



Synthesis of Asynchronous Reactive Programs from Temporal Specifications

Suguman Bansal^{1(✉)}, Kedar S. Namjoshi^{2(✉)}, and Yaniv Sa'ar^{3(✉)}

¹ Rice University, Houston, TX, USA
suguman@rice.edu

² Bell Labs, Nokia, Murray Hill, NJ, USA
kedar.namjoshi@nokia-bell-labs.com

³ Bell Labs, Nokia, Kfar Saba, Israel
yaniv.saar@nokia.bell-labs.com

Abstract. Asynchronous interactions are ubiquitous in computing systems and complicate design and programming. Automatic construction of asynchronous programs from specifications (“synthesis”) could ease the difficulty, but known methods are complex, and intractable in practice. This work develops substantially simpler synthesis methods. A direct, exponentially more compact automaton construction is formulated for the reduction of asynchronous to synchronous synthesis. Experiments with a prototype implementation of the new method demonstrate feasibility. Furthermore, it is shown that for several useful classes of temporal properties, automaton-based methods can be avoided altogether and replaced with simpler Boolean constraint solving.

1 Introduction

Modern software and hardware systems harness asynchronous interactions to improve speed, responsiveness, and power consumption: delay-insensitive circuits, networks of sensors, multi-threaded programs and interacting web services are all asynchronous in nature. Various factors contribute to asynchrony, such as unpredictable transmission delays, concurrency, distributed execution, and parallelism. The common result is that each component of a system operates with partial, out-of-date knowledge of the state of the others, which considerably complicates system design and programming. Yet, it is often easier to state the desired behavior of an asynchronous program. We therefore consider the question of automatically constructing (i.e., synthesizing) a correct reactive asynchronous program directly from its temporal specification.

The *asynchronous synthesis problem* was originally formulated by Pnueli and Rosner in 1989 on the heels of their work on synchronous synthesis [31, 32]. The task is that of constructing a (finite-state) program which interacts asynchronously with its environment while meeting a temporal specification on the actions at the interface between program and environment. Given a linear temporal specification φ , Pnueli-Rosner show that *asynchronous synthesis* can be

reduced to checking whether a derived specification φ' , specifying the required behavior of the scheduler, is *synchronously* synthesizable. That is, an asynchronous program can implement φ iff a synchronous program can implement φ' .

It may then appear straightforward to construct asynchronous programs using one of the many tools that exist for synchronous synthesis. However, the derived formula φ' embeds a nontrivial stutter quantification, which requires a complex intermediate automaton construction; it has not, to the authors' knowledge, ever been implemented. This situation is in stark contrast to that of synchronous synthesis, for which multiple tools and algorithms have been created.

Alternative methods have been proposed for asynchronous synthesis: Finkbeiner and Schewe reduce a bounded form of the problem to a SAT/SMT query [35], and Klein, Piterman and Pnueli show that some GR(1) specifications¹ can be transformed as above to an approximate synchronous GR(1) property [21,22]. These alternatives, however, have drawbacks of their own. The SAT/SMT reduction is exponential in the number of interface (input and output) bits, an important parameter; the GR(1) specifications amenable to transformation are limited and are characterized by semantic conditions that are not easily checked.

This work presents two key simplifications. First, we define a new property, $\text{PR}(\varphi)$ (named in honor of Pnueli-Rosner's pioneering work) which, like φ' , is synchronously realizable if, and only if, φ is asynchronously realizable. We then present an automaton construction for $\text{PR}(\varphi)$ that is direct and simpler, and results in an exponentially smaller automaton than the one for φ' . In particular, the automaton for $\text{PR}(\varphi)$ has only at most *twice* the states of the automaton for φ , as opposed to the *exponential blowup* of the state space (in the number of interface bits) incurred in the construction of the automaton for φ' . As almost all synchronous automaton-based synthesis tools use an explicit encoding for automaton states, this reduction is vital in practice.

We show how to implement the transformation PR symbolically (with BDDs), so that interface bits are always represented in symbolic form. One can then apply the modular strategy of Pnueli-Rosner: a symbolic automaton for φ is transformed to a symbolic automaton for $\text{PR}(\varphi)$ (instead of φ'), which is analyzed with a synchronous synthesis tool. We establish that PR is conjunctive and preserves safety². These are important properties, used by tools such as Acacia+ [8] and Unbeast [11] to optimize the synchronous synthesis task. The new construction has been implemented in a prototype tool, BAS, and experiments demonstrate feasibility in practice.

In addition, we establish that for several classes of temporal properties, which are easily characterized by syntax, the automaton-based method can be avoided entirely and replaced with Boolean constraint solving. The constraints are quantified Boolean formulae, with prefix $\exists\forall$ and a kernel that is derived from the original specification. This surprising reduction, which resolves a temporal prob-

¹ The GR(1) ("General Reactivity (1)") subclass has an efficient symbolic procedure for synchronous synthesis, formulated in [28] and implemented in several tools.

² I.e., $\text{PR}(\bigwedge_i f_i) = \bigwedge_i \text{PR}(f_i)$, and $\text{PR}(f)$ is a safety property if f is a safety property.

lem with Boolean reasoning, is a consequence of the highly adversarial role of the environment in the asynchronous setting.

These contributions turn a seemingly intractable synthesis task into one that is feasible in practice.

2 Preliminaries

Temporal Specifications. Linear Temporal Logic (LTL) [29] extends propositional logic with temporal operators. LTL formulae are defined as $\varphi ::= \text{True} \mid \text{False} \mid p \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid X\varphi \mid \varphi_1 U \varphi_2 \mid \diamond\varphi \mid \square\varphi \mid \boxplus\varphi$. Here p is a proposition, and X (Next), U (Until), \diamond (Eventually), \square (Always), and \boxplus (Always in the past) are temporal operators. The LTL semantics is standard, and is in the full version of the paper. For an LTL formula φ , let $\mathcal{L}(\varphi)$ denote the set of words (over subsets of propositions) that satisfy φ .

GR(1) is a useful fragment of LTL, where formulae are of the form $(\square S_e \wedge \bigwedge_{i=0}^m \square\diamond P_i) \Rightarrow (\square S_s \wedge \bigwedge_{i=0}^n \square\diamond Q_i)$, for propositional formulae S_e, S_s, P_i, Q_i . Typically, the left-hand side of the implication is used to restrict the environment, by requiring safety and liveness assumptions to hold, while the right-hand side is used to define the safety and liveness guarantees required of the system.

LTL specifications can be turned into equivalent Büchi automata, using standard constructions. A Büchi automaton, A , is specified by the tuple $(Q, Q_0, \Sigma, \delta, G)$, where Q is a set of states, $Q_0 \subseteq Q$ defines the initial states, Σ is the alphabet, $\delta \subseteq Q \times \Sigma \times Q$ is the transition relation, and $G \subseteq Q$ defines the “green” (also known as “accepting” or “final”) states. A *run* r of the automaton on an infinite word $\sigma = a_0, a_1, \dots$ over Σ is an infinite sequence $r = q_0, a_0, q_1, a_1, \dots$ such that q_0 is an initial state, and for each k , (q_k, a_k, q_{k+1}) is in the transition relation. Run r is accepting if a green state appears on it infinitely often; the language of A , denoted $\mathcal{L}(A)$, is the set of words that have an accepting run.

The Asynchronous Synthesis Model. The goal of synthesis is to construct an “open” program M meeting a specification at its interface. In the asynchronous setting, the program M interacts in a fair interleaved manner with its environment E . The fairness restriction requires that E and M are each scheduled infinitely often in all infinite executions. Let $E//M$ denote this composition. The interface between E and M is formed by the variables x and y . Variable x is written by E and is read-only for M , while y is written by M and is read-only for E . One can consider x (resp., y) to represent a vector of variables, i.e., $x = (x_1, \dots, x_n)$ (resp., $y = (y_1, \dots, y_m)$) which is read (resp., written) atomically. Many of our results also extend to non-atomic reads and writes, and are discussed in the full version of the paper.

The synthesis task is to construct a program M which satisfies a temporal property $\varphi(x, y)$ over the interface variables in the composition $E//M$, for *any* environment E . The most adversarial environment is the one which sets x to an arbitrary value at each scheduled step, we denote it by $\text{CHAOS}(x)$. The behaviors

of the composition $\text{CHAOS}(x)//M$ simulate those of $E//M$ for all E . Hence, it suffices to produce M which satisfies φ in the composition $\text{CHAOS}(x)//M$. One can limit the set of environments through an assumption in the specification.

This leads to the formal definition of an *asynchronous schedule*, given by a pair of functions, $r, w : \mathbb{N} \rightarrow \mathbb{N}$, which represent read and write points, respectively. The initial write point, $w(0) = 0$, and represents the choice of initial value for the variable y . Without loss of generality, the read-write points alternate, i.e., for all $i \geq 0$, $w(i) \leq r(i) < w(i+1)$ and $r(i) < w(i+1) \leq r(i+1)$. A *strict asynchronous schedule* does not allow read and write points to overlap, i.e., the constraints are strengthened to $w(i) < r(i) < w(i+1)$ and $r(i) < w(i+1) < r(i+1)$. A *tight asynchronous schedule* is the strict schedule without any non-read-write gaps, i.e., $r(k) = 2k + 1$ and $w(k) = 2k$, for all k . A *synchronous schedule* is the special non-strict schedule where $r(i) = i$ and $w(i) = i$, for all i .

Let D^v denote the binary domain $\{\text{True}, \text{False}\}$ for a variable v . A program M can be represented semantically as a function $f : (D^x)^* \rightarrow D^y$. For an asynchronous schedule (r, w) , a sequence $\sigma = (D^x \times D^y)^\omega$ is said to be an *asynchronous execution of f over (r, w)* if the value of y is changed only at writing points, in a manner that depends only on the values of x at prior reading points. Formally, for all $i \geq 0$, $y_{w(i+1)} = f(x_{r(0)} \dots x_{r(i)})$, and for all j such that $w(i) \leq j < w(i+1)$, $y_j = y_{w(i)}$. The initial value of y is the value it has at point $w(0) = 0$. The set of such sequences is denoted as $\text{asynch}(f)$. Over synchronous schedules, the set of such sequences is denoted by $\text{synch}(f)$. Function f is an asynchronous implementation of φ if all asynchronous executions of f over all possible schedules satisfy φ , i.e., if $\text{asynch}(f) \subseteq \mathcal{L}(\varphi)$.

This formulation agrees with that given by Pnueli and Rosner for strict schedules. For synchronous schedules (and other non-strict schedules), our formulation has a Moore-style semantics – the output depends on strictly earlier inputs – while Pnueli and Rosner formulate a Mealy semantics. A Moore semantics is more appropriate for modeling software programs, where the output variable is part of the state, and fits well with the theoretical constructions that follow.

Definition 1 (Asynchronous LTL Realizability). *Given an LTL property $\varphi(x, y)$ over the input variable x and output variable y , the asynchronous LTL realizability problem is to determine whether there is an asynchronous implementation for φ .*

Definition 2 (Asynchronous LTL Synthesis). *Given a realizable LTL-formula φ , the asynchronous LTL synthesis problem is to construct an asynchronous implementation of φ .*

Examples. Pnueli and Rosner give a number of interesting specifications. The specification $\Box(y \equiv Xx)$ (“the current output equals the next input”) is satisfiable but not realizable, as any implementation would have to be clairvoyant. On the other hand, the flipped specification $\Box(x \equiv Xy)$ (“the next output equals the current input”) is synchronously realizable by a Moore machine which replays

the current input as the next output. The specification $\diamond \square x \equiv \diamond \square y$ is synchronously realizable by the same machine, but is asynchronously unrealizable, as shown next. Consider two input (x) sequences, under a schedule where reads happen only at odd positions. In both, let $x = \text{true}$ at all reading points. Then any program must respond to both inputs with the same output sequence for y . Now suppose that in the first sequence x is **false** at all non-read positions, while in the second, x is **true** at all non-read positions. In the first case, the specification forces the output y -sequence to be **false** infinitely often; in the second, y is forced to be **true** from some point on, a contradiction.

The negated specification $\diamond \square x \not\equiv \diamond \square y$ is also asynchronously unrealizable, for the same reason. This “gap” illustrates an intriguing difference from the synchronous case, where either a specification is realizable for the system, or its negation is realizable for the environment. The two halves of the equivalence, i.e., $\diamond \square x \Rightarrow \diamond \square y$ and $\diamond \square y \Rightarrow \diamond \square x$ are individually asynchronously realizable, by strategies that fix the output to $y = \text{true}$ and to $y = \text{false}$, respectively.

From Asynchronous to Synchronous Synthesis. Pnueli and Rosner reduced asynchronous LTL synthesis to synchronous synthesis of Büchi objectives. Their reduction applied to LTL formulas with a single input and output variable [32]; it was later extended to the non-atomic case [30]. The original Rosner-Pnueli reduction deals exclusively with strict schedules, since they showed that it is sufficient to consider only strict schedules.

Two infinite sequences are said to be *stuttering equivalent* if one sequence can be obtained from the other by a finite duplication (“stretching”) of a given state or by deletion (“compressing”) of finitely many contiguous identical states retaining at least one of them. The *stuttering quantification* \exists^{\approx} is defined as follows: $\exists^{\approx} x. \varphi$ holds for sequence π if $\exists x. \varphi$ holds for a sequence π' that is stuttering equivalent to π . Pnueli-Rosner showed that an LTL-formula $\varphi(x, y)$ over input x and output y is asynchronously realizable iff a “kernel” formula (this is the precise formula referred to as φ' in the Introduction) $\mathcal{K}(r, w, x, y) = \alpha(r, w) \rightarrow \beta(r, w, x, y)$ over read sequence r , write sequence w , input sequence x and output sequence y is synchronously realizable:

$$\begin{aligned} \alpha(r, w) &= (\neg r \wedge \neg w \mathbf{U} r) \wedge \square \neg(r \wedge w) \wedge \square (r \Rightarrow (r \mathbf{U} (\neg r) \mathbf{U} w)) \\ &\quad \wedge \square (w \Rightarrow (w \mathbf{U} (\neg w) \mathbf{U} r)) \\ \beta(r, w, x, y) &= \varphi(x, y) \wedge \forall a. \square ((y = a) \Rightarrow ((y = a) \mathbf{U} (\neg w \wedge (y = a) \mathbf{U} w))) \\ &\quad \wedge \forall^{\approx} x'. (\square (\neg r \Rightarrow \neg r \mathbf{U} (x = x')) \Rightarrow \varphi(x', y)) \end{aligned}$$

Here, α encodes the strict scheduling constraints on read and write points, while β encodes conditions which assure a correct asynchronous execution over (r, w) . The \forall^{\approx} quantification, intuitively, quantifies over all adversarial schedules similar to the current (r, w) : it requires φ to hold over all sequences obtained from the current sequence σ by stretching or compressing the segments between read and write points, and choosing different values for x on those segments.

3 Symbolic Asynchronous Synthesis

Pnueli and Rosner’s procedure for asynchronous synthesis [32] is as follows: first, a Büchi automaton is built for the kernel formula $\neg\mathcal{K}$. This automaton is then determinized and complemented to form a deterministic word automaton for \mathcal{K} , which is then re-interpreted as a tree automaton and tested for non-emptiness. The transformations use standard constructions, except for the interpretation of the \exists^{\approx} operator in the formation of the Büchi automaton for $\neg\mathcal{K}$. For a Büchi automaton A , an automaton for $\exists^{\approx}\mathcal{L}(A)$ is constructed in two steps: first applying a “stretching” transformation on A , followed by a “compressing” transformation. Stretching introduces new automaton states of the form (q, a) , for each state q of A and each letter a .

When this general construction is applied to the formula $\neg\mathcal{K}$, the alphabet of the automaton A is formed of all possible valuations of the pair of variables (x, y) , which has size *exponential* in the number of interface bits. The stretching step introduces a copy of an automaton state for each letter, which results in an exponential blow-up of the state space of the constructed automaton. As all current tools for synchronous synthesis represent automaton states explicitly³, the exponential blowup introduced by the stuttering quantification is a significant obstacle to implementation.

In Pnueli-Rosner’s construction, the determinization and complementation steps are also complex, utilizing Safra’s construction. These steps are simplified by the “Safraless” procedure adopted in current tools for synchronous synthesis.

The other major issue with the Pnueli-Rosner construction is that the kernel formula \mathcal{K} introduces the scheduling variables r, w as input variables. However, the actions of a synthesized program should not rely on the values of these variables. Pnueli-Rosner ensure this by checking satisfiability over “canonical” tree models; it is unclear, however, how to realize this effect using a synchronous synthesis tool as a black box.

We define a new property, $\text{PR}(\varphi)$, that differs from \mathcal{K} but, similarly, is synchronously realizable if, and only if, φ is asynchronously realizable. We then present an automaton construction for $\text{PR}(\varphi)$ that bypasses the general construction for \exists^{\approx} , avoiding the exponential blowup and resulting in an automaton with *at most twice* the states of the original. Moreover, this construction refers only to x and y , avoiding the second issue as well. We then show that this construction can be implemented fully symbolically.

3.1 Basic Formulations and Properties

As formulated in Sect. 2, an asynchronous execution of f is determined by the schedule (r, w) . For a strict schedule, any infinite sequence representing an asynchronous behavior of f over (r, w) may be partitioned into a sequence of *blocks*, as follows. The start of the i ’th block is at the i ’th writing point, $w(i)$, and it

³ With one exception. BoSy’s DQBF procedure is fully symbolic but does not work as well as the default QBF procedure [12].

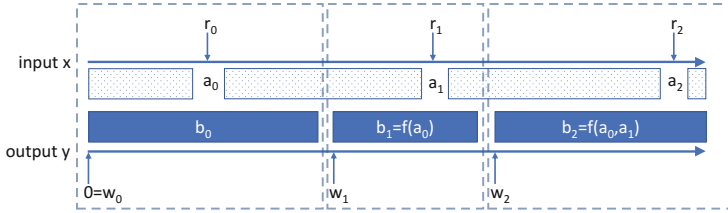


Fig. 1. A strict asynchronous computation for f . Values of x at non-reading points are shown as dotted. The y -value is constant between writing points, illustrated by a solid rectangle. Blocks are shown as dashed rectangles.

ends just before the $i + 1$ 'st writing point, $w(i + 1)$. The schedule ensures the i 'th block includes the i 'th reading point, $r(i)$, associated with the input-output value (x_i, y_i) . As the value of y changes only at writing points, y_i is constant in the i 'th block. Thus, the i 'th block follows the pattern $(\perp, y_i)^*(x_i, y_i)(\perp, y_i)^*$, where \perp denotes an arbitrary choice of x -value. Figure 1 illustrates a strict asynchronous computation and its decomposition into blocks.

Expansions. The set of *expansions* of sequence $\delta = (x_0, y_0)(x_1, y_1) \dots$ consists of all sequences obtained by simultaneously replacing each (x_i, y_i) in δ by a block with the pattern $(\perp, y_i)^*(x_i, y_i)(\perp, y_i)^*$. Formally, given sequences $\delta = (x_0, y_0)(x_1, y_1) \dots$ and $\sigma = (\bar{x}_0, \bar{y}_0)(\bar{x}_1, \bar{y}_1) \dots$, δ *expands to* σ , denoted as $\delta \text{ exp } \sigma$, if there exists an asynchronous schedule (\hat{r}, \hat{w}) for which σ is an execution that is a block pattern of δ , i.e., for all i , $x_i = \bar{x}_{\hat{r}(i)}$ and $y_i = \bar{y}_{\hat{w}(i)}$ and for all j , $\hat{w}(i) \leq j < \hat{w}(i + 1)$ it is the case that $\bar{y}_j = \bar{y}_{\hat{w}(i)}$. The inverse relation (read as *contracts to*) is denoted by exp^{-1} . Figure 2 shows the synchronous computation that contracts the computation shown in Fig. 1.

Relational Operators. For a relation R , the modal operators $\langle R \rangle$ and $[R]$ are defined as follows. For any set S ,

$$u \in \langle R \rangle S = (\exists v : uRv \wedge v \in S) \quad u \in [R]S = (\forall v : uRv \Rightarrow v \in S)$$

By definition, the operators are negation duals, i.e., $\neg \langle R \rangle (\neg S) = [R](S)$ for any R and any S . For an LTL formula φ and a relation R over infinite sequences, we let $\langle R \rangle \varphi$ abbreviate $\langle R \rangle (\mathcal{L}(\varphi))$, and similarly, let $[R] \varphi$ abbreviate $[R] (\mathcal{L}(\varphi))$.

Galois Connections. Given partial orders (A, \preceq_A) and (B, \preceq_B) , a pair of functions $g : A \rightarrow B$ and $h : B \rightarrow A$ form a Galois connection if, for all $a \in A, b \in B$: $g(a) \preceq_B b$ is equivalent to $a \preceq_A h(b)$. From the definitions, it is clear that the operators $(\langle R^{-1} \rangle, [R])$ form a Galois connection over the partial orders defined by the subset relation. I.e., for any sets S and T : $\langle R^{-1} \rangle S \subseteq T$ iff, $S \subseteq [R]T$.

We first establish that the asynchronous executions of f are precisely the synchronous executions of f under an inverse expansion.

Theorem 1. *For an implementation f , $\text{asynch}(f) = \langle \text{exp}^{-1} \rangle \text{synch}(f)$.*

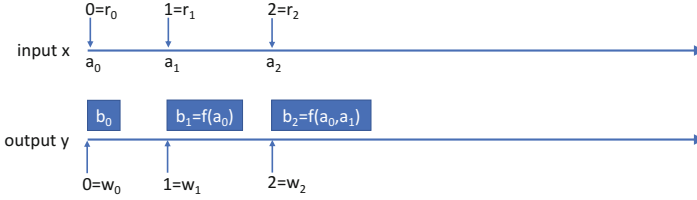


Fig. 2. The contracted synchronous (Moore) computation

Proof. (ping) Let σ be an execution in $\text{asynch}(f)$, generated for some schedule (r, w) . For any k , consider the k 'th block of σ . This is the set of positions from $w(k)$ to $w(k + 1) - 1$, which includes the k 'th reading point $r(k)$, say with the value (x_k, y_k) . Then the block follows the pattern $(\perp, y_k)^*(x_k, y_k)(\perp, y_k)^*$. So σ is an expansion of the sequence $\delta = (x_0, y_0)(x_1, y_1) \dots$. By the definition of an asynchronous execution, the value $y_{k+1} = f(x_0, \dots, x_k)$. This is precisely the requirement for δ to be a synchronous execution of f . Hence, we have that there is a δ such that $\delta \text{exp } \sigma$ and $\delta \in \text{synch}(f)$. Therefore, $\sigma \in \langle \text{exp}^{-1} \rangle \text{synch}(f)$.

(pong) Let σ be in $\langle \text{exp}^{-1} \rangle \text{synch}(f)$. By definition, there is a $\text{synch}(f)$ execution $\delta = (x_0, y_0)(x_1, y_1) \dots$ such that $\delta \text{exp } \sigma$. As δ is a synchronous execution of f , the value $y_{k+1} = f(x_0, x_1, \dots, x_k)$, for all k . Then σ is an asynchronous execution of f under the schedule where the k -th reading point is the point that the k 'th entry, (x_k, y_k) , from δ is mapped to in σ , and the $(k + 1)$ -th writing point is the first point of the $(k + 1)$ 'st block in the expansion. \square

We now use the Galois connection to show how asynchronous synthesis can be reduced to an equivalent synchronous synthesis task. Consider a property φ that must hold asynchronously for an implementation f .

Theorem 2. *Let f be an implementation function, and φ a property. Then $\text{asynch}(f) \subseteq \mathcal{L}(\varphi)$ if, and only if, $\text{synch}(f) \subseteq [\text{exp}] \varphi$.*

Proof. From Theorem 1, $\text{asynch}(f) \subseteq \mathcal{L}(\varphi)$ holds iff $\langle \text{exp}^{-1} \rangle \text{synch}(f) \subseteq \mathcal{L}(\varphi)$ does. By the Galois connection, this is equivalent to $\text{synch}(f) \subseteq [\text{exp}] \varphi$. \square

3.2 The Pnueli-Rosner Closure

We refer to the property $[\text{exp}] \varphi$ as the Pnueli-Rosner closure of φ , in honor of their pioneering work on this problem, and denote it by $\text{PR}(\varphi)$. This has interesting mathematical properties, which are useful in practice.

Theorem 3. *$\text{PR}(\varphi) = [\text{exp}] \varphi$ has the following properties.*

1. (Closure) PR is monotonic and a downward closure, i.e., $\text{PR}(\varphi) \subseteq \mathcal{L}(\varphi)$
2. (Conjunctivity) PR is conjunctive, i.e., $\text{PR}(\bigwedge_i \varphi_i) = \bigcap_i \text{PR}(\varphi_i)$
3. (Safety Preservation) If φ is a safety property, so is $\text{PR}(\varphi)$

The closure property relies on the reflexivity and transitivity of exp , and that $[R]$ is monotonic for every R . Conjunctivity follows from the conjunctivity of $[R]$ for any R . Safety preservation is based on the Alpern-Schneider [4] formulation of safety over infinite words. Proofs are in the full version of the paper.

Conjunctivity is exploited by the tools Acacia+ [8] and Unbeast [11] to optimize the synchronous synthesis procedure. The Unbeast tool also separates out safety from non-safety sub-properties to optimize the synthesis procedure. Thus, if a specification φ has the form $\varphi_1 \wedge \varphi_2$, where φ_1 is a safety property, then $\text{PR}(\varphi) = \text{PR}(\varphi_1) \cap \text{PR}(\varphi_2)$ also denotes the intersection of the safety property $\text{PR}(\varphi_1)$ with another property.

3.3 The Closure Automaton Construction

By negation duality, $\text{PR}(\varphi)$ equals $\neg\langle \text{exp} \rangle(\neg\varphi)$. We use this property to reduce asynchronous to synchronous synthesis, as follows.

1. Construct a non-deterministic Büchi automaton A for $\neg\varphi$,
2. Transform A to a non-deterministic Büchi automaton B for the negated Pnueli-Rosner closure of φ , i.e., the language of B is $\langle \text{exp} \rangle\mathcal{L}(A) = \langle \text{exp} \rangle(\neg\varphi)$,
3. Consider the structure of B as that of a *universal* co-Büchi automaton, which has language $\neg\mathcal{L}(B)$,
4. Synthesize an implementation f in the synchronous model which satisfies $\neg\mathcal{L}(B) = \neg\langle \text{exp} \rangle\mathcal{L}(A) = \neg\langle \text{exp} \rangle(\neg\varphi) = [\text{exp}]\varphi = \text{PR}(\varphi)$.

The new step is the second one, which constructs B from A ; the others use standard constructions and tools. This construction is as follows.

- The states and alphabet of B are the states and alphabet of A .
- The transitions of B are determined by a saturation procedure. For every pair of states q, q' , and letter (x, y) , let $\Pi(q, (x, y), q')$ be the set of paths in A from q to q' where the sequence of letters on the path matches the expansion pattern $(\perp, y)^*(x, y)(\perp, y)^*$. The transition $(q, (x, y), q')$ is in B if, and only if, this set is non-empty,
- If some path in $\Pi(q, (x, y), q')$ passes through a green (accepting) state of A , the transition $(q, (x, y), q')$ in B is colored “green” and that path is assigned as the witness to the transition in B . On the other hand, if none of the paths in $\Pi(q, (x, y), q')$ pass through a green state, this transition is not colored in B , and one of the paths in the set is chosen as the witness for this transition,
- The automaton B inherits the accepting (“green”) states of A and it may have, in addition, green transitions introduced as defined above,
- A sequence is accepted by B if there is a run of B on the sequence such that either there are infinitely many green states, or infinitely many green transitions on that run.

We establish that $\mathcal{L}(B) = \langle \text{exp} \rangle\mathcal{L}(A)$ through the following two lemmas.

Lemma 1. $\langle \text{exp} \rangle\mathcal{L}(A) \subseteq \mathcal{L}(B)$.

Proof. Let $\delta = (x_0, y_0)(x_1, y_1) \dots$ be a sequence in $\langle \text{exp} \rangle \mathcal{L}(A)$. By definition, there exists a sequence σ in $\mathcal{L}(A)$ such that $\delta \text{exp} \sigma$. The expansion σ follows the pattern $[(\perp, y_0)^*(x_0, y_0)(\perp, y_0)^*][(\perp, y_1)^*(x_1, y_1)(\perp, y_1)^*] \dots$, where $[\dots]$ are used merely to indicate the boundaries of a block. An accepting run of A on σ has the form $q_0[(\perp, y_0)^*(x_0, y_0)(\perp, y_0)^*]q_1[(\perp, y_1)^*(x_1, y_1)(\perp, y_1)^*]q_2 \dots$, where the states on the run inside a block have been elided. By the definition of B , the segment $q_0(\perp, y_0)^*(x_0, y_0)(\perp, y_0)^*q_1$ induces a transition from q_0 to q_1 in B on the letter (x_0, y_0) . Similarly, the following segment induces a transition from q_1 to q_2 on letter (x_1, y_1) , and so forth. These transitions together form a run $q_0(x_0, y_0)q_1(x_1, y_1)q_2 \dots$ of B on δ .

If one of the $\{q_i\}$ is green and appears infinitely often on the run on σ , the induced run on δ is accepting. Otherwise, as the run on σ is accepting, some green state of A occurs in the interior of infinitely many segments of that run. The transitions of B induced by those segments must be green, so the corresponding run on δ has infinitely many green edges, and is accepting for B . \square

Lemma 2. $\mathcal{L}(B) \subseteq \langle \text{exp} \rangle \mathcal{L}(A)$.

Proof. Let δ be accepted by B . We show that there is σ such that $\delta \text{exp} \sigma$ and σ is accepted by A . Let δ have the form $(x_0, y_0)(x_1, y_1) \dots$. Denote the accepting run of B on δ by $r = q_0(x_0, y_0)q_1(x_1, y_1) \dots$. From the construction of B , the transition from q_0 to q_1 on (x_0, y_0) has an associated witness path through A from q_0 to q_1 , which follows the expansion pattern $(\perp, y_0)^*(x_0, y_0)(\perp, y_0)^*$ on its edge labels. Stitching together the witness paths for each transition of r , we obtain both a sequence σ that is an expansion of δ and a run r' of A on σ .

As r is accepting for B , it must enter infinitely often either a green state or a green edge. If it enters a green state infinitely often, that state appears infinitely often on r' . If r enters a green edge infinitely often, the witness path for that edge contains a green state of A , say q ; as this path is repeated infinitely often on σ , q appears infinitely often on r' . In either case, a green state of A appears infinitely often on r' , which is therefore, an accepting run of A on σ . \square

Automaton B can be placed in standard form by converting its green edges to green states as follows, forming a new automaton, \hat{B} . Form a green copy of the state space, i.e., for each state q , form a green variant, $G(q)$, which is marked as an accepting state. Set up transitions as follows. If (q, a, q') is an original non-green transition, then (q, a, q') and $(G(q), a, q')$ are new transitions. If (q, a, q') is an original green transition, then $(q, a, G(q'))$ and $(G(q), a, G(q'))$ are new transitions. This at most doubles the size of the automaton. It is straightforward to establish that $\mathcal{L}(B) = \mathcal{L}(\hat{B})$.

3.4 Symbolic Construction

The symbolic construction of \hat{B} closely follows the definitions above. It is easily implemented with BDDs representing predicates on the input and output variables x and y . The crucial step is to use fixpoints to formulate the existence of paths in the set Π used in the definition of B . These definitions are similar to

the fixpoint definition of the CTL modality EF. We use $A(q, (x, y), q')$ to denote the predicate on (x, y) describing the transition from q to q' in automaton A .

Fixed Don't-Care Path. Let $\text{EfixedY}(q, y, q')$ hold if there is a path of length 0 or more from q to q' in A where the value of y is fixed. This is the least fixpoint (in Z) of the following implications:

- $(q' = q) \Rightarrow Z(q, y, q')$, and
- $(\exists x, r : A(q, (x, y), r) \wedge Z(r, y, q')) \Rightarrow Z(q, y, q')$

The predicate $A^\perp(q, y, r) = (\exists x : A(q, (x, y), r))$ is pre-computed. Then, the least fixpoint is computed iteratively as follows.

$$\text{EfixedY}^0(q, y, q') = (q = q')$$

$$\text{EfixedY}^{i+1}(q, y, q') = \text{EfixedY}^i(q, y, q') \vee (\exists r : A^\perp(q, y, r) \wedge \text{EfixedY}^i(r, y, q'))$$

Let predicate $\text{green}_A(r)$ be true for an accepting state r of A . The predicate $\text{Efixedgreen}(q, y, q')$ holds if there is a fixed y -path from q to q' where one of the states on it is green:

$$\text{Efixedgreen}(q, y, q') = (\exists r : \text{EfixedY}(q, y, r) \wedge \text{green}_A(r) \wedge \text{EfixedY}(r, y, q'))$$

Paths and Green Paths. Let $\text{Epath}(q, (x, y), q')$ hold if there is a path following the block pattern $(\perp, y)^*(x, y)(\perp, y)^*$ from q to q' in A . Then,

$$\text{Epath}(q, (x, y), q') = (\exists r, r' : \text{EfixedY}(q, y, r) \wedge A(r, (x, y), r') \wedge \text{EfixedY}(r', y, q'))$$

Similarly, let $\text{Egreenpath}(q, (x, y), q')$ hold if there is a path following the block pattern $(\perp, y)^*(x, y)(\perp, y)^*$ from q to q' in A , with an intermediate green state.

$$\text{Egreenpath}(q, (x, y), q') =$$

$$(\exists r, r' : \text{Efixedgreen}(q, y, r) \wedge A(r, (x, y), r') \wedge \text{EfixedY}(r', y, q')) \vee$$

$$(\exists r, r' : \text{EfixedY}(q, y, r) \wedge A(r, (x, y), r') \wedge \text{Efixedgreen}(r', y, q'))$$

State Space of \hat{B} . The state space of \hat{B} is formed by pairs (q, g) , where q is a state of A and g is a Boolean indicating whether it is a new green state. The accepting condition $\text{green}_{\hat{B}}(q, g)$ of \hat{B} is given by $\text{green}_A(q) \vee g$.

Initial States. The initial predicate $I_{\hat{B}}(q, g)$ is $I_A(q) \wedge \neg g$, where $I_A(q)$ is true for initial states of the input automata A .

Transition Relation of \hat{B} . The transition relation $\hat{B}((q, g), (x, y), (q', g'))$ is

$$\hat{B}((q, g), (x, y), (q', g')) = \text{Epath}(q, (x, y), q') \wedge (g' \equiv \text{Egreenpath}(q, (x, y), q'))$$

4 Implementation and Experiments

The PR algorithm has been implemented in a framework called BAS (Bounded Asynchronous Synthesis). It uses the LTL-to-automaton converter LTL3BA [3, 6], and follows the modular method, connecting to either of two solvers, BoSy [2, 12] and Acacia+ [1, 8] to solve the synchronous realizability of $\text{PR}(\varphi)$. The PR construction is implemented in about 1200 lines of OCaml, using an external BDD library. (The core construction requires only about 400 lines of code.) For an LTL specification φ , the BAS workflow for asynchronous synthesis is as follows:

1. Check whether φ is synchronously realizable; if not, return **UNREALIZABLE**,
2. Construct Büchi automata A for $\neg\varphi$, and \hat{A} for φ ,
3. Concurrently
 - (a) Construct $\text{PR}(\varphi)$ from A and check whether it is synchronously realizable; if so, return **REALIZABLE** and synthesize the implementation.
 - (b) Construct $\text{PR}(\neg\varphi)$ from \hat{A} and check whether it is synchronously realizable for the environment; if so, return **UNREALIZABLE**.

Upon termination of any, terminate the other execution as well.

The synchronous synthesis tools successively increase a bound until a limit (computed based on automaton structure) is reached. Thus, in theory, only the check in step 3(a) is needed. However, the checks in steps 1 and 3(b) may allow the tool to terminate early (before reaching the limit bound), if a winning strategy for the environment can be discovered.

To evaluate BAS we consider the list of examples presented in Table 1. The reported experiments were performed on a VM configured to have 8 CPU cores at 2.4 GHz, 8 GB RAM, running 64-bit Linux. The running times are reported in milliseconds. For each specification (presented in the second column) we report whether it is asynchronously realizable (third column), the time for the PR construction (our contribution), and the time for checking whether the specification is realizable using BoSy and Acacia+ solvers (resp., fifth and sixth columns).

The first set of examples (Specifications 1–11) list specifications discussed in this paper and in related works. As parameterized example we consider 2 variants of arbiter specifications. The arbiter has n inputs in which clients request permissions, and n outputs in which the clients are granted permissions. In both variants of the arbiter example, no two grants are allowed to be set simultaneously. The first arbiter example (Specification 12) requires that whenever an input request r_i is set, the corresponding output grant g_i must eventually be set. The second variant (Specification 13) also requires that a grant g_i is set only if request r_i is set as well. That is, in order for a client to be granted a permission, its corresponding request must be constantly set. Since the asynchronous case cannot observe the request in between read events, this variant of the arbiter is not realizable. The results are shown for $n = 2, 4, 6$. Note that the only comparable experimental evaluation is given in [18], where they report that asynchronous synthesis of the first arbiter example (Specification 12) takes over 8 h.

Table 1. BAS asynchronous synthesis runtime evaluation (times in milliseconds). We let BoSy run upto 2h, and Acacia+ upto 1000 iterations. “Na” denotes cases where the executions did not find a winning strategy within these boundaries.

	Specification	Asyn. Realizable?	PR constr.	Asyn. synthesis	
				BoSy	Acacia+
1	$\Box (x \equiv y)$	False	8	972	30
2	$\Diamond \Box x \equiv \Diamond \Box y$	False	9	Na	Na
3	$\Diamond \Box x \Rightarrow \Diamond \Box y$	True	8	899	Na
4	$\Diamond \Box y \Rightarrow \Diamond \Box x$	True	7	994	Na
5	$(\Diamond \Box x \vee \Diamond \Box \neg x) \Rightarrow \Diamond \Box x \equiv \Diamond \Box y$	True	13	1004	Na
6	$\Box (\neg x \Rightarrow (\neg x) \text{ U } (\neg y)) \Rightarrow \Diamond \Box x \equiv \Diamond \Box y$	True	10	Na	Na
7	$\Box \Diamond (x \wedge y) \Rightarrow (\Box \Diamond y \wedge \Box \Diamond \neg y)$	True	9	1053	30
8	$\Box \Diamond (x \vee y) \Rightarrow (\Box \Diamond y \wedge \Box \Diamond \neg y)$	True	9	995	40
9	$\Box \Diamond (x) \Rightarrow (\Box \Diamond y \wedge \Box \Diamond \neg y)$	True	8	934	30
10	$\Box (x \Rightarrow \Diamond y)$	True	8	960	30
11	$\Box (x \Rightarrow \Diamond y) \wedge \Box (\neg y \text{ U } x)$	False	10	1058	Na
Variants of parameterized arbiter (results shown are for $n = 2; 4; 6$)					
12	$\bigwedge_{i \neq j} \Box (\neg g_i \vee \neg g_j) \wedge \bigwedge_{i=1}^n \Box (r_i \Rightarrow \Diamond g_i)$	True	11; 13; 75	854; 1146; 4965	Na; Na; Na
13	$\bigwedge_{i \neq j} \Box (\neg g_i \vee \neg g_j) \wedge \bigwedge_{i=1}^n \Box (r_i \Rightarrow \Diamond g_i) \wedge \bigwedge_{i=1}^n \Box (g_i \Rightarrow r_i)$	False	17; 3124; 2024K	1129; 362K; Na	Na; Na; Na

The second specification φ is the one discussed in Sect. 2. It is surprisingly difficult to solve. Both φ and its negation are asynchronously unrealizable. Moreover, φ is synchronously realizable. Thus, the early detection tests (steps 1 and 3(b)) failed to discover a winning strategy for the environment; the bounded synthesis tools increase the considered bound monotonically without converging to an answer in a reasonable amount of time. This example highlights the need for better tests for unrealizability. The results in the following section provide simple QBF tests of unrealizability for subclasses of LTL.

5 Efficiently Solvable Subclasses of LTL

The high complexity of direct LTL (synchronous) synthesis has encouraged the search for general procedures that work well in practice, such as Safraless and bounded synthesis [24, 35]. Another useful direction has been to identify fragments of LTL with efficient synthesis algorithms [5]. Among the most noteworthy is the GR(1) subclass, for which there is an efficient, symbolic synthesis procedure ([28]). We explore this direction for *asynchronous* synthesis. Surprisingly, we show that synthesis for certain fragments of LTL can be reduced to

Boolean reasoning over properties in QBF. The results cover several types of GR(1) formulae, although the question of a reduction for all of GR(1) is open.

The QBF formulae that arise have the form $\exists y \forall x. p(x, y)$, where x and y are disjoint sets of variables, and p is a propositional formula over x, y . An assignment $y = b$ for which $\forall x. p(x, b)$ holds is called a *witness* to the formula. The first such reduction is for the property $\square \diamond P$.

Theorem 4. $\varphi = \square \diamond P$ is asynchronously realizable iff $\exists y \forall x P$ is True.

Proof. (ping) Let b be a witness to $\exists y \forall x. P$. The function that constantly outputs $y = b$ satisfies φ for any asynchronous schedule.

(pong) Let f be a candidate implementation function and suppose that $\forall y \exists x (\neg P)$ holds. Fix any schedule. For every value $y = b$ that function f outputs at a writing point, there exists an input value $x = a$ such that $\neg P(a, b)$ holds. Thus, the environment, by issuing $x = a$ in the interval from the current writing point (with $y = b$) up to the next one, can ensure that $\neg P$ holds throughout the execution. Thus the specification $\varphi = \square \diamond P$ does not hold on this execution. \square

The result in Theorem 4 applies to asynchronous synthesis, but does not apply to synchronous synthesis. For example, the property $\square \diamond (x \equiv y)$ is asynchronously unrealizable, as $\exists y \forall x (x \equiv y)$ is False. On the other hand, it is synchronously realizable with a Mealy machine that sets y to x at each point.

Theorem 4 extends easily to conjunction and disjunction of $\square \diamond$ properties.

Theorem 5. Specification $\varphi = \bigvee_{i=0}^m \square \diamond P_i$ is asynchronously realizable iff $\exists y \forall x. (\bigvee_{i=0}^m P_i)$ holds. Additionally, specification $\varphi = \bigwedge_{i=0}^m \square \diamond P_i$ is asynchronously realizable iff for all $i \in \{0, 1 \dots m\}$, $\exists y \forall x. P_i$ holds.

Proof. The first claim follows directly from the identity $\bigvee_{i=0}^m \square \diamond P_i \equiv \square \diamond (\bigvee_{i=0}^m P_i)$ and Theorem 4.

For the second, for each i , let $y = b_i$ be an assignment such that $\forall x. P_i(x, b_i)$ holds. The function that generates sequence b_0, b_1, \dots, b_m , ad infinitum, is an asynchronous implementation of $\bigwedge_{i=0}^m \square \diamond P_i$. On the other hand, suppose that for some i , $\forall y \exists x \neg P_i$ holds, then following the construction from Theorem 4, one can define an execution where P_i is always False. \square

Theorem 6. $\varphi = \diamond \square P$ is asynchronously realizable iff $\exists y \forall x. P$ is True.

The proof is similar to that for Theorem 4. Theorem 6 also extends to conjunctions and disjunctions of $\diamond \square$ properties, by arguments similar to those for Theorem 5. Namely, $\bigwedge_{i=0}^m \diamond \square P_i$ is asynchronously realizable iff $\exists y \forall x (\bigwedge_{i=0}^m P_i)$ is True, and, $\bigvee_{i=0}^m \diamond \square P_i$ is asynchronously realizable iff for some $i \in \{0, 1, \dots, m\}$, $\exists y \forall x. P_i$ is True. Theorems 4–6 apply to non-atomic reads and writes of multiple input and output variables. Proofs are in the full version of the paper.

We now consider a more general type of GR(1) formula. The *strict semantic* of GR(1) formula $\square S_e \wedge \square \diamond P \Rightarrow \square S_s \wedge \square \diamond Q$ is defined to be $\square (\boxplus S_e \Rightarrow S_s) \wedge (\square S_e \wedge \square \diamond P \Rightarrow \square \diamond Q)$ – i.e., S_s is required to hold so long as S_e has always held in the past; and if S_e holds always and P holds infinitely often, then Q holds infinitely often. This is the interpretation supported by GR(1) synchronous synthesis tools.

Theorem 7. *The strict semantics of GR(1) specification $\Box S_e \wedge \Box \Diamond P \Rightarrow \Box S_s \wedge \Box \Diamond Q$ is asynchronously realizable iff $\exists y \forall x. (S_e \Rightarrow (S_s \wedge (P \Rightarrow Q)))$ is True.*

Proof. (ping) If $y = b$ is a witness to $\exists y \forall x. (S_e \Rightarrow (S_s \wedge (P \Rightarrow Q)))$, let f be a function that always generates b . Suppose S_e holds up to point i , then as $y = b$, regardless of the x -value, S_s holds at point i . This shows that the first part of the specification holds. For the second, suppose that S_e holds always and P is true infinitely often. Then, by choice of $y = b$, $(P \Rightarrow Q)$ holds always, thus Q holds infinitely often as well.

(pong) To prove the other side of the implication, we proceed as in Theorem 4. Let f be a candidate implementation. Fix a schedule, and suppose that $\forall y \exists x. (S_e \wedge (\neg S_s \vee \neg(P \Rightarrow Q)))$ holds. Then for every step of the execution and for every value $y = b$ that function f outputs at a writing point, there exists a value $x = a$ which the environment can choose from that writing point to the next such that $S_e(a, b)$ is true, and one of $S_s(a, b)$ or $(P \Rightarrow Q)(a, b)$ is false at every point in that interval.

On this execution, S_e holds throughout. If S_s is false at some point, this violates the first part of the specification. If not, then $(P \Rightarrow Q)$ must be false everywhere; i.e., at every point P is true but Q is false. Thus, S_e holds everywhere and P holds infinitely often but Q does not hold infinitely often, violating the second part of the specification. \square

Theorem 7 applies to atomic reads and writes, showing that asynchronous synthesis of GR(1) specification can be reduced to Boolean reasoning over properties in QBF. For non-atomic reads and writes, safety in asynchronous systems is more nuanced, since there is a delay between the write points of the first and last outputs in each round. This is discussed in the full version of the paper. This proof strategy does not generalize easily to the full GR(1) format, where more than one $\Box \Diamond$ property can appear on either side of the implication.

These results establish that the asynchronous synthesis problem for such specifications is easily solvable—more easily than in the synchronous setting, surprisingly avoiding entirely the need for automaton constructions and bounded synthesis. From another, equally valuable, point of view, the results show that such types of specifications may be of limited interest for automated synthesis, as solvable cases have very simple solutions.

6 Conclusions and Related Work

This work tackles the task of asynchronous synthesis from temporal specifications. The main results are a new symbolic automaton construction for general temporal properties, and the reduction of the synthesis question for several classes of specifications to QBF. These are mathematically interesting, being substantial simplifications of prior methods. Moreover, they make it feasible to implement an asynchronous synthesis tool following the modular process suggested by Pnueli and Rosner in 1989, by reducing asynchronous synthesis to a synchronous synthesis question. To the best of our knowledge, this is the first

such tool. The prototype, which builds on tools for synchronous synthesis, is able to quickly synthesize asynchronous programs for several interesting properties. There are, undoubtedly, several challenges, one of which is the quick detection of unrealizable specifications.

Our work builds upon several earlier results, which we discuss here. The synthesis question for temporal properties originates from a question posed by Church in the 1950s (see [37]). The problem of synthesizing a synchronous reactive system from a linear temporal specification was formulated and studied by Pnueli and Rosner [31], who gave a solution based on non-emptiness of tree automata. There has been much progress on the synchronous synthesis question since. Key developments include the discovery of efficient symbolic (BDD-based) solutions for the GR(1) class [7, 28], the invention of “Safraless” procedures [24], the application of these ideas for bounded synthesis [15, 35], and their implementation in a number of tools, e.g. [8, 10, 11, 13, 20, 34]. These have been applied in many settings (cf. [9, 23, 25–27]).

The problem of synthesizing asynchronous programs was also formulated and studied by Pnueli and Rosner [32] but has proved to be much more challenging, with only limited progress. The original Pnueli-Rosner constructions are complex and were not implemented. Work by Klein, Piterman and Pnueli, nearly 20 years later [22], shows tractability for some GR(1) specifications. However, the class of specifications that can be so handled is characterized by semantic constraints such as stuttering-closure and memoryless-ness, which are difficult to recognize.

Finkbeiner and Schewe [18, 35] present an alternative method, based on bounded synthesis, that applies to all LTL properties: it encodes the existence of a deductive proof for a bounded program into SAT/SMT constraints. However, the encoding represents inputs and outputs explicitly and is, therefore, exponential in the number of input and output bits. The exponential blowup has practical consequences: an asynchronous arbiter specification requires over 8 h to synthesize [18]; the same specification can be synthesized by our method in seconds. (Note, however, that the method in [18] is not specialized to asynchronous synthesis, and this difference may not be solely due to the explicit state representation, as the specification has only 4 bits.) Recent work gives an alternative encoding of synchronous bounded synthesis into QBF constraints, retaining input and output bits in symbolic form [12]. We believe that a similar encoding applies to asynchronous bounded synthesis as well, this is a topic for future work.

Pnueli and Rosner’s model of interface communication is not the only choice. Other models for asynchrony could, for instance, be based on CCS/CSP-style rendezvous communication at the interface, or permit shared read-write variables with atomic lock/unlock actions. Petri net game models have also been suggested for distributed synthesis [16]. An orthogonal direction is to weaken the adversarial power of the environment through a probabilistic model which can be used to constrain unlikely, highly adversarial input patterns to have probability 0, thus turning the synthesis problem into one where programs satisfy their specifications with high probability. (The synthesis of multiple processes is known to be undecidable in most cases [17, 33].)

In the broader context of fully automatic program synthesis, there are various approaches to the synthesis of single-threaded, terminating programs from formal pre- and post-condition specifications and from examples, using type information and other techniques to prune the search space. (We will not attempt to survey this large field, some examples are [14, 19, 36].) An intriguing question is to investigate how the techniques developed in these distinct lines of work can be fruitfully combined to aid the development of asynchronous, reactive software.

Acknowledgements. Kedar Namjoshi was supported, in part, by NSF grant CCF-1563393 from the National Science Foundation. We would like to thank Michael Emmi for many helpful discussions during the early stages of this work.

References

1. Acacia+. <http://lit2.ulb.ac.be/acaciaplus/>
2. BoSy. <https://www.react.uni-saarland.de/tools/bosy/>
3. LTL3BA. <https://sourceforge.net/projects/ltl3ba/>
4. Alpern, B., Schneider, F.B.: Defining liveness. *Inf. Process. Lett.* **21**(4), 181–185 (1985)
5. Alur, R., La Torre, S.: Deterministic generators and games for LTL fragments. *ACM Trans. Comput. Log.* **5**(1), 1–25 (2004)
6. Babiak, T., Křetínský, M., Řehák, V., Strejček, J.: LTL to Büchi automata translation: fast and more deterministic. In: Flanagan, C., König, B. (eds.) TACAS 2012. LNCS, vol. 7214, pp. 95–109. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-28756-5_8
7. Bloem, R., Jobstmann, B., Piterman, N., Pnueli, A., Sa’ar, Y.: Synthesis of reactive (1) designs. *J. Comput. Syst. Sci.* **78**(3), 911–938 (2012)
8. Bohy, A., Bruyère, V., Filiot, E., Jin, N., Raskin, J.-F.: Acacia+, a tool for LTL synthesis. In: Madhusudan, P., Seshia, S.A. (eds.) CAV 2012. LNCS, vol. 7358, pp. 652–657. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-31424-7_45
9. D’Ippolito, N., Braberman, V., Piterman, N., Uchitel, S.: Synthesizing nonanomalous event-based controllers for liveness goals. *Trans. Softw. Eng. Methodol.* **22**(1), 9 (2013)
10. Ehlers, R.: Symbolic bounded synthesis. In: Touili, T., Cook, B., Jackson, P. (eds.) CAV 2010. LNCS, vol. 6174, pp. 365–379. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-14295-6_33
11. Ehlers, R.: Unbeast: symbolic bounded synthesis. In: Abdulla, P.A., Leino, K.R.M. (eds.) TACAS 2011. LNCS, vol. 6605, pp. 272–275. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-19835-9_25
12. Faymonville, P., Finkbeiner, B., Rabe, M.N., Tentrup, L.: Encodings of bounded synthesis. In: Legay, A., Margaria, T. (eds.) TACAS 2017. LNCS, vol. 10205, pp. 354–370. Springer, Heidelberg (2017). https://doi.org/10.1007/978-3-662-54577-5_20
13. Faymonville, P., Finkbeiner, B., Tentrup, L.: BoSy: an experimentation framework for bounded synthesis. In: Majumdar, R., Kunčák, V. (eds.) CAV 2017. LNCS, vol. 10427, pp. 325–332. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63390-9_17

14. Feng, Y., Martins, R., Wang, Y., Dillig, I., Reps, T.W.: Component-based synthesis for complex APIs. In: Castagna, G., Gordon, A.D. (eds.) *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, 18–20 January 2017*, pp. 599–612. ACM (2017)
15. Filiot, E., Jin, N., Raskin, J.-F.: Compositional algorithms for LTL synthesis. In: Bouajjani, A., Chin, W.-N. (eds.) *ATVA 2010. LNCS, vol. 6252*, pp. 112–127. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-15643-4_10
16. Finkbeiner, B., Olderog, E.-R.: Petri games: synthesis of distributed systems with causal memory. *Inf. Comput.* **253**, 181–203 (2017)
17. Finkbeiner, B., Schewe, S.: Uniform distributed synthesis. In: *20th IEEE Symposium on Logic in Computer Science (LICS 2005)*, 26–29 June 2005, Chicago, IL, USA, Proceedings, pp. 321–330. IEEE Computer Society (2005)
18. Finkbeiner, B., Schewe, S.: Bounded synthesis. *STTT* **15**(5–6), 519–539 (2013)
19. Frankle, J., Osera, P.M., Walker, D., Zdanczewic, S.: Example-directed synthesis: a type-theoretic interpretation. In: Bodík, R., Majumdar, R. (eds.) *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2016, St. Petersburg, FL, USA, 20–22 January 2016*, pp. 802–815. ACM (2016)
20. Jobstmann, B., Bloem, R.: Optimizations for LTL synthesis. In: *6th International Conference on Formal Methods in Computer-Aided Design, FMCAD 2006, San Jose, California, USA, 12–16 November 2006*, Proceedings, pp. 117–124. IEEE Computer Society (2006)
21. Klein, U.: *Topics in Formal Synthesis and Modeling*. Ph.D. thesis, New York University (2011)
22. Klein, U., Piterman, N., Pnueli, A.: Effective synthesis of asynchronous systems from GR(1) specifications. In: Kuncak, V., Rybalchenko, A. (eds.) *VMCAI 2012. LNCS, vol. 7148*, pp. 283–298. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-27940-9_19
23. Kress-Gazit, H., Pappas, G.J.: Automatic synthesis of robot controllers for tasks with locative prepositions. In: *International Conference on Robotics and Automation (ICRA)*, pp. 3215–3220 (2010)
24. Kupferman, O., Vardi, M.Y.: Safrless decision procedures. In: *Proceedings of FOCS*, pp. 531–540. IEEE (2005)
25. Liu, J., Ozay, N., Topcu, U., Murray, R.M.: Synthesis of reactive switching protocols from temporal logic specifications. *IEEE Trans. Autom. Control* **58**(7), 1771–1785 (2013)
26. Maoz, S., Sa’ar, Y.: AspectLTL: an aspect language for LTL specifications. In: Borba, P., Chiba, S. (eds.) *Proceedings of the 10th International Conference on Aspect-Oriented Software Development, AOSD 2011, Porto de Galinhas, Brazil, 21–25 March 2011*, pp. 19–30. ACM (2011)
27. Maoz, S., Sa’ar, Y.: Assume-guarantee scenarios: semantics and synthesis. In: France, R.B., Kazmeier, J., Breu, R., Atkinson, C. (eds.) *MODELS 2012. LNCS, vol. 7590*, pp. 335–351. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-33666-9_22
28. Piterman, N., Pnueli, A., Sa’ar, Y.: Synthesis of reactive(1) designs. In: Emerson, E.A., Namjoshi, K.S. (eds.) *VMCAI 2006. LNCS, vol. 3855*, pp. 364–380. Springer, Heidelberg (2005). https://doi.org/10.1007/11609773_24
29. Pnueli, A.: The temporal logic of programs. In: *Proceedings of FOCS*, pp. 46–57. IEEE (1977)

30. Pnueli, A., Klein, U.: Synthesis of programs from temporal property specifications. In: 2009 7th IEEE/ACM International Conference on Formal Methods and Models for Co-Design, MEMOCODE 2009, pp. 1–7. IEEE (2009)
31. Pnueli, A., Rosner, R.: On the synthesis of a reactive module. In: POPL, pp. 179–190 (1989)
32. Pnueli, A., Rosner, R.: On the synthesis of an asynchronous reactive module. In: Ausiello, G., Dezani-Ciancaglini, M., Della Rocca, S.R. (eds.) ICALP 1989. LNCS, vol. 372, pp. 652–671. Springer, Heidelberg (1989). <https://doi.org/10.1007/BFb0035790>
33. Pnueli, A., Rosner, R.: Distributed reactive systems are hard to synthesize. In: 31st Annual Symposium on Foundations of Computer Science, St. Louis, Missouri, USA, 22–24 October 1990, vol. II, pp. 746–757. IEEE Computer Society (1990)
34. Pnueli, A., Sa’ar, Y., Zuck, L.D.: JTLV: a framework for developing verification algorithms. In: Touili, T., Cook, B., Jackson, P. (eds.) CAV 2010. LNCS, vol. 6174, pp. 171–174. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-14295-6_18
35. Schewe, S., Finkbeiner, B.: Bounded synthesis. In: Namjoshi, K.S., Yoneda, T., Higashino, T., Okamura, Y. (eds.) ATVA 2007. LNCS, vol. 4762, pp. 474–488. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-75596-8_33
36. Srivastava, S., Gulwani, S., Foster, J.S.: From program verification to program synthesis. In: Hermenegildo, M.V., Palsberg, J. (eds.) Proceedings of the 37th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2010, Madrid, Spain, 17–23 January 2010, pp. 313–326. ACM (2010)
37. Thomas, W.: Facets of synthesis: revisiting Church’s problem. In: de Alfaro, L. (ed.) FoSSaCS 2009. LNCS, vol. 5504, pp. 1–14. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-00596-1_1

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter’s Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter’s Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

