

Compositional Model Checking

E. M. Clarke D. E. Long K. L. McMillan
April 19, 1989
CMU-CS-89-145

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

To appear in the *Proceedings of the Fourth IEEE Symposium on
Logic in Computer Science*, June 4–8, 1989, Asilomar, CA.

Abstract

We describe a method for reducing the complexity of temporal logic model checking in systems composed of many parallel processes. The goal is to check properties of the components of a system and then deduce global properties from these local properties. The main difficulty with this type of approach is that local properties are often not preserved at the global level. We present a general framework for using additional *interface processes* to model the environment for a component. These interface processes are typically much simpler than the full environment of the component. By composing a component with its interface processes and then checking properties of this composition, we can guarantee that these properties will be preserved at the global level. We give two example compositional systems based on the logic CTL*.

This research was sponsored in part by the Defense Advanced Research Projects Agency (DOD), ARPA Order No. 4976, Amendment 20, under Contract Number F33615-87-C-1499, monitored by the:

Avionics Laboratory

Air Force Wright Aeronautical Laboratories
Aeronautical Systems Division (AFSC)
United States Air Force
Wright-Patterson AFB, Ohio 45433-6543

This research was also partially supported by NSF grant CCR-87-226-33. The second author holds an NSF graduate fellowship.

1 Introduction

Temporal logic model checking procedures ([5, 6, 11, 18, 22, 23]) have been successful in finding subtle errors in relatively small finite state systems ([3, 4, 9]), but they all suffer from one apparently unavoidable problem: the state explosion problem. This problem arises in systems composed of many parallel processes. In general, the size of a parallel composition may grow as the product of the sizes of the components. Because of this phenomenon, a program with a relatively small number of processes may have far too many states for a model checking procedure to be directly useful. An obvious technique for avoiding the state explosion problem is to exploit the natural decomposition of a complex parallel program into processes. The goal of this approach is to verify the processes individually and then piece the results together to conclude that the original program is correct. The main difficulty is knowing when some property of a component process remains true in a parallel composition involving that process. It is easy to come up with examples where a critical property of some process is not preserved when the process is composed with another process. A similar technique may be used when the program is constructed in a modular or hierarchical fashion. In this case, lower level components may be simplified by hiding details that do not need to be visible externally and merging those states that become indistinguishable. If the original program is reconstructed from the simplified components, then the number of states will, in general, be much smaller and the program can be checked for correctness more easily. The problem this time is ensuring that the simplified program satisfies the same logical properties as the original program.

There have been a number of other papers on compositional techniques for reasoning about systems of processes. Milner's CCS calculus [19] is certainly compositional in nature, but it is only suitable for showing equivalence between processes and does not handle liveness properties. Barringer [1] has written several papers that show how to give a compositional temporal semantics for a parallel programming language like CSP, but it is not clear how such a semantics can be used in developing a compositional model checking algorithm. Pnueli [21] has developed a compositional proof system for temporal logic that is based on the assume-guarantee paradigm. The primitive formulas in his logic are triples of the form $\langle \varphi \rangle P \langle \psi \rangle$, where φ and ψ are temporal formulas and P is a process. A formula is true if P is guaranteed to behave according to ψ in any execution in which the environment behaves according to φ . One problem with this approach is the potential difficulty in expressing the necessary assumptions, since this can amount to giving the complete specification for an automaton in temporal logic. Josko [17] has developed a model checking algorithm in which the environment is modeled by a temporal logic formula, but the complexity of his procedure is exponential in worst case. Mishra and Clarke [20] have investigated a model checking algorithm for asynchronous circuits that can exploit the hierarchical structure of the circuit, but their approach is restricted to a particular fragment

of the logic CTL and is not as general as the one presented here.

In the present paper we use a different, although equally natural approach; we model the environment of a process by another process called an interface process. A rule of inference called the interface rule provides the basis for our compositional model checking technique. The interface rule allows us to deduce properties of a composition by checking properties of the individual processes. The complexity of showing that φ holds for $P_1 \parallel P_2$ by using this technique is roughly the same as the complexity of computing the parallel composition of P_1 and the interface process for P_2 . We present an algorithm for constructing the interface processes from P_1 and P_2 . We also give a general framework for the interface rule that is independent of any particular process model or logic. Within this framework we state four simple conditions that must be true of a process model and an associated logic in order for the inference rule to hold. These conditions can be easily checked to show that a new logic satisfies the inference rule.

We give two examples of compositional systems for which the inference rule is valid. Both systems use a branching time logic based on the temporal logic CTL* [10]. The first example uses an asynchronous process model and a notion of composition that is similar to the one used in theoretical CSP [15]. We define an equivalence relation on processes that allows for finite stuttering along computation paths. This definition is appropriate for reasoning about asynchronous processes since there is no notion of “next system state” in such cases. To illustrate the ideas, we prove the correctness of a tree arbiter. The proof is interesting since it shows how a simple inductive inductive argument can be combined with model checking to show that tree arbiters of arbitrary size are correct. The second system is designed for synchronous digital systems. There is an efficient algorithm ($O(n \log n)$) for deciding equivalence of processes in this model. As an illustration of this model, we consider a simple CPU with decoupled access and execution units.

2 The interface rule

In this section we give the basic rule of inference that is used in the remainder of the paper for obtaining compositional proofs for systems of finite state processes. We show that the rule is sound whenever the set of processes and the logic for reasoning about them satisfy four general and easily checked conditions. The importance of this section is not in the complexity of the soundness proof (which is quite simple) but in the generality of the conditions that we give. Let \mathcal{P} be a set of finite state processes, and assume that we know what it means for two processes P_1 and P_2 to be equivalent ($P_1 \equiv P_2$). Each process will have associated with it a certain set of atomic propositions that is used in distinguishing states and transitions. Σ_P will denote the set associated with process P . The set of propositions associated with the parallel composition

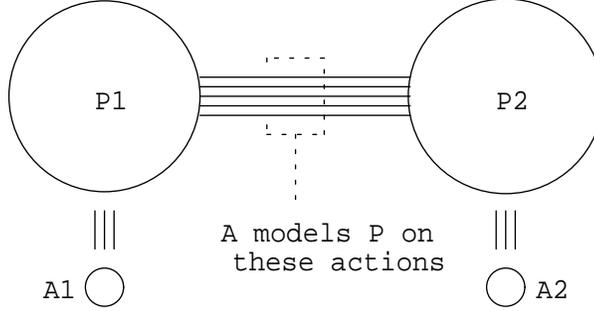


Figure 1: The interface rule

of two processes will be the union of the sets associated with the individual processes: $\Sigma_{P_1 \parallel P_2} = \Sigma_{P_1} \cup \Sigma_{P_2}$. $P \downarrow \Sigma_1$ will be the restriction of P to Σ_1 . This process is obtained by hiding all of the symbols in Σ_P that are not in Σ_1 ; consequently, $\Sigma_{P \downarrow \Sigma_1} = \Sigma_P \cap \Sigma_1$.

Suppose, in addition, that we have a logic \mathcal{L} for reasoning about the processes in \mathcal{P} and that we know what it means for a formula φ to be true of a process P ($P \models \varphi$). Each formula will be constructed from some set of atomic propositions, and we will write $\varphi \in \mathcal{L}(\Sigma)$ to indicate that the propositions used in φ are a subset of Σ . The *interface rule* deals with the parallel composition of two processes P_1 and P_2 . The reader might think of the processes as being connected by a set of wires as shown in figure 1, where the wires correspond to symbols in $\Sigma_{P_1} \cap \Sigma_{P_2}$. A_1 and A_2 are interface processes for P_2 and P_1 , respectively. Intuitively, A_1 is all that P_2 can observe of P_1 through the wires that connect them. An analogous relationship holds between processes A_2 and P_1 . There are two basic inference rules, which we state below:

$$\frac{P_1 \downarrow \Sigma_{P_2} \equiv A_1, \quad \varphi \in \mathcal{L}(\Sigma_{P_2}), \quad A_1 \parallel P_2 \models \varphi}{P_1 \parallel P_2 \models \varphi} \qquad \frac{P_2 \downarrow \Sigma_{P_1} \equiv A_2, \quad \psi \in \mathcal{L}(\Sigma_{P_1}), \quad P_1 \parallel A_2 \models \psi}{P_1 \parallel P_2 \models \psi}$$

If a state minimization procedure is known for \mathcal{P} , then A_1 can be obtained by running this algorithm on $P_1 \downarrow \Sigma_{P_2}$. If such a procedure is not available, the rule will still be useful as long as the size of A_1 is significantly less than the size of P_1 . A similar comment applies to A_2 . The soundness of the interface rule depends on the properties that we enumerate below:

- i. Suppose $\Sigma_{P_1} = \Sigma_{P_2}$, then $P_1 \equiv P_2$ implies $\forall \varphi \in \mathcal{L}(\Sigma_{P_1}) [P_1 \models \varphi \leftrightarrow P_2 \models \varphi]$.
- ii. If $P_1 \equiv P_2$ and Q is another process, then $P_1 \parallel Q \equiv P_2 \parallel Q$ and $Q \parallel P_1 \equiv Q \parallel P_2$.

- iii. $(P_1 \parallel P_2) \downarrow \Sigma_{P_1} \equiv P_1 \parallel (P_2 \downarrow \Sigma_{P_1})$ and $(P_1 \parallel P_2) \downarrow \Sigma_{P_2} \equiv (P_1 \downarrow \Sigma_{P_2}) \parallel P_2$.
- iv. If $\varphi \in \mathcal{L}(\Sigma)$ and $\Sigma \subseteq \Sigma_P$, then $P \models \varphi$ iff $P \downarrow \Sigma \models \varphi$.

It is easy to show that $P_1 \parallel P_2 \models \varphi$ follows from the three hypotheses of the theorem and the above four properties. We can also show the soundness of simple rules like:

$$\frac{\begin{array}{l} P_1 \downarrow \Sigma_{P_2} \equiv A_1, \quad P_2 \downarrow \Sigma_{P_1} \equiv A_2, \\ \varphi \in \mathcal{L}(\Sigma_{P_2}), \quad \psi \in \mathcal{L}(\Sigma_{P_1}), \\ A_1 \parallel P_2 \models \varphi, \quad P_1 \parallel A_2 \models \psi \end{array}}{P_1 \parallel P_2 \models \varphi \wedge \psi}$$

We can handle arbitrary boolean combinations of formulas in an analogous manner.

3 An asynchronous process model

We begin by defining a simple model of a communicating system. Processes in this model are asynchronous and communicate using shared synchronization actions. They are represented as a form of transition system [19].

Definition 1 *A finite transition system, or fts, is a quadruple $L = \langle K, q_0, \Sigma, \Delta \rangle$, where:*

- i. K is a finite set of states.
- ii. $q_0 \in K$ is an initial state.
- iii. Σ is a finite set of actions not containing τ .
- iv. $\Delta \subseteq K \times (\Sigma \cup \{\tau\}) \times K$ is a transition relation.

Here, τ represents an internal action of the fts. We can then define the operations of composition, hiding, and renaming on fts. The notation $\Delta(p, \sigma)$ is used to indicate $\{q \mid (p, \sigma, q) \in \Delta\}$ if $\sigma \in \Sigma$. By convention, $\Delta(p, \sigma) = \{p\}$ if $\sigma \notin \Sigma$.

Definition 2 $L'' = L \parallel L'$ (the composition of L and L') is given by:

- i. $K'' = K \times K'$.
- ii. $q_0'' = (q_0, q_0')$.
- iii. $\Sigma'' = \Sigma \cup \Sigma'$.
- iv. $\Delta''((p, p'), \sigma) = \Delta(p, \sigma) \times \Delta'(p', \sigma)$ for $\sigma \in \Sigma''$.
 $\Delta''((p, p'), \tau) = \Delta(p, \tau) \times \{p'\} \cup \{p\} \times \Delta'(p', \tau)$.

The notion of composition used here is in the style of CSP [15]; an action σ in the first fts synchronizes with an identical action in the second fts. In a similar fashion we define the notions of hiding (denoted $L \setminus \sigma$) and renaming (denoted $L[\sigma_0 / \sigma_1]$). We note that the operations as defined have certain obvious algebraic properties. For example, composition is commutative and associative up to isomorphism.

Associated with an fts is a set of runs through the structure. We use the following definitions and terminology. For a finite set S , the notation S^* denotes the set of finite sequences of elements of S , S^ω denotes the infinite sequences, and $S^\infty = S^* \cup S^\omega$.

Definition 3 *A run from $q \in K$ is a pair (q, π) , where $\pi = (p_0, \sigma_0, p_1)(p_1, \sigma_1, p_2) \dots \in \Delta^\infty$ and $p_0 = q$. A run from the initial state q_0 is often simply called a run.*

We write P_q^* and P_q^ω for the finite and infinite runs from q , and define $P_q^\infty = P_q^* \cup P_q^\omega$.

For the required equivalence between structures, we use the notion of a bisimulation [19]. Formally, for an fts L , we have the following.

Definition 4 *$F : 2^{K \times K} \rightarrow 2^{K \times K}$ is monotonic if $R_0 \subseteq R_1 \subseteq K \times K$ implies $F(R_0) \subseteq F(R_1)$. F is equivalence-preserving if $R \subseteq K \times K$ being an equivalence relation implies $F(R)$ is an equivalence relation.*

At this point we restrict our attention to F which are monotonic.

Definition 5 *R is an F -bisimulation if $R \subseteq F(R)$. We define $p \approx_F q$ iff there is an F -bisimulation R such that $p R q$.*

We also define a series of approximations $\approx_{n,F}$ to \approx_F by successively applying F to $K \times K$. The intersection of these approximations is denoted $\approx_{\omega,F}$ and is a fixpoint of F . We show that $\approx_{\omega,F} = \approx_F$ and that \approx_F is an equivalence relation if it is equivalence-preserving.

4 An asynchronous process logic

We illustrate the development of a compositional system within the framework of branching time temporal logic. The logic we use is essentially the computation tree logic (CTL) of Emerson and Clarke [5]. The process model is that defined in the previous section. We define a class of formulas $\mathcal{CTL}(\Sigma)$ specifying temporal properties over the alphabet Σ . In general $\mathcal{CTL}(\Sigma)$ formulas may express the existence of a potentially infinite computation. When dealing with infinite paths, it is often desirable to impose certain fairness constraints on the possible paths. For example, we would typically wish to insure that each element of a composition must make progress if possible. One method for doing this would be to extend our notion of an fts to contain some type of fairness

constraint. Here we take a simpler but less flexible approach. A run is fair if every transition which is enabled infinitely often occurs infinitely often. This condition is sufficient to guarantee progress by all components in a composition, and we show that it is suitable for use with $\mathcal{CTL}(\Sigma)$. It is *not* suitable for use with a linear temporal logic. We show how to define a labeling $\mathcal{L}(p)$ to each state p of an fts. The semantics of $\mathcal{CTL}(\Sigma)$ are defined in a straightforward fashion with respect to this labeling, with the provision that path quantifiers are restricted to fair runs. The equivalence we use, which is denoted by \approx_p , is defined by the following functional.

Definition 6 $p \xrightarrow{\sigma}_R q$, where $R \subseteq K \times K$, if there exists $(p, (p, \tau, p_1) \dots (p_{n-1}, \tau, p_n)(p_n, \sigma, q)) \in P_p^*$ such that $p R p_1 R \dots R p_n$. In this definition, we allow $\sigma = \tau$.

Definition 7 $F_p(R) = \{(p, q) \mid \forall \sigma \in \Sigma \cup \{\tau\} :$

- i. $\forall p' [p \xrightarrow{\sigma}_R p' \text{ implies } \exists q' (q \xrightarrow{\sigma}_R q' \wedge p' R q')]$
- ii. $\forall q' [q \xrightarrow{\sigma}_R q' \text{ implies } \exists p' (p \xrightarrow{\sigma}_R p' \wedge p' R q')]$

Lemma 1 Assume L and L' are fts, $\Sigma = \Sigma'$. Let $p \in K$, $p' \in K'$ be states, $\mathcal{L}(p) = \mathcal{L}(p') \wedge p \approx_p p'$, and let $\rho \in P_p^*$. Then there exist $\rho' \in P_{p'}^*$, and partitions $B_0 B_1 \dots B_n$ and $B'_0 B'_1 \dots B'_n$ of $\text{st}(\rho)$ and $\text{st}(\rho')$ such that for all $q \in B_i$ and $q' \in B'_i$, $\mathcal{L}(q) = \mathcal{L}(q') \wedge q \approx_p q'$.

Theorem 1 Assume L and L' are fts with $\Sigma = \Sigma'$. Let p and p' be states of L and L' with $\mathcal{L}(p) = \mathcal{L}(p') \wedge p \approx_p p'$. Then for all $\varphi \in \mathcal{CTL}(\Sigma)$:

$$L, p \models \varphi \text{ iff } L', p' \models \varphi.$$

Corollary 1 If $L \approx_p L'$, then for all $\varphi \in \mathcal{CTL}(\Sigma)$ we have $L, q_0 \models \varphi$ iff $L', q'_0 \models \varphi$.

We also show that the equivalence is a congruence with respect to the operations. For hiding and renaming, the result is trivial. For composition we have the following.

Theorem 2 Let L_0, L'_0, L_1 , and L'_1 be fts with $\Sigma_0 = \Sigma'_0$, $\Sigma_1 = \Sigma'_1$, $L_0 \approx_p L'_0$, and $L_1 \approx_p L'_1$. Then $L_0 \parallel L_1 \approx_p L'_0 \parallel L'_1$.

Using corollary 1 and theorem 2 it is easy to show that the first two conditions required for the interface rule are satisfied. The last two conditions have straightforward proofs.

As an example, we consider a tree arbiter used to control access to a shared resource. An arbiter cell has three communication channels which we denote by $C0$, $C1$, and Cp . Each channel consists of two signals, r and a , representing a request and an acknowledgement. A user request on one of the channels $C0$

or $C1$ initiates a request to a server on channel Cp . After an acknowledge is received on Cp , an acknowledge is passed on to the user. At this point the user is assumed to have access to the shared resource. The user initiates another request/acknowledge cycle when finished. By combining arbiter cells into a binary tree, we can form an arbiter for any number of users. An arbiter with three cells is shown in figure 2. The specification here is based on an example presented in [8]. We will represent an arbiter cell by the composition of the fts shown in figure 3. The fts for the users and the server are shown in figure 4.

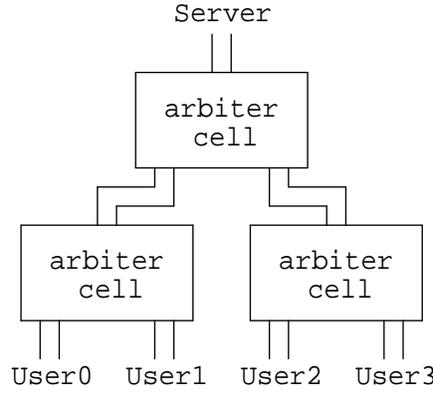


Figure 2: Three cell tree arbiter

We write *Arbiter* for the arbiter cell. To verify the class of tree arbiters, we begin by checking the following relations:

$$\begin{aligned}
& ((Arbiter \parallel User0 \parallel User1)[rp / r0'][ap / a0']) \\
& \quad \backslash r0 \backslash a0 \backslash r1 \backslash a1 \backslash t0 \backslash t1 \approx_p User0' \\
& ((Arbiter \parallel User0 \parallel User1)[rp / r1'][ap / a1']) \\
& \quad \backslash r0 \backslash a0 \backslash r1 \backslash a1 \backslash t0 \backslash t1 \approx_p User1' \\
& ((Arbiter \parallel User1 \parallel Server)[r0 / rp'][a0 / ap']) \\
& \quad \backslash r1 \backslash a1 \backslash rp \backslash ap \backslash t0 \backslash t1 \approx_p Server' \\
& ((Arbiter \parallel User0 \parallel Server)[r1 / rp'][a1 / ap']) \\
& \quad \backslash r0 \backslash a0 \backslash rp \backslash ap \backslash t0 \backslash t1 \approx_p Server'
\end{aligned}$$

The first relation here indicates that when we compose two users with an arbiter cell and restrict to the actions for the server port, the result is equivalent to another user. Thus, a user process can be used as an interface process for the two users and the arbiter cell. The other relations have similar interpretations. From these, we can perform an induction on the structure of a tree arbiter to deduce that each cell in an arbiter with users at the leaves and a single server at the root is equivalent to a cell in an environment of two users and a server, i.e., to:

$$Arbiter \parallel User0 \parallel User1 \parallel Server$$

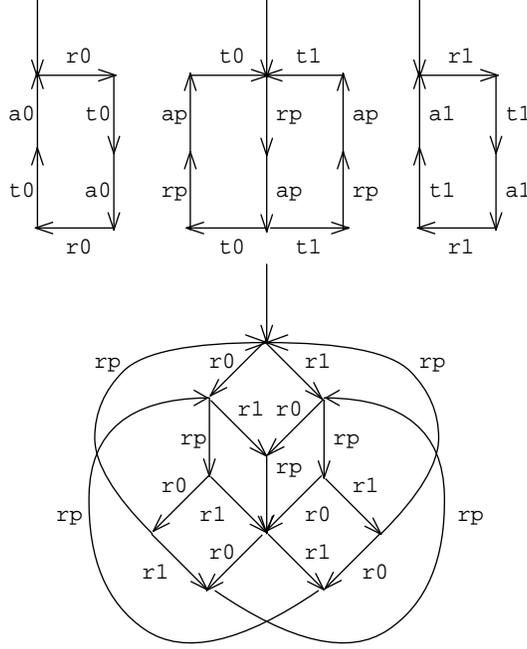


Figure 3: Tree arbiter cell

Properties of the entire arbiter are deduced from properties of these components. For example, liveness for the first user can be checked immediately by verifying:

$$\forall G(r0 \rightarrow \forall F(r0 \wedge a0)).$$

We can ensure mutual exclusion by checking:

$$\forall G(\neg(r0 \wedge a0) \vee \neg(r1 \wedge a1))$$

and:

$$\forall G((r0 \wedge a0) \rightarrow (rp \wedge ap)) \wedge \forall G((r1 \wedge a1) \rightarrow (rp \wedge ap)).$$

This example illustrates how it is sometimes possible to reason inductively about a system using the interface theorem.

5 A synchronous model and logic

In this section, we introduce a simple formal model of finite state machine composition which satisfies the conditions for the interface theorem and is applicable to synchronous systems. We use communicating Moore machines as our model.

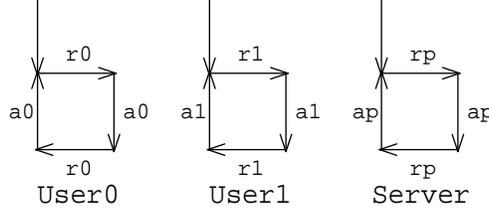


Figure 4: Users and server

A Moore machine has Boolean valued inputs and outputs and an internal state. The outputs are a function of the internal state only, while the next state is a function of the current state and the inputs.

Definition 8 A Moore machine is given by a structure $M = (K, q_0, \Sigma_i, \Sigma_o, \Gamma, \Theta)$ where

- i. K is the set of states
- ii. $q_0 \in K$ is the initial state
- iii. Σ_i and Σ_o are disjoint sets of identifiers (representing the inputs and outputs respectively)
- iv. Γ is a mapping $K \rightarrow \Sigma_o \rightarrow \{0, 1\}$ (the function mapping state to output)
- v. Θ is a mapping $K \rightarrow (\Sigma_i \rightarrow \{0, 1\}) \rightarrow K$ (the function mapping current state and input to next state).

The parallel composition of two Moore machines is depicted in figure 5. It is obtained connecting input and output signals with the same identifier. The outputs become outputs of the composition, while any inputs that remain unconnected to outputs become inputs of the composition. Since the result of connecting two outputs is undefined, the composition is undefined if any two outputs have the same identifier.

Definition 9 If M and M' are more machines, where Σ_o and Σ'_o are disjoint, $M'' = M \parallel M'$ is a Moore machine such that,

- i. $K'' =_{def} K \times K'$
- ii. $q''_0 =_{def} \langle q_0, q'_0 \rangle$
- iii. $\Sigma''_i =_{def} (\Sigma_i - \Sigma'_o) \cup (\Sigma'_i - \Sigma_o)$
- iv. $\Sigma''_o =_{def} \Sigma_o \cup \Sigma'_o$
- v. $\Gamma''\langle s, s' \rangle =_{def} \Gamma s \cup \Gamma' s'$

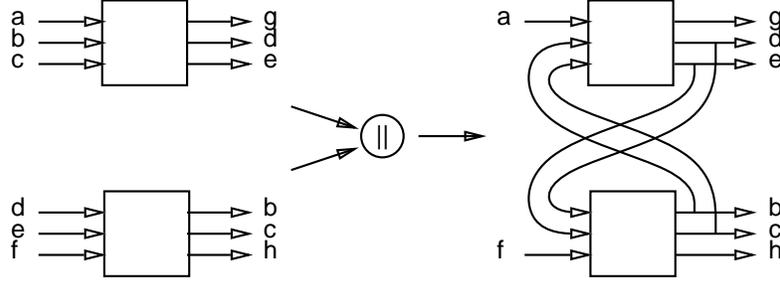


Figure 5: Illustration of Moore machine composition.

$$vi. \Theta'' \langle s, s' \rangle I =_{def} \langle \Theta s ((I \cup \Gamma' s') \downarrow \Sigma_i), \Theta' s' ((I \cup \Gamma s) \downarrow \Sigma'_i) \rangle.$$

That each of these items is of the appropriate type is easily verified. For example, if I is a mapping $\Sigma'_i \rightarrow \{0, 1\}$, then $I \cup \Gamma' s'$ is a mapping on $\Sigma_i \cup \Sigma'_i \cup \Sigma'_o$ and $(I \cup \Gamma' s') \downarrow \Sigma_i$ is a mapping on Σ_i . This expression gives the input valuation which determines the next state of M in the composition, in terms of the input to the composition and the output of M' .

Hiding of outputs of a Moore machine is accomplished using the restriction operator. Inputs can't be hidden, since Θ would cease to be a function in this case.

Definition 10 *If M is a Moore machine, and Σ is a set of identifiers containing Σ_i , then $M' = M \downarrow \Sigma$ is a Moore machine identical in all components to M except that $\Sigma'_o =_{def} \Sigma_o \cap \Sigma$ and for all $s \in K$, $\Gamma' s =_{def} \Gamma s \downarrow \Sigma$.*

We now define equivalence of Moore machines.

Definition 11 *The state equivalence relation \sim_M (written \sim where the context is unambiguous) of a Moore machine M is the unique greatest fixpoint of the functional $F_M : K \times K \rightarrow K \times K$ such that $\langle s_1, s_2 \rangle \in F_M(R)$ iff*

- i. $\langle s_1, s_2 \rangle \in R$
- ii. $\Gamma s_1 = \Gamma s_2$
- iii. $\forall I : \Sigma_i \rightarrow \{0, 1\}, \langle \Theta s_1 I, \Theta s_2 I \rangle \in R$.

Here, we have omitted the proof that F_M is monotonic and equivalence preserving.

Definition 12 *Given two Moore machines M and M' , $\Sigma_o = \Sigma'_o$, $\Sigma_i = \Sigma'_i$ and $K \cap K' = \emptyset$, their disjoint sum $M'' = M + M'$ is the Moore machine $(K \cup K', q''_0, \Sigma_i, \Sigma_o, \Gamma \cup \Gamma', \Theta \cup \Theta')$, where q''_0 is undefined.*

Definition 13 Given two Moore machines M and M' , let \sim be the state equivalence relation of $M+M'$. We say M and M' are equivalent, denoted by $M \sim M'$ iff $q_0 \sim q'_0$.

The equivalence classes of \sim on states of a Moore machine can be computed using the coarsest partitioning algorithm of [16]. Thus there is an algorithm for determining if $M \sim M'$ (i.e., for determining if the initial states of M and M' fall into the same equivalence class) in time $O((|M| + |M'|) \log(|M| + |M'|))$. Further, M can be minimized in time $O(|M| \log |M|)$.

The logic we use is the temporal logic CTL*, which subsumes both linear temporal logic (LTL) and computation tree logic (CTL). Since the semantics of this logic is defined with respect to Kripke structures [10], we associate by definition a Kripke structure K_M to each Moore machine M . This yields as a by-product a semantics for Moore machines.

Definition 14 Given a Moore machine M , let $K_M = (S, S_0, R, L)$, where

- i. S is the set of all pairs $\langle s, I \rangle$, where $s \in K$ and I is a mapping $\Sigma_i \rightarrow \{0, 1\}$ (each Kripke state is a Moore machine state coupled with a valuation of the Moore machine inputs),
- ii. S_0 is the set of all pairs $\langle q_0, I \rangle$, where I is a mapping $\Sigma_i \rightarrow \{0, 1\}$ (the initial states of the Kripke structure),
- iii. L is a map $S \rightarrow (\Sigma_i \cup \Sigma_o) \rightarrow \{0, 1\}$, where $L\langle s, I \rangle =_{def} I \cup \Gamma s$ (the atomic propositions of Kripke structure correspond to the inputs and outputs of the Moore machine),
- iv. $R \subseteq S \times S$ is such that $\langle s_1, I_1 \rangle R \langle s_2, I_2 \rangle$ iff $s_2 = \Theta s_1 I_1$ (the transition relation of the Kripke structure).

Definition 15 If M is a Moore machine and f is a CTL* formula, then $M \models f$ iff $K_M \models f$.

We now prove that the above definitions of \parallel , \downarrow , \sim and \models satisfy the preconditions of the interface rule. The first precondition is that equivalent Moore machines satisfy the same formulas in the logic. To prove this, we show that $M \sim M'$ implies $K_M \equiv K_{M'}$, where \equiv denotes strong bisimulation equivalence on Kripke structures. Kripke structures which are strongly bisimilar satisfy the same set of CTL* formulas. We first define a relation \approx on the states of K_M as follows.

Definition 16 Given a Kripke structure K_M , and two states $\langle s_1, I_1 \rangle, \langle s_2, I_2 \rangle \in S$, $(s_1, I_1) \approx (s_2, I_2)$ iff $s_1 \sim s_2$ and $I_1 = I_2$.

Lemma 2 The relation \approx is a strong bisimulation on K_M .

Proof The relation \approx is a strong bisimulation on K_M if $\langle s_1, I_1 \rangle \approx \langle s_2, I_2 \rangle$ implies

- i. $L\langle s_1, I_1 \rangle = L\langle s_2, I_2 \rangle$,
- ii. $\langle s_1, I_1 \rangle R\langle s'_1, I'_1 \rangle$ implies there exists $\langle s'_2, I'_2 \rangle$ such that $\langle s_2, I_2 \rangle R\langle s'_2, I'_2 \rangle$,
- iii. $\langle s_2, I_2 \rangle R\langle s'_2, I'_2 \rangle$ implies there exists $\langle s'_1, I'_1 \rangle$ such that $\langle s_1, I_1 \rangle R\langle s'_1, I'_1 \rangle$.

The first item is trivial, since $\langle s_1, I_1 \rangle \approx \langle s_2, I_2 \rangle$ implies $I_1 = I_2$ and $\Gamma s_1 = \Gamma s_2$. As to the second item, $\langle s_1, I_1 \rangle R\langle s'_1, I'_1 \rangle$ implies $\Theta s_1 I_1 = s'_1$ (definition 14) and $s_1 \sim s_2$ (definition 16). Let $s'_2 = \Theta s_2 I_2 = \Theta s_2 I_1$. Since \sim is a fixed point of F_M in definition 11, $s'_1 \sim s'_2$, satisfying item two. Proof of the third item is similar. \square

Corollary 2 *If $M \sim M'$ then $K_M \equiv K_{M'}$.*

Proof $M \sim M'$ implies a state equivalence \sim on $M + M'$ such that $q_0 \sim q'_0$. By lemma 2, this implies a strong bisimulation \equiv on $K_{M+M'} = K_M + K_{M'}$ such that, for all $I : \Sigma_i \rightarrow \{0, 1\}$, $\langle q_0, I \rangle \equiv \langle q'_0, I \rangle$, hence $K_M \equiv K_{M'}$. \square

The second precondition of the interface rule is that $M_1 \sim M'_1$ and $M_2 \sim M'_2$ imply $(M_1 \parallel M_2) \sim (M'_1 \parallel M'_2)$. To prove this, we define a relation \approx (different from \approx above) on the states of $(M_1 \parallel M_2) + (M'_1 \parallel M'_2)$ as follows:

Definition 17 $\langle s_1, s_2 \rangle \approx \langle s'_1, s'_2 \rangle$ iff $s_1 \sim s'_1$ and $s_2 \sim s'_2$.

We note that the relation \approx is a fixed point of $F_{(M_1 \parallel M_2) + (M'_1 \parallel M'_2)}$. Hence we have

Proposition 1 $M_1 \sim M'_1$ and $M_2 \sim M'_2$ imply $M_1 \parallel M_2 \sim M'_1 \parallel M'_2$

Proof Since the state equivalence relation \sim on $(M_1 \parallel M_2) + (M'_1 \parallel M'_2)$ is the greatest fixed point of $F_{(M_1 \parallel M_2) + (M'_1 \parallel M'_2)}$, $\langle s_1, s_2 \rangle \approx \langle s'_1, s'_2 \rangle$ implies $\langle s_1, s_2 \rangle \sim \langle s'_1, s'_2 \rangle$. By definition 17, $\langle q_{01}, q_{02} \rangle \approx \langle q'_{01}, q'_{02} \rangle$, therefore $\langle q_{01}, q_{02} \rangle \sim \langle q'_{01}, q'_{02} \rangle$, satisfying the definition of equivalence for $M_1 \parallel M_2$ and $M'_1 \parallel M'_2$. \square

We next prove that our definitions satisfy precondition three of the interface rule. Although this requires equivalence, we show a stronger condition of equality.

Lemma 3 *If M and M' are Moore machines, and Σ_M is a set of identifiers containing Σ_i and Σ'_i , then $(M \parallel M') \downarrow \Sigma_M = (M \downarrow \Sigma_M) \parallel (M' \downarrow \Sigma_M)$*

Proof We consider the Γ and Θ components. Let $P = (M \parallel M') \downarrow \Sigma_M$ and $Q = (M \downarrow \Sigma_M) \parallel (M' \downarrow \Sigma_M)$. Then we have

$$\begin{aligned}
& \Gamma_Q \langle s, s' \rangle \\
&= (\Gamma s \downarrow \Sigma_M) \cup (\Gamma' s' \downarrow \Sigma_M) \\
&= (\Gamma s \cup \Gamma' s') \downarrow \Sigma_M \\
&= \Gamma_P \langle s, s' \rangle
\end{aligned}$$

$$\begin{aligned}
& \Theta_Q \langle s, s' \rangle I \\
&= \langle \Theta s((I \cup \Gamma' s' \downarrow \Sigma_M) \downarrow \Sigma_i), \Theta' s'((I \cup \Gamma s \downarrow \Sigma_M) \downarrow \Sigma'_i) \rangle \\
&= \langle \Theta s((I \cup \Gamma' s') \downarrow \Sigma_i), \Theta' s'((I \cup \Gamma s) \downarrow \Sigma'_i) \rangle \\
&= \Theta_P \langle s, s' \rangle I
\end{aligned}$$

The other components are trivially equal. \square

The last precondition of the interface rule is satisfied, since the Kripke structure $K_{M \downarrow \Sigma}$ corresponding to $M \downarrow \Sigma$ is just K_M with the state labelling function L restricted to Σ . This Kripke structure trivially satisfies the same set of formulas on Σ as K_M .

Clarke, *et al.*, describe an application of the interface rule to reduce the complexity of verifying a system of communicating Moore machines [7]. The example is the modular controller of a simple CPU with decoupled access and execution units. By hiding some outputs of one module and applying the Hopcroft minimization algorithm to produce a reduced interface process, the number of states in the composition is reduced by roughly a factor 6 in verifying the specification of the other module.

6 Directions for Future Research

The most important open question is, of course, whether the compositional techniques described in this paper will permit verification of much more complicated finite state concurrent systems than has previously been the case. This can only be determined by further experimentation. It is clear, however, that our technique is most suitable for loosely coupled systems. When this relationship does not hold, the interface processes may be large, and we do not get a significant state reduction by using the interface theorem. Fortunately, the parallel composition of two tightly coupled processes does not seem to generate as many states as the parallel composition of two loosely coupled processes of comparable size—there are simply not as many possible interleavings. Consequently, compositional reasoning may not be as important in this case as in the case of loosely coupled processes.

It will be interesting to see how well the results of this paper apply to other process models. For systems like finite transition systems with propositional dynamic logic [13] or CCS [19] with Hennessy-Milner logic, the results should be straightforward and will be given in the full version of the paper. Our techniques should work quite well for the Caesar system of Sifakis [22]. It should also be possible to apply our ideas to Berry's Esterel [2] and Harel's Statecharts [14] if we use a logic like CTL.

Finally, the techniques that we describe in this paper also have some limitations. For example, the interface rule allows us to handle formulas that are boolean combinations of temporal properties of the individual processes in a

parallel composition. We are currently unable to handle more general properties involving temporal assertions about several processes. Developing more general rules seems like an important direction for research but also a very hard one. O. Grumberg [12] has obtained some negative results which indicate that it may be impossible to develop a fully general system of inference rules that will handle arbitrary temporal properties. Furthermore, in some cases it seems likely that it will be necessary to combine the use of the interface rule and model checker with proofs of validity for certain CTL formulas. In order to use the interface rule it may be necessary to prove an implication of the form $(\phi \wedge \psi) \rightarrow \delta$ where δ is another CTL formula that expresses a global property. We believe that in many cases it will be possible to use informal reasoning to establish such implications.

References

- [1] H. Barringer. Using temporal logic in compositional specification of concurrent systems. In *Conference on Temporal Logic and Its Applications*. Leeds University, January 1986.
- [2] G. Berry and L. Cosserat. The esterel synchronous programming language and its mathematical semantics. Technical report, Ecole Nationale Supérieure des Mines de Paris, 1984.
- [3] M. C. Browne, E. M. Clarke, and D. L. Dill. Automatic circuit verification using temporal logic: Two new examples. In *Formal Aspects of VLSI Design*. Elsevier Science Publishers, 1986.
- [4] M. C. Browne, E. M. Clarke, D. L. Dill, and B. Mishra. Automatic verification of sequential circuits using temporal logic. *IEEE Transactions on Computers*, C-35(12), December 1986.
- [5] E. M. Clarke and E. A. Emerson. Synthesis of synchronization skeletons for branching time temporal logic. In *Proceedings of the Workshop on Logic of Programs*, pages 52–71. Springer-Verlag, 1981.
- [6] E. M. Clarke, E. A. Emerson, and A. P. Sistla. Automatic verification of finite-state concurrent systems using temporal logic specifications. *ACM Transactions on Programming Languages and Systems*, 8(2):244–263, 1986.
- [7] E. M. Clarke, D. E. Long, and K. L. McMillan. A language for compositional specification and verification of finite state hardware controllers. In *Proceedings of the Conference on Hardware Description Languages*, 1989. To appear.

- [8] D. L. Dill. *Trace Theory for Automatic Heirarchical Verification of Speed-Independent Circuits*. PhD thesis, Carnegie Mellon University, Pittsburgh, PA 15213, 1987.
- [9] D. L. Dill and E. M. Clarke. Automatic verification of asynchronous circuits using temporal logic. *IEE Proceedings*, 133, part E(5), September 1986.
- [10] E. A. Emerson and J. Y. Halpern. “sometimes” and “not never” revisited: On branching versus linear time temporal logic. *Journal of the Association of Computing Machinery*, 33(1):151–178, January 1986.
- [11] E.A. Emerson and Chin Laung Lei. Modalities for model checking: Branching time strikes back. In *Conference Record of the Twelfth Annual ACM Symposium on Principles of Programming Languages*, New Orleans, La., January 1985.
- [12] O. Grumberg. Personal communication.
- [13] D. Harel. Dynamic logic. In D. Gabby and F. Guenthner, editors, *Handbook of Philosophical Logic II*, pages 498–544. Reidel, 1984.
- [14] D. Harel. Statecharts: A visual approach to complex systems. Technical Report CS84-05, The Weizmann Institute of Science, February 1984.
- [15] C. A. R. Hoare. Communicating sequential processes. *Communications of the Association of Computing Machinery*, 21(8):666–677, August 1978.
- [16] J. E. Hopcroft. An $n \log n$ algorithm for minimizing the states in a finite automaton. In *The Theory of Machines and Computation*, pages 189–196. Academic Press, New York, N.Y., 1971.
- [17] B. Josko. MCTL - an extension of CTL for modular verification of concurrent systems. In H. Barringer, editor, *Workshop on Temporal Logic*. University of Manchester, April 1987. To appear in LNCS.
- [18] O. Lichtenstein and A. Pnueli. Checking that finite state concurrent programs satisfy their linear specification. In *Conference Record of the Twelfth Annual ACM Symposium on Principles of Programming Languages*, New Orleans, La., January 1985.
- [19] R. Milner. *A Calculus of Communicating Systems*, volume 92 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin, 1980.
- [20] B. Mishra and E.M. Clarke. Hierarchical verification of asynchronous circuits using temporal logic. *Theoretical Computer Science*, 38:269–291, 1985.
- [21] A. Pnueli. In transition for global to modular temporal reasoning about programs. In K. R. Apt, editor, *Logics and Models of Concurrent Systems*, pages 123–144. Springer-Verlag, 1984.

- [22] J. P. Quielle and J. Sifakis. Specification and verification of concurrent systems in CESAR. In *Proceedings of the Fifth International Symposium in Programming*, 1981.
- [23] M. Vardi and P. Wolper. An automata-theoretic approach to automatic program verification. In *Proceedings of the Conference on Logic in Computer Science*, Boston, Mass., June 1986.