

TRANSITION LOGIC
how to reason about temporal properties
of programs in a compositional way

Rob Gerth

RUU-CS-83-17

November 1983

Februari 1984



Rijksuniversiteit Utrecht

Vakgroep informatica

Budapestlaan 6 3584 CD Utrecht
Corr. adres: Postbus 80.012 3508 TA Utrecht
Telefoon 030-53 1454
The Netherlands

TRANSITION LOGIC

how to reason about temporal properties
of programs in a compositional way

Rob Gerth

Technical Report-RUU-CS-83-17

November 1983

Februari 1984

Dspartment of Computer Science

University of Utrecht

P.O. Box 80.012

3508 TA Utrecht

the Netherlands

TRANSITION LOGIC

how to reason about temporal
properties in a compositional way

ROB GERTH

Department of Computer Science, University of Utrecht
P.O. Box 80.012, 3508TA Utrecht, the Netherlands

ABSTRACT.

This paper addresses the problem of obtaining formal proof systems that support reasoning about temporal properties of parallel programs in a way that is *compositional* or 'syntax directed' - the distinctive feature of e.g. Hoare style proof systems. Now, temporal properties of programs express properties about execution traces of these programs. In the presence of concurrency, compositionality is obtained by concentrating exclusively on the execution traces associated with the atomic actions of programs - the so-called *transitions* of those programs. To reason about such transitions in a compositional way, 'Transition Logic' is proposed, expressing properties of the form "every transition of a program, say α , that starts in a state satisfying assertion p , must end in a state validating q ": $[p]\alpha[q]$. Essential is that these assertions can express properties of control locations of programs, too. For this logic, a compositional proof system is obtained which is proved to be sound and relatively complete. An interesting feature of the proof system is the axiomatization of the flow-of-control of programs. The relevance of this logic is supported by the fact that the temporal behaviour of a program is ultimately provable in terms of properties of its transitions, as shown by the work of Z. Manna and A. Pnueli [16].

TO APPEAR AT THE 16th ACM STOC

1. INTRODUCTION.

The motivation for this research originates from the principle that the development of a program and the verification that it meets its specification should proceed simultaneously and should guide each other.

This paradigm has found forceful proponents in, e.g., E.W. Dijkstra [6] and D. Gries [8]. Whether it is a practical one is debatable. However, the principle is of obvious methodological value and it makes sense to study its consequences for verification techniques.

What are those consequences? Firstly, as the verification effort should guide program development, it should reflect the way in which programs are composed from its (sub-) programs. I.e., the method should be syntax directed. Secondly, high-level languages allow (sub-) programs to be viewed as black boxes: Only the behaviour is relevant, not the way this behaviour is achieved. Hence, verification should be based on specifications of programs only. Consequently, the following principle is implied:

The specification of a program should be verifiable in terms of the specifications of its (syntactic) subprograms.

It is not difficult to discern Frege's notion of *semantic compositionality* in this statement ([11]).

The principle neatly demarcates various approaches to program verification:

(1) For input-output or partial correctness properties of sequential programs, one may contrast Floyd's non-compositional 'inductive assertion method' [7],

These investigations were supported by the Foundation for Computer Science Research in the Netherlands (SION) with financial aid from the Netherlands Organization for the Advancement of Pure Research (ZWO).

with Hoare's compositional 'Hoare logic' [10].

(2) For similar properties of concurrent programs, one can oppose Owicki's non-compositional notion of 'interference freedom of proofs' [20], against Lamport's compositional 'Concurrent Hoare Logic' [12].

(3) For general temporal properties of concurrent programs, there is the work of Manna and Pnueli [16]. They reduce such properties to properties of so-called transitions - traces associated with atomic actions. To verify such transition properties, Floyd's inductive assertion method is extended. This contrasts with the current paper which proposes Transition Logic, supporting compositional reasoning about such properties.

The work of Z. Manna and A. Pnueli shows that the temporal behaviour of programs is provable in terms of transition properties. The motivation, for the present paper, to concentrate on transition properties goes deeper, as they are essential to obtaining compositionality. In the presence of concurrency, execution of a program is, in general, influenced by the (concurrent) actions of its environment, i.e., by the actions of the program(s) executing in parallel. Only for atomic, indivisibly executing, programs does it make sense to consider them in isolation from the environment.

This is obvious for the interleaving model of concurrency - adopted in this paper - in which atomic statements are executed one at the time. It is defendable for other models of concurrency, too, such as A. Salwicki's maximal parallelism [22] or the partial order semantics of W. Reisig [21], both grounded in Petri-net theory: If moves are in conflict, their execution is mutually exclusive. If they are conflict free, this intuitively means that execution of one move does not influence the execution of the others. In that case, one can view such moves as being executed in a strict but unspecified order; no change in the computed values will ensue.

The execution trace of a program within an environment is construed as being formed from the transitions of the program and the environment. The ordering of the transitions in the trace reflects the flow-of-control during execution. Consequently, in Transition Logic a program is viewed as offering a set of transitions which it can perform, subject to constraints imposed by the syntactic structure of

the program, on the order in which the transitions can be taken; i.e., as a set of transitions constrained by the flow-of-control of the program.

Program specifications in Transition Logic take the form $\{p\} \alpha \{q\}$: Any transition of the program α that starts in a state satisfying the formula p , must end in a state satisfying q . To specify the flow-of-control, locations within programs will be labeled and these labels may appear in the formulae of the logic. Such labels also allow specifying the behaviour of each *individual* transition.

The next section introduces Transition Logic in more detail and defines the syntax and semantics of the logic w.r.t. a simple programming language - essentially while programs with parallel composition and variable sharing. Section 3 presents a formal system for this logic. Because validity depends on control locations too, the proof system contains an axiomatization of the flow-of-control in the programs discussed. The connection with temporal logic is described in section 4. Section 5 addresses issues of soundness and completeness. To render the rôle of the control-flow axiomatization explicit, first the basic properties on which completeness is based are formulated. Then, relative completeness of the logic is proved w.r.t. these properties. Finally, section 6 contains a conclusion and discusses some directions for further research.

This paper is clearly influenced by Lamport's work on invariance properties of concurrent programs, [12]. Lamport was the first to propose axiomatizing the flow-of-control, in the context of program verification. His 'Concurrent Hoare logic' was formalized by P. and R. Cousot in [5]. That paper contains a careful and abstract, language independent, analysis of the assumptions upon which completeness of the logic is based, that is of relevance to this paper, too. In [2], K. Apt and C. Delporte use similar ideas to prove certain eventuality properties of sequential programs.

The idea that control-flow forms an important aspect of concurrency is wellknown. That the formal study of flow-of-control is the fundamental principle behind compositional reasoning about properties of concurrent programs, seems to be less often realized.

2. TRANSITION LOGIC

Concurrency poses problems for compositional reasoning.

This is already apparent on a semantical level: Consider programs $\alpha \equiv x:=2$ and $\beta \equiv x:=1; x:=x+1$. For simplicity's sake, assume that the states of both programs consist of the single variable x , and let any execution start in the state satisfying $x=0$, i.e., in state 0. Denote the execution traces of α and β by $(0,2)$ and $(0,1,2)$, respectively. Now, consider the parallel execution of α and β : $\alpha \parallel \beta$. In any reasonable semantics, the moves of α and β will be interleaved. Hence, for $\alpha \parallel \beta$ the trace set $\{(0,2,1,2), (0,1,2,3), (0,1,2,2)\}$ is obtained. If the semantics of a program is defined as the set of its traces, then the principle of compositionality entails that the traces of $\alpha \parallel \beta$ should be obtainable from the traces of α and β . But how is one to construct, e.g., the trace $(0,1,2,3)$ from the traces $(0,2)$ and $(0,1,2)$? Such traces simply provide too little information, as state-changes caused by actions of the environment (β) of program α are ignored.

A similar observation can be made on the level of program specifications: The assertion $x \leq 2$ is an invariant property of the trace(s) of both α and β . This is expressed in a temporal logic like notation by: $\alpha: \Box(x \leq 2)$ and $\beta: \Box(x \leq 2)$. But although both α and β satisfy the same specification, $\alpha \parallel \beta: \Box(x \leq 2)$ is not valid, whereas $\alpha \parallel \alpha: \Box(x \leq 2)$ is. Again, too little information is provided to support compositionality.

The semantics of an atomic program can be defined as the set of its traces, because no concurrent action can change the program's behaviour. This suggests that the semantics of a non-atomic program may be caught in terms of the sequences of atomic programs it can execute. More precisely: The semantics of a program will be a set consisting of sequences of traces; each such trace corresponds to the execution of an atomic action within the program. The ordering of the traces in a sequence, reflects the control-flow within the program. *As there is no information about the actions of the environment that may interleave, no relation will be assumed between the values of the program variables as recorded in the end-state of a trace and their values as recorded in the begin-state of the next trace in any trace sequence.*

As an example, the trace sequences of the programs α and β above, are $(a,2)$ and $(b,1)(c,c+1)$, where a , b , and c denote arbitrary values. The set of trace sequences of $\alpha \parallel \beta$ is obtained, simply by interleaving the traces in the sequences of its components:

$\{(a,2)(b,1)(c,c+1), (b,1)(a,2)(c,c+1), (b,1)(c,c+1)(a,2)\}$. Once it is decided that $\alpha \parallel \beta$ executes in an empty environment, the execution traces can be reconstructed: As no actions will interleave, in the second sequence necessarily $a=1$ and $c=2$. The corresponding trace is obtained by ignoring duplicate states; hence $(b,1,2,3)$. These ideas on compositional semantics for concurrency are not original; see, e.g., [1, 4] and for early references relating these ideas to the notion of resumptions, [3, 18].

Transition Logic concentrates on properties of trace sequences. This means that for any program, the logic must be able to specify the behaviour of its atomic programs and the (syntactic) constraints on the flow-of-control that the program imposes. For the latter, every statement in a program must be labeled and these labels can be used as propositions in the (state-) formulae of the logic. The program state is extended to provide truth-values for the program labels. Predictably, if a label, l , is true in a state this means that in this state control within the program resides at the location labeled l . A vital observation is that, whereas the environment can alter the values of the variables of a program, it cannot change the values of the labels in the program; i.e., it cannot modify the location at which control resides in the program. It is this fact that makes it possible to axiomatize the flow-of-control, necessary to obtain a complete proof system.

During execution of an atomic program, no actions interleave. Hence, two atomic programs are for all purposes equivalent, whenever they have the same input-output behaviour¹. This is reflected in the basic formulae of Transition Logic, $[p]\alpha[q]$, in which atomic programs are viewed as transitions, transforming input-states into output-states: Any transition in α transforms a state satisfying p into a state satisfying q (remember that p and q may contain label-propositions).

One last point needs to be considered. If one wants to prove, say, $[p]\alpha;\beta[q]$ it would seem sufficient to show that $[p]\alpha[q]$ and $[p]\beta[q]$: The transitions of $\alpha;\beta$ are just the transitions of α plus the ones of β . For that matter, the same observation applies to a

¹Having no output (-state) counts as behaviour, too. Of course, if one wants to distinguish between, e.g., failure and divergence, a special divergence (or failure) state is needed.

proof of $[p]_{\alpha} \parallel [q]_{\beta}$. Now a problem arises, because p and q will, in general, refer to labels in α and in β . However, in the transition sequences of α , nothing can be assumed about the truth-values of β 's labels (and vice versa), as this entirely depends on the context in which α and β occur. In a context $\alpha; \beta$, no label of β can be true when control remains within α ; this is not so in a context $\alpha \parallel \beta$. These connections between the control locations of α and β are quite important. They form the reason that, e.g., the formula $[\text{true}]((x:=1; x:=2) \ell :) [\ell \rightarrow x=2]$ is a valid one, whereas a change of sequential into parallel composition in this program, invalidates it. Hence, Transition Logic allows formulae to appear in context: $\langle \beta \mid [p]_{\alpha} [q] \rangle$, where β is a context in which α occurs. Such contexts enforce conditions on the label-valuations of states occurring (in this case) in the transition sequences of α . Such states must have label-valuations which are consistent with the syntactic structure of the context. I.e., a formula $\langle \alpha; \beta \mid \text{in}(\alpha) \rightarrow \neg \text{in}(\beta) \rangle$ will be valid, whereas the formula $\langle \alpha \parallel \beta \mid \text{in}(\alpha) \rightarrow \neg \text{in}(\beta) \rangle$ will not be valid ($\text{in}(\alpha)$ expresses that control resides within α).

Introducing contexts solves the problem, but it is not claimed to be the only feasible solution. Another solution might be to disallow formulae $[p]_{\alpha} [q]$, in which p or q contains labels that do not appear in α . There is a trade-off here, between simple proof rules and simple semantics. This paper opts for the simpler proof rules.

syntax.

Start with some 1st order similarity type t , which defines the predicate, function and constant symbols that appear in the programs and formulae, and a second similarity type l , which defines a countable set of 0-ary predicate letters (i.e., propositions) ℓ_0, ℓ_1, \dots , disjoint from t , and functioning as labels.

Definition 2.0. The set of programs over t and l , $\text{PROG}(t, l)$, is inductively defined as the smallest set X such that

- (1) $(\ell.x:=e.\ell') \in X$ and $(\ell.b?.\ell') \in X$ for all distinct labels ℓ and ℓ' , variables x , and terms e and formulae b (atomic, if one wishes) over t
- (2) $\alpha, \beta \in X \Rightarrow (\ell.\alpha \circ \beta.\ell') \in X$ where (a) \circ is of the form $U, ;, * \text{ or } \parallel$, (b) no label in α (β) may appear in β (α), and (c) ℓ and ℓ' are distinct labels not appearing in α or β .

$\text{PROG}(t, l)$ is the set of regular programs with merge over assignments and tests. The program $\alpha * \beta$ is equivalent to $\alpha; (\beta; \alpha)^*$, and is introduced for technical reasons as explained in the next section. Note that $\alpha^* \equiv (\text{true}? * \alpha)$.

Definition 2.1. Transition Logic, $\text{TL}(t, l)$, is defined as the smallest set X such that

- (1) $\langle \beta \mid p \rangle \in X$ for programs β and formulae p over t, l
- (2) $\langle \beta \mid [p]_{\alpha} [q] \rangle \in X$ for programs β and subprograms α of β and formulae p and q over t, l .

There is no constraint on the labels that appear in p and q . Some notation:

- (1) $\text{LAB}(\alpha), \text{LAB}(p)$ - the set of labels appearing in program α respectively formula p ,
- (2) $F(S)$ - the 1st order formulae over S ,
- (3) $\langle \beta \mid \alpha \rangle$ - α is a subprogram of β ,
- (4) $\langle \beta \mid \ell \parallel \alpha \rangle$ - ℓ occurs in a subprogram of β that can be executed parallel to α . Formally: There are programs $\bar{\alpha}$ and γ and labels m and m' such that $\langle \bar{\alpha} \mid \alpha \rangle, \ell \in \text{LAB}(\gamma)$ and either $\langle \beta \mid (m.\gamma \parallel \bar{\alpha}.m') \rangle$ or $\langle \beta \mid (m.\bar{\alpha} \parallel \gamma.m') \rangle$
- (5) if $\alpha \equiv (\ell.\beta.\ell')$ then $\alpha \stackrel{D}{=} \ell, \alpha^* \stackrel{D}{=} \ell^*$ and $\bar{\alpha} \stackrel{D}{=} (\vee \{ \bar{\ell} \mid \bar{\ell} \in \text{LAB}(\alpha) \}) \wedge \neg \ell'$.

In (5), $\alpha, \bar{\alpha}$ and α^* , correspond to the location-predicates $\text{at}(\alpha), \text{in}(\alpha)$ and $\text{after}(\alpha)$ of [12].

semantics.

A model for $\text{TL}(t, l)$ is a tuple, $\langle M, S \rangle$, where M is some t -structure (i.e., a model for $F(t)$) and S is the class of states over M ; each state providing a valuation of the variables and of the label-propositions. Such models define

- (1) a partial function $\text{Trs}: \text{PROG}(t, l)^2 \rightarrow P((S \times S)^\omega)$, where the range of Trs is the powerset of finite and infinite sequences of state-pairs, and
- (2) a satisfaction relation $\models \subseteq \langle M, S \rangle \times \text{TL}(t, l)$.

With any pair of programs, (β, α) , such that $\langle \beta \mid \alpha \rangle$, $\text{Trs}(\beta, \alpha)$ associates the set of transition sequences obtained by executing α within the context β . First, an auxiliary function $\text{Tr}: \text{PROG}(t, l) \rightarrow P((S \times S)^\omega)$ is defined, that associates transition sequences with programs in an empty environment.

The definition of Tr is standard, but for the label-valuations. The strategy in the inductive definition below will be to take states (for the induction base) and transition sequences (for the induction step) in which all labels whose truth-value will be affected, are set initially to false. Labels whose

value must change in a state, will be set explicitly.

There is one tricky point. If l is some label that does not appear in a program α , nothing can be assumed about its truth-value during execution of α . Whether a transition in α will change its value or not, entirely depends on the context of α and l , of which nothing is known. Compare, e.g., the contexts $\alpha; (l.b.l')$ and $\alpha \parallel (l.b.l')$. This has the curious effect that any transition in a program must be allowed to change arbitrarily the truth-value of any label not appearing in the program.

In the definition below, state-variants are indicated as usual by $\sigma(\cdot/\cdot)$. This notation is (1) extended to denote variants of transitions: $\{\cdot/\cdot\}$, $\{\cdot/\cdot\}_1$ - denoting variants of the 1st and 2nd component of the state pair - and $\{\cdot/\cdot\}_{1,2}$ - denoting a variant of both components; and (2) elementwise extended to denote variants of sequences. Satisfaction of $F(tl)$ -formulae, $M, \sigma \models p$, is assumed to be understood. To define the semantics of $(l.\alpha*\beta.l')$, the syntax is (temporarily) extended to allow for any integer $n \geq 0$, programs $(l.\alpha(n)\beta.l')$. Abusing notation, we have $(l.\alpha*\beta.l') \stackrel{D}{=} \cup \{(l.\alpha(n)\beta.l') \mid n \geq 0\}$.

Definition 2.3. For any $\alpha \in \text{PROG}(tl)$, inductively define $\text{Tr}(\alpha)$ by the following clauses.

- (1) $\alpha = (l.b?.l')$ $\{ (\sigma_1\{tt/l\}, \sigma_2\{tt/l'\}) \mid \sigma_i \in S, \sigma_i(l) = \sigma_i(l') = ff \ i=1,2, \sigma_1 \upharpoonright \text{Var} = \sigma_2 \upharpoonright \text{Var}, M, \sigma_1 \models b \} = \text{Tr}(\alpha)$
- (2) $\alpha = (l.x := e.l')$ $\{ (\sigma_1\{tt/l\}, \sigma_2\{tt/l'\}) \mid \sigma_i \in S, \sigma_i(l) = \sigma_i(l') = ff \ i=1,2, \sigma_1(\sigma_1(e)/x) \upharpoonright \text{Var} = \sigma_2 \upharpoonright \text{Var} \} = \text{Tr}(\alpha)$

For any $\alpha \in \text{PROG}(tl)$, $\tau_i \in \text{Tr}(\alpha_i)$ $i=1,2$ and labels l, l' such that $\alpha = (l.\alpha_1 \alpha_2.l')$ and for $i=1,2$ $\text{len} \tau_i = k_i$, $L_i = \text{LAB}(\alpha_i)$, $\tau_i[k_i] = (\sigma_i, \sigma'_i)$, $\tau_i[1] = (\bar{\sigma}_i, \bar{\sigma}'_i)$ and $\forall \bar{l} \in L_{i \bmod 2 + 1} \tau_i(\bar{l}) = \tau_i(\bar{l}) = \tau_i(l') = ff$:

- (3) $\alpha = \cup \{ \tau_i[1] \{tt/l\} \wedge \tau_i[2:k_i-1] \wedge \tau_i[k_i] \{tt/l'\} \mid i=1,2 \} \subseteq \text{Tr}(\alpha)$
- (4) $\alpha = ; \tau_1[1] \{tt/l\} \wedge \tau_1[2:k_1-1] \wedge \tau_1[k_1] \{ \bar{\sigma}_2(\bar{l})/2 \bar{l} \in L_2 \} \wedge \tau_2[1] \{ \sigma'_1(\bar{l})/1 \bar{l} \in L_1 \} \wedge \tau_2[2:k_2-1] \wedge \tau_2[k_2] \{tt/l'\} \in \text{Tr}(\alpha)$
- (5) $\alpha = (0) \tau_1[1] \{tt/l\} \wedge \tau_1[2:k_1-1] \wedge \tau_1[k_1] \{tt/l'\} \in \text{Tr}(\alpha)$
- (6) $\alpha = (n+1) \{ \tau_1[1:k-1] \wedge \tau_1[k] \{ff, \bar{\sigma}_2(\bar{l})/2 \bar{l}' \in L_2 \} \wedge \tau_2[1] \{ \sigma'_1(\bar{l})/1 \bar{l} \in L_1 \} \wedge \tau_2[2:k_2-1] \wedge \tau_2[k_2] \{tt, \bar{\sigma}_1(\bar{l})/2 \bar{l} \in L_1 \} \wedge \tau_1[1] \{tt, \sigma'_2(\bar{l})/1 \bar{l} \in L_2 \} \wedge \tau_1[2:k_1-1] \wedge \tau_1[k_1] \{tt/l'\} \mid \tau \in \text{Tr}((l.\alpha_1(n)\alpha_2.l')) \} \subseteq \text{Tr}(\alpha)$
- (7) $\alpha = * \text{Tr}((l.\alpha_1(n)\alpha_2.l')) \subseteq \text{Tr}(\alpha)$ for all $n \geq 0$
- (8) $\alpha = \parallel \{ \tau[1] \{tt/l\} \wedge \tau[2:k-1] \wedge \tau[k] \{tt/l'\} \mid \tau \in \text{merge}(\bar{\tau}_1, L_1, \bar{\tau}_2, L_2), k = \text{len} \tau, \text{for } i=1,2 \bar{\tau}_{i \bmod 2 + 1} = \tau_{i \bmod 2 + 1} \{ \bar{\sigma}_i(\bar{l})/i_2 \bar{l} \in L_i \} \} \subseteq \text{Tr}(\alpha)$
where $\text{merge}(\tau, \Gamma, \tau', \Gamma') \stackrel{D}{=} \text{lmerge}(\tau, \Gamma, \tau', \Gamma') \cup \text{lmerge}(\tau', \Gamma', \tau, \Gamma)$
and $\text{lmerge}(\tau, \Gamma, \tau', \Gamma') \stackrel{D}{=} \begin{cases} \{\tau'\} & \text{if } \text{len} = 0, \\ \{ \tau[1] \wedge \bar{\tau} \mid \bar{\tau} \in \text{merge}(\tau[2:], \Gamma, \tau'(\bar{\sigma}'(l)/i_2 \bar{l} \in \Gamma), \Gamma') \} & \text{otherwise} \end{cases}$

¹The value of l, \bar{l} and l' in any state of τ_i is false.

From this definition, it follows that for any program α , the possible valuations of the labels appearing in α , only depend on the structure of α and (hence) do not change if α is embedded in some context $\bar{\alpha}$; provided α remains reachable in $\bar{\alpha}$. More precisely, the following *encapsulation* property holds:

For any programs $\alpha, \bar{\alpha}$ and any label $l \in \text{LAB}(\alpha)$ such that $\langle \bar{\alpha} \mid \alpha \rangle$ and $\exists \tau \in \text{Tr}(\bar{\alpha}) \exists \bar{\sigma} \in \text{int} \bar{\sigma}(\cdot) = tt$:
 $(\exists \tau \in \text{Tr}(\alpha) \exists \sigma \in \text{int} \sigma(l) = w) \Leftrightarrow (\exists \bar{\tau} \in \text{Tr}(\bar{\alpha}) \exists \bar{\sigma} \in \text{int} \bar{\sigma}(l) = w)$,
where w denotes either tt or ff .

It is this property that makes the control-flow axiomatization in section 3 possible.

Next, to define TrS , the restrictions, imposed by a context (say, of α), on the values of the labels which do not occur in α , have to be enforced. This requires some additional machinery. A context β of a program α restricts the label valuations of α 's transition sequences, so as to make them consistent with executing α in that context (β). Hence, it would seem that a trace $\tau \in \text{TrS}(\beta, \alpha)$ is obtained by "projecting" a trace $\bar{\tau} \in \text{Tr}(\beta)$ on the transitions in α . Such a trace, however, need not exist; consider, e.g., a

context $\beta = \text{false?}; \alpha$. The fact that α cannot be reached semantically, is of no concern for the consistency of α 's label-valuations, which should depend on the syntactic structure of the context only. In fact, the straightforward axiomatization of the consistency of label-valuations in section 3 (LC_α) is rendered incomplete by taking semantic reachability into account. This suggests, defining for each program α , its "companion" α^c , obtained by replacing each test $(l.b?.l')$ in α by $(l.\text{true?}.l')$ (; in fact, it suffices to replace it by a test that can be passed from

at least one state). In α^c , any execution path that is syntactically possible (in α), is semantically possible.

Now, auxiliary predicates, β -sgs (called β -state good) and β -tgs (called β -transition good) can be defined, to express that a state or a transition has valuation(s) that are consistent with β 's syntax:

Definition 2.4. For any $\beta \in \text{PROG}(t1)$, $\sigma \in S$, $L \stackrel{D}{=} \text{LAB}(\beta)$

- (1) β -sg(σ) iff $\exists \tau \in \text{Tr}(\beta^c) \exists i \geq 1 \tau[i] = (\bar{\sigma}, \bar{\sigma}')$
 $\bar{\sigma}|L = \sigma|L \vee \bar{\sigma}'|L = \sigma|L$
 (2) β -tg(σ, σ') iff $\exists \tau \in \text{Tr}(\beta^c) \exists i \geq 1 \tau[i] = (\bar{\sigma}, \bar{\sigma}')$
 $\bar{\sigma}|L = \sigma|L \wedge \bar{\sigma}'|L = \sigma'|L$

Note the disjunction in (1) and the conjunction in (2).

Definition 2.5. For any $\alpha, \beta \in \text{PROG}(t1)$ such that $\langle \beta | \alpha \rangle$ and $\tau \in (S \times S)^\omega$:

$\tau \in \text{Trs}(\beta, \alpha)$ iff $\tau \in \text{Tr}(\alpha) \wedge \forall i (\tau[i] = (\sigma, \sigma') \Rightarrow \beta\text{-tg}(\sigma, \sigma'))$

Finally, validity and satisfaction in a model can be defined:

Definition 2.6. Let $p, q \in F(t1)$, $\phi \in \text{TL}(t1)$, $\alpha, \beta \in \text{PROG}(t1)$ and $\langle \beta | \alpha \rangle$. Then

- (1) $\models \phi$ iff \forall models $\langle M, S \rangle$ of $\text{TL}(t1)$ $\langle M, S \rangle \models \phi$
 (2) $\langle M, S \rangle \models \langle \beta | p \rangle$ iff $\forall \sigma \in S$ β -sg(σ) $\Rightarrow M, \sigma \models p$
 (3) $\langle M, S \rangle \models \langle \beta | [p] \alpha [q] \rangle$ iff $\forall \tau \in \text{Tr}(\beta, \alpha) \forall i \geq 1$
 $\tau[i] = (\sigma, \sigma') \Rightarrow (M, \sigma \models p \Rightarrow M, \sigma' \models q)$

3. THE PROOF SYSTEM, PS(M).

As notation for proof rules, $\langle \beta | A_1, \dots, A_n \Rightarrow C \rangle$ is used, with the usual interpretation: To infer $\langle \beta | C \rangle$, prove $\langle \beta | A_1 \rangle, \dots, \langle \beta | A_n \rangle$.

The system consists of three parts. The first part concerns the composition of program (-proofs) and consists of the 4 Composition rules, the Test and the Assignment rule:

For α of the form $U, ;, * \text{ or } \parallel, C_0$.

$\langle \beta | [p] \alpha_1 [q], [p] \alpha_2 [q] \Rightarrow [p] (\ell.a_1 \alpha_2 . \ell') [q] \rangle$

T. $\langle \beta | (p \wedge b \wedge \alpha \rightarrow q) \Rightarrow [p] \alpha [q] \rangle$ provided $\alpha \equiv (\ell.b?.\ell')$ and $\text{LAB}(p) = \emptyset$

As. $\langle \beta | (\alpha \wedge \exists y (p[y/x] \wedge x = e[y/x]) \rightarrow q) \Rightarrow [p] \alpha [q] \rangle$ provided $\alpha \equiv (\ell.x := e.\ell')$, $\text{LAB}(p) = \emptyset$ and $y \notin \text{Fvar}(p, e) \cup \{x\}$

Next, follow some auxiliary rules and axioms. The Invariance, Strengthen and the 2 Extension rules, the Oracle rule and the Propositional tautology axiom:

Inv. $\langle \beta | p \Rightarrow [\text{true}] \alpha [p] \rangle$

S. $\langle \beta | [p \wedge q] \alpha [r], \alpha \rightarrow q, [q] \alpha [q \vee \alpha'] \Rightarrow [p] \alpha [r] \rangle$, $q \in F(1)$

E₁. prove $\langle \beta | p \rangle$ to infer $\langle \beta' | p \rangle$ provided $\langle \beta' | \beta \rangle$ and $p \in F(1)$

E₂. prove $\langle \beta | [p] \alpha [q] \rangle$ to infer $\langle \beta' | [p] \alpha [q] \rangle$ provided $\langle \beta' | \beta \rangle$

O. $\langle \beta | p \rangle$ for any $p \in F(t1)$ such that $M \models p$

Pt. All instances of propositional tautologies

Finally, the control flow axioms (Flow and Label Consistency). The formulae $\text{LC}(\alpha)$ are defined below:

F_{1,2}. $\langle \beta | [\ell] \alpha [\ell] \rangle$ provided $\langle \beta | \ell \parallel \alpha \rangle$ and α is atomic
 LC _{α} . $\langle \alpha | \text{LC}(\alpha) \rangle$ for any $\alpha \in \text{PROG}(t1)$

The recursively enumerable set of label consistency axioms, LC_α , axiomatizes the consistency of label valuations. The $F(1)$ -formulae $\text{LC}(\alpha)$ are recursively defined as follows ($\bar{\vee}$ denotes the exclusive or)

- (1) $\alpha \equiv (\ell.b?.\ell')$ or $(\ell.x := e.\ell')$ $\Rightarrow \text{LC}(\alpha) \stackrel{D}{=} \ell \bar{\vee} \ell'$
 if $\alpha \equiv (\ell.a_1 \alpha_2 . \ell')$ then $\text{LC}(\alpha) \stackrel{D}{=} \text{LC}(\alpha_1) \wedge \text{LC}(\alpha_2) \wedge C$,
 where C is defined by the following clauses:
 (2) $\alpha \equiv U \Rightarrow \ell \bar{\vee} (\alpha_1 \bar{\vee} \alpha_2) \wedge \ell \bar{\vee} (\alpha_1 \bar{\vee} \alpha_2) \wedge \neg (\hat{\alpha}_1 \vee \alpha_1) \wedge (\hat{\alpha}_2 \vee \alpha_2)$
 (3) $\alpha \equiv ; \Rightarrow \ell \bar{\vee} \alpha_1 \wedge \alpha_1 \bar{\vee} \alpha_2 \wedge \alpha_2 \bar{\vee} \ell' \wedge \neg (\hat{\alpha}_1 \wedge (\hat{\alpha}_2 \vee \alpha_2))$
 (4) $\alpha \equiv * \Rightarrow \ell \bar{\vee} \alpha_1 \wedge \alpha_1 \bar{\vee} (\alpha_2 \bar{\vee} \ell') \wedge \alpha_2 \bar{\vee} \ell \wedge \neg (\hat{\alpha}_1 \wedge \hat{\alpha}_2)$
 (5) $\alpha \equiv \parallel \Rightarrow \ell \bar{\vee} (\alpha_1 \wedge \alpha_2) \wedge \ell' \bar{\vee} (\alpha_1 \wedge \alpha_2) \wedge (\hat{\alpha}_1 \vee \alpha_1 \bar{\vee} \alpha_2 \vee \hat{\alpha}_2 \vee \alpha_2)$

$\text{LC}(\alpha)$ expresses that a valuation of the labels of α is consistent with α 's syntactic structure. Such a recursive definition is possible, because this consistency does not depend on α 's environment, i.e., because the encapsulation property holds. This would have not been true, had the usual Kleene-* been used instead of the binary *: $\text{LC}((\ell.a^*.\ell'))$ necessarily will allow the inference $\alpha' \rightarrow \alpha$ (or $\alpha \rightarrow \alpha'$) in order to model looping. As $\alpha, \alpha' \in \text{LAB}(\alpha)$, this means that now, a context of α can induce new relations between α 's labels.

Some useful derived rules and axioms are:

- I. $\langle \beta | p \rightarrow p', [p'] \alpha [q], q' \rightarrow q \Rightarrow [p] \alpha [q] \rangle$
 C _{\wedge, \vee} . $\langle \beta | [p] \alpha [q], [p'] \alpha [q'] \Rightarrow [p \wedge p'] \alpha [q \vee q'] \rangle$
 F₃. $\langle \beta | [\text{true}] \alpha [\alpha'] \rangle$ provided α is atomic
 F₄. $\langle \beta | [\text{true}] \alpha [\hat{\alpha} \vee \alpha'] \rangle$

Axiom F₃ easily follows from T and AS. Proofs for the other axiom and the proof rules are constructed using structural induction w.r.t. α and, for F₄, applying the trivial observation that $\models \forall \gamma \gamma \bar{\vee} \hat{\alpha} \vee \alpha'$ if $\langle \hat{\alpha} | \gamma \rangle$ (cf. the definition of $\hat{\alpha}$ and α').

example.

Consider Lamport's Concurrent Hoare logic [12]. His sequential composition rule is a derived rule of PS(M). Lamport deals with safety properties and his

specifications, $\{p\}\alpha\{q\}$, have the interpretation: "If execution is begun anywhere in α with the predicate p true, then executing α will leave p true while control is inside α , and will make q true if and when α terminates". Consequently, $\{p\}\alpha\{q\}$ translates into Transition Logic as: $\{p\}\alpha\{\delta \rightarrow p \wedge \alpha' \rightarrow q\}$ (in [10] the context is always left implicit).

In this paper's notation, Lamport's sequential composition rule becomes:

$$SC. \langle \delta \mid [t]\alpha[\delta \rightarrow t \wedge \alpha' \rightarrow u], [v]\beta[\beta \rightarrow v \wedge \beta' \rightarrow w], u \wedge \beta \rightarrow v \Rightarrow \\ \delta \rightarrow t \wedge \beta \rightarrow w \mid \gamma[\delta \rightarrow t \wedge \beta \rightarrow v \wedge \gamma' \rightarrow w] \rangle,$$

where $\gamma = (\delta. \alpha; \beta. \delta')$ and $\langle \delta \mid \gamma \rangle$.

To derive SC, it will be convenient to generalize it to:

$$SC'. \langle \delta \mid [p]\alpha[q], [r]\beta[s] \Rightarrow [A]\gamma[C] \rangle, \text{ where } \gamma \text{ and } \delta \\ \text{are as in SC, } A = \delta \rightarrow p \wedge \beta \rightarrow r \text{ and} \\ C = (\delta \vee \alpha' \rightarrow q) \wedge ((\neg \beta \wedge (\beta \vee \beta')) \rightarrow s)$$

It is not difficult to see that SC' is sound. To derive SC' requires some more effort:

We want to use the C₁ rule. So, first derive $\vdash \langle \delta \mid [A]\alpha[C] \rangle$. In the derivation, the context is suppressed.

$$\begin{array}{l} [true]\alpha[\delta \vee \alpha'] \quad LC(\gamma) \rightarrow (\delta \vee \alpha' \rightarrow (\neg \beta \vee \neg(\beta \vee \beta'))) \\ \delta \rightarrow true \quad LC(\gamma) \\ \hline [\delta] \alpha [\delta \vee \alpha'] \quad [p] \alpha [q] \quad \text{Pt} \\ \hline [\delta \wedge p] \alpha [(\delta \vee \alpha') \wedge q] \quad C \\ \delta \wedge A \rightarrow \delta \wedge p \quad (\delta \vee \alpha') \wedge q \rightarrow C \\ \hline [\delta \wedge A] \alpha [C] \\ \alpha \rightarrow i \quad [true] \alpha [\delta \vee \alpha'] \\ \hline [A] \alpha [C] \quad S \end{array}$$

The proof of $\langle \delta \mid [A]\beta[C] \rangle$ is completely analogous. Now, SC' is obtained by a final application of C₁.

It remains to show that SC is a derived rule, too.

To prove this, apply SC' with $p = t$, $q = (\delta \rightarrow t \wedge \alpha' \rightarrow u)$, $r = v$ and $s = \beta \rightarrow v \wedge \beta' \rightarrow w$. Formula A has the required form, C not yet. Observe that

$$\vdash C \leftrightarrow (\delta \rightarrow t \wedge \alpha' \rightarrow u \wedge \neg \beta \wedge \beta \rightarrow v \wedge \neg \beta \wedge \beta' \rightarrow w) \text{ and that} \\ \vdash LC(\delta) \rightarrow (\alpha \leftrightarrow \beta \wedge \beta' \rightarrow \gamma \wedge \beta' \rightarrow \neg \beta).$$

Hence, using O, Pt and LC₀, we obtain

$$\langle \delta \mid [A]\gamma[C] \rangle \vdash \langle \delta \mid [A]\gamma[\delta \rightarrow t \wedge \beta \rightarrow u \wedge \neg \beta \wedge \beta' \rightarrow v \wedge \beta' \rightarrow w] \rangle.$$

To obtain the consequent of SC, use the assumption $\langle \delta \mid u \wedge \beta \rightarrow w \rangle$, to replace $\beta \rightarrow u$ by $\beta \rightarrow w$.

4. THE CONNECTION WITH TEMPORAL LOGIC.

Transition Logic cannot reason about and specify temporal properties in a direct way, like the logics of Lamport [13] and Manna and Pnueli [16] can. This ability was sacrificed to obtain compositionality. However, this does not imply a loss in reasoning power. To wit: Any temporal property, expressible in Manna and Pnueli's Temporal Program Logic (TPL), that is true of a program α ($\alpha \in \text{PROG}(t1)$), can be proved hence reasoned about in a suitable extension of Transition Logic.

This claim follows from a straightforward translation of results from [15, 16, 17]: TPL reduces proofs of temporal properties to proofs of classical properties without temporal modalities. These classical properties are readily translated into Transition Logic formulae. In [17], rules are introduced for certain classes of temporal formulae, which are complete in the sense that to prove a formula from any of these classes, one application of the respective rule suffices, the rest is classical reasoning (which can be done in PS(t1)). The claim follows by extending PS(t1) with these rules and by observing that (for programs in PROG(t1)) any temporal formula can be reduced to a formula in one of the above classes.

Note that no justness and fairness constraints are imposed upon the executions of PROG(t1)-programs. Hence, although TPL is capable of dealing with such constraints, this aspect is ignored here. The rest of the section only states the facts and (hence) the reader should be familiar with [15, 16, 17].

In TPL, (temporal) formulae are interpreted over a fixed set of state-sequences; that is, over the execution traces of a fixed program. Consequently, there is no notion of composition of programs in TPL. Hence, with each program $\alpha \in \text{PROG}(t1)$, associate its set of traces, Et(α) (obtained from Tr(α)). As in TPL, if ϵ is a trace of α ($\epsilon \in \text{Et}(\alpha)$), then by definition any non-empty suffix of ϵ is a trace of α , too.

Temporal logic, Tel(t1), is obtained by extending classical 1st order logic F(t1), with 3 temporal modalities: \bigcirc (next state), \diamond (eventuality) and U (weak until or unless). The formulae of Tel(t1) are defined as usual, by inductively applying these modalities, starting with F(t1)-formulae. It is a fact that any temporal modality can be expressed using

\circ, \diamond and U . The interpretation of $\text{Tel}(t1)$ (over $\text{Et}(\alpha)$) is inductively defined as follows:

- $\langle M, S \rangle \vDash \alpha : \phi$ iff $\forall \epsilon \in \text{Et}(\alpha) M, \epsilon \vDash \phi$
- $M, \epsilon \vdash p$ iff $M, \epsilon[1] \vdash p$ (for $p \in F(t1)$)
- $M, \epsilon \vdash \circ$ iff $\text{len} \epsilon \geq 1$ and $M, \epsilon[2:] \vdash \phi$
- $M, \epsilon \vdash \diamond$ iff $\exists i \leq \text{len} \epsilon M, \epsilon[i:] \vdash \phi$
- $M, \epsilon \vdash U \psi$ iff $\forall i \leq \text{len} \epsilon M, \epsilon[i:] \vdash \phi$ or
 $\exists j \leq \text{len} \epsilon \forall 1 \leq i < j M, \epsilon[i:] \vdash \phi \wedge M, \epsilon[j:] \vdash \psi$

In [17] three rules are introduced for proving properties of the form $p \rightarrow q$, $p \rightarrow \circ q$ and $p \rightarrow q \cup r$, where p, q and r are state-formulae, i.e., $p, q, r \in F(t1)$. Provided $F(t1)$ meets certain requirements, these rules are shown to be complete in the above sense that any of these properties can be proved, using one single application of the respective rule as the only temporal step. I.e., no temporal modalities appear in the premisses of these rules. In this paper's terminology, the rules are:

$$\begin{array}{c}
 \begin{array}{l}
 [p \wedge \bar{p}] \alpha [q] , p \rightarrow \bar{b} \\
 \alpha : p \rightarrow q
 \end{array}
 \quad U \quad
 \begin{array}{l}
 [\bar{q} \wedge \bar{p}] \alpha [\bar{r} \vee \bar{q}] , \alpha \rightarrow \bar{p} , [\bar{p}] \alpha [\bar{p}] \\
 p \wedge \bar{p} \rightarrow \bar{q} \vee \bar{r} , \bar{q} \wedge \bar{p} \rightarrow q , \bar{r} \wedge \bar{p} \rightarrow r \\
 \alpha : p \rightarrow q \cup r
 \end{array}
 \\
 \\
 \begin{array}{l}
 [\bar{q}(n)] \alpha [q \vee (\exists \lambda \exists m < n. \bar{q}(m))] \\
 \alpha : p \rightarrow q
 \end{array}
 \quad \text{where } n \text{ and } m \\
 \begin{array}{l}
 \alpha \rightarrow \bar{p} , [\bar{p}] \alpha [\bar{p}] , p \wedge \bar{p} \rightarrow q \vee (\exists \lambda \exists n. \bar{q}(n)) \\
 \text{are natural numbers.}
 \end{array}
 \end{array}$$

In these rules, the TL-formula $[p] \alpha [q]$ translates the TPL-expression " α leads from p to q ", with interpretation [17]: For any atomic action τ of α and for any states s and s' , $p(s) \wedge s' \in M[\tau]s \rightarrow q(s')$. Note that as a consequence, $\text{PS}(t1)$ allows formal reasoning about such TPL-expressions; something that is missing in [16, 17].

Rules E and N have been changed w.r.t the original rules of [17]: In TPL, traces of programs are always infinite (if necessary, the terminal state of such a trace is repeated); not so for the traces in $\text{Et}(\alpha)$. This accounts for the additional premiss $p \rightarrow \bar{b}$ in N and for the appearance of \bar{a} in the first and last premiss of E. Moreover, the E-rule could be simplified, because justness and fairness issues are ignored.

The proofs in [15, 17] indicate that, in order to obtain completeness of these rules, the following predicates and relation must be $F(t1)$ -definable:

- (1) " $\bar{\tau}$ is a finite prefix of some $\tau \in \text{Et}(\alpha)$ ",
- (2) " $M, \sigma \vdash p$ " (for $p \in F(t1)$),
- (3) " n is a standard natural number" and " $n < m$ ".

Rule E uses a well-foundedness argument to show eventuality. This explains the third requirement. Also, the formula $\bar{q}(n)$ in rule E, will have in general to express something like " α will establish q in at most n more execution steps". Hence, requirements 1 and 2.

Now, what about the proof of general temporal properties? Observe that no transition sequence in $\text{Tr}(\alpha)$, hence no trace in $\text{Et}(\alpha)$, has infinite length. This implies that " $\tau \in \text{Et}(\alpha)$ " is $F(t1)$ -definable, too. Consequently, " $M, \tau \vdash \phi$ " is definable for any $\phi \in \text{Tel}(t1)$ and $\tau \in \text{Et}(\alpha)$, so that any temporal formula can be reduced to one of the standard forms $p \rightarrow \circ q$, $p \rightarrow \diamond q$ or $p \rightarrow q \cup r$.

We agree that the use of such coding tricks is unsatisfactory but observe that such tricks have to be used in TPL, too.

As an example, let p_e be the signature of Peano-arithmetic, extended with a unary predicate (symbol) nat . Let M be a class of models in which nat is interpreted as the defining predicate of the standard natural numbers (and in which the symbols in p receive their usual interpretation). Clearly, $F(p_e 1)$ satisfies the definability requirements 1, 2 and 3. Hence, assuming for the moment that $\text{PS}(M)$ is complete (w.r.t. M), then

$$M \vDash \alpha : \phi \Leftrightarrow \text{PS}(M) \cup \{N, U, E\} \vdash \phi, \alpha \in \text{PROG}(p_e 1), \phi \in \text{Tel}(p_e 1).$$

This can of course be generalized to any signature that allows an arithmetical model (cf. Harel's [9]).

5. SOUNDNESS AND COMPLETENESS.

$\text{PS}(M)$ is an ordinary Hilbert-style proof system. Consequently, to establish soundness, it suffices to show that each individual proof rule is sound and that every axiom scheme is valid.

In this section, references to models, M , are usually suppressed. Also, if ℓ is a label then $\bar{\ell} \in F(1)$ denotes uniformly either the formula ℓ or $\neg \ell$, in expressions.

Lemma 5.0. $E_1, O, Pt, F_{1,2}$ and LC_α are valid.

Proof.

Rule E_1 . Remember that $\langle M, S \rangle \vDash \langle B' | p \rangle \Leftrightarrow (\forall \sigma \in S \beta' \text{-sg}(\sigma) \Rightarrow M, \sigma \vdash p)$. Now, observe that $\beta' \text{-sg} \subseteq \beta \text{-sg}$. This is a consequence of the encapsulation property and the fact that for $\ell \notin \text{LAB}(\beta)$, $\beta \text{-sg}$ imposes no con-

¹Strictly speaking, an additional rule is needed, allowing the substitution of equivalent formulae in temporal formulae, so as to be able to reduce to standard form.

straint on their values.

0, Pt. Trivial and uninteresting.

$F_{1,2}$. By definition of $\langle \beta | \ell | \alpha \rangle$, there is a subprogram of β , $\rho = (m, \gamma | \delta, m')$ such that $\langle \gamma | \alpha \rangle$ and $\ell \in \text{LAB}(\delta)$ or vice versa. W.l.o.g., assume that $\langle \gamma | \alpha \rangle$. The encapsulation property implies that it suffices to prove that for any transition (σ, σ') caused by an action in γ , $\rho - \text{tg}(\sigma, \sigma') = \sigma(\ell) = \sigma'(\ell)$. This follows directly from the definition of merge.

LC_α . This is proved using structural induction. If α is atomic, LC_α is easily seen to hold. The induction step amounts to a depressing analysis of how composite transition sequences are constructed and is left to the reader.

Lemma 5.1. C_0, T, As, Inv, S and E_2 are sound rules.

Proof.

As. Choose any $(\sigma, \sigma') \in \text{Trs}(\beta, \alpha)$ such that $\sigma \models p$. To show: $\sigma' \models q$.

Assume that $\sigma' \models \alpha \wedge \exists y(\dots)$. As $\sigma' \models \alpha$, this implies $\sigma'(\bar{v}/y) \models (\dots)$ for every $\bar{v} \in |M|$; in particular $\sigma'(\bar{v}/y) \models (\dots)$, where $\bar{v} = \sigma(x)$. Because $\sigma'(\bar{v}/y) \models x = e[y/x]$ ($y \notin \text{Fvar}(e) \cup \{x\}$), also $\sigma'(\bar{v}/y) \models p[y/x]$. Consequently, $\sigma'(\bar{v}/x) \models p$ ($y \notin \text{Fvar}(p)$). As $\text{LAB}(p) = \emptyset$, this implies $\sigma \models p$; a contradiction. Hence $\sigma' \models \alpha \wedge \exists y(\dots)$, so that $\sigma' \models q$.

C_0, T, Inv, S, E_2 . These rules are even more trivial to prove sound and are left as exercises.

Theorem 5.0. $\text{PS}(M) \vdash \phi \Rightarrow M \models \phi$ for any $\phi \in \text{TL}(t1)$.

Proof. An easy induction w.r.t. the complexity of the proof.

completeness.

$\text{PS}(M)$ is complete relative to the theory of M (cf. the Oracle axiom). Notably, no further restrictions on models M are needed, such as Cook's notion of expressiveness. The completeness proof does rely on the definability of the strongest post condition (spc) for atomic programs. However, in TL this is definable in any (1st order) model. Of course, section 4 indicated that in order to prove interesting properties, models do have to meet additional criteria.

The following two properties are essential for obtaining completeness:

- (A) $LC(\beta)$ defines $\beta\text{-sg}$: $\forall \sigma \in S, M, \sigma \models LC(\beta) \Leftrightarrow \beta\text{-sg}(\sigma)$
- (B) Let $\alpha, \beta \in \text{PROG}(t1)$, α atomic, $\langle \beta | \alpha \rangle$ and $\ell \in \text{LAB}(\beta)$. If $\langle \beta | \ell | \alpha \rangle$ then the value of ℓ remains unaffected by executing α : $\forall (\sigma, \sigma') \in \text{Trs}(\beta, \alpha) \sigma(\ell) = \sigma'(\ell)$.

If $\neg \langle \beta | \ell | \alpha \rangle$ then the value of ℓ in a state produced by α is independent of the state in which α was executed: For w either tt or ff,

$$(\exists (\sigma, \sigma') \in \text{Trs}(\beta, \alpha) \forall \sigma'' (\sigma, \sigma'') \in \text{Trs}(\beta, \alpha) \Rightarrow \sigma''(\ell) = w) \Rightarrow \forall (\bar{\sigma}, \bar{\sigma}') \in \text{Trs}(\beta, \alpha) \bar{\sigma}'(\ell) = w$$

The proofs of these properties are straightforward but very lengthy analyses of the construction of the label valuations in the definition of Tr and will not be given in this paper.

The completeness proof splits into a number of cases. The easiest case is the subject of

Theorem 5.1. $M \models \langle \beta | p \rangle \Rightarrow \text{PS}(M) \vdash \langle \beta | p \rangle$ ($p \in \text{EF}(t1)$).

Proof. Property (A) implies that $M \models LC(\beta) \Rightarrow p$. Now, use $0, LC_\beta$ and simple propositional reasoning to obtain $\vdash \langle \beta | p \rangle$.

Completeness for transition formulae is based on the definability of spc. Hence:

Lemma 5.2. Let $\alpha, \beta \in \text{PROG}(t1)$, α atomic, $\langle \beta | \alpha \rangle$ and $p \in \text{EF}(t1)$. Then the spc of p w.r.t. $\langle \beta | \alpha \rangle$, $\text{sp}(p, \langle \beta | \alpha \rangle)$, is $F(t1)$ -definable.

Proof. As usual $M, \bar{\sigma} \models \text{sp}(p, \langle \beta | \alpha \rangle) \Leftrightarrow \exists (\sigma, \sigma') \in \text{Trs}(\beta, \alpha) M, \sigma \models p \wedge \sigma' = \bar{\sigma}$. Define $L \stackrel{D}{=} \bigwedge \{ \ell \vee \neg \ell \mid \ell \in \text{LAB}(p) \}$ and bring L into disjunctive normal form¹ $L_1 \vee \dots \vee L_n$. Each L_i is a conjunction of atomic formulae of the form ℓ or $\neg \ell$ and $\text{LAB}(L_i) = \text{LAB}(p)$. Obtain assertions $p_i \in \text{EF}(t)$ by substituting in p , true or false for any $\ell \in \text{LAB}(p)$ depending on whether ℓ or $\neg \ell$ appears in L_i . By definition $\text{LAB}(p_i) = \emptyset$. Clearly, $\vdash \langle \beta | L_i \rightarrow (p \leftrightarrow p_i) \rangle$ hence $\vdash \langle \beta | p \leftrightarrow (p_i \wedge L_i \mid i=1..n) \rangle$. Now, for any p_i , let \bar{p}_i denote the formula $p_i \wedge b$ in case $\alpha = (\ell.b?.\ell')$ and the formula $\exists y(p_i[y/x] \wedge x = e[y/x])$ in case $\alpha = (\ell.x := e.\ell')$ ($y \notin \text{Fvar}(p, e) \cup \{x\}$).

Let for any L_i , \bar{L}_i denote the $F(1)$ -formula $\bigwedge \{ \ell \mid \vdash \langle \beta | L_i \wedge \ell \rangle \} \wedge \bigwedge \{ \neg \ell \mid \vdash \langle \beta | L_i \wedge \neg \ell \rangle \}$. Then it is a simple exercise to show that $\text{sp}(p, \langle \beta | \alpha \rangle) \stackrel{D}{=} \alpha \wedge LC(\beta) \wedge \bigvee \{ L_i \wedge \bar{p}_i \mid i=1..n \}$.

This representation will be used in the sequel.

Lemma 5.3. $\text{PS}(M) \vdash \langle \beta | [L] \alpha [\text{sp}(L, \langle \beta | \alpha \rangle)] \rangle$, where α is atomic and $L \in \text{EF}(1)$.

Proof. Let \bar{L} denote the spc. There are two cases.

¹ Strictly speaking, true is a d.n.f. of L , too. What is meant, here and in the sequel, is to apply the well-known algorithm which (syntactically) transforms a formula into normal form.

(1) $\models \langle \beta | \alpha \rightarrow \neg L \rangle$: In the following derivation, the context is suppressed.

$$\frac{\begin{array}{c} [\text{false}] \alpha [L] \\ L \wedge \neg L \rightarrow \text{false} \\ \text{-----} \\ [L \wedge \neg L] \alpha [L] \end{array} \quad \alpha \rightarrow \neg L \quad \begin{array}{c} [\text{true}] \alpha [\alpha'] \\ \alpha' \rightarrow \alpha \vee L \\ \text{-----} \\ [\text{true}] \alpha [\alpha \vee L] \end{array}}{[L] \alpha [L]} S$$

(2) $\models \langle \beta | \alpha \rightarrow \neg L \rangle$: If $\models \langle \beta | L \rightarrow \text{false} \rangle$, then α has no output-state. This is only possible if $\alpha = (\lambda . \text{false} ? . \lambda')$. But then, \top can be used to obtain $\models \langle \beta | [L] \alpha [\text{false}] \rangle$. Hence, assume $\models \langle \beta | L \rightarrow \text{false} \rangle$. This means that there exists at least one transition $(\sigma, \sigma') \in \text{Trs}(\beta, \alpha)$ such that $\sigma' \models L$. Consequently, property (B) implies the following facts:

- (1) if $\langle \beta | \ell \models \alpha \rangle$ then $\models \langle \beta | [L] \alpha [\ell] \rangle \Rightarrow \models \langle \beta | [\ell] \alpha [\ell] \rangle$, and
- (2) if $\neg \langle \beta | \ell \models \alpha \rangle$ then $\models \langle \beta | [L] \alpha [\ell] \rangle \Rightarrow \models \langle \beta | \alpha' \rightarrow \ell \rangle$.

We intend to use the representation of \bar{L} as in the proof of lemma 5.2. So, bring L into disjunctive normal form $L_1 \vee \dots \vee L_n$ and let S denote either the formula b or $\exists y. x = e[y/x]$ ¹ depending on whether α is a test or an assignment. Then, using the notation of lemma 5.2: $\bar{L} = S \wedge \alpha' \wedge \text{ALC}(\beta) \wedge (\bar{L}_1 \vee \dots \vee \bar{L}_n)$. Now, it is easy to obtain $\models \langle \beta | [L] \alpha [S \wedge \text{ALC}(\beta)] \rangle$. Facts (1) and (2) above, imply that the conjunction \bar{L}_i can be restricted to the labels ℓ such that $\langle \beta | \ell \models \alpha \rangle$. But for such labels, if $\models \langle \beta | [L] \alpha [\bar{\ell}] \rangle$ then this can actually be derived, using F_1 or F_2 and the fact that $\models \langle \beta | L \rightarrow \bar{\ell} \rangle$ (a consequence of property (B)). This implies that $\models \langle \beta | [L] \alpha [\bar{L}_1 \vee \dots \vee \bar{L}_n] \rangle$ holds and hence that $\models \langle \beta | [L] \alpha [\bar{L}] \rangle$

Theorem 5.2. $M \models \langle \beta | [p] \alpha [q] \rangle \Rightarrow \text{PS}(M) \models \langle \beta | [p] \alpha [q] \rangle$, where $p, q \in F(1)$ and α is atomic.

Proof. Lemma 5.3 and theorem 5.1 (and an application of I).

Definition. Let $p \in F(1)$, $\alpha, \beta \in \text{PROG}(t1)$, α atomic and $\langle \beta | \alpha \rangle$. Then $L^0(p, \langle \beta | \alpha \rangle)$ is the set of labels whose truth-values are not determined by p ; i.e., it is the set $\{ \ell \mid \models \langle \beta | [p] \alpha [\ell] \rangle, \models \langle \beta | [p] \alpha [\neg \ell] \rangle \}$.

The following lemma states the basic property of such labels:

Lemma 5.4. Let $K \in F(1)$, $p, q \in F(t1)$, $\alpha, \beta \in \text{PROG}(t1)$, α atomic and $\langle \beta | \alpha \rangle$. Assume that $M \models \langle \beta | \alpha \rightarrow \neg K \rangle$, $\text{LAB}(p) = \emptyset$ and $\text{LAB}(q) \subseteq L^0(p, \langle \beta | \alpha \rangle)$. Then $M \models \langle \beta | [K \wedge p] \alpha [q] \rangle$ implies $M \models \langle \beta | [p] \alpha [q] \rangle$.

¹Note, that the pre-assertion L states nothing about the values of the program variables.

Proof. Let $L = L^0(p, \langle \beta | \alpha \rangle)$. Suppose $\models \langle \beta | [p] \alpha [q] \rangle$.

Then, there is a transition $(\sigma, \sigma') \in \text{Trs}(\beta, \alpha)$ such that $\sigma \models p$ and $\sigma' \models q$. By assumption, there is a state $\bar{\sigma}$ such that $\beta\text{-sg}(\bar{\sigma})$ and $\bar{\sigma} \models \alpha \wedge K$. As $\text{LAB}(p) = \emptyset$, $\bar{\sigma} \models \alpha \wedge K \wedge p$ holds, too. Now consider the states $\bar{\sigma}'$ such that $(\bar{\sigma}, \bar{\sigma}') \in \text{Trs}(\beta, \alpha)$. By definition of L , there is a $\bar{\sigma}'$ such that $\bar{\sigma}' \models (L \vee \text{Var}) \rightarrow \sigma' \models (L \vee \text{Var})$. Because $\text{LAB}(s) \subseteq L$, this means that $\bar{\sigma}' \models q$ and hence that $\models \langle \beta | [K \wedge p] \alpha [q] \rangle$; contradiction.

Finally, completeness for general transition formulae, $\langle \beta | [p] \alpha [q] \rangle$ ($p, q \in F(t1)$) can be shown. The intuition behind the proof is simple: Consider all (consistent) control locations in α . For each control location, L , derive from p the constraints on the variable values at this control point. Next, derive from q the constraints on the variable values in any state that can be reached by a single transition, when control resides at L . Finally, show that these constraints are met, if the transition is taken from a state satisfying p , too. Observe that this strategy corresponds to the way in which analogous program properties are proved in the temporal logic proof system of Manna and Pnueli [16].

Theorem 5.3. $M \models \langle \beta | [p] \alpha [q] \rangle \Rightarrow \text{PS}(M) \models \langle \beta | [p] \alpha [q] \rangle$, $p, q \in F(t1)$.

Proof. Because of the C_0 -rules, it suffices to show completeness for atomic α . Let $L \stackrel{D}{=} LD(\beta) \wedge \alpha \wedge \lambda (\ell \vee \bar{\ell} \mid \ell \in \text{LAB}(p) \setminus \text{LAB}(\beta))$. By definition $\models \langle \beta | \alpha \rightarrow L \rangle$. It suffices to show that $\models \langle \beta | [p \wedge L] \alpha [q] \rangle$, because:

$$\frac{\begin{array}{c} [\text{true}] \alpha [\alpha'] \\ \alpha' \rightarrow (\alpha \vee L) \\ \text{-----} \\ [\text{true}] \alpha [\alpha \vee L] \end{array} \quad \alpha \rightarrow L \quad [p \wedge L] \alpha [q]}{[p] \alpha [q]} S$$

Now, bring L into disjunctive normal form $L_1 \vee \dots \vee L_n$. Because of the C_V -rule, concentrate on $\models \langle \beta | [p \wedge L_i] \alpha [q] \rangle$. By definition, $\text{LAB}(p) = \text{LAB}(L_i)$. Hence, as in the proof of lemma 5.2, formulae $p_i \in F(t)$ can be (effectively) found, that satisfy $\text{LAB}(p_i) = \emptyset$ and $\models \langle \beta | L_i \rightarrow (p \wedge p_i) \rangle$. Consequently, it suffices to derive $\models \langle \beta | [p_i \wedge L_i] \alpha [q] \rangle$.

As the next step, define $M_i \stackrel{D}{=} \text{sp}(L_i, \langle \beta | \alpha \rangle)$. By theorem 5.2 (and I), $\models \langle \beta | [p_i \wedge L_i] \alpha [M_i] \rangle$. Bring M_i into conjunctive normal form $M_{i1} \wedge \dots \wedge M_{im}$. Let \bar{M}_i be the conjunction of those M_{ij} that are of the form ℓ or $\neg \ell$ (or true if no such M_{ij} exist). Clearly $\models \langle \beta | [p_i \wedge L_i] \alpha [\bar{M}_i] \rangle$. Again, construct formulae $q_i \in F(t1)$

that satisfy $\models \langle \beta | \bar{M}_i \rightarrow (q \leftrightarrow q_i) \rangle$ and $LAB(q_i) \cap LAB(\bar{M}_i) = \emptyset$. The following derivation justifies our concentrating on proving (I): $\models \langle \beta | [p_i \wedge L_i] \alpha [q_i] \rangle$.

$$\frac{\frac{[p_i \wedge L_i] \alpha [q_i] \quad [p_i \wedge L_i] \alpha [\bar{M}_i]}{[p_i \wedge L_i] \alpha [q_i \wedge \bar{M}_i]} \text{C} \wedge \quad \bar{M}_i \wedge q_i \rightarrow q}{[p_i \wedge L_i] \alpha [q]} \text{I}$$

If $\models \langle \beta | \delta \rightarrow L_i \rangle$ then (I) reduces to $\langle \beta | [false] \alpha [q_i] \rangle$ which is easily proved using AS or T. Next, assume that $\models \langle \beta | \delta \rightarrow L_i \rangle$. Moreover, suppose that (II) $LAB(q_i) \subseteq L_i^\circ(L_i, \langle \beta | \alpha \rangle)$. Then, lemma 5.4 applies: As $\models \langle \beta | [p_i \wedge L_i] \alpha [q_i] \rangle$, we obtain $\models \langle \beta | [p_i] \alpha [q_i] \rangle$ and hence $\models \langle \beta | sp(p_i, \langle \beta | \alpha \rangle) \rightarrow q_i \rangle$. Because $LAB(p_i) = \emptyset$, the representation of the spc in lemma 5.2, simplifies to $\alpha' \wedge LC(\beta) \wedge \bar{p}_i$, so that $\models \langle \beta | \alpha' \wedge \bar{p}_i \rightarrow q_i \rangle$ holds. This is the premiss of the AS or T rule. So, finally we obtain $\models \langle \beta | [p_i] \alpha [q_i] \rangle$, hence (using I) $\models \langle \beta | [p_i \wedge L_i] \alpha [q_i] \rangle$.

It remains to show that (II) actually holds; i.e., that $K \stackrel{D}{=} LAB(q) \setminus LAB(\bar{M}_i) \subseteq L_i^\circ(L_i, \langle \beta | \alpha \rangle)$: Take any $\mathcal{R} \in K$. Because $\mathcal{R} \in LAB(\bar{M}_i)$, $\models \langle \beta | \bar{M}_i \rightarrow \mathcal{R} \rangle$ and $\models \langle \beta | \bar{M}_i \rightarrow \mathcal{R} \rangle$ hold and therefore $\models \langle \beta | M_i \rightarrow \mathcal{R} \rangle$ and $\models \langle \beta | M_i \rightarrow \mathcal{R} \rangle$. By definition of M_i , this means that $\mathcal{R} \in L_i^\circ(L_i, \langle \beta | \alpha \rangle)$.

6. DISCUSSION.

The paper introduces compositionality into proofs of temporal properties of concurrent programs. The key-observation to view programs as sets of sequences of atomic actions, leads to the proposal of Transition Logic as supporting compositional reasoning. Essential features of the deductive system for this logic - as exemplified in the completeness proof - are (1) the axiomatization of the flow-of-control and (2) the expressibility of the strongest post-condition of atomic programs.

Transition Logic is developed here, for a rudimentary language, i.e., for regular programs with merge over assignments and tests. It would be interesting to consider languages with more intricate syntactic structure and introduce, e.g., recursion, synchronization commands or a construct, $\langle \alpha \rangle$, to execute arbitrary programs α , as indivisible actions. The logic can be straightforwardly extended to deal with the latter two constructs; no basic difficulties are envisaged in treating recursion.

Another possibility is introducing communication actions. Of course, there is an easy way to deal with such commands, as is illustrated by Lamport and Schneider [13] in the context of CSP. However, that

solution ignores the interesting fact that for such languages (without variable sharing) there is a clear distinction between the internal behaviour of a program and its external behaviour, i.e., the sequences of communications it is willing to participate in. A distinction that should be reflected in Transition Logic, thus making a connection with [19] and [23] on partial correctness of networks of communicating processes.

Transition Logic concentrates on (input-output) properties of atomic actions. While this suffices for proving any temporal property, it is at the same time unsatisfactory because one would like to reason about temporal properties in a more direct way. The basic problem is to define the execution sequences of programs when interleaving occurs. In the current set-up, this is impossible to do in a compositional way. Hence, satisfaction of temporal formulae cannot be defined. A possible extension would be to first formulate a set of assumptions, specifying the way the environment of a program may alter the program state. Relative to such assumptions, execution sequences of programs can be defined and temporal properties can be reasoned about using temporal modalities. Such an extension of Transition Logic, would be a first step in providing Lamport's modular specification method [14] with a compositional deductive system.

Finally, an obvious question is, whether the logic can be extended to deal with just and fair executions of programs. A problem arises here: For a transition to be justly taken, it should eventually become continuously enabled. Hence, it should be certain that no move of the environment - of which nothing is known - can disable this transition from some time onwards. Similar reasoning applies to fairness, as a move need not be fairly taken, if that move eventually becomes continuously disabled. It follows that justness and fairness, too, only make sense relative to assumptions about the behaviour of the environment.

ACKNOWLEDGEMENTS.

I thank Willem Paul de Roever for keeping me on course while developing these ideas and for the unfailing way in which he, time and again, homed in on the weak spots in both the proofs and the presentation. This paper has benefited, too, from some remarks by Amir Pnueli.

- [1] K. R. ABRAHAMSON. Modal logic of concurrent nondeterministic programs. LNCS70, G. Kahn ed., pp.21-33, Springer Verlag, New York, 1979.
- [2] K. R. APT, C. DELPORTE. An axiomatization of the intermittent assertion method using temporal logic (extended abstract). LNCS154, J. Diaz ed., pp.15-27, Springer Verlag, New York, 1983.
- [3] H. BEKIC. Towards a mathematical theory of processes. Report TR25.125, IBM Laboratory, Vienna, 1971.
- [4] E. BEST. A relational framework for concurrent programs using atomic actions. Proc. IFIP TC2 Conference, D. Bjørner ed., North-Holland, 1982.
- [5] P. and R. COUSOT. On the soundness and completeness of generalized Hoare logic. Report CRIN-82-P093, Nancy, 1982.
- [6] E. W. DIJKSTRA. A discipline of programming. Prentice Hall, 1979.
- [7] R. W. FLOYD. Assigning meaning to programs. Proc. Symp. in Appl. Math.-19, J. T. Schwartz ed., pp.19-32, AMS, 1967.
- [8] D. GRIES. The science of programming. Springer Verlag, New York, 1982.
- [9] D. HAREL. First order dynamic logic. LNCS68, Springer Verlag, New York, 1979.
- [10] C. A. R. HOARE. An axiomatic basis for computer programming. CACM12-10, pp.576-580, 1969.
- [11] T. JANSEN, P. van EMDE BOAS. Some observations on compositional semantics. LNCS131, D. Kozen ed., pp.137-170, Springer Verlag, New York, 1982.
- [12] L. LAMPORT. The "Hoare logic" of concurrent programs. Acta Inf. 14, pp.21-37, 1980.
- [13] L. LAMPORT, F. B. SCHNEIDER. The "Hoare logic" of CSP, and all that. Comp. Sci. Lab., SRI, 1982.
- [14] L. LAMPORT. Specifying concurrent program modules. TOPLAS 5-2, pp.190-223, 1983.
- [15] D. LEHMAN, A. PNUELI, J. STAVI. Impartiality, justice and fairness: the ethics of concurrent termination. LNCS115, pp.264-278, Springer Verlag, New York, 1981.
- [16] Z. MANNA, A. PNUELI. How to cook a temporal proof system for your pet language. Proc. POPL, Austin, ACM, 1983.
- [17] Z. MANNA, A. PNUELI. Proving precedence properties - the temporal way. LNCS154, J. Diaz ed., Springer Verlag, New York, 1983.
- [18] R. MILNER. An approach to the semantics of parallel programs. Proc. of the convegno di Informatica Teorica, Pisa, 1973.
- [19] J. MISRA, K. M. CHANDY. Proofs of networks of processes. IEEE SE-N, SE-7, no.4, pp.417-427, 1981.
- [20] S. OWICKI, D. GRIES. An axiomatic proof technique for parallel programs I. Acta Inf. 6, pp.319-340, 1976.
- [21] W. REISIG. Partial order semantics versus interleaving semantics for CSP-like languages and its impact on fairness. Proc. ICALP84, LNCS, Springer Verlag, New York, 1984.
- [22] A. SALWICKI, T. MULDER. On the algorithmic properties of concurrent programs. LNCS125, E. Engeler ed., pp.169-177, Springer Verlag, New York, 1981.
- [23] J. ZWIERS, A. de BRUIN, W. P. de ROEVER. A proof system for partial correctness of dynamic networks of processes (extended abstract). Proc. 2nd Workshop on Logics of Programs, LNCS, D. Kozen, E. Clarke eds., Springer Verlag, New York, 1984.