Reasoning about Rings*

E. Allen Emerson[1] and Kedar S. Namjoshi[2]

[1] Department of Computer Sciences, The University of Texas at Austin
emerson@cs.utexas.edu
[2] Bell Laboratories, Lucent Technologies
kedar@research.bell-labs.com

Abstract. Distributed processes are often composed of similar processes connected in a unidirectional ring network. Processes communicate by passing a token in a fixed direction; the process that holds the token is allowed to perform certain actions. Usually, correctness properties are expected to hold irrespective of the size of the ring. We show that the question of checking many useful correctness properties for rings of *all* sizes can be reduced to checking them on ring of sizes up to a small *cutoff* size. We apply our results to the verification of a mutual exclusion protocol and Milner's scheduler protocol.

1 Introduction

The ring is a particularly useful way of structuring a system of concurrently executing processes. Well known examples include protocols for mutual exclusion, leader election, scheduling, and the dining philosophers problem. These have two features in common: the individual processes of the ring are isomorphic (i.e., the "code" of one can be transformed into that of another by a simple renaming), and the desired correctness properties are expected to be satisfied by instances of arbitrary size. The protocols are thus *parameterized* by the number of processes. The usual method of verifying that such a parameterized system satisfies a specification is by a deductive, manual proof [MC 88,MP 94], which requires considerable ingenuity, and can in practice be done only for reasonably simple protocols.

In this paper, we show that, for *any* system of many isomorphic processes organized in a ring which communicate through a token used as a signal, a property holds for *every* instance of the system iff it holds for instances up to a small *cutoff* size. Thus, for systems composed of isomorphic finite state processes, a fully automated technique, such as Model Checking [CE 81,CES 86] may be applied for the verification of the infinite-state parameterized system.

Correctness properties are expressed in a stuttering-insensitive sub-logic, $\mathsf{CTL}^*\backslash\mathsf{X}$ [BCG 89], of the branching time logic CTL^* [EH 86]. $\mathsf{CTL}^*\backslash\mathsf{X}$ allows all formulas of CTL^* that do not contain the next-state operator X. Formulas in $\mathsf{CTL}^*\backslash\mathsf{X}$ are insensitive to "stuttering", i.e., repeated occurrences of the same state label or non-observable actions. As the correctness properties are to hold of rings of various sizes, it seems reasonable to make them stuttering-insensitive, since the exact timing requirements will, in general, vary amongst rings of various sizes.

Many correctness properties of parameterized systems can be expressed in an indexed version of $\mathsf{CTL}^*\backslash\mathsf{X}$. Let $g(\boldsymbol{i})$, for a vector of variables $\boldsymbol{i}$, be a formula in which every proposition and action is indexed by one of the variables $i_k, k < |\boldsymbol{i}|$. Thus, for example, the instances $g(1)$ and $g(2)$ of the formula $g(i)$ are isomorphic up to re-indexing. We consider correctness properties of the form $(\forall \boldsymbol{i} :: g(\boldsymbol{i}))$. For instance, $(\forall i :: g(i))$ (every process satisfies $g(i)$), and $(\forall i, j : i{\neq}j : g(i, j))$ (every distinct pair of processes satisfies $g(i, j)$). The

* This work was supported in part by NSF grants CCR-9415496 and CCR 980-4736 and SRC contract 99-TJ-685. A preliminary version of this paper was presented at the ACM Symposium on Principles of Programming Languages (POPL), 1995. The work of the second author was done while at the University of Texas at Austin.

quantified variables range over $[0 \ldots n-1]$, where n is the size of the ring instance over which the formula is evaluated. This range is left implicit for conciseness. For instance, mutual exclusion can be expressed as $(\forall i, j : i \neq j : \mathsf{AG}\neg(\mathit{critical}_i \wedge \mathit{critical}_j))$, and eventual access by $(\forall i :: \mathsf{AG}(\mathit{trying}_i \Rightarrow \mathit{critical}_i))$. Our specific results are as follows:

- $(\forall i :: g(i))$ has a cutoff of 2.
- $(\forall i :: g(i, i+1))$ has a cutoff of 3.
- $(\forall i, j : i \neq j : g(i, j))$ has a cutoff of 4.
- $(\forall i, j : i \neq j : g(i, i+1, j))$ has a cutoff of 5.

Similar results for other quantifier combinations may be derived by the application of our proof techniques. The rotational symmetry of the ring plays an important role in the proofs of these results. It allows us to reduce the check of a formula such as $(\forall i :: g(i))$, which ranges over all possible process indices in an instance of the system, to that of one with a particular index, say $g(0)$. We then establish the existence of a correspondence, with respect to the behavior of process 0, between the state transition graphs of an instance of the system with n processes, and one with the number of processes equal to the cutoff. The symmetry then allows us to establish the result. The key to the correspondence proof is the observation that a segment of the ring not containing any observable processes (e.g(, for $(\forall i :: g(i))$, one not containing process 0) has behaviour similar to that of a single process.

The paper is organized as follows. Section 2 defines the notation and constructions we need. Section 3 contains the definition of the system model. We prove the main results in Section 4. Section 5 contains applications of these results to two protocols. In Section 6 we show a sharp boundary for the decidability of parameterized verification: the question is undecidable for ring networks where the token may carry a single bit of mutable information. Section 7 concludes the paper with a discussion of related work and future extensions. The Appendix contains detailed proofs of the main results.

2 Preliminaries

This section contains definitions of the syntax of the temporal logics considered in the paper, the types of structures over which the logics are interpreted, and the semantics of these logics.

2.1 Notation

The notation follows the format introduced by Dijkstra and Scholten in [DS 90]. Quantified expressions are written in the format $(\mathbf{Q}x : r : p)$, where $\mathbf{Q}$ is the quantifier (e.g(, $\exists, \forall, \mathit{min}, \mathit{max}$), x the bound variable, r the range, and p the expression being quantified. When the range of x is clear from the context, it may be dropped, in which case the quantified expression has the form $(\mathbf{Q}x :: p)$. Sets defined by set comprehension are written in the format $\{x|P(x)\}$, which represents the set of all elements x for which the predicate P holds. The powerset of a set S is denoted by 2^S. For a binary relation R, we often write $s\ R\ t$ instead of $(s,t) \in R$ for readability. Disjoint set-union is represented by $\uplus$. For $W = S \uplus T$, $W|_S$ $(W|_T)$ is the projection of W on to S (T).

Proofs are often given in the calculational style [DS 90]. A proof in this format is constructed out of the following basic unit:

$$\begin{array}{l} \quad P \\ op\ (\ \text{hint why}\ P\ op\ Q\ ) \\ \quad Q \end{array}$$

The binary operator op is reflexive and transitive, which allows the chaining together (using transitivity) of a number of similar proof units. Typical operators are $\Rightarrow$ (implies), $\Leftarrow$ (follows-from) and $\Leftrightarrow$ (if-and-only-if).

2.2 Labeled Transition Systems

The types of structures over which temporal logics are interpreted are called Labeled Transition Systems (LTS's) [Kel 76]. Informally, a LTS is a directed graph, where each vertex (a state) is labeled with the atomic propositions true at that state, and each edge (an action) is labeled with an action symbol.

Definition 1 (LTS) *A Labeled Transition System (or LTS) A is a tuple $(Q, \Sigma, \delta, \lambda, L, I)$ where*

- *Q is a nonempty set of states,*
- *Σ is a transition (or action) alphabet, possibly containing the distinguished symbol τ (the silent action),*
- *$\delta \subseteq Q \times \Sigma \times Q$ is the transition relation,*
- *$\lambda : Q \to L$ is a labelling function on the states,*
- *L is a nonempty set of labels,*
- *$I \subseteq Q$ is the set of initial states.*

For clarity, we write $s \stackrel{a}{\to}_A t$ instead of $(s, a, t) \in \delta$. We write $s \to_A t$ for $(\exists a : a \in \Sigma : s \stackrel{a}{\to}_A t)$. The subscript A is often elided when a specific LTS is under consideration. A *path* p of A is a pair (σ, α), where σ and α are sequences (finite or infinite) of states and actions respectively such that $1 + |\alpha| = |\sigma|$, where for every i such that $i + 1 < |\sigma|$, $\sigma_i \stackrel{\alpha_i}{\to}_A \sigma_{i+1}$. A *fullpath* is a maximal path; i.e., either it is infinite or the last state has no successor. A *computation* is a path where the state sequence starts at an initial state of A. Let the length of a path $p = (\sigma, \alpha)$, $|p|$, equal $|\sigma|$.

The semantics of most temporal logics do not take edge labels into account. They are defined over a special kind of LTS known as a Kripke structure.

Definition 2 (Kripke Structure) *A Kripke Structure is an LTS of the form $(Q, \{\tau\}, \delta, \lambda, 2^{AP}, I)$, where AP is a set of* atomic propositions.

2.3 LTS Composition

The parallel composition of two LTS's $A = (Q_A, \Sigma_A, \delta_A, \lambda_A, L_A, I_A)$ and $B = (Q_B, \Sigma_B, \delta_B, \lambda_B, L_B, I_B)$ is defined with respect to a partial "synchronizing" function $\Gamma : \Sigma_A \to \Sigma_B$, which is a bijection on its domain. The composition, denoted by $A \parallel_\Gamma B$, defines the LTS $AB = (Q_{AB}, \Sigma_{AB}, \delta_{AB}, \lambda_{AB}, L_{AB}, I_{AB})$:

1. $Q_{AB} = Q_A \times Q_B$,
2. $\Sigma_{AB} = (\Sigma_A \backslash dom(\Gamma)) \uplus (\Sigma_B \backslash rng(\Gamma))$,
3. δ_{AB} is defined by: $(s, t) \stackrel{a}{\to}_{AB} (u, v)$ iff
   (a) $a \notin dom(\Gamma) \wedge s \stackrel{a}{\to}_A u \wedge t = v$ (a move by A only), or
   (b) $a \notin rng(\Gamma) \wedge t \stackrel{a}{\to}_B v \wedge s = u$ (a move by B only), or
   (c) $a = \tau \wedge (\exists b : b \in dom(\Gamma) : s \stackrel{b}{\to}_A u \wedge t \stackrel{\Gamma(b)}{\to}_B v)$ (a synchronized move)
4. $\lambda_{AB}(s, t) = \lambda_A(s) \uplus \lambda_B(t)$,
5. $L_{AB} = L_A \uplus L_B$, and
6. $I_{AB} = I_A \times I_B$.

It is easily shown that $\parallel_\Gamma$ is associative. It is also symmetric up to strong bisimulation.

2.4 LTS Projection

For an LTS $C = A \|_\Gamma B$, the projection of C on A elides the labels on actions performed by the B-component, replacing them with the silent action τ. Formally, the projection is defined as the LTS $C|_A = (Q, \Sigma, \delta, \lambda, L, I)$, where

1. $Q = Q_C$,
2. $\Sigma = \Sigma_A$,
3. δ is defined by: $s \xrightarrow{a} t$ iff
   (a) $a \in \Sigma_A \ \wedge \ s \xrightarrow{a}_C t$, or
   (b) $a = \tau \ \wedge \ (s \xrightarrow{\tau}_C t \ \vee \ (\exists b : b \in (\Sigma_B \backslash \Sigma_A) : s \xrightarrow{b}_C t))$
4. $\lambda(s) = (\lambda_C(s))|_{L_A}$,
5. $L = L_A$, and
6. $I = I_C$

2.5 Temporal Logics and their Semantics

We will define the syntax and semantics of the *Full Branching Time Logic* CTL^*. Other logics such as $\mathsf{CTL}^* \backslash \mathsf{X}$, CTL and LTL may be defined as sub-logics of CTL^*.

CTL^* (read as "Computation Tree Logic - star") [EH 86] derives its expressive power from the freedom of combining modalities which quantify over paths and the modalities which quantify states along a particular path. These modalities are $\mathsf{A}, \mathsf{E}, \mathsf{F}, \mathsf{G}, \mathsf{X}_s$, and U_w ("for all futures", "for some future", "sometime", "always", "strong next-time", and "weak until", respectively), and they are allowed to appear in virtually arbitrary combinations. Usually, CTL^* is interpreted over LTS's where there are no visible actions. For our purposes, however, it is more convenient to consider LTS's with visible actions; hence, we add the operators X_s^a, for each visible action a, obtaining a logic that we call CTL^*_Σ. Formally, we inductively define a class of state formulas (true or false of states) and a class of path formulas (true or false of paths), which is the least set of formulas satisfying:

(S1) Any atomic proposition P is a state formula.
(S2) If p, q are state formulas, then so are $p \ \wedge \ q$, $\neg p$.
(S3) If p is a path formula then $\mathsf{E}p$ is a state formula.
(P1) Any state formula p is also a path formula.
(P2) If p, q are path formulas, then so are $p \ \wedge \ q$, $\neg p$.
(P3) If p, q are path formulas then so are $\mathsf{X}_s p$, $\mathsf{X}_s^a p$ and $p \mathsf{U}_w q$.

The semantics of a formula is defined with respect to a LTS $M = (S, \Sigma, \delta, \lambda, 2^{AP}, I)$. We write $M, s \models p$ ($M, x \models p$) to mean that state formula p (path formula p) is true in LTS M at state s (of fullpath x, resp.). When M is understood, we write simply $s \models p$ ($x \models p$). We define $\models$ inductively using the convention that x denotes a fullpath and x^i denotes the suffix fullpath starting at the ith state in x, provided $i < |x|$. For a state s:

(S1) $s \models P$ iff $P \ \in \ \lambda(s)$ for atomic proposition P
(S2) $s \models p \ \wedge \ q$ iff $s \models p$ and $s \models q$,
     $s \models \neg p$ iff not $(s \models p)$
(S3) $s \models \mathsf{E}p$ iff for some fullpath x starting at s, $x \models p$

For a fullpath x:

(P1) $x \models p$ iff $s \models p$, for any state formula p, where s is the first state on p,

(P2) $x \models p \ \wedge \ q$ iff $x \models p$ and $x \models q$,
$x \models \neg p$ iff not $(x \models p)$
(P3) $x \models \mathsf{X}_s p$ iff x^1 is defined and $x^1 \models p$,
$x \models \mathsf{X}_s^a p$ iff there exists $i < |x|$ such that for every $j < i$, the jth action on x is a τ-action, and the ith action is a, and x^{i+1} is defined and $x^{i+1} \models p$,
$x \models (p \, \mathsf{U}_w \, q)$ iff for all $i < |x|$, either $x^j \models \neg q$, for some $j \leq i$, or $x^i \models p$

We say that state formula p is *valid*, and write $\models p$, if, for every LTS M and every state s in M, $M, s \models p$. We say that state formula p is satisfiable iff for some LTS M and some state s in M, $M, s \models p$. In this case we also say that M defines a *model* of p. We define validity and satisfiability for path formulas similarly.

Other connectives can then be defined as abbreviations in the usual way: $p \vee q \equiv \neg(\neg p \wedge \neg q)$, $p \Rightarrow q \equiv \neg p \vee q$, $p \Leftrightarrow q \equiv (p \Rightarrow q) \wedge (q \Rightarrow p)$, $\mathsf{A}p \equiv \neg \mathsf{E} \neg p$, $\mathsf{G}p \equiv p \, \mathsf{U}_w \, false$, and $\mathsf{F}p \equiv \neg \mathsf{G} \neg p$. Further operators may also be defined as follows:
$\mathsf{X}_w p \equiv \neg \mathsf{X}_s \neg p$ is the weak next-time,
$p \mathsf{U}_s q \equiv (p \mathsf{U}_w q) \wedge \mathsf{F} q$ is the strong until,
$\overset{\infty}{\mathsf{F}} p \equiv \mathsf{GFX}_s p$ means infinitely often p,
$\overset{\infty}{\mathsf{G}} p \equiv \mathsf{FGX}_s p$ means almost everywhere p.

Sublogics of CTL^*_Σ are defined either by restrictions on the operators allowed, or by restrictions on the ways in which the operators can be combined.

LTL (Linear Temporal Logic) is obtained by leaving out (S2) and (S3), and redefining $M, s \models f$ to be true iff for every fullpath x that starts at s, $M, x \models f$.

$\mathsf{CTL}^*_\Sigma \backslash \mathsf{X}$ is obtained by leaving out the operator X_s, while retaining the X_s^a operators. This makes the logic insensitive to *stuttering* (repetition of the same state or of unobservable actions).

The paper by Emerson [Em 90] contains a detailed survey of these logics and a comparison of their expressive powers.

2.6 Block Bisimulation

In order to compare different instances of the parameterized ring system, we define a form of bisimulation that ignores stuttering (τ-actions), by dividing a computation into "blocks" w.r.t. an equivalence relation $\approx$. It can be shown that this is equivalent to the notion of stuttering bisimulation introduced by [BCG 88].

Let $A = (Q, \Sigma, \delta, \lambda, L, I)$ be a LTS. For an equivalence relation $\approx$ on Q, let $Fin_\approx(s)$ be the set of finite paths $(\sigma; v, \alpha; a)$, where α is a sequence of τ-actions, σ begins at s, every state in σ is in the same $\approx$-class as s, and either v is in a different class, or a is a non-τ action. Let $Inf_\approx(s)$ be the set of all fullpaths $(\sigma; \alpha)$ where σ begins with s, α is a sequence of τ-actions, and every state on σ is in the same $\approx$-class as s. Essentially, $Fin_\approx(s)$ and $Inf_\approx(s)$ are the maximal "blocks" or paths starting at s for which there is no observable change. For a path $p = (\sigma; v, \alpha; a)$ in $Fin_\approx()$, let $endS(p) = v$ and $endA(p) = a$.

Definition 3 (Block Bisimulation) *Let $h = (h^L, h^\Sigma)$ be a pair of bijections on L and Σ respectively. A relation R on Q is a block bisimulation on the LTS A w.r.t. h iff*

1. *R is an equivalence relation,*
2. *If $s \ R \ t$ then*
   (a) $h(\lambda(s)) = \lambda(t)$,
   (b) $(\forall p : p \in Fin_R(s) : (\exists q : q \in Fin_R(t) : h^\Sigma(endA(p)) = endA(q) \ \wedge \ endS(p) \ R \ endS(q)))$
   (c) $(\forall p : p \in Inf_R(s) : (\exists q : q \in Inf_R(t) : |p| = \omega \ \equiv \ |q| = \omega))$

□

The term "block" bisimulation is used as the definition essentially looks for matching (w.r.t. R) blocks (i.e., paths) from states related by R. In order to compare two LTS's A and B, one can form the disjoint union C of A and B and define a block bisimulation on C.

For Kripke Structures, where $L = 2^{AP}$, for any formula f, let $h(f)$ be the formula obtained by replacing each occurrence of each proposition P of f by $h^L(p)$ and of each action a by $h^\Sigma(a)$. The following theorem follows from the equivalence of block bisimulation and stuttering bisimulation and the corresponding result in [BCG 88].

Theorem 1 *Let R be a block bisimulation on LTS A w.r.t. h. For s R t and any formula f of $\mathsf{CTL}^*_\Sigma\backslash\mathsf{X}$, $A, s \models f$ iff $A, t \models h(f)$.* □

We write $s \sim t$ (read as "s and t are block equivalent") if there is a block bisimulation on A that includes the pair (s, t). We write $A \sim B$, for LTS's A and B iff there is a block bisimulation on the disjoint union of A and B such that every initial state in A is related to some initial state of B and vice-versa.

2.7 Group Theoretic notions

Some simple notation from Group Theory is needed to define the symmetries of LTS's. For an natural number n, let $[n]$ represent the set $\{x \mid 0 \leq x \leq n-1\}$. Let $Sym\ I$ be the full symmetry group on the index set $I \subseteq \mathbf{N}$, and $\mathcal{C}_n$ be the permutation group of rotations on a ring of size n. For groups $\mathcal{C}$ and $\mathcal{D}$, we write $\mathcal{C} \preceq \mathcal{D}$ to mean that $\mathcal{C}$ is a subgroup of $\mathcal{D}$. Clearly, $\mathcal{C}_n \preceq Sym\ [n]$.

For a formula f indexed over a set I, $Aut\ f = \{\pi \in Sym\ I |\ \pi(f) \equiv f\}$. For a LTS M composed of n processes in parallel, the states are indexed by $[n]$, and $Aut\ M$ is the group of permutations in $Sym\ [n]$ that when applied to every state and transition of M, map M to an isomorphic copy [ES 93]. Abusing notation slightly, we define $Aut\ s$, for a state s, to be the set of permutations that, when applied to s, map it to itself.

3 System Model

Informally, the token passing model is defined by providing a process "template" T which has the following properties:

- The template process has two types of transitions: those that are enabled only if the process has the token, and others which can be enabled without the process possessing the token. We let the system evolve according to pure nondeterminism, i.e., no fairness is assumed. The actions are divided into token transfer actions snd, rcv, and the other actions Σ.
- Send and Recieve actions must alternate in every computation of the template process, and the first non-Σ action must be a recieve.

In the parameterized system, a non-deterministically chosen process will be initially assigned the unique token through its initial receive action. These requirements are easily checked. For a finite-state system, the last requirement may be expressed in $\mathsf{CTL}^*_\Sigma\backslash\mathsf{X}$ as $\mathsf{AG}(recieve \Rightarrow \mathsf{AF}\, send)$ and model checked for a finite state process, taking any local fairness constraints into account.

Formally, individual processes of a ring are constructed from a template process T, which is an LTS defined by:

1. The set of states, $Q \times \mathbf{B}$. The boolean component indicates possession of the token.
2. The set of actions, Σ, which is partitioned into Σ_f, the set of "free" actions, Σ_d, the set of "token dependent" actions, and $\{snd, rcv\}$, the set of token transfer actions.

3. The transition relation δ, where
   - For every $(q,b) \xrightarrow{a} (q',b') \in \delta$,
     (a) $a \in \Sigma_f \Rightarrow b \equiv b'$ (A free transition cannot change possession.)
     (b) $a \in \Sigma_d \Rightarrow b \ \wedge \ b'$ (A token dependent transition can execute only if the process possesses the token.)
     (c) $a = rcv \Rightarrow \neg b \ \wedge \ b'$ (A receive establishes possession of the token.)
     (d) $a = snd \Rightarrow b \ \wedge \ \neg b'$ (A send revokes possession of the token.)
   - For every (q,b) such that $(q^0, false) \xrightarrow{a} (q,b) \in \delta$, $a = rcv$. (The only possible initial action is a receive.)
   - rcv and snd actions alternate along every path in T.
   - δ is *total* in the first component. (The process is nonterminating)
4. The set of propositions is $Q \times \mathbf{B}$, and the labelling function is the identity function.
5. The initial state is $(q^0, false)$.

For a finite-state template process, it is possible to check the conditions on δ automatically, by model-checking a certain CTL formula.

An individual process K_i $(i \in [n], n \geq 1)$ in an instance of the system with n processes is defined by $K_i = \gamma_i(T)$, where γ_i is the re-labelling defined by the action renaming $\{rcv_i/rcv, snd_{i+1}/snd\} \cup \{a_i/a : \ a \in \Sigma\}$ and the state renaming s_i/s for every state s. The instance is defined by $Ring_n = K_0 \parallel K_1 \parallel \ \ldots \ \parallel K_{n-1} \parallel D_n$, where D_n is the initial token distribution process that synchronizes initially with a non-deterministically chosen K_i to pass it the token. For each composition $K_i \parallel K_{i+1}$, the synchronized actions are snd_{i+1} and rcv_{i+1} (all arithmetic in a ring of size n is done modulo n). Let M_n denote the LTS induced by $Ring_n$. For a set of indices $I \subseteq [n]$, we let $M_n|_I$ denote the projection of M_n on to the processes indexed by I. For a set P of propositions indexed by $[n]$, P_I denotes the subset that is indexed with elements from I.

Definition 4 (Ring Intervals) *Let $[i:j]_n$ denote the indices on the clockwise segment from i to j on a ring of size n. Let $[i:j)_n = [i:j]_n - \{j\}$, $(i:j]_n = [i:j]_n \backslash \{i\}$, and $(i:j)_n = [i:j]_n \backslash \{i,j\}$. Note that $(i:i)_n$, $[i:i)_n$, and $(i:i]_n$ are all empty.* □

4 Property specific abstractions

We show here that for specific types of properties, it suffices to consider instances of size at most a small *cutoff* size in order to show that the property holds for *all* instances.

The properties that we consider are of the forms $(\forall i :: g(i))$, $(\forall i,j : i \neq j : g(i,j))$, and $(\forall i,j : i \neq j : g(i,i+1,j))$. In each case, we first show by symmetry arguments, that in each instance of size n, it is sufficient to instantiate the quantifier $(\forall i)$ with some process index, say 0. We then prove that an instance of size greater than the cutoff size satisfies this smaller formula iff the instance of cutoff size does so. This is shown by exhibiting a block bisimulation between the large and the small system, using the following key theorem.

Theorem 2 (Reduction Theorem) *Let $I \subseteq [n]$, and $J \subseteq [k]$ be sets of indices such that there is a bijection $h : I \to J$ such that for any $i,j \in I$,*

1. *$i \leq j$ iff $h(i) \leq h(j)$, and*
2. *$(i:j)_n \neq \emptyset$ iff $(h(i):h(j))_k \neq \emptyset$.*

Then $M_n|_I \sim M_k|_J$.

Proof Sketch. From h, one can derive the renaming function h^A for indexed actions defined by $h^A(a_i) = a_{h(i)}$, and the renaming function h^P for indexed propositions, defined by $h^P(p_i) = P_{h(i)}$. We define a relation that is a block bisimulation on M_n and M_k w.r.t. (h^P, h^A).

For a state s of M_i, let $tok_i.s$ be the index of the process that possesses the token in state s. Let R_{nk} be the relation between states of M_n and M_k, defined as follows:

$s\ R_{nk}\ t$ iff

1. The local states of D_n in s and D_k in t are identical (recall that D_n is the initial token distribution process),
2. The state of the process in I and J is identical up to h: $h^P(\lambda(s)|_I) = \lambda(t)|_J$,
3. The token is at a process in I in s iff it is at the corresponding process $h(i)$ in t: $(\forall i : i \in I : tok_n(s) = i \equiv tok_k(t) = h(i))$, and
4. The token is in-between processes i and j of I in s iff it is in-between the corresponding processes in J: $(\forall i, j : i, j \in I : tok_n(s) \in (i : j)_n \equiv tok_k(t) \in (h(i) : h(j))_k)$.

Define the relation $B = R_{nn} \cup R_{nk} \cup R_{kk} \cup R_{kn}$. It is straightforward to check that this is an equivalence relation. The Appendix contains the proof that B is a block bisimulation w.r.t. $H = (h^P, h^A)$. □

4.1 Properties of the form $(\forall i :: g(i))$

Properties of the form $(\forall i :: g(i))$ refer to properties that every individual process in a system of processes must satisfy. They are typically used to specify progress requirements, as in absence of starvation $(\forall i :: \mathsf{AG}(trying_i \Rightarrow critical_i))$. We show first that for a symmetric system, it is possible to reduce this property to its instantiation with a single index. The following lemma is a refinement of one in [ES 93].

Lemma 1 *If A is the LTS of a system with n isomorphic processes, $\mathcal{C}_n \preceq Aut\ \ A$, and the start state q^0 is symmetric (i.e $Aut\ \ q^0 = Sym\ [n]$) then, $A, q^0 \models (\forall i :: g(i))$ iff $A, q^0 \models g(0)$.*

Proof.

$(LHS \Rightarrow RHS)$ This is immediate from the definition of $\models$.

$(LHS \Leftarrow RHS)$

$$
\begin{array}{ll}
 & A, q^0 \models g(0) \\
\Rightarrow & (\ \pi \text{ is a permutation }\ ) \\
 & (\forall \pi :\ \pi \in Aut\ \ A :\ \pi(A), \pi(q^0) \models \pi(g(0))) \\
\Leftrightarrow & (\ \pi(A) = A \text{ as } \pi \in Aut\ \ A;\ \pi(q^0) = q^0, \text{ as } q^0 \text{ is symmetric }\ ) \\
 & (\forall \pi :\ \pi \in Aut\ \ A :\ A, q^0 \models g(\pi(0)) \\
\Rightarrow & (\ \text{as } \mathcal{C}_n \leq Aut\ \ A\ ) \\
 & (\forall \pi :\ \pi \in \mathcal{C}_n :\ A, q^0 \models g(\pi(0)) \\
\Rightarrow & (\ \text{logic}\ ) \\
 & A, q^0 \models (\forall \pi :\ \pi \in \mathcal{C}_n : g(\pi(0)) \\
\Rightarrow & (\ \text{The cyclic shifts in } \mathcal{C}_n \text{ drive 0 to } i \text{ for every } i \in [n]\ ) \\
 & A, q^0 \models (\forall i :: g(i))
\end{array}
$$

□

Theorem 3 *Let f be a formula of the form $(\forall i :: g(i))$, where $g(i)$ is a $\mathsf{CTL}^*_\Sigma\backslash\mathsf{X}$ formula that refers only to propositions indexed by i. Then, $(\forall n : n \geq 2 :\ M_n, q^0_n \models f \equiv M_2, q^0_2 \models f)$.*

Proof.

For any $n \geq 2$,
$\quad M_n, q_n^0 \models f$
$\Leftrightarrow$ (the initial state is symmetric, Lemma 1)
$\quad M_n, q_n^0 \models g(0)$
$\Leftrightarrow$ (as $g(0)$ refers only to propositions indexed by 0)
$\quad M_n|_{\{0\}}, q_n^0 \models g(0)$
$\Leftrightarrow$ (By Theorem 2, $M_n|_{\{0\}} \sim M_2|_{\{0\}}$ w.r.t. $h : 0 \mapsto 0$)
$\quad M_2|_{\{0\}}, q_2^0 \models g(0)$
$\Leftrightarrow$ (as $g(0)$ refers only to propositions indexed by 0)
$\quad M_2, q_2^0 \models g(0)$
$\Leftrightarrow$ (by Lemma 1)
$\quad M_2, q_2^0 \models f$
□

4.2 Properties of the form $(\forall i :: g(i, i+1))$

Formulas with the form $(\forall i :: g(i, i+1))$ can express properties of neighboring pairs of processes. For instance, the exclusion property for the Dining Philosophers problem may be expressed as $(\forall i :: \mathsf{AG}\neg(eating_i \ \wedge \ eating_{i+1}))$. By a proof analogous to that for Lemma 1, we have the following lemma.

Lemma 2 *If A is the global state transition graph of a system with n isomorphic processes, $\mathcal{C}_n \preceq Aut \ \ A$, and the start state q^0 is symmetric, then $A, q^0 \models (\forall i :: g(i, i+1))$ iff $A, q^0 \models g(0, 1)$.* □

We have to take into account the various situations that process K_0 and K_1 can be in. Intuitively, either K_0 or K_1 can be given the token initially, or some other process could be assigned the token. This suggests that a three process system may be sufficient. And in fact we have, by an argument analogous to that for Theorem 3,

Theorem 4 *Let f be a formula of the form $(\forall i :: g(i, i+1))$, where $g(i, i+1)$ is a $\mathsf{CTL}^*_{\Sigma}\backslash\mathsf{X}$ formula. Then,* $(\forall n : n \geq 3 : M_n, q_n^0 \models f \ \equiv \ M_3, q_3^0 \models f)$ □

4.3 Properties of the form $(\forall i, j : i \neq j : g(i, j))$

We consider next formulas of the form $(\forall i, j : i \neq j : g(i, j))$. These are used to express properties that involve distinct pairs of processes. Mutual exclusion, for instance, can be written as $(\forall i, j : i \neq j : \mathsf{AG}\neg(critical_i \ \wedge \ critical_j))$.

Lemma 3 *If A is the LTS of a system with n isomorphic processes, $\mathcal{C}_n \preceq Aut \ \ A$, and the start state q^0 is symmetric then, $A, q^0 \models (\forall i, j : i \neq j : g(i, j))$ iff $A, q^0 \models (\forall j : j \neq 0 : g(0, j))$.*

Proof.

$(LHS \Rightarrow RHS)$ This is immediate from the definition of $\models$.

$(LHS \Leftarrow RHS)$

$\quad A, q^0 \models (\forall j : j \neq 0 : g(0, j))$
$\Rightarrow$ (π is a permutation)
$\quad (\forall \pi : \ \pi \in Aut \ \ A : \ \pi(A), \pi(q^0) \models \pi(\forall j : j \neq 0 : g(0, j)))$
$\Leftrightarrow$ ($\pi(A) = A$ as $\pi \in Aut \ \ A$, $\pi(q^0) = q^0$, as q^0 is symmetric)
$\quad (\forall \pi : \ \pi \in Aut \ \ A : \ A, q^0 \models (\forall j : j \neq \pi(0) : g(\pi(0), j)))$

$\Rightarrow$ (as $\mathcal{C}_n \leq Aut\ \ A$)
$\quad (\forall \pi :\ \pi \in \mathcal{C}_n :\ A, q^0 \models (\forall j : j \neq \pi(0) : g(\pi(0), j)))$
$\Rightarrow$ (logic)
$\quad A, q^0 \models (\forall \pi :\ \pi \in \mathcal{C}_n :\ (\forall j : j \neq \pi(0) : g(\pi(0), j)))$
$\Rightarrow$ (The cyclic shifts in $\mathcal{C}_n$ drive 0 to i for every $i \in [n]$)
$\quad A, q^0 \models (\forall i, j : i{\neq}j : g(i, j))$
$\square$

From this lemma, we need to consider only the pairs of processes $(0, j)$, such that $j{\neq}0$. From the Reduction Theorem, we obtain the following theorem.

Theorem 5 *For any $n \geq 4$,*

1. $M_n|_{(0,1)} \sim M_4|_{(0,1)}$,
2. $M_n|_{(0,n-1)} \sim M_4|_{(0,3)}$, *and*
3. $M_n|_{(0,j)} \sim M_4|_{(0,2)}$, *for $j \notin \{0, 1, n-1\}$.*

Proof. Let $h : [n] \to [4]$ be defined by $h(0) = 0$, $h(1) = 1$, $h(n-1) = 3$, and $h(j) = 2$, for $j \notin \{0, 1, n-1\}$. Each of the claims above follows from Theorem 2 by observing that h restricted to the pairs mentioned in each claim is a bijection that satisfies both preconditions of the theorem. The function h may be represented pictorially as shown in Figure 1.

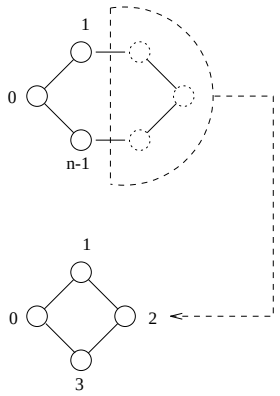

Fig. 1. The abstraction used in the proof

$\square$

Theorem 6 *Let f be a formula of the form $(\forall i, j : i{\neq}j : g(i, j))$, where $g(i, j)$ is a $\mathsf{CTL}^*_\Sigma\backslash\mathsf{X}$ formula. Then, $(\forall n : n \geq 4 : M_n, q^0_n \models f \ \equiv\ M_4, q^0_4 \models f)$.*

Proof.

For any $n \geq 4$,
$\quad M_n, q^0_n \models f$
iff (the initial state is symmetric, Lemma 3)
$\quad M_n, q^0_n \models (\forall j : j{\neq}0 : g(0, j))$

iff (splitting up the formula)
$M_n, q_n^0 \models g(0,1) \ \wedge \ M_n, q_n^0 \models g(0,n-1) \wedge$
$M_n, q_n^0 \models (\forall j : j \notin \{0,1,n-1\} : g(0,j))$
iff (by Theorems 5 and 1)
$M_4, q_4^0 \models g(0,1) \ \wedge \ M_4, q_4^0 \models g(0,3) \wedge$
$M_4, q_4^0 \models g(0,2)$
iff (recombining the formula)
$M_4, q_4^0 \models (\forall j : j \neq 0 : g(0,j))$
iff (by Lemma 3)
$M_4, q_4^0 \models f$
□

As a particular case, the formula $(\forall i,j : i \neq j : \mathsf{AG}\neg(critical_i \wedge critical_j))$, which expresses mutual exclusion, can be checked in a 4-process system.

4.4 Properties of the form $(\forall i,j : i \neq j : g(i,i+1,j))$

Properties of the form $(\forall i,j : i \neq j : g(i,i+1,j))$ are used to express global properties that indicate the relationship between a pair of neighboring processes and an arbitrary process. An example of such a property is presented in Section 5. By proofs similar to those for Theorems 3 and 6, we obtain the following theorems.

Lemma 4 *If A is the LTS of a system with n isomorphic processes, $n > 1$. $\mathcal{C}_n \preceq Aut \ \ A$, and the start state q^0 is symmetric, then, $A, q^0 \models (\forall i,j : i \neq j : g(i,i+1,j))$ iff $A, q^0 \models (\forall j : j \neq 0 : g(0,1,j))$.* □

Theorem 7 *Let f be a formula of the form $(\forall i,j : i \neq j : g(i,i+1,j))$, where $g(i,i+1,j)$ is a $\mathsf{CTL}^*_{\Sigma}\backslash\mathsf{X}$ formula. Then, $(\forall n : n \geq 5, M_n, q_n^0 \models f \ \equiv \ M_5, q_5^0 \models f)$.* □

5 Applications

To illustrate the use of the results, we look at two protocols, one for distributed mutual exclusion [WL 89] and Milner's scheduler protocol [Mil 90].

5.1 Distributed Mutual Exclusion

The template process for the protocol in [WL 89] is given below.

Initially, every process is in state $S0$. Here ϵ is used to indicate a local action. It is easily checked that the process description satisfies the conditions for the token passing model (cf. Section 3). So properties such as

- If a process requests the token, it will eventually receive it. $(\forall i :: \mathsf{AG}(S0_i \Rightarrow \mathsf{AF}(S1_i)))$
- Every trying process eventually enters its critical section. $(\forall i :: \mathsf{AG}(S0_i \Rightarrow \mathsf{AF}(S2_i \ \vee \ S4_i)))$

can be checked using a 2-process system by Theorem 3. Mutual exclusion, i.e $(\forall i,j : i \neq j : \mathsf{AG}\neg(S3_i \wedge S3_j))$ can be checked in a 4-process system by Theorem 6.

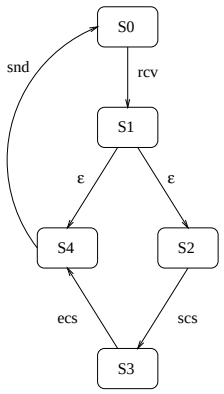

Fig. 2. Template process for the Mutual Exclusion Protocol

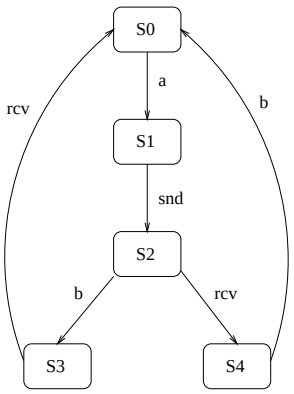

Fig. 3. Template process for Milner's Scheduler Protocol

5.2 Milner's scheduler protocol

The template process for the protocol in [Mil 90] is given below:

The initial state is $S3$. To express the correctness properties concisely, we abbreviate $\mathsf{X}_s^{a_i} true$ by $\hat{a}_i$. The correctness properties are as follows:

- Along every computation, $a_0, a_1 \ldots a_{n-1}$ must be performed cyclically, starting with whichever process is enabled first. The following formula expresses the cyclic order. This has a cutoff of 5, by theorem 7.

  $(\forall i, j : i \neq j : \mathsf{AG}\ (\hat{a_i} \Rightarrow \mathsf{X}_s^{a_i}((\neg \hat{a}_i\ \wedge\ \neg \hat{a}_j)\ \mathsf{U}\ \hat{a}_{i+1})))$

- Every process performs a_i and b_i alternately. The following formula expresses this property and can be checked in a 2-process system by Theorem 3.

  $(\forall i ::\ \mathsf{AG}((\hat{a}_i\ \Rightarrow\ \mathsf{X}_s^{a_i}(\neg \hat{a}_i\ \mathsf{U}\ \hat{b}_i))\ \wedge\ (\hat{b}_i \Rightarrow \mathsf{X}_s^{b_i}(\neg \hat{b}_i\ \mathsf{U}\ \hat{a}_i))))$

6 Undecidability Results

Apt and Kozen [AK 86] show that determining if a temporal property holds for all instances of a parameterized network is undecidable, by a reduction from the non-halting problem for Turing machines. The essential idea is to have a ring of size n simulate the computation of a Turing machine for n steps. Suzuki [Su 88] proved a sharper result by showing that the undecidability holds even when the individual programs are finite-state and the ring is unidirectional.

We give a simpler proof of this result, that also delineates sharply the boundary between decidable and undecidable token-ring systems. The proof is by a reduction from the non-halting problem for 2-counter machines. We show that undecidability arises even if the token takes values from a binary domain. The decidability results in this chapter hold for a token with a single value. Thus the information carrying capacity of the token defines the boundary between decidability and undecidability.

A 2-counter machine [Min 62] has four types of instructions: an increment and decrement for each counter, a zero-test, and a halt instruction. The halting problem for 2-counter machines is known to be undecidable. We reduce this problem to the parameterized model checking problem by using a ring of size n to simulate n steps of a 2-counter machine. Without loss of generality, we may suppose that the 2-counter machine ignores its input tape, and initially has both counter values equal to zero.

Each counter is implemented in unary by a single bit in each node of the ring. The bit takes values $\{up, down\}$ with the number of *up*'s in the global state giving the value of the counter. Each node runs a copy of the same program. Initially a single node is chosen nondeterministically, by giving it the token. That node executes the program of the 2-counter machine. A counter increment is done by circulating a token with an increment instruction. The node with a *down* bit for that counter that first receives this token changes its bit to *up*, then forwards the token with a "neutral" value. The decrement of a counter value occurs in a similar way (by changing the first *up* to *down*). The test for zero is done by circulating the token with value "zero"; if it is received by a process which has the corresponding bit set to *up*, then the token value is changed to "non-zero". If the increment token returns without the neutral value, this implies that the value of that counter is n, so the process simulating the 2-counter machine enters a "dead" state, which has no outgoing transitions.

This simple simulation uses seven values for the token: increment and decrement instructions for each counter, zero/non-zero values, and a neutral value. We can reduce is to a binary token by encoding each value in unary. A processes passes one of these values to its neighbor by sending its unary code to the neighbor as a sequence of tokens with value 1, terminated by a token with value 0. The neighbor accumulates the unary count on receiving 1-tokens, but retransmits them with value 0, which return to the sender, without effecting any change on the local states of the processes in between.

Consider the property $(\forall i :: initially_token_i \Rightarrow \mathsf{AG}\neg halt_i)$, which expresses the condition that the node executing the 2-counter machine program does not reach a halt state. If the 2-counter machine halts, then it must do so in n steps, for some n, so by the simulation, this property is false for a ring of size at least n. If it does not halt, then this property is true for rings of all sizes. Thus, the property is true for all instances iff the machine does not halt. Hence, parameterized verification is undecidable, specifically *co-RE*.

7 Related Work and Conclusions

Among related work, [AK 86,Su 88] show that the problem of automatically checking a specification for every instance of a parameterized system is in general undecidable. Positive results include those of Clarke, Grumberg and Browne [CG 87,BCG 89]; however, their method requires the manual construction of bisimulations or that of a *closure process* which represents computations of an arbitrary number of processes. [KM 89] and [WL 89] introduce the related notion of a *process invariant*. All these methods rely on human ingenuity

to manually construct a suitable process closure or invariant. [GS 92] use automata-theoretic methods to construct process closures for processes connected in a complete network, and use them to establish single index properties. Multi-index properties can be indirectly catered for, but the complexity then becomes multi-exponential. In any case, this does not provide an algorithm for ring networks. Vernier [V 93] has developed a model-checking algorithm for a class of parameterized systems, however, the complexity is high.

The closest results are those of Shtadler and Grumberg [SG 89], who use a network grammar to specify a communication topology, and those of [LSY 94] which deal with ring networks of Petri nets. [SG 89] show that if certain sufficient conditions are satisfied, then every network generated by the grammar satisfies specifications written in linear-time temporal logic. The sufficiency check, however, may require time exponential in the size of an individual process. [LSY 94] contains another (exponential-time) sufficiency test which is used to show that certain parameterized protocols on rings satisfy a single-index linear-time specification. This has recently been extended to show global safety properties in [CGJ 95]. Our results for token rings are more general than those that may be derived from these earlier results, as they allow multi-index, branching-time properties to be checked with polynomial cost.

The advantage of the approach presented here is that, to check whether a ring comprised of isomorphic processes satisfies a specification for all instances, it is both necessary and sufficient to check only the small rings of size less than or equal to the cutoff. This result is independent of the actual program executing on each process in the ring, provided that it follows the token-passing discipline. The ring of size equal to the cutoff is analogous to a closure process, but is trivial to construct.

In addition, the result holds even if the transition graph of the template process is not finite, provided that it is finite-branching. If it is finite, then an automated tool such as SMV [McM92], or the Concurrency Workbench [CPS 89] can be used to model check the desired property for the small ring. This check can be done in time polynomial in the size of a process.

An interesting problem to consider is whether there are conditions under which a similar result holds for systems with multiple-valued tokens. The simplicity of the 2-counter machine simulation suggests that such conditions should be semantic instead of syntactic, perhaps in the form of invariance constraints on the communication patterns.

References

[AK 86] Apt, K., Kozen, D. Limits for automatic verification of finite-state concurrent systems. IPL 15, pp. 307-309.

[BCG 88] Browne, M. C., Clarke, E. M., Grumberg, O. Characterizing Finite Kripke Structures in Propositional Temporal Logic, Theor. Comp. Sci., vol. 59, pp. 115–131, 1988.

[BCG 89] Browne, M. C., Clarke, E. M., Grumberg, O. Reasoning about Networks with Many Identical Finite State Processes, Information and Computation, vol. 81, no. 1, pp. 13–31, April 1989.

[Br 89] Browne, M. Automatic Verification of Finite State Machines using Temporal Logic. PhD thesis, Carnegie-Mellon Univ., CMU-CS-89-117.

[C 93] Cleaveland, R. Analyzing Concurrent Systems using the Concurrency Workbench. in *Functional Programming, Concurrency, Simulation, and Automated Reasoning*, Springer-Verlag LNCS 693.

[CE 81] Clarke, E. M., Emerson, E. A. Design and Synthesis of Synchronization Skeletons using Branching Time Temporal Logic, in *Workshop on Logics of Programs*, Springer-Verlag LNCS 131.

[CES 86] Clarke, E. M., Emerson, E. A., and Sistla, A. P., Automatic Verification of Finite-State Concurrent Systems using Temporal Logic, *ACM Trans. Prog Lang and Sys.*, vol. 8, no. 2, pp. 244-263, April 1986.

[CFJ 93] Clarke, E. M., Filkorn, T., Jha, S. Exploiting Symmetry in Temporal Logic Model Checking, 5th CAV, Springer-Verlag LNCS 697.

[CG 87] Clarke, E. M., Grumberg, O. Avoiding the State Explosion Problem in Temporal Logic Model Checking Algorithms, PODC 1987.

[CGJ 95] Clarke, E.M., Grumberg, O., Jha, S. Verifying Parameterized Networks using Abstraction and Regular Languages. CONCUR 95.

[CPS 89] Cleaveland, R., Parrow, J., Steffen, B. The Concurrency Workbench. In J.Sifakis (ed), *Automatic Verification Methods for Finite State Systems*, Springer-Verlag, LNCS 407.

[D 85] Dijkstra, E. W. Invariance and non-determinacy, in *Mathematical Logic and Programming Languages*, Prentice-Hall International Series in Computer Science. eds. C.A.R Hoare and J.C Shepherdson.

[DS 90] Dijkstra, E. W., Scholten. C. S. *Predicate Calculus and Program Semantics*, Springer-Verlag, 1990.

[dNV 90] De Nicola, R., Vaandrager, F. Three logics for Branching Bisimulation, 5th Annual IEEE Symp. on Logic in Computer Science, pp. 118-129, 1990.

[Em 90] Emerson, E. A. Temporal and Modal Logic, in Handbook of Theoretical Computer Science, (J. van Leeuwen, ed.), Elsevier/North-Holland, 1991.

[EH 86] Emerson, E.A., Halpern, J. Y. "Sometimes" and "Not Never" Revisited: On Branching versus Linear Time Temporal Logic Journal of the ACM, Vol. 33, pp. 151-178, 1986.

[EN 95] Emerson, E.A., Namjoshi, K.S. Reasoning about Rings. *Proc. ACM Symposium on Principles of Programming Languages*, 1995.

[ES 93] Emerson, E. A., Sistla, A. P. Symmetry and Model Checking, 5th CAV, Springer-Verlag LNCS 697.

[ES 95] Emerson, E.A., Sistla, A.P. Utilizing Symmetry when Model Checking under Fairness Assumptions: An Automata-theoretic approach. CAV 1995.

[GS 92] German, S. M., Sistla, A. P. Reasoning about Systems with Many Processes. J.ACM, Vol. 39, Number 3, July 1992.

[GW 89] van Glabbeek, R. J., Weijland, W. P. Branching time and abstraction in bisimulation semantics. in *Information Processing 89*, Elsevier Science Publishers, North-Holland, 1989.

[ID 93] Ip, C., Dill, D. Better verification through symmetry. Proc. 11th Intl. Symp. on Computer Hardware Description Languages and their Applications.

[Kel 76] Keller, R. M. Formal verification of parallel programs. CACM, July 1976.

[KM 89] Kurshan, R. P., McMillan, K. A Structural Induction Theorem for Processes, PODC 1989.

[La 80] Lamport, L. "Sometimes" is Sometimes "Not Never". in *POPL*, 1980.

[LSY 94] Li, J., Suzuki, I., Yamashita, M. A New Structural Induction Theorem for Rings of Temporal Petri Nets. IEEE Trans. Soft. Engg, vol. 20, No. 2, February 1994.

[LP85] Litchtenstein, O., and Pnueli, A., Checking That Finite State Concurrent Programs Satisfy Their Linear Specifications, POPL 85, pp. 97-107.

[Lo 93] Long, D. Model Checking, Abstraction, and Compositional Verification. Ph.D. Thesis, Carnegie-Mellon University, 1993.

[Lu 84] Lubachevsky, B. An Approach to Automating the Verification of Compact Parallel Coordination Programs I. *Acta Informatica* 21, 1984.

[Mil 90] Milner, R. *Communication and Concurrency*, Prentice-Hall International Series in Computer Science. ed. C.A.R Hoare.

[Min 62] Minsky, M. Computation : Finite and Infinite Machines, Prentice-Hall, 1962.

[McM92] McMillan, K., Symbolic Model Checking: An Approach to the State Explosion Problem, Ph.D. Thesis, Carnegie-Mellon University, 1992.

[MC 88] Misra, J., Chandy, K. M. *Parallel Program Design : A Foundation*, Addison-Wesley Publishers, 1988.

[MP 94] Manna, Z., Pnueli, A. Verification of Parameterized Programs. in *Specification and Validation Methods* (E. Borger, ed.), Oxford University Press, pp. 167-230, 1994.

[dNV 90] de Nicola, R., Vaandrager, F. Three logics for branching bisimulation. in *LICS*, 1990. Full version in *Journal of the ACM*, 42(2):458-487, 1995.

[Pn 77] Pnueli, A. The Temporal Logic of Programs. FOCS 1977.

[SG 89] Shtadler, Z., Grumberg, O. Network Grammars, Communication Behaviours and Automatic Verification. In J.Sifakis (ed), *Automatic Verification Methods for Finite State Systems*, Springer-Verlag, LNCS 407.

[Su 88] Suzuki, I. Proving properties of a ring of finite state machines. IPL 28, pp. 213-214.

[Va 87] Vardi, M. Verification of Concurrent Programs - The Automata Theoretic Framework. in *LICS*, 1987. Full version in *Annals of Pure and Applied Logic*, 51:79-98, 1991.

[Va9?] Vardi, M. An Automata-theoretic Approach to Linear Temporal Logic, Proceedings of Banff Higher Order Workshop on Logics for Concurrency, F. Moller, ed., Springer-Verlag LNCS, to appear.

[VW 86] Vardi, M., Wolper, P. An Automata-theoretic Approach to Automatic Program Verification, Proc. IEEE LICS, pp. 332-344, 1986.

[V 93] Vernier, I. Specification and Verification of Parameterized Parallel Programs. Proc. 8th International Symposium on Computer and Information Sciences, Istanbul, Turkey, pp. 622-625.

[WL 89] Wolper, P., Lovinfosse, V. Verifying Properties of Large Sets of Processes with Network Invariants. In J.Sifakis (ed), *Automatic Verification Methods for Finite State Systems*, Springer-Verlag, LNCS 407.

A Appendix

A.1 Proof of Theorem 2

Theorem 2 *Let $I \subseteq [n]$, and $J \subseteq [k]$ be sets of indices such that there is a bijection $h : I \to J$ such that for any $i, j \in I$,*

1. *$i \leq j$ iff $h(i) \leq h(j)$, and*
2. *$(i : j)_n \neq \emptyset$ iff $(h(i) : h(j))_k \neq \emptyset$.*

Then $M_n|_I \sim M_k|_J$.

Proof. From h, one can derive the renaming function h^A for indexed actions defined by $h^A(a_i) = a_{h(i)}$, and the renaming function h^P for indexed propositions, defined by $h^P(p_i) = P_{h(i)}$. We define a relation that is a block bisimulation on M_n and M_k w.r.t. (h^P, h^A).

For a state s of M_i, let $tok_i.s$ be the index of the process that possesses the token in state s. Let R_{nk} be the relation between states of M_n and M_k, defined as follows:

$s\ R_{nk}\ t$ iff

1. The local states of D_n in s and D_k in t are identical (recall that D_n is the initial token distribution process),
2. The state of the process in I and J is identical up to h: $h^P(\lambda(s)|_I) = \lambda(t)|_J$,
3. The token is at a process in I in s iff it is at the corresponding process $h(i)$ in t: $(\forall i : i \in I : tok_n(s) = i \equiv tok_k(t) = h(i))$, and
4. The token is in-between processes i and j of I in s iff it is in-between the corresponding processes in J: $(\forall i, j : i, j \in I : tok_n(s) \in (i : j)_n \equiv tok_k(t) \in (h(i) : h(j))_k)$.

Define the relation $B = R_{nn} \cup R_{nk} \cup R_{kk} \cup R_{kn}$. It is straightforward to check that this is an equivalence relation. Suppose sBt and s is a state in M_n, and t is a state in M_k. By definition of B, $h^P(\lambda(s)|_I) = \lambda(t)|_J$.

1. s is the initial state of M_n.

The only enabled action is the token transfer to a process, say process i. As sBt, t is the initial state of M_k, hence the token transfer action is enabled at t. If $i \in I$, let the transfer from t be to process $h(i)$. Otherwise, there is a pair k, l of indices in I such that $i \in (k : l)$, and $(k : l)$ does not contain an index in I. As sBt, $(h(k) : h(l)) \neq \emptyset$, so let the transfer from t be to some process with index in $(h(k) : h(l))$. It is straightforward to check that the resulting states are related. These single step transitions are the only sequences in $Fin_B(s)$.

2. s is not the initial state of M_n.

(a) Let p be a path in $Fin_B(s)$. Since the only non-τ actions in $M_k|_I$ are those of processes in I, the last action in p is a move by some process i in I. This action can be either a local move, or a token send or receive.

- The action is an internal move by process i. As process $h(i)$ has the same local state, the internal move can be performed at that process, and the resulting states are related by B.
- The action is a token send by process i. Note that the states of process i and $h(i)$ are identical. Hence, the same token send can be performed at $h(i)$, leading to a state related by B to the one resulting from the process i action.

– The action is a receipt of the token by process i from process $i-1$, for $i \in I$ (arithmetic is modulo n). Hence, the token is in the range $(k:i)$ for some (least) k in I. As s and t are related by B, the token in in the range $(h(k):h(i))$ in t. If process $h(i)-1$ is ready to send the token, the receive action may be executed at process $h(i)$, leading to a related state.
Otherwise, the token is at some process j in $(h(k):h(i))$. By the structure of the processes (cf. Section 3), there is a path where the process j eventually reaches a blocking send action and processes without the token reach a blocking receive action. At the end of this path, the actions that transfer the token from process j to process $j+1$ are enabled. Iteratively appending such paths for all processes l from j to $h(i)-1$, one obtains a path where the token moves from process j to a state in process $h(i)-1$ where the token transfer to $h(i)$ is enabled, by a sequence of unobservable actions in $M_k|_J$ and through a sequence of states that are all related by B to the state t. By the earlier argument, the token transfer to $h(i)$ can now be executed, leading to a state related by B to the state resulting from the token transfer to process i in M_n.

(b) The other case, that of infinite sequences from s that are in $Inf_B(s)$, does not arise because every process alternates between send and receive actions. This induces the token to move from one process to the next. Thus, the token eventually either reaches or leaves a process in I, which induces a change to a state not equivalent to s.

The proof above is for the case where s is a state in M_n and t a state in M_k. A similar proof shows the other case. It follows that B is a block bisimulation w.r.t. h^P between $M_n|_I$ and $M_k|_J$. Notice that this proof shows also that the non-τ actions are matched by B up to h^A. □

A.2 Block Bisimulation

Let A be an LTS, $A = (Q, \Sigma, \delta, \lambda, L, I)$. [BCG 88] define stuttering bisimulation over finite state Kripke Structures with a total transition relation. Their definition can be generalized to an LTS as follows:

Definition 5 (Stuttering Bisimulation) *A relation B on A is a stuttering bisimulation iff B is symmetric, and for any s, t such that sBt,*

1. *$\lambda(s) = \lambda(t)$, and*
2. *For every fullpath p starting at s, there is a fullpath q starting at t such that p matches q by B.*

where a *partition* of a fullpath $p = (\sigma, \alpha)$ is a division of p into *segments*, where each segment is a sub-path of the form $(\sigma_i \ldots \sigma_{i+k}, \alpha_i \ldots \alpha_{i+k})$, for some non-negative i, k, and fullpaths p and q *match* w.r.t. B iff they can be partitioned into an equal number of segments, such that for each segment k, every state in the kth segment of p is related by B to every state in the kth segment of q, and the sequence of observable actions is the same in both segments.

It can be shown, using the Knaster-Tarski theorem with the above definition (as in [Mil 90]), that the greatest stuttering bisimulation exists and is an equivalence relation. Two states are stuttering equivalent iff there is a stuttering bisimulation that includes the pair.

Theorem 3 *Let B be a block bisimulation w.r.t. id. Then B is also a stuttering bisimulation.*

Proof. B is an equivalence relation by (1) of Defn. 3. Let (s, t) be an arbitrary pair in B. Then, $\lambda(s) = \lambda(t)$ by 2(a). Consider any fullpath p starting at s. If $p \in Inf_B(s)$, then by 2(c) of Defn. 3, there is a fullpath q starting at t such that $q \in Inf_B(t)$, which is infinite if and only if p is infinite. So p and q may be partitioned into the same number of segments, which by 2(a), consist of states related by B. If p is not in $Inf_B(s)$, then there is either an observable transition, or a state on the path that is in the same B-class as

s. Hence, there is be a maximal prefix p' such that $p' \in Fin_B(s)$. By 2(b), there is a finite path q' in $Fin_B(t)$ such that $endA(p') = endA(q')$ and $endS(p')\ B\ endS(q')$. p' and q' form a pair of matching segments, and $(endS(p'), endS(q'))$ form the start of a new pair of corresponding segments. Continuing in this manner, one can inductively define a fullpath q starting at t, and partitions of p and q such that p and q match w.r.t. B.

Hence, B is a stuttering bisimulation. □

Theorem 4 *Let B be the greatest stuttering bisimulation. Then B is a block bisimulation w.r.t. id.*

Proof. The greatest stuttering bisimulation is an equivalence relation. Let (s,t) be an arbitrary pair in B. Then, $\lambda(s) = \lambda(t)$, hence, condition 2(a) is satisfied.

Consider a path p in $Inf_B(s)$. As p is a fullpath, by (2) of Defn. 5, there is a fullpath q starting at t such that p matches q w.r.t. B. This implies that q is in $Inf_B(t)$, so 2(c) holds.

Consider a finite path p in $Fin_B(s)$. Let γ be a path such that $p;\gamma$ is a fullpath starting at s. By condition 2(b) of Defn. 5, there is a matching fullpath q starting at t. Since p ends with either a state in a different B-class than s, or the first observable action from s, by the definition of matching paths, there is an initial segment q of q which is in $Fin_B(t)$, such that $endA(p) = endA(q)$ and $endS(p)\ R\ endS(q)$. q is non-empty as it includes t. Hence, condition 2(b) of the block bisimulation definition is satisfied.

Thus, B is a block bisimulation w.r.t. id. □

A.3 Proof of Theorem 1

Theorem 1 *Let B be a block bisimulation on LTS A w.r.t. h. For $(s,t) \in B$ and any formula f of $\mathsf{CTL}^*_\Sigma\backslash\mathsf{X}$, $A, s \models f$ iff $A, t \models h(f)$.* □

Proof.

The proof is by structural induction on formulas of $\mathsf{CTL}^*_\Sigma\backslash\mathsf{X}$.

Basis: $f \in AP$.

$M, s \models f$
iff (definition)
$f \in \lambda(s)$
iff (h^{AP} is a bijection on AP)
$h(f) \in h(\lambda(s))$
iff (condition 2(a) of Defn. 3)
$h(f) \in \lambda(t)$
iff (definition)
$M, t \models h(f)$

Inductive case: If f has the form $g_0 \wedge g_1$ or $\neg g$, then the proof follows directly from the inductive hypothesis. Consider the case where f has the form $\mathsf{E}g$, for a path formula g.

$M, s \models \mathsf{E}g$
iff (definition)
$(\exists p : fp(p,s) : M, p \models g)$
iff (by following proof)
$(\exists q : fp(q,t) : M, q \models h(g))$
iff (definition)
$M, t \models \mathsf{E}h(g)$
iff (definition)
$M, t \models h(f)$

To justify the second step of the proof above, suppose that $p = (\sigma, \alpha)$ is a fullpath starting at s. As $(s,t) \in B$, by Defn. 3 and a simple inductive argument involving 2(b) and 2(c), there is a fullpath $q = (\delta, \beta)$ starting at t such that p and q match w.r.t. B. By the inductive hypothesis, states related by B agree up to h on the truth value of sub-formulas of the path formula g that are state formulas. The inductive argument where g has the form $g_0 \ \wedge \ g_1$ or $\neg g_0$ is straightforward. For the case where g has the form $\mathsf{X}^a_s g_0$,

$M, p \models \mathsf{X}^a g_0$
iff (definition)
$(\exists i : \alpha[0 \ldots i-1] \in \tau^* \ \wedge \ \alpha_i = a : p^{i+1} \models g_0)$
iff (p and q match w.r.t. B; observation below, inductive hypothesis)
$(\exists k : \beta[0 \ldots k-1] \in \tau^* \ \wedge \ \beta_k = a : p^{k+1} \models h(g_0))$
iff (definition)
$M, q \models X^a_s h(g_0)$
iff (definition)
$M, q \models h(X^a_s g_0)$

Consider the case where g has the form $g_0 \mathsf{U} g_1$.

$M, p \models g_0 \mathsf{U} g_1$
iff (definition)
$(\exists i : M, p^i \models g_1 \ \wedge \ (\forall j : j < i : M, p^j \models g_0))$
iff (p and q match w.r.t. B; observation below, inductive hypothesis)
$(\exists k : M, q^k \models h(g_1) \ \wedge \ (\forall l : l < k : M, p^l \models h(g_0)))$
iff (definition)
$M, q \models h(g_0) \mathsf{U} h(g_1)$
iff (definition)
$M, q \models h(g_0 \mathsf{U} g_1)$

In both proofs, the second step follows from the following observation: for matching paths, every position on one path has a corresponding position (in the corresponding segment) from which the suffix paths also match. $\square$