# Optimized L\*-Based Assume-Guarantee Reasoning⋆

Sagar Chaki[1] and Ofer Strichman[2]

[1] Software Engineering Institute, Pittsburgh, USA
chaki@sei.cmu.edu
[2] Information Systems Engineering, IE, Technion, Israel
ofers@ie.technion.ac.il

**Abstract.** In this paper, we suggest three optimizations to the L\*-based automated Assume-Guarantee reasoning algorithm for the compositional verification of concurrent systems. First, we use each counterexample from the model checker to supply multiple strings to L\*, saving candidate queries. Second, we observe that in existing instances of this paradigm, the learning algorithm is coupled weakly with the teacher. Thus, the learner ignores completely the details about the internal structure of the system and specification being verified, which are available already to the teacher. We suggest an optimization that uses this information in order to avoid many unnecessary – and expensive, since they involve model checking – membership and candidate queries. Finally, and most importantly, we develop a method for minimizing the alphabet used by the assumption, which reduces the size of the assumption and the number of queries required to construct it. We present these three optimizations in the context of verifying trace containment for concurrent systems composed of finite state machines. We have implemented our approach and experimented with real-life examples. Our results exhibit an average speedup of over 12 times due to the proposed improvements.

## 1 Introduction

Formal reasoning about concurrent programs is particularly hard due to the number of reachable states in the overall system. In particular, the number of such states can grow exponentially with each added component. Assume-Guarantee (AG) is a method for compositional reasoning that can be helpful in such cases. Consider a system with two components $M_1$ and $M_2$ that need to synchronize on a given set of shared actions, and a property $\varphi$ that the system should be verified against. In its simplest form, AG requires checking one of the components, say $M_1$, separately, while making some assumption on the behaviors permitted by $M_2$. The assumption should then be discharged when checking $M_2$ in order to conclude the conformance of the product machine with the property. This idea is formalized with the following AG rule:

---

$$A \times M_1 \preceq \varphi$$
$$\frac{M_2 \preceq A}{M_1 \times M_2 \preceq \varphi} \quad \text{(AG-NC)} \qquad\qquad (1)$$

where $\preceq$ stands for some conformance relation[1]. For trace containment, simulation and some other known relations, AG-NC is a sound and complete rule. In this paper, we consider the case in which $M_1, M_2$ and $\varphi$ are non-deterministic finite automata, and interpret $\preceq$ as the trace containment (i.e., language inclusion) relation.

Recently, Cobleigh et al. proposed [1] a completely automatic method for finding the assumption $A$, using Angluin's L* algorithm [2]. L* constructs a minimal Deterministic Finite Automaton (DFA) that accepts an unknown regular language $U$. L* interacts iteratively with a Minimally Adequate Teacher (MAT). In each iteration, L* queries the MAT about *membership* of strings in $U$ and whether the language of a specific *candidate* DFA is equal to $U$. The MAT is expected to supply a "Yes/No" answer to both types of questions. It is also expected to provide a counterexample along with a negative answer to a question of the latter type. L* then uses the counterexample to refine its candidate DFA while enlarging it by at least one state. L* is guaranteed to terminate within no more than $n$ iterations, where $n$ is the size of the minimal DFA accepting $U$.

In this paper we suggest three improvements to the automated AG procedure. The first improvement is based on the observation that counterexamples can sometimes be reused in the refinement process, which saves candidate queries.

The second improvement is based on the observation that the core L* algorithm is completely unaware of the internal details of $M_1, M_2$ and $\varphi$. With a simple analysis of these automata, most queries to the MAT can in fact be avoided. Indeed, we suggest to allow the core L* procedure access to the internal structure of $M_1, M_2$ and $\varphi$. This leads to a tighter coupling between the L* procedure and the MAT, and enables L* to make queries to the MAT in a more intelligent manner. Since each MAT query incurs an expensive model checking run, overall performance is improved considerably.

The last and most important improvement is based on the observation that the alphabet of the assumption $A$ is fixed conservatively to be the entire interface alphabet between $M_1$ and $\varphi$ on the one hand, and $M_2$ on the other. While the full interface alphabet is always sufficient, it is often possible to complete the verification successfully with a much smaller assumption alphabet. Since the overall complexity of the procedure depends on the alphabet size, a smaller alphabet can improve the overall performance. In other words, while L* guarantees the minimality of the learned assumption DFA with respect to a given alphabet, our improvement reduces the size of the alphabet itself, and hence also the size of the learned DFA. The technique we present is based on an automated abstraction/refinement procedure: we start with the empty alphabet and keep refining it based on an analysis of the counterexamples, using a pseudo-Boolean solver. The procedure is guaranteed to terminate with a minimal assumption alphabet that

---

[1] Clearly, for this rule to be effective, $A \times M_1$ must be easier to compute than $M_1 \times M_2$.

suffices to complete the overall verification. This technique effectively combines the two paradigms of automated AG reasoning and abstraction-refinement.

Although our optimizations are presented in the context of a non-circular AG rule, they are applicable for circular AG rules as well, although for lack of space we do not cover this topic in this paper. We implemented our approach in the COMFORT [3] reasoning framework and experimented with a set of benchmarks derived from real-life source code. The improvements reduce the overall number of queries to the MAT and the size of the learned automaton. While individual speedup factors exceeded 23, an average speedup of a factor of over 12 was observed, as reported in Section 6.

**Related Work.** The L* algorithm was developed originally by Angluin [2]. Most learning-based AG implementations, including ours, use a more sophisticated version of L* proposed by Rivest and Schapire [4]. Machine learning techniques have been used in several contexts related to verification [5,6,7,8,9]. The use of L* for AG reasoning was first proposed by Cobleigh et al. [1]. A symbolic version of this framework has also been developed by Alur et al. [10]. The use of learning for automated AG reasoning has also been investigated in the context of simulation checking [11] and deadlock detection [12]. The basic idea behind the automated AG reasoning paradigm is to learn an assumption [13], using L*, that satisfies the two premises of AG-NC. The AG paradigm was proposed originally by Pnueli [14] and has since been explored (in manual/semi-automated forms) widely. The third optimization we propose amounts to a form of counterexample-guided abstraction refinement (CEGAR). The core ideas behind CEGAR were proposed originally by Kurshan [15], and CEGAR has since been used successfully for automated hardware [16] and software [17] verification. An approach similar to our third optimization was proposed independently by Gheorghiu et. al [18]. However, they use polynomial (greedy) heuristics aimed at minimizing the alphabet size, whereas we find the optimal value, and hence we solve an NP-hard problem.

## 2      Preliminaries

Let $\lambda$ and $\cdot$ denote the empty string and the concatenation operator respectively. We use lower letters ($\alpha, \beta$, etc.) to denote actions, and higher letters ($\sigma, \pi$, etc.) to denote strings.

**Definition 1 (Finite Automaton).** *A finite automaton (FA) is a 5-tuple* $(S, Init, \Sigma, T, F)$ *where (i) $S$ is a finite set of states, (ii) $Init \subseteq S$ is the set of initial states, (iii) $\Sigma$ is a finite alphabet of actions, (iv) $T \subseteq S \times \Sigma \times S$ is the transition relation, and (v) $F \subseteq S$ is a set of accepting states.*

For any FA $M = (S, Init, \Sigma, T, F)$, we write $s \xrightarrow{\alpha} s'$ to mean $(s, \alpha, s') \in T$. Then the function $\delta$ is defined as follows: $\forall \alpha \in \Sigma . \forall s \in S . \delta(\alpha, s) = \{s' | s \xrightarrow{\alpha} s'\}$. We extend $\delta$ to operate on strings and sets of states in the natural manner. Thus, for any $\sigma \in \Sigma^*$ and $S' \subseteq S$, $\delta(\sigma, S')$ denotes the set of states of $M$ reached

by simulating $\sigma$ on $M$ starting from any $s \in S'$. The language accepted by $M$, denoted $\mathcal{L}(M)$, is defined as follows: $\mathcal{L}(M) = \{\sigma \in \Sigma^* \mid \delta(\sigma, Init) \cap F \neq \emptyset\}$.

**Determinism.** An FA $M = (S, Init, \Sigma, T, F)$ is said to be a deterministic FA, or DFA, if $|Init| = 1$ and $\forall \alpha \in \Sigma \textbf{.} \forall s \in S \textbf{.} |\delta(\alpha, s)| \leq 1$. Also, $M$ is said to be complete if $\forall \alpha \in \Sigma \textbf{.} \forall s \in S \textbf{.} |\delta(\alpha, s)| \geq 1$. Thus, for a complete DFA, we have the following: $\forall \alpha \in \Sigma \textbf{.} \forall s \in S \textbf{.} |\delta(\alpha, s)| = 1$. Unless otherwise mentioned, all DFA we consider in the rest of this paper are also complete. It is well-known that a language is regular iff it is accepted by some FA (or DFA, since FA and DFA have the same accepting power). Also, every regular language is accepted by a unique (up to isomorphism) minimal DFA.

**Complementation.** For any regular language $L$, over the alphabet $\Sigma$, we write $\overline{L}$ to mean the language $\Sigma^* - L$. If $L$ is regular, then so is $\overline{L}$. For any FA $M$ we write $\overline{M}$ to mean the (unique) minimal DFA accepting $\overline{\mathcal{L}(M)}$.

**Projection.** The projection of any string $\sigma$ over an alphabet $\Sigma$ is denoted by $\sigma{\downarrow}_\Sigma$ and defined inductively on the structure of $\sigma$ as follows: (i) $\lambda{\downarrow}_\Sigma = \lambda$, and (ii) $(\alpha \cdot \sigma'){\downarrow}_\Sigma = \alpha \cdot (\sigma'{\downarrow}_\Sigma)$ if $\alpha \in \Sigma$ and $\sigma'{\downarrow}_\Sigma$ otherwise. The projection of any regular language $L$ on an alphabet $\Sigma$ is defined as: $L{\downarrow}_\Sigma = \{\sigma{\downarrow}_\Sigma \mid \sigma \in L\}$. If $L$ is regular, so is $L{\downarrow}_\Sigma$. Finally, the projection $M{\downarrow}_\Sigma$ of any FA $M$ on an alphabet $\Sigma$ is the (unique) minimal DFA accepting the language $\mathcal{L}(M){\downarrow}_\Sigma$.

For the purpose of modeling systems with components that need to synchronize, it is convenient to distinguish between *local* and *global* actions. Specifically, local actions belong to the alphabet of a single component, while global actions are shared between multiple components. As defined formally below, components synchronize on global actions, and execute asynchronously on local actions.

**Definition 2 (Parallel Composition).** *Given two finite automata* $M_1 = (S_1, Init_1, \Sigma_1, T_1, F_1)$ *and* $M_2 = (S_2, Init_2, \Sigma_2, T_2, F_2)$, *their parallel composition* $M_1 \times M_2$ *is the FA* $(S_1 \times S_2, Init_1 \times Init_2, \Sigma_1 \cup \Sigma_2, T, F_1 \times F_2)$ *such that* $\forall s_1, s_1' \in S_1 \textbf{.} \forall s_2, s_2' \in S_2$, $(s_1, s_2) \xrightarrow{\alpha} (s_1', s_2')$ *iff for* $i \in \{1, 2\}$ *either* $\alpha \notin \Sigma_i \wedge s_i = s_i'$ *or* $s_i \xrightarrow{\alpha} s_i'$.

**Trace Containment.** For any FA $M_1$ and $M_2$, we write $M_1 \preceq M_2$ to mean $\mathcal{L}(M_1 \times \overline{M_2}) = \emptyset$. A counterexample to $M_1 \preceq M_2$ is a string $\sigma \in \mathcal{L}(M_1 \times \overline{M_2})$.

## 3   The L* Algorithm

The L* algorithm for learning DFAs was developed by Angluin [2] and later improved by Rivest and Schapire [4]. In essence, L* learns an unknown regular language $U$, over an alphabet $\Sigma$, by generating the minimal DFA that accepts $U$. In order to learn $U$, L* requires "Yes/No" answers to two types of queries:

1. *Membership query*: for a string $\sigma \in \Sigma^*$, 'is $\sigma \in U$ ?'
2. *Candidate query*: for a DFA $C$, 'is $\mathcal{L}(C) = U$ ?'

If the answer to a candidate query is "No", L* expects a counterexample string $\sigma$ such that $\sigma \in U - \mathcal{L}(C)$ or $\sigma \in \mathcal{L}(C) - U$. In the first case, we call $\sigma$ a *positive* counterexample, because it should be added to $\mathcal{L}(C)$. In the second case, we call $\sigma$ a *negative counterexample* since it should be removed from $\mathcal{L}(C)$. As mentioned before, L* uses the MAT to obtain answers to these queries.

**Observation Table.** L* builds an observation table $(S, E, T)$ where: (i) $S \subseteq \Sigma^*$ is the set of rows, (ii) $E \subseteq \Sigma^*$ is the set of columns (or experiments), and (iii) $T : (S \cup S \cdot \Sigma) \times E \to \{0, 1\}$ is a function defined as follows:

$$\forall s \in (S \cup S \cdot \Sigma). \forall e \in E. T(s, e) = \begin{cases} 1 & s \cdot e \in U \\ 0 & \text{otherwise} \end{cases} \tag{2}$$

**Consistency and Closure.** For any $s_1, s_2 \in (S \cup S \cdot \Sigma)$, $s_1$ and $s_2$ are equivalent (denoted as $s_1 \equiv s_2$) if $\forall e \in E. T(s_1, e) = T(s_2, e)$. A table is *consistent* if $\forall s_1, s_2 \in S. s_1 \neq s_2 \Rightarrow s_1 \not\equiv s_2$. L* always maintains a consistent table. In addition, a table is *closed* if $\forall s \in S. \forall \alpha \in \Sigma. \exists s' \in S. s' \equiv s \cdot \alpha$.

**Candidate Construction.** Given a closed and consistent table $(S, E, T)$, L* constructs a candidate DFA $C = (S, \{\lambda\}, \Sigma, \Delta, F)$ such that: (i) $F = \{s \in S \mid T(s, \lambda) = 1\}$, and (ii) $\Delta = \{(s, \alpha, s') \mid s' \equiv s \cdot \alpha\}$. Note that $C$ is deterministic and complete since $(S, E, T)$ is consistent and closed. Since a row corresponds to a state of $C$, we use the terms "row" and "candidate state" synonymously.

| | | $\lambda$ | $e_2$ | $e_3$ |
|---|---|---|---|---|
| $S$ | $\lambda$ | 0 | 1 | 0 |
| | $\alpha$ | 1 | 1 | 0 |
| $S \cdot \Sigma$ | $\beta$ | 0 | 1 | 0 |
| | $\alpha\alpha$ | 1 | 1 | 0 |
| | $\alpha\beta$ | 0 | 1 | 0 |



**Fig. 1.** An Observation Table and the Corresponding Candidate DFA

*Example 1.* Consider Figure 1. On the left is an observation table with the entries being the $T$ values. Let $\Sigma = \{\alpha, \beta\}$. From this table we see that $\{e_2, \alpha, \alpha \cdot e_2, \beta \cdot e_2, \alpha\alpha, \ldots\} \in U$. On the right is the corresponding candidate DFA. □

**L* Step-By-Step.** We now describe L* in more detail, using line numbers from its algorithmic description in Figure 2. This conventional version of L* is used currently in the context of automated AG reasoning. We also point out the specific issues that are addressed by the improvements we propose later on in this paper. Recall that $\lambda$ denotes the empty string. After the initialization at Line 1, the table has one cell corresponding to $(\lambda, \lambda)$. In the top-level loop, the table entries are first computed (at Line 2) using membership queries.

Next, L* closes the table by trying to find (at Line 3) for each $s \in S$, some *uncovered* action $\alpha \in \Sigma$ such that $\forall s' \in S . s' \not\equiv s \cdot \alpha$. If such an uncovered action $\alpha$ is found for some $s \in S$, L* adds $s \cdot \alpha$ to $S$ at Line 4 and continues with the closure process. Otherwise, it proceeds to the next Step. Note that each $\alpha \in \Sigma$ is considered when attempting to find an uncovered action.

(1)  let $S = E = \{\lambda\}$
     loop {
(2)     Update $T$ using queries
        while $(S, E, T)$ is not closed {
(3)        Find $(s, \alpha) \in S \times \Sigma$ such that $\forall s' \in S . s' \not\equiv s \cdot \alpha$
(4)        Add $s \cdot \alpha$ to $S$
        }
(5)     Construct candidate DFA $C$ from $(S, E, T)$
(6)     Make the conjecture $C$
(7)     if $C$ is correct return $C$
(8)     else Add $e \in \Sigma^*$ that witnesses the counterexample to E
     }

**Fig. 2.** The L* algorithm for learning an unknown regular language

Once the table is closed, L* constructs (at Line 5) a candidate DFA $C$ using the procedure described previously. Next, at Line 6, L* conjectures that $\mathcal{L}(C) = U$ via a candidate query. If the conjecture is wrong L* extracts from the counterexample $CE$ (returned by the MAT) a suffix $e$ that, when added to $E$, causes the table to cease being closed. The process of extracting the feedback $e$ has been presented elsewhere [4] and we do not describe it here. Once $e$ has been obtained, L* adds $e$ to $E$ and iterates the top-level loop by returning to line 2. Note that since the table is no longer closed, the subsequent process of closing it strictly increases the size of $S$. It can also be shown that the size of $S$ cannot exceed $n$, where $n$ is number of states of the minimal DFA accepting $U$. Therefore, the top-level loop of L* executes no more than $n$ times.

**Non-uniform Refinement.** It is interesting to note that the feedback from $CE$ does not refine the candidate in the abstraction/refinement sense: refinement here does not necessarily add/eliminate a positive/negative $CE$; this occurs eventually, but not necessarily in one step. Indeed, the first improvement we propose leverages this observation to reduce the number of candidate queries. It is also interesting to note that the refinement does not work in one direction: it may remove strings that are in $U$ or add strings that are not in $U$. The only guarantee that we have is that in each step at least one state is added to the candidate and that eventually L* learns $U$ itself.

**Complexity.** Overall, the number of membership queries made by L* is $\mathcal{O}(kn^2 + n \log m)$, where $k = |\Sigma|$ is the size of the alphabet of $U$, and $m$ is the length of the longest counterexample to a candidate query returned by the MAT [4]. The

dominating fragment of this complexity is $kn^2$ which varies directly with the size of $\Sigma$. As noted before, the $\Sigma$ used in the literature is sufficient, but often unnecessarily large. The third improvement we propose is aimed at reducing the number of membership queries by minimizing the size of $\Sigma$.

## 4    AG Reasoning with L*

In this section, we describe the key ideas behind the automated AG procedure proposed by Cobleigh et al. [1]. We begin with a fact that we use later on.

**Fact 1.** *For any FA $M_1$ and $M_2$ with alphabets $\Sigma_1$ and $\Sigma_2$, $\mathcal{L}(M_1 \times M_2) \neq \emptyset$ iff $\exists \sigma \in \mathcal{L}(M_1) . \sigma\!\downarrow_{(\Sigma_1 \cap \Sigma_2)} \in \mathcal{L}(M_2)\!\downarrow_{(\Sigma_1 \cap \Sigma_2)}$.*

Let us now restate AG-NC to reflect our implementation more accurately:

$$\frac{\begin{array}{c} A \times (M_1 \times \bar{\varphi}) \preceq \bot \\ M_2 \preceq A \end{array}}{(M_1 \times \bar{\varphi}) \times M_2 \preceq \bot} \tag{3}$$

where $\bot$ denotes a DFA accepting the empty language. The unknown language to be learned is

$$U = \mathcal{L}(\overline{(M_1 \times \bar{\varphi})\!\downarrow_\Sigma}) \tag{4}$$

over the alphabet $\Sigma = (\Sigma_1 \cup \Sigma_\varphi) \cap \Sigma_2$ where $\Sigma_1, \Sigma_2$ and $\Sigma_\varphi$ are the alphabets of $M_1, M_2$ and $\varphi$ respectively[2]. The choice of $U$ and $\Sigma$ is significant because, by Fact 1, the consequence of Eq. 3 does not hold iff the intersection between $\overline{U} = \mathcal{L}((M_1 \times \bar{\varphi})\!\downarrow_\Sigma)$ and $\mathcal{L}(M_2\!\downarrow_\Sigma)$ is non-empty. This situation is depicted in Fig. 3(a). Hence, if $A$ is the DFA computed by L* such that $\mathcal{L}(A) = U$, any counterexample to the second premise $M_2 \preceq A$ is guaranteed to be a real one. However, in practice, the process terminates after learning $U$ itself only in the worst case. As we shall see, it usually terminates earlier by finding either a counterexample to $M_1 \times M_2 \preceq \varphi$, or an assumption $A$ that satisfies the two premises of Eq. 3. This later case is depicted in Fig. 3(b).

**MAT Implementation.** The answer to a membership query with a string $\sigma$ is "Yes" iff $\sigma$ cannot be simulated on $M_1 \times \bar{\varphi}$ (see Eq. 4). A candidate query with some candidate $A$, on the other hand, is more complicated, and is described step-wise as follows (for brevity, we omit a diagram and refer the reader to the non-dashed portion of Figure 4):

**Step 1.** Use model checking to verify that $A$ satisfies the first premise of Eq. 3. If the verification of the first premise fails, obtain a counterexample trace $\pi \in \mathcal{L}(A \times M_1 \times \bar{\varphi})$ and proceed to Step 2. Otherwise, go to Step 3.

---

[2] Note that we do not compute $U$ directly because complementing $M_1$, a non-deterministic automaton, is typically intractable.
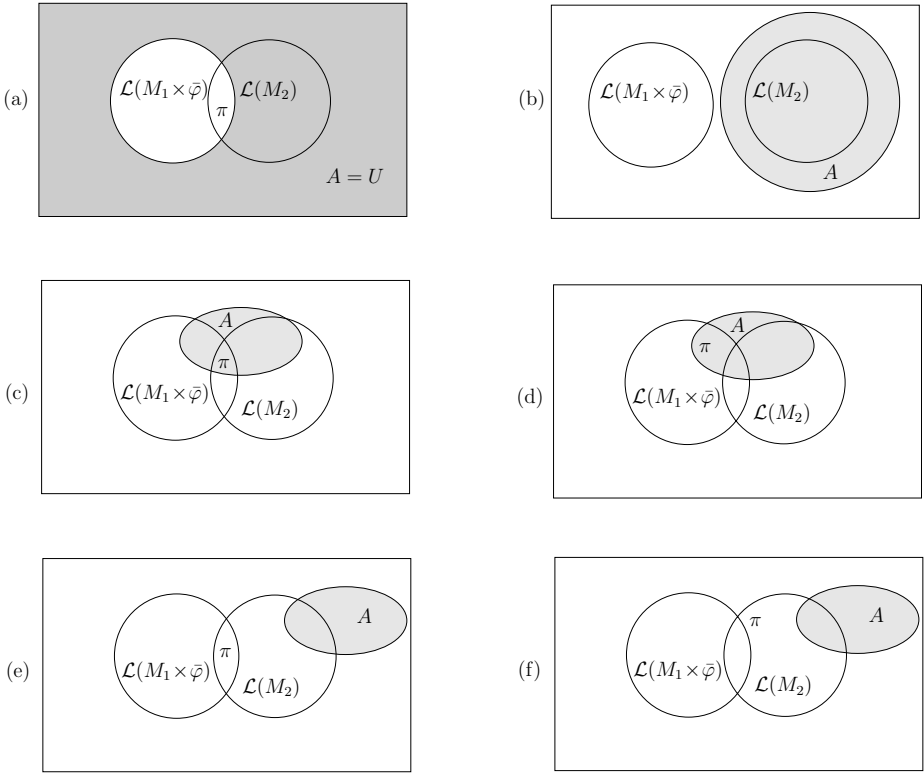
**Fig. 3.** Different L* scenarios. The gray area represents the candidate assumption $A$.

**Step 2.** Denote $\pi\!\downarrow_\Sigma$ by $\pi'$. Check via simulation if $\pi' \in \mathcal{L}(M_2\!\downarrow_\Sigma)$. If so, then by Fact 1, $\mathcal{L}(M_1\times\bar\varphi\times M_2) \neq \emptyset$ (i.e., $M_1 \times M_2 \not\preceq \varphi$) and the algorithm terminates. This situation is depicted in Fig. 3(c). Otherwise $\pi' \in \mathcal{L}(A) - U$ is a negative counterexample, as depicted in Fig. 3(d). Control is returned to L* with $\pi'$.

**Step 3.** At this point $A$ is known to satisfy the first premise. Proceed to model check the second premise. If $M_2 \preceq A$ holds as well, then by Eq. 3 conclude that $M_1 \times M_2 \preceq \varphi$ and terminate. This possibility was already shown in Fig. 3(b). Otherwise obtain a counterexample $\pi \in \mathcal{L}(M_2 \times \overline{A})$ and proceed to Step 4.

**Step 4.** Once again denote $\pi\!\downarrow_\Sigma$ by $\pi'$. Check if $\pi' \in \mathcal{L}((M_1\times\bar\varphi)\!\downarrow_\Sigma)$. If so, then by Fact 1, $\mathcal{L}(M_1\times\bar\varphi\times M_2) \neq \emptyset$ (i.e., $M_1 \times M_2 \not\preceq \varphi$) and the algorithm terminates. This scenario is depicted in Fig. 3(e). Otherwise $\pi' \in U - \mathcal{L}(A)$ is a positive counterexample, as depicted in Fig. 3(f) and we return to L* with $\pi'$.

Note that Steps 2 and 4 above are duals obtained by interchanging $M_1 \times \bar\varphi$ with $M_2$ and $U$ with $\mathcal{L}(A)$. Also, note that Fact 1 could be applied in Steps 2 and 4 above only because $\Sigma = (\Sigma_1 \cup \Sigma_\varphi) \cap \Sigma_2$. In the next section, we propose

an improvement that allows $\Sigma$ to be varied. Consequently, we also modify the procedure for answering candidate queries so that Fact 1 is used only in a valid manner.

## 5   Optimized L*-Based AG Reasoning

In this section we list three improvements to the algorithm described in Section 4. The first two improvements reduce the number of candidate and membership queries respectively. The third improvement is aimed at completing the verification process using an assumption alphabet that is smaller than $(\Sigma_1 \cup \Sigma_\varphi) \cap \Sigma_2$.

### 5.1   Reusing Counterexamples

Recall from Section 3 that every candidate query counterexample $\pi$ returned to L* is used to find a suffix that makes the table not closed, and hence adds at least one state (row) to the current candidate $C$ (observation table). Let $C'$ denote the new candidate constructed in the next iteration of the top-level loop (see Figure 2). We say that $C'$ is obtained by refining $C$ on $\pi$. However, the refinement process does not guarantee the addition/elimination of a positive/negative counterexample from $C'$. Thus, a negative counterexample $\pi \in L(C) - U$ may still be accepted by $C'$, and a positive counterexample $\pi \in U - L(C)$ may still be rejected by $C'$. This leads naturally to the idea of reusing counterexamples. Specifically, for every candidate $C'$ obtained by refining on a negative counterexample $\pi$, we check, via simulation, whether $\pi \in \mathcal{L}(C')$. If this is the case, we repeat the refinement process on $C'$ using $\pi$ instead of performing a candidate query with $C'$. The same idea is applied to positive counterexamples as well. Thus, if we find that $\pi \notin \mathcal{L}(C')$ for a positive counterexample $\pi$, then $\pi$ is used to further refine $C'$. This optimization reduces the number of candidate queries.

### 5.2   Selective Membership Queries

Recall the operation of closing the table (see Lines 3 and 4 of Figure 2) in L*. For every row $s$ added to $S$, L* must compute $T$ for every possible extension of $s$ by a single action. Thus L* must decide if $s \cdot \alpha \cdot e \in U$ for each $\alpha \in \Sigma$ and $e \in E$ — a total of $|\Sigma| \cdot |E|$ membership queries. To see how a membership query is answered, for any $\sigma \in \Sigma^*$, let $Sim(\sigma)$ be the set of states of $M_1 \times \bar\varphi$ reached by simulating $\sigma$ from an initial state of $M_1 \times \bar\varphi$ and by treating actions not in $\Sigma$ as $\epsilon$ (i.e., $\epsilon$-transitions are allowed where the actions are local to $M_1 \times \bar\varphi$). Then, $\sigma \in U$ iff $Sim(\sigma)$ does not contain an accepting state of $M_1 \times \bar\varphi$.

Let us return to the problem of deciding if $s \cdot \alpha \cdot e \in U$. Let $En(s) = \{\alpha' \in \Sigma \mid \delta(\alpha', Sim(s)) \neq \emptyset\}$ be the set of enabled actions from $Sim(s)$ in $M_1 \times \bar\varphi$. Now, for any $\alpha \notin En(s)$, $Sim(s \cdot \alpha \cdot e) = \emptyset$ and hence $s \cdot \alpha \cdot e$ is guaranteed to belong to $U$. This observation leads to our second improvement. Specifically, for every $s$ added to $S$, we first compute $En(s)$. Note that $En(s)$ is computed by simulating $s\!\downarrow_{\Sigma_1}$ on $M_1$ and $s\!\downarrow_{\Sigma_\varphi}$ on $\varphi$ separately, without composing $M_1$ and $\varphi$. We then make membership queries with $s \cdot \alpha \cdot e$, but only for $\alpha \in En(s)$. For all $\alpha \notin En(s)$ we directly set $T(s \cdot \alpha, e) = 1$ since we know that in this case

$s \cdot \alpha \cdot e \in U$. The motivation behind this optimization is that $En(s)$ is usually much smaller that $\Sigma$ for any $s$. The actual improvement in performance due to this tactic depends on the relative sizes of $En(s)$ and $\Sigma$ for the different $s \in S$.

### 5.3   Minimizing the Assumption Alphabet

As mentioned before, existing automated AG procedures use a constant assumption alphabet $\Sigma = (\Sigma_1 \cup \Sigma_\varphi) \cap \Sigma_2$. There may exist, however, an assumption $A$ over a smaller alphabet $\Sigma_c \subset \Sigma$ that satisfies the two premises of Eq. 3. Since Eq. 3 is sound, the existence of such an $A$ would still imply that $M_1 \times M_2 \preceq \varphi$. However, recall that the number of L* membership queries varies directly with the alphabet size. Therefore, the benefit, in the context of learning $A$, is that a smaller alphabet leads to fewer membership queries.

In this section, we propose an abstraction-refinement scheme for building an assumption over a minimal alphabet. During our experiments, this improvement led to a 6 times reduction in the size of the assumption alphabet. The main problem with changing $\Sigma$ is of course that AG-NC is no longer complete. Specifically, if $\Sigma_C \subset \Sigma$, then there might not exist any assumption $A$ over $\Sigma_C$ that satisfies the two premises of AG-NC even though the conclusion of AG-NC holds. The following theorem characterizes this phenomenon precisely.

**Theorem 1 (Incompleteness of AG-NC).** *Suppose there exists a string $\pi$ and an alphabet $\Sigma_C$ such that: (INC) $\pi{\downarrow}_{\Sigma_C} \in \mathcal{L}((M_1 \times \bar{\varphi}){\downarrow}_{\Sigma_C})$ and $\pi{\downarrow}_{\Sigma_C} \in \mathcal{L}(M_2{\downarrow}_{\Sigma_C})$. Then no assumption $A$ over $\Sigma_C$ satisfies the two premises of AG-NC.*

*Proof.* Suppose there exists a $\pi$ satisfying INC and an $A$ over $\Sigma_C$ satisfying the two premises of AG-NC. This leads to a contradiction as follows:

- *Case 1*: $\pi{\downarrow}_{\Sigma_C} \in \mathcal{L}(A)$. Since $A$ satisfies the first premise of AG-NC, we have $\pi{\downarrow}_{\Sigma_C} \notin \mathcal{L}((M_1 \times \bar{\varphi}){\downarrow}_{\Sigma_C})$, a contradiction with INC.
- *Case 2*: $\pi{\downarrow}_{\Sigma_C} \notin \mathcal{L}(A)$. Hence $\pi{\downarrow}_{\Sigma_C} \in \mathcal{L}(\overline{A})$. Since $A$ satisfies the second premise of AG-NC, we have $\pi{\downarrow}_{\Sigma_C} \notin \mathcal{L}(M_2{\downarrow}_{\Sigma_C})$, again contradicting INC.   □

We say that an alphabet $\Sigma_C$ is incomplete if $\Sigma_C \neq \Sigma$ and there exists a trace $\pi$ satisfying condition INC above. Therefore, whenever we come across a trace $\pi$ that satisfies INC, unless $\Sigma_C = \Sigma$, we know that the current $\Sigma_C$ is incomplete and must be refined. We now describe our overall procedure which incorporates testing $\Sigma_C$ for incompleteness and refining an incomplete $\Sigma_C$ appropriately.

**Detecting Incompleteness.** Our optimized automated AG procedure is depicted in Fig. 4. Initially $\Sigma_c = \emptyset^3$. Let us write $\pi'$ and $\pi''$ to mean $\pi{\downarrow}_{\Sigma_C}$ and $\pi{\downarrow}_{\Sigma}$ respectively. The process continues as in Section 4, until one of the following two scenarios occur while answering a candidate query:

- *Scenario 1*: We reach Step 2 with a trace $\pi \in \mathcal{L}(A \times M_1 \times \bar{\varphi})$. Note that this implies $\pi' \in \mathcal{L}((M_1 \times \bar{\varphi}){\downarrow}_{\Sigma_C})$. Now we first check if $\pi' \in \mathcal{L}(M_2{\downarrow}_{\Sigma_C})$. If not,

---

<sup>3</sup> We could also start with $\Sigma_c = \Sigma_\varphi$ since it is very unlikely that $\varphi$ can be proven or disproven without controlling the actions that define it.
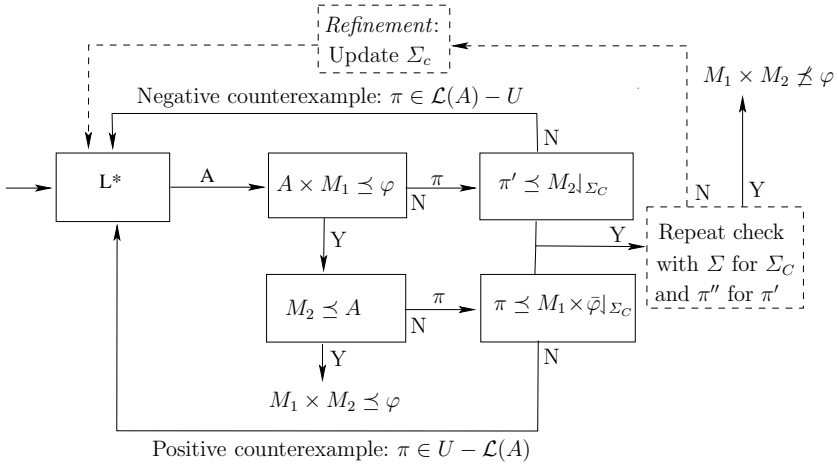
**Fig. 4.** Generalized AG with L*, with an abstraction-refinement loop (added with dashed lines) based on the assumption alphabet $\Sigma_c \subseteq \Sigma$. Strings $\pi'$ and $\pi''$ denote $\pi\!\downarrow_{\Sigma_C}$ and $\pi\!\downarrow_{\Sigma}$ respectively.

we return $\pi'$ as a negative counterexample to L* exactly as in Section 4. However, if $\pi' \in \mathcal{L}(M_2\!\downarrow_{\Sigma_C})$, then $\pi$ satisfies the condition INC of Theorem 1, and hence $\Sigma_C$ is incomplete. Instead of refining $\Sigma_C$ at this point, we first check if $\pi'' \in \mathcal{L}(M_2\!\downarrow_{\Sigma})$. If so, then as in Section 4, by a valid application of Fact 1, $M_1 \times M_2 \not\preceq \varphi$ and the algorithm terminates. Otherwise, if $\pi'' \notin \mathcal{L}(M_2\!\downarrow_{\Sigma})$, we refine $\Sigma_C$.

- *Scenario 2*: We reach Step 4 with $\pi \in \mathcal{L}(M_2 \times \overline{A})$. Note that this implies $\pi' \in \mathcal{L}(M_2\!\downarrow_{\Sigma_C})$. We first check if $\pi' \in \mathcal{L}((M_1 \times \bar\varphi)\!\downarrow_{\Sigma_C})$. If not, we return $\pi'$ as a positive counterexample to L* exactly as in Section 4. However, if $\pi' \in \mathcal{L}((M_1 \times \bar\varphi)\!\downarrow_{\Sigma_C})$, then $\pi$ satisfies INC, and hence by Theorem 1, $\Sigma_C$ is incomplete. Instead of refining $\Sigma_C$ at this point, we first check if $\pi'' \in \mathcal{L}((M_1 \times \bar\varphi)\!\downarrow_{\Sigma})$. If so, then as in Section 4, by a valid application of Fact 1, $M_1 \times M_2 \not\preceq \varphi$ and we terminate. Otherwise, if $\pi'' \notin \mathcal{L}((M_1 \times \bar\varphi)\!\downarrow_{\Sigma})$, we refine $\Sigma_C$.

Note that the checks involving $\pi''$ in the two scenarios above correspond to the concretization attempts in a standard CEGAR loop. Also, Scenarios 1 and 2 are duals (as in the case of Steps 2 and 4 in Section 4) obtained by interchanging $M_1 \times \bar\varphi$ with $M_2$ and $U$ with $\mathcal{L}(A)$. In essence, while solving a candidate query, an incomplete $\Sigma_C$ results in a trace (specifically, $\pi$ above) that satisfies INC and leads neither to an actual counterexample of $M_1 \times M_2 \preceq \varphi$, nor to a counterexample to the candidate query being solved. In accordance with the CEGAR terminology, we refer to such traces as *spurious* counterexamples and use them collectively to refine $\Sigma_C$ as described next. In the rest of this section, all counterexamples we mention are spurious unless otherwise specified.

**Refining the Assumption Alphabet.** A counterexample arising from Scenario 1 above is said to be negative. Otherwise, it arises from Scenario 2 and is said to be positive. Our description that follows unifies the treatment of these two types of counterexamples, with the help of a common notation for $M_1 \times \bar{\varphi}$ and $M_2$. Specifically, let

$$M^{(\pi)} = \begin{cases} M_1 \times \bar{\varphi} & \pi \text{ is positive} \\ M_2 & \pi \text{ is negative} \end{cases}$$

We say that an alphabet $\Sigma'$ eliminates a counterexample $\pi$, and denote this with $Elim(\pi, \Sigma')$, if $\pi|_{\Sigma'} \notin \mathcal{L}(M^{(\pi)}|_{\Sigma'})$. Therefore, any counterexample $\pi$ is eliminated if we choose $\Sigma_C$ such that $Elim(\pi, \Sigma_C)$ holds since $\pi$ no longer satisfies the condition INC. Our goal, however, is to find a minimal alphabet $\Sigma_C$ with this property. It turns out that finding such an alphabet is computationally hard.

**Theorem 2.** *Finding a minimal eliminating alphabet is NP-hard in $|\Sigma|$.*

*Proof.* The proof relies on a reduction from the minimal hitting set problem.

**Minimal Hitting Set.** A Minimal Hitting Set (MHS) problem is a pair $(U, T)$ where $U$ is a finite set and $T \subseteq 2^U$ is a finite set of subsets of $U$. A solution to $(U, T)$ is a minimal $X \subseteq U$ such that $\forall T' \in T \boldsymbol{.} X \cap T' \neq \emptyset$. It is well-known that MHS is NP-complete in $|U|$.

Now we reduce MHS to finding a minimal eliminating alphabet. Let $(U, T)$ be any MHS problem and let $<$ be a strict order imposed on the elements of $U$. Consider the following problem $\mathcal{P}$ of finding a minimal eliminating alphabet. First, let $\Sigma = U$. Next, for each $T' \in T$ we introduce a counterexample $\pi(T')$ obtained by arranging the elements of $U$ according to $<$, repeating each element of $T'$ twice and the remaining elements of $U$ just once. For example suppose $U = \{a, b, c, d, e\}$ such that $a < b < c < d < e$. Then for $T' = \{b, d, e\}$ we introduce the counterexample $\pi(T') = a \cdot b \cdot b \cdot c \cdot d \cdot d \cdot e \cdot e$. Also, for each counterexample $\pi(T')$ introduced, let $M(\pi(T'))$ accept a single string obtained by arranging the elements of $U$ according to $<$, repeating each element of $U$ just once. Thus, for the example $U$ above, $M(\pi(T'))$ accepts the single string $a \cdot b \cdot c \cdot d \cdot e$.

Let us first show the following result: for any $T' \in T$ and any $X \subseteq U$, $X \cap T' \neq \emptyset$ iff $Elim(\pi(T'), X)$. In other words, $X \cap T' \neq \emptyset$ iff $\pi(T')|_X \notin \mathcal{L}(M(\pi(T'))|_X)$. Indeed suppose that some $\alpha \in X \cap T'$. Then $\pi(T')|_X$ contains two consecutive occurrences of $\alpha$ and hence cannot be accepted by $M(\pi(T'))|_X$. By the converse implication, if $M(\pi(T'))|_X$ does not accept $\pi(T')|_X$, then $\pi(T')|_X$ must contain two consecutive occurrences of some action $\alpha$. But then $\alpha \in X \cap T'$ and hence $X \cap T' \neq \emptyset$. The above result implies immediately that any solution to the MHS problem $(U, T)$ is also a minimal eliminating alphabet for $\mathcal{P}$. Also, the reduction from $(U, T)$ to $\mathcal{P}$ described above can be performed using logarithmic space in $|U| + |T|$. Finally, $|\Sigma| = |U|$, which completes our proof. □

As we just proved, finding the minimal eliminating alphabet is NP-hard in $|\Sigma|$. Yet, since $|\Sigma|$ is relatively small, this problem can still be feasible in practice (as

our experiments have shown: see Section 6). We propose a solution based on a reduction to Pseudo-Boolean constraints. Pseudo-Boolean constraints have the same modeling power as 0-1 ILP, but solvers for this logic are typically based on adapting SAT engines for linear constraints over Boolean variables, and geared towards problems with relatively few linear constraints (and a linear objective function) and constraints in CNF.

**Optimal Refinement.** Let $\Pi$ be the set of all (positive and negative) counterexamples seen so far. We wish to find a minimal $\Sigma_C$ such that: $\forall \pi \in \Pi.Elim(\pi, \Sigma_C)$. To this end, we formulate and solve a Pseudo-Boolean constraint problem with an objective function stating that we seek a solution which minimizes the chosen set of actions. The set of constraints of the problem is $\Phi = \bigcup_{\pi \in \Pi} \Phi(\pi)$. In essence, if $M^{[\pi]}$ is the minimal DFA accepting $\{\pi\}$, then $\Phi(\pi)$ represents symbolically the states reachable in $M^{[\pi]} \times M^{(\pi)}$, taking into account all possible values of $\Sigma_C$. Henceforth, we continue to use square brackets when referring to elements of $M^{[\pi]}$, and regular parenthesis when referring to elements of $M^{(\pi)}$.

We now define $\Phi(\pi)$ formally. Let $M^{[\pi]} = (S^{[\pi]}, Init^{[\pi]}, \Sigma^{[\pi]}, T^{[\pi]}, F^{[\pi]})$ and $M^{(\pi)} = (S^{(\pi)}, Init^{(\pi)}, \Sigma^{(\pi)}, T^{(\pi)}, F^{(\pi)})$. Let $\delta^{[\pi]}$ and $\delta^{(\pi)}$ be the $\delta$ functions of $M^{[\pi]}$ and $M^{(\pi)}$ respectively. We define a state variable of the form $(s, t)$ for each $s \in S^{[\pi]}$ and $t \in S^{(\pi)}$. Intuitively, the variable $(s, t)$ indicates whether the product state $(s, t)$ is reachable in $M^{[\pi]} \times M^{(\pi)}$. We also define a choice variable $s(\alpha)$ for each action $\alpha \in \Sigma$, indicating whether $\alpha$ is selected to be included in $\Sigma_C$. Now, $\Phi(\pi)$ consists of the following clauses:

*Initialization and Acceptance*: Every initial and no accepting state is reachable:

$$\forall s \in Init^{[\pi]}.\forall t \in Init^{(\pi)}.(s,t) \qquad \forall s \in F^{[\pi]}.\forall t \in F^{(\pi)}.\neg(s,t)$$

*Shared Actions*: Successors depend on whether an action is selected or not:

$$\forall \alpha \in \Sigma.\forall s \in S^{[\pi]}.\forall s' \in \delta^{[\pi]}(\alpha, s).\forall t \in S^{(\pi)}.\forall t' \in \delta^{(\pi)}(\alpha, t).(s,t) \Rightarrow (s',t')$$
$$\forall \alpha \in \Sigma.\forall s \in S^{[\pi]}.\forall s' \in \delta^{[\pi]}(\alpha, s).\forall t \in S^{(\pi)}.\neg s(\alpha) \wedge (s,t) \Rightarrow (s',t)$$
$$\forall \alpha \in \Sigma.\forall s \in S^{[\pi]}.\forall t \in S^{(\pi)}.\forall t' \in \delta^{(\pi)}(\alpha, t).\neg s(\alpha) \wedge (s,t) \Rightarrow (s,t')$$

*Local Actions*: Asynchronous interleaving:

$$\forall \alpha \in \Sigma^{[\pi]} - \Sigma.\forall s \in S^{[\pi]}.\forall s' \in \delta^{[\pi]}(\alpha, s).\forall t \in S^{(\pi)}.(s,t) \Rightarrow (s',t)$$
$$\forall \alpha \in \Sigma^{(\pi)} - \Sigma.\forall s \in S^{[\pi]}.\forall t \in S^{(\pi)}.\forall t' \in \delta^{(\pi)}(\alpha, t).(s,t) \Rightarrow (s,t')$$

As mentioned before, the global set of constraints $\Phi$ is obtained by collecting together the constraints in each $\Phi(\pi)$. Observe that any solution $\nu$ to $\Phi$ has the following property. Let $\Sigma_C = \{\alpha \mid \nu(s(\alpha)) = 1\}$. Then we have $\forall \pi \in \Pi.\mathcal{L}((M^{[\pi]}\!\downarrow_{\Sigma_C}) \times (M^{(\pi)}\!\downarrow_{\Sigma_C})) = \emptyset$. But since $\mathcal{L}(M^{[\pi]}) = \{\pi\}$, the above statement is equivalent to $\forall \pi \in \Pi.(\pi\!\downarrow_{\Sigma_C}) \notin \mathcal{L}(M^{(\pi)}\!\downarrow_{\Sigma_C})$, which is further equivalent to $\forall \pi \in \Pi.Elim(\pi, \Sigma_C)$. Thus, $\Sigma_C$ eliminates all counterexamples. Finally, since we want the minimal such $\Sigma_C$, we minimize the number of chosen actions via the following objective function: $\min \sum_{\alpha \in \Sigma} s(\alpha)$.

**Fig. 5.** A positive counterexample $\pi$ and $M^{(\pi)} = M_1 \times \bar{\varphi}$

*Example 2.* Consider Fig. 5, in which there is one counterexample $\pi$, and an FA $M^{(\pi)} = M_1 \times \bar{\varphi}$ on which $\pi$ can be simulated if $\Sigma_c = \emptyset$. The state variables are $(s_i, t_j)$ for $i, j \in [0..2]$ and the choice variables are $s(\alpha), s(\beta)$. The constraints are:

$$
\begin{aligned}
Initialization : & \; (s_0, t_0) \\
SharedActions : & \; (s_0, t_1) \rightarrow (s_1, t_2) \\
& (s_0, t_0) \wedge \neg s(\alpha) \rightarrow (s_1, t_0) \\
& (s_0, t_1) \wedge \neg s(\alpha) \rightarrow (s_1, t_1) \\
& (s_0, t_2) \wedge \neg s(\alpha) \rightarrow (s_1, t_2) \\
& (s_0, t_0) \wedge \neg s(\beta) \rightarrow (s_0, t_1) \\
& (s_1, t_0) \wedge \neg s(\beta) \rightarrow (s_1, t_1) \\
& (s_2, t_0) \wedge \neg s(\beta) \rightarrow (s_2, t_1)
\end{aligned}
\qquad
\begin{aligned}
Acceptance : & \; \neg(s_2, t_2) \\
& (s_1, t_0) \rightarrow (s_2, t_1) \\
& (s_1, t_0) \wedge \neg s(\beta) \rightarrow (s_2, t_0) \\
& (s_1, t_1) \wedge \neg s(\beta) \rightarrow (s_2, t_1) \\
& (s_1, t_2) \wedge \neg s(\beta) \rightarrow (s_2, t_2) \\
& (s_0, t_1) \wedge \neg s(\alpha) \rightarrow (s_0, t_2) \\
& (s_1, t_1) \wedge \neg s(\alpha) \rightarrow (s_1, t_2) \\
& (s_2, t_1) \wedge \neg s(\alpha) \rightarrow (s_2, t_2)
\end{aligned}
$$

Since there are no local actions, these are all the constraints. The objective is to minimize $s(\alpha) + s(\beta)$. The optimal solution is $s(\alpha) = s(\beta) = 1$, corresponding to the fact that both actions need to be in $\Sigma_C$ in order to eliminate $\pi$.  □

## 6  Experiments

We implemented our technique in COMFORT and experimented with a set of benchmarks derived from real-life source code. All our experiments were carried out on quad 2.4 GHz machine with 4 GB RAM running RedHat Linux 9. We used PBS version $2.1^4$ to solve the Pseudo-Boolean constraints. The benchmarks were derived from the source code of OpenSSL version 0.9.6c. Specifically, we used the code that implements the handshake between a client and a server at the beginning of an SSL session. We designed a suite of 10 examples, each aiming a specific property (involving a sequence of message-passing events) that a correct handshake should exhibit. For instance, the first example (SSL-1) was aimed at verifying that a handshake is always initiated by a client and never by a server.

The experiments were aimed at evaluating our proposed improvements separately, and in conjunction with each other in the context of AG-NC. The results are described in Figure 6. The columns labeled MemQ and CandQ contain the total number of membership and candidate queries respectively. The columns labeled with $T_i$ and $\neg T_i$ contain results with/without the $i^{th}$ improvement for

---

[4] http://www.eecs.umich.edu/~faloul/Tools/pbs

| Name | CandQ | | MemQ | | Alph | | Time $\neg T_1$ | | | | Time $T_1$ | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | | | | | | $\neg T_2$ | | $T_2$ | | $\neg T_2$ | | $T_2$ | |
| | $\neg T_1$ | $T_1$ | $\neg T_2$ | $T_2$ | $\neg T_3$ | $T_3$ | $\neg T_3$ | $T_3$ | $\neg T_3$ | $T_3$ | $\neg T_3$ | $T_3$ | $\neg T_3$ | $T_3$ |
| SSL-1 | 2.2 | **2.0** | 37.5 | **4.5** | 12 | **1** | 25.4 | 19.7 | 12.3 | 20.0 | 23.8 | 20.1 | **10.5** | 20.5 |
| SSL-2 | **5.0** | 5.2 | 101.5 | **11.5** | 12 | **4** | 31.5 | 40.0 | **12.6** | 30.0 | 32.4 | 44.6 | 13.7 | 30.2 |
| SSL-3 | 8.5 | **7.5** | 163.0 | **28.0** | 12 | **4** | 43.8 | 49.1 | **14.5** | 35.3 | 45.6 | 48.9 | 15.6 | 35.5 |
| SSL-4 | 13.0 | **10.5** | 248.0 | **56.5** | 12 | **4** | 63.0 | 67.5 | **17.4** | 58.1 | 61.5 | 67.7 | 18.6 | 48.4 |
| SSL-5 | 3.2 | **3.0** | 73.0 | **9.5** | 12 | **1** | 33.8 | 22.3 | **13.6** | 24.1 | 36.2 | 22.2 | 13.8 | 22.2 |
| SSL-6 | **6.8** | 7.2 | 252.0 | **36.5** | 12 | **2** | 102.8 | 30.6 | 24.2 | 29.0 | 102.2 | 43.3 | **23.1** | 29.8 |
| SSL-7 | 9.8 | **8.0** | 328.8 | **52.5** | 12 | **2** | 139.9 | 44.4 | **27.8** | 43.9 | 138.2 | 38.6 | 28.2 | 40.6 |
| SSL-8 | 15.0 | **13.0** | 443.0 | **77.5** | 12 | **3** | 183.3 | 73.6 | 37.1 | 67.9 | 184.0 | 73.2 | **35.8** | 64.2 |
| SSL-9 | 23.5 | **18.2** | 568.0 | **109.5** | 12 | **3** | 234.1 | 120.5 | 44.1 | 133.7 | 236.2 | 133.4 | **41.0** | 109.3 |
| SSL-10 | 25.5 | **22.0** | 689.5 | **128.5** | 12 | **3** | 293.9 | 188.6 | 48.4 | 168.1 | 297.0 | 179.9 | **45.9** | 169.7 |
| Avg. | 10.8 | **9.2** | 290.0 | **51.0** | 12 | **2** | 115.1 | 65.6 | 25.2 | 61.0 | 115.7 | 67.2 | **24.6** | 57.1 |

**Fig. 6.** Experimental Results for Non-Circular Rule AG-NC

$i \in \{1, 2, 3\}$. The row labeled "Avg." contains the arithmetic mean for the rest of the column. Best figures are highlighted. Note that entries under MemQ and CandQ are fractional since they represent the average over the four possible values of the remaining two improvements. Specifically, these are improvements 2 and 3 for CandQ, and improvements 1 and 3 for MemQ.

We observe that the improvements lead to the expected results in terms of reducing the number of queries and the size of assumption alphabets. The second and third improvements also lead to significant reductions in overall verification time, by a factor of over 12 on an average. Finally, even though the first improvement entails fewer candidate queries, it is practically ineffective for reducing overall verification time.

## References

1. Cobleigh, J.M., Giannakopoulou, D., Păsăreanu, C.S.: Learning assumptions for compositional verification. In: Proc. of TACAS. (2003)
2. Angluin, D.: Learning Regular Sets from Queries and Counterexamples. Information and Computation (2) (1987)
3. Chaki, S., Ivers, J., Sharygina, N., Wallnau, K.: The ComFoRT Reasoning Framework. In: Proc. of CAV. (2005)
4. Rivest, R.L., Schapire, R.E.: Inference of Finite Automata Using Homing Sequences. Information and Computation (2) (1993)
5. Peled, D., Vardi, M., Yannakakis, M.: Black box checking. In: Proc. of FORTE. (1999)
6. Groce, A., Peled, D., Yannakakis, M.: Adaptive Model Checking. In: Proc. of TACAS. (2002)
7. Alur, R., Cerny, P., Gupta, G., Madhusudan, P., Nam, W., Srivastava, A.: Synthesis of Interface Specifications for Java Classes. In: POPL. (2005)
8. Habermehl, P., Vojnar, T.: Regular model checking using inference of regular languages. In: Proc. of INFINITY. (2005)

9. Ernst, M., Cockrell, J., Griswold, W., Notkin, D.: Dynamically Discovering Likely Program Invariants to Support Program Evolution. In: Proc. of ICSE. (1999)
10. Alur, R., Madhusudan, P., Nam, W.: Symbolic compositional verification by learning assumptions. In: Proc. of CAV. (2005)
11. Chaki, S., Clarke, E., Sinha, N., Thati, P.: Automated Assume-Guarantee Reasoning for Simulation Conformance. In: Proc. of CAV. (2005)
12. Chaki, S., Sinha, N.: Assume-guarantee reasoning for deadlock. In: Proc. of FMCAD. (2006)
13. Giannakopoulou, D., Păsăreanu, C., Barringer, H.: Assumption Generation for Software Component Verification. In: Proc. of ASE. (2002)
14. Pnueli, A.: In Transition from Global to Modular Temporal Reasoning About Programs. Logics and Models of Concurrent Systems (1985)
15. Kurshan, R.: Computer-Aided Verification of Coordinating Processes: The Automata-Theoretic Approach. Princeton University Press (1994)
16. Clarke, E., Grumberg, O., Jha, S., Lu, Y., Veith, H.: Counterexample-Guided Abstraction Refinement for Symbolic Model Checking. Journal of the ACM (JACM) (5) (2003)
17. Ball, T., Rajamani, S.: Generating Abstract Explanations of Spurious Counterexamples in C Programs. Technical Report MSR-TR-2002-09, Microsoft (2002)
18. Gheorghiu, M., Giannakopoulou, D., Păsăreanu, C.: Refining Interface Alphabets for Compositional Verification. In: Proc. of TACAS. (2007)