Utilizing Symmetry when Model Checking under Fairness Assumptions: An Automata-theoretic Approach

E. A. Emerson¹ and A. P. Sistla²

¹ Department of Computer Science, University of Texas at Austin, Austin, TX 78712 USA.

² Department of Electrical Engineering and Computer Science, University of Illinois at Chicago, Chicago, IL 60680 USA.

Abstract: One useful technique for combating the state explosion problem is to exploit symmetry [ID93, CFJ93, ES93] when performing temporal logic model checking [CE81, CES86]. In [CFJ93] [ES93] it is shown how, using some basic notions of group theory, symmetry may be exploited for the full range of correctness properties expressible in the very expressive temporal logic CTL*. Surprisingly, while fairness properties are readily expressible in CTL*, these methods are not powerful enough to admit any amelioration of state explosion, when fairness assumptions are involved. We show that it is nonetheless possible to handle fairness efficiently by trading some group theory for automata theory. Our automata-theoretic approach [VW86] depends on detecting fair paths subtly encoded in a permutation annotated quotient structure, using a threaded structure to “physically” reflect coordinate permutations.

1 Introduction

Recently there has been much interest in using various techniques to combat the state explosion problem in the automatic verification of finite state concurrent systems [CI93, DGG93]. One valuable technique (cf. [ID93, CFJ93, ES93]) is to exploit the symmetry inherent in systems with many similar processes or subcomponents when performing temporal logic model checking [CE81, CES86]. In [ID93] the focus is on reasoning about a simple but basic type of correctness, viz., safety properties expressible in the temporal logic CTL by an assertion of the form $AG \neg error$. The works of Clarke, Filkorn, & Jha [CFJ93] and Emerson & Sistla [ES93] show how, using some basic notions of group theory, symmetry may be exploited for the full range of correctness properties expressible in the very expressive temporal logic CTL*. Surprisingly, while fairness properties are readily expressible in CTL*, the methods of [CFJ93] and [ES93] are not powerful enough to admit any amelioration of state explosion, when fairness assumptions are involved.

Fairness is necessary for faithfully modeling concurrency with the usual interleaving semantics. Here, concurrent execution of multiple sequential processes is treated as the nondeterministic interleaving of the execution of sufficiently small atomic steps of the individual processes. In order to reflect the basic requirement that each process should be running at some positive but indefinite speed, fairness in its simplest form amounts to assuming that each process is scheduled infinitely often. Much work in the literature has been devoted to the topic of reasoning under fairness constraints (cf. [Fr86]), because of its importance in proving liveness and progress properties.

In this paper, we will explain why fairness is unexpectedly problematic. We will then present a solution that permits us to efficiently handle the full range of fairness properties, including strong fairness, weak fairness, and unconditional fairness.

¹ This author's research is supported in part by NSF grant CCR-9415496 and Semiconductor Research Corporation Contract 94-DP-388.

² This author's research is supported in part by NSF grant CCR-9212183.

In [CFJ93, ES93] model checking a CTL* specification f over a large structure M is reduced to model checking f over a (usually) smaller quotient structure $\bar{M}$. The structure M is potentially intractably large, of size exponential in that of the underlying program whose state graph it represents. The quotient structure $\bar{M}$ is often exponentially smaller than M .

$\bar{M}$ is derived from (the program defining) M and f by identifying “ $\mathcal{G}$ -symmetric” states, where $\mathcal{G}$ is a subgroup of permutations on process indices. For example, states (N_1, C_2) and (C_1, N_2) might be identified in a solution to the Critical section problem. The group $\mathcal{G}$ reflects the global symmetry of M but must also respect the internal symmetry of the specification f . The latter requirement of so respecting f is crucial but is also the source of the problem (as explained in detail in Section 2.6).

Interestingly, our solution depends on trading some group theory for automata theory. We give a new automata-theoretic method (cf. [VW86]) that offers the additional advantage of allowing use of a single quotient structure $\bar{M}$ that depends only on M and and that can be used with any specification f . Moreover, the dependence of $\bar{M}$ on f in the “purely group-theoretic” approaches of [CFJ93] and [ES93] only serves to limit the amount of compression obtained. We can now exploit all the symmetry inherent in M , irrespective of f . An earlier automata-theoretic approach to exploit symmetry was described in [ES93]. But that earlier method could *not* handle fairness efficiently.

Let Φ be a fairness constraint. The crux of our new, automata-theoretic method is showing how to efficiently detect the existence of fair paths in the large global state graph M , i.e., testing $M, s_0 \models E\Phi$, using an *annotated* quotient structure $\bar{M}$ defined with respect to a group $\mathcal{G}$ that respects the symmetry of M but does not depend on Φ . In the annotated $\bar{M}$ arcs are labeled by permutations indicating how the meaning of coordinates shift as $\bar{M}$ is uncompressed to obtain M . These permutations are what make it possible for $\bar{M}$ to succinctly encode M and yet provide enough information to model check over $\bar{M}$ (even though $\bar{M}$ does not respect the symmetry of the specification f or Φ). But they also scramble the meaning of the propositional labeling of states in $\bar{M}$ making it difficult to check for the existence of fair paths. Nonetheless, we show how to efficiently search $\bar{M}$ for any possible “subtly fair” strongly connected subgraphs $\bar{C}$ of $\bar{M}$ that can be unwound into strongly connected subgraphs C of M that contain a path satisfying Φ . To facilitate this search process, we effectively resolve $\bar{M}$ into a *threaded* structure M^* which in essence physically reflects the coordinate shifts in $\bar{M}$ caused by the permutations on arcs. When we find the appropriate “plainly fair” strongly connected subgraphs of M^* then we can conclude that $M, s_0 \models E\Phi$. We can now check whether $M, s_0 \models E(\Phi \wedge f)$ where f is a linear time formula using the automata-theoretic approach of [ES93]. This in turn makes it possible to efficiently model check $M, s \models g$ where g is a Fair Indexed CTL* formula.

This paper is organized as follows. Section 2 presents preliminaries. The main results are described in Section 3. Some concluding remarks are given in Section 4. Technical proofs of the main lemmas and theorems are given in the Appendix.

2 Preliminaries

2.1 Systems and Structures. Let $I = [1 : n]$ be a set of n process *indices*. We have a system $\mathcal{P} = //_{i \in I} K_i$ of n processes K_i running in parallel, and using a set V of (local and/or shared) *variables* including location counters, and other data objects. We assume also that there is a set AP of *atomic propositions*, which amount to boolean variables. Letting $V' = V \cup AP$, then a *global state* s of the system is an assignment of values (over appropriate domains) to each variable in V' . We write $s(v')$ for the value of variable v' in state s .

We assume that all the variables $u, v, \dots$ in V are indexed, and the indices of each variable indicate the processes that share the variable. For example, $u_{i,j}$ denotes a variable shared between processes i and j , whereas v_i is a local variable of process i . In practice, the local variables can be used to denote “location counters” of processes. For example, in a solution to the mutual exclusion problem we may have a state of the form (C_1, T_2) indicating that process 1 is in its critical section while process 2 is in its trying region. if

The formal semantics of such a system is given by a structure $M = (S, R, L, S_0)$ where S is the set of global states, R is a (total) binary transition relation specifying single steps of the system, $L(s) = \{P \in AP : s(P) = \text{true}\}$ is the labeling of each state s with true propositions, and S_0 is the set of initial states.

2.2 Automorphisms. Let $Sym\ I$ be the set of all permutations π on I . $Sym\ I$ forms a group with functional composition ($\circ$) being the group operation. Our convention is that $\pi_b \circ \pi_a$ is evaluated right-to-left: first apply π_a , then π_b . Let Id denote the identity permutation and π^{-1} the inverse of π .

For any indexed object b , such as a state, a variable, or a formula, whose definition depends on I , we can define the notion of permutation π acting on b , by simultaneously replacing each occurrence of index $i \in I$ by $\pi(i)$ in b to get the result $\pi(b)$.

For a structure M , the set of (permutations on I defining graph) automorphisms of M , is $Aut\ M = \{\pi \in Sym\ I : \pi(M) = M\}$. Less tersely, this means that $\pi \in Aut\ M$ precisely when (a) the mapping $\pi : S \rightarrow S : s \mapsto \pi(s)$ is one-one and onto; (b) $s \rightarrow t \in R$ iff $\pi(s) \rightarrow \pi(t) \in R$; and (c) for each $s, s \in S_0$ iff $\pi(s) \in S_0$.

For example, in Figure 1, $Aut\ M = \{Id, Flip\}$ where $Flip$ is the transposition of 1 and 2. For state (C_1, T_2) , we have $Flip((C_1, T_2)) = (T_1, C_2)$.

$Aut\ M$ forms a subgroup of $Sym\ I$. Let $\mathcal{G}$ be any subgroup of $Aut\ M$. Two states s and t in S are $\mathcal{G}$ -equivalent, written $s \equiv_{\mathcal{G}} t$ iff there exists a $\pi \in \mathcal{G}$ such that $t = \pi(s)$. Since $\mathcal{G}$ is a group, $\equiv_{\mathcal{G}}$ is an equivalence relation.

2.3 Quotients. Given structure $M = (S, R, L, S_0)$ and subgroup $\mathcal{G}$ of $Aut\ M$, an annotated quotient structure for M modulo $\mathcal{G}$, $\overline{M} = M/\mathcal{G}$, is a structure of the form $(\overline{S}, \overline{R}, \overline{L}, \overline{S}_0)$ where

- $\overline{S}$ – is a set of designated representative states, exactly one representative $\overline{s}$ for each $\mathcal{G}$ -equivalence class $[s]$ of states in S ,
- $\overline{R}$ – is a set of triples $\overline{s} \xrightarrow{\pi} \overline{t}$ denoting edges between representative states annotated with permutations from $\mathcal{G}$. To define $\overline{R}$ formally, with each state $s \in S$, we associate a canonical permutation π_s such that $\pi_s(s) = \overline{s}$. Now, we define $\overline{R} = \{\overline{s} \xrightarrow{\pi_s^{-1}} \overline{t} : \overline{s} \rightarrow \overline{t} \in R\}$. Intuitively, the edge $\overline{s} \rightarrow \overline{t}$ in $\overline{M}$ is represented by the annotated edge $\overline{s} \xrightarrow{\pi_s^{-1}} \overline{t}$,
- $\overline{L}(\overline{s}) = L(\overline{s})$ for each $\overline{s} \in \overline{S}$, and
- $\overline{S}_0$ – the set of representatives of states in S_0 .

The intuition is that $\overline{M}$ encodes M , often succinctly.³ M itself maybe recovered by “unwinding” $\overline{M}$. (See Figure 2, the annotated quotient of Figure 1.) $\overline{M}$ can be constructed incrementally from the program text in essentially the same fashion as that discussed in [ES93, CFJ93, ID93] for the unannotated quotient.⁴

³ $\overline{M}$ can have an exponentially reduced number of states and transitions than M , but it may also have multiple edges with different permutation annotations between the same pair of nodes. This multiplicity of edges between the same pair of nodes can be limited as follows. If the underlying program defining M is the parallel composition of n sequential (or deterministic) processes,

then the maximal outdegree of each node in M is n . As each annotated edge $\overline{s} \xrightarrow{\pi} \overline{t} \in \overline{M}$ is induced by some edge $\overline{s} \rightarrow \overline{t} \in M$, there can be at most n distinct edges in $\overline{M}$ from $\overline{s}$ to $\overline{t}$.

⁴ The incremental algorithm needs to check if the next generated state is equivalent to a previously

2.4 Temporal Logics. PLTL is the standard propositional linear temporal logic built up from atomic propositions, boolean connectives, and the usual linear time operators G (always), F (sometime), X (next time), and U (until). CTL* extends PLTL by also allowing the path quantifiers A (for all fullpaths) and E (for some fullpath) to also be used. The *basic modalities* of CTL* are formulae of the form Ef where f is a pure PLTL formula. All CTL* formulae can be obtained by taking boolean combinations and nestings of the basic modalities (e.g., $AGEFP$ has basic modalities of the form AG and EF). CTL is a restricted version of CTL* with basic modalities of the form: A or E , followed by a single F, G, X , or U .

Indexed CTL* is built up from basic modalities of the form $\forall_i Ef_i$ and $\wedge_i Ef_i$ where f_i is a PLTL formula that uses only singly indexed *local* atomic propositions such as P_i for $i \in I$ and/or *global* propositional formulas q .⁵ The latter are unindexed (global) atomic propositions such as Q , or propositional formulae of the form $\forall_i p_i, \wedge_i p_i, \forall_{i \neq j} p_i, j, \wedge_{i \neq j} p_i, j$. Here, $\forall_i, \wedge_i$ act as existential and universal process quantifiers ranging over all individual process indices, $\forall_{i \neq j}, \wedge_{i \neq j}$ range over all distinct pairs of process indices, p_i is a propositional formula involving only global atomic propositions and local propositions of index i , and $p_{i,j}$ is a propositional formula involving only global atomic propositions and local propositions of index i and/or j . Formulas of Indexed CTL* are inductively built up from the basic modalities using boolean connectives and nesting (an Indexed CTL* formula may be substituted for a global proposition in another Indexed CTL* formula).⁶ The logic can specify that each (or some) individual processes satisfies a local property such absence of individual starvation, and also specify many important global properties, such as mutual exclusion, by using global propositional formulas appropriately (see [CFJ93]).

Fair Indexed CTL* is just like Indexed CTL* but uses the path quantifiers E_Φ and A_Φ where path quantification ranges only over fair paths [EL87]. Thus $\wedge_i E_\Phi f_i$ means that, for each $i \in I$, the CTL* formula $E(\Phi \wedge f_i)$ holds, where Φ is (expanded to be) a fairness constraint as (specified in PLTL) below.

The formal semantics of these logics is defined in the usual way [Em90] and we write, e.g., $M, s \models g$ to mean that in structure M at state s formula g holds true. Observe that for any $\pi \in \text{Aut } M$ and any $s \in S$, $M, s \models q$ iff $M, \pi(s) \models q$ for global q , while $M, s \models P_i$ iff $M, \pi(s) \models P_{\pi(i)}$.

2.5 Fairness. We use special local propositions, en_i and ex_i , for each $i \in I$. Proposition en_i holds in state s iff process i is enabled in s ; proposition ex_i holds in state s iff all transitions leading to s are due to the execution of a single step of process i . Plainly, for any $\pi \in \text{Aut } M$, $M, s \models ex_i$ iff $M, \pi(s) \models ex_{\pi(i)}$, and similarly for en_i .

An infinite path x of M is *strongly fair* if it satisfies $\Phi = \wedge_{i=1 \dots n} (GFen_i \Rightarrow GFex_i)$, meaning that each process that is enabled infinitely often is executed infinitely often. Similarly a path x is *weakly fair* if it satisfies $\Phi = \wedge_{i=1 \dots n} (FGen_i \Rightarrow GFex_i)$ meaning that any process that is continuously enabled is executed infinitely often. A path is *unconditionally fair* if it satisfies $\Phi = \wedge_{i=1 \dots n} GFex_i$, meaning simply that each process is executed infinitely often.

For example, Figure 1 is the global state graph of a two process solution to the mutual exclusion problem. Each process cycles from its Noncritical to its Trying to its Critical region and back. It is free from individual starvation under strong fairness,

generated state under the group of permutations $\mathcal{G}$. The complexity of performing this check depends on the group $\mathcal{G}$. It is, in general, as hard as graph isomorphism [CFJ93], a problem for which there is no known polynomial time algorithm, but which is not typically viewed as intractable in practice. Moreover, for many commonly occurring $\mathcal{G}$ associated with real systems such as full symmetry or rotational symmetry, it can be done very simply and efficiently [ES93, CFJ93, ID93].

⁵ We assume f_i, f_j , etc. are isomorphic up to reindexing.

⁶ This formulation of Indexed CTL* is slightly different from ICTL* in [CGB89].

satisfying, e.g., $\wedge_i A_\Phi(G(T_i \Rightarrow FC_i))$; but not under pure nondeterministic scheduling, i.e., violating $\wedge_i A(G(T_i \Rightarrow FC_i))$. Mutual exclusion is maintained, so it also satisfies $\wedge_i A_\Phi G(\wedge_{i \neq j} \neg(C_i \wedge C_j))$.

2.6 Background. In [ES93] and [CFJ93] it is shown that

$$(*) \quad M, s_0 \models f \text{ iff } \overline{M}, \overline{s}_0 \models f$$

where f is any CTL* formula and $\overline{M}$ is the (unannotated) quotient $M/\mathcal{G}$ – provided $\mathcal{G}$ is a subgroup of $\text{Aut } M \cap \text{Auto } f$, where $\text{Auto } f$ is the group of permutations⁷ that leave semantically invariant all maximal propositional subformulas of f . For example, if $I = [1 : 4]$ and $f = E(GFex_1 \wedge GFex_2)$, then $\text{Auto } f$ consists of those permutations that leave fixed both 1 and 2.

If we consider $f = E\Phi = E(GFex_1 \wedge \dots \wedge GFex_n)$ where $I = [1 : n]$, then $\text{Auto } f$ is the set of permutations leaving invariant each of $ex_1, \dots, ex_n$. The only such permutation is the identity Id . Hence, $\text{Auto } f = \mathcal{G} = \{Id\}$, and $\overline{M} = M$. Thus, no compression is possible under fairness. Since the uncompressed $\overline{M} = M$ is in general of size exponential in n , the “purely group-theoretic” methods of [CFJ93] and [ES93] are not adequate for dealing with fairness.

3 Checking Correctness under Fairness Assumptions

In this section we give efficient algorithms for model checking formulas of Fair Indexed CTL* under strong fairness assumption $\Phi = \wedge_{i=1 \dots n} (GFen_i \Rightarrow GFex_i)$. This also permits us to handle unconditional and weak fairness since they are special cases of strong fairness [EL87].

The crucial step in the model checking problem is to handle the basic modalities [EL87]. We will describe in detail how to handle $\vee_i E_\Phi f_i$ and $\wedge_i E_\Phi f_i$. By the following lemma, it is therefore enough to check satisfaction of such basic modalities at representative states. The proof of the lemma depends on the “top-level” symmetry inherent in Fair Indexed CTL* formulas due to the process quantifiers.

Lemma 3.1. Two ($\mathcal{G}$)-equivalent states in M satisfy the same set of Fair Indexed CTL* formulas.

We initially consider the formula $E_\Phi f_i$, abbreviating the CTL* formula $E(\Phi \wedge f_i)$. The classical automata-theoretic approach [VW86] to model checking can be used. Apply the usual tableau construction to form the (Buchi) ω -string automaton $\mathcal{A}$ equivalent to f_i which accepts an infinite sequence of (labels of) states of M by guessing a run which goes through GREEN (accepting) automaton states infinitely often. Next construct essentially the product graph (automaton) B of the structure M with automaton $\mathcal{A}$,

⁷ For the interested reader, we explain in detail the crucial role played by $\text{Auto } f$ in the proof of (*). (We also emphasize that $\text{Auto } f$ is required, rather than $\text{Aut } f$, despite 3 typo’s near the beginning of [ES93].) Assume that $M, s_0 \models f$, where $f = E(GFex_1 \wedge GFex_2)$. We wish to show that $\overline{M}, \overline{s}_0 \models f$. Then there is a path $x = s_0, s_1, s_2, \dots$ in M such that $M, x \models GFex_1 \wedge GFex_2$. Moreover, there is (it turns out) a corresponding path $\overline{x} = \overline{s}_0, \overline{s}_1, \overline{s}_2, \dots$ in $\overline{M}$ such that, for each i , $s_i \equiv_{\mathcal{G}} \overline{s}_i$. Thus, s_i and $\overline{s}_i$ are not in general identical, but they do match up to a permutation in $\mathcal{G}$. Since $s_i \equiv_{\mathcal{G}} \overline{s}_i$ and in particular since $\mathcal{G}$ is a subgroup of $\text{Auto } f$, it follows from the definitions that $M, s_i \models p$ iff $\overline{M}, \overline{s}_i \models p$ for the maximal propositional subformulas ex_1 and ex_2 . Intuitively, this means that the indexing of the representatives in $\overline{M}$ preserves the truth of ex_1 and ex_2 . Hence, $\overline{x}$ will be “decorated” with ex_1 and ex_2 exactly as is x . We conclude $\overline{M}, \overline{x} \models GFex_1 \wedge GFex_2$ and $\overline{M}, \overline{s}_0 \models f$.

If it happened that $\mathcal{G}$ did not leave invariant some such propositional p , the proof would break down. For example, if representatives had been chosen so that in $\overline{M}$ index 1 were uniformly replaced by 3, say, then $\overline{M}, \overline{s}_0 \not\models E(GFex_1 \wedge GFex_2)$ since 1 does not even appear in the quotient. This problem is, in general, unavoidable.

which, intuitively, guesses a path through M that is also accepted by $\mathcal{A}$. Then test B for nonemptiness by finding reachable and accepting strongly connected components in the transition graph of B . Using complemented pairs acceptance for B [St81], we get the standard result (Theorem 3.2) below, which effectively solves our problem in time $O(|M| \cdot |\mathcal{A}| \cdot n^2)$ (cf. [LP85], [EL87]). But it requires construction of the global state graph M , which may very well be prohibitively large, and is often of size exponential in the size of the original program text.

Definition. Given a graph B' , an *scsg* C' of B' is a nonempty set of nodes of B' such that the subgraph of B' induced by C' is strongly connected and total. An *mscsg* C' of graph B' is a maximal scsg, i.e., an scsg such that no proper superset is an scsg of B' .⁸

Definition. An scsg C of B is *obviously fair* provided that, for each $j \in [1 : n]$ and each node s of C , if s is labeled with en_j then there is some node t in C labeled with ex_j . An scsg C is *green* iff it contains a node with automaton component in the GREEN accepting set for the Buchi automaton $\mathcal{A}$.

Theorem 3.2. $M, \bar{s}_0 \models E(\Phi \wedge f_i)$ iff B contains an obviously fair, green scsg C reachable from its start state.

We now develop a method to work with the annotated quotient structure $\bar{M}$, which is typically much smaller than M . An automata-theoretic approach based on annotated quotients – that cannot handle fairness efficiently – was described in [ES93].⁹ The idea is that, to check (the “unfair”) $M, s_0 \models E f_i$, we form a product graph (automaton) $\bar{B}$ of the form $\bar{M} \times \mathcal{A} \times I$ that succinctly encodes B , just as $\bar{M}$ encodes M . (In particular, each node u of B has a representative node $\bar{u}$ in $\bar{B}$; the technical definition of $\bar{B}$ is given in the Appendix.) Intuitively, $\bar{B}$ guesses a path $\bar{x}$ through $\bar{M}$ and simulates the automaton $\mathcal{A}$ for formula f_i along the path x in M obtained by “unwinding” $\bar{x}$, using the permutations along $\bar{x}$ to keep track of the shifting position of index i . $M, \bar{s}_0 \models E f_i$ iff the language accepted by $\bar{B}$ is nonempty; we can test nonemptiness of $\bar{B}$ in time $O(|\bar{B}|) = O(|\bar{M}| \cdot |\mathcal{A}| \cdot n)$. We thus have an efficient model checking algorithm that avoids construction of the global state graph M .

This approach can be easily extended to check $M, \bar{s}_0 \models E g$, where g is any linear time formula in which k process indices $h_1, \dots, h_k$ are used. Construct the automaton $\hat{B}$ of the form $\bar{M} \times \mathcal{A} \times I^k$, where $\mathcal{A}$ is the ω -automaton for g . Each state of $\hat{B}$ is of the form $(\bar{s}, q, j_1, \dots, j_k)$ where $\bar{s} \in \bar{M}$, $q \in \mathcal{A}$, and $j_1, \dots, j_k \in I$ store a list of “current values” for indices $h_1, \dots, h_k$, respectively. $\hat{B}$ is nonempty iff $M, s_0 \models E g$. We can test nonemptiness of $\hat{B}$ in time $O(|\hat{B}|) = O(|\bar{M}| \cdot |\mathcal{A}| \cdot n^k)$. However, when $g = \Phi \wedge f_i$, since the fairness constraint Φ involves n coordinates, we have that $k = n$ and the state space of $|\hat{B}|$ is of size exponential in n . As M itself is typically of size exponential in n , this approach defeats our purpose. *It is, thus, not possible to efficiently handle fairness by applying the automata-theoretic approach of [ES93] directly.*

We now present a nontrivial generalization of the automata-theoretic approach that does efficiently cater for fairness. We use the key result (Theorem 3.3) below which will ultimately lead to an efficient algorithm for $E(\Phi \wedge f_i)$, based on the notion of a “subtly” fair scsg $\bar{C}$ of $\bar{B}$ that encodes an obviously fair scsg C of B .

Definition. An scsg $\bar{C}$ of $\bar{B}$ is *subtly fair* provided that, for each $j \in [1 : n]$ and $\bar{u} \in \bar{C}$, if $\bar{u}$ is labeled with en_j , then there is a path in $\bar{C}$ starting at $\bar{u}_0 = \bar{u}$, $\bar{x} = \bar{u}_0 \xrightarrow{\pi_1} \bar{u}_1 \dots \xrightarrow{\pi_k} \bar{u}_k$, such that $\bar{u}_k$ is labeled with $ex_{\pi_k^{-1}(j)}$, where $\pi_{\bar{x}} = \pi_1 \circ \dots \circ \pi_k$ (and $\bar{u}_k$ there by represents a reachable state labeled with ex_j). An scsg $\bar{C}$ is *green* iff its automaton component is in the set GREEN.

⁸ Such mscsgs C' are total, maximal, strongly connected components of B' , which will facilitate their calculation.

⁹ This is in addition to the “purely” group-theoretic approach shown not to work in Section 2.

Theorem 3.3. $M, \bar{s}_0 \models E(\Phi \wedge f_i)$ iff

$\bar{B}$ contains a subtly fair, green scsg $\bar{C}$ reachable from start state $(\bar{s}_0, q_0, i)$.

The proof of Theorem 3.3 depends on establishing precise correspondences between B and $\bar{B}$, and between their respective scsgs.

Lemma 3.4. (Correspondence of paths in B and $\bar{B}$)

Let N be a nonempty subgraph of B , and $\bar{N}$ the subgraph of $\bar{B}$ induced by the representatives of N .

(a) If $\bar{x} = \bar{u}_0 \xrightarrow{\pi_1} \bar{u}_1 \xrightarrow{\pi_2} \bar{u}_2 \dots$ is a path in $\bar{N}$, then the path obtained by unwinding $\bar{x}$,

$$unwind(\bar{x}) \stackrel{\text{def}}{=} \bar{u}_0 \rightarrow \pi_1(\bar{u}_1) \rightarrow \pi_1 \circ \pi_2(\bar{u}_2) \dots$$

is a path in B .

(b) If $x = \bar{u}_0 \rightarrow u_1 \rightarrow u_2 \dots$ is a path in N starting at representative node $\bar{u}_0$, then there exists a path in $\bar{N}$

$$\bar{x} \stackrel{\text{def}}{=} \bar{u}_0 \xrightarrow{\pi_1} \bar{u}_1 \xrightarrow{\pi_2} \bar{u}_2 \dots$$

such that $x = unwind(\bar{x})$.

Definition. If $\bar{C}$ is an scsg of $\bar{B}$ and $\bar{c}$ is any node of $\bar{C}$, then the subgraph of B induced by exhaustively unwinding $\bar{C}$ starting at $\bar{c}$ is called $unwind(\bar{C}, \bar{c})$. It is technically defined to be the set of nodes in B appearing along paths of the form $unwind(\bar{x})$ where $\bar{x}$ is a path in $\bar{C}$ starting at $\bar{c}$.

Lemma 3.5. If C is an obviously fair, green scsg of B , then the subgraph of $\bar{B}$ induced by the representatives of nodes in C , $\bar{C}$, is a subtly fair, green scsg of $\bar{B}$.

Lemma 3.6. If $\bar{C}$ is a subtly fair, green scsg of $\bar{B}$ and $\bar{c}$ is any node of $\bar{C}$, then $unwind(\bar{C}, \bar{c})$ is an obviously fair, green scsg of $\bar{B}$.

At this point, we have, by Theorem 3.3, reduced the problem to that of detecting subtly fair scsgs $\bar{C}$ of $\bar{B}$, which may not be straightforward due to their scrambled nature. However, by resolving $\bar{C}$ into a “threaded” subgraph C^* which physically realizes the permutation annotations, the difficulty is overcome easily. See Figures 3 and 4.

Definition. The *threaded graph* associated with (the subgraph of $\bar{B}$ induced by) scsg $\bar{C}$ is the graph $C^* = (V_C^*, R_C^*, L_C^*)$ where

$$V_C^* = \{(\bar{u}, j) : \bar{u} \in \bar{C} \text{ and } j \in [1 : n]\},$$

$$R_C^* = \{(\bar{u}, j) \rightarrow (\bar{v}, k) : \bar{u} \xrightarrow{\pi} \bar{v} \in \bar{C} \text{ and } k = \pi^{-1}(j)\}, \text{ and}$$

$L_C^*((\bar{u}, j))$ contains en (respectively, ex) provided en_j (respectively, ex_j) is in the label of $\bar{u}$.

Note. π is the “decoding” function, while π^{-1} is the “encoding” function. In unwinding or decoding $\bar{M}$ to get M we use π , while encoding $\bar{C}$ into C^* requires π^{-1} .

Definition. Threaded graph C^* is *plainly fair* provided that, for each node $(\bar{s}, j)$ in C^* , if $(\bar{s}, j)$ is labeled with en , then there is in C^* a path to some node $(\bar{t}, k)$ labeled with ex . Threaded graph C^* is *green* iff it contains a node whose automaton component is in the GREEN set of automaton states.

Lemma 3.7. (Correspondence of paths in $\bar{C}$ and C^*)

Let $\bar{C}$ be a scsg of $\bar{B}$ and C^* the threaded graph associated with $\bar{C}$.

(a) If $\bar{x} = \bar{u}_0 \xrightarrow{\pi_1} \bar{u}_1 \xrightarrow{\pi_2} \bar{u}_2 \dots$ is a path in $\bar{C}$, then for each $j \in [1 : n]$, the path

$$(\bar{u}_0, j) \rightarrow (\bar{u}_1, \pi_1^{-1}(j)) \rightarrow (\bar{u}_2, \pi_2^{-1} \circ \pi_1^{-1}(j)) \dots$$

is a corresponding *thread* in C^* .

(b) If $x^* = (\bar{u}_0, j_0) \rightarrow (\bar{u}_1, j_1) \rightarrow (\bar{u}_2, j_2) \dots$ is a thread in C^* , then there is a corresponding path in $\bar{C}$

$$\bar{x} = \bar{u}_0 \xrightarrow{\pi_1} \bar{u}_1 \xrightarrow{\pi_2} \bar{u}_2 \dots \text{ one of whose threads is } x^*.$$

Using this correspondence we can now prove the following crucial lemma relating subtle fairness of $\overline{C}$ to plain fairness of C^* .

Lemma 3.8. Let $\overline{C}$ be a scsg of $\overline{B}$ and C^* the threaded graph associated with $\overline{C}$. Then,

$\overline{C}$ is subtly fair and green iff C^* is plainly fair and green.

Putting it all together, we can get an efficient algorithm that determines (“marks”) all representative states $\overline{s}$ such that $M, \overline{s} \models E(\Phi \wedge f_i)$. We assume that an earlier algorithm has constructed the annotated quotient structure $\overline{M}$.

Algorithm.

1. Construct the automaton $\mathcal{A}$ corresponding to the formula f_i . Construct $\overline{B}$ from $\overline{M}$ and $\mathcal{A}$. Save a copy of $\overline{B}$ as the graph $\overline{H}$.
2. Repeat the following procedure $n + 1$ times:
 - Compute the mscsgs of $\overline{B}$
 - For each mscsg $\overline{C}$ of $\overline{B}$
 - Check if $\overline{C}$ is subtly fair:
 - Resolve $\overline{C}$ into C^*
 - Check if C^* is plainly fair
 - if so, mark all nodes of $\overline{C}$
 - if not, delete from $\overline{B}$ each node $\overline{s}$ of $\overline{C}$ for which some node $(\overline{s}, j)$ in C^* is “bad”, meaning that it is labelled with en , but it cannot reach in C^* a node labeled ex
3. Using $\overline{H}$, propagate the marking to all nodes that can reach an already marked node.

Theorem 3.9. The above algorithm correctly marks exactly those representative states $\overline{s}$ such that $M, \overline{s} \models E(\Phi \wedge f_i)$. It runs in time $O(|\overline{B}| \cdot n^2) = O(|\overline{M}| \cdot |\mathcal{A}| \cdot n^3)$.

Note that the above claim actually applies to any index $i \in [1 : n]$ and that the states of $\overline{B}$ include $(\overline{s}, q_0, i)$ for all $i \in [1 : n]$. Thus, the above algorithm is, in effect, checking *simultaneously* correctness of all n formulas $E(\Phi \wedge f_1), \dots, E(\Phi \wedge f_n)$, yielding trivially an algorithm for the basic modalities:

Corollary 3.10. (Model checking for basic modalities $\wedge_i E_\Phi f_i$ and $\vee_i E_\Phi f_i$)

- (a) $M, \overline{s} \models \wedge_i E_\Phi f_i$ iff all n states $(\overline{s}, q_0, 1), \dots, (\overline{s}, q_0, n)$ are marked.
- (b) $M, \overline{s} \models \vee_i E_\Phi f_i$ iff one or more of the n states $(\overline{s}, q_0, 1), \dots, (\overline{s}, q_0, n)$ are marked.

Theorem 3.11. (Model Checking for Fair Indexed CTL*) The above procedure provides an algorithm which checks $M, \overline{s}_0 \models g$ for any Fair Indexed CTL* formula g under strong fairness, in time $O(|\overline{M}| \cdot n^3 \cdot |g| \cdot a)$ where $|\overline{M}|$ is the size of $\overline{M}$, $n = |I|$ is the number of processes, $|g|$ is the length of g , and a is the maximum size of the automaton $\mathcal{A}$ for any basic modality of g .

It is interesting to compare the complexity bound of $O(|\overline{M}| \cdot n^3 \cdot |g| \cdot a)$ for our algorithm with the complexity bound of the traditional algorithm that works directly on the original structure M , which can be shown to be $O(|M| \cdot n^3 \cdot |g| \cdot a)$ in the terminology of theorem 3.11.¹⁰ Thus, the complexity bounds for our algorithm and the traditional algorithm have the same form except that our bound is over the likely succinct annotated quotient $\overline{M}$, while that of the traditional algorithm is over the entire structure M which may be intractably large. The “Lichtenstein-Pnueli Thesis” [LP85] suggests that it is

¹⁰ This is seen by noting that for any fixed i , the complexity of the traditional algorithm for model checking the formula $E_\Phi f_i$ on M is $O(|M| \cdot |\mathcal{A}| n^2)$ where $\mathcal{A}$ is the automaton referred to in theorem 3.9. The complexity when we use this algorithm for model checking a basic modality $\wedge_i E_\Phi f_i$ is $O(|M| \cdot |\mathcal{A}| n^3)$, since we need to repeat the previous step for each $i = 1, \dots, n$. The claimed bound follows by noting that there are at most $|g|$ basic modalities in g .

the linear complexity in the structure size that is most important for applications, since structures tend to be extremely large while correctness properties tend to be short.

For Fair Indexed CTL the factor a is a small constant and the bound simplifies to $O(|\overline{M}| \cdot n^3 \cdot |g|)$. Moreover, under weak and unconditional fairness, it turns out that there is no need to perform the iteration in step 2 of the core algorithm. Hence, the bound simplifies to $O(|\overline{M}| \cdot n^2 \cdot |g| \cdot a)$ for Fair Indexed CTL* and $O(|\overline{M}| \cdot n^2 \cdot |g|)$ for Fair Indexed CTL.

4 Conclusions

Symmetry reduction has previously been shown to be a powerful tool for reasoning about reactive systems in a number of contexts. For example, there has been a good deal of work in the Petri-net community using symmetry reduction to ameliorate state explosion when reasoning about reachability, boundedness and other Petri-net related properties (cf. [JR91], [Je94]). Symmetry reduction has been successfully applied to facilitate protocol verification in [APS83] and [Ku86] (cf. [Ku94]), and to facilitate hardware verification in [ID93]. These works did not cater for the full range of temporal correctness properties, but symmetry reduction techniques for arbitrary CTL* properties were developed in [CFJ93] and [ES93]. Unfortunately, as explained in sections 2.6 and 3 of this paper, these techniques are, in effect, inapplicable when trying to reason under fairness assumptions.

Our contribution here is to broaden the range of applicability of symmetry reduction techniques by showing how to efficiently handle Fair Indexed CTL* formulas. This permits efficient reasoning about many important liveness properties that depend on fairness assumptions. It is also worth noting that this could not be done using the essentially group-theoretic approaches of [CFJ93, ES93], but depends crucially on the automata-theoretic framework. This seems to testify to the power of automata [VW84, VW86] (cf. [Ku94]).

Finally, it would be interesting to compare the symmetry reduction approach of collapsing M according to symmetry group G to get quotient $\overline{M}$ and the approach of computing $\hat{M}$, the quotient of M modulo its "coarsest bisimulation". Actually, there are several unsettled points regarding this latter notion. First, a precise definition of "the" coarsest bisimulation in this context must be given; it would seem to inherently involve itself in some group-theoretic considerations relating to symmetries, considering that, for example, states (T_1, N_2) and (N_1, T_2) should be bisimilar. Thus, such a bisimulation would itself be a type of symmetry reduction. Another, more important, problem may be computation of the coarsest bisimulation incrementally. Although there are some recent results for incremental computation of the standard coarsest bisimulation [BFHRR92, LY92], they require symbolic representation of equivalence classes and "oracles" for performing certain operations on these equivalence classes efficiently. As a consequence, they are of uncertain general applicability, and matters seem even less certain for the notion of a symmetry based bisimulation. In any event, this is a topic that might merit further attention.

References

- [APS83] Aggarwal S., Kurshan R. P., Sabnani K. K., "A Calculus for Protocol Specification and Validation", in Protocol Specification, Testing and Verification III, H. Ruden, C. West (ed's), North-Holland 1983, 19-34.
- [BFHRR92] Bouajjani, A., Fernandez, J., Halbwachs, N., Raymond, P., and Ratel, C., Minimal State Graph Generation, *Science of Computer Programming*, 1992.

- [CE81] Clarke, E. M., and Emerson, E. A., Design and Verification of Synchronization Skeletons using Branching Time Temporal Logic, Logics of Programs Workshop 1981, Springer LNCS no. 131.
- [CES86] Clarke, E. M., Emerson, E. A., and Sistla, A. P., Automatic Verification of Finite State Concurrent Programs using Temporal Logic: A Practical Approach, *ACM TOPLAS*, April 1986
- [CFJ93] Clarke, E. M., Filkorn, T., Jha, S. Exploiting Symmetry in Temporal Logic Model Checking, 5th International Conference on Computer Aided Verification, Crete, Greece, June 1993.
- [CGB88] Clarke, E. M., Grumberg, O., and Brown, M., Characterizing Kripke Structures in Temporal Logic, *Theor. Comp. Sci.*, 1988
- [CGB89] Clarke, E. M., Grumberg, O., and Brown, M., Reasoning about Many Identical Processes, *Inform. and Comp.*, 1989
- [CI93] Cleaveland, R., Analyzing Concurrent Systems using the Concurrency Workbench, Functional Programming, Concurrency, Simulation, and Automated Reasoning Springer LNCS no. 693, pp. 129-144, 1993.
- [DGG93] Dams, D., Grumberg, O., and Gerth, R., Generation of Reduced Models for checking fragments of CTL, CAV93, Springer LNCS no. 697, 1993.
- [EL87] Emerson, E. A. and Lei, C.-L., Modalities for Model Checking: branching Time Strikes Back, *Science of Computer Programming*, v. 8, pp. 275-306, 1987
- [ES93] Emerson, E. A., and Sistla, A. P., Symmetry and Model Checking, 5th International Conference on Computer Aided Verification, Crete, Greece, June 1993
- [Em90] Emerson, E. A., Temporal and Modal Logic, in *Handbook of Theoretical Computer Science*, (J. van Leeuwen, ed.), Elsevier/North-Holland, 1991.
- [Fr86] Francez, N., Fairness, Springer-Verlag, New York, 1986
- [GS92] German, S. M. and Sistla, A. P. Reasoning about Systems with many Processes, *Journal of the ACM*, July 1992, Vol 39, No 3, pp 675-735.
- [ID93] Ip, C-W. N., Dill, D. L., Better Verification through Symmetry, CHDL, April 1993.
- [Je94] Jensen, K., Colored Petri Nets: Basic Concepts, Analysis Methods, and Practical Use, vol. 2: Analysis Methods, EATCS Monographs, Springer-Verlag, 1994.
- [JR91] Jensen, K., and Rozenberg, G. (eds.), High-level Petri Nets: Theory and Application, Springer-Verlag, 1991.
- [Ku86] Kurshan, R. P., "Testing Containment of omega-regular Languages", Bell Labs Tech. Report 1121-861010-33 (1986); conference version in R. P. Kurshan, "Reducibility in Analysis of Coordination", LNCIS 103 (1987) Springer-Verlag 19-39.
- [Ku94] Kurshan, R. P., *Computer-Aided Verification of Coordinating Processes: The Automata-Theoretic Approach* Princeton University Press, Princeton, New Jersey 1994.
- [LY92] Lee, D., and Yannakakis, M., On-Line Minimization of Transition Systems, STOC92.
- [LP85] Lichtenstein, O., and Pnueli, A., Checking That Finite State Concurrent Programs Satisfy Their Linear Specifications, POPL85
- [MP92] Manna, Z. and Pnueli, A., Temporal Logic of Reactive and Concurrent Systems: Specification, Springer-Verlag, 1992
- [SG87] Sistla, A. P. and German, S. M., Reasoning with many Processes, Proceedings of the Symposium on Logic in Computer Science, Ithaca, New York, 1987
- [St81] Streett, R., Propositional Dynamic Logic of Looping and Converse, PhD Thesis, MIT, 1981.
- [VW84] Vardi, M., and Wolper, P., An Automata-Theoretic Framework for Modal Logics of Programs, STOC84
- [VW86] Vardi, M., and Wolper, P., An Automata-Theoretic Framework for Automatic Program Verification, LICS86.

FIG 1 : Global Transition Graph

FIG 2 : The Reduced Graph

FIG 3 : A scsg $\bar{C}$

FIG 4 : The threaded graph C^*

5 Appendix

In this appendix, we summarize the formal definitions and proofs omitted from the main sections.

Lemma 3.1 Two ($\mathcal{G}$ -)equivalent states in M satisfy the same set of Fair Indexed CTL* formulas.

Proof. We can argue by induction on formula structure that for any formula f of Fair Indexed CTL*, and any permutation π in $Sym\ I$, f and $\pi(f)$ are equivalent, i.e. express the same property. Now, assume $s \equiv_{\mathcal{G}} t$ so that $t = \pi(s)$ for some $\pi \in \mathcal{G}$. We have $M, s \models f$ iff $M, \pi(s) \models \pi(f)$ (since $\pi \in \mathcal{G} \subseteq Aut\ M$) iff $M, t \models f$ (since $t = \pi(s)$ and $\pi(f)$ is equivalent to f). $\square$

For each $i \in I$ there is a Buchi automaton $\mathcal{A}_i$ corresponding to f_i , which has input symbols that are sets of global propositions Q and local propositions of the form P_i . All of these various $\mathcal{A}_i$ are isomorphic upto reindexing as are the f_i . Hence, it is convenient to define a single, *generic automaton* $\mathcal{A}$ which is, intuitively, the same as any $\mathcal{A}_i$ but the indexes have been stripped off the local propositions P_i . Technically, there is a transition from state q to state r in $\mathcal{A}$ on input $\{Q, \dots, Q', P, \dots, P'\}$ iff there is a transition from q to r in $\mathcal{A}_i$ on input $\{Q, \dots, Q', P_i, \dots, P'_i\}$ (any i will do). The generic automaton $\mathcal{A}$ can be used for any coordinate j , in lieu of $\mathcal{A}_j$, using the operator $s|j$ which, in effect, projects the label of state s onto coordinate j and then strips off the subscripts. Note that $\mathcal{G}$ does not act on the unindexed $\mathcal{A}$.

We will now define the product graph B to be essentially $M \times \mathcal{A} \times I$, which will be used in the proof but is not actually constructed by the algorithm. Intuitively, B denotes the simulation of the automaton $\mathcal{A}$ on the execution of different processes indicated by the process index j in each node. It is to be noted that all edges in B are between nodes with the same process index.

Definition: Formally, $B = (V_B, E_B)$ where $V_B = S \times Q \times I$ is the set of nodes (s, q, j) and E_B is the set of edges of B ; edge $(s, q, j) \rightarrow (s', q', j') \in E_B$ iff $j = j'$, $(s, s') \in R$, and $q' \in \delta_{\mathcal{A}}(q, s|j)$ (i. e., there is a transition of the automaton $\mathcal{A}_j$ from state q to the state q' on the input symbol comprised of the local propositions of index j as well as the global propositions that label state $\bar{s}$).

For any permutation π on I and any node (s, q, j) in V , $\pi((s, q, j))$ is the node $(\pi(s), q, \pi(j))$. For every $\pi \in Aut\ M$, we have $(s, q, j) \rightarrow (t, r, j) \in B$ iff $\pi((s, q, j)) \rightarrow \pi((t, r, j)) \in B$. Hence every $\pi \in Aut\ M$ is also an automorphism of B .

We now define the annotated product graph $\bar{B}$ which is of the form $\bar{M} \times \mathcal{A} \times I$; however, it is the product of the annotated quotient structure $\bar{M}$ with the automaton $\mathcal{A}$ for *permuted* process indices. Thus, in general, the edges in $\bar{B}$ can be between nodes with *different* process indices, while this is not the case in B .¹¹

Definition: Formally, $\bar{B} = (\bar{V}, \bar{E})$, where $\bar{V} = \bar{S} \times Q \times I$ is the set of nodes $(\bar{s}, q, j)$, and $\bar{E}$ is the set of edges. Each edge is of the form (x, π, y) , also denoted $x \xrightarrow{\pi} y$, where $x, y \in \bar{V}$, and $\pi \in \mathcal{G}$; edge $(\bar{s}, q, j) \xrightarrow{\pi} (\bar{t}, r, k) \in \bar{E}$ iff $\bar{s} \xrightarrow{\pi} \bar{t} \in \bar{M}$, $\pi^{-1}(j) = k$, and $r \in \delta_{\mathcal{A}}(q, \bar{s}|j)$ (i.e., there is a transition of the automaton $\mathcal{A}_j$ from state q to state r on the input symbol comprised of the local propositions of index j and the global propositions that label state $\bar{s}$).

Recall that for any state $s \in S$, $\bar{s}$ denotes the unique representative of the $\equiv_{\mathcal{G}}$ -equivalence class containing s . We extend the notion of representative to nodes of B .

¹¹ $\bar{B}$ would be built by a straightforward implementation of the algorithm. The space requirements may be reduced by constructing the threaded graph B^* directly, rather than $\bar{B}$. The idea is that (the salient information from the) collection of annotated edges $\{\bar{u} \xrightarrow{\pi} \bar{v}, \dots, \bar{u} \xrightarrow{\pi} \bar{v}\}$ from $\bar{u}$ to $\bar{v}$ in $\bar{B}$ is represented by a collection of at most n^2 edges in B^* of the form $\{(\bar{u}, i) \rightarrow (\bar{v}, \pi^{-1}(i)) : i \in I \text{ and } \pi \text{ ranges over } \pi_a, \dots, \pi_c\}$.

Definition: For any node $u = (s, q, j) \in V_B$, let $\bar{u}$ denote the *representative node* $(\bar{s}, q, \pi_s(j))$ where π_s is the designated permutation, associated with $\bar{M}$, that maps s to $\bar{s}$, i.e. $\pi_s(s) = \bar{s}$. For any set of nodes $C \subseteq V_B$, let $\bar{C} = \{\bar{u} : u \in C\}$. The subgraph of $\bar{B}$ induced by C has $\bar{C}$ as its set of nodes and $\{\bar{c} \xrightarrow{\pi} \bar{d} : c, d \in C \text{ and } \bar{c} \xrightarrow{\pi} \bar{d} \in E_B\}$ as its set of edges.

Definition: For any finite path $\bar{x}$ in $\bar{B}$, we define $\pi_{\bar{x}}$ to be the product of permutations appearing on the path from left to right. Let $\bar{C}$ be a scsg in $\bar{B}$ and let $\bar{c}$ be a node in $\bar{C}$. We define the subgraph of $\bar{B}$, denoted by $\text{unwind}(\bar{C}, \bar{c})$, obtained by unwinding $\bar{C}$ starting from $\bar{c}$ as follows. The set of nodes in $\text{unwind}(\bar{C}, \bar{c})$ is exactly the set of nodes $\pi_{\bar{x}}(\bar{u})$ where $\bar{x}$ is a path (entirely) in (the subgraph induced by) $\bar{C}$ starting from $\bar{c}$ to a node $\bar{u}$. The set of edges are of the form $(\pi_{\bar{x}}(\bar{u}), \pi_{\bar{y}}(\bar{v}))$ where $\bar{x}$ and $\bar{y}$ are paths starting in $\bar{C}$ starting from $\bar{c}$ and ending with $\bar{u}$ and $\bar{v}$ respectively, and such that $\bar{y}$ is an extension of $\bar{x}$ by a single edge from $\bar{u}$ to $\bar{v}$ in $\bar{C}$.

Theorem 3.2 $M, \bar{s}_0 \models E(\Phi \wedge f_i)$ iff B contains an obviously fair, green scsg C reachable starting from node $(\bar{s}_0, q_0, i)$ where q_0 is the start state of the automaton $\mathcal{A}$.

Proof. To prove the forward direction, assume that $M, \bar{s}_0 \models E(\Phi \wedge f_i)$. Then there is an infinite path $p = \bar{s}_0, s_1, s_2, \dots$ in M starting from the state $\bar{s}_0$, which satisfies the fairness constraint Φ and the formula f_i . Hence, there exists an accepting run $\rho = q_0, q_1, q_2, \dots$ of the automaton $\mathcal{A}$ along the input string obtained by projecting the path p onto coordinate i . When we combine p with this run ρ , we get an infinite path $(\bar{s}_0, q_0, i), (s_1, q_1, i), (s_2, q_2, i), \dots$ in B and the set of nodes appearing infinitely often on this path defines an obviously fair green scsg C reachable from $(\bar{s}_0, q_0, i)$.

To prove the reverse direction, assume that an obviously fair green scsg C is reachable from $(\bar{s}_0, q_0, i)$ in B . >From this we can get a path in B starting from $(\bar{s}_0, q_0, i)$ and visiting each node of C infinitely often; this path satisfies the fairness constraint Φ and contains an accepting node of $\mathcal{A}$ infinitely often. This path, when projected on to the state components of M , gives us a path p in M starting from $\bar{s}_0$ which satisfies the formula f_i and the fairness constraint Φ . $\square$

Theorem 3.3 $M, \bar{s}_0 \models E(\Phi \wedge f_i)$ iff $\bar{B}$ contains a subtly fair, green scsg $\bar{C}$ reachable starting from node $(\bar{s}_0, q_0, i)$.

Proof. To prove the forward direction, assume that $M, \bar{s}_0 \models E(\Phi \wedge f_i)$. By theorem 3.2, we see that B contains an obviously fair, green scsg C reachable from $(\bar{s}_0, q_0, i)$. Using lemma 3.5 below, we see that $\bar{C}$ is a subtly fair, green scsg of $\bar{B}$. Furthermore, using lemma 3.4 below, we see that there is a path in $\bar{B}$ from $(\bar{s}_0, q_0, i)$ to the scsg $\bar{C}$.

To prove the reverse direction, assume that $\bar{B}$ contains a subtly fair, green scsg $\bar{C}$ reachable from $(\bar{s}_0, q_0, i)$. Using lemma 3.6 below, we see that $\text{unwind}(\bar{C}, \bar{c})$ is an obviously fair scsg of B where $\bar{c}$ is any node in $\bar{C}$. >From lemma 3.4 we see that, if p is a path in $\bar{B}$ from the node $(\bar{s}_0, q_0, i)$ to $\bar{c}$ in $\bar{C}$ then $\text{unwind}(p)$ is a path in B from $(\bar{s}_0, q_0, i)$ to some node in $\text{unwind}(\bar{C}, \bar{c})$. Hence $\text{unwind}(\bar{C}, \bar{c})$ is a green obviously fair scsg in B which is reachable from $(\bar{s}_0, q_0, i)$. By theorem 3.2, we conclude that $M, \bar{s}_0 \models E(\Phi \wedge f_i)$. $\square$

Lemma 3.4 (Correspondence of paths in B and $\bar{B}$)

Let N be a nonempty subgraph of B , and $\bar{N}$ the subgraph of $\bar{B}$ induced by N .

(a) If $\bar{x} = \bar{u}_0 \xrightarrow{\pi_1} \bar{u}_1 \xrightarrow{\pi_2} \bar{u}_2 \dots$ is a path in $\bar{N}$, then the path obtained by unwinding $\bar{x}$,

$x = \text{unwind}(\bar{x}) \stackrel{\text{def}}{=} \bar{u}_0 \rightarrow \pi_1(\bar{u}_1) \rightarrow \pi_1 \circ \pi_2(\bar{u}_2) \dots$ is a path in B .

(b) If $x = \bar{u}_0 \rightarrow u_1 \rightarrow u_2 \dots$ is a path in N starting at representative state $\bar{u}_0$, then there exists a path in $\bar{N}$

$\bar{x} \stackrel{\text{def}}{=} \bar{u}_0 \xrightarrow{\pi_1} \bar{u}_1 \xrightarrow{\pi_2} \bar{u}_2 \dots$
 such that $x = \text{unwind}(\bar{x})$.

Proof. To prove (a), it is enough if we show that, for each $j \geq 1$, $\pi_1 \circ \pi_2 \circ \dots \circ \pi_j(\bar{u}_j) \rightarrow \pi_1 \circ \pi_2 \circ \dots \circ \pi_{j+1}(\bar{u}_{j+1})$ is a transition in B . Since $\mathcal{G}$ is a group, we know that $\pi_1 \circ \pi_2 \circ \dots \circ \pi_j$ is a permutation in $\mathcal{G}$, and an automorphism of B . Furthermore, from the definition of $\bar{B}$, we know that $\bar{u}_j \rightarrow \pi_{j+1}(\bar{u}_{j+1})$ is a transition of B . The desired result follows.

To prove (b), we note that we can write any such x in the form $\bar{u}_0 \rightarrow \phi_1(\bar{u}_1) \rightarrow \phi_2(\bar{u}_2) \rightarrow \dots$ where, for each $j \geq 1$, ϕ_i is any permutation in $\mathcal{G}$ such that $\phi_i(\bar{u}_i) = u_i$. We will argue by induction on j that we can take ϕ_j to be a permutation of the form $\pi_1 \circ \dots \circ \pi_j$ as in the statement of the lemma. For $j = 1$, since $\bar{u}_0 \rightarrow \phi_1(\bar{u}_1) \in B$, by definition of $\bar{B}$, there is some $\pi_1 \in \mathcal{G}$ and $\bar{u}_0 \xrightarrow{\pi_1} \bar{u}_1 \in \bar{B}$ such that $\pi_1(\bar{u}_1) = \phi_1(\bar{u}_1)$. Thus, we can take ϕ_1 to be π_1 . Inductively, we can take $\phi_j = \pi_1 \circ \dots \circ \pi_j$. Because $\phi_j(\bar{u}_j) \rightarrow \phi_{j+1}(\bar{u}_{j+1}) \in B$, $\bar{u}_j \rightarrow (\pi_1 \circ \dots \circ \pi_j)^{-1} \circ \phi_{j+1}(\bar{u}_{j+1}) \in B$, by induction hypothesis and since ϕ_j^{-1} is an automorphism of B . Hence, there is some $\pi_{j+1} \in \mathcal{G}$ and some $\bar{u}_j \xrightarrow{\pi_{j+1}} \bar{u}_{j+1} \in \bar{B}$ such that $\pi_{j+1}(\bar{u}_{j+1}) = (\pi_1 \circ \dots \circ \pi_j)^{-1} \circ \phi_{j+1}(\bar{u}_{j+1})$. Thus, we can take $\phi_{j+1} = \pi_1 \circ \dots \circ \pi_j \circ \pi_{j+1}$, thereby completing the induction step. Then, the path $\bar{x} = \bar{u}_0 \xrightarrow{\pi_1} \bar{u}_1 \xrightarrow{\pi_2} \bar{u}_2 \dots$ is such that $\text{unwind}(\bar{x}) = x$. Moreover, since x is in N , then all the representatives $\bar{u}_j$ are in $\bar{N}$, and $\bar{x}$ is itself in $\bar{N}$.

Lemma 3.5 If C is an obviously fair, green scsg of B , then the subgraph $\bar{C}$ of $\bar{B}$, induced by the representatives of nodes in C , is a subtly fair, green scsg of $\bar{B}$.

Proof. Assume that C is an obviously fair, green scsg of B . Let u', v' be any two nodes in $\bar{C}$. There exist nodes u, v in C such that $u' = \bar{u}$ and $v' = \bar{v}$. Because C is an scsg of B , there exists a path in C from u to v . Now, since u' is the representative of u , $u' \equiv_{\mathcal{G}} \bar{u}$, and there is a permutation $\pi_u \in \mathcal{G}$ such that $\pi_u(u) = u'$. As π_u is an automorphism of M and also of B , there also exists a path from $\bar{u} = u' = \pi_u(u)$ to $\pi_u(v)$ in $\pi_u(C)$, which is an isomorphic copy of C in B whose set of representatives coincides with that of C thereby inducing the same subgraph $\bar{C}$ of B as does C . Using lemma 3.4, it follows that there exists a path in $\bar{C}$ from $u' = \bar{u}$ to $v' = \bar{v}$. Hence $\bar{C}$ is strongly connected and is a scsg.

To see that $\bar{C}$ is subtly fair, consider any node u' in $\bar{C}$ and any index j such that u' satisfies en_j .¹² There exists a node u in C such that u' is a representative of u . Let $\pi \in \mathcal{G}$ be a permutation such that $\pi(u) = u'$. Since u' satisfies en_j , it is the case that u satisfies $en_{\pi^{-1}(j)}$. Since C is obviously fair, there exists another node v in C such that v satisfies $ex_{\pi^{-1}(j)}$. Let y be a path in C from u to v . Since π is an automorphism of B , the sequence $x = \pi(y)$ is also a path in B , from $u' = \pi(u)$ to $\pi(v)$. By lemma 3.4, there exists a path $\bar{x}$ in $\bar{C}$ from u' to $\bar{v}$, such that $\text{unwind}(\bar{x}) = x$. Hence $\pi(v) = \pi_{\bar{x}}(\bar{v})$ (Recall that $\pi_{\bar{x}}$ is the product of permutations appearing on the edges in the path $\bar{x}$). Since v satisfies $ex_{\pi^{-1}(j)}$, it is the case that $\pi(v)$, and hence $\pi_{\bar{x}}(\bar{v})$, satisfies ex_j .¹³ Since $\pi_{\bar{x}}(\bar{v})$ satisfies ex_j , it is the case that $\bar{v}$ satisfies $ex_{\pi_{\bar{x}}^{-1}(j)}$. From this, it follows that $\bar{C}$ is subtly fair. Since $\mathcal{G}$ leaves states of $\mathcal{A}$ invariant, it should be clear that $\bar{C}$ is green.

¹² We say that node $u' = (s', q', i')$ of B satisfies a proposition P provided that its state component s' satisfies P , i.e., contains P in its label. Here P may be a global proposition or a local one Q_j ; in the latter case there need be no relationship between j and i' .

¹³ Here, we use this principle. If u satisfies P then $\pi(u)$ satisfies $\pi(P)$ for any π in $\mathcal{G}$. The justification is that we are simply reindexing both sides of "satisfies". The restriction to $\mathcal{G}$ a subgroup of $\text{Aut } M$ is to ensure that $\pi(u)$ is a well-defined node of B .

Lemma 3.6 If $\bar{C}$ is a subtly fair, green scsg of $\bar{B}$ and $\bar{c}$ is any node of $\bar{C}$, then $\text{unwind}(\bar{C}, \bar{c})$ is an obviously fair, green scsg of B .

Proof. Assume that $\bar{C}$ is a subtly fair, green scsg of $\bar{B}$ and $\bar{c}$ is any node of $\bar{C}$. We first show that $\text{unwind}(\bar{C}, \bar{c})$ is a scsg in B . By definition $\bar{c}$ is in $\text{unwind}(\bar{C}, \bar{c})$. Let u be any node in $\text{unwind}(\bar{C}, \bar{c})$. We will show that there is a cycle in $\text{unwind}(\bar{C}, \bar{c})$ which contains both $\bar{c}$ and u . This will ensure that all the nodes in $\text{unwind}(\bar{C}, \bar{c})$ are strongly connected. >From the definition of $\text{unwind}(\bar{C}, \bar{c})$, we see that there exists a path $\bar{x}$ from $\bar{c}$ to $\bar{u}$ in $\bar{C}$ such that $\pi_{\bar{x}}(\bar{u}) = u$. Since $\bar{C}$ is a scsg, it follows that there exists a cycle $\bar{y}$ which is an extension of $\bar{x}$ (i.e., $\bar{x}$ is a prefix of $\bar{y}$), such that $\bar{y}$ starts and ends with $\bar{c}$. The permutation $\pi_{\bar{y}}$ is in $\mathcal{G}$. Hence there exists an integer $l > 0$ such that $(\pi_{\bar{y}})^l$ is the identity permutation. Now consider the cycle $\bar{y}^l$, i.e. the cycle obtained by repeating $\bar{y}$ exactly l times. Using lemma 3.4, it is easy to see that $\text{unwind}(\bar{y}^l)$ is a cycle in $\text{unwind}(\bar{C}, \bar{c})$ that contains both $\bar{c}$ and u .

Now, we prove that $\text{unwind}(\bar{C}, \bar{c})$ is obviously fair. Let u be any node in $\text{unwind}(\bar{C}, \bar{c})$ and j be any index such that u satisfies en_j . Then, there is a path $\bar{x}$ from $\bar{c}$ to $\bar{u}$ in $\bar{C}$ such that $\text{unwind}(\bar{x})$ is a path from $\bar{c}$ to u in C and $u = \pi_{\bar{x}}(\bar{u})$. Let $k = \pi_{\bar{x}}^{-1}(j)$. Since u satisfies en_j , it is the case that $\bar{u}$ satisfies en_k . Since $\bar{C}$ is subtly fair, there exists a path $\bar{y}$ from $\bar{u}$ to a node $\bar{v}$ in $\bar{C}$ such that $\bar{v}$ satisfies $ex_{\pi_{\bar{y}}^{-1}(k)}$. Now $\bar{x} \cdot \bar{y}$, the fusion of $\bar{x}$ with $\bar{y}$, is a path from $\bar{c}$ to $\bar{v}$ in $\bar{C}$. Hence $\text{unwind}(\bar{x} \cdot \bar{y})$ is a path from $\bar{c}$ to u to $\pi_{\bar{x}\bar{y}}(\bar{v})$ in $\text{unwind}(\bar{C}, \bar{c})$. Since $\bar{v}$ satisfies $ex_{\pi_{\bar{y}}^{-1}(k)}$, $k = \pi_{\bar{x}}^{-1}(j)$, and $\pi_{\bar{x}\bar{y}} = \pi_{\bar{x}}\pi_{\bar{y}}$, we calculate that $\pi_{\bar{x}\bar{y}}(\bar{v})$ satisfies ex_j . This establishes that $\text{unwind}(\bar{C}, \bar{c})$ is obviously fair. As $\mathcal{G}$ does not act on the states of $\mathcal{A}$, it is clear that $\text{unwind}(\bar{C}, \bar{c})$ is green.

Proof of lemma 3.7 is straightforward from the definition of C^* and lemma 3.4, and is left to the reader.

Lemma 3.8 Let $\bar{C}$ be a scsg of $\bar{B}$ and C^* the threaded graph associated with $\bar{C}$. Then, $\bar{C}$ is subtly fair and green iff C^* is plainly fair and green.

Proof. To prove the forward direction, assume that $\bar{C}$ is subtly fair and green. Since $\mathcal{G}$ does not act on $\mathcal{A}$, it is clear that C^* is green. We prove that C^* is plainly fair. Consider any node $(\bar{u}, j)$ in C^* which is labeled with en . This means that the state $\bar{u}$ is labeled with en_j in $\bar{C}$. Since $\bar{C}$ is subtly fair, there exists a path $\bar{x}$ from $\bar{u}$ to a node $\bar{v}$ such that $\bar{v}$ is labeled with $ex_{\pi_{\bar{x}}^{-1}(j)}$. From lemma 3.7, it is seen that there is a thread in C^* from $(\bar{u}, j)$ to $(\bar{v}, \pi_{\bar{x}}^{-1}(j))$. Clearly, the node $(\bar{v}, \pi_{\bar{x}}^{-1}(j))$ in C^* is labeled with ex . Hence C^* is plainly fair. Now we prove the lemma in the other direction. Assume that C^* is plainly fair and green. It is trivial to see that $\bar{C}$ is green. Now consider any node $\bar{u}$ which is labeled with en_j for some j . In C^* the node $(\bar{u}, j)$ is labeled with en . Since C^* is plainly fair, there exists a node $(\bar{v}, k)$, labeled with ex , which is reachable from $(\bar{u}, j)$. >From lemma 3.7, it should be easy to see that there exists a path $\bar{x}$ in $\bar{C}$ from $\bar{u}$ to $\bar{v}$ such that $k = \pi_{\bar{x}}^{-1}(j)$. Clearly, $\bar{v}$ is labeled with ex_k . This proves that C^* is plainly fair.

Theorem 3.9 The algorithm correctly marks exactly those representative states $\bar{s}$ such that $M, \bar{s} \models E(\Phi \wedge f_i)$. It runs in time $O(|\bar{B}| \cdot n^2) = O(|\bar{M}| \cdot |\mathcal{A}| \cdot n^3)$.

Proof. To prove the theorem, it is enough if we show that in the second step of the algorithm, the node $\bar{u}$ is marked iff it belongs to a subtly fair green scsg in $\bar{B}$. >From all our previous lemmas, it is easy to see that if a node $\bar{u}$ is marked then it must belong to a subtly fair green scsg in $\bar{B}$. To prove the other direction we need the following claim.

Claim. If $\bar{C}$ is a scsg in $\bar{B}$ then, C^* is plainly fair iff all the maximal scsgs in C^* are

plainly fair (i.e. if each maximal scsg in C^* contains a node labeled with en then it also contains a node labeled with ex).

Proof. We prove the claim by simply showing that all the maximal scsgs in C^* are disconnected, i.e. C^* does not have edges that connect nodes in different scsgs. It is enough to show that if there is an edge connecting two nodes in C^* then both the nodes belong to the same scsg. Assume that there is an edge in C^* from some node $(\bar{u}, j)$ to some other node $(\bar{v}, k)$. This means that there is an edge e in $\bar{B}$ from $\bar{u}$ to $\bar{v}$ which is labeled with a permutation π such that $\pi^{-1}(j) = k$. Since $\bar{u}, \bar{v}$ belong to the scsg $\bar{C}$, there exists a path $\bar{x}$ starting and ending with the node $\bar{u}$ and using the edge e for the first transition. Clearly, $\pi_{\bar{x}}$, which is the product of permutations along $\bar{x}$, is in $\mathcal{G}$. Since $\mathcal{G}$ is a group, it is easily seen that there exists an integer $l > 0$ such that $(\pi_{\bar{x}})^l$ is the identity permutation. Clearly, $(\pi_{\bar{x}})^l$ is the permutation associated with the path $\bar{x}^l$. Using lemma 3.8, it is easily seen that there is a thread in C^* , corresponding to the path $\bar{x}^l$, which is a cycle and which uses the edge from $(\bar{u}, j)$ to $(\bar{v}, k)$. Hence both the nodes $(\bar{u}, j)$ and $(\bar{v}, k)$ belong to the same scsg in C^* .

Now, assume that $\bar{u}$ is a node in a subtly fair green scsg $\bar{D}$ in $\bar{B}$. It is not difficult to see that in step 2 of the algorithm none of the nodes in $\bar{D}$ are ever going to be deleted. Now, we show that within the $n + 1$ iterations of the outermost loop in step 2, $\bar{u}$ (in fact, all the nodes in $\bar{D}$) is going to be marked. In an outermost iteration at step 2, assume that some nodes in $\bar{B}$ belonging to the mscsg containing $\bar{u}$ are deleted; assume that this deletion occurs due to the presence of bad nodes in the mscsg of the threaded graph containing $(\bar{u}, j)$ for some j (in this case, the mscsg in the threaded graph containing $(\bar{u}, j)$ does not have any nodes labeled with ex and the bad nodes are all the nodes in it that are labeled with en). It is to be noted that, in all future instances after this iteration, the mscsg in the threaded graph containing $(\bar{u}, j)$ will have no more bad nodes, and it will be plainly fair. Hence within at most n iterations, for each $j = 1, \dots, n$, all the mscsgs in the threaded graph containing $(\bar{u}, j)$ will be plainly fair; at this time, the mscsg in $\bar{B}$ containing $\bar{u}$, call it $\bar{C}$, will be subtly fair; this is because C^* will be plainly fair (due to the above claim). It should be now be clear that in the next iteration of step 2 all the nodes in $\bar{D}$ will be marked.

It should be easy to see that step 2 dominates the time complexity. Each iteration of the outermost loop of step 2 can be implemented in time proportional to the size of the threaded version B^* of $\bar{B}$, i.e., $O(|\bar{B}| \cdot n)$ time. Hence, step 2 can be implemented in time $O(|\bar{B}| \cdot n^2) = O(|\bar{M}| \cdot |\mathcal{A}| \cdot n^3)$.

Theorem 3.11. (Model Checking for Fair Indexed CTL*) The above procedure provides an algorithm which checks $M, \bar{s}_0 \models g$ for any Fair Indexed CTL* formula g under strong fair semantics, in time $O(|\bar{M}| \cdot n^3 \cdot |g| \cdot a)$ where $|\bar{M}|$ is the size of $\bar{M}$, $n = |I|$ is the number of processes, $|g|$ is the length of g , and a is the maximum size of the automaton $\mathcal{A}$ for any basic modality of g .

Proof. We can handle any single basic modality in g in time $O(|\bar{M}| \cdot a \cdot n^3)$ by Theorem 3.9 and the definition of a . Then we can handle all of g , which is composed of boolean combinations and nesting of basic modalities, by recursive descent in time at most $O(|\bar{M}| \cdot n^3 \cdot |g| \cdot a)$ (cf. [EL87]).