

# A Dash of Fairness for Compositional Reasoning

Ariel Cohen<sup>1</sup>, Kedar S. Namjoshi<sup>2</sup>, and Yaniv Sa'ar<sup>3</sup>

<sup>1</sup> New York University, New York, NY. Email: arielc@cs.nyu.edu

<sup>2</sup> Bell Labs, Alcatel-Lucent, Murray Hill, NJ. Email: kedar@research.bell-labs.com

<sup>3</sup> Weizmann Institute of Science, Rehovot, Israel. Email: yaniv.saar@weizmann.ac.il

**Abstract.** Proofs of progress properties often require fairness assumptions. Incorporating global fairness assumptions in a compositional method is a challenge, however, given the local flavor of such reasoning. We present a fully automated local reasoning algorithm which handles fairness assumptions. As local reasoning is inherently incomplete, the algorithm incorporates a refinement step, which strengthens local proofs by adding auxiliary shared variables. Experiments show that verification with the new algorithm is often significantly faster than with standard model checking.

## 1 Introduction

Model checking is fundamentally constrained by state explosion [6]. For concurrent programs with several independent processes, the state space can grow exponentially with the number of processes. A promising approach to ameliorating state explosion, referred to as compositional or modular reasoning, is to decompose a verification task so that the reasoning is as localized as possible. In this work, we propose and evaluate a new algorithm for compositional reasoning with fairness constraints.

Fairness assumptions are often required for proofs of progress properties. While it has long been understood how to incorporate fairness in standard model checking [5, 14], this is a challenge for compositional methods. The difficulty is that such assumptions are typically global in nature. For instance, a common (strong) fairness assumption is that “*for every process: if the process is enabled infinitely often, it is infinitely often executed*”, which is stated using a number of shared and local variables (to express enabledness). This conflicts with the localized view taken in compositional reasoning.

This work develops a new algorithm for compositional model checking with fairness assumptions, and a new compositional proof rule. The model checking algorithm can be instrumented to generate a valid instantiation of the proof rule.

The new algorithm is the continuation of a line of research on mechanizing assertional (i.e., predicate-based) compositional verification. The starting point is an algorithm from [21] which computes the strongest *split invariant*, which is a vector of interference-free, per-process invariants. (A set of per-process invariants is free of interference [23, 19] if the action of one process does not invalidate the invariant of another process.) The term ‘split invariant’ is used as the conjunction of the local invariants

forms a globally inductive invariant. The strongest split invariant may be weaker than the set of reachable states, and therefore not strong enough to prove a safety property. In [8], we solve this problem by formulating a complete verification procedure which strengthens the split invariant by adding auxiliary shared variables which track local predicates. In [9], we use split invariants to construct a compositional algorithm for proving LTL properties expressed on the shared variables. A number of experimental results in these papers show that assertional local reasoning can be significantly faster than monolithic (i.e., non-compositional) model checking.

The local liveness method of [9] does not immediately extend to fairness. Incorporating fairness into the specification is allowed by the identity  $M \models A_{Fair}(Spec) \equiv M \models A(Fair \Rightarrow Spec)$ . This, however, results in a new specification which names local variables (due to *Fair*), which violates the requirement that the specification be expressed only on the shared variables. One can, of course, turn all local variables from *Fair* into shared variables, but this defeats the purpose of local reasoning.

The new algorithm gets around this difficulty by working always with a *weakened* form of the original fairness assumption, which is stated over just the shared variables. Weakening is done by quantifying out local variables from the original fairness statement, in a manner which depends monotonically on the split invariant. A process of successive strengthening is then followed. If verification succeeds under a weaker fairness assumption, the property is proved. Otherwise, if a bogus counter-example is produced, the trace is analyzed to discover local predicates which are exposed as auxiliary variables. Exposing local state strengthens the split invariant in the next round of computation, and therefore strengthens the abstracted fairness assumption. This process terminates—and is thus complete—for finite-state programs: eventually, enough of the local state is exposed to either prove a correct property or to disprove an incorrect one. Moreover, it is possible to disprove a property *without* building up the entire state space.

The algorithm, being predicate-based, has a simple implementation using BDDs. We carry out an evaluation with several parameterized protocols, where each instance of the protocol is finite-state. The experimental results show promise: the compositional verification is faster in almost all cases, sometimes by one or two orders of magnitude. Exposing a limited amount of local state suffices for both proofs and disproofs of properties, validating the basic premise behind compositional reasoning.

## 2 Related Work

The question of handling fairness in compositional verification is a natural and important one. The comprehensive book on compositional methods by de Roever et. al. [11], however, does not mention a compositional proof rule directly incorporating fairness. Compositional proof rules for general LTL properties (e.g., [1, 20, 22]) handle fairness only by compiling it into the specification. To the best of our knowledge, this is the first compositional algorithm and proof rule to directly incorporate fairness.

The methods used here are assertional; i.e., they are based on computing state predicates. The “thread-modular” reasoning method [16] computes a split invariant using

explicit-state representations, but is limited to safety properties. An alternative line of work on automated compositional reasoning is based on representing interface behavior, and is thus behavioral in nature. One instance of this method uses the following complete proof rule: to show  $M_1 // M_2 \models Spec$ , find an interface automaton  $A$  such that  $M_1 \models A$  and  $M_2 // A \models Spec$ . (Here,  $\models$  is read as language inclusion.) The procedures developed in [17, 26] employ a combination of model checking and finite-automaton learning via variants of the  $L^*$  method [2] to construct an appropriate  $A$ . Standard learning algorithms compute automata on finite words, and hence can be used only for proofs of safety properties. An algorithm is developed in [15] for learning a Büchi automaton, but it has not yet been applied to verification of progress properties. Although an automaton is a powerful and compact representation object, current implementations of behavioral methods have difficulty showing an improvement over non-compositional model checking [7]. Part of the reason may be in the complexity of the automaton learning procedures.

### 3 Background: Local Reasoning and Liveness Properties

This section defines the system model and split invariance, and gives a short summary of the method for local liveness checking. Part of this material is taken from [21, 8, 9], and is repeated here for convenience.

*A Note on Notation.* Throughout the paper, we use notation based on that of Dijkstra and Scholten [12]. Sets of program states are represented by first-order formula on program variables. Existential quantification of formula  $\xi$  by a set of variables  $X$  is written as  $(\exists X : \xi)$ . The notation  $[\xi]$  stands for “ $\xi$  is valid”. The successor operation is denoted by  $sp$  (for strongest postcondition):  $sp(\tau, \xi)$  represents the set of states reachable from states satisfying  $\xi$  in one  $\tau$ -transition. The notation  $sp_i(\tau, \xi)$  is used for successors computed within the state space of process  $P_i$ .

#### 3.1 Model: Asynchronous, Shared-Memory Composition

A *process* is given by a tuple  $(V, I, T)$ , where  $V$  is a set of (typed) variables,  $I(V)$  is a predicate over  $V$  defining an initial condition, and  $T(V, V')$  is a predicate defining a transition condition, where  $V'$  is a fresh set of variables in 1-1 correspondence with  $V$ . The semantics of a process is given by a *transition system* in the standard way.

The *asynchronous composition* of processes  $\{P_i\}$ , written as  $\parallel_i P_i$ , is the process  $P = (V, I, T)$ , where the components are defined as follows. Let  $V = \bigcup_i V_i$  and  $I = \bigwedge_i I_i$ . The *shared variables*, denoted  $X$ , are those that belong to  $V_i \cap V_j$ , for some distinct pair  $(i, j)$ . The *local variables* of process  $P_i$ , denoted  $L_i$ , are the variables in  $V_i$  that are not shared (i.e.,  $L_i = V_i \setminus X$ ). The set of local variables is  $L = \bigcup_i L_i$ . Let  $\hat{T}_i = T_i(V_i, V'_i) \wedge (\forall j: j \neq i \Rightarrow pres(L_j))$ , where  $pres(W)$  is short for  $\bigwedge_{w \in W} (w' = w)$ . Thus,  $\hat{T}_i$  behaves like  $T_i$ , but leaves local variables of other processes unchanged. The transition relation of the composition,  $T$ , is defined as  $\bigvee_i \hat{T}_i$ .

### 3.2 Split-Invariance: Definition and Calculation

Let  $P = \parallel_k P_k$  be an  $N$ -process composition. For localized reasoning about invariance, the shape of invariance assertions is restricted to a conjunction of local (i.e., per-process) assertions. A *local* assertion is one that is based on the variables of a single process. A *split assertion* is a vector of local assertions,  $\theta = (\theta_1, \theta_2, \dots, \theta_N)$ , one for each process, so that  $\theta_i$  is defined on  $V_i$  (equivalently, on  $X$  and  $L_i$ ). Split assertion  $\theta$  is a *split invariant* if the conjunction of its components, i.e.,  $\bigwedge_k \theta_k$ , is an inductive invariant for the full program  $P$ . Split-invariance can equivalently be defined as in Figure 1.

**Definition 0** The notation  $[T_k]^\theta(X, X')$  denotes  $(\exists L_k, L'_k : T_k \wedge \theta_k)$ . This is a “summary transition”, representing the effect that a move of  $P_k$  from a state satisfying its local invariant has on the shared variables.

For each process index  $i$ :

1. **[initiality]**  $\theta_i$  includes all initial states of process  $P_i$ . I.e.,  $[(\exists L \setminus L_i : I) \Rightarrow \theta_i]$
2. **[step]**  $\theta_i$  is closed under transitions of  $P_i$ . I.e.,  $[sp_i(T_i, \theta_i) \Rightarrow \theta_i]$
3. **[non-interference]**  $\theta_i$  is closed under transitions by processes other than  $P_i$ . I.e., for all  $k$  different from  $i$ ,  $[sp_i([T_k]^\theta \wedge pres(L_i), \theta_i) \Rightarrow \theta_i]$

**Fig. 1.** Split Invariance Conditions

These conditions are a simple instance of (syntactically circular) *assume-guarantee reasoning*:  $\theta_i$  is the invariance guarantee provided by process  $i$ , based on assumptions  $\{\theta_j : j \neq i\}$  about the other processes. The constraints can be gathered into the set of simultaneous implications: for each  $i$ ,

$$[(\exists L \setminus L_i : I) \vee sp_i(T_i, \theta_i) \vee (\forall k : k \neq i : sp_i([T_k]^\theta \wedge pres(L_i), \theta_i)) \Rightarrow \theta_i] \quad (1)$$

**Theorem 0** (Namjoshi [21]) *The strongest split invariant exists. It is the simultaneous least fixpoint of equations (1), which exists by the Knaster-Tarski fixpoint theorem.*

### 3.3 Incompleteness and Auxiliary Variables

Local reasoning is inherently incomplete. This is illustrated by the mutual exclusion protocol from Figure 2. The strongest split invariant for 2 processes is  $(true, true)$ , which is too weak to prove mutual exclusion. A general mechanism for overcoming incompleteness, proposed by Owicki-Gries and Lamport [19], is to add auxiliary shared variables which expose portions of the local state or execution history. In Figure 3, an auxiliary variable records the last process to enter the critical section. The strongest split invariant for the augmented protocol is given by  $\theta_i \equiv (L_2(i) \equiv (x = 0) \wedge (last = i))$ , which suffices to prove mutual exclusion as  $[\theta_i \wedge \theta_j \wedge (i \neq j) \Rightarrow \neg(L_2(i) \wedge L_2(j))]$ . The discovery of auxiliary predicates can be effectively automated [8].

$$\begin{array}{l}
x: \text{boolean initially } x = 1 \\
\parallel_{i=1}^N P[i] :: \left[ \begin{array}{l} \text{loop forever do} \\ \left[ \begin{array}{l} l_0: \text{Non-Critical} \\ l_1: \text{request } x; \\ l_2: \text{Critical} \\ l_3: \text{release } x \end{array} \right] \end{array} \right]
\end{array}$$

Fig. 2. MUXSEM

$$\begin{array}{l}
\text{last}: 0..N \text{ initially last} = 0 \\
x: \text{boolean initially } x = 1 \\
\parallel_{i=1}^N P[i] :: \left[ \begin{array}{l} \text{loop forever do} \\ \left[ \begin{array}{l} l_0: \text{Non-Critical} \\ l_1: \langle \text{request } x; \text{last} := i; \rangle \\ l_2: \text{Critical} \\ l_3: \text{release } x \end{array} \right] \end{array} \right]
\end{array}$$

Fig. 3. MUXSEM with auxiliary variable

### 3.4 Local Verification of Liveness Properties

Owicki and Gries also developed compositional proof rules for termination. In [9], a related proof rule is turned into a compositional algorithm for checking general linear-time temporal properties. This “local liveness” method, referred to subsequently as the LL algorithm, is shown in Figure 4. We give a sketch of the soundness proof, as it is important for the extension to fairness. The LL algorithm requires that the LTL property is expressed only by shared variables. With this method, one can show that the property “infinitely often ( $x = 0$ )” holds for the protocol in Figure 2—i.e., that some process is in the critical section infinitely often. Starvation freedom, however, holds only under a strong fairness assumption, and its compositional proof requires the new method.

#### Local Liveness (LL) Algorithm

1. Compute the strongest split invariant,  $\theta$ .
2. For each  $i$ : build an abstract form of process  $i$ , called  $[P_i]^\theta$ , with initial states those of process  $i$ , and two kinds of transitions:
  - the transition  $T_i$  of process  $i$ , and
  - summary transitions  $[T_j]^\theta$  (see Defn. 0) for all other processes  $P_j$  ( $j \neq i$ )
3. Form a Büchi automaton for the *negated* specification. For each  $i$ , form the synchronous product of this automaton with  $[P_i]^\theta$  and check that *there is no computation where infinitely often there is a process  $i$  transition from a Buchi accepting state*
4. Declare success if the check succeeds **for each** abstract process

Fig. 4. Local Liveness (LL) Algorithm.

**Theorem 1** (Cohen-Namjoshi [9]) *The LL proof method is sound.*

**Proof Sketch.** The soundness proof shows the following: if a property does not hold, any global counter-example can be projected to a counter-example for **some** abstract process. Let  $\sigma$  be a global counter-example. Then (1) each state of  $\sigma$  must satisfy the split invariant and (2) the Büchi automaton must accept infinitely often along  $\sigma$ . As there is a fixed number of processes, by (2), there is a process, say  $P_i$ , whose transition is executed infinitely often from a Büchi accepting state along  $\sigma$ . Consider the abstract process  $[P_i]^\theta$  formed out of  $P_i$ . As shown in detail in [9], the computation  $\sigma$  can be

projected, transition-by-transition, to an execution of  $[P_i]^\theta$ . A transition by process  $P_i$  is kept as is; a transition by another process, say  $P_k$ , is replaced by its summary transition,  $[T_k]^\theta$ . Thus, in the projected computation, there are infinitely many positions where there is a transition by  $P_i$  from an accepting automaton state. Any summary transition preserves the change to shared variables made by the original; hence, the sequence of shared-variable values is identical in the original and the projected computations. As the automaton checks properties defined only over shared variables, its accepting run carries over to the projected computation. Thus, the check in Step 3 fails for process  $[P_i]^\theta$ .  $\square$

## 4 Fairness

We describe the modifications necessary to incorporate fairness assumptions into the local liveness method. We begin with a simple but useful kind of fairness, called *unconditional fairness*.

### 4.1 Unconditional Fairness

This fairness notion is a foundational concept in the UNITY programming language and proof system [3], and it suffices for many interesting distributed protocols. Under unconditional fairness, every process is scheduled infinitely often in an infinite computation. The statement uses “scheduled” rather than “executed”—a process may be scheduled but do nothing (i.e., behave as *skip*) because its transition is not enabled. To analyze a protocol under unconditional fairness, Step 3 of the local liveness method is modified to check that, for each  $[P_i]^\theta$ , *there is no unconditionally fair computation where infinitely often there is a process  $i$  transition from a Buchi accepting state*.

**Theorem 2** *The LL method modified for unconditional fairness is sound.*

**Proof Sketch.** The proof sketch for Theorem 1 shows that the sequence of process identifiers associated with the transitions is identical in the original and the projected computations. As the original error computation is unconditionally fair by assumption, the projected error computation must also be unconditionally fair. This argument shows that the modified check is sound.  $\square$

### 4.2 Strong Fairness

The strong fairness algorithm is based on iterated refinement. The idea is to start with a weakened form of the strong fairness assumption, and use the refinement mechanism which adds auxiliary variables to strengthen this assumption with each iteration, until a conclusive result is obtained. To keep the notation simple, we consider a common form of strong fairness, represented by  $\Phi \equiv (\bigwedge i : \text{FG}(p_i) \vee \text{GF}(q_i))$ , where  $p_i$  and  $q_i$  are assertions over the variables of process  $P_i$ . Recall that the proof of soundness of the local liveness method projects a global counter-example,  $\sigma$ , on to a local computation of abstract process  $[P_k]^\theta$ , for some  $k$ . In the presence of fairness, there are two key properties of  $\sigma$ :

1. Every state on  $\sigma$  satisfies  $\bigwedge_i \theta_i$ , as  $\theta$  is a split invariant, and
2.  $\sigma$  satisfies the fairness assumption  $\bar{\Phi}$

Taken together, this implies—crucially—that  $\sigma$  must also satisfy the *stronger* fairness assertion,  $\bar{\Phi}^*$ , given by  $(\bigwedge_i i : \text{FG}(\theta_i \wedge p_i) \vee \text{GF}(\theta_i \wedge q_i))$ . The fact that  $\bar{\Phi}^*$  is stronger than  $\bar{\Phi}$  for any  $\theta$  follows from the monotonicity of  $G$  and  $F$ . The fact that  $\bar{\Phi}^*$  holds for  $\sigma$  follows by the first property: as every state on  $\sigma$  satisfies  $(\bigwedge_j j : \theta_j)$ , assertions  $\text{FG}(\theta_i \wedge p_i)$  and  $\text{FG}(p_i)$  are equivalent on  $\sigma$ , as are assertions  $\text{GF}(\theta_i \wedge q_i)$  and  $\text{GF}(q_i)$ .

The abstract fairness property is formed by quantifying out local variables from  $\bar{\Phi}^*$ , as follows.

$$[\bar{\Phi}]^\theta = (\bigwedge_i i : \text{FG}((\exists L_i : \theta_i \wedge p_i)) \vee \text{GF}((\exists L_i : \theta_i \wedge q_i)))$$

Subsequently, we refer to the term  $(\exists L_i : \theta_i \wedge p_i)$  as  $[p_i]^\theta$  and to  $(\exists L_i : \theta_i \wedge q_i)$  as  $[q_i]^\theta$ . The transformed fairness property is weaker than  $\bar{\Phi}^*$ , but not necessarily weaker than  $\bar{\Phi}$ , and it is defined over the shared variables only.

It is important that  $[\bar{\Phi}]^\theta$  depends on  $\theta$ , and does so in a monotonic manner. This enables refinement: as the split invariant is strengthened by adding auxiliary variables, the abstract fairness assumption also becomes stronger. The new method is shown in Figure 5; other than a modified check at Step 3, it is identical to the LL method from Figure 4.

#### Fair Local Liveness (FLL) Algorithm

1. Compute the strongest split invariant,  $\theta$ .
2. For each  $i$ : build an abstract form of process  $i$ ,  $[P_i]^\theta$ , as defined in Figure 4
3. Form a Büchi automaton for the *negation* of the specification. For each  $i$ , form the synchronous product of this automaton with  $[P_i]^\theta$  and check that *there is no computation which is strongly fair according to  $[\bar{\Phi}]^\theta$  and on which infinitely often there is a process  $i$  transition from a Büchi accepting state*
4. Declare success if the check succeeds **for each** abstract process

**Fig. 5.** Fair Local Liveness (FLL) Algorithm.

**Theorem 3** *The FLL method is sound.*

**Proof.** This proof is an extension of the proof of Theorem 1. Consider a global counter-example  $\sigma$  which is fair according to  $\bar{\Phi}$ . By the proof of Theorem 1, the projection of  $\sigma$  on  $[P_i]^\theta$  satisfies the second part of the condition of Step 3: i.e., infinitely often there is a process  $i$  transition from a Büchi accepting state. It remains to be shown that the projected computation also satisfies  $[\bar{\Phi}]^\theta$ .

As  $\sigma$  is a counter-example based on the fairness assumption, it satisfies  $\bar{\Phi}$ ; as it is a program computation, it satisfies the split invariant,  $\theta$ . Hence, by the reasoning above, it satisfies  $\bar{\Phi}^*$  and therefore the weaker property  $[\bar{\Phi}]^\theta$ . As  $[\bar{\Phi}]^\theta$  is a property over shared state only, and the sequence of values for shared variables is preserved by the projection,  $[\bar{\Phi}]^\theta$  holds also of the projected computation.  $\square$



**Remark 0** Our implementation uses a stronger abstraction of the fairness property. In Step 3 of the FLL algorithm, instead of the uniform assumption  $[\Phi]^\theta$ , the implementation uses a fairness assumption for  $[P_i]^\theta$  where all terms from  $\Phi$  are abstracted relative to  $\theta$  as described above, *except* the term  $(FG(p_i) \vee GF(q_i))$ , which is used as is, since it refers only to variables of process  $P_i$ .

### 4.3 Refinement for Fairness

As local reasoning is approximate, it is possible for the FLL method to fail even though the property is true of the whole program. One can analyze the failure, though, to suggest auxiliary Boolean variables which expose local state predicates, as shown in Figure 6, which extends the refinement procedure used for the LL method.

#### Refinement for Fair Local Liveness

1. Check if every summary transition in the abstract counter-example  $\sigma$  is a **MUST** transition for the process which makes it. If not, expose a local predicate for the **MUST** condition, as defined in [9] for the LL method, and **REPEAT** the full verification.
2. Inductively construct a global computation  $\delta$  which matches  $\sigma$
3. Check if  $\delta$  satisfies the original fairness condition,  $\Phi$ . If so, **HALT** with  $\delta$  as the valid global counter-example.
4. Use a fairness term  $(FG(p_j) \vee GF(q_j))$  which is *not* satisfied by  $\delta$  to discover and expose a local predicate, and **REPEAT** full verification.

**Fig. 6.** Refinement for the FLL method, given a counter-example  $\sigma$  in the abstract process  $[P_i]^\theta$ .

Step 1 is the refinement step for LL. Recall that a transition in  $\sigma$  by a process  $P_k$  where  $k \neq i$  only modifies the shared variables. A change of shared state from  $X = a$  to  $X' = b$  is considered a **MUST** transition if this change is possible no matter what the local state of process  $P_k$  may be, so long as it is consistent with  $\theta_k$ . The predicate  $m(L_k) \equiv \theta_k(a, L_k) \wedge \neg(\exists L'_k : T_k(a, L_k, b, L'_k))$  expresses this succinctly: the transition from  $X = a$  to  $X' = b$  is **MUST** transition if, and only if,  $m$  is unsatisfiable. If  $m$  is satisfiable, it is “exposed” by adding an auxiliary shared variable  $x_m$ . The constraint  $x'_m \equiv m(L'_k)$  is added to the transition relation of  $P_k$ , and the constraint  $x'_m \equiv x_m$  to that for all other processes. Together with the initialization of  $x_m$  to  $m(L_k)$ , these constraints maintain the global invariant  $(x_m \equiv m)$ .

Regarding Step 2, if each summary transition in  $\sigma$  is a **MUST** transition, it is possible to inductively construct a global computation  $\delta$  which matches  $\sigma$ . The initial values for the local variables for processes other than  $P_i$  can be chosen arbitrarily, consistent with the initial condition. Inductively, the **MUST** property guarantees that a concrete transition can be found for each process making a summary move such that the change to the shared state is preserved. Although  $\sigma$  satisfies  $[\Phi]^\theta$ , it need not be the case that  $\delta$  satisfies  $\Phi$ . If  $\Phi$  fails to hold on the computed  $\delta$  (Step 3), the proof of Theorem 5 shows how a new predicate can be derived by analyzing this failure.



**Theorem 4** (*Soundness for failures*) *If the FLL refinement procedure declares a counter-example trace to be real, the trace is a valid counter-example under strong fairness.*

**Proof.** Follows from the reasoning given for Steps 2 and 3.  $\square$

**Theorem 5** (*Completeness for finite-state programs*) *For finite-state programs, the combination of the FLL procedure with refinement terminates.*

**Proof.** It suffices to show that a new predicate—one that is not a Boolean combination of existing predicates—is added at each refinement step. Termination follows, as there is a finite number of distinct predicates. Theorems 3 and 4 show that each termination outcome is correct; thus, the method is complete. For Step 1, the fact that a new predicate is added was shown for the LL method in [9]. For the predicate added at Step 4, it can be shown as follows.

If the check at Step 3 fails, there is a term,  $(FG(p_j) \vee GF(q_j))$ , for some  $j$ , which fails to hold for  $\delta$ . Thus, from some point on, all states on  $\delta$  fail  $q_j$ , and infinitely often, there is a state failing  $p_j$ . Depending on which sub-term is used to satisfy  $(FG([p_j]^\theta) \vee GF([q_j]^\theta))$  on  $\sigma$ , there is a state  $s$  that is on  $\sigma$  and its corresponding state  $t$  on  $\delta$  such that either (i)  $s$  satisfies  $[p_j]^\theta$  and  $t$  does not satisfy  $p_j$  or (ii)  $s$  satisfies  $[q_j]^\theta$  while  $t$  does not satisfy  $q_j$ .

Consider the first case, the proof of the second is similar. By the definition of  $[p_j]^\theta$  as  $(\exists L_j : \theta_j \wedge p_j)$ , there is a valuation  $c$  for  $L_j$  such that for  $u = (s(X), c)$  it is the case that  $\theta_j(u)$  and  $p_j(u)$  both hold. For  $v = (t(X), t(L_j))$ , however, while  $\theta_j(v)$  holds by the invariance of  $\theta$  for the concrete computation  $\delta$ ,  $p_j(v)$  does not hold by the assumption. From the correspondence of  $s$  and  $t$ , they agree on the value of shared variables. Thus,  $u$  and  $v$  must differ only on the valuation of  $L_j$ . Let  $q$  be a predicate expressing this difference (e.g.,  $q(L_j) = (L_j = c)$ ). We have to show that  $q$  is a new predicate; i.e., it cannot be expressed as a function of the already-exposed predicates.

A property of the split invariant, which can be shown by induction, is that  $[\theta_j \Rightarrow (x_m \equiv m)]$  for each shared refinement variable  $x_m$  that is added for a predicate  $m$  exposed for process  $P_j$ . As  $u$  and  $v$  agree on all shared variables, including refinement variables, and as both satisfy  $\theta_j$ , it follows from the property that all prior predicates exposed for  $P_j$  have identical values on  $c$  and  $t(L_j)$ . As this is not true for  $q$ , it cannot be expressed as a function of the already exposed predicates.  $\square$

#### 4.4 Weak and Generalized Fairness

Weak Fairness, also called “justice”, has the normal form  $GF(p)$  (“infinitely often  $p$ ”). It is often used to express the constraint that a continuously enabled transition cannot be forever ignored; i.e.,  $FG(enabled) \Rightarrow GF(executed)$ . As its normal form is a special case of strong fairness, the algorithm developed for strong fairness can be applied to it. Thus, the common weak fairness specification  $\Phi \equiv (\bigwedge i : GF(p_i))$ , where  $p_i$  is an assertion over the variables of process  $P_i$ , is abstracted to  $[\Phi]^\theta \equiv (\bigwedge i : GF(\exists L_i : \theta_i \wedge p_i))$  for use in the FLL algorithm.

	Example	Property	N	JTLV		SPLIT		
				Nodes	Time	Ref.	Nodes	Time
1	BAKERY	$G(wait(1) \Rightarrow F(crit(1)))$	3	300k	0.3	2	1.2M	2.5
		– Valid –	4	11.6M	93	2	14.6M	52
2	MUXSEM	$G(wait(1) \Rightarrow F(crit(1)))$	5	58k	0.2	1	48k	0.3
		– Counter example –	10	21M	24	2	371k	1.1
			20	over 20 minutes	2	2.1M	9	

**Table 1.** Experimental results when assuming only unconditional fairness

Emerson and Lei consider a general fairness criterion in [13], which is a disjunction of strong fairness conditions. This can be handled by abstracting each disjunct separately and re-forming the disjunction.

For simplicity, the development of the algorithm considered fairness assertions  $(\bigwedge i : FG(p_i) \vee GF(q_i))$  where  $p_i$  and  $q_i$  are expressed in terms of the variables of process  $P_i$ . In a more general setting, these predicates may be expressed over the local state of more than one process. The analysis method extends easily, with each predicate being abstracted by quantifying out the relevant local variables. Thus, the general abstraction function is  $[p]^\theta \equiv (\exists L : (\bigwedge i : \theta_i) \wedge p)$ .

## 5 Experimental Results

We implemented our method as part of SPLIT [10] – a compositional LTL verifier, and tested it on several parameterized examples which require fairness assumptions. We also compared it with the LTL model checker implemented on top of JTLV [25], and with the model checker NUSMV [4]. The latter, however, is optimized for verifying synchronous systems and even after disabling the conjunctive partitioning the results obtained by it were considerably inferior to those obtained by JTLV and SPLIT. We therefore do not include in this paper the results obtained by NUSMV. The experiments were conducted on a Intel Core 2 Duo 2.4 GHz with 4 GB RAM running 64-bit Linux. Both SPLIT and JTLV were configured to use the CUDD BDD library. We set a timeout of 20 minutes for the experiments.

As mentioned in Subsection 4.1, unconditional fairness is sufficient to guarantee various properties in selected protocols. For example, in algorithm BAKERY [18] ensuring that a process that wishes to enter the critical section eventually does, i.e.,  $\forall i : G(wait(i) \Rightarrow F(crit(i)))$ , does not require to assume any weak or strong fairness conditions. For

	Example	Property	N	JTLV		SPLIT		
				Nodes	Time	Ref.	Nodes	Time
3	MUXSEM	$G(wait(1) \Rightarrow F(crit(1)))$	5	24k	0	1	61k	0.2
		– Valid –	10	1.2M	3.8	1	259k	0.7
			20	over 20 minutes	1	1.2M	3	

**Table 2.** Experimental results when assuming only strong fairness only over P[0]

$f$ : array  $[0..N]$  of boolean initially  $f = 1$

$$P[1] ::= \left[ \begin{array}{l} \text{loop forever do} \\ l_0: \text{Non-Critical} \\ l_1: \text{request } f[2] \\ l_2: \text{request } f[1] \\ l_3: \text{Critical} \\ l_4: \text{release } f[1] \\ l_5: \text{release } f[2] \end{array} \right] \quad \parallel_{i=2}^N \quad P[i] ::= \left[ \begin{array}{l} \text{loop forever do} \\ l_0: \text{Non-Critical} \\ l_1: \text{request } f[i] \\ l_2: \text{request } f[i \oplus_N 1] \\ l_3: \text{Critical} \\ l_4: \text{release } f[i \oplus_N 1] \\ l_5: \text{release } f[i] \end{array} \right]$$

Fig. 7. Program DINING-PHIL: the dining philosophers

other protocols, such as MUXSEM, the same property is not valid when assuming only unconditional fairness, and both JTLV and SPLIT generate valid counter examples when attempting to verify it. The results for checking the eventual access property of  $P[1]$  for the two protocols are provided in Table 1. Note that since the property should be over global variables, the location variable of  $P[1]$  was exposed to all processes. “N” is the number of processes, “Nodes” is the peak number of BDD nodes generated, “Time” is the runtime in seconds, and “Ref.” is number of refinements had to be executed by SPLIT. For both examples the run-times are better for SPLIT; for MUXSEM, where counter examples had to be constructed, they are better by several orders of magnitude. Both SPLIT and JTLV required more than 20 minutes for verifying BAKERY for  $N = 5$ .

Assuming the strong fairness  $\text{GF}(P[1].at\_loc_1 \wedge x) \Rightarrow \text{GF}(P[1].at\_loc_2)$  only for  $P[1]$  is sufficient to prove the correctness of  $\text{G}(wait(1) \Rightarrow \text{F}(crit(1)))$  for MUXSEM. Both model checkers indeed validated the property under this condition and the results are provided in Table 2; they are again in favor of our method by a few orders of magnitude.

Most interesting and challenging test cases with respect to fairness are those that require to assume weak or strong fairness conditions for *all* the processes. The first such example is DINING-PHIL (a simple solution to the dining philosophers problem using semaphores), presented in Fig. 7. The eventual access property is valid only when assuming that  $\text{GF}(P[i].at\_loc_1 \wedge f[i]) \Rightarrow \text{GF}(P[i].at\_loc_2)$  and  $\text{GF}(P[i].at\_loc_2 \wedge f[i \oplus_N 1]) \Rightarrow \text{GF}(P[i].at\_loc_3)$  for  $1 < i \leq N$  and assuming the symmetric con-

	Example	Property	N	JTLV		SPLIT		
				Nodes	Time	Ref.	Nodes	Time
4	DINING-PHIL	$\text{G}(wait(1) \Rightarrow \text{F}(crit(1)))$ – Valid –	8	3M	13	0	1.9M	4
			9	9.1M	63	0	4.1M	8.6
			10	25M	421	0	8.6M	18
5	COND-TERM	$\text{F}(termination(1))$ – Valid –	4	91k	0.4	3	389k	1
			6	537k	1.6	3	2.1M	6.7
			8	4M	10	3	19M	101
6	MUX-SEM-N-D	$\text{G}(wait(1) \Rightarrow \text{F}(crit(1)))$ – Valid –	8	262k	0.6	1	172k	0.5
			12	5.3M	32.6	1	393k	1
			16	over 20 minutes		1	720k	1.8

Table 3. Results for properties that require to assume general fairness over *all* processes

ditions for  $i = 1$ . Namely, for each philosopher, if she can (enabled) infinitely often pick the first fork and subsequently pick the second fork then she should eat Spaghetti infinitely often. Example 4 in Table 3 presents the run-time results for verifying this example, that are again in favor of our method.

$$\begin{array}{l}
 y: \text{integer where } y = M \\
 \parallel_{i=1}^N P[i] :: \left[ \begin{array}{l}
 l_0: \text{while } y \neq 0 \text{ do} \\
 \quad \left[ \begin{array}{l}
 l_2: y := y - 1 \\
 \text{or} \\
 l_3: y := y + 1
 \end{array} \right] \\
 : l_4
 \end{array} \right]
 \end{array}$$

**Fig. 8.** COND-TERM

The second example is proving termination of COND-TERM. The protocol is presented in Fig. 8. A process terminates only if the strong fairness condition  $\text{GF}(P[i].at.l_3) \Rightarrow \text{GF}(false)$  is assumed over *all* processes. This condition permits only computations where  $y$  is increased a finite number of times. We verified the termination of COND-TERM for  $M = 15$ . The results are provided as example 5 in Table 3. This time they are in favor of the monolithic model checking as SPLIT requires three refinements to prove the property.

$$\begin{array}{l}
 x: \text{boolean where } x = 1 \\
 \parallel_{i=1}^N P[i] :: \left[ \begin{array}{l}
 \text{loop forever do} \\
 \quad \left[ \begin{array}{l}
 l_0: \text{Non-Critical} \\
 l_1: \text{request } x \\
 l_2: \langle \text{Critical}; \text{await } (false) \text{ or skip} \rangle \\
 l_3: \text{release } x
 \end{array} \right]
 \end{array} \right]
 \end{array}$$

**Fig. 9.** MUX-SEM-N-D: mutual exclusion with a nondeterministic stay in critical section

The last example that requires to assume general fairness over all the processes is MUX-SEM-N-D, presented in figure Fig. 9. This example is a variation of MUXSEM that allows each of the processes to stay non-deterministically, possibly forever, in the critical section. Thus,  $\text{G}(wait(1) \Rightarrow \text{F}(crit(1)))$  is valid only when assuming  $\bigwedge i : \text{GF}(P[i].at.l_2) \Rightarrow \text{GF}(P[i].at.l_3)$ . Namely, for each process, if it can (enabled) infinitely often leave the critical section then it should leave it infinitely often. Example 6 in Table 3 presents the run-time results for verifying this example, that are again in favor of our method.

## 6 Deductive Compositional Proofs under Fairness

The LL method was derived in [9] from a proof rule for verifying linear-time properties expressed by a Büchi automaton for their negation. That proof rule has two parts: the

first part expresses that  $\theta$  is a split invariant, while the second part shows that a Büchi accepting state occurs only finitely often on any joint computation of the program and the automaton, using rank functions which are local to each process.

1. Find a vector of local assertions,  $\theta = (\theta_1, \dots, \theta_N)$ , which meets the split invariance conditions from Figure 1
2. Form a fairness assertion,  $\Xi$ , out of the abstract assertions in  $[\Phi]^\theta$  and the acceptance condition of the Büchi automaton for the negated property. For each  $i$ , instantiate the strong fairness proof rule of [24] for the synchronous composition of the automaton and the abstract process  $[P_i]^\theta$ , with the fairness assertion  $\Xi$  and specification  $G(true \Rightarrow Ffalse)$ .

**Fig. 10.** Local Proof Rule for LTL properties

This structure can be modified to accommodate fairness, as shown in Figure 10. The proof rule of [24] is used with the conclusion being *false*. Hence, a valid proof shows the absence of any joint computation which is fair and which represents an accepting Büchi automaton run. All assertions and rank functions are local by definition. Moreover, as shown in [24], one can generate these components by instrumenting the model checking algorithms used in FLL.

## 7 Conclusions and Future Work

The algorithm presented here enables fully automated and compositional verification of progress properties under fairness and is, we believe, the first algorithm to do so. It deals with the main difficulty, that of handling the global nature of fairness, by a process of refinement: the fairness assumption is initially weakened relative to a split invariant, and is then strengthened in subsequent iterations until a decisive result is obtained. The algorithm has a simple implementation. Experiments with several parameterized protocols are promising, showing a clear advantage for the compositional method over a non-compositional one.

One aspect that merits further exploration is the choice of counter-example trace for refinement; currently, the algorithm uses whichever trace is provided by the model checking procedure. It would help, for instance, if the trace generation is biased to generate a trace which satisfies as many **MUST** requirements as possible.

## References

1. R. Alur and T. Henzinger. Reactive modules. In *IEEE LICS*, 1996.
2. D. Angluin. Learning regular sets from queries and counterexamples. *Inf. Comput.*, 75(2):87–106, 1987.
3. K.M. Chandy and J. Misra. *Parallel Program Design: A Foundation*. Addison-Wesley, 1988.

4. A. Cimatti, E. M. Clarke, F. Giunchiglia, and M. Roveri. NUSMV: a new symbolic model verifier. In *CAV*, pages 495–499, 1999.
5. E. M. Clarke, E. A. Emerson, and A. P. Sistla. Automatic verification of finite-state concurrent systems using temporal logic. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 8(2), 1986.
6. E. M. Clarke and O. Grumberg. Avoiding the state explosion problem in temporal logic model checking. In *PODC*, pages 294–303, 1987.
7. J. M. Cobleigh, G. S. Avrunin, and L. A. Clarke. Breaking up is hard to do: an investigation of decomposition for assume-guarantee reasoning. In *ISSTA*, pages 97–108, 2006.
8. A. Cohen and K. S. Namjoshi. Local proofs for global safety properties. In *CAV*, volume 4590 of *LNCS*, pages 55–67. Springer, 2007.
9. A. Cohen and K. S. Namjoshi. Local proofs for linear-time properties of concurrent programs. In *CAV*, volume 5123 of *LNCS*, pages 149–161. Springer, 2008.
10. A. Cohen, K. S. Namjoshi, and Y. Sa’ar. SPLIT: a compositional LTL verifier. submitted to *CAV* 2010.
11. W-P. de Roever, F. de Boer, U. Hannemann, J. Hooman, Y. Lakhnech, M. Poel, and J. Zwiers. *Concurrency Verification: Introduction to Compositional and Noncompositional Proof Methods*. Cambridge University Press, 2001.
12. E.W. Dijkstra and C.S. Scholten. *Predicate Calculus and Program Semantics*. Springer Verlag, 1990.
13. E.A. Emerson and C-L. Lei. Efficient model checking in fragments of the propositional mu-calculus (extended abstract). In *LICS*, 1986.
14. E.A. Emerson and C-L. Lei. Modalities for model checking: Branching time logic strikes back. *Sci. of Comp. Programming*, 8(3), 1987.
15. A. Farzan, Y. Chen, E. M. Clarke, Y. Tsan, and B. Wang. Extending automated compositional verification to the full class of omega-regular languages. In *TACAS*, LNCS, 2008.
16. C. Flanagan and S. Qadeer. Thread-modular model checking. In *SPIN*, volume 2648 of *LNCS*, pages 213–224, 2003.
17. D. Giannakopoulou, C. S. Pasareanu, and H. Barringer. Assumption generation for software component verification. In *ASE*, pages 3–12, 2002.
18. L. Lamport. A new solution of dijkstra’s concurrent programming problem. *Commun. ACM*, 17(8):453–455, 1974.
19. L. Lamport. Proving the correctness of multiprocess programs. *IEEE Trans. Software Eng.*, 3(2), 1977.
20. K.L. McMillan. A compositional rule for hardware design refinement. In *CAV*, volume 1254 of *LNCS*, 1997.
21. K.S. Namjoshi. Symmetry and completeness in the analysis of parameterized systems. In *VMCAI*, volume 4349 of *LNCS*, 2007.
22. K.S. Namjoshi and R.J. Treffler. On the completeness of compositional reasoning. In *CAV*, volume 1855 of *LNCS*. Springer-Verlag, 2000.
23. S. S. Owicki and D. Gries. Verifying properties of parallel programs: An axiomatic approach. *Commun. ACM*, 19(5):279–285, 1976.
24. A. Pnueli and Y. Sa’ar. All you need is compassion. In *VMCAI*, volume 4905 of *Lecture Notes in Computer Science*, pages 233–247. Springer, 2008.
25. A. Pnueli, Y. Sa’ar, and L. D. Zuck. JTLV: A framework for developing verification algorithms. web: <http://jtlv.sourceforge.net/>.
26. O. Tkachuk, M. B. Dwyer, and C. S. Pasareanu. Automated environment generation for software model checking. In *ASE*, pages 116–129, 2003.