

Promptness and Bounded Fairness in Concurrent and Parameterized Systems [★]

Sven Jacobs¹[0000–0002–9051–4050], Mouhammad Sakr^{1,2}[0000–0002–5160–0327],
and Martin Zimmermann³[0000–0002–8038–2453]

¹ CISA Helmholtz Center for Information Security, Saarbrücken, Germany

² Saarland University, Saarbrücken, Germany

³ University of Liverpool, Liverpool, United Kingdom

Abstract. We investigate the satisfaction of specifications in Prompt Linear Temporal Logic (Prompt-LTL) by concurrent systems. Prompt-LTL is an extension of LTL that allows to specify parametric bounds on the satisfaction of eventualities, thus adding a quantitative aspect to the specification language. We establish a connection between bounded fairness, bounded stutter equivalence, and the satisfaction of Prompt-LTL $\setminus \mathbf{X}$ formulas. Based on this connection, we prove the first cutoff results for different classes of systems with a parametric number of components and quantitative specifications, thereby identifying previously unknown decidable fragments of the parameterized model checking problem.

1 Introduction

Concurrent systems are notoriously hard to get correct, and are therefore a promising application area for formal methods like model checking or synthesis. However, these methods usually give correctness guarantees only for systems with a given, fixed number of components, and the state explosion problem prevents us from using them for systems with a large number of components. To ensure that desired properties hold for systems with a very large or even an *arbitrary* number of components, methods for *parameterized* model checking and synthesis have been devised.

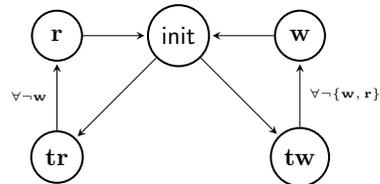
While parameterized model checking is undecidable even for simple safety properties and systems with uniform finite-state components [33], there exist a number of methods that decide the problem for specific classes of systems [2, 10, 12–15, 17, 20, 29], some of which have been collected in surveys of the literature recently [8, 16]. Additionally, there are semi-decision procedures that are successful in many interesting cases [9, 11, 25, 28, 30]. However, most of these approaches only support safety properties, or their support for progress or liveness properties is limited, e.g., because global fairness properties are not considered and cannot be expressed in the supported logic (cp. Außerlechner et al. [5]).

[★] Partially funded by grant EP/S032207/1 from the Engineering and Physical Sciences Research Council (EPSRC).

In this paper, we investigate cases in which we can guarantee that a system with an arbitrary number of components satisfies strong liveness properties, including a quantitative version of liveness called *promptness*. The idea of promptness is that a desired event should not only happen at *some* time in the future, but there should exist a *bound* on the time that can pass before it happens. We consider specifications in Prompt-LTL, an extension of LTL with an operator that expresses prompt eventualities [27], i.e., the logic puts a symbolic bound on the satisfaction of the eventuality, and the model checking problem asks if there is a value for this symbolic bound such that the property is guaranteed to be satisfied with respect to this value. In many settings, adding promptness comes for free in terms of asymptotic complexity [27], e.g., model checking and synthesis [24].⁴ Hence, here we study *parameterized* model checking for Prompt-LTL and show that in many cases adding promptness is also free for this problem.

More precisely, as is common in the analysis of concurrent systems, we abstract concurrency by an interleaving semantics and consider the satisfaction of a specification *up to stuttering*. Therefore, we limit our specifications to Prompt-LTL $\setminus\mathbf{X}$, an extension of the stutter-insensitive logic LTL $\setminus\mathbf{X}$ that does not have the next-time operator. Determining satisfaction of Prompt-LTL $\setminus\mathbf{X}$ specifications by concurrent systems brings new challenges and has not been investigated in detail before.

Motivating Example. For instance, consider the reader-writer protocol on the right which models access to shared data between processes. If a process wants to “read”, it enters the state **tr** (“try-read”) that has a direct transition to the reading state **r**. However, this transition is guarded by $\forall\neg\mathbf{w}$, which stands for the set of all states except **w**, the “writing” state. That is, the transition is only enabled if no other process is currently in state **w**.



Likewise, if a process wants to enter **w** it has to go through **tw**, but the transition to **w** is enabled only if no other process is reading or writing.

For such systems, previous results [5, 14] provide cutoff results for parameterized verification of properties from LTL $\setminus\mathbf{X}$, e.g.,

$$\forall i. \mathbf{G}((tr_i \rightarrow \mathbf{F}r_i) \wedge (tw_i \rightarrow \mathbf{F}w_i)),$$

In this paper we investigate whether the same cutoffs still hold if we consider specifications in Prompt-LTL $\setminus\mathbf{X}$, e.g., if we substitute the LTL eventually operator **F** above with the prompt-eventually operator **F_P**, while imposing a bounded fairness assumption on the scheduler.

⁴ Prompt-LTL can be seen as a fragment of parametric LTL, a logic introduced by Alur et al. [1]. However, since most decision problems for parametric LTL, including model checking, can be reduced to those for Prompt-LTL, we can restrict our attention to the simpler logic.

Contributions. As a first step, we note that Prompt-LTL $\setminus\mathbf{X}$ is not a stutter-insensitive logic, since unbounded stuttering could invalidate a promptness property. This leads us to define the notion of *bounded stutter equivalence*, and proving that Prompt-LTL $\setminus\mathbf{X}$ is *bounded stutter insensitive*.

This observation is then used in an investigation of existing approaches that solve parameterized model checking by the *cutoff* method, which reduces problems from systems with an arbitrary number of components to systems with a fixed number of components. More precisely, these approaches prove that for every trace in a large system, a stutter-equivalent trace in the cutoff system exists, and vice versa. We show that in many cases, modifications of these constructions allow us to obtain traces that are *bounded stutter equivalent*, and therefore the cutoff results extend to specifications in Prompt-LTL $\setminus\mathbf{X}$. The types of systems for which we prove these results include *guarded protocols*, as introduced by Emerson and Kahlon [14], and *token-passing systems*, as introduced by Emerson and Namjoshi [13] for uni-directional rings, and by Clarke et al. [10] for arbitrary topologies. Parameterized model checking for both of these system classes has recently been further investigated [2,3,5,22,31,32], but thus far not in a context that includes promptness properties. Due to space constraints, we omit detailed proofs in some cases. They can be found in the full version of the paper [23].

2 Prompt-LTL $\setminus\mathbf{X}$ and Bounded Stutter Equivalence

We assume that the reader is aware of standard notions such as finite-state transition systems and linear temporal logic (LTL) [6].

We consider concurrent systems that are represented as an interleaving composition of finite-state transition systems, possibly with synchronizing transitions where multiple processes take a step at the same time. In such systems, a process may stay in the same state for many global transitions while other processes are moving. From the perspective of that process, these are *stuttering steps*.

Stuttering is a well-known phenomenon, and temporal languages that include the next-time operator \mathbf{X} are *stutter sensitive*: they can require some atomic proposition to hold at the next moment in time, and the insertion of a stuttering step may change whether the formula is satisfied or not. On the other hand, LTL $\setminus\mathbf{X}$, which does not have the \mathbf{X} operator, is *stutter-insensitive*: two words that only differ in stuttering steps cannot be distinguished by the logic [6].

In the following, we introduce Prompt-LTL $\setminus\mathbf{X}$, an extension of LTL $\setminus\mathbf{X}$, and investigate its properties with respect to stuttering.

2.1 Prompt-LTL $\setminus\mathbf{X}$

Let AP be the set of atomic propositions. The syntax of Prompt-LTL $\setminus\mathbf{X}$ formulas over AP is given by the following grammar:

$$\varphi ::= a \mid \neg a \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \mathbf{F}_p \varphi \mid \varphi \mathbf{U} \varphi \mid \varphi \mathbf{R} \varphi, \text{ where } a \in AP$$

The semantics of Prompt-LTL\X formulas is defined over infinite words $w = w_0w_1 \dots \in (2^{AP})^\omega$, positions $i \in \mathbb{N}$, and bounds $k \in \mathbb{N}$. The semantics of the prompt-eventually operator \mathbf{F}_p is defined as follows:

$$(w, i, k) \models \mathbf{F}_p \varphi \text{ iff there exists } j \text{ such that } i \leq j \leq i + k \text{ and } (w, j, k) \models \varphi.$$

All other operators ignore the bound k and have the same semantics as in LTL, moreover we define \mathbf{F} and \mathbf{G} in terms of \mathbf{U} and \mathbf{R} as usual.

2.2 Prompt-LTL and Stuttering

Our first observation is that Prompt-LTL\X is stutter sensitive: to satisfy the formula $\varphi = \mathbf{GF}_p q$ with respect to a bound k , q has to appear at least once in every k steps. Given a word w that satisfies φ for some bound k , we can construct a word that does not satisfy φ for any bound k by introducing an increasing (and unbounded) number of stuttering steps between every two appearances of q . In the following, we show that Prompt-LTL\X is stutter insensitive if and only if there is a bound on the number of consecutive stuttering steps.

Bounded Stutter Equivalence. A finite word $w \in (2^{AP})^+$ is a *block* if $\exists \alpha \subseteq AP$ such that $w = \alpha^{|w|}$. Two blocks $w, w' \in (2^{AP})^+$ are *d-compatible* if $\exists \alpha \subseteq AP$ such that $w = \alpha^{|w|}, w' = \alpha^{|w'|}, |w| \leq d \cdot |w'|$ and $|w'| \leq d \cdot |w|$. Two infinite sequences of blocks $w_0w_1w_2 \dots, w'_0w'_1w'_2 \dots$ are *d-compatible* if w_i, w'_i are d-compatible for all i .

Two words $w, w' \in (2^{AP})^\omega$ are *d-stutter equivalent*, denoted $w \equiv_d w'$, if they can be written as *d-compatible* sequences of blocks. They are *bounded stutter equivalent* if they are *d-stutter equivalent* for some d . We denote by \hat{w} a sequence of blocks that corresponds to a word w .

Given an infinite sequence of blocks $\hat{w} = w_0, w_1, w_2 \dots$, let $N_i^{\hat{w}} = \{\sum_{l=0}^{i-1} |w_l|, \dots, \sum_{l=0}^{i-1} |w_l| + |w_i| - 1\}$ be the set of positions of the i th block. Given a position n , there is a unique i such that $n \in N_i^{\hat{w}}$.

To prove that Prompt-LTL\X is *bounded stutter insensitive*, i.e., it cannot distinguish two words that are bounded stutter equivalent, we define a function that maps between positions in two *d-compatible* sequences of blocks: given two infinite *d-stutter equivalent* words w, w' such that \hat{w}, \hat{w}' are *d-compatible*, define the function $f : \mathbb{N} \rightarrow 2^{\mathbb{N}}$ where: $f(j) = N_i^{\hat{w}'} \Leftrightarrow j \in N_i^{\hat{w}}$. Note that $\forall j' \in f(j)$ we have $w(j) = w'(j')$, where $w(i)$ denotes the i th symbol in w . For an infinite word w , let $w[i, \infty)$ denote its suffix starting at position i , and $w[i : j]$ its infix starting at i and ending at j . Then we can state the following.

Remark 1. Given two words w and w' , if $w \equiv_d w'$, then $\forall j \in \mathbb{N} \forall j' \in f(j) : w[j, \infty) \equiv_d w'[j', \infty)$.

Now, we can state our first theorem.

Theorem 1 (Prompt-LTL\X is Bounded Stutter Insensitive). *Let w, w' be d-stutter equivalent words, φ a Prompt-LTL\X formula, and f as defined above. Then $\forall i, k \in \mathbb{N}$:*

$$(w, i, k) \models \varphi \Rightarrow \forall j \in f(i) : (w', j, d \cdot k) \models \varphi$$

Proof. The proof works inductively over the structure of φ . Let $\hat{w} = w_0, w_1, \dots$ and $\hat{w}' = w'_0, w'_1, \dots$ be two d -compatible sequences of w and w' . We denote by n_i, m_i the number of elements inside $N_i^w, N_i^{w'}$ respectively. We consider two cases, the other cases are trivial or similar to Case 2:

Case 1: $\varphi = \mathbf{F}_p\varphi$. $(w, i, k) \models \mathbf{F}_p\varphi \Leftrightarrow \exists e, x : i \leq e \leq i + k, e \in N_x^{\hat{w}}$, and $(w, e, k) \models \varphi$ where $(\sum_{l=0}^{x-1} n_l) \leq e < (\sum_{l=0}^x n_l)$. Then by induction hypothesis we have: $\forall j \in f(e) (w', j, d \cdot k) \models \varphi$. Let s be the smallest position in $f(e)$, then $s = \sum_{l=0}^{x-1} m_l$. There exists $y \in \mathbb{N}$ s.t. $i \in N_y^{\hat{w}}$ then $s = \sum_{l=0}^{y-1} m_l + \sum_{l=y}^{x-1} m_l \leq \sum_{l=0}^{y-1} m_l + \sum_{l=y}^{x-1} n_l \cdot d \leq \sum_{l=0}^{y-1} m_l + d \cdot (\sum_{l=y}^{x-1} n_l) \leq \sum_{l=0}^{y-1} m_l + k \cdot d$ (note that $i \in N_y^{\hat{w}}$ and $(w, i, k) \models \mathbf{F}_p\varphi$). As $\sum_{l=0}^{y-1} m_l$ is the smallest position in $f(i)$, then $\forall j \in f(i) : (w', j, d \cdot k) \models \mathbf{F}_p\varphi$.

Case 2: $\varphi = \varphi_1 \mathbf{U} \varphi_2$. $(w, i, k) \models \varphi_1 \mathbf{U} \varphi_2 \Leftrightarrow \exists j \geq i : (w, j, k) \models \varphi_2$ and $\forall e < j : (w, e, k) \models \varphi_1$. Then, by induction hypothesis we have: $\forall e < j \forall l \in f(e) : (w', l, d \cdot k) \models \varphi_1$ and $\forall l \in f(j) : (w', l, d \cdot k) \models \varphi_2$, therefore $\forall j \in f(i) : (w', j, d \cdot k) \models \varphi_1 \mathbf{U} \varphi_2$. \square

Our later proofs will be based on the existence of counterexamples to a given property, and will use the following consequence of Theorem 1.

Corollary 1 *Let w, w' be d -stutter equivalent words, φ a Prompt-LTL\(\mathbf{X}\) formula, and f as defined above. Then $\forall k \in \mathbb{N}$:*

$$(w, i, k) \not\models \varphi \Rightarrow \forall j \in f(i) : (w', j, k/d) \not\models \varphi$$

3 Guarded Protocols and Parameterized Model Checking

In the following, we introduce a system model for concurrent systems, called guarded protocols. However, we will see that some of our results are of interest for other classes of concurrent and parameterized systems, e.g., the token-passing systems that we investigate in Section 6.

3.1 System Model: Guarded Protocols

We consider systems of the form $A \parallel B^n$, consisting of one copy of a process template A and n copies of a process template B , in an interleaving parallel composition. We distinguish objects that belong to different templates by indexing them with the template. E.g., for process template $U \in \{A, B\}$, Q_U is the set of states of U . For this section, fix a finite set of states $Q = Q_A \cup Q_B$ and a positive integer n , and let $\mathcal{G} = \{\exists, \forall\} \times 2^Q$ be the set of guards.

Processes. A *process template* is a transition system $U = (Q_U, \text{init}_U, \delta_U)$ where

- $Q_U \subseteq Q$ is a finite set of states including the initial state init_U ,
- $\delta_U \subseteq Q_U \times \mathcal{G} \times Q_U$ is a guarded transition relation.

Guarded Protocols. The semantics of $A\|B^n$ is given by the transition system $(S, \text{init}_S, \Delta)$, where ⁵

- $S = Q_A \times (Q_B)^n$ is the set of (global) states,
- $\text{init}_S = (\text{init}_A, \text{init}_B, \dots, \text{init}_B)$ is the global initial state, and
- $\Delta \subseteq S \times S$ is the global transition relation. Δ will be defined by local guarded transitions of the process templates A and B in the following.

We distinguish different copies of process template B in $A\|B^n$ by subscript, and each B_i is called a B -process. We denote the set $\{A, B_1, \dots, B_n\}$ as \mathcal{P} , and a process in \mathcal{P} as p . For a global state $s \in S$ and $p \in \mathcal{P}$, let the *local state of p in s* be the projection of s onto that process, denoted $s(p)$.

Then a local transition (q, g, q') of process $p \in \mathcal{P}$ is *enabled* in global state s if $s(p) = q$ and either

- $g = (\exists, G)$ and $\exists p' \in \mathcal{P} \setminus \{p\} : s(p') \in G$, or
- $g = (\forall, G)$ and $\forall p' \in \mathcal{P} \setminus \{p\} : s(p') \in G$.

Finally, $(s, s') \in \Delta$ if there exists $p \in \mathcal{P}$ such that $(s(p), g, s'(p)) \in \delta_p$ is enabled in s , and $s(p') = s'(p')$ for all $p' \in \mathcal{P} \setminus \{p\}$. We say that the transition (s, s') is *based* on the local transition $(s(p), g, s'(p))$ of p .

Disjunctive and Conjunctive Systems. We distinguish disjunctive and conjunctive systems, as defined by Emerson and Kahlon [14]. In a *disjunctive process template*, every guard is of the form (\exists, G) for some $G \subseteq Q$. In a *conjunctive process template*, every guard is of the form (\forall, G) , and $\{\text{init}_A, \text{init}_B\} \subseteq G$, i.e., initial states act as neutral states for all transitions. A *disjunctive (conjunctive) system* consists of only disjunctive (conjunctive) process templates. For conjunctive systems we additionally assume that processes are *initializing*, i.e., any process that moves infinitely often visits its initial state infinitely often.⁶

Runs. A *path* of a system $A\|B^n$ is a sequence $x = s_0 s_1 \dots$ of global states such that for all $i < |x|$ there is a transition $(s_i, s_{i+1}) \in \Delta$ based on a local transition of some process $p \in \mathcal{P}$. We say that p *moves at moment i* . A path can be finite or infinite, and a *maximal path* is a path that cannot be extended, i.e., it is either infinite or ends in a global state where no local transition is enabled, also called a *deadlock*. A *run* is a maximal path starting in init_S . We write $x \in A\|B^n$ to denote that x is a run of $A\|B^n$.

Given a path $x = s_0 s_1 \dots$ and a process p , the *local path* of p in x is the projection $x(p) = s_0(p) s_1(p) \dots$ of x onto local states of p . It is a *local run* of p if x is a run. Additionally we denote by $x(p_1, \dots, p_k)$ the projection $s_0(p_1, \dots, p_k) s_1(p_1, \dots, p_k) \dots$ of x onto the processes $p_1, \dots, p_k \in \mathcal{P}$.

⁵ By similar arguments as in Emerson and Kahlon [14], our results can be extended to systems with an arbitrary (but fixed) number of process templates. The same holds for *open* process templates that can receive inputs from an environment, as considered by Außerlechner et al. [5].

⁶ This restriction has already been considered by Außerlechner et al. [5], and was necessary to support global fairness assumptions.

Fairness. We say a process p is *enabled* in global state s if at least one of its transitions is enabled in s , otherwise it is *disabled*. Then, an infinite run x of a system $A\|B^n$ is

- *strongly fair* if for every process p , if p is enabled infinitely often, then p moves infinitely often.
- *unconditionally fair*, denoted $u\text{-fair}(x)$, if every process moves infinitely often.
- *globally b -bounded fair*, denoted $b\text{-gfair}(x)$, for some $b \in \mathbb{N}$, if

$$\forall p \in \mathcal{P} \forall m \in \mathbb{N} \exists j \in \mathbb{N} : m \leq j \leq m + b \text{ and } p \text{ moves at moment } j.$$

- *locally b -bounded fair* for $E \subseteq \mathcal{P}$, denoted $b\text{-lfair}(x, E)$, if it is unconditionally fair and

$$\forall p \in E \forall m \in \mathbb{N} \exists j \in \mathbb{N} : m \leq j \leq m + b \text{ and } p \text{ moves at moment } j.$$

Bounded-fair System. We consider systems that keep track of bounded fairness explicitly by running in parallel to $A\|B^n$ one counter for each process. In a step of the system where process p moves, the counter of p is reset, and all other counters are incremented. If one of the counters exceeds the bound b , the counter goes into a failure state from which no transition is enabled. We call such a system a *bounded-fair system*, and denote it $A\|_b B^n$.

A *path* of a bounded-fair system $A\|_b B^n$ is given as $x = (s_0, b_0)(s_1, b_1) \dots$, and extends a path of $A\|B^n$ by valuations $b_i \in \{0, \dots, b\}^{n+1}$ of the counters. Note that a run (i.e., a maximal path) of $A\|_b B^n$ is finite iff either it is deadlocked (in which case also its projection to a run of $A\|B^n$ is deadlocked) or a failure state is reached. Thus, the projection of all infinite runs of $A\|_b B^n$ to $A\|B^n$ are exactly the globally b -bounded fair runs of $A\|B^n$.

3.2 Parameterized Model Checking and Cutoffs

Prompt-LTL\X Specifications. Given a system $A\|B^n$, we consider specifications over $AP = Q_A \cup (Q_B \times \{1, \dots, n\})$, i.e., states of processes are used as atomic propositions. For $i_1, \dots, i_c \in \{1, \dots, n\}$, we write $\varphi(A, B_{i_1}, \dots, B_{i_c})$ for a formula that contains only atomic propositions from $Q_A \cup (Q_B \times \{i_1, \dots, i_c\})$.

In the absence of fairness considerations, we say that $A\|B^n$ *satisfies* φ if

$$\exists k \in \mathbb{N} \forall x \in A\|B^n : (x, 0, k) \models \varphi.$$

We say that $A\|B^n$ *satisfies* $\varphi(A, B_1, \dots, B_c)$ *under global bounded fairness*, written $A\|B^n \models_{gb} \varphi(A, B_1, \dots, B_c)$, if

$$\forall b \in \mathbb{N} \exists k \in \mathbb{N} \forall x \in A\|B^n : b\text{-gfair}(x) \Rightarrow (x, 0, k) \models \varphi(A, B_1, \dots, B_c).$$

Finally, for local bounded fairness we usually require bounded fairness for all processes that appear in the formula $\varphi(A, B_1, \dots, B_c)$. Thus, we say that

$A\|B^n$ satisfies $\varphi(A, B_1, \dots, B_c)$ under local bounded fairness, written $A\|B^n \models_{lb} \varphi(A, B_1, \dots, B_c)$, if

$$\forall b \in \mathbb{N} \exists k \in \mathbb{N} \forall x \in A\|B^n : b\text{-lfair}(x, \{1, \dots, c\}) \Rightarrow (x, 0, k) \models \varphi(A, B_1, \dots, B_c).$$

Parameterized Specifications. A *parameterized specification* is a Prompt-LTL $\setminus\mathbf{X}$ formula with quantification over the indices of atomic propositions. A *h -indexed formula* is of the form $\forall i_1, \dots, \forall i_h. \varphi(A, B_{i_1}, \dots, B_{i_h})$. Let $f \in \{gb, lb\}$, then for given $n \geq h$,

$$A\|B^n \models_f \forall i_1, \dots, \forall i_h. \varphi(A, B_{i_1}, \dots, B_{i_h})$$

\Leftrightarrow

$$\forall j_1 \neq \dots \neq j_h \in \{1, \dots, n\} : A\|B^n \models_f \varphi(A, B_{j_1}, \dots, B_{j_h}).$$

By symmetry of guarded protocols, this is equivalent (cp. [14]) to $A\|B^n \models_f \varphi(A, B_1, \dots, B_h)$. The latter formula is denoted by $\varphi(A, B^{(h)})$, and we often use it instead of the original $\forall i_1, \dots, \forall i_h. \varphi(A, B_{i_1}, \dots, B_{i_h})$.

(Parameterized) Model Checking Problems. For $n \in \mathbb{N}$, a specification $\varphi(A, B^{(h)})$ with $n \geq h$, and $f \in \{gb, lb\}$:

- the *model checking problem* is to decide whether $A\|B^n \models_f \varphi(A, B^{(h)})$,
- the *parameterized model checking problem* (PMCP) is to decide whether $\forall m \geq n : A\|B^m \models_f \varphi(A, B^{(h)})$.

Cutoffs and Decidability. We define cutoffs with respect to a class of systems (either disjunctive or conjunctive), a class of process templates P , e.g., templates of bounded size, and a class of properties, e.g. satisfaction of h -indexed Prompt-LTL $\setminus\mathbf{X}$ formulas under a given fairness notion.

A *cutoff* for a given class of systems with processes from P , a fairness notion $f \in \{lb, gb\}$ and a set of Prompt-LTL $\setminus\mathbf{X}$ formulas Φ is a number $c \in \mathbb{N}$ such that

$$\forall A, B \in P \forall \varphi \in \Phi \forall n \geq c : A\|B^n \models_f \varphi \Leftrightarrow A\|B^c \models_f \varphi.$$

Note that the existence of a cutoff implies that the PMCP is *decidable* iff the model checking problem for the cutoff system $A\|B^c$ is decidable. Decidability of model checking for finite transition systems with specifications in Prompt-LTL $\setminus\mathbf{X}$ and bounded fairness follows from the fact that bounded fairness can be expressed in Prompt-LTL $\setminus\mathbf{X}$, and from results on decidability of assume-guarantee model checking for Prompt-LTL (cf. Kupferman et al. [27] and Faymonville and Zimmermann [19][Lemmas 8, 9]).

4 Cutoffs for Disjunctive Systems

In this section, we prove cutoff results for disjunctive systems under bounded fairness and stutter-insensitive specifications with or without promptness. To this end, in Section 4.1 we prove two lemmas that show how to simulate, up to

bounded stuttering, local runs from a system of given size n in a smaller or larger disjunctive system. We then use these two lemmas in Subsections 4.2 and 4.3 to obtain cutoffs for specifications in $LTL \setminus \mathbf{X}$ and Prompt- $LTL \setminus \mathbf{X}$, respectively.

Moreover for the proofs of these two lemmas we utilize the same construction techniques that were used in [4, 5, 14], but in addition we analyze their effects on bounded fairness and bounded stutter equivalence. Note that we will only consider formulas of the form $\varphi(A, B^{(1)})$, however, as in previous work [4, 14], our results extend to specifications over an arbitrary number h of B -processes.

Table 1 summarizes the results of this section: for specifications in $LTL \setminus \mathbf{X}$ and Prompt- $LTL \setminus \mathbf{X}$ we obtain a cutoff that depends on the size of process template B , as well as on the number h of quantified index variables. The table states generalizations of Theorems 2 and 3 from the 2-indexed case to the h -indexed case for arbitrary $h \in \mathbb{N}$. For one of the cases we were not able to obtain a cutoff result (as explained in the full version [23]).

Simple Reader-Writer Example. Consider the disjunctive system $W \parallel R^n$, where W is a writer process (Figure 2), and R is a reader process (Figure 1). Let the specification φ be $\forall i \mathbf{G}(\mathbf{w} \rightarrow \mathbf{F}_p[(\mathbf{w} \wedge \mathbf{nr}_i)])$, i.e., if process W is in state \mathbf{w} , then eventually all the R processes will be in state \mathbf{nr} , while W is in \mathbf{w} . According to Table 1, the cutoff for checking whether $W \parallel R^n \models_{lb} \varphi$ is 5.

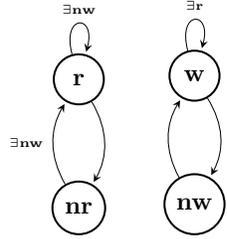


Fig. 1:
Reader

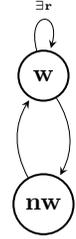


Fig. 2:
Writer

Table 1: Cutoffs for Disjunctive Systems

| | Local Bounded Fairness | Global Bounded Fairness |
|---|------------------------|-------------------------|
| h -indexed $LTL \setminus \mathbf{X}$ | $2 Q_B + h$ | $2 Q_B + h$ |
| h -indexed Prompt- $LTL \setminus \mathbf{X}$ | $2 Q_B + h$ | - |

4.1 Simulation up to Bounded Stutter Equivalence

Definitions. Fix a run $x = x_0x_1\dots$ of the disjunctive system $A \parallel B^n$. Our constructions are based on the following definitions, where $q \in Q_B$:

- $\text{appears}^{B_i}(q)$ is the set of all moments in x where process B_i is in state q : $\text{appears}^{B_i}(q) = \{m \in \mathbb{N} \mid x_m(B_i) = q\}$.
- $\text{appears}(q)$ is the set of all moments in x where at least one B -process is in state q : $\text{appears}(q) = \{m \in \mathbb{N} \mid \exists i \in \{1, \dots, n\} : x_m(B_i) = q\}$.
- f_q is the first moment in x where q appears: $f_q = \min(\text{appears}(q))$, and $\text{first}_q \in \{1, \dots, n\}$ is the index of a B -process where q appears first, i.e., with $x_{f_q}(B_{\text{first}_q}) = q$.
- if $\text{appears}(q)$ is finite, then $l_q = \max(\text{appears}(q))$ is the last moment where q appears, and $\text{last}_q \in \{1, \dots, n\}$ is a process index with $x_{l_q}(B_{\text{last}_q}) = q$

- let $\text{Visited}^{inf} = \{q \in Q_B \mid \exists B_i \in \{B_2, \dots, B_n\} : \text{appears}^{B_i}(q) \text{ is infinite}\}$
and $\text{Visited}^{fin} = \{q \in Q_B \mid \forall B_i \in \{B_2, \dots, B_n\} : \text{appears}^{B_i}(q) \text{ is finite}\}$.
- $\text{Set}(x_i)$ is the set of all state that are visited by some process at moment i :
 $\text{Set}(x_i) = \{q \mid q \in (Q_A \cup Q_B) \text{ and } \exists p \in \mathcal{P} : x_i(p) = q\}$.

Our first lemma states that any behavior of processes A and B_1 in a system $A \parallel B^n$ can be simulated up to bounded stuttering in a system $A \parallel B^{n+1}$. This type of lemma is called a *monotonicity lemma*.

Lemma 1 (Monotonicity Lemma for Bounded Stutter Equivalence).

Let A, B be disjunctive process templates, $n \geq 2$, $b \in \mathbb{N}$ and $x \in A \parallel B^n$ with b -lfair($x, \{A, B_1\}$). Then there exists $y \in A \parallel B^{n+1}$ with $2b$ -lfair($y, \{A, B_1\}$) and $x(A, B_1) \equiv_2 y(A, B_1)$.

Proof. Let x be a run of $A \parallel B^n$ where b -lfair($x, \{A, B_1\}$). Let $y(A) = x(A)$ and $y(B_j) = x(B_j)$ for all $B_j \in \{B_1, \dots, B_n\}$ and let the new process B_{n+1} copy one of the B -processes of $A \parallel B^n$, i.e., $y(B_{n+1}) = x(B_i)$ for some $i \in \{1, \dots, n\}$. Copying a local run violates the interleaving semantics as two processes will be moving at the same time. To solve this problem, we split every transition (y_l, y_{l+1}) where the interleaving semantics is violated by B_i and B_{n+1} executing local transitions (q_i, g, q'_i) and (q_{n+1}, g, q'_{n+1}) , respectively. To do this, replace (y_l, y_{l+1}) with two consecutive transitions $(y_l, u)(u, y_{l+1})$, where (y_l, u) is based on the local transition (q_i, g, q'_i) and (u, y_{l+1}) is based on the local transition (q_{n+1}, g, q'_{n+1}) . Note that both of these local transitions are enabled in the constructed run y since the transition (q_i, g, q'_i) is enabled in the original run x . Moreover, run y inherits unconditional fairness from x . Finally, it is easy to see that for every local transition of process B_i in x , establishing interleaving semantics has added one additional stuttering step to every local run in y including processes A and B_1 . Therefore we have that $2b$ -lfair($y, \{A, B_1\}$) and $x(A, B_1) \equiv_2 y(A, B_1)$. \square

Reader-Writer Example. Consider the run x of the system $W \parallel R^2$ in Figure 3 where W and R are as defined in Figures 1 and 2. We construct a run y of the system $W \parallel R^3$ (see Figure 4) such that $x(W, R_1) \equiv_2 y(W, R_1)$. The local run of process R_3 is obtained by (i) copying the run of R_2 , and (ii) establishing the interleaving semantics as in the proof of Lemma 1.

| t | W | R_1 | R_2 |
|-----|-----------|----------|----------|
| 0 | nw | nr | nr |
| 1 | nw | r | nr |
| 2 | nw | r | r |
| 3 | w | r | r |
| 4 | nw | r | r |
| 5 | nw | r | r |
| 6 | nw | r | r |

| t | W | R_1 | R_2 | R_3 |
|-----|-----------|----------|----------|----------|
| 0 | nw | nr | nr | nr |
| 1 | nw | r | nr | nr |
| 2 | nw | r | r | nr |
| 3 | nw | r | r | r |
| 4 | w | r | r | r |
| 5 | nw | r | r | r |
| 6 | nw | r | r | r |
| 7 | nw | r | r | r |
| 8 | nw | r | r | r |

Fig. 3: Run: $W \parallel R^2$ Fig. 4: Run: $W \parallel R^3$

As mentioned in the above construction, if a local run of x is d -bounded fair for some $d \in \mathbb{N}$, then it will be $2d$ -bounded fair in the constructed run y . This observation leads to the following corollary.

Corollary 2 Let A, B be disjunctive process templates, $n \geq 2$, $b \in \mathbb{N}$ and $x \in A \parallel B^n$ with b -gfair(x). Then there exists $y \in A \parallel B^{n+1}$ with $2b$ -gfair(y) and $x(A, B_1) \equiv_2 y(A, B_1)$.

Our second lemma is a *bounding lemma* which states that any behavior of processes A and B_1 in a disjunctive system $A\|B^n$ can be simulated up to bounded stuttering in a system $A\|B^c$, if c is chosen to be sufficiently large and $n \geq c$.

Lemma 2 (Bounding Lemma for Bounded Stutter Equivalence). *Let A, B be disjunctive process templates, $c = 2|Q_B| + 1$, $n \geq c$, $b \in \mathbb{N}$ and $x \in A\|B^n$ with b -lfair($x, \{A, B_1\}$). Then there exists $y \in A\|B^c$ with $(b \cdot c)$ -lfair($y, \{A, B_1\}$) and $x(A, B_1) \equiv_c y(A, B_1)$.*

Proof. Let x be a run of $A\|B^n$ where b -lfair($x, \{A, B_1\}$). We show how to construct a run y of $A\|B^c$ where $(b \cdot c)$ -lfair($y, \{A, B_1\}$) and $x(A, B_1) \equiv_c y(A, B_1)$.

The basic idea is that, in order to ensure that all transitions in y are enabled at the time they are taken, we “flood” every state q that is visited in x with one or more processes that enter q and stay there. Additionally, we need to take care of fairness, which requires a more complicated construction that allows every such process to move infinitely often. Therefore, some processes have to leave the state they have flooded (if that state only appears finitely often in the original run), and every process needs to eventually enter a loop that allows it to move infinitely often. In the following, we construct such runs formally.

Construction:

1. **(Flooding with evacuation):** To every $q \in \text{Visited}^{fin}(x)$, devote one process B_{i_q} that copies B_{first_q} until the time f_q , then stutters in q until time l_q where it starts copying B_{last_q} forever. Formally:

$$y(B_{i_q}) = x(B_{\text{first}_q})[0 : f_q].(q)^{l_q - f_q}.x(B_{\text{last}_q})[l_q + 1 : \infty]$$

2. **(Flooding with fair extension):** For every $q \in \text{Visited}^{inf}(x)$, let B_q^{inf} be a process that visits q infinitely often in x . We devote to q two processes $B_{i_{q_1}}$ and $B_{i_{q_2}}$ that both copy B_{first_q} until the time f_q , and then stutter in q until B_q^{inf} reaches q for the first time. After that, let $B_{i_{q_1}}$ and $B_{i_{q_2}}$ copy B_q^{inf} in turns as follows: $B_{i_{q_1}}$ copies B_q^{inf} until it reaches q while $B_{i_{q_2}}$ stutters in q , then $B_{i_{q_2}}$ copies B_q^{inf} until it reaches q while $B_{i_{q_1}}$ stutters in q and so on.
3. Establish interleaving semantics as in the proof of Lemma 1.

After steps 1 and 2, the following property holds: at any time t we have that $\text{Set}(x_t) \subseteq \text{Set}(y_t)$, which guarantees that every transition along the run is enabled. Note that establishing the interleaving semantics preserves this property.

Finally, establishing interleaving semantics could introduce additional stuttering steps to the local runs of processes A and B_1 whenever steps 1 or 2 of the construction use the same local run from x more than once (e.g. if $\exists q_i, q_j \in Q_B$ with $\text{first}_{q_i} = \text{first}_{q_j}$). A local run of x can be used in the above construction at most $2|Q_B|$ times, therefore we have $x(A, B_1) \equiv_c y(A, B_1)$. Moreover, since the upper bound of consecutive stuttering steps in A or B_1 is $(2|Q_B| + 1) \cdot b$, we get $(b \cdot c)$ -lfair($y, \{A, B_1\}$). \square

Reader-Writer Example. Consider again the reader-writer system in Figures 1 and 2. For any run x of $W\|R^n$, using the construction above we obtain a run y of $W\|R^5$ (or even a smaller system) with $x(W, R_1) \equiv_5 y(W, R_1)$.

4.2 Cutoffs for Specifications in $LTL \setminus \mathbf{X}$ under Bounded Fairness

The PMCP for disjunctive systems with specifications from $LTL \setminus \mathbf{X}$ has been considered in several previous works [5, 14, 22]. In the following we extend these results by proving cutoff results under bounded fairness.

Theorem 2 (Cutoff for $LTL \setminus \mathbf{X}$ with Global Bounded Fairness). *Let A, B be disjunctive process templates, $c = 2|Q_B| + 1$, $n \geq c$, and $\varphi(A, B^{(1)})$ a specification with $\varphi \in LTL \setminus \mathbf{X}$. Then:*

$$\left(\forall b \in \mathbb{N} : A \parallel_b B^n \models \varphi(A, B^{(1)}) \right) \Leftrightarrow \left(\forall b' \in \mathbb{N} : A \parallel_{b'} B^c \models \varphi(A, B^{(1)}) \right)$$

We prove the theorem by proving two lemmas, one for each direction of the equivalence.

Lemma 3 (Monotonicity Lemma for $LTL \setminus \mathbf{X}$). *Let A, B be disjunctive process templates, $n \geq 1$, and $\varphi(A, B^{(1)})$ a specification with $\varphi \in LTL \setminus \mathbf{X}$. Then:*

$$\left(\exists b \in \mathbb{N} : A \parallel_b B^n \not\models \varphi(A, B^{(1)}) \right) \Rightarrow \left(\exists b' \in \mathbb{N} : A \parallel_{b'} B^{n+1} \not\models \varphi(A, B^{(1)}) \right)$$

Proof. Assume $\exists b \in \mathbb{N} : A \parallel_b B^n \not\models \varphi(A, B^{(1)})$. Then there exists a run x of $A \parallel B^n$ where x is b -gfair(x) and $x \not\models \varphi(A, B^{(1)})$. According to Corollary 2 there exists y of $A \parallel B^{n+1}$ where $2b$ -gfair(y) and $x(A, B_1) \equiv_2 y(A, B_1)$, which guarantees that $y \not\models \varphi(A, B^{(1)})$. \square

For the corresponding bounding lemma, our construction is based on that of Lemma 2. However, the local runs resulting from that construction might stutter in some local states for an unbounded time (e.g. local runs devoted for states in Visited_F^{fin}). To bound stuttering in such constructions, given an arbitrary run of a system $A \parallel B^n$, we first show that whenever there exists a bounded-fair run that violates a specification in $LTL \setminus \mathbf{X}$, then there also exists an ultimately periodic run with the same property.

A (non-deterministic) Büchi automaton is a tuple $\mathcal{A} = (\Sigma, Q_{\mathcal{A}}, \delta, a_0, \alpha)$, where Σ is a finite alphabet, $Q_{\mathcal{A}}$ is a finite set of states, $\delta : Q_{\mathcal{A}} \times \Sigma \rightarrow 2^{Q_{\mathcal{A}}}$ is a transition function, $a_0 \in Q_{\mathcal{A}}$ is an initial state, and $\alpha \subseteq Q_{\mathcal{A}}$ is a Büchi acceptance condition. Given an LTL specification φ , we denote by A_{φ} the Büchi automaton that accepts exactly all words that satisfy φ [34].

Lemma 4 (Ultimately Periodic Counter-Example). *Let $\varphi \in LTL$ and $b \in \mathbb{N}$. If $A \parallel_b B^n \not\models \varphi$ then there exists a run $x = uv^{\omega}$ of $A \parallel B^n$ with b -gfair(x), and $x \not\models \varphi$, where u, v are finite paths, and $|u|, |v| \leq 2 \cdot |Q_{\mathcal{A}}| \cdot |Q_B|^n \cdot b^{n+1} \cdot |Q_{A_{-\varphi}}|$.*

Now, we have all the ingredients to prove the bounding lemma for the case of $LTL \setminus \mathbf{X}$ specifications and (global) bounded fairness.

Lemma 5 (Bounding Lemma for $LTL \setminus \mathbf{X}$). *Let A, B be disjunctive process templates, $c = 2|Q_B| + 1$, $n \geq c$, and $\varphi(A, B^{(1)})$ a specification with $\varphi \in LTL \setminus \mathbf{X}$. Then:*

$$\left(\exists b \in \mathbb{N} : A \parallel_b B^n \not\models \varphi(A, B^{(1)}) \right) \Rightarrow \left(\exists b' \in \mathbb{N} : A \parallel_{b'} B^c \not\models \varphi(A, B^{(1)}) \right)$$

Proof. Assume $\exists b \in \mathbb{N} : A \parallel_b B^n \not\models \varphi(A, B^{(1)})$. Then by Lemma 4 there is a run $x = uv^\omega$ of $A \parallel B^n$, where $b\text{-gfair}(x)$ and $|u|, |v| \leq 2 \cdot |Q_A| \cdot |Q_B|^n \cdot b^{n+1} \cdot |Q_{A \setminus \varphi}|$. According to Lemma 2, we can construct out of x a run y of $A \parallel B^c$ where $b''\text{-lfair}(y, \{A, B_1\})$, and $x(A, B_1) \equiv_d y(A, B_1)$ with $d = 2|Q_B| + 1$ and $b'' = b \cdot d$. The latter guarantees that $y \not\models \varphi(A, B^{(1)})$. We still need to show that $b'\text{-gfair}(y)$ for some $b' \in \mathbb{N}$. As $x = uv^\omega$, we observe that the construction of Lemma 2 ensures the following:

- The number of consecutive stuttering steps per process introduced in step 1 is bounded by $|u|$.
- The number of consecutive stuttering steps introduced in step 2 for a given process is bounded by $|u| + 2|v|$ because B_q^{inf} needs up to $|u| + |v|$ steps to reach q , and one of the processes has to wait for up to $|v|$ additional global steps before it can move.

In addition to the stuttering steps introduced in step 1 and 2, if more than one of the constructed processes simulate the same local run of x then establishing the interleaving semantics would be required, which in turn introduces additional stuttering steps. Therefore the upper bound of consecutive stuttering steps introduced in step 3 of the construction is $(2|Q_B| + 1) \cdot b$. Therefore $b'\text{-gfair}(y)$ where $b' = (2|Q_B| + 1) \cdot b + 6 \cdot |Q_A| \cdot |Q_B|^n \cdot b^{n+1} \cdot |Q_{A \setminus \varphi}|$. \square

Remark 2. With a more complex construction that uses a stutter-insensitive automaton \mathcal{A} [18] to represent the specification and considers runs of the composition of system and automaton, we can obtain a much smaller b' that is also independent of n . This is based on the observation that if in y some process is consecutively stuttering for more than $|A \parallel B^c \times \mathcal{A}|$ steps, then there must be a repetition of states from the product in this time, and we can simply cut the infix between the repeating states from the constructed run y .

4.3 Cutoffs for Specifications in Prompt-LTL \setminus \mathbf{X}

LTL specifications cannot enforce boundedness of the time that elapses before a liveness property is satisfied. Prompt-LTL solves this problem by introducing the prompt-eventually operator explained in Section 2.1. Since we consider concurrent asynchronous systems, the satisfaction of a Prompt-LTL formula can also depend on the scheduling of processes. If scheduling can introduce unbounded delays for a process, then promptness can in general not be guaranteed. Hence, non-trivial Prompt-LTL specifications can *only* be satisfied under the assumption of bounded fairness, and therefore this is the only case we consider here.

Theorem 3 (Cutoff for Prompt-LTL \setminus \mathbf{X} with Local Bounded Fairness). *Let A, B be disjunctive process templates, $c = 2|Q_B| + 1$ $n \geq c$, and $\varphi(A, B^{(1)})$ a specification with $\varphi \in \text{Prompt-LTL} \setminus \mathbf{X}$. Then:*

$$A \parallel B^c \models_{lb} \varphi(A, B^{(1)}) \Leftrightarrow A \parallel B^n \models_{lb} \varphi(A, B^{(1)}).$$

Again, we prove the theorem by proving a monotonicity and a bounding lemma. Note that $A \parallel B^n \not\models_{lb} \varphi(A, B^{(1)})$ iff

$$\exists b \in \mathbb{N} \forall k \in \mathbb{N} \exists x \in A \parallel B^n : b\text{-lfair}(x, \{A, B^{(1)}\}) \wedge (x, 0, k) \not\models \varphi(A, B^{(1)}).$$

Lemma 6 (Monotonicity Lemma for Prompt-LTL\X). *Let A, B be disjunctive process templates, $n \geq 2$, and $\varphi(A, B^{(1)})$ a specification with $\varphi \in \text{Prompt-LTL}\setminus\mathbf{X}$. Then:*

$$A\|B^n \not\models_{lb} \varphi(A, B^{(1)}) \Rightarrow A\|B^{n+1} \not\models_{lb} \varphi(A, B^{(1)}).$$

Proof. Assume $A\|B^n \not\models_{lb} \varphi(A, B^{(1)})$. Then there exists $b \in \mathbb{N}$ such that $\forall k \in \mathbb{N}$ there is a run x of $A\|B^n$ where b -lfair($x, \{A, B^{(1)}\}$), and $(x, 0, 2 \cdot k) \not\models \varphi(A, B^{(1)})$. Then according to Lemma 1 there exists y of $A\|B^{n+1}$ where $2b$ -lfair($y, \{A, B^{(1)}\}$) and $x(A, B_1) \equiv_2 y(A, B_1)$, which guarantees, according to Corollary 1, that $(y, 0, k) \not\models \varphi(A, B^{(1)})$. As a consequence there exists $b \in \mathbb{N}$ such that $\forall k \in \mathbb{N}$ there is a run y of $A\|B^c$ where $2b$ -lfair($y, \{A, B^{(1)}\}$) and $(y, 0, k) \not\models \varphi(A, B^{(1)})$, thus $A\|B^c \not\models_{lb} \varphi(A, B^{(1)})$. \square

Using the same argument of the above proof but by using Corollary 2 instead of Lemma 1 to construct the globally bounded fair counter example, we obtain the following:

Corollary 3 *Let A, B be disjunctive process templates, $n \geq 2$, and $\varphi(A, B^{(1)})$ a specification with $\varphi \in \text{Prompt-LTL}\setminus\mathbf{X}$. Then:*

$$A\|B^n \not\models_{gb} \varphi(A, B^{(1)}) \Rightarrow A\|B^{n+1} \not\models_{gb} \varphi(A, B^{(1)}).$$

Lemma 7 (Bounding Lemma for Prompt-LTL\X). *Let A, B be disjunctive process templates, $c = 2|Q_B| + 1$, $n \geq c$, and $\varphi(A, B^{(1)})$ a specification with $\varphi \in \text{Prompt-LTL}\setminus\mathbf{X}$. Then:*

$$A\|B^n \not\models_{lb} \varphi(A, B^{(1)}) \Rightarrow A\|B^c \not\models_{lb} \varphi(A, B^{(1)}).$$

Proof. Assume $A\|B^n \not\models_{lb} \varphi(A, B^{(1)})$. Then there exists $b \in \mathbb{N}$ such that $\forall k \in \mathbb{N}$ there is a run x of $A\|B^n$ where b -lfair($x, \{A, B^{(1)}\}$) and $(x, 0, d \cdot k) \not\models \varphi(A, B^{(1)})$ with $d = (2|Q_B| + 1)$. According to Lemma 2 we can construct for every such x a run y of $A\|B^c$ where $(d \cdot b)$ -lfair($y, \{A, B^{(1)}\}$), and $x(A, B_1) \equiv_d y(A, B_1)$, which guarantees that $(y, 0, k) \not\models \varphi(A, B^{(1)})$ (see Corollary 1). Thus, there exists $b \in \mathbb{N}$ such that $\forall k \in \mathbb{N}$ there is a run y of $A\|B^c$ where $(d \cdot b)$ -lfair($y, \{A, B^{(1)}\}$) and $(y, 0, k) \not\models \varphi(A, B^{(1)})$, thus $A\|B^c \not\models_{lb} \varphi(A, B^{(1)})$. \square

5 Cutoffs for Conjunctive Systems

In this section we investigate cutoff results for conjunctive systems under bounded fairness and specifications in Prompt-LTL\X. Table 2 summarizes the results of this section, as generalizations of Theorems 4 and 5 to h -indexed specifications. Note that for results marked with a * we require processes to be *bounded initializing*, i.e., that every cycle in the process template contains the initial state.⁷

⁷ This is only slightly more restrictive than the assumption that they are initializing, as stated in the definition of conjunctive systems in Section 3.1.

Table 2: Cutoffs for Conjunctive Systems

| | Local Bounded Fairness | Global Bounded Fairness |
|--|------------------------|-------------------------|
| h -indexed $LTL \setminus \mathbf{X}$ | $h + 1$ | $h + 1^*$ |
| h -indexed $Prompt-LTL \setminus \mathbf{X}$ | $h + 1$ | $h + 1^*$ |

5.1 Cutoffs under Local Bounded Fairness

Theorem 4 (Cutoff for Prompt-LTL $\setminus \mathbf{X}$ with Local Bounded Fairness).

Let A, B be conjunctive process templates, $n \geq 2$, and $\varphi(A, B^{(1)})$ a specification with $\varphi \in \text{Prompt-LTL} \setminus \mathbf{X}$. Then:

$$A \parallel B^2 \models_{lb} \varphi(A, B^{(1)}) \Leftrightarrow A \parallel B^n \models_{lb} \varphi(A, B^{(1)}).$$

We prove the theorem by proving two lemmas, one for each direction of the equivalence. Note that $A \parallel B^n \not\models_{lb} \varphi(A, B^{(1)})$ iff $\exists b \in \mathbb{N} \forall k \in \mathbb{N} \exists x \in A \parallel B^n : b\text{-gfair}(x) \wedge (x, 0, k) \not\models \varphi(A, B^{(1)})$.

Lemma 8 (Monotonicity Lemma, Prompt-LTL $\setminus \mathbf{X}$ with Local Bounded Fairness). Let A, B be conjunctive process templates, $n \geq 2$, and $\varphi(A, B^{(1)})$ a specification with $\varphi \in \text{Prompt-LTL} \setminus \mathbf{X}$. Then:

$$A \parallel B^n \not\models_{lb} \varphi(A, B^{(1)}) \Rightarrow A \parallel B^{n+1} \not\models_{lb} \varphi(A, B^{(1)}).$$

Proof. Assume $A \parallel B^n \not\models_{lb} \varphi(A, B^{(1)})$. Then there exists $b \in \mathbb{N}$ such that $\forall k \in \mathbb{N}$ there is a run x of $A \parallel B^n$ where $b\text{-gfair}(x)$ and $(x, 0, k) \not\models \varphi(A, B^{(1)})$. For every such x , we construct a run y of $A \parallel B^{n+1}$ with $b\text{-lfair}(y)$ and $(y, 0, k) \not\models \varphi(A, B^{(1)})$. Let $y(A) = x(A)$ and $y(B_j) = x(B_j)$ for all $B_j \in \{B_1, \dots, B_n\}$ and let the new process B_{n+1} "share" a local run $x(B_i)$ with an existing process B_i of $A \parallel B^{n+1}$ in the following way: one process stutters in $init_B$ while the other makes transitions from $x(B_i)$, and whenever $x(B_i)$ enters $init_B$ the roles are reversed. Since this changes the behavior of B_i , B_i cannot be a process that is mentioned in the formula, i.e. we need $n \geq 2$ for a formula $\varphi(A, B^{(1)})$. Then we have $b\text{-lfair}(y, \{A, B_1\})$ as the run of B_{n+1} inherits the unconditional fairness behavior from the local run of the process B_i in x . Note that it is not guaranteed that the local runs $y(B_i)$ and $y(B_{n+1})$ are bounded fair as the time between two occurrences of $init_B$ in $x(B_i)$ is not bounded. Moreover we have $x(A, B_1) \equiv_1 y(A, B_1)$, which according to Corollary 1 implies $(y(A, B_1), k) \not\models \varphi(A, B^{(1)})$. \square

Lemma 9 (Bounding Lemma, Prompt-LTL $\setminus \mathbf{X}$, Local Bounded Fairness). Let A, B be conjunctive process templates, $n \geq 1$, and $\varphi(A, B^{(1)})$ a specification with $\varphi \in \text{Prompt-LTL} \setminus \mathbf{X}$. Then:

$$A \parallel B^n \not\models_{lb} \varphi(A, B^{(1)}) \Rightarrow A \parallel B^1 \not\models_{lb} \varphi(A, B^{(1)}).$$

Proof. Assume $A \parallel B^n \not\models_{lb} \varphi(A, B^{(1)})$. Then there exists $b \in \mathbb{N}$ such that $\forall k \in \mathbb{N}$ there is a run x of $A \parallel B^n$ where $b\text{-gfair}(x)$, and $(x, 0, b \cdot k) \not\models \varphi(A, B^{(1)})$. For every such x , we construct a run y in the cutoff system $A \parallel B^1$ by copying the local runs of processes A and B_1 in x and deleting stuttering steps. It is easy to see that $b\text{-gfair}(y)$ then we have $x(A, B_1) \equiv_b y(A, B_1)$, and by Corollary 1 $(y(A, B_1), k) \not\models \varphi(A, B^{(1)})$. \square

Note that this is the same proof construction as in Außerlechner et al. [5], and we simply observe that this construction preserves bounded fairness.

5.2 Cutoffs under Global Bounded Fairness

As mentioned before, to obtain a result that preserves global bounded fairness, we need to restrict process template B to be bounded initializing.

Theorem 5 (Cutoff for Prompt-LTL\X with Global Bounded Fairness). *Let A, B be conjunctive process templates, where B is bounded initializing, $n \geq 2$, and $\varphi(A, B^{(1)})$ a specification with $\varphi \in \text{Prompt-LTL}\backslash\mathbf{X}$. Then:*

$$A\|B^2 \models_{gb} \varphi(A, B^{(1)}) \Leftrightarrow A\|B^n \models_{gb} \varphi(A, B^{(1)}).$$

Again, the theorem can be separated into two lemmas.

Lemma 10 (Monotonicity Lemma, Prompt-LTL\X, Global Bounded Fairness). *Let A, B be conjunctive process templates, where B is bounded initializing, $n \geq 2$, and $\varphi(A, B^{(1)})$ a specification with $\varphi \in \text{Prompt-LTL}\backslash\mathbf{X}$. Then:*

$$A\|B^n \not\models_{gb} \varphi(A, B^{(1)}) \Rightarrow A\|B^{n+1} \not\models_{gb} \varphi(A, B^{(1)}).$$

Proof. Assume $A\|B^n \not\models_{gb} \varphi(A, B^{(1)})$. Then there exists $b \in \mathbb{N}$ such that $\forall k \in \mathbb{N}$ there is a run x of $A\|B^n$ where b -gfair(x), and $(x, 0, (b + |Q_B|) \cdot k) \not\models \varphi(A, B^{(1)})$. For every such x , we construct a run y of $A\|B^{n+1}$ in the same way we did in the proof of Lemma 8. Then we have b' -gfair(y) with $b' = b + |Q_B|$ as $init_B$ is on every cycle of the process template B . Moreover we have $x(A, B_1) \equiv_1 y(A, B_1)$ which according to Corollary 1 implies that $(y(A, B_1), k) \not\models \varphi(A, B^{(1)})$. \square

Lemma 11 (Bounding Lemma, Prompt-LTL\X, Global Bounded Fairness). *Let A, B be conjunctive process templates, where B is bounded initializing, $n \geq 1$, and $\varphi(A, B^{(1)})$ a specification with $\varphi \in \text{Prompt-LTL}\backslash\mathbf{X}$. Then:*

$$A\|B^n \not\models_{gb} \varphi(A, B^{(1)}) \Rightarrow A\|B^1 \not\models_{gb} \varphi(A, B^{(1)}).$$

Proof. Under the given assumptions, we can observe that the construction from Lemma 9 also preserves global bounded fairness.

6 Token Passing Systems

In this section, we first introduce a system model for token passing systems and then show how to obtain cutoff results for this class of systems.

6.1 System Model

Processes. A *token passing process* is a transition system $T = (Q_T, I_T, \Sigma_T, \delta)$ where

- $Q_T = \overline{Q_T} \times \{0, 1\}$ is a finite set of states. $\overline{Q_T}$ is a finite non-empty set. The boolean component $\{0, 1\}$ indicates the possession of the token.
- I_T is the set of initial states with $I_T \cap (\overline{Q_T} \times \{0\}) \neq \emptyset$ and $I_T \cap (\overline{Q_T} \times \{1\}) \neq \emptyset$.
- $\Sigma_T = \{\epsilon, rcv, snd\}$ is the set of actions, where ϵ is an asynchronous action, and $\{rcv, snd\}$ are the actions to receive and send the token.
- $\delta_T = Q_T \times \Sigma_T \times Q_T$ is a transition relation, such that $((q, b), a, (q', b')) \in \delta_T$ iff all of the following hold:
 - $a = \epsilon \Rightarrow b = b'$.
 - $a = snd \Rightarrow b = 1$ and $b' = 0$
 - $a = rcv \Rightarrow b = 0$ and $b' = 1$

Token Passing System. Let $G = (V, E)$ be a finite directed graph without self loops where $V = \{1, \dots, n\}$ is the set of vertices, and $E \subseteq V \times V$ is the set of edges. A *token passing system* T_G^n is a concurrent system containing n instances of process T where the only synchronization between the processes is the sending/receiving of a token according to the graph G . Formally, $T_G^n = (S, init_S, \Delta)$ with:

- $S = (Q_T)^n$.
- $init_S = \{s \in (I_T)^n \text{ such that exactly one process holds the token }\}$,
- $\Delta \subseteq S \times S$ such that $((q_1, \dots, q_n), (q'_1, \dots, q'_n)) \in \Delta$ iff:
 - **Asynchronous Transition.** $\exists i \in V$ such that $(q_i, \epsilon, q'_i) \in \delta_T$, and $\forall j \neq i$ we have $q_j = q'_j$.
 - **Synchronous Transition.** $\exists (i, j) \in E$ such that $(q_i, snd, q'_i) \in \delta_T$, $(q_j, rcv, q'_j) \in \delta_T$, and $\forall z \in V \setminus \{i, j\}$ we have $q_z = q'_z$.

Runs. A *configuration* of a system T_G^n is a tuple (s, ac) where $s \in S$, and either $ac = a_i$ with $a \in \Sigma_T$, and $i \in V$ is a process index, or $ac = (snd_i, rcv_j)$ where $i, j \in V$ are two process indices with $i \neq j$. A run is an infinite sequence of configurations $x = (s_0, ac_0)(s_1, ac_1) \dots$ where $s_0 \in init_S$ and s_{i+1} results from executing action ac_i in s_i . Additionally we denote by $x(i, \dots, j)$ the projection $(s_0(i, \dots, j), ac_0(i, \dots, j))(s_1(i, \dots, j), ac_1(i, \dots, j)) \dots$ where $s_e(i, \dots, j)$ is the projection of s_e on the local states of (T_i, \dots, T_j) and

$$ac(i, \dots, j) = \begin{cases} \perp & \text{if } ac = a_m \text{ and } m \notin \{i, \dots, j\} \\ \perp & \text{if } ac = (snd_m, rcv_n) \text{ and } m, n \notin \{i, \dots, j\} \\ ac & \text{otherwise} \end{cases}$$

Bounded Fairness. A run x of a token passing system T_G^n is b -gfair(x) if for every moment m and every process T_i , T_i receives the token at least once between moments m and $m + b$.

Cutoffs for Complex Networks. In the presence of different network topologies, represented by the graph G , we define a cutoff to be a bound on the size of G that is sufficient to decide the PMCP. Note that, in order to obtain a decision

procedure for the PMCP, we not only need to know the size of the graphs, but also which graphs of this size we need to investigate. This is straightforward if the graph always falls into a simple class, such as rings, cliques, or stars, but is more challenging if the graph can become more complex with increasing size.

6.2 Cutoff Results for Token Passing Systems

Table 3 summarizes the results of this section, generalizing Theorem 6 to the case of h -indexed specifications. Similar to previous sections, the specifications are over states of processes. The results for local bounded fairness follow from the results for global bounded fairness.

To prove the results of this section, we need some additional definitions.

Table 3: Cutoff Results for Token Passing Systems

| | Local Bounded Fairness | Global Bounded Fairness |
|---|------------------------|-------------------------|
| h -indexed LTL $\setminus\mathbf{X}$ | $2h$ | $2h$ |
| h -indexed Prompt-LTL $\setminus\mathbf{X}$ | $2h$ | $2h$ |

Connectivity vector [10]. Given two indices $i, j \in V$ in a finite directed graph G , we define the connectivity vector $v(G, i, j) = (u_1, u_2, u_3, u_4, u_5, u_6)$ as follows:

- $u_1 = 1$ if there is a non-empty path from i to i that does not contain j .
 $u_1 = 0$ otherwise.
- $u_2 = 1$ if there is a path from i to j via vertices different from i and j . $u_2 = 0$ otherwise.
- $u_3 = 1$ if there is a direct edge from i to j . $u_3 = 0$ otherwise.
- u_4, u_5, u_6 are defined like u_1, u_2, u_3 , respectively where i is replaced by j and vice versa.

Immediately Sends. Given a token passing process T , we fix two local states q^{snd} and q^{rcv} , such that there is (i) a local path q^{init}, \dots, q^{rcv} where $q^{init} \in I_T \cap (\overline{Q_T} \times \{0\})$, (ii) a local path q^{rcv}, \dots, q^{snd} that starts with a receive action, and (iii) a local path q^{snd}, \dots, q^{rcv} that starts with a send action.

When constructing a local run for a process T_i that is currently in local state q^{rcv} , we say that T_i *immediately sends the token* if and only if:

1. T_i executes consecutively all the actions on a simple path q^{rcv}, \dots, q^{snd} , then sends the token, and then executes consecutively all the actions on a simple path q^{snd}, \dots, q^{rcv} .
2. All other processes remain idle until T_i reaches q^{rcv} .

Note that, when T_i *immediately sends the token*, it executes at most $|Q_T|$ actions, since the two paths cannot share any states except q^{rcv} and q^{snd} .

Theorem 6 (Cutoff for Prompt-LTL\X). *Let T_G^n be a token-passing system, $g, h \in V$, and $\varphi(T_g, T_h)$ a specification with $\varphi \in \text{Prompt-LTL}\setminus\mathbf{X}$. Then there exists a system $T_{G'}^A$ with $G' = (V', E')$ and $i, j \in V'$ such that $v(G, g, h) = v(G', i, j)$, and*

$$T_G^n \models_{gb} \varphi(T_g, T_h) \Leftrightarrow T_{G'}^A \models_{gb} \varphi(T_i, T_j).$$

We prove the theorem by proving two lemmas, one for each direction of the equivalence. Note that $T_G^n \not\models_{gb} \varphi(T_g, T_h)$ iff $\exists b \in \mathbb{N} \forall k \in \mathbb{N} \exists x \in T_G^n : b\text{-gfair}(x) \wedge (x, 0, k) \not\models \varphi(T_g, T_h)$.

Lemma 12 (Monotonicity Lemma). *Let T_G^n be a token-passing system with $n \geq 3$ and $g, h \in V$, and $\varphi(T_g, T_h)$ a specification with $\varphi \in \text{Prompt-LTL}\setminus\mathbf{X}$. Then there exists a system $T_{G'}^{n+1}$ with $G' = (V', E')$ and $i, j \in V'$ such that $v(G, g, h) = v(G', i, j)$ and*

$$T_G^n \not\models_{gb} \varphi(T_g, T_h) \Rightarrow T_{G'}^{n+1} \not\models_{gb} \varphi(T_i, T_j).$$

Proof. Let a be a vertex of G with $a \notin \{g, h\}$. Then we construct G' from G as follows: Let $V' = V \cup \{n+1\}$, and $E' = (E \cup \{(n+1, m) \mid (a, m) \in E \text{ for some } m \in V\} \cup \{(a, n+1)\}) \setminus \{(a, m) \mid (a, m) \in E \text{ for some } m \in V\}$, i.e. we copy all the outgoing edges of a to the vertex $n+1$, and replace all the outgoing edges of a by one outgoing edge to $n+1$.

Assume $T_G^n \not\models_{gb} \varphi(T_g, T_h)$. Then there exists $b \in \mathbb{N}$ such that $\forall k' \in \mathbb{N}$ there is a run x of T_G^n where $b\text{-gfair}(x)$, and $(x, 0, |Q_T| \cdot k') \not\models \varphi(T_g, T_h)$. Let $b' = b + (b - n + 2) \cdot |Q_T|$, and $d = |Q_T| + 1$. We will construct for every such run x a run y of $T_{G'}^{n+1}$ where $b'\text{-gfair}(y)$, and $x(T_g, T_h) \equiv_d y(T_i, T_j)$ which guarantees that $(y, 0, k') \not\models \varphi(T_i, T_j)$ (see Corollary 1).

Construction. The construction is such that we keep the local paths of the n existing processes up to bounded stuttering, and we add a process T_{n+1} that always immediately sends the token after receiving it, with q^{rcv}, q^{snd} and the corresponding paths as defined above. In the following, as a short-hand notation, if $s = (q_1, \dots, q_n)$ is a global state of T_G^n and $q \in Q_T$, we write (s, q) for (q_1, \dots, q_n, q) .

Let $x = (s_0, ac_0)(s_1, ac_1) \dots$ and $y' = ((s_0, q^{rcv}), ac_0)((s_1, q^{rcv}), ac_1) \dots$. Note that y' is a sequence of configurations of $T_{G'}^{n+1}$, but not a run. To obtain a run, first let $y'' = ((s_0, q^{\text{init}}), \epsilon) \dots ((s_0, q^{rcv}), ac_0)((s_1, q^{rcv}), ac_1) \dots$. Finally, replace every occurrence of a pair of consecutive configurations $((s, q^{rcv}), (snd_a, rcv_z))$, $((s', q^{rcv}), ac')$, where $s, s' \in Q_T^n, z \in V, ac' \in \Sigma$, with the sequence $((s, q^{rcv}), (snd_a, rcv_{n+1})) \dots ((s, q^{snd}), (snd_{n+1}, rcv_z)) \dots ((s', q^{rcv}), ac')$.

In other words, instead of sending the token to T_z , T_a sends the token to T_{n+1} , and T_{n+1} sends the token immediately to T_z . Furthermore, in x between moments t and $t+b$, T_a can send the token at most $b-n+1$ times, and whenever T_{n+1} receives the token, it takes at most $|Q_T|$ steps before reaching q^{rcv} again. Finally, note that the number of steps T_{n+1} takes to reach q^{rcv} for the first time is also bounded by $|Q_T|$. Therefore we have $b'\text{-gfair}(y)$ and $x(T_g, T_h) \equiv_d y(T_i, T_j)$ (as $b' \leq b \cdot d$) which by Corollary 1 implies that $(y, 0, k') \not\models \varphi(T_i, T_j)$. \square

Lemma 13 (Bounding Lemma). *Let T_G^n be a system with $n \geq 4$ and $g, h \in V$, and $\varphi(T_g, T_h)$ a specification with $\varphi \in \text{Prompt-LTL} \setminus \mathbf{X}$. Then there exists a system $T_{G'}^4$ with $G' = (V', E')$ and $i, j \in V'$ such that $v(G, g, h) = v(G', i, j)$ and*

$$T_G^n \not\models_{gb} \varphi(T_g, T_h) \Rightarrow T_{G'}^4 \not\models_{gb} \varphi(T_i, T_j).$$

Proof (Proof idea). First, note that the existence of G' and $i, j \in V'$ with $v(G, g, h) = v(G', i, j)$ follows directly from Proposition 1 in Clarke et al. [10]. As usual, assuming that $T_G^n \not\models_{gb} \varphi(T_g, T_h)$, we need to construct counterexample runs of $T_{G'}^4$ for some $b' \in \mathbb{N}$ and all $k' \in \mathbb{N}$.

The construction is based on the same ideas as in the proof of Lemma 12, with the following modifications: i) instead of keeping all local runs of a run $x \in T_G^n$, we only keep the local runs of T_g and T_h (now assigned to T_i and T_j), ii) instead of constructing one local run for the new process, we now construct local runs for two new processes T_k and T_l (basically, each of them is responsible for passing the token to T_i or T_j , respectively), and iii) the details of the construction of these runs depend on the connectivity vector $v(G, g, h)$, which essentially determines which of the new processes holds the token when neither T_i nor T_j have it.

As usual, the construction ensures that y is globally bounded fair and that $y(T_i, T_j) \equiv_d x(T_g, T_h)$ for some d , which by Corollary 1 implies that $(y, 0, k') \not\models \varphi(T_i, T_j)$. \square

7 Conclusions

We have investigated the behavior of concurrent systems with respect to promptness properties specified in $\text{Prompt-LTL} \setminus \mathbf{X}$. Our first important observation is that $\text{Prompt-LTL} \setminus \mathbf{X}$ is not stutter insensitive, so the standard notion of stutter equivalence is insufficient to compare traces of concurrent systems if we are interested in promptness. Based on this, we have defined *bounded stutter equivalence*, and have shown that $\text{Prompt-LTL} \setminus \mathbf{X}$ is *bounded stutter insensitive*.

We have shown how this allows us to obtain cutoff results for guarded protocols and token-passing systems, and have obtained cutoffs for $\text{Prompt-LTL} \setminus \mathbf{X}$ (with locally or globally bounded fairness) that are the same as those that were previously shown for $\text{LTL} \setminus \mathbf{X}$ (with unbounded fairness). This implies that, for the cases where we do obtain cutoffs, the PMCP for $\text{Prompt-LTL} \setminus \mathbf{X}$ has the same asymptotic complexity as the PMCP for $\text{LTL} \setminus \mathbf{X}$.

One case that we investigated remains open: disjunctive systems with global bounded fairness. In future work, we will try to solve this open problem, and investigate whether other cutoff results in the literature can also be lifted from $\text{LTL} \setminus \mathbf{X}$ to $\text{Prompt-LTL} \setminus \mathbf{X}$.

Finally, we note that together with methods for distributed synthesis from $\text{Prompt-LTL} \setminus \mathbf{X}$ specifications, our cutoff results enable the synthesis of parameterized systems based on the *parameterized synthesis* approach [21] that has been used to solve challenging synthesis benchmarks by reducing them to systems with a small number of components [7, 26].

References

1. Alur, R., Etessami, K., La Torre, S., Peled, D.A.: Parametric temporal logic for "model measuring". *ACM Trans. Comput. Log.* **2**(3), 388–407 (2001). <https://doi.org/10.1145/377978.377990>, <https://doi.org/10.1145/377978.377990>
2. Aminof, B., Jacobs, S., Khalimov, A., Rubin, S.: Parameterized model checking of token-passing systems. In: *VMCAI. LNCS*, vol. 8318, pp. 262–281. Springer (2014). https://doi.org/10.1007/978-3-642-54013-4_15
3. Aminof, B., Kotek, T., Rubin, S., Spegni, F., Veith, H.: Parameterized model checking of rendezvous systems. *Distributed Computing* **31**(3), 187–222 (2018). <https://doi.org/10.1007/s00446-017-0302-6>
4. Außerlechner, S., Jacobs, S., Khalimov, A.: Tight cutoffs for guarded protocols with fairness. *CoRR abs/1505.03273* (2015), <http://arxiv.org/abs/1505.03273>
5. Außerlechner, S., Jacobs, S., Khalimov, A.: Tight cutoffs for guarded protocols with fairness. In: *VMCAI. LNCS*, vol. 9583, pp. 476–494. Springer (2016). https://doi.org/10.1007/978-3-662-49122-5_23
6. Baier, C., Katoen, J.P.: *Principles of model checking*, vol. 26202649. MIT press Cambridge (2008)
7. Bloem, R., Jacobs, S., Khalimov, A.: Parameterized synthesis case study: AMBA AHB. In: *SYNT. EPTCS*, vol. 157, pp. 68–83 (2014). <https://doi.org/10.4204/EPTCS.157.9>
8. Bloem, R., Jacobs, S., Khalimov, A., Konnov, I., Rubin, S., Veith, H., Widder, J.: *Decidability of Parameterized Verification. Synthesis Lectures on Distributed Computing Theory*, Morgan & Claypool Publishers (2015). <https://doi.org/10.2200/S00658ED1V01Y201508DCT013>
9. Bouajjani, A., Jonsson, B., Nilsson, M., Touili, T.: Regular model checking. In: *CAV. LNCS*, vol. 1855, pp. 403–418. Springer (2000). https://doi.org/10.1007/10722167_31
10. Clarke, E.M., Talupur, M., Touili, T., Veith, H.: Verification by network decomposition. In: *CONCUR. LNCS*, vol. 3170, pp. 276–291. Springer (2004). https://doi.org/10.1007/978-3-540-28644-8_18
11. Clarke, E.M., Talupur, M., Veith, H.: Proving ptolemy right: The environment abstraction framework for model checking concurrent systems. In: *TACAS. LNCS*, vol. 4963, pp. 33–47. Springer (2008). https://doi.org/10.1007/978-3-540-78800-3_4
12. Emerson, E.A., Kahlon, V.: Model checking guarded protocols. In: *LICS*. pp. 361–370. IEEE Computer Society (2003). <https://doi.org/10.1109/LICS.2003.1210076>
13. Emerson, E.A., Namjoshi, K.S.: On reasoning about rings. *Foundations of Computer Science* **14**(4), 527–549 (2003). <https://doi.org/10.1142/S0129054103001881>
14. Emerson, E.A., Kahlon, V.: Reducing model checking of the many to the few. In: *CADE. LNCS*, vol. 1831, pp. 236–254. Springer (2000). https://doi.org/10.1007/10721959_19
15. Esparza, J., Finkel, A., Mayr, R.: On the verification of broadcast protocols. In: *LICS*. pp. 352–359. IEEE Computer Society (1999). <https://doi.org/10.1109/LICS.1999.782630>
16. Esparza, J.: Keeping a crowd safe: On the complexity of parameterized verification (invited talk). In: *STACS. LIPIcs*, vol. 25, pp. 1–10. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik (2014). <https://doi.org/10.4230/LIPIcs.STACS.2014.1>
17. Esparza, J., Ganty, P., Majumdar, R.: Parameterized verification of asynchronous shared-memory systems. *J. ACM* **63**(1), 10:1–10:48 (2016). <https://doi.org/10.1145/2842603>

18. Etesami, K.: Stutter-invariant languages, ω -automata, and temporal logic. In: International Conference on Computer Aided Verification. pp. 236–248. Springer (1999)
19. Faymonville, P., Zimmermann, M.: Parametric linear dynamic logic. *Inf. Comput.* **253**, 237–256 (2017). <https://doi.org/10.1016/j.ic.2016.07.009>, <https://doi.org/10.1016/j.ic.2016.07.009>
20. German, S.M., Sistla, A.P.: Reasoning about systems with many processes. *J. ACM* **39**(3), 675–735 (1992). <https://doi.org/10.1145/146637.146681>
21. Jacobs, S., Bloem, R.: Parameterized synthesis. *Logical Methods in Computer Science* **10**, 1–29 (2014). [https://doi.org/10.2168/LMCS-10\(1:12\)2014](https://doi.org/10.2168/LMCS-10(1:12)2014)
22. Jacobs, S., Sakr, M.: Analyzing guarded protocols: Better cutoffs, more systems, more expressivity. In: VMCAI. *Lecture Notes in Computer Science*, vol. 10747, pp. 247–268. Springer (2018). https://doi.org/10.1007/978-3-319-73721-8_12
23. Jacobs, S., Sakr, M., Zimmermann, M.: Promptness and bounded fairness in concurrent and parameterized systems. *CoRR* **abs/1911.03122** (2019), <http://arxiv.org/abs/1911.03122>
24. Jacobs, S., Tentrup, L., Zimmermann, M.: Distributed synthesis for parameterized temporal logics. *Information and Computation* **262**, 311–328 (2018). <https://doi.org/10.1016/j.ic.2018.09.009>
25. Kaiser, A., Kroening, D., Wahl, T.: Dynamic cutoff detection in parameterized concurrent programs. In: CAV. *LNCS*, vol. 6174, pp. 645–659. Springer (2010). https://doi.org/10.1007/978-3-642-14295-6_55
26. Khalimov, A., Jacobs, S., Bloem, R.: Towards efficient parameterized synthesis. In: VMCAI. *LNCS*, vol. 7737, pp. 108–127. Springer (2013). https://doi.org/10.1007/978-3-642-35873-9_9
27. Kupferman, O., Piterman, N., Vardi, M.Y.: From liveness to promptness. *Formal Methods in System Design* **34**(2), 83–103 (2009)
28. Kurshan, R.P., McMillan, K.L.: A structural induction theorem for processes. *Inf. and Comp.* **117**(1), 1–11 (1995). <https://doi.org/10.1006/inco.1995.1024>
29. Namjoshi, K.S.: Symmetry and completeness in the analysis of parameterized systems. In: VMCAI. *Lecture Notes in Computer Science*, vol. 4349, pp. 299–313. Springer (2007). https://doi.org/10.1007/978-3-540-69738-1_22
30. Pnueli, A., Ruah, S., Zuck, L.D.: Automatic deductive verification with invisible invariants. In: TACAS. *Lecture Notes in Computer Science*, vol. 2031, pp. 82–97. Springer (2001). https://doi.org/10.1007/3-540-45319-9_7
31. Spalazzi, L., Spegni, F.: Parameterized model-checking of timed systems with conjunctive guards. In: VSTTE. *Lecture Notes in Computer Science*, vol. 8471, pp. 235–251. Springer (2014). https://doi.org/10.1007/978-3-319-12154-3_15
32. Spalazzi, L., Spegni, F.: On the existence of cutoffs for model checking disjunctive timed networks. In: ICTCS/CILC. *CEUR Workshop Proceedings*, vol. 1949, pp. 174–185. CEUR-WS.org (2017)
33. Suzuki, I.: Proving properties of a ring of finite state machines. *Inf. Process. Lett.* **28**(4), 213–214 (1988). [https://doi.org/10.1016/0020-0190\(88\)90211-6](https://doi.org/10.1016/0020-0190(88)90211-6)
34. Vardi, M.Y., Wolper, P.: An automata-theoretic approach to automatic program verification. In: LICS. pp. 322–331. IEEE Computer Society (1986)