') \rightarrow \bigvee_{0 < c' \leq m} \text{BEDGE}((t, c), (t', c'))$	Every possible blue edge must be taken.
$\bigwedge_{r \in T}$	$\bigwedge_{t \leq r} \neg \text{VISITED}((r, 1), t) \wedge \bigwedge_{t > r} \text{VISITED}((r, 1), t)$	Only non-roots of the corresponding sub-graph can be successors of a root.
$\bigwedge_{(t,c) \in S, s \in N}$	$\text{BEDGE}((t, c), s) \rightarrow \neg \text{VISITED}(s, t) \wedge (\text{VISITED}(s, t') \leftrightarrow \text{VISITED}((t, c), t'))$	Every vertex appears at most once on a path from the root to a leaf.
$\bigwedge_{s \in S, s' \in M}$	$\text{REDGE}(s, s') \rightarrow \bigvee_{0 < c \leq m} \text{RBOUND}(c) = f(s)$	The list of red edges is complete. ( $f(s)$ maps each state of $S$ to a unique number in $\{1, \dots, n \cdot m\}$ )
$\bigwedge_{0 < c \leq m}$	$\text{RBOUND}(c) < \text{RBOUND}(c + 1)$	Red edges are strictly ordered.

## 5 Experimental Results

We have implemented the Bounded Cycle Synthesis approach in our tool *BoWSer*, the Bounded Witness Synthesizer, and compared it against standard Bounded Synthesis and *Acacia+* (v2.3) [6, 7]. To ensure a common encoding, we used *BoWSer* for both, the Bounded Synthesis and the Bounded Cycle Synthesis approach. Our tool uses *LTL3BA* (v1.0.2) [3] to convert specifications to universal

**Table 2.** Results of the tools *LTL3BA*, *Aca(cia)+* and *BoWSer*. The *LTL3BA* tool was used to generate the universal co-Büchi tree automata  $\mathcal{A}_{uct}$ . The Bo(unded) Sy(nthesis) and Bo(unded) Cy(cle Synthesis) encodings were generated by *BoWSer*.

Benchmark	Size			Cycles			Time (s)				
	$\mathcal{A}_{uct}$	Aca+	BoSy/ BoCy	Aca+	BoSy	BoCy	Aca+	SAT		UNSAT	
								BoSy	BoCy	BoSy	BoCy
ARBITER[2]	6	26	2	5439901	3	3	0.261	0.847	0.868	0.300	0.836
ARBITER[3]	20	111	3	> 9999999	8	4	0.511	9.170	9.601	3.916	9.481
ARBITER[4]	64	470	4	> 9999999	8	5	12.981	105.527	109.180	56.853	106.803
LOCK[2]	12	4	3	12	6	5	0.459	0.395	0.522	0.165	0.487
LOCK[3]	20	4	3	12	5	5	55.917	1.037	1.245	0.433	1.107
LOCK[4]	36	–	3	–	6	5	> 999	4.419	4.761	1.407	3.726
ENCODE[2]	3	6	2	41	3	3	0.473	0.071	0.089	0.048	0.084
ENCODE[3]	5	16	3	90428	8	8	1.871	0.292	0.561	0.200	0.503
ENCODE[4]	5	20	4	> 9999999	24	24	4.780	1.007	16.166	0.579	> 999
DECODE	1	4	1	8	1	1	0.328	0.055	0.051	–	–
SHIFT	3	6	2	31	3	3	0.387	0.060	0.072	0.041	0.071
TBURST4	103	14	7	61	19	7	0.634	8.294	206.604	6.261	> 999
TINCR	43	5	3	7	5	2	0.396	2.262	2.279	0.845	2.221
TSINGLE	22	8	4	12	5	4	0.372	1.863	2.143	1.165	2.067

co-Büchi automata. The created SAT queries are solved by *MiniSat* (v.2.2.0) [5] and *clasp* (v.3.1.4) [10], where the result of the faster solver is taken.

The benchmarks are given in TLSF [11] and represent a decomposition of ARM’s *Advanced Microcontroller Bus Architecture* (AMBA) [2]. They are created from the assumptions and guarantees presented in [12], which were split into modules, connected by new signals. A detailed description of the benchmarks is given in [11].

All experiments were executed on a Unix machine, operated by a 64-bit kernel (v4.1.12) running on an Intel Core i7 with 2.8GHz and 8GB RAM. Each experiment had a time limit of 1000 seconds and a memory limit of 8GB. When counting cycles of a solution, the limit was set to 10000000 cycles.

The results of the evaluation are shown in Table 2, which displays the sizes of the intermediate universal co-Büchi tree automata  $\mathcal{A}_{uct}$ , the sizes of the synthesized implementations  $\mathcal{M}$ , the number of cycles of each implementation  $\mathcal{M}$ , and the overall synthesis time. Thereby, for each instance, we guessed the minimal number of states for the Bounded Synthesis approach and, additionally, the minimal number of cycles for the Bounded Cycle Synthesis approach, to obtain a single satisfiable instance. Further, to verify the result, we also created the unsatisfiable instance, where the state bound was decreased by one in the case of Bounded Synthesis and the cycle bound was decreased by one in the case of Bounded Cycle Synthesis. Note that these two instances already give an almost complete picture, since for increased and decreased bounds the synthesis times



behave monotonically. Hence, increasing the bound beyond the first realizable instance increases the synthesis time. Decreasing it below the last unsatisfiable instance decreases the synthesis time. The results for the TBURST4 component are additionally depicted in Figure 1.

On most benchmarks, Acacia+ solves the synthesis problem the fastest, followed by Bounded Synthesis and our approach. (On some benchmarks, Bounded Synthesis outperforms Acacia+.) Comparing the running times of Bounded Synthesis and Bounded Cycle Synthesis, the overhead for bounding the number of cycles is insignificant on most benchmarks. The two exceptions are ENCODE, which requires a fully connected implementation, and TBURST4, where the reduction in the number of cycles is substantial. In terms of states and cycles, our tool outperforms Bounded Synthesis on half of the benchmarks and it outperforms Acacia+ on all benchmarks.

The results of Acacia+ show that the number of cycles is indeed an explosive factor. However, they also show that this explosion can be avoided effectively.

## 6 Conclusions

We have introduced the Bounded Cycle Synthesis problem, where we limit the number of cycles in an implementation synthesized from an LTL specification. Our solution is based on the construction of a witness structure that limits the number of cycles. The existence of such a witness can be encoded as a SAT problem. Our experience in applying Bounded Cycle Synthesis to the synthesis of the AMBA bus arbiter shows that the approach leads to significantly better implementations. Furthermore, the performance of our prototype implementation is sufficient to synthesize the components (in a natural decomposition of the specification) in reasonable time.

Both Bounded Synthesis and Bounded Cycle Synthesis can be seen as the introduction of structure into the space of implementations. Bounded Synthesis structures the implementations according to the number of states, Bounded Cycle Synthesis additionally according to the number of cycles. The double exponential blow-up between the size of the specification and the number of states, and the triple exponential blow-up between the size and the number of cycles indicate that, while both parameters provide a fine-grained structure, the number of cycles may even be the superior parameter. Formalizing this intuition and finding other useful parameters is a challenge for future work.

Our method does not lead to a synthesis algorithm in the classical sense, where just a specification is given and an implementation or an unsatisfiability result is returned. In our setting, the bounds are part of the input, and have to be determined beforehand. In Bounded Synthesis, the bound is usually eliminated by increasing the bound incrementally. With multiple bounds, the choice which parameter to increase becomes non-obvious. Finding a good strategy for this problem is a challenge on its own and beyond the scope of this paper. We leave it open for future research.

## References

1. Alur, R., La Torre, S.: Deterministic Generators and Games for LTL Fragments. *ACM Trans. Comput. Log.* 5(1), 1–25 (2004), <http://doi.acm.org/10.1145/963927.963928>
2. ARM Ltd.: AMBA Specification (rev. 2) (1999), available at [www.arm.com](http://www.arm.com)
3. Babiak, T., Kretínský, M., Reháč, V., Strejcek, J.: LTL to Büchi Automata Translation: Fast and More Deterministic. In: Flanagan, C., König, B. (eds.) *Tools and Algorithms for the Construction and Analysis of Systems - 18th International Conference, TACAS 2012, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2012, Tallinn, Estonia, March 24 - April 1, 2012. Proceedings. Lecture Notes in Computer Science*, vol. 7214, pp. 95–109. Springer (2012), [http://dx.doi.org/10.1007/978-3-642-28756-5\\_8](http://dx.doi.org/10.1007/978-3-642-28756-5_8)
4. Bloem, R., Chatterjee, K., Henzinger, T.A., Jobstmann, B.: Better Quality in Synthesis through Quantitative Objectives. In: Bouajjani, A., Maler, O. (eds.) *Computer Aided Verification, 21st International Conference, CAV 2009, Grenoble, France, June 26 - July 2, 2009. Proceedings. Lecture Notes in Computer Science*, vol. 5643, pp. 140–156. Springer (2009), [http://dx.doi.org/10.1007/978-3-642-02658-4\\_14](http://dx.doi.org/10.1007/978-3-642-02658-4_14)
5. Eén, N., Sörensson, N.: An Extensible SAT-solver. In: Giunchiglia, E., Tacchella, A. (eds.) *Theory and Applications of Satisfiability Testing, 6th International Conference, SAT 2003, Santa Margherita Ligure, Italy, May 5-8, 2003 Selected Revised Papers. Lecture Notes in Computer Science*, vol. 2919, pp. 502–518. Springer (2003), [http://dx.doi.org/10.1007/978-3-540-24605-3\\_37](http://dx.doi.org/10.1007/978-3-540-24605-3_37)
6. Filiot, E., Jin, N., Raskin, J.: Antichains and Compositional Algorithms for LTL Synthesis. *Formal Methods in System Design* 39(3), 261–296 (2011), <http://dx.doi.org/10.1007/s10703-011-0115-3>
7. Filiot, E., Jin, N., Raskin, J.: Exploiting structure in LTL synthesis. *STTT* 15(5-6), 541–561 (2013), <http://dx.doi.org/10.1007/s10009-012-0222-5>
8. Finkbeiner, B., Klein, F.: Bounded Cycle Synthesis. *CoRR abs/1605.01511* (2016), <http://arxiv.org/abs/1605.01511>
9. Finkbeiner, B., Schewe, S.: Bounded synthesis. *STTT* 15(5-6), 519–539 (2013), <http://dx.doi.org/10.1007/s10009-012-0228-z>
10. Gebser, M., Kaufmann, B., Neumann, A., Schaub, T.: *clasp* : A Conflict-Driven Answer Set Solver. In: Baral, C., Brewka, G., Schlipf, J.S. (eds.) *Logic Programming and Nonmonotonic Reasoning, 9th International Conference, LPNMR 2007, Tempe, AZ, USA, May 15-17, 2007, Proceedings. Lecture Notes in Computer Science*, vol. 4483, pp. 260–265. Springer (2007), [http://dx.doi.org/10.1007/978-3-540-72200-7\\_23](http://dx.doi.org/10.1007/978-3-540-72200-7_23)
11. Jacobs, S., Klein, F.: A High-Level LTL Synthesis Format: TLSF v1.1 (Extended Version). *CoRR abs/1604.02284* (2016), <http://arxiv.org/abs/1604.02284>
12. Jobstmann, B.: Applications and Optimizations for LTL Synthesis. Ph.D. thesis, Graz University of Technology (Mar 2007)
13. Johnson, D.B.: Finding All the Elementary Circuits of a Directed Graph. *SIAM J. Comput.* 4(1), 77–84 (1975), <http://dx.doi.org/10.1137/0204007>
14. Kupferman, O.: Recent Challenges and Ideas in Temporal Synthesis. In: Bieliková, M., Friedrich, G., Gottlob, G., Katzenbeisser, S., Turán, G. (eds.) *SOFSEM 2012: Theory and Practice of Computer Science - 38th Conference on Current Trends in Theory and Practice of Computer Science, Špindlerův Mlýn, Czech Republic, January 21-27, 2012. Proceedings. Lecture Notes in Computer Science*, vol. 7147, pp. 88–98. Springer (2012), [http://dx.doi.org/10.1007/978-3-642-27660-6\\_8](http://dx.doi.org/10.1007/978-3-642-27660-6_8)

15. Kupferman, O., Vardi, M.Y.: Safraless Decision Procedures. In: 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2005), 23-25 October 2005, Pittsburgh, PA, USA, Proceedings. pp. 531–542. IEEE Computer Society (2005), <http://dx.doi.org/10.1109/SFCS.2005.66>
16. Piterman, N.: From Nondeterministic Büchi and Streett Automata to Deterministic Parity Automata. *Logical Methods in Computer Science* 3(3) (2007), [http://dx.doi.org/10.2168/LMCS-3\(3:5\)2007](http://dx.doi.org/10.2168/LMCS-3(3:5)2007)
17. Tiernan, J.C.: An Efficient Search Algorithm to Find the Elementary Circuits of a Graph. *Commun. ACM* 13(12), 722–726 (1970), <http://doi.acm.org/10.1145/362814.362819>
18. Weinblatt, H.: A New Search Algorithm for Finding the Simple Cycles of a Finite Directed Graph. *J. ACM* 19(1), 43–56 (1972), <http://doi.acm.org/10.1145/321679.321684>