Decision procedures for Time and Chance
(Extended abstract)

Sarit Kraus and Daniel Lehmann

Institute of Mathematics and Computer Science,
Hebrew University, Jerusalem 9104 (Israel)

Abstract:

Decision procedures are provided for checking the satisfiability of a formula in each of the three systems TCg, TCb and TCf defined in [LS]. The procedures for TCg and TCf run in non-deterministic time $2^{2^{cn}}$ where n is the size of the formula and c is a constant. The procedure for TCb runs in non-deterministic time $2^{2^{cn^2}}$. A deterministic exponential lower bound is proved for the three systems. All three systems are also shown to be PSPACE-hard using results of [SC]. Those decision procedures are not as efficient as the deterministic (one or two)-exponential time procedures proposed in [BMP] and [EH1] for different logics of branching time that are weaker than ours in expressive power. No elementary decision procedure is known for a logic of branching time that is as expressive as ours. The decision procedures of the probabilistic logics of [HS] run in deterministic exponential time but their language is essentially less expressive than ours.

1. Introduction

In [LS] an extension of the propositional logic of linear time was described and claimed useful for reasoning in an uncertain world, i.e. a world in which time flows and in which the transition from one instant of time to the next is probabilistic in nature. Three logical systems of axioms and rules of inference were proposed and proved deductively complete with respect to different classes of models. One of these systems, TCb, was claimed to be most adequate for stating and proving properties about a large interesting family of probabilistic algorithms, such as those of [CLP], [LR] and [Ra3]. All three systems had impressive expressive power, since they are similar in expressive power to the system CTL^* described, for branching time, in [EH2]. This last system is the most expressive of those defined in [EH2] and no elementary decision procedure is known for it. It is even possible that methods similar to ours could yield efficient decision procedure for CTL^*. We assume familiarity with the basic notions and notations of [LS] (see ICALP 83, Springer Verlag Lecture Notes in C.S. for an extended abstract): the models, the language (Γ) and validity.

2. The logical systems

In [LS] three different logical systems: TCg, TCb and TCf were proposed, each one of them corresponding to one of the notions of γ-validity defined there (γ is g, b or f). The symbol $\underset{\gamma}{\vdash}$ denotes provability in the system corresponding to γ. The following completeness result was proved.

Theorem 1: For any $\gamma \in \{g,b,f\}$ and for any $\alpha \in \Gamma$, $\underset{\gamma}{\models} \alpha \Leftrightarrow \underset{\gamma}{\vdash} \alpha$.

3. The lower bounds

Our first result is that all three systems are PSPACE-hard. In [SC] the temporal logic of linear time is shown to be PSPACE-hard. Let α be a formula of the temporal logic of linear time. It is also a formula of Γ. If α is valid in the logic of linear time then it holds for every path of a g-model and it is g-valid. Any g-valid formula is b-valid and any b-valid formula is f-valid. Let us call a g-model deterministic if for any state s there is a state t such that $p(s,t) = 1$. If a formula α is f-valid then it is, in particular, satisfied by all finite deterministic models. But finite deterministic models are exactly finite linear models. Since the logic of linear time has the finite model property any formula of the logic of linear time that is f-valid is valid.

Our second result is that no decision procedure for any of our three systems can work in less than exponential time.

Theorem 2: There is a constant $c > 1$ such that the validity problems for TCg, TCb and TCf are not members of $DTIME(c^n)$, where n is the size (#) of the formula.

The proof is by reduction to one letter PDL. Let us call PDL^- the restriction of PDL obtained by restricting oneself to a single basic program ($\vdash$) and not allowing sequential composition (;) or non-deterministic choice ($\cup$). Notice that PDL^- has essentially only two programs: $\vdash$ and $\vdash^*$. In [FL], it is shown that the validity problem for PDL^- requires exponential time. Define UB^- as the subset of UB (see [BMP]) defined when one allows $\forall G$ and $\forall X$ as the only temporal connectives (i.e. $\forall F$ and $\exists G$ are disallowed, classical connectives are allowed). The simple translation of PDL^- to UB^- that translates $[\vdash]$ by $\forall X$ and $[\vdash^*]$ by $\forall G$ shows that the validity problem for UB^- requires exponential time (this result seems to be part of the folklore of the subject). Our remark is essentially that UB^- is so restricted that the branching time interpretation is equivalent to the probabilistic interpretation.

More precisely, we translate UB^- into a subset Γ' of Γ by translating $\forall X$ by $\nabla \bigcirc$ and $\forall G$ by $\nabla \Box$. If α is a formula of UB^-, its translation in Γ' will be denoted by $\bar{\alpha}$. Let now $\mathcal{U}$ be a g-model (probabilistic), we shall denote by $\mathcal{U}'$ the branching time model obtained from $\mathcal{U}$ by allowing those transitions that have positive probability in $\mathcal{U}$ (and forgetting the probabilities).

Lemma 1: If α is a formula of UB^- and $\mathcal{U}$ is a g-model, then $\mathcal{U} \not\models \bar{\alpha} \implies \mathcal{U} \models \neg\bar{\alpha}$.

Proof: An easy induction on the structure of α, reasoning at the same time on a formula and its negation.

Lemma 2: Let α be a formula of UB^- and $\mathcal{U}$ be a g-model, then $\mathcal{U}' \models \alpha \iff \mathcal{U} \models \bar{\alpha}$.

Proof: An easy induction on the structure of α.

Lemma 3: Let α be a formula of UB^-, then the following propositions are equivalent:
a) α is valid for branching time logic
b) $\bar{\alpha}$ is g-valid
c) $\bar{\alpha}$ is b-valid
d) $\bar{\alpha}$ is f-valid

Proof: a) $\implies$ b) by lemma 2. b) $\implies$ c) and c) $\implies$ d) are obvious. Suppose d) is true. By lemma 2 for any f-model $\mathcal{U}$ we have $\mathcal{U}' \models \alpha$. Given a finite model for branching time $\mathcal{V}$, one may always give numeric positive probabilities to the transitions of $\mathcal{V}$ to obtain an f-model $\mathcal{U}$ such that $\mathcal{U}' = \mathcal{V}$. We conclude that α is satisfied in all finite models for branching time. Since UB^- enjoys the finite model property, we may conclude that α is valid. **Q.E.D.**

4. Upper bounds

All three completeness results of [LS] proceed by building a "universal" model, i.e. a model in which points are labeled by traces of a suitable size and every trace of that size labels some point of the model that is built in such a way that the standard generic paths out of a point satisfy exactly those formulas that appear in their label. In all three cases those models, though infinite in the cases g and b, may be finitely (and succinctly) described. The basic idea is to collect a set of properties enjoyed by the universal model that is rich enough to guarantee that the standard generic paths out of a point satisfy exactly those formulas that appear in their label. We may know guess a model satisfying those properties and look out for a label containing formula $\nabla\alpha$. If one is found then α is satisfiable. If no model contains such a label, then the universal model contains no such label and α is not satisfiable.

5. A decision procedure for TCg

First we shall define the small set of formulas to be considered.

Definition 1: If $\alpha \in \Gamma$, we define $\mathcal{C}l^-(\alpha)$ to be the smallest subset of Γ satisfying:
a) $\nabla\alpha \in \mathcal{C}l^-(\alpha)$
b) $\mathcal{C}l^-(\alpha)$ is closed under sub-formulas
c) $\mathcal{C}l^-(\alpha)$ is closed under negation
d) $p \in \mathcal{C}l^-(\alpha)$ (p stands for a propositional variable) $\implies \nabla p \in \mathcal{C}l^-(\alpha)$
e) $\Box b \in \mathcal{C}l^-(\alpha) \implies \bigcirc\Box b \in \mathcal{C}l^-(\alpha)$
f) $b\, \mathcal{U}\mathbf{ntil}\, c \in \mathcal{C}l^-(\alpha) \implies \bigcirc\left(b\, \mathcal{U}\mathbf{ntil}\, c\right) \in \mathcal{C}l^-(\alpha)$ and $\neg\Box\neg c \in \mathcal{C}l^-(\alpha)$.

We shall also assume that double negations are automatically removed. It is easy to see that the size of $\mathcal{C}l^-(\alpha)$ is linear in the length of α.

Similarly to what must be done in the completeness proof we need to define a set of formulas that is larger than $\mathcal{C}l^-(\alpha)$. The labels considered will be subsets of this larger set of formulas.

Definition 2: A formula b is in $\mathcal{C}l(\alpha)$ iff it is either in $\mathcal{C}l^-(\alpha)$ or of the form $\nabla\left(b_1 \vee \cdots b_k\right)$ where $b_i \in \mathcal{C}l(\alpha)$, $i = 1, \cdots, k$.

Notice that $\mathcal{C}l(\alpha)$ contains only 2^{cn} formulas. Now, we want to restrict our attention to those subsets of $\mathcal{C}l(\alpha)$ that are reasonable candidates for labels. Notice that our requirements deal only with the formulas of $\mathcal{C}l^-(\alpha)$.

Definition 3: $D \subseteq \mathcal{C}l(\alpha)$ is a standard set (for α) if it satisfies:
a) if $b \in \mathcal{C}l^-(\alpha)$, $\neg b \in D \iff b \notin D$
b) if $b \vee c \in \mathcal{C}l^-(\alpha)$, then $b \vee c \in D \iff b \in D$ or $c \in D$
c) $\Box b \in D \iff b \in D$ and $\bigcirc\Box b \in D$
d) if $\nabla b \in \mathcal{C}l^-(\alpha)$, then $\nabla b \in D \implies b \in D$
e) $p \in D \implies \nabla p \in D$
f) $b\, \mathcal{U}\mathbf{ntil}\, c \in D \implies \neg\Box\neg c \in D$
g) $b\, \mathcal{U}\mathbf{ntil}\, c \in D \iff c \in D$ or $b \in D$ and $\bigcirc\left(b\, \mathcal{U}\mathbf{ntil}\, c\right) \in D$

We define $\mathfrak{D}$ to be the set of all the standard subsets of $\mathcal{C}l(\alpha)$
Remark: $\mathfrak{D}$ contains at most $2^{2^{c\#(\alpha)}}$ elements, for some constant c
We define, now, the successor relation R on $\mathfrak{D}$. Notice again that only formulas of $\mathcal{C}l^-(\alpha)$ are considered.

Definition 4: Let D_1 and D_2 be standard sets (i.e $D_i \in \mathfrak{D}$) we say that $D_1 R D_2$ iff for all $b \in \Gamma$ such that $\bigcirc b \in \mathcal{C}l^-(\alpha)$, we have $\bigcirc b \in D_1 \iff b \in D_2$.

Similarly we define a relation E on $\mathfrak{D}$. It is crucial here that our requirement is only for formulas of $\mathcal{C}l^-(\alpha)$.

Definition 5: Let D_1 and D_2 be in $\mathfrak{D}$. We say that $D_1 E D_2$ iff for all $b \in \Gamma$ such that $\nabla b \in \mathcal{C}l^-(\alpha)$, we have $\nabla b \in D_1 \iff \nabla b \in D_2$.

Now we may define the pseudo-models for the case g.

Definition 6: A pseudo-model (for α) is a triple $\langle \mathcal{W}, r, \sim \rangle$ satisfying:
1) $\mathcal{W}$ is a set of standard sets (i.e $\mathcal{W} \subseteq \mathfrak{D}$)
2) r is a binary relation on $\mathcal{W}$ that is contained in R, i.e. $D_1 r D_2 \implies D_1 R D_2$
3) $\sim$ is an equivalence relation on $\mathcal{W}$ that is contained in E
4) $(r \sim) \subseteq (\sim R)$ (i.e if $D_1, D_2, D_3 \in \mathcal{W}$ such that $D_1 r D_2$, and $D_2 \sim D_3$, then there exists $D_4 \in \mathcal{W}$ such that $D_1 \sim D_4$, and $D_4 R D_3$).
5) for all $D \in \mathcal{W}$, there exists a $D' \in \mathcal{W}$ such that $D r D'$
6) for all $D \in \mathcal{W}$ and for all $\Box b \in \mathcal{C}l^-(\alpha)$, if $\Box b \notin D$ then there exists a $D' \in \mathcal{W}$ such that $D r^* D'$ and $b \notin D'$ where r^* denotes the reflexive and transitive closure of r
7) for all $D \in \mathcal{W}$ and for all $\nabla b \in \mathcal{C}l^-(\alpha)$, if $\nabla b \notin D$ then there exists a $D' \in \mathcal{W}$ such that $D \sim D'$ and $b \notin D'$

Lemma 4: If α is a g-satisfiable formula, then there are a pseudo-model $\langle \mathcal{W}, r, \sim \rangle$ for α and a $D \in \mathcal{W}$ such that $\nabla\alpha \in D$.

Proof: Let $\mathcal{A}$ be the set of all D's for which there is a g-consistent and complete theory T such that $D = T \cap Cl(\alpha)$. Notice that here we consider $Cl(\alpha)$ not $Cl^{-}(\alpha)$. One may see that $\langle \mathcal{A}, s, e^{*} \rangle$, is a pseudo-model for α, where s and e are defined below and e^{*} is the reflexive and transitive closure of e. The relation s is defined by $D_1 s D_2$ iff there is a g-consistent and complete theory T_1 such that $D_1 = T_1 \cap Cl(\alpha)$ and $D_2 = T_1^{+} \cap Cl(\alpha)$ where + is defined as in [LS]. The relation e is defined by $D_1 e D_2$ iff there are g-consistent and complete theories T_1 and T_2 such that $T_1 \equiv T_2$, $D_i = T_i \cap Cl(\alpha)$ for $i=1,2$, where $\equiv$ is defined as in [LS]. The only delicate points are verifying conditions 4 and 6 of definition 6. Condition 6 is proved exactly as in section 11.6 of [LS]. We shall sketch the verification of condition 4. We should show that $s e^{*} \subseteq e^{*} R$. Indeed we shall show that $s e^{*} \subseteq e R$. Suppose $D_1 s D_2$ and $D_2 e^{*} D_3$. We know that there is a g-consistent and complete theory T_1 such that $D_1 = T_1 \cap Cl(\alpha)$, $D_2 = T_1^{+} \cap Cl(\alpha)$, and for any $\nabla b \in Cl(\alpha)$ (the larger set, not only $Cl^{-}(\alpha)$), $\nabla b \in D_2 \Longleftrightarrow \nabla b \in D_3$. Let $U = \{b \mid \nabla b \in T_1\} \cup \{\bigcirc d \mid d \in D_3 \cap Cl^{-}(\alpha)\}$. One easily sees that U is g-consistent. Now one may extend U to a g-consistent and complete theory T_4 ($T_4 \equiv T_1$) and take $D_4 = T_4 \cap Cl(\alpha)$.

One also sees that if α is satisfiable so is $\nabla \alpha$ and that there is a g-consistent and complete theory that contains $\nabla \alpha$ and therefore a $D \in \mathcal{A}$ such that $\nabla \alpha \in D$. **Q.E.D.**

We must now show that in a pseudo-model, the relations r and $\sim$ have all the properties needed to perform the construction of a g-model that to prove that it enjoys properties similar to those of the universal model built in [LS] (we are able to carry out the proof of lemma 21 of [LS]). First we prove the equivalent of lemma 8 of [LS].

Lemma 5: Let $\langle \mathcal{W}, r, \sim \rangle$ be a pseudo-model and let $k \geq 0$ and $D_0, D_1, \cdots, D_k \in \mathcal{W}$ such that $\forall i, 0 \leq i < k$, $D_i \sim r D_{i+1}$. There are S_i, for $i = 0, \cdots, k$, such that $S_i \in \mathcal{W}$ and

1) $S_i \sim D_i, \forall i,\ 0 \leq i \leq k$
2) $S_i R S_{i+1}, \forall i,\ 0 \leq i < k$

Proof: By induction on k. For $k = 0$ the result is clear. For $k > 0$ use the induction hypothesis for the sequence $D_1, \ldots D_k$ to find $S_1, \cdots, S_k$. Now $D_0 \sim r \sim S_1$ and by definition 6 parts 4 and 3 $D_0 \sim R S_1$. **Q.E.D.**

We may know prove the main lemma.

Lemma 6: If there exist a pseudo-model $\langle \mathcal{W}, r, \sim \rangle$ and a $C \in \mathcal{W}$ such that $\nabla \alpha \in C$, then there exists a model for α.

Proof: The model $\mathcal{U} = \langle \mathcal{S}, u, l, p \rangle$, is defined the following way:

1) $\mathcal{S} = \mathbb{N} \times \mathcal{W}$
2) $u = \langle 0, C \rangle$
3) $l(\langle i, D \rangle) = \{p \mid p \in D\}$
4) Let us say, first, that the only transitions with non-zero probability are those that increase the first coordinate by one and use $\sim r$ to move along the second coordinate. In other terms, $p(\langle i, D \rangle, \langle j, F \rangle) \neq 0 \Longleftrightarrow j = i+1$ and $D \sim r F$. If $D r F$ let us call the transition from $\langle i, D \rangle$ to $\langle i+1, F \rangle$ a normal transition. A transition of positive probability that is not normal, will be called exceptional.
Our goal is to give increasing (with the first coordinate) weight to the normal transitions, and ensure that, with probability one, after a certain time, only normal transitions occur. Therefore we choose a sequence α_i of real numbers between 0 and 1, such that $\prod_{i=0}^{\infty} \alpha_i > 0$. From state $\langle i, D \rangle$, we give equal probability to all normal transitions, so as to give them total weight α_i, and equal probability to all exceptional transitions, so as to give them total weight $1 - \alpha_i$.

If $\sigma = \langle j_i, D_i \rangle$ is a sequence of states (any sequence) we say that σ is standard iff $D_i R D_{i+1}$ for any $i \geq 0$. A sequence of normal transitions is standard. We say that σ is ultimately standard if for any $i \geq 0$ $D_i \sim R D_{i+1}$ and σ has a standard tail. Our goal is to show that the model $\mathcal{U}$ satisfies α. With the changes just made lemma 19 of [LS] still holds. The proof is unchanged.

If $\sigma = \{\langle k_i, D_i \rangle\}$ and $\tau = \{\langle j_i, F_i \rangle\}$ are sequences of states, we say that σ and τ are equivalent and write $\sigma E \tau$ iff for every $i \in \mathbb{N}$, $D_i E F_i$. Generic sequences are defined as in [LS], replacing ρ by r. Lemma 20 of [LS] still holds, and for the same reasons.
Notice that, if a sequence of states $\sigma = \{\langle k_i, D_i \rangle\}$ is generic and if $F \in \mathcal{W}$ appears an infinite number of times in the sequence (as a second component) then every $G \in \mathcal{W}$ such that $F r^{*} G$ also appears an infinite number of times.

Our basic result concerning $\mathcal{U}$ is the following.

Lemma 7: Let $b \in Cl^{-}(\alpha)$, σ a generic standard sequence of states (of $\mathcal{U}$), and τ and τ' two equivalent sequences of states, then

a) $b \mid_{\mathcal{U}}^{\sigma} = \mathit{true} \Longleftrightarrow b \in D_0$, where $\sigma_i = \langle k_i, D_i \rangle$

b) $b \mid_{\mathcal{U}}^{\tau} = b \mid_{\mathcal{U}}^{\tau'}$.

Proof: The proof is by induction on the size of b, i.e. $\#(b)$, at each induction step, we prove a) first, and then b).

$b = p$

a) $p \mid_{\mathcal{U}}^{\sigma} = \mathit{true} \Longleftrightarrow p \in l(\sigma_0) \Longleftrightarrow p \in D_0$

b) $p \mid_{\mathcal{U}}^{\tau} = \mathit{true} \Longleftrightarrow p \in F$, where $\tau_0 = \langle i, F \rangle \Longleftrightarrow \nabla p \in F$, by conditions d) and e) of definition 3 $\Longleftrightarrow \nabla p \in F'$, where $\tau'_0 = \langle i, F' \rangle$, since $F E F'$ by definition 5 $\Longleftrightarrow p \in F'$, by condition e of definition 3 $\Longleftrightarrow p \mid_{\mathcal{U}}^{\tau'} = \mathit{true}$.

$b = \neg c$

By condition a) of definition 3.

$b = c \vee d$

By condition b) of definition 3.

$b = \bigcirc c$

a) $\bigcirc c \mid_{\mathcal{U}}^{\sigma} = \mathit{true} \Longleftrightarrow c \mid_{\mathcal{U}}^{\sigma^1} = \mathit{true} \Longleftrightarrow c \in D_1 \Longleftrightarrow \bigcirc c \in D_0$, since $D_0 R D_1$ and by definition 4 (since σ is a standard sequence).

b) Obvious.

$b = \Box c$

a) $\Box c \in D_0 \Longrightarrow \forall i \in \mathbb{N}, \Box c \in D_i$, using condition c) of definition 3 and definition 4 since σ is a standard sequence $\Longrightarrow \forall i \in \mathbb{N}, c \in D_i$ by condition c) of definition 3 $\Longrightarrow \forall i \in \mathbb{N}, c \mid_{\mathcal{U}}^{\sigma^i} = \mathit{true}$, by the induction hypothesis and since the end part of a generic sequence is generic $\Longrightarrow \Box c \mid_{\mathcal{U}}^{\sigma} = \mathit{true}$.

Suppose now that $\Box c \notin D_0$. It follows that

$\exists i \in \mathbb{N}, c \notin D_i$ or $\forall i \in \mathbb{N}, \Box c \notin D_i$ (by induction on i using condition c of definition 3 and definition 4). We want to show that the first alternative is true. In the last case, there is a member of $\mathcal{W}$, say F, that appears an infinite number of times, and $\Box c \notin F$. Since the sequence σ is generic, any $G \in \mathcal{W}$ such that $Fr^* G$ appears an infinite number of times in the sequence. By condition 6 of definition 6 then, there is an $i \in \mathbb{N}$, for which $c \notin D_i$. We conclude that $\exists i \in \mathbb{N}$ such that $c \notin D_i$ and by the induction hypothesis $c \mid_{\mathcal{U}}^{\sigma^i} = \textbf{\textit{false}}$ and $\Box c \mid_{\mathcal{U}}^{\sigma} = \textbf{\textit{false}}$.

b) Obvious.

$b = c\, \mathcal{U}\mathbf{ntil}\, d$

a) Suppose $c\, \mathcal{U}\mathbf{ntil}\, d \in D_0$. By condition f) of definition 3 $\neg\Box\neg d \in D_0$ and by the induction hypothesis there exists an index k for which $d \in D_k$. Let i be the smallest such k. Using definition 4 and property g) of definition 3 one sees that for any j, $j < i$, we have $c \in D_j$. We conclude by the induction hypothesis.

Suppose now that $c\, \mathcal{U}\mathbf{ntil}\, d \notin D_0$. If for every k, we have $d \notin D_k$ we conclude straightforwardly by the induction hypothesis. Suppose i is the smallest index for which $d \in D_i$. Using definition 5 and property g) of definition 3, one sees that there is an index j, $j < i$ for which $b \notin D_j$ and we conclude by the induction hypothesis.

b) Obvious.

$b = \nabla c$

a) Suppose $\nabla c \in D_0$. We want to show that, for almost all sequences $\tau \in P_{\sigma_0}$, we have $c \mid_{\mathcal{U}}^{\tau} = \textbf{\textit{true}}$.
Since almost all sequences of P_{σ_0} are generic (by lemma 20 of [LS]), it is enough, by lemma 19 of [LS] to show that if τ is generic and ultimately standard, then $c \mid_{\mathcal{U}}^{\tau} = \textbf{\textit{true}}$. Let τ be standard from index i on. By lemma 5, there are $S_m \in \mathcal{W}$, $0 \le m < i$ such that $S_m \sim D_m$, for $m = 0, \cdots, i-1$ and $S_m\, R\, S_{m+1}$, for $m = 0, \cdots, i-2$. Let us define the sequence τ' by: $\tau'_m = \langle k_m, S_m \rangle$, for $m = 0, \cdots, i-1$ and $\tau'_m = \tau_m$, for $m \ge i$. The sequence τ' is equivalent to τ. Therefore, by the induction hypothesis, part b), $c \mid_{\mathcal{U}}^{\tau} = c \mid_{\mathcal{U}}^{\tau'}$. It is generic since it is identical with τ from index i on and since τ is generic. It is also standard. Since τ' is standard and generic, we conclude, by the induction hypothesis part a), that: $c \mid_{\mathcal{U}}^{\tau'} = \textbf{\textit{true}} \iff c \in D'$, where $\tau'_0 = \langle k, D' \rangle$. But since $\tau\, E\, \tau'$, and $\nabla c \in D_0$, we conclude that $c \in D'$, by definition 5 and condition d) of definition 3.

Suppose now that $\nabla c \notin D_0$. We must find a set Q of sequences that begin at σ_0 and do not satisfy c, such that Q has a positive measure. Remember that $\sigma_0 = \langle k_0, D_0 \rangle$. By condition 7 of definition 6, there is a $F \in \mathcal{W}$ such that $D_0 \sim F$ and $c \notin F$. Let F' be any member of $\mathcal{W}$ such that $F\, r\, F'$. Condition 5 of definition 6 ensures the existence of such an F'. We have $D_0 \sim r\ F'$, and by the definition of our model $p(\langle k_0, D_0 \rangle, \langle k_0+1, F' \rangle) > 0$. Let us define Q as the set of all sequences τ such that: $\tau_0 = \sigma_0$, $\tau_1 = \langle k_1, F' \rangle$ and the sequence τ^1 is standard and generic. By lemma 19 of [LS], we have $\bar{p}_{\sigma_0}(Q) > 0$. It is left to us to show that no sequence of Q satisfies c. Let τ be any sequence of Q. Let τ' be the sequence $\langle k_0, F \rangle, \tau_1, \tau_2, \cdots$. It is a generic standard sequence. By the induction hypothesis, part a) we have $c \mid_{\mathcal{U}}^{\tau'} = \textbf{\textit{false}}$. But since $D_0 \sim F$, $\tau\, E\, \tau'$. We conclude, by the induction hypothesis, part b), that $c \mid_{\mathcal{U}}^{\tau} = \textbf{\textit{false}}$.

b) Let σ be any generic standard sequence starting at τ_0. Since the truth value of ∇c depends only on the first state of the sequence we have $\nabla c \mid_{\mathcal{U}}^{\tau} = \nabla c \mid_{\mathcal{U}}^{\sigma}$. By the induction hypothesis, part a) just above: $\nabla c \mid_{\mathcal{U}}^{\sigma} = \textbf{\textit{true}} \iff \nabla c \in D$, where $\tau_0 = \sigma_0 = \langle k_0, D \rangle$. Similarly $\nabla c \mid_{\mathcal{U}}^{\tau'} = \textbf{\textit{true}} \iff \nabla c \in D'$, where $\tau'_0 = \langle k_0, D' \rangle$. But, since $D\, E\, D'$, $\nabla c \in D \iff \nabla c \in D'$. **Q.E.D.**

We may now conclude the proof of lemma 6. Since $\nabla \alpha \in D_0$ and D_0 is the initial state of our model $\mathcal{U}$, we conclude from lemmas 19 of [LS] and lemma 7 that $\bar{p}_u(\alpha) = 1$ and $\mathcal{U} \models \alpha$. **Q.E.D.**

Theorem 3: Satisfiability in TCg is decidable in $NTIME(2^{2^{cn}})$ for some constant $c \ge 0$ where n is the size (#) of the formula .

Proof: Algorithm 1.
To test if α is satisfiable:
(1) Guess a pseudo-model $\langle \mathcal{W}, r, \sim \rangle$.
(2) Test if $\exists\, D \in \mathcal{W}$ such that $\nabla \alpha \in D$.

Notice that the task of checking if a triplet is a pseudo-model is polynomial in the size of $\mathcal{W}$ and that the correctness of the algorithm is clear from lemmas 4 and 6. **Q.E.D.**

6. A decision procedure for TCf

Since we must consider terminal theories and their traces we must enlarge somewhat the set of formulas to be considered, therefore we replace definition 1 by the following.

Definition 7: Let $\alpha \in \Gamma$, we define $\mathcal{Cl}^-(\alpha)$ to be the smallest subset of Γ satisfying conditions a)-d) of definition 1 and
e) if $b \in \mathcal{Cl}^-(\alpha)$ that does not begin by $\Box$ then $\Box\neg\Box b \in \mathcal{Cl}^-(\alpha)$.
We shall assume that two simplifications are made automatically: $\Box\Box$ is simplified to $\Box$ and $\neg\Box\neg\Box\neg\Box$ is simplified to $\Box\neg\Box$ With those assumptions one may see that the size of $\mathcal{Cl}^-(\alpha)$ remains linear in the length (#) of α. Let α be a formula in TCf ,the closure of α, ($\mathcal{Cl}(\alpha)$) and standard sets for α are defined as in the case of TCg.

We need to define terminal subsets of $\mathcal{Cl}(\alpha)$.

Definition 8: $D \subseteq \mathcal{Cl}(\alpha)$ is a terminal set (for α) if it satisfies:
if $b \in \mathcal{Cl}^-(\alpha)$, $b \notin D \implies \Box\neg\Box b \in D$

We may now define pseudo-models.

Definition 9: An f-pseudo-model for α is a triplet $\langle \mathcal{W}, r, \sim \rangle$ satisfying:
1) $\mathcal{W} \subseteq \{0,1\} \times \mathfrak{D}$ such that $\langle 1, D \rangle \in \mathcal{W} \implies D$ is terminal
2) r is a binary relation on $\mathcal{W}$ such that $\langle i, D_1 \rangle\, r\, \langle j, D_2 \rangle \implies D_1\, R\, D_2$ and such that $\langle 1, D_1 \rangle\, r\, \langle j, D_2 \rangle \implies j = 1$
3) $\sim$ is an equivalence relation on $\mathcal{W}$ such that $\langle i, D_1 \rangle \sim \langle j, D_2 \rangle \implies D_1\, E\, D_2 \wedge i = j$ and such that $\langle 1, D_1 \rangle \sim \langle 1, D_2 \rangle \implies$ for any $\Box b \in \mathcal{Cl}(\alpha)$, $\Box b \in D_1 \iff \Box b \in D_2$
4) if w_1, w_2 and $w_3 = \langle i, D_3 \rangle$ are members of $\mathcal{W}$ such

that $w_1 r w_2$, and $w_2 \sim w_3$, then there exists an $w_4 = <j, D_4> \in \mathcal{W}$ such that $w_1 \sim w_4$, and $D_4 R D_3$.
5) for all $w \in \mathcal{W}$, there exists an $w' \in \mathcal{W}$ such that $w r w'$ and an $w'' \in \mathcal{W}$ of first coordinate 1 such that $w r^* w''$ where r^* denotes the reflexive and transitive closure of r
6) for all $w = <i, D> \in \mathcal{W}$ and for all $\Box b \in \mathcal{C}l^-(a)$, if $\Box b \notin D$ then there exists an $w' = <j, D'> \in \mathcal{W}$ such that $w r^* w'$ and $b \notin D'$
7) for all $w = <i, D> \in \mathcal{W}$ and for all $\nabla b \in \mathcal{C}l^-(a)$, if $\nabla b \notin D$ then there exists an $s' = <i, D'> \in \mathcal{W}$ such that $s \sim s'$ and $b \notin D'$

Now we may proceed as in the case of TCg.

Lemma 8: If a is an f-satisfiable formula, then there are an f-pseudo-model for a, $\langle \mathcal{W}, r, \sim \rangle$, and a $D \in \mathcal{W}$ such that $\nabla a \in D$.

Proof: One sees that $\langle \mathcal{A}, s, e^* \rangle$ is a suitable f-pseudo-model where
$\mathcal{A}$ contains the pairs $<0, T \cap \mathcal{C}l(a)>$ where T is an f-consistent and complete theory and the pairs $<1, T \cap \mathcal{C}l(a)>$ where T is such a terminal theory. The relation s is defined by:
a) $<0, D_1> s <i, D_2>$ iff there is an f-consistent and complete theory T such that $D_1 = T \cap \mathcal{C}l(a)$ and $D_2 = T^+ \cap \mathcal{C}l(a)$
b) $<1, D_1> s <1, D_2>$ iff there is a terminal, f-consistent and complete T as above.

The relation e is defined by:
a) $<0, D_1> e <0, D_2>$ iff there are f-consistent and complete theories $T_k, k = 1,2$ such that $T_1 \equiv T_2$, $D_k = T_k \cap \mathcal{C}l(a)$ for $k = 1,2$
b) $<1, D_1> e <1, D_2>$ iff there are terminal such theories.
Q.E.D.

Lemma 9: If there are an f-pseudo-model $\langle \mathcal{W}, r, \sim \rangle$ and a $w = <i, C> \in \mathcal{W}$ such that $\nabla a \in C$ then there exists an f-model for a.

Proof: The model $\mathcal{U} = \langle \mathcal{S}, u, l, p \rangle$, is defined the following way:
1) $\mathcal{S} = \mathcal{W}$
2) $u = C$
3) $l(D) = \{p \mid p \in D\}$
4) We decide that $p(w_1, w_2) \neq 0$ iff $w_1 \sim r\ w_2$ and that all transitions of positive probability from w_1 have equal probabilities.

Let m be a natural number. Let σ be a sequence of states of $\mathcal{U}$ (any sequence). Suppose $\sigma = <z_i, D_i>$. We shall say that σ is an m-standard sequence if there exists an $n \in \mathbb{N}$ such that: for any i such that $0 \le i < n$ we have $z_i = 0$, for any $i \ge n$, $z_i = 1$, for any i such that $0 \le i < n+m$, we have $D_i R D_{i+1}$ and for any $i \ge n+m$, we have $\sigma_i \sim r\ \sigma_{i+1}$. A sequence is ultimately m-standard if it has some m-standard tail.

We see, using conditions 2 and 5 of definition 9, that the set of m-standard sequences beginning at a state s has positive weight and that the set of ultimately m-standard sequences beginning at a state s has weight 1. Equivalent (E) sequences are defined as in the case of TCg.

We shall now need a stronger definition of generic sequences (it is needed in the $\Diamond$ case below). Let σ be a sequence of states of $\mathcal{U}$ ($\sigma \in \mathcal{S}^{\mathbb{N}}$), σ is said to be generic iff for any $s_0 \in \mathcal{S}$ that appears an infinite number of times in σ and for any finite sequence of $s_i \in \mathcal{S}$ $s_0, s_1, \cdots, s_m$ such that $s_i r s_{i+1}$, for every i such that $0 \le i < m$ the sequence above appears in σ (in this order) an infinite number of times. Clearly the weight of generic sequences is 1.

Our basic result concerning $\mathcal{U}$ is the following.

Lemma 10: Let $b \in \mathcal{C}l^-(a)$, σ an $\Omega(b)$-standard generic sequence of states (of $\mathcal{U}$) (and let the second coordinate of σ_i be D_i) and τ and τ' two equivalent sequences of states, then
a) $b \mid^{\sigma}_{\mathcal{U}} = \textit{true} \iff b \in \sigma_0$
b) $b \mid^{\tau}_{\mathcal{U}} = b \mid^{\tau'}_{\mathcal{U}}$.

Proof: The proof is very similar to that of lemma 7, and we signal only the differences.
$b = \bigcirc c$
a) $\bigcirc c \mid^{\sigma}_{\mathcal{U}} = \textit{true} \iff c \mid^{\sigma^1}_{\mathcal{U}} = \textit{true}$. The sequence σ^1 is generic since σ is. Since σ is $\Omega(b)$-standard, σ^1 is $(\Omega(b)-1)$-standard. But $\Omega(b)-1 = \Omega(c)$. Therefore, by the induction hypothesis: $c \mid^{\sigma^1}_{\mathcal{U}} = \textit{true} \iff c \in D_1$. Since $\Omega(b) \ge 1$, the first transition of σ is an R-transition and $D_0 R D_1$. We conclude that $c \in D_1 \iff \bigcirc c \in D_0$.
$b = \Box c$
a) Suppose $\Box c \in D_0$. Since σ is 0-standard, we have $\forall i \in \mathbb{N}, \Box c \in D_i$. Our goal is to use the induction hypothesis on c. We notice that $\forall i \in \mathbb{N}$, σ^i is generic. Let i be given. In general σ^i is not $\Omega(c)$-standard. Let $m = \Omega(c)$. Now we have to distinguish between two cases following whether the first coordinate of σ_i is 0 or 1. If it is 0, σ^i is clearly $\Omega(c)$-standard and we may use the induction hypothesis part a) to conclude: $c \mid^{\sigma^i}_{\mathcal{U}} = \textit{true}$.
On the other hand, suppose the first coordinate of σ_i is 1. By lemma 5 (mutatis mutandis), we may find $w_n \in \mathcal{W}$, for $n = i, \cdots, i+m-1$, (let F_n be the second coordinate of w_n) such that
1) $w_n \sim \sigma_n$, $i <= n \le i+m-1$
2) $F_n R F_{n+1}$, $i \le n < i+m-1$. By condition 3 of definition 9, the first coordinates of the w_n's are 1, and therefore the F_n's are terminal. By the same condition $\Box c \in F_i$. Let the sequence τ be defined by: $\tau^m = \sigma^{i+m}$, and $\tau_n = <1, F_{i+n}>$, $\forall n\ 0 \le n \le m-1$. By construction σ^i and τ are equivalent and therefore by the induction hypothesis b) $c \mid^{\sigma^i}_{\mathcal{U}} = c \mid^{\tau}_{\mathcal{U}}$. But τ is a generic m-standard sequence and by the induction hypothesis a) $c \mid^{\tau}_{\mathcal{U}} = \textit{true} \iff c \in F_i$. We conclude that $\Box c \mid^{\sigma}_{\mathcal{U}} = \textit{true}$.

$\Box c \notin D_0$ Following the line of reasoning used in lemma 7 one may see that $\exists i \in \mathbb{N}$ such that $c \notin D_i$. Now, as just above, we must distinguish two cases. If the first coordinate of σ_i is 0, σ^i must be m-standard and one may use the induction hypothesis part a). If the first coordinate of σ_i is 1, the proof is more delicate. In this case D_i is terminal and does not contain c, therefore it does contain $\Box\neg\Box c$ (this last formula is in $\mathcal{C}l^-(a)$). We may then show that for all indexes $j \ge i$, $\Box\neg\Box c \in D_j$, since σ^i contains only transitions of "type" R and $\sim r$. Therefore there is an $F \in \mathfrak{D}$ that contains $\neg\Box c$ and appears an infinite number of times in σ (as a second component). Since σ is generic, we conclude, by condition 6 of definition 9 that there is a $G \in \mathfrak{D}$ that does not contain c and that appears an infinite number of times in σ (as a second component). Since σ is generic, G must appear in σ an infinite number of times followed by at least $\Omega(c)$ r-transitions. Let j be such a point in σ, with $j \ge i$. At j, we may apply the induction hypothesis part a) and see that $c \,|_{\mathcal{U}}^{\sigma^j} = \boldsymbol{false}$. We conclude that $\Box c \,|_{\mathcal{U}}^{\sigma^i} = \boldsymbol{false}$.

$b = \nabla c$

a) Suppose $\nabla c \in D_0$. We want to show that, for almost all sequences $\tau \in P_{\sigma_0}$, we have $c \,|_{\mathcal{U}}^{\tau} = \boldsymbol{true}$. Since almost all sequences of P_{σ_0} are generic and $\Omega(c)$-ultimately standard it is enough to show that if τ is generic and $\Omega(c)$-ultimately standard, then $c \,|_{\mathcal{U}}^{\tau} = \boldsymbol{true}$. Let $m = \Omega(c)$. Since τ is m-ultimately standard it has an m-standard tail. Using lemma 5 (mutatis mutandis) as we did in the corresponding part of lemma 7 we may build a sequence τ' that is equivalent to τ, generic and m-standard. By the induction hypothesis, part b), $c \,|_{\mathcal{U}}^{\tau} = c \,|_{\mathcal{U}}^{\tau'}$. We conclude by using the induction hypothesis, part a).

Suppose now that $\neg\nabla c \in D_0$. We must find a set Q of sequences that begin at σ_0 and do not satisfy c, such that Q has a positive measure. By the condition 7 of definition 9, there exists a $w \sim \sigma_0$ (say $w = \langle z, F\rangle$ such that $c \notin F$. Let F' be any member of $\mathcal{W}$ such that $F r F'$. Notice that, by condition 5 definition 9 there is such an F'. By the construction of our model, we have $p(\langle z_0, D_0\rangle, \langle z_0, F'\rangle) > 0$. Let us define Q as the set of all sequences τ such that: $\tau_0 = \sigma_0$, $\tau_1 = \langle z_0, F'\rangle$ and the sequence τ^1 is generic and $\Omega(c)$-standard. We have $\tilde{p}_{\sigma_0}(Q) > 0$. It is left to us to show that no sequence of Q satisfies c. Let τ be any sequence of Q. Let τ' be the sequence $w, \tau_1, \tau_2, \cdots$. It is a generic $\Omega(c)$-standard sequence. By the induction hypothesis, part a) we have $c \,|_{\mathcal{U}}^{\tau'} = \boldsymbol{false}$. But since $\sigma_0 \sim w$, $\tau E \tau'$. We conclude, by the induction hypothesis, part b), that $c \,|_{\mathcal{U}}^{\tau} = \boldsymbol{false}$. **Q.E.D.**

The proof of lemma 9 is completed as in the previous case.

Theorem 4: Satisfiability in TCf is decidable in $NTIME(2^{2^{cn}})$ for some constant $c \ge 0$ where n is the size ($\#$) of the formula .

Proof: The proof is like in the previous case. **Q.E.D.**

7. A decision procedure for TCb

We need to enlarge the set of formulas to be considered. The definition of $\mathcal{Cl}^-(a)$ is as in definition 7. For the definition of $\mathcal{Cl}(a)$ we replace definition 2 by the following.

Definition 10: A formula b is in $\mathcal{Cl}(a)$ iff it is in $\mathcal{Cl}^-(a)$, of the form $\nabla\left[b_1 \vee \cdots b_k\right]$ where $b_i \in \mathcal{Cl}(a)$, $i = 1, \cdots, k$ or of the form $\Delta\left[\bigwedge_{k=0}^{m} \bigcirc^{(l_k)} \nabla a_k\right]$, for $m \in \mathbb{N}$, ∇a_k in $\mathcal{Cl}^-(a)$ and $l_k \le \#(a)$ for any $k \le m$.

One may see that the size of $\mathcal{Cl}(a)$ is 2^{cn^2} where $n = \#(a)$.

Our new notion of a pseudo-model is the following.

Definition 11: A b-pseudo-model for a is a triplet $\langle \mathcal{W}, r, \sim \rangle$ satisfying:

1) $\mathcal{W} \subseteq \{0,1\} \times \mathcal{D}$ such that $\langle 1, D\rangle \in \mathcal{W} \implies D$ is terminal

2) r is a binary relation on $\mathcal{W}$ such that $\langle i, D_1\rangle r \langle j, D_2\rangle \implies D_1 R D_2$

3) $\sim$ is an equivalence relation on $\mathcal{W}$ such that $\langle i, D_1\rangle \sim \langle j, D_2\rangle \implies D_1 E D_2$ and such that $\langle 1, D_1\rangle \sim \langle 1, D_2\rangle \implies$ for any $\Box b \in \mathcal{Cl}(a)$, $\Box b \in D_1 \iff \Box b \in D_2$

4) if w_1, w_2 and $w_3 = \langle i, D_3\rangle$ are members of $\mathcal{W}$ such that $w_1 r w_2$, and $w_2 \sim w_3$, then there exists an $w_4 = \langle j, D_4\rangle \in \mathcal{W}$ such that $w_1 \sim w_4$, and $D_4 R D_3$.

5) for all $w \in \mathcal{W}$, there exists an $w' \in \mathcal{W}$ such that $w r w'$, for all w's of first coordinate 1 there is such a w' of first coordinate 1, and for all w's of first coordinate 0 there is a $w'' \in \mathcal{W}$ of first coordinate 1 such that $w r^* w''$ where r^* denotes the reflexive and transitive closure of r

6) for all $w = \langle i, D\rangle \in \mathcal{W}$ and for all $\Box b \in \mathcal{Cl}^-(a)$, if $\Box b \notin D$ then there exists an $w' = \langle j, D'\rangle \in \mathcal{W}$ such that $w r^* w'$ and $b \notin D'$

7) for all $w = \langle i, D\rangle \in \mathcal{W}$ and for all $\nabla b \in \mathcal{Cl}^-(a)$, if $\nabla b \notin D$ then there exists an $s' = \langle 0, D'\rangle \in \mathcal{W}$ such that $s \sim s'$ and $b \notin D'$

8) for any $w_0, w_1, .. w_k \in \mathcal{W}$ where $k \le \#(a)$ such that w_j has first coordinate 1, for $j = 0, \cdots, k$ and $w_j \sim r\, w_{j+1}$, for $j = 0, \cdots, k-1$ there exist $v_j \in \mathcal{W}$ of first coordinate 1, $j = 0, \cdots k$ satisfying

a) $v_0 \sim w_0$

b) $v_j E w_j$, for $j = 1, \cdots, k$, where E is interpreted on the second coordinate only

c) $v_j R v_{j+1}$, for $j = 0, \cdots, k-1$, where R is interpreted on the second coordinate only.

9) for any $v, w \in \mathcal{W}$ such that $v \sim w$ and for any $k \le \#(a)$ there are v_j, w_j, for $j = 0, \cdots, k$ such that

a) $v_0 = v$ and $w_0 = w$

b) $v_j E w_j$ for $j = 0, \cdots, k$ where E is interpreted on the second coordinate only

c) $v_j \sim r v_{j+1}$, for $j = 0, \cdots, k-1$

d) $w_j R w_{j+1}$, for $j = 0, \cdots, k-1$.

Lemma 11: If a is a b-satisfiable formula, then there are a b-pseudo-model $\langle \mathcal{W}, r, \sim \rangle$ for a and a $D \in \mathcal{W}$ such that $\nabla a \in D$.

Proof: One sees that $\langle \mathcal{A}, s, e^* \rangle$ is a suitable b-pseudo-model where

$\mathcal{A}$ contains the pairs $\langle 0, T \cap \mathcal{Cl}(a)\rangle$ where T is an b-consistent and complete theory and the pairs $\langle 1, T \cap \mathcal{Cl}(a)\rangle$ where T is such a terminal theory. The relation s is defined by:

a) if $i=0$ or $j=0$ $\langle i, D_1\rangle s \langle j, D_2\rangle$ iff there is an b-consistent and complete theory T such that $D_1 = T \cap \mathcal{Cl}(a)$ and $D_2 = T^+ \cap \mathcal{Cl}(a)$

b) $\langle 1, D_1\rangle s \langle 1, D_2\rangle$ iff there is a terminal, b-consistent and complete T as above.

The relation e is defined by:

a) if $i=0$ or $j=0$ $\langle i, D_1\rangle e \langle j, D_2\rangle$ iff there are b-consistent and complete theories T_k, $k = 1,2$ such that $T_1 = T_2$, $D_k = T_k \cap \mathcal{Cl}(a)$ for $k = 1,2$

b) $\langle 1, D_1\rangle e \langle 1, D_2\rangle$ iff there are terminal theories T_k for $k=1,2$ such that $T_1 = T_2$, $T_1 \le T_2$, $T_2 \le T_1$ and $D_k = T_k \cap \mathcal{Cl}(a)$ for $k=1,2$. **Q.E.D.**

Lemma 12: If there are a b-pseudo-model $\langle \mathcal{W}, r, \sim, \rangle$ for a , and a $w = \langle i, D\rangle$ in $\mathcal{W}$ such that $\nabla a \in D$ then a is b-satisfiable.

Proof: The model $\mathcal{U} = \langle \mathcal{S}, u, l, p \rangle$, that satisfies α, is defined the following way:

1) $\mathcal{S}$ is a subset of $\mathbb{N} \times \mathfrak{D}$ defined by:
   $<0, D> \in \mathcal{S}$ iff $<0, D> \in \mathcal{W}$
   for $i > 0$, $<i, D> \in \mathcal{S}$ iff $<1, D> \in \mathcal{W}$
2) $u = w'$
3) $l(<i, D>) = \{p \mid p \in D\}$
4) Choose some number α: $\frac{1}{2} < \alpha < 1$. We distinguish here between the states of first coordinate 0 and the other ones. For states whose first coordinate is zero, we give a positive probability to a move from $<0, D>$ to $<0, D'>$ iff $<0, D> \sim r <0, D'>$. We give a positive probability to a move from $<0, D>$ to $<1, D'>$ iff $<0, D> \sim r <1, D'>$. All other transitions from $<0, D>$ have probability zero. We give equal probabilities to all moves of positive probability. For states whose first coordinate is positive, we allow to increase or decrease by one the first coordinate. If the first coordinate is 1, we give a positive probability to transitions from $<1, D>$ to $<0, D'>$ iff $<1, D> \sim r <0, D'>$ and to transitions from $<1, D>$ to $<2, D'>$ iff there is a $w \in \mathcal{W}$ in the pseudo-model, such that w has first coordinate 1, $<1, D> \sim w$ and $w\, r <1, D'>$. If $i > 1$, we give a positive probability to a transition from $<i, D>$ to $<j, D'>$ iff (either $j = i-1$ or $j = i+1$) and there is a $w \in \mathcal{W}$ in the pseudo-model, such that w has first coordinate 1, $<1, D> \sim w$ and $w\, r <1, D'>$. Moreover, at each state s of positive first coordinate we give a combined weight of α ($\alpha > \frac{1}{2}$) to those moves that increase the first coordinate.

Let σ be a sequence and let $\sigma_i = <k_i, D_i>$. Let m be a natural number. We define σ to be m-standard iff there exists a j, such that $k_i > 0$, for $i \geq j$, such that for any i, $i \leq j + m - 1$ we have $D_i\, R\, D_{i+1}$ and such that for any i, $i \geq j + m$ we have $\sigma_i \sim r\, \sigma_{i+1}$.

A sequence is said to be m-ultimately standard if it has a tail that is m-standard.

We see, using condition 5 of definition 11, that the set of m-standard sequences beginning at a state s has positive weight and that the set of ultimately m-standard sequences beginning at a state s has weight 1. Equivalent (E) sequences are defined as in the case of TCg. Generic sequences are defined as in the case of TCf. Clearly the weight of generic sequences is 1.

Our basic result concerning $\mathcal{U}$ is the following.

Lemma 13: Let $b \in \mathcal{Cl}(\alpha)$, σ a generic $\Omega(b)$-standard sequence of states (of $\mathcal{U}$) and τ and τ' any two equivalent (E) sequences of states, then

a) $b \mid_{\mathcal{U}}^{\sigma} = \textit{true} \iff b \in D_0$, where $\sigma_i = <k_i, D_i>$

b) $b \mid_{\mathcal{U}}^{\tau} = b \mid_{\mathcal{U}}^{\tau'}$.

Proof: The proof is very similar to the proof of lemma 10, and therefore we shall only highlight the changes to be made.

$b = \Box c$

a) Suppose $\Box c \in D_0$. Since σ is 0-standard, we have $\forall i \in \mathbb{N}, \Box c \in D_i$. Our goal is to use the induction hypothesis on c. We notice that $\forall i \in \mathbb{N}$, σ^i is generic. Let i be given. In general σ^i is not $\Omega(c)$-standard. Let $m = \Omega(c)$. Since σ is m-standard, there is an index j such that $k_l > 0$, for $l \geq j$, such that for any l, $l \leq j + m - 1$ we have $D_l\, R\, D_{l+1}$ and such that for any l, $l \geq j + m$ we have $\sigma_l \sim r\, \sigma_{l+1}$. Now we have to distinguish between two cases following whether $i < j$ or not. If $i < j$, σ^i is clearly $\Omega(c)$-standard and we may use the induction hypothesis part a) to conclude: $c \mid_{\mathcal{U}}^{\sigma^i} = \textit{true}$.

On the other hand, suppose $i >= j$. By condition 8 of definition 11 we may find $w_n \in \mathcal{W}$, for $n = i, \cdots, i+m-1$, (let $w_n = <z_n, F_n>$) such that

a) $z_n = k_n$, $n = i, \cdots, i+m-1$
b) $w_i \sim \sigma_i$
c) $F_n\, E\, D_n$, $n = i+1, \cdots, i+m-1$
d) $F_n\, R\, F_{n+1}$, $n = i, \cdots, i+m-1$.

By condition 3 of definition 11, $\Box c \in F_i$. Let the sequence τ be defined by: $\tau^m = \sigma^{i+m}$, and $\tau_n = <z_{i+n}, F_{i+n}>$, $n = 0, \cdots, m-1$. By construction σ^i and τ are equivalent and therefore by the induction hypothesis b) $c \mid_{\mathcal{U}}^{\sigma^i} = c \mid_{\mathcal{U}}^{\tau}$. But τ is a generic m-standard sequence and by the induction hypothesis a) $c \mid_{\mathcal{U}}^{\tau} = \textit{true} \iff c \in F_i$. We conclude that $\Box c \mid_{\mathcal{U}}^{\sigma} = \textit{true}$.

$\Box c \notin D_0$ Following the line of reasoning used in lemma 7 one may see that $\exists i \in \mathbb{N}$ such that $c \notin D_i$. Now, as just above, we must distinguish two cases. If the first coordinate of σ_i is 0, σ^i must be m-standard and one may use the induction hypothesis part a). If the first coordinate of σ_i is positive, the proof is more delicate. In this case D_i is terminal and does not contain c, therefore it does contain $\Box \neg \Box c$ (this last formula is in $\mathcal{Cl}^-(\alpha)$. We may then show that for all indexes $j \geq i$, $\Box \neg \Box c \in D_j$, since σ^i contains only transitions of "type" R and $\sim r$. Therefore there is an $F \in \mathfrak{D}$ that contains $\neg \Box c$ and appears an infinite number of times in σ (as a second component). Since σ is generic, we conclude, by condition 6 of definition 11 that there is a $G \in \mathfrak{D}$ that does not contain c and that appears an infinite number of times in σ (as a second component). Since σ is generic, G must appear in σ an infinite number of times followed by at least $\Omega(c)$ r-transitions. Let j be such a point in σ, with $j \geq i$. At j, we may apply the induction hypothesis part a) and see that $c \mid_{\mathcal{U}}^{\sigma^j} = \textit{false}$. We conclude that $\Box c \mid_{\mathcal{U}}^{\sigma^i} = \textit{false}$.

$b = \nabla c$ The first half proceeds as in lemma 10.

Suppose now that $\neg \nabla c \in D_0$. We must find a set Q of sequences that begin at σ_0 and do not satisfy c, such that Q has a positive measure. Remember that $\sigma_0 = <k_0, D_0>$. By condition 7 of definition 11 there is an $s = <0, F>$ such that $c \notin F$ and $\sigma_0 \sim t$. By condition 9 of definition 11 we may find s_i, t_i for $i = 0, \cdots, m$ (where $m = \Omega(c)$) such that

a) $\sigma_0 = s_0$ and $t = t_0$
b) $s_i\, E\, t_i$
c) $s_i \sim r\, s_{i+1}$
d) $t_i\, R\, t_{i+1}$.

Let Q be the set of all sequences τ such that: $\tau_i = s_i$, for $i = 0, \cdots, m$ and τ^m is generic and 0-standard. By construction the set Q has a positive measure. It is left to us to show that no sequence of Q satisfies c. Let τ be a sequence of the set Q. Define the sequence τ' by: $\tau'^{m+1} = \tau^{m+1}$ and $\tau'_i = t_i$, for $i = 0, \cdots, m$. By construction $\tau\, E\, \tau'$ and therefore, by the induction hypothesis part b) $c \mid_{\mathcal{U}}^{\tau} = c \mid_{\mathcal{U}}^{\tau'}$. But the sequence τ' is generic since it has a generic tail. It is also m-standard since it consists of a prefix of m R-transitions and a 0-standard tail. By the induction hypothesis part a) we conclude $c \mid_{\mathcal{U}}^{\tau'} = \textit{false}$.

The proof of lemma 12 is completed as in the previous case. **Q.E.D.**

Theorem 5: Satisfiability in TCb is decidable in $NTIME(2^{2^{cn^2}})$ for some constant $c \geq 0$ where n is the size (#) of the formula .

Proof: The proof is like in the previous case. **Q.E.D.**

8. Open problems There is a gap between the deterministic exponential lower bound and the non-deterministic double- exponential upper bound. It is our feeling that this gap may be reduced. A natural idea is to try and obtain a deterministic double-exponential time decision procedure by an iterative method in the spirit of [SH], but a too simple-minded effort fails. The idea seems applicable, though, for a sublanguage containing only "state formulas", large enough to include the language of [HS]. One should also look for methods to check whether a given formula is satisfied in a given model (described in some finite way).

9. Acknowledgements

J. Halpern suggested the exponential-time lower bound.

References

[A] Abrahamson, K. Decidability and expressiveness of logics of processes, PhD. Thesis, Univ. of Washington (1980).

[BMP] Ben-Ari, M., Manna, Z. and Pnueli, A. The temporal logic of branching time, Conf. Record 8th Annual ACM Symposium on Principles of Programming Languages, Williamsburg, Va. (Jan. 1981), pp. 164-176.

[CE] Clarke, E.M. and Emerson, E.A. Design and synthesis of synchronization skeletons using branching time temporal logic, Proc. Workshop on Logics of Programs, Kozen ed., Springer-Verlag (1982) (to appear).

[CLP] Cohen, S., Lehmann D. and Pnueli, A. Symmetric and economical solutions to the mutual exclusion problem in a distributed system Proc. of 10th International Conference on Automata, Languages and Programming, Barcelona (1983).

[EH1] Emerson, E. A. and Halpern, J. Y. Decision procedures and expressiveness in the temporal logic of branching time, Conf. Record 14th Annual ACM Symposium on Theory of Computing, San Francisco, CA (May 1982), pp. 169-179.

[EH2] Emerson, E.A. and Halpern, J.Y. "Sometimes" and "not never" revisited: on branching versus linear time (preliminary report), Proc. of 10th Annual ACM Symposium on Principles of Programming Languages, Austin, Texas (Jan. 1983), pp. 127-139.

[FL] Fischer, M.J. and Ladner, R.E. Propositional dynamic logic of regular programs, Journal of Computer and System Sciences, Vol. 18 (1979), pp. 194-211.

[GPSS] Gabbay, D., Pnueli, A., Shelah, S. and Stavi, J. On the temporal analysis of fairness, Conf. Record of 7th Annual ACM Symposium on Principles of Programming Languages, Las Vegas, Nevada (Jan. 1980), pp. 163-173.

[HS] Hart, S. and Sharir, M. Probabilistic Propositional Temporal Logics, manuscript, June 83.

[La] Lamport, L. "Sometimes" is sometimes "not never", Conf. Record of 7th Annual ACM Symposium on Principles of Programming Languages, Las Vegas, Nevada (Jan. 1980), pp. 174-183.

[LR] Lehmann, D. and Rabin, M. O. On the advantages of free choice: a symmetric and fully distributed solution to the dining philosophers problem (extended abstract), Conf. Record of 8th Annual ACM Symposium on Principles of Programming Languages, Williamsburg, Va. (Jan. 1981), pp. 133-138.

[LS] Lehmann, D. and Shelah, S. Reasoning with Time and Chance, Information and Control (in press) (Extended Abstract in Proc. 10th International Conference on Automata, Languages and Programming, Barcelona, 1983, pp. 445-457.

[Pn] Pnueli, A. The temporal semantics of concurrent programs, Theoretical Computer Science Vol. 13 (1981), pp. 45-60.

[Ra1] Rabin, M. O. Decidability of second order theories and automata on infinite trees, Trans. AMS Vol. 141 (1969), pp. 1-35.

[Ra2] Rabin, M.O. N-process mutual exclusion with bounded waiting by 4.log N-valued shared variable, Journal of Computer and System Sciences Vol.25 (1982) pp. 66-75.

[Ra3] Rabin, M.O. The choice coordination problem, Acta Informatica Vol. 17 (1982), pp. 121-134.

[SH] Sherman, R. and Harel D. A combined proof of one-exponential decidability and completeness for PDL, Proc. First International Workshop on Foundations of Theoretic Computer Science G.T.I., Univ. of Paderborn, West Germany (March 1983) pp.221-233.

[SC] Sistla, A. P. and Clarke, E. M. The complexity of propositional linear temporal logics, Proc. 14th Annual ACM Symposium on Theory of Computing, San Francisco, CA (May 1982), pp. 159-168.