

Assumption Generation for Software Component Verification

Dimitra Giannakopoulou Corina S. Păsăreanu
RIACS/USRA Kestrel Technologies LLC
NASA Ames Research Center
Moffett Field, CA 94035-1000, USA
{dimitra, pcorina}@email.arc.nasa.gov

Howard Barringer*
Department of Computer Science
University of Manchester
Oxford Road, Manchester, M13 9PL, UK
howard@cs.man.ac.uk

Abstract

Model checking is an automated technique that can be used to determine whether a system satisfies certain required properties. The typical approach to verifying properties of software components is to check them for all possible environments. In reality, however, a component is only required to satisfy properties in specific environments. Unless these environments are formally characterized and used during verification (assume-guarantee paradigm), the results returned by verification can be overly pessimistic. This work defines a framework that brings a new dimension to model checking of software components. When checking a component against a property, our model checking algorithms return one of the following three results: the component satisfies a property for any environment; the component violates the property for any environment; or finally, our algorithms generate an assumption that characterizes exactly those environments in which the component satisfies its required property. Our approach has been implemented in the LTSA tool and has been applied to the analysis of a NASA application.

1. Introduction

Our work is motivated by an ongoing project at NASA Ames Research Center on the verification of autonomous software. Autonomous software involves complex concurrent behaviours for reacting to external stimuli without human intervention. Extensive verification is a pre-requisite for the deployment of missions that involve autonomy.

Model checking is an automated verification technique that can be used to determine whether a concurrent system satisfies certain properties by exhaustively exploring all its possible executions. Software model checking is typically

applied to *components* of a larger system for several reasons. For example: a software component may be embedded as is the case for autonomous software; one would typically ignore the details of the operating system in which a component operates; a system may be partially specified; finally, given the fact that the state explosion problem [9] is particularly acute in software systems, one realistically needs to “divide and conquer”, that is, to break up the verification task in smaller tasks.

In order to model check a component in isolation one needs to incorporate a model of the environment interacting with the component. By default, this is the “most general environment”, an environment that can invoke, in any order, any action of the interface between the two, or that may refuse any service that the component requires. We believe that the above approach to component checking is overly *pessimistic*; the underlying assumption is that the environment is free to behave as it pleases, and that the component will satisfy the required property for any environment. A similar observation is made by De Alfaro and Henzinger in the context of interface compatibility checking [10, 11].

In the world of model checking, this problem has given rise to the assume-guarantee style of reasoning [36], where the model of the environment is restricted by assumptions provided by the developer. This style of reasoning is typically performed in an interactive fashion. Developers first check a component with the most general environment. If a counterexample is returned that is unrealistic for the system under analysis, they make several attempts at defining an assumption that is strong enough to eliminate false violations, but that also reflects appropriately the remaining system.

In this paper we propose and describe a novel framework for model checking of components that provides more useful user feedback than the usual counter-example generation for property violations. When model checking a component against a property, our algorithms return one of the following three results: (i) the component satisfies the property for any environment; (ii) the component violates the property for any environment; or finally, (iii) an automatically

*This author is most grateful for the partial support received from RIACS/USRA to undertake this research whilst on leave at NASA Ames Research Center.

generated assumption that characterizes exactly those environments in which the component satisfies the property.

Let us illustrate this with a small example. A multi-threaded component uses a mutex to coordinate accesses to a shared variable, which may also be accessed by the environment. The requirement is that race violations should not occur in the system. If some thread within the component performs unprotected accesses to the variable, the requirement may be violated irrespective of the environment. Our approach reports this fact, together with a counterexample illustrating it. Now assume that all accesses to the variable within the component are protected by the mutex. Model checking under the most general environment would return a violation. Our algorithms would return an assumption, reflecting the fact that all accesses to the shared variable by the environment must be protected by the lock.

In fact, our approach generates the *weakest* environment assumption that enables the property to hold. Therefore, in selecting an appropriate environment for a component, one can safely reject any environment that does not satisfy the assumption generated. Assumption generation may also be seen as a way of providing extra automated support for assume-guarantee reasoning. Finally, for systems like the ones we study, the environment is often unpredictable. Some assumptions are typically made about it, but loss of mission must be avoided even if the environment falls outside these assumptions. For such cases, assumptions can be used as runtime monitors of the actual environment [20]. Monitors can generate appropriate warnings when the environment falls outside expected behaviour and trigger special system behaviour, if necessary.

We have implemented our approach in the Labelled Transition Systems Analyzer (LTSA) tool [31, 30], which provides good support for incremental system design and verification. It implements such features as component abstraction and minimization that make the integration of our approach straightforward.

The problem of assumption generation can be associated with such problems as submodule construction, controller synthesis and model matching. To our knowledge, such work has not been directly applied to model checking before; the relation of our approach with these domains is further discussed in Section 6. The remainder of the paper is organized as follows. Section 2 briefly discusses the LTSA tool and the theory that underlies our approach. It is followed by the presentation of our approach in Section 3. Section 4 describes our experience with analyzing the Executive module of an experimental Mars Rover developed at NASA Ames. We discuss the applicability of our approach in practice and extensions that we are considering in Section 5. Finally, Section 6 presents related work, and Section 7 concludes the paper.

2. Background

In this section, we describe the LTSA framework in which our approach has been introduced. We provide formal definitions for those aspects of the tool that we have used and/or modified.

2.1. The LTSA Tool

The LTSA [31, 30] is an automated tool that supports Compositional Reachability Analysis (CRA) of a software system based on its architecture. In general, the software architecture of a concurrent system has a hierarchical structure [29]. CRA incrementally computes and abstracts the behaviour of composite components based on the behaviour of their immediate children in the hierarchy. Abstraction consists of hiding the actions that do not belong to the interface of a component, and minimizing with respect to observational equivalence [17].

The input language “FSP” of the tool is a process-algebra style notation with Labelled Transition Systems (LTS) semantics. A property is also expressed as an LTS, but with extended semantics, and is treated as an ordinary component during composition. Properties are combined with the components to which they refer. They do not interfere with system behaviour, unless they are violated. In the presence of violations, the properties introduced may reduce the state space of the (sub)systems analyzed.

As in our approach, the LTSA framework treats components as open systems that may only satisfy some requirements in specific contexts. By composing components with their properties, it postpones analysis until the system is closed, meaning that all contextual behaviour that is applicable has been provided. We extend this framework by performing useful analysis at the component level.


The LTSA tool also features graphical display of LTSs, interactive simulation and graphical animation of behaviour models [32], all helpful aids in both design and verification of system models.

2.2. Program Model

We use labelled transition systems (LTSs) to model the behaviour of communicating components in a concurrent system. Let Act be the universal set of observable actions, and $Act_\tau = Act \cup \{\tau\}$, where τ denotes a local action *unobservable* to a component’s environment. We use π to denote a special error state, which models the fact that a safety violation has occurred in the associated system.

A *labelled transition system* T is a quadruple $\langle S, \alpha T, R, s_0 \rangle$, where S is a set of states, $\alpha T \subseteq Act$ is a set of actions called the *alphabet* of T , $R \subseteq S \times \alpha T \cup \{\tau\} \times S$

Mutex:



Writer:




Figure 1. LTSs for a Mutex and a Writer




Figure 2. Mutual exclusion property

is a transition relation and $s_0 \in S$ is the initial state. We use Π to denote the LTS $\langle \{\pi\}, Act, \emptyset, \pi \rangle$.

For example, Figure 1 illustrates LTSs for a *Writer* component and a *Mutex*. In all illustrations of LTSs in this paper, state 0 is the initial state. The *Writer* acquires the mutex (action $W.acquire$), enters and subsequently exists a critical section ($W.enterCS$, $W.exitCS$) used to model the fact that the *Writer* updates some shared variable, and then releases the mutex $W.release$. The *Mutex* component can be acquired and released by the *Writer* ($W.acquire$, $W.release$) or its environment ($E.acquire$, $E.release$), but a single component can hold it at any one time.

We call an LTS well formed if the error state π has no outgoing transitions. We only consider well formed LTSs in this work. An LTS $T = \langle S, \alpha T, R, s_0 \rangle$ is *non-deterministic* if $\exists (s, a, s'), (s, a, s'') \in R$ such that $s' \neq s''$ (otherwise T is *deterministic*).

A *trace* σ of an LTS T is a sequence of observable actions that T can perform starting at its initial state. For example, $\langle W.acquire \rangle$ and $\langle W.acquire, W.enterCS, W.exitCS \rangle$ are both traces of the *Writer* component of Figure 1. The set of traces of T is denoted as $Tr(T)$. For $\mathcal{A} \subseteq Act$, we use $\sigma \upharpoonright \mathcal{A}$ to denote the trace obtained by removing from σ all occurrences of actions $a \notin \mathcal{A}$. We denote as $errTr(T)$ the set of traces that lead to the error state π , which we call *error traces* of T .

Operators

In the following, we provide semantics for the operators defined on LTSs that are used in our work. Although we provide transitional semantics in a typical process algebra style, our aim here is not to define an algebra.

Let $T = \langle S, \alpha T, R, s_0 \rangle$ and $T' = \langle S', \alpha T', R', s'_0 \rangle$. We say that T *transits* into T' with action a , denoted $T \xrightarrow{a} T'$, iff $(s_0, a, s'_0) \in R$ and: either $\alpha T = \alpha T'$ and $R = R'$ for $s'_0 \neq \pi$, or, in the special case where $s'_0 = \pi$, $T' = \Pi$.

The *interface* operator \uparrow is used to make unobservable those actions in the LTS of a component that are not part of its interface. Formally, given an LTS T and a set of observable actions $\mathcal{A} \subseteq Act$, $T \uparrow \mathcal{A}$ is defined as follows. For $T = \Pi$, $T \uparrow \mathcal{A} = \Pi$. For $T \neq \Pi$, $T \uparrow \mathcal{A}$ is an LTS with the same set of states and initial state as T . The alphabet of $T \uparrow \mathcal{A}$ is $\alpha T \cap \mathcal{A}$, and its transition relation is described by the following rules:

$$\frac{T \xrightarrow{a} T', a \in \mathcal{A}}{T \uparrow \mathcal{A} \xrightarrow{a} T' \uparrow \mathcal{A}} \quad \frac{T \xrightarrow{a} T', a \notin \mathcal{A}}{T \uparrow \mathcal{A} \xrightarrow{\tau} T' \uparrow \mathcal{A}}$$

Parallel composition “ \parallel ” is a *commutative* and *associative* operator that combines the behaviour of two components by synchronization of the actions common to their alphabets and interleaving of the remaining actions. For example, in computing the parallel composition of components *Writer* and *Mutex* of Figure 1, actions $W.acquire$ and $W.release$ will each be synchronized.

Formally, let $T_1 = \langle S^1, \alpha T_1, R^1, s_0^1 \rangle$ and $T_2 = \langle S^2, \alpha T_2, R^2, s_0^2 \rangle$ be two LTSs. If either $T_1 = \Pi$ or $T_2 = \Pi$, then $T_1 \parallel T_2 = \Pi$. Otherwise, $T_1 \parallel T_2$ is an LTS $T = \langle S, \alpha T, R, s_0 \rangle$, where $S = S^1 \times S^2$, $s_0 = (s_0^1, s_0^2)$, $\alpha T = \alpha T_1 \cup \alpha T_2$, and R is defined as follows:

$$\frac{T_1 \xrightarrow{a} T_1', a \notin \alpha T_2}{T_1 \parallel T_2 \xrightarrow{a} T_1' \parallel T_2} \quad \frac{T_1 \xrightarrow{a} T_1', T_2 \xrightarrow{a} T_2', a \neq \tau}{T_1 \parallel T_2 \xrightarrow{a} T_1' \parallel T_2'}$$

Properties

A safety property is specified as a *deterministic* LTS that contains no τ transitions, and no π state. The set of traces $Tr(P)$ of property P defines the set of acceptable behaviours over αP . An LTS T satisfies P , denoted as $T \models P$ iff $Tr(T \uparrow \alpha P) \subseteq Tr(P)$.

The LTS automatically derives from a property P an *error LTS* denoted P_{err} , which traps possible violations with the π state. Formally, the error LTS of a property $P = \langle S, \alpha P, R, s_0 \rangle$ is $P_{err} = \langle S \cup \{\pi\}, \alpha P_{err}, R', s_0 \rangle$, where $\alpha P_{err} = \alpha P$ and $R' = R \cup \{(s, a, \pi) \mid a \in \alpha P \text{ and } \neg \exists s' \in S : (s, a, s') \in R\}$. Note that the error automaton is *complete*, i.e., each state (other than the error state) has outgoing transitions for every action in the alphabet.

For example, Figure 2 illustrates a mutual exclusion property for a system consisting of the LTSs of Figure 1.

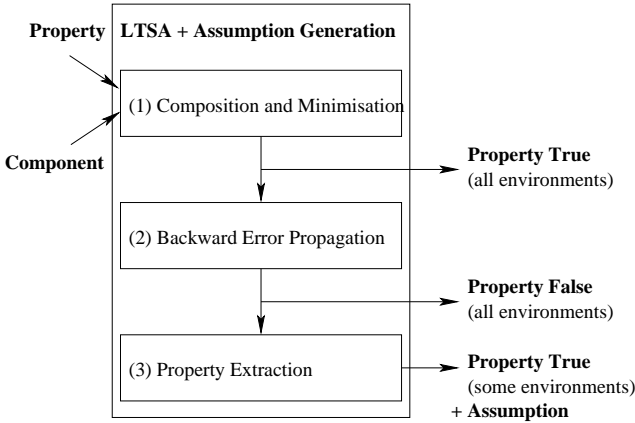


Figure 3. Model Checking with Assumption Generation

The property comprises states 0, 1, 2 and the transitions denoted by solid arrows. It expresses the fact that the component and its environment should never be in their critical sections at the same time. In other words, the intervals defined by their mutual *enterCS* and *exitCS* actions should never overlap. The dashed arrows illustrate the transitions to the error state that are added to the property to obtain its error LTS.

Let T be an LTS that has no error traces. To detect violations of property P by component T , the LTSA computes $T||P_{err}$. It has been proven in [8] that T violates P iff the π state is reachable in $T||P_{err}$, or equivalently, iff $errTr(T||P_{err}) \neq \emptyset$. The error state has special treatment during minimization, so that a violation does not disappear as a result of abstraction. In fact, an error state within a component can only disappear with composition, i.e., if a component with which it interacts blocks the erroneous behaviour.

3. Assumption Generation

In this section we describe in detail our proposed extensions to traditional model checking, and their implementation in LTSA. We also provide a formal proof of correctness.

3.1. General Method

The traditional approach to verifying a property of an *open system* (i.e., a software component that interacts with an environment, represented by other components) is to check it for all the possible environments. The result of verification is either **true**, if the property holds for *all* the possible environments, or **false**, if there exists *some* environment that can lead the component to falsify the property. We believe that this approach is overly pessimistic and only

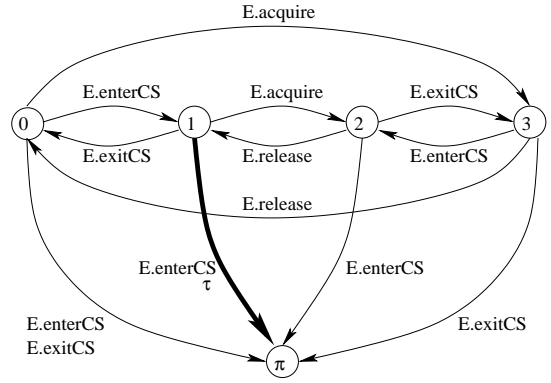


Figure 4. Composite LTS

appropriate for the analysis of *closed* systems, where no further interaction with the environment is expected. When analyzing open systems, an optimistic view, which assumes a *helpful* environment, is more appropriate. Usually, software components are required to satisfy properties in specific environments, so it is natural to accept a component if there are *some* environments in which the component does not violate the property.

In our approach, the result of component verification is also **true**, if the property holds for *all* environments. However, the result is **false** only if the property is falsified in *all* environments. If there exist *some* environments in which the component satisfies the property, the result of verification is not false, as in the traditional approach, but rather **true** in environments that satisfy a specific assumption. This assumption, i.e. a *property* LTS, is automatically generated and characterizes exactly those environments. Intuitively, this environment assumption encodes all possible “winning strategies” of the environment in a game between the system, which attempts to get to the error state, and the environment, which attempts to prevent this. Figure 3 illustrates our approach together with the steps we follow to build the assumptions (that are described below).

Step 1: Composition and Minimization

Given an *open system* and a *property* LTS that may relate the behaviour of the system with the behaviour of the environment, our first step is to compute all the violating traces of the system for unrestricted environments, and turn into τ all actions in these traces over which the environment has no control, i.e., the internal actions of the system. We perform this step by building the *composition* of the system with the *error* LTS of the property, and subsequently hiding the internal actions of the system. The resulting LTS can be minimized with respect to observational equivalence, since such minimization preserves traces.

For example, Figure 4 depicts the result of composing the components depicted in Figure 1 with the mutual exclu-




Figure 5. The Result after Backward Error Propagation

sion property of Figure 2, after minimization. The internal actions of the system, i.e. the “W” labelled transitions, were abstracted to τ .

If the error state is not reachable in this composition, the property is **true** in any environment, and this is reported to the user. Otherwise, we identify whether there exist environments that can help the system avoid the error in all circumstances; this is achieved through the following steps.

Step 2: Backward Error Propagation

This step first performs *backward propagation* of the error state over τ transitions, thus pruning the states where the environment cannot prevent the error state from being entered via one or more τ steps. Since we are interested only in the error traces, we also eliminate the states that are not backward reachable from the error state. If, as a result of this transformation, the initial state becomes an error state, it means that no environment can prevent the system from possibly reaching the error state, so the property is **false** (for all environments) and this is reported to the user.

Consider again the composite system in Figure 4. The thicker line marks the only τ transition that remains in the system after minimization. As a result of backward propagation, we identify state 1 with the error state; the result is shown in Figure 5. The intuition here is that, if the component is in a state from which it can violate the property by some number of internal moves, then no environment can prevent the violation from occurring.

Step 3: Property Extraction

This step builds the *property LTS* that is our assumption. It performs this in two stages; first it builds the *error LTS* for the assumption, from which it extracts the corresponding property LTS. Note that the LTS resulting from Step 2 might not be an error LTS, although it contains an error state. Recall from the background section that the error LTS is deterministic and complete.

In order to get an error LTS we make the LTS obtained from step 2 deterministic by applying to it τ elimination and the subset construction [3], but by taking special care of the π state as follows. During subset construction, the states




Figure 6. Generated Assumption

of the deterministic LTS that is being generated are *sets of states* in the original non-deterministic LTS. In our context, if any one of the states in the set is π , the entire set becomes π . Intuitively, a trace that non-deterministically may or may not lead to an error has to be considered as an error trace. Such non-determinism reflects the fact that, by performing a particular sequence of actions, the environment cannot guarantee that the component will avoid error states.

For example, consider again the composite system in Figure 4. There are two outgoing transitions from the initial state 0 that are labelled by the same environment action $E.enterCS$: one leads to the error state, while the other one leads to state 1. This means that if the environment performs action $E.enterCS$, it can not prevent the system from getting to the error, so we would like to identify state 1 with π . In our example in Figure 4, this was achieved during Step 2, but this may not be the case in general.

What remains to be performed at this stage is to make the resulting LTS complete. *Completion* is performed by adding a new “sink” state to the LTS, and adding a transition to this state for each missing transition in the “incomplete” LTS. The missing transitions in the incomplete LTS represent behaviour of the environment that is never exercised by the open system under analysis. As a result, no assumptions need to be made about these behaviours. The sink state reflects exactly this fact, since it poses no implementation restrictions to the environment.

Once we have the error LTS, we obtain the assumption by deleting the error state and the transitions that lead to it. Figure 6 depicts the assumption generated for our example. Since the result from Step 2 is already deterministic, we get the assumption by completing it with the sink state, denoted by θ , and deleting the π state. The assumption expresses the fact that the environment should only access its critical section protected by the mutex, $E.acquire$ and $E.release$ actions of the environment can only alternate, and therefore any different behaviour is inconsequential. Notice for example that from state 0, action $E.release$ leads to state θ .

Implementation in the LTSA

As mentioned, the LTSA provides a framework that facilitates the introduction of the extensions we have presented. For example, we took advantage of its support for composition, abstraction, minimization and determinization. The extra features that our approach required are:

- Special treatment of the error state, π , during determinization. The special semantics of this state were not previously taken into account.
- Backwards reachability and error propagation as required by step 2. We believe that error propagation should be performed during CRA for increased efficiency, irrespective of our approach.
- Completion with the sink state, θ , and property extraction from the error LTS.

3.2. Correctness of Approach

Let T denote an *open* system with alphabet αT and let E denote another system representing an arbitrary environment for T , whose alphabet is αE . Let P denote a *property* LTS with alphabet $\alpha P \subseteq \alpha T \cup \alpha E$ (a property may refer to actions in both T and E).

Let $\mathcal{C} = \alpha T \cap \alpha E$ be the set of *common* actions between T and E , and let $\mathcal{I} = \alpha T - \mathcal{C}$ denote the *internal* actions of the system.

Our tool generates the *property* LTS A with alphabet $\alpha A = \mathcal{C} \cup (\alpha P - \mathcal{I})$, representing the weakest assumption characterizing all the environments that, composed with the system, satisfy the property, i.e., $E \models A$ if and only if $E||T \models P$.

The following proposition says that the error traces of A_{err} are obtained from the traces in $T||P_{err}$ that *may* lead to an error state, from which we remove the actions not present in αA .

Proposition 3.1 $errTr(A_{err}) = \{\sigma \in \alpha A^* \mid \exists \sigma' \in errTr(T||P_{err}) \wedge \sigma = \sigma' \upharpoonright \alpha A\}$.

The following theorem makes precise the claim that A is the weakest assumption about the environment E of T that ensures property T .

Theorem 3.2 $\forall E, E \models A$ if and only if $E||T \models P$.

Proof.

- $\forall E$ such that $E \models A$, we have to show that $E||T \models P$. The proof is by contradiction.

Assume $E||T \not\models P$. Then, there is a trace σ in $E||T||P_{err}$ that leads to the error state (i.e., $\sigma \in errTr(E||T||P_{err})$). We use σ to build a trace $\sigma' \in$

$Tr(E)$ such that $\sigma' \upharpoonright \alpha A \in errTr(A_{err})$, thus contradicting $E \models A$.

Since σ is an error trace in $E||T||P_{err}$, it follows that $\sigma \upharpoonright \alpha E \in Tr(E)$ and $\sigma \upharpoonright (\alpha T \cup \alpha P) \in errTr(T||P_{err})$. From proposition 3.1, it follows that $(\sigma \upharpoonright (\alpha T \cup \alpha P)) \upharpoonright \alpha A \in errTr(A_{err})$.

Since $\alpha A \subseteq \alpha E$ and $\alpha A \subseteq \alpha T \cup \alpha P$, we also have that $(\sigma \upharpoonright \alpha E) \upharpoonright \alpha A = (\sigma \upharpoonright (\alpha T \cup \alpha P)) \upharpoonright \alpha A$. Let $\sigma' = \sigma \upharpoonright \alpha E$. We then have that $\sigma' \upharpoonright \alpha A = (\sigma \upharpoonright (\alpha T \cup \alpha P)) \upharpoonright \alpha A \in errTr(A_{err})$, and thus we have a contradiction.

- $\forall E$ such that $E||T \models P$, we have to show that $E \models A$. Again, we prove this by contradiction.

Assume $E \not\models A$. Then, there is a trace $\sigma \in Tr(E)$ such that $\sigma \upharpoonright \alpha A \in errTr(A_{err})$. From proposition 3.1, it follows that there is a trace $\sigma' \in errTr(T||P_{err})$ such that $\sigma \upharpoonright \alpha A = \sigma' \upharpoonright \alpha A$. We use σ and σ' to build a trace σ'' in $E||T||P_{err}$ such that $\sigma'' \upharpoonright \alpha P \in errTr(P_{err})$, thus reaching the contradiction of $E||T \models P$.

Since σ is a trace of E , σ' is a trace of $T||P_{err}$, $\sigma \upharpoonright \alpha A = \sigma' \upharpoonright \alpha A$ and $\mathcal{C} \subseteq \alpha A$ it follows that σ and σ' may differ only on non-common actions. It follows that there exists a trace σ'' in $E||T||P_{err}$ such that $\sigma'' \upharpoonright \alpha E = \sigma$ and $\sigma'' \upharpoonright (\alpha T \cup \alpha P) = \sigma'$. (we build σ'' by “composing” σ and σ' using the same rules as for parallel composition of systems).

Since $\sigma' \in errTr(T||P_{err})$, it follows that $\sigma' \upharpoonright \alpha P \in errTr(P_{err})$. We also have $\sigma'' \upharpoonright \alpha P = \sigma' \upharpoonright \alpha P$, since σ may introduce in σ'' only actions that are not present in αA or αP . It follows that $\sigma'' \upharpoonright \alpha P \in errTr(P_{err})$, and thus we have a contradiction. \square

4. Application: the Rover Executive

We experimented with our approach in the context of the verification of the executive subsystem for the K9 Mars Rover, developed at NASA Ames. The executive receives flexible plans from a Planner, which it executes according to the plan language semantics. A plan is a hierarchical structure of actions that the Rover must perform. Traditionally, plans are deterministic sequences of actions. However, increased Rover autonomy requires added flexibility. The plan language therefore allows for branching based on state or temporal conditions that need to be checked, and also for flexibility with respect to the starting time of an action. The plan language allows the association of each action with a number of state or temporal pre-, maintenance, and post-conditions, which must hold before, during, and on completion of the action execution, respectively.

The executive has been implemented as a multi-threaded system, made up of a main coordinating component named “Executive”, components for monitoring the state conditions “ExecCondChecker”, and temporal conditions “ExecTimerChecker” - each further decomposed into two threads - and finally an “ActionExecution” thread that is responsible for issuing the commands to the Rover. Synchronization between these threads is performed through mutexes and condition variables. The developers provided some design documents to us, which described the synchronization between these components in an add-hoc flowchart-style language. They looked very much like LTSs, which allowed us to translate them in a straightforward and systematic, albeit manual, way into FSP for the LTSA.

We first checked the occurrence of race conditions for the case of a variable of the ExecCondChecker shared with the Executive. We checked the property on the ExecCondChecker (that consists internally of two threads) together with the mutexes it uses, since mutexes constitute the synchronization mechanism in this system. The ExecCondChecker with mutexes and the property had 426 states but minimized to 18 states. The propagation of the error state produced an LTS of just 10 states, and the final assumption generated had 12 states (one being the sink state). We were surprised to see that our approach did not generate the expected assumption, i.e. that accesses to the shared variable by the environment must be protected by the appropriate mutex, as in the example of Section 3. In fact, the assumption obtained was weaker. It reflected the knowledge that, once the environment holds the mutex, the values that the environment reads reflect changes that only the environment may have made. For example, assume that, while holding the mutex, the environment assigns value x to the variable. Reading any value $x' \neq x$ would lead the environment to the sink state, because this behaviour will never actually be exercised in the context of the ExecCondChecker.

The second property that we checked in this fashion was one that the developer thought might be violated by the code, but could actually not produce an execution that would demonstrate this fact. For a specific variable of the ExecCondChecker shared with the Executive, the property stated the following: if the Executive reads the value of the variable, then the ExecCondChecker should not read this value until the Executive clears it first. Again, we used the ExecCondChecker together with mutexes and the property to generate an assumption on the behaviour of the Executive. The result had 524 states, minimized to 9 states, reduced to 7 states with error propagation, and to 6 states with determinization. The resulting assumption had 7 states (including the sink state). It stated that the environment of the component should read the variable after acquiring a mutex, and should hold on to that mutex until it clears the variable. Note that, again, there were transitions to the sink state, ex-

pressing the fact that some behaviour of the environment is never exercised. For example, the assumption made clear that the ExecCondChecker only updates the variable with values larger than the one it currently holds.

The assumption generated was satisfied by the design level Executive. Our result gave confidence to the developers about the correctness of their design and implementation. They also found it useful to be able to understand how the property decomposes across modules of the system.

5. Discussion

The complexity bottleneck of our approach is the determinization step, which, in the worst case, is exponential in the number of the states of the given LTS. There are several reasons that lead us to believe that this may not be the case often in practice. In our experiments such as the Rover study reported in Section 4, non-determinism almost disappears by propagation of the error state. As we only study modules of a larger system, we expect that the state space of these modules will be relatively small. This will be the case in particular when they interact through limited interfaces with their environment, which will allow the minimization step to considerably reduce their behaviour. Note also that, if we extend our results to other frameworks, the assumption may not be required to be deterministic. Admittedly, however, deterministic assumptions tend to be clearer to understand.

From our extensive experience with compositional reachability analysis (CRA) techniques, we are only too aware of the potential intermediate state explosion associated with them [18]. This problem describes the fact that, in lack of a context, a component may exhibit an excessively large state-space. However, this does not occur in the general case for well-designed software architectures. Moreover, several approaches have been proposed in the literature [18, 7, 26] for addressing the problem.

Our approach extends the LTSA tool in several useful ways. First of all, it achieves further reduction of component behavior by applying propagation of the error states, a computationally inexpensive but efficient step. Moreover, our approach generates the *weakest* environment assumptions. As such, these assumptions may be used for runtime monitoring, or for component retrieval, capabilities that were not formerly provided by the tool.

As far as component retrieval is concerned, we would like to stress the following observation from our experiments (Section 4). The sink state that our assumptions contain, reflects the fact that some services that a component provides will never be used in the context of a system. Our assumptions allow free implementations for these services, and simply ensure that the used services comply with the requirements.

The ability to generate assumptions also opens up a number of other interesting research directions: we mention a few to give some flavour.

- Assumptions may be further analyzed. Assume that a component does not violate a property in any environment. However, it may be that, when put in the context of the weakest assumption that our algorithms generate, no useful behavior is obtained as a result. In the extreme, the component with the assumption may result in a single deadlock state. Or the assumption may be that the environment will hold on to a specific lock for ever. All these are indications that there is something inherently wrong with the behavior of the component under analysis.
- Our work has been performed with a limited but important set of properties (safety) expressed within a specific framework that facilitates the development of our algorithms. However, we believe our approach has application in other frameworks. In particular, we are investigating the extension of our approach for the case of fairness and/or liveness properties, which requires a more expressive formalism.
- When the behavior of the environment, or part thereof, is provided, we wish to find effective ways of discharging assumptions on the environment. One way would be to use the assumption as a property, and model check in the same fashion components in the environment. This process can be seen as a way of decomposing, automatically, a property across components of a system. Indeed, an assumption reflects those aspects of the property that have not been satisfied by the component and that remain to be satisfied by its environment. Property decomposition is an extremely difficult problem, and our approach may be seen a helpful step in its facilitation. Of course, such decomposition will not be effective in all cases. It is easy to imagine that there will be cases where assumptions may gradually grow in size during this process, a problem referred to in the literature as “property explosion”.
- Our approach to assumption generation can straightforwardly be used for submodule construction, where the submodule is placed as an interacting component in parallel with the given one. Generalization to other forms of composition is a natural step, e.g. find a program context for the given component to achieve the desired property in some appropriate environment.

6. Related Work

For over three decades now, there has been research effort focused on finding tractable approaches to the formal

specification, design and development of complex systems. Significant early progress occurred with techniques and tools for sequential, non-interacting or transformational, systems. However, the quest for obtaining effective methods and tools for the formal support of compositional and/or modular development and reasoning for *reactive* systems still remains, in our view, a major challenge. As there is insufficient space to do justice to the work that has been undertaken, we refer the interested reader to the proceedings [14] - its introductory chapter in particular [12] - and the recent book [13].

In more recent years with the development and take-up of OO-design technology, formal techniques for support of component-based design is also gaining prominence, see for example [10, 11], for which *modular*-based reasoning is key. The work of Inverardi and colleagues, see [22] and [21] for example, has also been aimed at providing support for the modular checking of certain properties, as deadlock freedom, but is somewhat limited in the checks performed for compatibility between components.

In order to make progress in any of these areas, some form of assumption (either implicit or explicit) about the interaction with, or interference from, the environment has to be made, [23, 2]. Even though we have sound and complete reasoning systems for such rely-guarantee (or assumption-commitment) style of reasoning, see for example [24, 39] and most recently [41], it is always a mental challenge to obtain the most appropriate assumption (if there is such). It is even more of a challenge to find automated techniques to support this reasoning style - the thread modular reasoning underlying the Calvin tool [16] is one start in this direction. In the framework of temporal logic, the work on Alternating time Temporal Logic ATL (and transition systems) [5] was proposed for the specification and verification of open systems together with automated support via symbolic model checking procedures, albeit of rather high complexity; the Mocha toolkit [4] provides support for modular verification of components with requirement specification based on the ATL. It goes without saying that if tool support is lacking, take-up of these techniques will be rather low.

The underlying approach to automated assumption generation that we’ve adopted and implemented in LTSA has similarity to a number of other problems that have been considered by a number of researchers over the past two decades. Closest to our our work in the software engineering and concurrency theory are the “sub-module construction problem”, “scheduler synthesis” and “interface equation solving” problems. In the discrete event community, it appears as the “supervisory control” problem, in control theory there is the “model matching” problem and in the logic synthesis world there is the “interacting FSM synthesis”. Of course, the particular frameworks in which these problems are considered makes all the difference to their

solution(s) and as such it would be quite inappropriate to claim they are solving the same problem. However, in very general terms, each can be seen as an instance of the following problem, given a component, C , and a desired behaviour, B , find a context for C , X , such that $X(C) \equiv B$, for some appropriate notion of equivalence.

Merlin and Bochmann [33] were probably the first to address the above as submodule construction in the world of communication protocol specification and synthesis. In a setting of labelled transition systems, given a module specification M_0 and a submodule specification M_1 , they outlined and exemplified a manual approach to construct an interacting submodule M_2 such that M_1 and M_2 together achieve the desired specification of M_0 . Their construction has much in common with ours although some significant aspects of the construction were left to the reader's imagination. The later work of [38, 19] has revisited the Merlin-Bochmann approach and provided new, detailed, algorithms for the sub-module construction and implemented an automated tool. One recognized limitation of the Merlin-Bochmann is that the notion of correctness, namely just trace equivalence, does not capture a number of behavioural properties, e.g. potential deadlock.

The work of Shields [37], over a decade later, introduces the "Interface Equation" in the setting of the process algebra, CCS [34], under observational equivalence. In order to solve $(C|X)\setminus_L = B$ for the process X , he restricts to cases where B is deterministic, with some minor restrictions on the sorts of C and B , and provides necessary and sufficient conditions for a solution to exist and then in such situations presents an explicit construction. Parrow [35] also addressed the interface equation and presented a procedure for solving the equations via successive transformation of the CCS equations to simpler ones, generating a solution along the way; his approach is based upon a tableau method. Parrow's method attempts to find a *most general* solution, but even if it exists, it is not necessarily appropriate for implementation. Continuing in the process algebra framework, Larsen and Xinxin [28] consider the more general problem of solving a system of equations $C_i(X) \simeq P_i$, for $0 < i \leq n$, where the C_i are arbitrary contexts, P_i are arbitrary processes and X is the process to be found - the equivalence is taken as bisimulation. They considered the problem in the context of disjunctive modal transition systems, [27] and implemented an automated tool for solving the equations (in the finite state case) when a solution exists.

As stated above, there is a further body of work in supervisory control synthesis, discrete event systems, and logic synthesis areas, see for example [1, 15, 40, 25, 6]. However, we should stress that whilst these approaches are in general set in a FSM/DFA context, the principal goal is quite different in comparison with ours.

7. Conclusions

We presented an approach to model checking components as open, rather than closed systems. Our approach reports whether there is something inherently wrong with the component behaviour, or whether satisfying a requirement is simply a matter of providing the right environment. Moreover, it characterizes exactly all helpful environments.

The possibility of generating assumptions provides increased flexibility in model checking, and opens up a number of interesting research topics. It allows, for example, the discharge of assumptions at run-time for unpredictable environments, the retrieval of components focused on only relevant aspects of their behaviour, or the decomposition of properties across components. It remains to further investigate how useful our approach is in practice. Open research issues include optimizations and extensions for fairness/liveness properties and other frameworks. However, our early experiments with real case studies provide strong evidence in favour of this line of research.

References

- [1] A. Aziz, F. Balarin, R. K. Brayton, M. D. Dibenedetto, A. Sladanha, and A. L. Sangiovanni-Vincentelli. Supervisory control of finite state machines. In *7th Int. Conference on Computer Aided Verification*, volume 939 of *Lecture Notes in Computer Science*, Liège, Belgium. Springer Verlag.
- [2] M. Abadi and L. Lamport. Composing specifications. *ACM Transactions on Programming Languages and Systems*, 15(1):73–132, January 1992.
- [3] A. V. Aho, J. E. Hopcroft, and J. D. Ullman. *Introduction to Automata Theory, Languages, and Computation*. Addison-Wesley, 2000.
- [4] R. Alur, T. Henzinger, F. Mang, S. Qadeer, S. Rajamani, and S. Tasiran. Mocha: Modularity in model checking. In *Proceedings of 10th International Conference on Computer Aided Verification*, volume 1427 of *Lecture Notes in Computer Science*, pages 521–525. Springer Verlag, 1998.
- [5] R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. In de Roever et al. [14], pages 23–60.
- [6] S. Balemi, G. Hoffmann, P. Gyugyi, H. Wong-Toi, and G. Franklin. Supervisory control of a rapid thermal multiprocessor. *IEEE Transactions on Automatic Control*, 38(7):1040–1059, July 1993.
- [7] S. Cheung and J. Kramer. Context constraints for compositional reachability analysis. *ACM Transactions on Software Engineering and Methodology*, 5(4):334–377, 1996.
- [8] S. Cheung and J. Kramer. Checking safety properties using compositional reachability analysis. *ACM Transactions on Software Engineering and Methodology*, 8(1):49–78, 1999.
- [9] E. M. Clarke, O. Grumberg, and D. Peled. *Model Checking*. MIT Press, 2000.
- [10] L. de Alfaro and T. Henzinger. Interface automata. In *Proc. of the Joint 8th European Software Engineering Conference and 9th ACM SIGSOFT International Symposium on Foundations of Software Engineering*. ACM Press, 2001.

- [11] L. de Alfaro and T. Henzinger. Interface theories for component-based design. In *Proceedings of EMSOFT 01: Embedded Software*, volume 2211 of *Lecture Notes in Computer Science*, pages 148–165. Springer Verlag, 2001.
- [12] W.-P. de Roever. The need for compositional proof systems: A survey. In de Roever et al. [14], pages 1–22.
- [13] W.-P. de Roever, F. de Boer, U. Hanneman, J. Hooman, Y. Lakhnech, M. Poel, and J. Zwiers. *Concurrency Verification: Introduction to Compositional and Non-compositional Methods*. Cambridge University Press, 2001.
- [14] W.-P. de Roever, H. Langmaack, and A. Pnueli, editors. *Compositionality: The Significant Difference - An International Symposium, COMPOS'97*, volume 1536 of *Lecture Notes in Computer Science*. Springer-Verlag, 1997.
- [15] M. di Benedetto and A. Sangiovanni-Vincentelli. Model matching for finite-state machines. *IEEE Transactions on Automatic Control*, 46(11):1726–1743, November 2001.
- [16] C. Flanagan, S. Freund, and S. Qadeer. Thread-modular verification for shared-memory programs. In *Proceedings of the European Symposium on Programming*, 2002.
- [17] D. Giannakopoulou, J. Kramer, and S. Cheung. Analysing the behaviour of distributed systems using Tracta. *Journal of Automated Software Engineering, special issue on Automated Analysis of Software*, 6(1):7–35, 1999.
- [18] S. Graf, B. Steffen, and G. Lüttgen. Compositional minimisation of finite state systems using interface specifications. *Formal Aspects of Computation*, 8, 1996.
- [19] E. Haghverdi and H. Ural. Submodule construction from concurrent system specifications. *Information and Software Technology*, 41:499–506, 1999.
- [20] K. Havelund and G. Rosu. Monitoring Java programs with Java PathExplorer. In *First Workshop on Runtime Verification (RV'01)*, volume 55(2) of *Electronic Notes in Theoretical Computer Science*, Paris, France, 2001.
- [21] P. Inverardi and S. Scriboni. Connectors synthesis for deadlock-free component based architectures. In *Proceedings of 16th IEEE Annual International Conference on Automated Software Engineering*, pages 174–181, 2001.
- [22] P. Inverardi, A. Wolf, and D. Yankelevich. Static checking of system behaviors using derived component assumptions. *ACM Transactions on Software Engineering Methods*, 9(3):239–272, July 2000.
- [23] C. Jones. Specification and design of (parallel) programs. In R. Mason, editor, *Information Processing 83: Proceedings of the IFIP 9th World Congress*, pages 321–332. IFIP: North Holland, 1983.
- [24] C. Jones. Tentative steps towards a development method for interfering programs. *ACM Transactions on Programming Languages and Systems*, 5(4):596–619, 1983.
- [25] S. Khatri, A. Narayan, S. Krishnan, K. McMillan, R. Brayton, and A. Sangiovanni-Vincentelli. Engineering change in a non-deterministic FSM setting. In *Proceedings of 33rd IEEE/ACM Design Automation Conference*, 1996.
- [26] J.-P. Krimm and L. Mounier. Compositional state space generation from LOTOS programs. In E. Brinksma, editor, *3rd International Workshop on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'97)*, volume 1217 of *Lecture Notes in Computer Science*, Enschede, The Netherlands, 1997. Springer.
- [27] K. Larsen and B. Thomsen. A modal process logic. In *Proceedings of the IEEE/ACM Conference on Logic in Computer Science, LICS'88*, 1988.
- [28] K. Larsen and L. Xinxin. Equation solving using modal transition systems. In *Proceedings of the IEEE/ACM Conference on Logic in Computer Science, LICS'90*, 1990.
- [29] J. Magee, N. Dulay, and J. Kramer. Regis: A constructive development environment for parallel and distributed programs. *Distributed Systems Engineering Journal, Special Issue on Configurable Distributed Systems*, 1(5):304–312, 1994.
- [30] J. Magee and J. Kramer. *Concurrency: State Models & Java Programs*. John Wiley & Sons, 1999.
- [31] J. Magee, J. Kramer, and D. Giannakopoulou. Behaviour analysis of software architectures. In *1st Working IFIP Conference on Software Architecture (WICSA1)*, San Antonio, TX, USA, 1999.
- [32] J. Magee, N. Pryce, D. Giannakopoulou, and J. Kramer. Graphical animation of behavior models. In *22d International Conference on Software Engineering (ICSE 2000)*, Limerick Ireland, 2000. ACM.
- [33] P. Merlin and G. V. Bochmann. On the construction of submodule specification and communication protocols. *ACM Transactions on Programming Languages and Systems*, 5:1–25, 1983.
- [34] R. Milner. *Communication and Concurrency*. Prentice-Hall, 1989.
- [35] J. Parrow. Submodule construction as equation solving CCS. *Theoretical Computer Science*, 68:175–202, 1989.
- [36] C. Păsăreanu, M. Dwyer, and M. Huth. Assume-guarantee model checking of software: A comparative case study. In D. Dams, R. Gerth, S. Leue, and M. Massink, editors, *Theoretical and Practical Aspects of SPIN Model Checking*, volume 1680 of *Lecture Notes in Computer Science*, pages 168–183. Springer-Verlag, 1999.
- [37] M. Shields. A note on the simple interface equation. *The Computer Journal*, 32(5):399–412, 1989.
- [38] D. P. Sidhu and J. Aristizabal. Constructing submodule specifications and network protocols. *IEEE Transactions on Software Engineering*, 14(11):1565–1577, November 1988.
- [39] K. Stølen. A method for the development of totally correct shared-state parallel programs. In J. Baeten and J. Groote, editors, *Proceedings of Concur'91*, volume 527 of *Lecture Notes in Computer Science*. Springer Verlag, 1991.
- [40] E. Tronci. Automatic synthesis of controllers from formal specifications. In *Proc. of 2nd IEEE Int. Conf. on Formal Engineering Methods, Brisbane, Australia*, 1998.
- [41] Q. Xu, W.-P. de Roever, and J. He. The rely-guarantee method for verifying shared variable concurrent programs. *Formal Aspects of Computing*, 9(2):149–174, 1997.