

Interactive Verification of Concurrent Systems using Symbolic Execution

Michael Balsler, Simon Bäumler, Wolfgang Reif, and Gerhard Schellhorn
University of Augsburg, Germany

Abstract

This paper presents an interactive proof method for the verification of temporal properties of concurrent systems based on symbolic execution. Symbolic execution is a well known and very intuitive strategy for the verification of sequential programs. We have carried over this approach to the interactive verification of arbitrary linear temporal logic properties of (infinite state) parallel programs. The resulting proof method is very intuitive to apply and can be automated to a large extent. It smoothly combines first order reasoning with reasoning in temporal logic. The proof method has been implemented in the interactive verification environment KIV and has been used in several case studies.

1 Introduction

Compared to sequential programs, both the design and the verification of concurrent systems is more difficult, mainly because the control flow is more complex. Particularly, for reactive systems, not only the final result of execution but the sequence of output over time is relevant for system behavior. Finding errors by means of testing strategies is limited, because, especially for interleaved systems, an exponential amount of possible executions must be considered. The execution is nondeterministic, making it difficult to reproduce errors. An alternative to testing is the use of formal methods to specify and verify concurrent systems with mathematical rigor. Automatic methods – especially model checking – have been applied successfully to discover flaws in the design and implementation of systems. Starting from systems with finite state spaces of manageable size, research in model checking aims at mastering ever more complex state spaces and to reduce infinite state systems by abstraction. In general, systems must be manually abstracted to ensure that formal analysis terminates. An alternative approach to large or infinite state spaces are interactive proof calculi. They directly address the problem of infinite state spaces. Here, the challenge is to achieve a high degree of automation. Existing interactive calculi to reason in temporal logic about concurrent systems are generally difficult to apply. The strategy of symbolic execution, on the other hand, has been successfully applied to the interactive verification of sequential programs (e.g. Dynamic Logic [9, 11]). Symbolic execution gives intuitive proofs with a high degree of automation.

Combining Dynamic Logic and temporal logic has already been investigated, e.g. Process Logic [17] and Concurrent Dynamic Logic [16] and more recently [8, 10]. These works focus on combinations of logic, while we are interested in an interactive proof method based on symbolic execution. Symbolic execution of parallel programs has been investigated in [1], however, this approach has been restricted to the verification of pre/post conditions. Other approaches are often restricted to the verification of certain types of temporal properties. Our approach presents an interactive proof method to verify arbitrary temporal properties for parallel programs with the strategy of symbolic execution and thus promises to be intuitive and highly automatic.

Our logic is based on Interval Temporal Logic (ITL) [15]. The chop operator $\varphi; \psi$ of ITL corresponds to sequential composition and programs are just a special case of temporal formulas. In addition, we have defined an interleaving operator $\varphi || \psi$ to interleave arbitrary temporal formulas. Our logic explicitly considers arbitrary environment steps after each system transition, which is similar to Temporal Logic of Actions (TLA) [13]. This ensures that systems are compositional and proofs can be decomposed. In total, we have defined a compositional logic which includes a rich programming language with interleaved parallel processes similar to the one of STeP [7]. Important for interactive proofs, system descriptions need not be translated to flat transition systems. The calculus directly reasons about programs.

A short overview of our logic is given in Section 2. The calculus for symbolic execution is described in Section 3. The strategy has been implemented in KIV (see Section 4). Section 6 concludes. For more details on the logic and calculus, especially on induction and double primed variables to decompose proofs, we refer to [2].

2 Logic

Similar to [15], we have defined a first order interval temporal logic with static variables a , dynamic variables A , functions f , and predicates p . Let v be a static (i.e. constants - written in small letters) or dynamic variable. Then, the syntax of (a subset of) our logic is defined

$$\begin{aligned}
 e & ::= a \mid A \mid A' \mid A'' \mid f(e_1, \dots, e_n) \\
 \varphi & ::= p(e_1, \dots, e_n) \mid \neg \varphi \mid \varphi_1 \wedge \varphi_2 \mid \exists v. \varphi \\
 & \quad \mid \mathbf{step} \mid \varphi_1; \varphi_2 \mid \varphi^* \\
 & \quad \mid [A_1, \dots, A_n] \mid \varphi_1 \parallel^< \varphi_2
 \end{aligned}$$

Dynamic variables can be primed and double primed. It is possible to quantify both static and dynamic variables. The chop operator $\varphi_1; \varphi_2$ directly corresponds to the sequential composition of programs. The star operator φ^* is similar to a loop. We have added an operator $[A_1, \dots, A_n]$ to define an explicit frame assumption. Furthermore, operator $\varphi_1 \parallel^< \varphi_2$ can be used to interleave two “processes”. The basic operator $\parallel^<$ gives precedence to the left process, i.e., a transition of φ_1 is executed first.

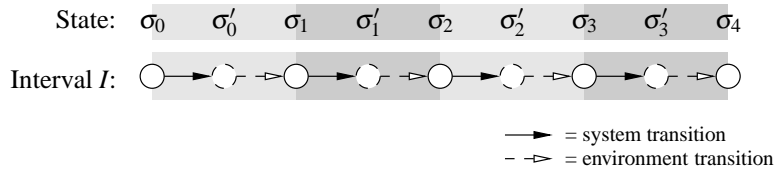


Figure 1: An interval as sequence of states

In ITL, an interval is considered to be a finite or infinite sequence of states, where a state is a mapping from variables to their values. In our setting, we introduce additional intermediate states σ'_i to distinguish between system and environment transitions. For intuition, compare the following to Figure 1. Let $\bar{n} \in \mathbb{N}^\infty$. An *interval*

$$I = (\sigma_0, \sigma'_0, \sigma_1, \dots, \sigma'_{\bar{n}-1}, \sigma_{\bar{n}})$$

consists of an initial state σ_0 , and a finite or infinite and possibly empty sequence of transitions $(\sigma'_i, \sigma_{i+1})_{i=0}^{\bar{n}-1}$. In the intermediate states σ'_i the values of the variables after a system transition are stored. The following states σ_{i+1} reflect the states after an environment transition. In this manner, system and environment transitions alternate. An empty interval consists only of an initial state σ_0 .

Given an interval I , variables are evaluated as follows:

$$\begin{aligned}
 \llbracket v \rrbracket_I & ::= \sigma_0(v) \\
 \llbracket A' \rrbracket_I & ::= \begin{cases} \sigma'_0(A) & \text{if } |I| > 0 \\ \sigma_0(A) & \text{otherwise} \end{cases} \\
 \llbracket A'' \rrbracket_I & ::= \begin{cases} \sigma_1(A) & \text{if } |I| > 0 \\ \sigma_0(A) & \text{otherwise} \end{cases}
 \end{aligned}$$

Static and unprimed dynamic variables are evaluated in the initial state. Primed variables are evaluated in σ'_0 and double primed variables in σ_1 . In the last state, i.e. if I is empty, the value of a primed or

double primed variable is equal to the unprimed variable. It is assumed that after a system has terminated, the variables do not change.

The semantics of the standard ITL operators can be carried over one-to-one to our notion of an interval, and is therefore omitted here. In [15], however, assignments only restrict the assigned variables, whereas program assignments leave all other dynamic variables unchanged. This is known as a frame assumption. In our logic, we have defined an explicit frame assumption $[A_1, \dots, A_n]$ which states that a system transition leaves all but a selection of dynamic variables unchanged.

$$I \models [A_1, \dots, A_n] \quad \text{iff} \quad \sigma'_0(A) = \sigma_0(A) \text{ for all } A \notin \{A_1, \dots, A_n\}$$

Details on frame assumptions can be found in [2].

The interleaving operator $\varphi \parallel^< \psi$ interleaves arbitrary formulas φ and ψ . The two formulas represent sets of intervals. We have therefore defined the semantics of $\varphi \parallel^< \psi$ relative to the interleaving $\llbracket I_1 \parallel^< I_2 \rrbracket$ of two concrete intervals I_1 and I_2 .

$$I \models \varphi \parallel^< \psi \quad \text{iff} \quad \begin{array}{l} \text{there exist } I_1, I_2 \\ \text{with } I \in \llbracket I_1 \parallel^< I_2 \rrbracket \text{ and } I_1 \models \varphi \text{ and } I_2 \models \psi \end{array}$$

For nonempty intervals $I_1 = (\sigma_0, \sigma'_0, \sigma_1, \dots)$, and $I_2 = (\tau_0, \tau'_0, \tau_1, \dots)$, interleaving of intervals adheres to the following recursive equations.

$$\begin{aligned} \llbracket I_1 \parallel^< I_2 \rrbracket &= \begin{cases} (\sigma_0, \sigma'_0) \oplus \llbracket (\sigma_1, \dots) \parallel I_2 \rrbracket, & \text{if } I_1 \text{ is not blocked} \\ (\tau_0, \tau'_0) \oplus \llbracket (\sigma_1, \dots) \parallel (\tau_1, \dots) \rrbracket, & \text{if } I_1 \text{ is blocked, } \sigma_0 = \tau_0 \\ \emptyset, & \text{otherwise} \end{cases} \\ \llbracket I_1 \parallel I_2 \rrbracket &= \llbracket I_1 \parallel^< I_2 \rrbracket \cup \llbracket I_2 \parallel^< I_1 \rrbracket \end{aligned}$$

Interval I_1 is blocked, if $\sigma'_0(\text{blk}) \neq \sigma_0(\text{blk})$ for a special dynamic variable blk . The system transition toggles the variable to signal that the process is currently blocked (see **await** statement below). If I_1 is not blocked, then the first transition of I_1 is executed and the system continues with interleaving the remaining interval with I_2 . (Function \oplus prefixes all of the intervals of a given set with the two additional states.) If I_1 is blocked, then a transition of I_2 is executed instead. However, the blocked transition of I_1 is also consumed. A detailed definition of the semantics can be found in [2].

Additional common logical operators can be defined as abbreviations. A list of frequently used abbreviations is contained in Table 1, where most of the abbreviations are common in ITL. The next operator comes in two flavors. The strong next $\circ \varphi$ requires that there is a next step satisfying φ , the weak next $\bullet \varphi$ only states that if there is a next step, it must satisfy φ .

As can be seen in Table 1, the standard constructs for sequential programs can be derived. Executing an assignment requires exactly one step. The value of e is “assigned” to the primed value of A . Other variables are unchanged ($[A]$). Note that for conditionals and loops, the condition ψ evaluates in a single step. For local variable definitions, the local variable is quantified ($\exists A$), in addition, the environment cannot access the local variable ($\square A'' = A'$). The global value of the variable is unchanged ($\square A' = A$).

Parallel programs communicate using shared variables. In order to synchronize execution, the operator **await** ψ **do** φ can be used. A special dynamic variable blk is used to mark whether a parallel program is blocked. The **await** operator behaves like a loop waiting for the condition to be satisfied. While the operator waits, no variable is changed except for variable blk which is toggled. In other words, a process guarded with an **await** operator actively waits for the environment to satisfy its condition. Immediately after condition ψ is satisfied, construct φ is executed.

more \equiv step ; true	$\diamond \varphi$ \equiv finite ; φ
last \equiv \neg more	$\square \varphi$ \equiv $\neg \diamond \neg \varphi$
inf \equiv true; false	$\circ \varphi$ \equiv step ; φ
finite \equiv \neg inf	$\bullet \varphi$ \equiv $\neg \circ \neg \varphi$
$A := e$ \equiv $A' = e \wedge [A] \wedge$ step	
skip \equiv $\square \wedge$ step	
if ψ then φ_1 else φ_2 \equiv $\psi \wedge (\mathbf{skip}; \varphi_1) \vee \neg \psi \wedge (\mathbf{skip}; \varphi_2)$	
while ψ do φ \equiv $((\psi \wedge (\mathbf{skip}; \varphi))^* \wedge \square (\mathbf{last} \rightarrow \neg \psi)); \mathbf{skip}$	
var A in φ \equiv $\square A' = A \wedge \exists A. \varphi \wedge \square A'' = A'$	
bskip \equiv $blk' \neq blk \wedge [blk] \wedge$ step	
await φ do ψ \equiv $((\neg \varphi \wedge \mathbf{bskip})^* \wedge \square (\mathbf{last} \rightarrow \varphi)); \psi$	
await φ \equiv await φ do skip	
$\varphi \parallel \psi$ \equiv $\varphi \parallel^< \psi \vee \psi \parallel^< \varphi$	

Table 1: Frequently used temporal abbreviations

3 Calculus

Our proof method is based on a sequent calculus with calculus rules of the following form:

$$\frac{\Gamma_1 \vdash \Delta_1 \quad \dots \quad \Gamma_n \vdash \Delta_n}{\Gamma \vdash \Delta}$$

Rules are applied bottom-up. Rule *name* refines a given conclusion $\Gamma \vdash \Delta$ with n premises $\Gamma_i \vdash \Delta_i$. We also rely heavily on rewriting with theorems of the form $\varphi \leftrightarrow \psi$. These allow to replace instances of φ with instances of ψ anywhere in a sequent.

3.1 Normal form

Our proof strategy is symbolic execution of temporal formulas, parallel programs being just a special case thereof. In Dynamic Logic, the leading assignment of a DL formula $\langle v := e; \alpha \rangle \varphi$ is executed as follows:

$$\frac{\frac{\Gamma_v^{v_0}, v = e_v^{v_0} \vdash \langle \alpha \rangle \varphi}{\Gamma \vdash \langle v := e \rangle \langle \alpha \rangle \varphi} \text{ asg } r}{\Gamma \vdash \langle v := e; \alpha \rangle \varphi} \text{ normalize}$$

The sequential composition is normalized and is replaced with a succession of diamond operators. Afterwards, the assignment is “executed” to compute the strongest postcondition $\Gamma_v^{v_0}, v = e_v^{v_0}$ (The substitution $\Gamma_v^{v_0}$ means that variable v is replaced with variable v_0 in Γ). In our setting, we normalize all the temporal formulas of a given sequent by rewriting the formulas to a so called normal form which separates the possible first transitions and the corresponding temporal formulas describing the system in the next state. Afterwards, an overall step for the whole sequent is executed (see below). More formally, a program (or temporal formula) is rewritten to a formula of the following type

$$\tau \wedge \circ \varphi$$

where τ is a predicate logic formula that describes a transition as a relation between unprimed, primed and double primed variables while φ describes what happens in the rest of the interval. A program may also terminate, i.e., under certain conditions, the current state may be the last. Furthermore, the

$$\begin{array}{c}
\frac{\varphi, \Gamma \vdash \Delta \quad \psi, \Gamma \vdash \Delta}{\varphi \vee \psi, \Gamma \vdash \Delta} \text{ dis } l \quad \frac{\varphi_a^{a_0}, \Gamma \vdash \Delta}{\exists a. \varphi, \Gamma \vdash \Delta} \text{ ex } l (a_0 \notin \text{free}(\varphi) \setminus \{a\} \cup \text{free}(\Gamma, \Delta)) \\
\frac{\tau_{A,A',A''}^{a,a,a} \vdash}{\tau, \mathbf{last} \vdash} \text{ lst } (a \notin \text{free}(\tau)) \quad \frac{\tau_{A,A',A''}^{a_1,a_2,A}, \varphi}{\tau, \circ \varphi \vdash} \text{ stp } (a_1, a_2 \notin \text{free}(\tau, \varphi))
\end{array}$$

Table 2: Rules for executing an overall step

$$\begin{array}{l}
\text{alw:} \quad \square \varphi \quad \leftrightarrow \quad \varphi \wedge \bullet \square \varphi \\
\text{ev:} \quad \diamond \varphi \quad \leftrightarrow \quad \varphi \vee \circ \diamond \varphi
\end{array}$$

Table 3: Rules for executing temporal logic operators \square and \diamond

next transition can be nondeterministic, i.e., different τ_i with corresponding φ_i may exist describing the possible transitions and corresponding next steps. Finally, there may exist a link between the transition τ_i and system φ_i which cannot be expressed as a relation between unprimed, primed, and double primed variables in the transition alone. This link is captured in existentially quantified static variables a_i which occur in both τ_i and φ_i . The general pattern to separate the first transitions of a given temporal formula is

$$\tau_0 \wedge \mathbf{last} \vee \bigvee_{i=1}^n (\exists a_i. \tau_i \wedge \circ \varphi_i).$$

We will refer to this general pattern as normal form.

3.2 Executing an overall step

Assume that the antecedent of a sequent has been rewritten to normal form. Further assume – to keep it simple – that the succedent is empty. (This can be assumed as formulas in the succedent are equivalent to negated formulas in the antecedent. Furthermore, several formulas in the antecedent can be combined to a single normal form.)

$$\tau_0 \wedge \mathbf{last} \vee \bigvee_{i=1}^n (\exists a_i. \tau_i \wedge \circ \varphi_i) \vdash$$

With the two rules *dis l* and *ex l* of Table 2, disjunction and quantification can be eliminated. For the remaining premises,

$$\tau_0 \wedge \mathbf{last} \vdash \quad \tau_i \wedge \circ \varphi_i \vdash$$

the two rules *lst* and *stp* can be applied. If execution terminates, all free dynamic variables A – no matter, if they are unprimed, primed or double primed – are replaced by fresh static variables a . The result is a formula in pure predicate logic with static variables only, which can be proven with standard first order reasoning. Rule *stp* advances a step in the trace. The values of the dynamic variables A and A' in the old state are stored in fresh static variables a_1 and a_2 . Double primed variables are unprimed variables in the next state. Finally, the leading next operators are discarded. The proof method now continues with the execution of φ_i .

3.3 Executing temporal logic

The idea of symbolic execution can be applied to formulas of temporal logic. For example, operator $\square \varphi$ is similar to a loop in a programming language: formula φ is executed in every step. An appropriate rewrite rule is *alw* of Table 3. Formula φ must hold now, and in the next step $\square \varphi$ holds again. To arrive at a formula in normal form, the first conjunct of the resulting formula $\varphi \wedge \bullet \square \varphi$ must be further rewritten.

$$\frac{\vdash \varphi_1 \rightarrow \varphi_2}{\vdash \varphi_1; \psi \rightarrow \varphi_2; \psi} \text{ chp lem}$$

<i>chp dis</i> :	$(\varphi_1 \vee \varphi_2); \psi$	\leftrightarrow	$\varphi_1; \psi \vee \varphi_2; \psi$
<i>chp ex</i> :	$(\exists a. \varphi); \psi$	\leftrightarrow	$\exists a_0. \varphi_a^{a_0}; \psi$
<i>chp lst</i> :	$(\tau \wedge \mathbf{last}); \psi$	\leftrightarrow	$\tau_{A', A''}^{A, A} \wedge \psi$
<i>chp stp</i> :	$(\tau \wedge \circ \varphi); \psi$	\leftrightarrow	$\tau \wedge \circ (\varphi; \psi)$

Table 4: Rules for executing sequential composition

The rewrite rule above corresponds to the recursive definition of $\square \varphi$. Other temporal operators can be executed similarly.

3.4 Executing sequential composition

The execution of sequential composition of programs $\varphi; \psi$ is more complicated as we cannot give a simple equivalence which rewrites a composition to normal form. The problem is that the first formula φ could take an unknown number of steps to execute. Only after φ has terminated, we continue with executing ψ .

Rules for the execution of $\varphi; \psi$ are given in Table 4. In order to execute composition, the idea is to first rewrite formula φ to normal form

$$\left(\tau_0 \wedge \mathbf{last} \vee \bigvee (\exists a_i. \tau_i \wedge \circ \varphi_i) \right); \psi$$

The first sub-formula φ can be rewritten with rule *chp lem* of Table 4. If it is valid that φ_1 implies φ_2 , then $\varphi_1; \psi$ also implies $\varphi_2; \psi$. (This rule is a so-called congruence rule.) After rewriting the first sub-formula to normal form, we rewrite the composition operator. According to rules *chp dis* and *chp ex*, composition distributes over disjunction and existential quantification. If we apply these rules to the formula above, we receive a number of cases

$$(\tau_0 \wedge \mathbf{last}); \psi \vee \bigvee \exists a_{i,0}. (\tau_{i,0} \wedge \circ \varphi_{i,0}); \psi$$

In the first case, program φ terminates, in the other cases, the program takes a step τ_i and continues with program φ_i . Rules *chp lst* and *chp stp* can be used to further rewrite the composition. Dynamic variables A stutter in the last step, and therefore the primed and double primed variables A' and A'' of τ are replaced by the corresponding unprimed variables if the first sub-formula terminates. The two rules give

$$\tau_{0, A', A''}^{A, A} \wedge \psi \vee \bigvee \exists a_{i,0}. \tau_{i,0} \wedge \circ (\varphi_{i,0}; \psi)$$

In the first case, φ has terminated and we still need to execute ψ to arrive with a formula in normal form. In the other cases, we have successfully separated the formula into the first transition $\tau_{i,0}$ and the corresponding rest of the program $\varphi_{i,0}; \psi$.

3.5 Interleaving

As with sequential composition, the interleaving of programs cannot be executed directly, but the sub-formulas need to be rewritten to normal form first. The basic operator for interleaving is the left interleaving operator $\varphi \parallel^< \psi$ which gives precedence to the left process. In order to execute $\parallel^<$, the first sub-formula must be rewritten to normal form before the operator itself can be rewritten.

$$\frac{\vdash \varphi_1 \rightarrow \varphi_2}{\vdash \varphi_1 \parallel^< \psi \rightarrow \varphi_2 \parallel^< \psi} \text{ ilvl lem}$$

$$\begin{aligned} \text{ilvl dis:} \quad & (\varphi_1 \vee \varphi_2) \parallel^< \psi \leftrightarrow \varphi_1 \parallel^< \psi \vee \varphi_2 \parallel^< \psi \\ \text{ilvl ex:} \quad & (\exists a. \varphi) \parallel^< \psi \leftrightarrow \exists a_0. \varphi_a^{a_0} \parallel^< \psi \\ & a_0 \text{ fresh with respect to } (\text{free}(\varphi) \setminus \{a\}) \cup \text{free}(\psi) \\ \text{ilvl lst:} \quad & (\tau \wedge \mathbf{last}) \parallel^< \psi \leftrightarrow \tau_{A',A''}^{A,A} \wedge \psi \\ \text{ilvl stp:} \quad & (\tau \wedge \neg \mathbf{blocked} \wedge \circ \varphi) \parallel^< \psi \\ & \leftrightarrow \exists a_2. (\tau_{A''}^{a_2} \wedge \neg \mathbf{blocked} \wedge \circ ((A = a_2 \wedge \varphi) \parallel \psi)) \\ \text{ilvl blk:} \quad & (\tau \wedge \mathbf{blocked} \wedge \circ \varphi) \parallel^< \psi \\ & \leftrightarrow \exists a_{2,1}. (\exists a_1. \tau_{A',A''}^{a_1,a_{2,1}}) \wedge (A = a_{2,1} \wedge \varphi) \parallel_b^< \psi \end{aligned}$$

Table 5: Rules for executing left interleaving

For left interleaving, the rules of Table 5 are similar to the rules for sequential composition. Congruence rule *ilvl lem* makes it possible to rewrite the first sub-formula to normal form. Similar to *chop*, left interleaving also distributes over disjunction (*ilvl dis*) and existential quantifiers (*ilvl ex*). If the first formula terminates, execution continues with the second (*ilvl lst*). This is similar to rule *chp lst*. Otherwise, execution depends on the first process being blocked. If it is not blocked, rule *ilvl stp* executes the transition and continues with interleaving the remaining φ with ψ . Note that the double primed variables of τ are replaced by static variables a_2 which must be equal to the unprimed variables the next time a transition of the first process is executed. This is to establish the environment transition of the first process as a relation which also includes transitions of the second. If the first process is blocked, then rule *ilvl blk* executes the blocked transition; the process actively waits while being blocked. However, the primed variables of τ are replaced by static variables a_1 : the blocked transition of the first process does not contribute to the transition of the overall interleaving. Instead, a transition of the other process is executed.

$$\frac{\vdash \varphi_1 \rightarrow \varphi_2}{\vdash \psi \parallel_b^< \varphi_1 \rightarrow \psi \parallel_b^< \varphi_2} \text{ ilvlb lem}$$

$$\begin{aligned} \text{ilvlb dis:} \quad & \psi \parallel_b^< (\varphi_1 \vee \varphi_2) \leftrightarrow \psi \parallel_b^< \varphi_1 \vee \psi \parallel_b^< \varphi_2 \\ \text{ilvlb ex:} \quad & \psi \parallel_b^< (\exists a. \varphi) \leftrightarrow \exists a_0. \psi \parallel_b^< \varphi_a^{a_0} \\ & a_0 \text{ fresh with respect to } (\text{free}(\varphi) \setminus \{a\}) \cup \text{free}(\psi) \\ \text{ilvlb lst:} \quad & \psi \parallel_b^< (\tau \wedge \mathbf{last}) \leftrightarrow \text{false} \\ \text{ilvlb stp:} \quad & \psi \parallel_b^< (\tau \wedge \circ \varphi) \leftrightarrow \exists a_{2,2}. \tau_{A''}^{a_{2,2}} \wedge \circ (\psi \parallel (A = a_{2,2} \wedge \varphi)) \end{aligned}$$

Table 6: Rules for executing blocked left interleaving

The situation where the first process is blocked and it remains to execute a transition of the second process is represented by a derived operator $\varphi \parallel_b^< \psi$ which is defined

$$\varphi \parallel_b^< \psi \quad ::= \quad (\mathbf{blocked} \wedge \circ \varphi) \parallel^< \psi$$

and for which the rules of Table 6 are applicable. Again, the rules are very similar to the rules above. A congruence rule *ilvlb lem* ensures that the second sub-formula can be rewritten to normal form. With rules *ilvlb dis* and *ilvlb ex*, the operator distributes over disjunction and existential quantification. If the second process terminates, *ilvlb lst* is applicable, otherwise *ilvlb stp* can be applied.

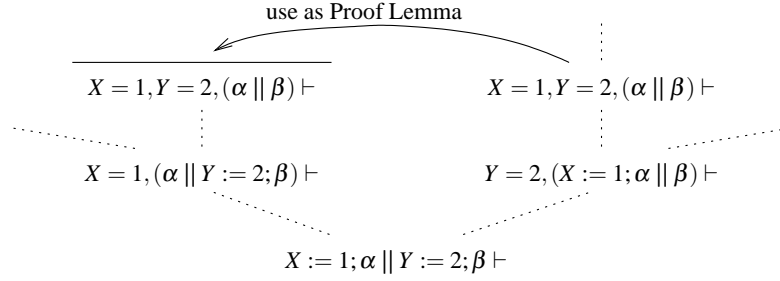


Figure 2: Example of proof lemma rule

Similar rules have been defined for all the operators of our logic [2]. In summary, every temporal formula can be rewritten to normal form, an overall step can be executed, and the process of symbolic execution can be repeated.

3.6 Induction and Proof Lemmas

To prove programs with loops, an inductive argument is necessary. Just as for sequential programs, we use well-founded and structural induction over datatypes. A specific rule for temporal induction over the interval is unnecessary: if a liveness property $\diamond \varphi$ is known, the equivalence

$$\diamond \varphi \leftrightarrow \exists N. N'' = N - 1 \text{ until } (N = 0 \wedge \varphi)$$

can be used to induce over the number of steps N it takes to reach the first state satisfying φ . To prove a safety property such a liveness property can be derived by using the equivalence $\square \varphi \leftrightarrow \neg \diamond \neg \varphi$. The proof is then by contradiction, assuming there is a number N of steps, after which φ is violated.

Execution of interleaved programs often leads to the same sequent occurring in different branches of the proof tree. Normally, the user would specify a lemma which can be applied to all these sequents. To simplify this, we allow that a premise of the proof tree can automatically be used as lemma to close another goal. An example is shown in Figure 2. This generalisation of proof trees to (acyclic) proof graphs allows the dynamic construction of verification diagrams [18].

4 Implementation

The interactive proof method has been implemented in KIV [4], an interactive theorem prover which is available at [12]. KIV supports algebraic specifications, predicate logic, dynamic logic, and higher order logic. Especially, reasoning in predicate logic and dynamic logic is very elaborate. Support for concurrent systems and temporal logic has been added. Considerable effort has been spent to ensure that the calculus rules are automatically applied to a large extent. Almost all of the rules are invertible ensuring that, if the conclusion is provable, the resulting premises remain valid. The overall strategy is

- to symbolically execute a given proof obligation,
- to simplify the PL formulas describing the current state,
- to combine premises with the same sequent, and
- to use induction, if a system loop has been executed.

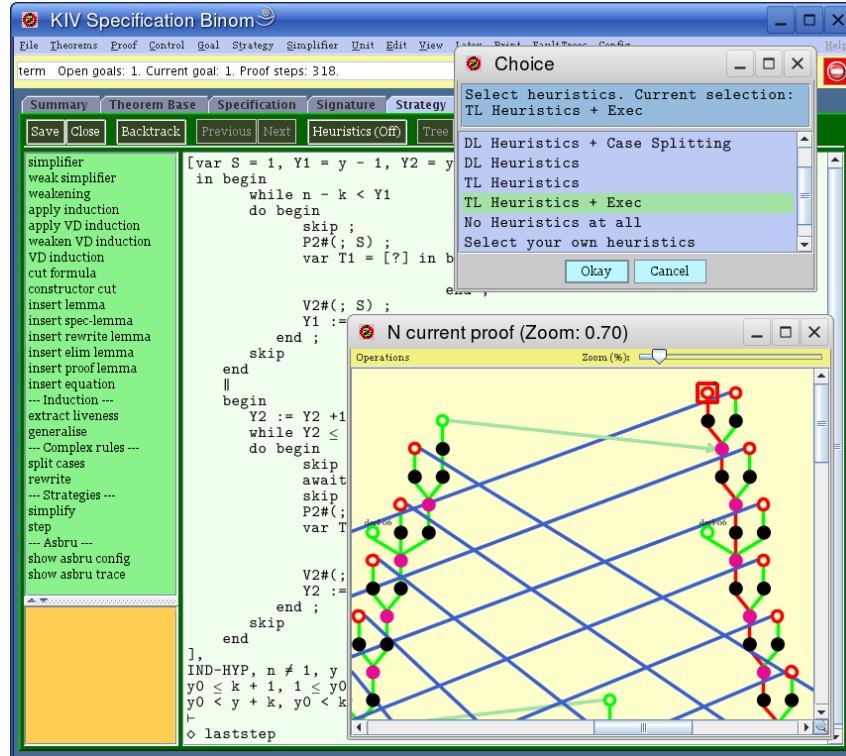


Figure 3: Verification in KIV

5 Verification in KIV

Figure 3 contains a screen-shot to illustrate how to construct proofs in KIV. In the lower right, the proof graph is displayed. The main area in the large window contains the sequent under verification with the partially executed parallel program and the first order formulas describing the current state. A list of applicable proof rules is displayed to the left. Heuristics are used to automatically apply these rules. Two set of heuristics TL Heuristics and TL Heuristics + Exec implement the overall strategy of Section 4.

6 Conclusion

We have successfully carried over the strategy of symbolic execution to verify arbitrary temporal properties for concurrent systems. The proof method is based on symbolic execution, sequencing, and induction. How to symbolically execute arbitrary temporal formulas – parallel programs being just a special case thereof – has been explained in this paper.

Our proof method is easily extendable to other temporal operators. For every operator, a set of rules must be provided to rewrite the operator to normal form. With rules similar to the ones of Tables 4 and 5, we support in KIV operators for Dijkstra's choice, synchronous parallel execution, and interrupts. Furthermore, we have integrated STATEMATE state charts [3] as well as UML state charts [5] as alternative formalisms to define concurrent systems. For all of our extensions, the strategy of sequencing and induction has remained unchanged and arbitrary temporal formulas can be verified.

Using double primed variables, we have defined a compositional semantics for every operator in-

cluding interleaving of formulas. This allowed us to find a suitable assumption-guarantee theorem to apply compositional reasoning [6]. This technique is very important to verify large case studies.

The implementation in KIV has shown that symbolic execution can be automated to a large extent. We have applied the strategy to small and medium size case studies. Currently, the strategy is applied in a European project called Protocure to verify medical guidelines which can be seen as yet another form of concurrent system [14]. Overall, we believe that the strategy has the potential to make interactive proofs in (linear) temporal logic in general more intuitive and automatic.

References

- [1] K.R. Apt and E.-R. Olderog. *Verification of Sequential and Concurrent Programs*. Springer-Verlag, 1991.
- [2] M. Balsler. *Verifying Concurrent System with Symbolic Execution – Temporal Reasoning is Symbolic Execution with a Little Induction*. PhD thesis, University of Augsburg, Augsburg, Germany, 2005.
- [3] M. Balsler and A. Thums. Interactive verification of statecharts. In *Integration of Software Specification Techniques (INT’02)*, 2002. <http://tfs.cs.tu-berlin.de/~mgr/int02/proceedings.html>.
- [4] Michael Balsler, Wolfgang Reif, Gerhard Schellhorn, Kurt Stenzel, and Andreas Thums. Formal system development with KIV. In T. Maibaum, editor, *Fundamental Approaches to Software Engineering*. Springer LNCS 1783, 2000.
- [5] S. Bäumlner, M. Balsler, A. Knapp, W. Reif, and A. Thums. Interactive verification of uml state machines. In *Formal Methods and Software Engineering*, number 3308 in LNCS. Springer, 2004.
- [6] Simon Bäumlner, Florian Nafz, Michael Balsler, and Wolfgang Reif. Compositional proofs with symbolic execution. In Bernhard Beckert and Gerwin Klein, editors, *Proceedings of the 5th International Verification Workshop*, volume 372 of *Ceur Workshop Proceedings*, 2008.
- [7] N. Bjørner, A. Brown, M. Colon, B. Finkbeiner, Z. Manna, H. Sipma, and T. Uribe. Verifying temporal properties of reactive systems: A step tutorial. In *Formal Methods in System Design*, pages 227–270, 2000.
- [8] D. Harel and D. Peleg. Process logic with regular formulas. *Theoretical Computer Science*, 38:307–322, 1985.
- [9] David Harel. Dynamic logic. In D. Gabbay and F. Guenther, editors, *Handbook of Philosophical Logic*, volume 2, pages 496–604. Reidel, 1984.
- [10] David Harel and Eli Singerman. Computation path logic: An expressive, yet elementary, process logic. *Annals of Pure and Applied Logic*, pages 167–186, 1999.
- [11] M. Heisel, W. Reif, and W. Stephan. A Dynamic Logic for Program Verification. In A. Meyer and M. Taitslin, editors, *Logical Foundations of Computer Science*, LNCS 363, pages 134–145, Berlin, 1989. Logic at Botik, Pereslavl-Zalessky, Russia, Springer.
- [12] KIV homepage. <http://www.informatik.uni-augsburg.de/swt/kiv>.
- [13] L. Lamport. The temporal logic of actions. *ACM Transactions on Programming Languages and Systems*, 16(3):872–923, May 1994.
- [14] M. Balsler, J. Schmitt, and W. Reif. Verification of medical guidelines with KIV. In *Proceedings of Workshop on AI techniques in healthcare: evidence-based guidelines and protocols (ECAI’06)*, 2006.
- [15] B. Moszkowski. A temporal logic for multilevel reasoning about hardware. *IEEE Computer*, 18(2):10–19, 1985.
- [16] D. Peleg. Concurrent dynamic logic. *J. ACM*, 34(2):450–479, 1987.
- [17] R. Sherman, A. Pnueli, and D. Harel. Is the interesting part of process logic uninteresting?: a translation from pl to pdl. In *POPL ’82: Proceedings of the 9th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 347–360, New York, NY, USA, 1982. ACM Press.
- [18] Henny Sipma. *Diagram-based verification of reactive, real-time and hybrid systems*. PhD thesis, Stanford University, 1999.