CHARACTERIZING FINITE KRIPKE STRUCTURES IN PROPOSITIONAL TEMPORAL LOGIC *

M.C. BROWNE, E.M. CLARKE and O. GRÜMBERG

Department of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213, U.S.A.

Abstract. We show that if two *finite* Kripke structures can be distinguished by some CTL* formula that contains both branching-time and linear-time operators, then the structures can be distinguished by a CTL formula that contains only branching-time operators. Our proof involves showing that, for any finite Kripke structure M, it is possible to construct a CTL formula F_M that uniquely characterizes M. Since one Kripke structure may be a trivial unrolling of another, we use a notion of equivalence between Kripke structures that is similar to the notion of *bisimulation* studied by Milner [15].

Our first construction of F_M requires the use of the *nexttime* operator. We also consider the case in which the nexttime operator is disallowed in CTL formulas. The proof, in this case, requires another notion of equivalence—*equivalence with respect to stuttering* and is much more difficult since it is possible for two inequivalent states to have exactly the same finite behaviors (modulo stuttering), but different infinite behaviors. We also give a polynomial algorithm for determining if two structures are stuttering equivalent and discuss the relevance of our results for temporal logic model checking and synthesis procedures.

1. Introduction

The question of whether branching-time temporal logic or linear-time temporal logic is best for reasoning about concurrent programs is one of the most controversial issues in logics of programs. Concurrent programs are usually modelled by labelled state-transition graphs in which some state is designated as the initial state. For historical reasons such graphs are called Kripke structures [11]. In linear temporal logic, operators are provided for describing events along a single time path (i.e., along a single path in a Kripke structure). In a branching-time logic the temporal operators quantify over the futures that are possible from a given state (i.e., over the possible paths that lead from a state). It is well known that the two types of temporal logic have different expressive powers [6, 12]. Linear temporal logic, for example, can express certain *fairness* properties that cannot be expressed in branching-time temporal logic. On the other hand, certain practical decision problems like *model checking* [5, 20] are easier for branching-time temporal logic than for linear temporal logic.

In this paper we provide further insight on which type of logic is best. We show that if two *finite* Kripke structures can be distinguished by some formula that contains both branching-time and linear-time operators, then the structures can be distinguished by a formula that contains only branching-time operators. Specifically, we show that if two finite Kripke structures can be distinguished by some formula of the logic CTL* (i.e., if there is some CTL* formula that is true in one but not in the other), then they can be distinguished by some formula of the logic CTL. The logic CTL* [5, 6] is a very powerful temporal logic that combines both branching-time and linear-time operators; a path quantifier, either **A** ("for all paths") or **E** ("for some paths") can prefix an assertion composed of arbitrary combinations of the usual linear-time operators **G** ("always"), **F** ("sometimes"), **X** ("nexttime"), and **U** ("until"). CTL [2, 4] is a restricted subset of CTL* that permits only branching-time operators—each path quantifier must be immediately followed by exactly one of the operators **G**, **F**, **X**, or **U**.

Our goal is to show that, for any finite Kripke structure M, it is possible to construct a CTL formula F_M that uniquely characterizes M. Since one Kripke structure may be a trivial unrolling of another, we use a notion of equivalence between Kripke structures that is similar to the notion of *bisimulation* studied by Milner [15]. We say that states s and s' are *equivalent* if they have the same labelling of atomic propositions and for each transition from one of the two states to some state t there is a corresponding transition from the other state to a state t' that is equivalent to t. Two Kripke structures are equivalent if their initial states are equivalent. It is not difficult to prove that if two Kripke structures are equivalent, then their initial states must satisfy the same CTL* formulas.

An obvious first attempt to construct F_M is simply to write a CTL formula that specifies the transition relation of M. For each state s in M we include in F_M a conjunct of the form

$$\mathbf{AG}\left(\mathscr{L}(s) \Rightarrow \bigwedge_i \mathbf{EX}\mathscr{L}(s_i) \wedge \mathbf{AX}\left(\bigvee_i \mathscr{L}(s_i) \right) \right)$$

where $s_1, \ldots, s_n$ are the successors of s and $\mathscr{L}(t)$ is the labelling of atomic propositions associated with state t. It is easy to see, however, that this simple approach cannot work in general: several states in M may have exactly the same labelling of atomic propositions.

Instead, we first show that it is possible to write a CTL formula that will distinguish between two states in the same structure that are not equivalent according to the above definition. Two inequivalent states may have exactly the same labelling of atomic propositions, they may even have corresponding successors, but the computation trees rooted at those states must differ at some finite depth. The difference in the computation trees can be exploited to give a CTL formula that distinguishes between the states. Since equivalent states satisfy the same CTL* formulas, it follows that if two states can be distinguished by a CTL* formula, they can be distinguished by a CTL formula. Once we can distinguish between inequivalent states in the same

structure, we can write a single CTL formula that encodes the entire Kripke structure; this formula is the F_M that we seek.

The above construction requires the use of the nexttime operator in specifying F_M. In reasoning about concurrent systems, however, the nexttime operator may be dangerous since it refers to the *global* next state instead of the *local* next state within a process [13]. What happens if we disallow the nexttime operator in CTL formulas? The proof, in this case, requires another notion of equivalence— *equivalence with respect to stuttering.* We say that two state sequences *correspond* if each can be partitioned into finite blocks of identically labelled states such that each state in the *i*th block in one sequence is equivalent to each state in the *i*th block of the other sequence. Thus, duplicating some state in a sequence any finite number of times will always result in a corresponding sequence. We say that two states are *equivalent* if for each state sequence starting at one there is a corresponding state sequence that starts at the other. Under this second notion of equivalence the proof of the characterization theorem becomes much more complicated since it is possible for two inequivalent states to have exactly the same finite behaviors (modulo stuttering), but different infinite behaviors.

Equivalence under stuttering turns out to be quite useful for reasoning about hierarchically constructed concurrent systems. In determining the correctness of such a system by using a technique like temporal logic model checking [1, 3, 4, 5, 14, 16, 19, 20, 21], it is often desirable to replace a low-level module by an equivalent structure with fewer states. Our results show how this can be done while preserving all of those properties that are invariant under stuttering. We give polynomial algorithms both for determining if two structures are equivalent with respect to stuttering and for minimizing the number of states in a given structure under this notion of equivalence.

Finally, our results have some interesting implications for the problem of synthesizing finite-state concurrent systems from temporal logic specifications [4, 17]. In order to guarantee that any Kripke structure can be synthesized from a specification in linear temporal logic, Wolper [22] was forced to introduce more complicated operators based on regular expressions. Our results show that (at least in theory) no such extension is necessary for branching-time temporal logic. Any Kripke structure can be specified directly by a formula of branching-time logic.

The expressive power of various temporal logics has been discussed in several papers; see [6, 12], for example. Hennessy and Milner [9], Hennessy and Stirling [10], Graf and Sifakis [8], and Pnueli [18] have all discussed the relationship between temporal logic and various notions of equivalence between models of concurrent programs. However, we believe that we are the first to show that it is possible to characterize Kripke models within branching-time logic and to investigate the consequences of this result.

Our paper is organized as follows: In Section 2 we describe the logics CTL and CTL*. In Section 3, we state formally what it means for two states in a Kripke structure to be equivalent and prove that equivalent states satisfy exactly the same

CTL* formulas. Section 3 also contains the first of the two main results of the paper: we show how to characterize Kripke structures using CTL formulas with the nexttime operator. Section 4 introduces the second notion of equivalence (equivalence with respect to stuttering) and shows that if the nexttime operator is disallowed, then equivalent states again satisfy exactly the same CTL* formulas. We also extend the characterization theorem of Section 3 to Kripke structures with the new notion of equivalence. In Section 5 we give a polynomial algorithm for determining if two states are equivalent up to stuttering. The paper concludes in Section 6 with a discussion of some remaining open problems.

2. The logics CTL and CTL*

There are two types of formulas in CTL*: *state formulas* (which are true in a specific state) and *path formulas* (which are true along a specific path). Let **AP** be the set of atomic proposition names. A state formula is either:

- A if $A \in$ **AP**;
- if f and g are state formulas, then $\neg f$ and $f \vee g$ are state formulas;
- if f is a path formula, then $\mathbf{E}(f)$ is a state formula.

A path formula is either

- a state formula;
- if f and g are path formulas, then $\neg f$, $f \vee g$, $\mathbf{X}f$, and $f\mathbf{U}g$ are path formulas.

  CTL* is the set of state formulas generated by the above rules.

  CTL is a subset of CTL* in which we restrict the path formulas to be

- if f and g are state formulas, then $\mathbf{X}f$ and $f\mathbf{U}g$ are path formulas;
- if f is a path formula, then so is $\neg f$.

  We define the semantics of both logics with respect to a structure $M = \langle S, R, \mathscr{L} \rangle$, where

- S is a set of states;
- $R \subseteq S \times S$ is the transition relation, which must be total; we write $s_1 \to s_2$ to indicate that $(s_1, s_2) \in R$;
- $\mathscr{L} : S \to \mathscr{P}(\mathbf{AP})$ is the proposition labelling.

  Unless otherwise stated, all of our results apply only to finite Kripke structures.

  We only consider transition relations where every state is reachable from the initial state. We define a *path in M* to be a sequence of states, $\pi = s_0, s_1, \dots$ such that, for every $i \geqslant 0$, $s_i \to s_{i+1}$. π^i will denote the *suffix* of π starting at s_i.

  We use the standard notation to indicate that a state formula f holds in a structure: $M, s \vDash f$ means that f holds at state s in structure M. Similarly, if f is a path formula, $M, \pi \vDash f$ means that f holds along path π in structure M. The relation $\vDash$ is inductively defined as follows (assuming that f_1 and f_2 are state formulas and g_1 and g_2 are path formulas):

(1) $s \vDash A \Leftrightarrow A \in \mathscr{L}(s)$.

(2) $s \vDash \neg f_1 \Leftrightarrow s \nvDash f_1$.

(3) $s \vDash f_1 \vee f_2 \Leftrightarrow s \vDash f_1$ or $s \vDash f_2$.

(4) $s \vDash \mathbf{E}(g_1) \Leftrightarrow$ there exists a path π starting with s such that $\pi \vDash g_1$.

(5) $\pi \vDash f_1 \Leftrightarrow s$ is the first state of π and $s \vDash f_1$.

(6) $\pi \vDash \neg g_1 \Leftrightarrow \pi \nvDash g_1$.

(7) $\pi \vDash g_1 \vee g_2 \Leftrightarrow \pi \vDash g_1$ or $\pi \vDash g_2$.

(8) $\pi \vDash \mathbf{X} g_1 \Leftrightarrow \pi^1 \vDash g_1$.

(9) $\pi \vDash g_1 \mathbf{U} g_2 \Leftrightarrow$ there exists a $k \geqslant 0$ such that $\pi^k \vDash g_2$ and, for all $0 \leqslant j < k$, $\pi^j \vDash g_1$.

We will also use the following abbreviations in writing CTL* (and CTL) formulas:

- $f \wedge g \equiv \neg(\neg f \vee \neg g)$,
- $\mathbf{F} f \equiv true\ \mathbf{U} f$,
- $\mathbf{A}(f) \equiv \neg \mathbf{E}(\neg f)$,
- $\mathbf{G} f \equiv \neg \mathbf{F} \neg f$.

3. Equivalence of Kripke structures

Given two structures M and M' with the same set of atomic propositions $\mathbf{AP}$, we define a sequence of equivalence relations $E_0, E_1, \ldots$ on $S \times S'$ as follows:

- $s E_0 s'$ if and only if $\mathscr{L}(s) = \mathscr{L}(s')$;
- $s E_{n+1} s'$ if and only if
  - $\mathscr{L}(s) = \mathscr{L}(s')$,
  - $\forall s_1[s \to s_1 \Rightarrow \exists s_1'[s' \to s_1' \wedge s_1 E_n s_1']]$, and
  - $\forall s_1'[s' \to s_1' \Rightarrow \exists s_1[s \to s_1 \wedge s_1 E_n s_1']]$.

Now, we define our notion of *equivalence between states*: $s E s'$ if and only if $s E_i s'$ for all $i \geqslant 0$. Furthermore, we say that M with initial state s_0 is equivalent to M' with initial state s_0' iff $s_0 E s_0'$. Two paths $s, s_1, \ldots$ and $s', s_1', \ldots$ *correspond* if $s E s'$ and $\forall i[s_i E s_i']$. Also note that if two paths correspond, so do all the respective pairs of tails.

Lemma 3.1. *Let $s E s'$; then for every path starting from s there exists a corresponding path starting from s', and for every path starting from s' there exists a corresponding path starting from s.*

Proof. Note first that $E_{n+1} \subseteq E_n$ for all $n \geqslant 0$. Since E_0 is finite, there must be a k such that $E_{k+1} = E_k = E$. Thus, we can substitute E for E_k in the definition of E_{k+1} giving $s E s'$ if and only if

- $\mathscr{L}(s) = \mathscr{L}(s')$,
- $\forall s_1[s \to s_1 \Rightarrow \exists s_1'[s' \to s_1' \wedge s_1 E s_1']]$, and
- $\forall s_1'[s' \to s_1' \Rightarrow \exists s_1[s \to s_1 \wedge s_1 E s_1']]$.

The remainder of the proof is a straightforward induction on the position in the path. $\square$

Theorem 3.2. *If $s \, E \, s'$, then $\forall f \in \mathrm{CTL}^*[s \models f \Leftrightarrow s' \models f]$.*

This theorem is a consequence of the following lemma.

Lemma 3.3. *Let h be either a state formula or a path formula. Let $\pi = s, s_1, \ldots$ be a path in M and $\pi' = s', s'_1, \ldots$ be a corresponding path in M'. Then*
  (i) *$s \models h \Leftrightarrow s' \models h$ if h is a state formula, and*
  (ii) *$\pi \models h \Leftrightarrow \pi' \models h$ if h is a path formula.*

Proof. We prove the lemma by induction on the structure of h.
  Base: $h = A$. By the definition of E, $s \models A \Leftrightarrow s' \models A$.
  Induction: There are several cases.
  (1) $h = \neg h_1$, a state formula.

$$s \models h \;\Leftrightarrow\; s \not\models h_1$$

$$\Leftrightarrow\; s' \not\models h_1 \quad \text{(induction hypothesis)}$$

$$\Leftrightarrow\; s' \models h$$

The same reasoning holds if h is a path formula.
  (2) $h = h_1 \vee h_2$, a state formula.

$$s \models h \;\Leftrightarrow\; s \models h_1 \text{ or } s \models h_2$$

$$\Leftrightarrow\; s' \models h_1 \text{ or } s' \models h_2 \quad \text{(induction hypothesis)}$$

$$\Leftrightarrow\; s' \models h$$

We can also use this argument if h is a path formula.
  (3) $h = \mathbf{E}(h_1)$, a state formula. Suppose that $s \models h$. Then there is a path, π_1 starting with s such that $\pi_1 \models h_1$. By Lemma 3.1, there is a corresponding path π'_1 in M' starting with s'. So, by the induction hypothesis, $\pi_1 \models h_1 \Leftrightarrow \pi'_1 \models h_1$. Therefore, $s \models \mathbf{E}(h_1) \Rightarrow s' \models \mathbf{E}(h_1)$. We can use the same argument in the other direction.
  (4) $h = h_1$, where h is a path formula and h_1 is a state formula. Although the lengths of h and h_1 are the same, we can imagine that $h = \mathbf{path}(h_1)$, where **path** is an operator which converts a state formula into a path formula. Therefore, we are simplifying h by dropping this **path** operator. So now

$$\pi \models h \;\Leftrightarrow\; s \models h_1$$

$$\Leftrightarrow\; s' \models h_1 \quad \text{(induction hypothesis)}$$

$$\Rightarrow\; \pi' \models h.$$

The reverse direction is similar.
  (5) $h = \mathbf{X}h_1$, a path formula. By the definition of the nexttime operator, $\pi^1 \models h_1$. Since π and π' correspond, so do π^1 and π'^1. Therefore, by the inductive hypothesis, $\pi'^1 \models h_1$, so $\pi' \models h$.

We can use the same argument in the other direction.

(6) $h = h_1 U h_2$, a path formula. Suppose that $\pi \models h_1 U h_2$. By the definition of the until operator, there is a k such that $\pi^k \models h_2$ and for all $0 \leq j < k$, $\pi^j \models h_1$. Since π and π' correspond, so do π^j and π'^j for any j. Therefore, by the inductive hypothesis, $\pi'^k \models h_2$ and $\pi'^j \models h_1$ for all $0 \leq j < k$. Therefore, $\pi' \models h$.

We can use the same argument in the other direction. $\square$

Another property of two equivalent states is that they both have corresponding *computation trees*. For every $s \in S$, $\mathrm{Tr}_n(s)$ is the computation tree of depth n rooted at s. Formally, $\mathrm{Tr}_0(s)$ consists of a single node which has the same label as s. $\mathrm{Tr}_{n+1}(s)$ has as its root a node m with the same label as s. If s has successors $s_1, \ldots, s_p$ in the Kripke structure, then node m will have subtrees $\mathrm{Tr}_n(s_1), \ldots, \mathrm{Tr}_n(s_p)$.

Two trees $\mathrm{Tr}_n(s)$ and $\mathrm{Tr}_n(s')$ *correspond* (denoted $\mathrm{Tr}_n(s) \equiv \mathrm{Tr}_n(s')$) if and only if both of their roots have the same label and for every subtree of depth $n-1$ of the root of one, it is possible to find a corresponding subtree of the root of the other.

Lemma 3.4. *$s E_n s'$ if and only if $\mathrm{Tr}_j(s) \equiv \mathrm{Tr}_j(s')$ for all $j \leq n$.*

Lemma 3.5. *Given a finite set of states $s_1, \ldots, s_n$, there exists a c such that if two states s_i and s_j are not E-equivalent, then $\mathrm{Tr}_c(s_i)$ and $\mathrm{Tr}_c(s_j)$ will not correspond.*

We will call the minimal such value of c for S the *characteristic number* of the structure.

We associate a CTL formula with a tree $\mathrm{Tr}_n(s)$ as follows:

- $\mathscr{F}[\mathrm{Tr}_0(s)] = (p_1 \wedge \cdots \wedge p_u) \wedge (\neg q_1 \cdots \neg q_v)$, where $\mathscr{L}(s) = \{p_1, \ldots, p_u\}$ and $\mathbf{AP} - \mathscr{L}(s) = \{q_1, \ldots, q_v\}$;

- $\mathscr{F}[\mathrm{Tr}_{n+1}(s)] = (\bigwedge_i \mathbf{EX} \mathscr{F}[\mathrm{Tr}_n(s_i)]) \wedge \mathbf{AX}(\bigvee_i \mathscr{F}[\mathrm{Tr}_n(s_i)]) \wedge \mathscr{F}[\mathrm{Tr}_0(s)]$, where s_i is a successor of s.

Lemma 3.6. *$s \models \mathscr{F}[\mathrm{Tr}_n(s)]$ for all $n \geq 0$.*

Lemma 3.7. *If $s \models \mathscr{F}[\mathrm{Tr}_n(s')]$, then $\mathrm{Tr}_n(s) \equiv \mathrm{Tr}_n(s')$.*

Proof. The proof is by induction on n. The basis case is trivial. Thus, we assume that $n > 0$. Let $s_1, s_2, \ldots, s_p$ be the sons of s in $\mathrm{Tr}_n(s)$ and $s'_1, s'_2, \ldots, s'_q$ be the sons of s' in $\mathrm{Tr}_n(s')$.

It is easy to see that s and s' have the same labelling of atomic propositions.

We must show that $\mathrm{Tr}_{n-1}(s_{i_0})$ corresponds to some $\mathrm{Tr}_{n-1}(s'_j)$. Since $s \models \mathscr{F}[\mathrm{Tr}_n(s')]$, $s \models \mathbf{AX}(\bigvee_j \mathscr{F}[\mathrm{Tr}_{n-1}(s'_j)])$. Since s_{i_0} is a successor of s, $s_{i_0} \models \mathscr{F}[\mathrm{Tr}_{n-1}(s'_{j_0})]$ for some j_0. Hence, $\mathrm{Tr}_{n-1}(s_{i_0}) \equiv \mathrm{Tr}_{n-1}(s'_{j_0})$ by our inductive hypothesis.

Finally, we must show that $\mathrm{Tr}_{n-1}(s'_{j_0})$ corresponds to some $\mathrm{Tr}_{n-1}(s_i)$. Since $s \models \mathscr{F}[\mathrm{Tr}_n(s')]$, $s \models \bigwedge_j \mathbf{EX} \mathscr{F}[\mathrm{Tr}_{n-1}(s'_j)]$. Since s'_{j_0} is a successor of s', $s \models \mathbf{EX} \mathscr{F}[\mathrm{Tr}_{n-1}(s'_{j_0})]$. Therefore, there exists an i_0 such that $s_{i_0} \models \mathscr{F}[\mathrm{Tr}_{n-1}(s'_{j_0})]$. Hence, $\mathrm{Tr}_{n-1}(s_{i_0}) \equiv \mathrm{Tr}_{n-1}(s'_{j_0})$ by our inductive hypothesis. $\square$

Lemma 3.8. *If s is a state in a Kripke structure M, then there is a CTL formula $\mathscr{C}(M, s)$ that determines s up to E-equivalence within M, i.e., $\mathscr{C}(M, s)$ is true in s and in every state in M that is E-equivalent to s, but false in every state in M that is not equivalent to s.*

Proof. We choose $\mathscr{C}(M, s) = \mathscr{F}[\mathrm{Tr}_c(s)]$ where c is the characteristic number of M. $\mathscr{C}(M, s)$ is true in s and hence in all states E-equivalent to s. Let s' be a state that is not E-equivalent to s; then $\mathrm{Tr}_c(s) \not\equiv \mathrm{Tr}_c(s')$. Hence, by Lemma 3.7, $s' \not\models \mathscr{C}(M, s)$. □

Theorem 3.9. *Given a Kripke structure M with initial state s_0, there is a CTL formula $F(M, s_0)$ that characterizes that structure up to E-equivalence, i.e., $M', s_0' \models F(M, s_0) \Leftrightarrow s_0 \, E \, s_0'$.*

Proof. For any state s in M, let $s_1, \ldots, s_p$ be the successors of s. We define

$$G(M, s) = \mathbf{AG}\left(\mathscr{C}(M, s) \Rightarrow \bigwedge_i \mathbf{EX}\,\mathscr{C}(M, s_i) \wedge \mathbf{AX} \bigvee_i \mathscr{C}(M, s_i) \right).$$

$G(M, s)$ describes all of the possible transitions from s. $F(M, s_0)$ is the formula $\mathscr{C}(M, s_0) \wedge \bigwedge_s G(M, s)$. If two structures M, s_0 and M', s_0' are equivalent, then, by Theorem 3.2, they satisfy the same formulas. Since $M, s_0 \models F(M, s_0)$, so does M', s_0'.

For the other direction we show by induction on n that if $M', s_0' \models F(M, s_0)$, then $\mathrm{Tr}_n(s_0) \equiv \mathrm{Tr}_n(s_0')$ for all $n \geq 0$. By Lemma 3.4, the two structures are then E-equivalent. □

Corollary 3.10. *Given two structures M and M' with initial states s_0 and s_0' respectively, $s_0 \, E \, s_0'$ if and only if $\forall f \in \mathrm{CTL}^*[M, s_0 \models f \Leftrightarrow M', s_0' \models f]$.*

Corollary 3.11. *Given two structures M and M' with initial states s_0 and s_0' respectively, if there is a formula of CTL* that is true in one and false in the other, then there is also a formula of CTL that is true in the one and false in the other.*

We will illustrate our method of characterizing Kripke structures with the example in Fig. 1. The characteristic number of this structure is 1 since $\mathrm{Tr}_0(s_0) \not\equiv \mathrm{Tr}_0(s_2)$, $\mathrm{Tr}_0(s_1) \not\equiv \mathrm{Tr}_0(s_2)$, and $\mathrm{Tr}_1(s_0) \not\equiv \mathrm{Tr}_1(s_1)$. Let

- $\mathscr{C}(M, s_0) = a \wedge \neg b \wedge \mathbf{EX}(a \wedge \neg b) \wedge \mathbf{EX}(\neg a \wedge b) \wedge \mathbf{AX}(a \wedge \neg b \vee \neg a \wedge b)$,
- $\mathscr{C}(M, s_1) = a \wedge \neg b \wedge \mathbf{EX}(a \wedge \neg b) \wedge \mathbf{AX}(a \wedge \neg b)$,
- $\mathscr{C}(M, s_2) = \neg a \wedge b \wedge \mathbf{EX}(a \wedge \neg b) \wedge \mathbf{AX}(a \wedge \neg b)$.

We can now state the formula that characterizes this structure:

$$F(M, s_0) = \mathscr{C}(M, s_0)$$
$$\wedge \mathbf{AG}(\mathscr{C}(M, s_0) \Rightarrow \mathbf{EX}\,\mathscr{C}(M, s_1) \wedge \mathbf{EX}\,\mathscr{C}(M, s_2)$$
$$\wedge \mathbf{AX}(\mathscr{C}(M, s_1) \vee \mathscr{C}(M, s_2)))$$
$$\wedge \mathbf{AG}(\mathscr{C}(M, s_1) \Rightarrow \mathbf{EX}\,\mathscr{C}(M, s_0) \wedge \mathbf{AX}\,\mathscr{C}(M, s_0))$$
$$\wedge \mathbf{AG}(\mathscr{C}(M, s_2) \Rightarrow \mathbf{EX}\,\mathscr{C}(M, s_0) \wedge \mathbf{AX}\,\mathscr{C}(M, s_0)).$$

Fig. 1. A Kripke structure in which every other state is labelled *A*.

4. Equivalence with respect to stuttering

We first define what it means for two Kripke structures to be equivalent with respect to stuttering. Given two structures M and M' with the same set of atomic propositions, we define a sequence of equivalence relations $E_0, E_1, \ldots$ on $S \times S'$ as follows:

- $s E_0 s'$ if and only if $\mathscr{L}(s) = \mathscr{L}(s')$.
- $s E_{n+1} s'$ if and only if

  (1) for every path π in M that starts in s there is a path π' in M' that starts in s', a partition $B_1 B_2 \ldots$ of π, and a partition $B'_1 B'_2 \ldots$ of π' such that, for all $j \in \mathbb{N}$, B_j and B'_j are both nonempty and finite, and every state in B_j is E_n-related to every state in B'_j, and

  (2) for every path π' in M' starting in s' there is a path π in M starting in s that satisfies the same condition as in (1).

We will say that two paths π and π' *s-correspond* if they satisfy condition (1) above.

Our notion of *equivalence with respect to stuttering* is defined as follows: $s E s'$ if and only if $s E_i s'$ for all $i \geq 0$. Furthermore, we say that M with initial state s_0 is equivalent to M' with initial state s'_0 if $s_0 E s'_0$.

Lemma 4.1. *Given two Kripke structures M and M', there exists an l such that $\forall s \forall s' [s E_l s' \text{ iff } s E s']$.*

Proof. By the definition of E_{l+1}, $s E_{l+1} s' \Rightarrow s E_l s'$, so $E_0 \supseteq E_1 \supseteq E_2 \supseteq \cdots$. Since M and M' are both finite, E_0 must be finite as well, so only a finite number of these containments can be proper. Let E_l be the last relation that is properly included in E_{l-1}. By the definition of proper containment, $\forall m \geq l [E_l = E_m]$, so $s E_l s' \Rightarrow s E_m s'$ for $m \geq l$. Since

$$s E_l s' \Rightarrow s E_{l-1} s' \Rightarrow s E_{l-2} s' \Rightarrow \cdots,$$

we have $s E_l s' \Rightarrow \forall m [s E_m s']$, so $s E_l s' \Rightarrow s E s'$. The other direction is trivial. $\square$

Theorem 4.2. *If $s E s'$, then for every CTL* formula f without the nexttime operator, $s \models f$ iff $s' \models f$.*

The proof is similar to that of Theorem 3.2.

Lemma 4.3. *Given a Kripke structure M, for every state $s \in M$ there is a CTL formula $\mathcal{C}(M, s)$ such that $\forall t \in M[t \models \mathcal{C}(M, s)$ iff $s E t]$.*

Proof. We will prove this by induction on l:

- If $\neg(s E_l t)$, then there is a CTL formula $d_l(s, t)$ such that $\forall v \in M[s E_l v \Rightarrow v \models d_l(s, t)]$ and $t \not\models d_l(s, t)$.
- For every state $s \in M$, there is a CTL formula $\mathcal{C}_l(M, s)$ such that for every $t \in M$, $t \models \mathcal{C}_l(M, s)$ iff $s E_l t$.

$d_l(s, t)$ is a formula that distinguishes between t and states equivalent to s within the structure M, and $\mathcal{C}_l(M, s)$ is a formula that characterizes E_l-equivalence to state s within M.

If we let $\mathcal{C}_l(M, s)$ be a conjunction of $\mathcal{C}_{l-1}(M, s)$ and $d_l(s, t)$ for every t that is not E_l-related to s, the second assertion follows easily. By Lemma 4.1, this condition implies that the lemma is true. Now it is necessary to show how to construct $d_l(s, t)$ by induction on l.

Basis $(l = 0)$: Let $\{p_i\}$ be the set of atomic propositions in $\mathcal{L}(s)$ and $\{q_i\}$ be the set of atomic propositions in $AP - \mathcal{L}(s)$. Now, let

$$\mathcal{C}_0(M, s) = d_0(s, t) = \bigwedge_i p_i \wedge \bigwedge_j \neg q_j.$$

It is clear that this formula is only true in states with the same labelling of atomic propositions as s. Therefore, the base case is established.

Induction: Assume that the result is true for l. We will show it for $l + 1$.

Since $\neg(s E_{l+1} t)$, either there is a path from s without an s-corresponding path from t, or vice versa. In the latter case, we will use the argument below to find a $d_{l+1}(t, s)$ such that $t \models d_{l+1}(t, s)$ and $s \not\models d_{l+1}(t, s)$. We can negate this formula to get the desired $d_{l+1}(s, t)$.

If there is a path from s without an s-corresponding path from t, we can divide this path into blocks $(B_1 B_2 \ldots)$ such that

$$\forall i[x \in B_i \Rightarrow x \models \mathcal{C}_l(M, \text{first}(B_i)) \quad \text{and} \quad \text{first}(B_{i+1}) \not\models \mathcal{C}_l(M, \text{first}(B_i))].$$

Now, there are two cases: either there is a finite path from one state without an s-corresponding path from the other, or there is an infinite path without an s-corresponding path, but every finite prefix of this path has an s-corresponding path.

In the first case, the path from s is finite, so the blocks are finite and there are only a finite number of them (say n). Consider the CTL formula:

$$d_{l+1}(s, t) = \mathcal{C}_l(M, \text{first}(B_1)) \wedge \mathbf{E}[\mathcal{C}_l(M, \text{first}(B_1))\mathbf{U}\mathcal{C}_l(M, \text{first}(B_2))$$

$$\wedge \mathbf{E}[\ldots \mathbf{U}\mathcal{C}_l(M, \text{first}(B_n))] \ldots].$$

It is clear that $s \models d_{l+1}(s, t)$ along the path $B_1 B_2 \ldots B_n$. However, if $t \models d_{l+1}(s, t)$, then there is a path that can be partitioned into blocks $B'_1 B'_2 \ldots B'_n$ such that $\forall i[v \in B'_i \Rightarrow v \models \mathcal{C}_l(M, \text{first}(B_i))]$. Since every state in B_i satisfies $\mathcal{C}_l(M, \text{first}(B_i))$, the inductive

hypothesis and the definition of E_l gives $B_i E_l B'_i$. Therefore, this path from t s-corresponds to the path from s, a contradiction. We conclude that $t \nvDash d_{l+1}(s, t)$.

In the second case, we start by showing that the path from s has only a finite number of blocks by using an argument based on König's Lemma. We can construct a tree rooted at t such that $tt_1 \ldots t_n$ is a path through the tree if and only if there is a path in the Kripke structure

$$tu_1 \ldots u_p t_1 v_1 \ldots v_q t_2 \ldots t_n$$

that s-corresponds to a prefix of the path from s with $B'_1 = \langle tu_1 \ldots u_p \rangle$, $B'_2 = \langle t_1 v_1 \ldots v_q \rangle$, and so on. Now, if the path from s has an infinite number of blocks, this tree must have an infinite number of nodes. Otherwise, if the tree had n nodes, there could be no path of length $n+1$, so the first $n+1$ blocks of the path from s would have no s-corresponding path from t. Since the Kripke structure is finite, we also know that this tree must be finitely branching. Therefore, by König's Lemma, there must be an infinite path through the tree. But this implies that there is an infinite path from t that can be divided into an infinite number of blocks that correspond to the blocks of the path from s, so there is a path from t s-corresponding to the path from s, violating our assumption. Therefore, the path from s has only a finite number of blocks.

So, suppose that there are n blocks, all of which are finite except the last. Consider the CTL formula:

$$d_{l+1}(s, t) = \mathscr{C}_l(M, \text{first}(B_1)) \wedge \mathbf{E}[\, \mathscr{C}_l(M, \text{first}(B_1)) \mathbf{U} \mathscr{C}_l(M, \text{first}(B_2))$$

$$\wedge \mathbf{E}[\ldots \mathbf{U} \, \mathbf{EG} \, \mathscr{C}_l(M, \text{first}(B_n))] \ldots]$$

It is clear that $s \vDash d_{l+1}(s, t)$ along the path $B_1 B_2 \ldots B_n$. However, if $t \vDash d_{l+1}(s, t)$, then there is a path that can be partitioned into blocks $B'_1 B'_2 \ldots B'_n$ such that all of the blocks are finite except B'_n and $\forall i [v \in B'_i \Rightarrow v \vDash \mathscr{C}_l(M, \text{first}(B_i))]$. Since every state in B_i satisfies $\mathscr{C}_l(M, \text{first}(B_i))$, the inductive hypothesis and the definition of E_l gives $B_i E_l B'_i$. We can also divide the infinite blocks B_n and B'_n into an infinite set of blocks containing one state each. Therefore, this path from t s-corresponds to the path from s, so we have a contradiction. We conclude that $t \nvDash d_{l+1}(s, t)$.

Now, these $d_{l+1}(s, t)$ describe the existence or nonexistence of a single path along which some $\mathscr{C}_l$ formulas hold. By the definition of $s E_{l+1} v$, every path from s has an s-corresponding path from v along which the same $\mathscr{C}_l$ formulas hold and vice versa. Therefore, $s E_{l+1} v \Rightarrow v \vDash d_{l+1}(s, t)$.

Therefore, the lemma is true. $\square$

Theorem 4.4. *Given a Kripke structure M with initial state s_0, there is a CTL formula $F(M, s_0)$ that characterizes that structure up to E-equivalence with respect to stuttering, i.e., $M', s'_0 \vDash F(M, s_0) \Leftrightarrow s_0 E s'_0$.*

Proof. For any state s in M, let $s_1, \ldots, s_p$ be the *extended successors* of s, where an extended successor is a state that is not E-related to s and is reachable from s along

a path consisting entirely of states that are E-equivalent to s. Next, we construct $G(M, s)$, which describes all of the transitions from s in M. In this construction, it is convenient to use the *weak until* operator, $\mathbf{A}[f\mathbf{W}g] = \neg\mathbf{E}[\neg g\mathbf{U}\neg f \wedge \neg g]$, which differs from the ordinary until in that it permits an infinite path along which every state satisfies the first argument. So now:

$$G(M, s) = \begin{cases} \bigwedge_i \mathbf{E}[\,\mathscr{C}(M, s)\mathbf{U}\,\mathscr{C}(M, s_i)] \wedge \mathbf{A}\left[\,\mathscr{C}(M, s)\mathbf{W}\bigvee_i \mathscr{C}(M, s_i)\right] \wedge \mathbf{EG}\,\mathscr{C}(M, s) \\ \qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \text{if } s\models\mathbf{EG}\,\mathscr{C}(M, s), \\[2mm] \bigwedge_i \mathbf{E}[\,\mathscr{C}(M, s)\mathbf{U}\,\mathscr{C}(M, s_i)] \wedge \mathbf{A}\left[\,\mathscr{C}(M, s)\mathbf{W}\bigvee_i \mathscr{C}(M, s_i)\right] \wedge \neg\mathbf{EG}\,\mathscr{C}(M, s) \\ \qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \text{otherwise.} \end{cases}$$

Let $F(M, s_0)$ be the formula $\mathscr{C}(M, s_0) \wedge \bigwedge_s \mathbf{AG}(\mathscr{C}(M, s)\Rightarrow G(M, s))$. The correctness of $F(M, s_0)$ is an easy consequence of the next two lemmas and Theorem 4.2. $\square$

Lemma 4.5. $s\models F(M, s)$.

Proof. Since every state is trivially equivalent to itself, $s\models\mathscr{C}(M, s)$ is true by Lemma 4.3. Therefore, if $s\not\models F(M, s)$, then there is a $t \in M$ such that $s\models\mathbf{EF}(\mathscr{C}(M, t)\wedge \neg G(M, t))$. Let v be a state reachable from s that satisfies $\mathscr{C}(M, t)\wedge\neg G(M, t)$. By Lemma 4.3, this condition implies $t\,E\,v$, so t and v must satisfy the same CTL formulas (Theorem 4.2). We will show that $t\not\models\neg G(M, t)$, giving a contradiction. There are four cases.

(1) $t\not\models\mathbf{E}[\,\mathscr{C}(M, t)\mathbf{U}\,\mathscr{C}(M, w)]$, for some extended successor of t, w. By the definition of extended successor, there is a path from t to w and the states on this path are E-related to t. By Lemma 4.3, these states must satisfy $\mathscr{C}(M, t)$. Since $w\models\mathscr{C}(M, w)$ is trivial, this path satisfies $\mathscr{C}(M, t)\mathbf{U}\,\mathscr{C}(M, w)$, which is a contradiction.

(2) $t\not\models\mathbf{EG}\,\mathscr{C}(M, t)$. Since $\mathbf{EG}\,\mathscr{C}(M, t)$ is a conjunct of $G(M, t)$ if and only if $t\models\mathbf{EG}\,\mathscr{C}(M, t)$, we have an immediate contradiction.

(3) $t\not\models\neg\mathbf{EG}\,\mathscr{C}(M, t)$. Since $\neg\mathbf{EG}\,\mathscr{C}(M, t)$ is a conjunct of $G(M, t)$ if and only if $t\not\models\mathbf{EG}\,\mathscr{C}(M, t)$, we have an immediate contradiction.

(4) $t\not\models\mathbf{A}[\,\mathscr{C}(M, t)\mathbf{W}\bigvee_i \mathscr{C}(M, w_i)]$. In this case

$$t\models\mathbf{E}\left[\,\mathscr{C}(M, t)\mathbf{U}(\neg\mathscr{C}(M, t)\wedge\bigwedge_i \neg\mathscr{C}(M, w_i))\right].$$

Let $tt_1\ldots t_n$ be this path, where $t_n\models\neg\mathscr{C}(M, t)\wedge\bigwedge_i \neg\mathscr{C}(M, w_i)$ and $\forall i < n$ $[t_i\models\mathscr{C}(M, t)]$. By Lemma 4.3, $\neg(t_n\,E\,t)$ and $\forall i < n[t_i\,E\,t]$. Therefore, t_n is an extended successor of t. But since $t_n\models\mathscr{C}(M, t_n)$ is trivially true, $t_n\models\bigwedge_i \neg\mathscr{C}(M, w_i)$ cannot be true, so we have a contradiction.

Therefore, the lemma is true. $\square$

Lemma 4.6. *If* $s \models F(M, t)$ *and* $s' \models F(M, t)$, *then* $s \, E \, s'$.

Proof. Since $s \, E \, s'$ if and only if $s \, E_l \, s'$ for all $l \geqslant 0$, we will prove $s \models F(M, t)$ and $s' \models F(M, t)$ implies $s \, E_l \, s'$ by induction on l.

Basis $(l = 0)$: Since $s \models F(M, t)$, $s \models \mathcal{C}(M, t)$ and therefore $s \models \mathcal{C}_0(M, t)$. Similarly, $s' \models \mathcal{C}_0(M, t)$, so $\mathcal{L}(s) = \mathcal{L}(t) = \mathcal{L}(s')$. Therefore, $s \, E_0 \, s'$.

Induction: Assume that the result is true for l. We will now show it for $l + 1$.

We want to show that every path π from s has an s-corresponding path π' from s'. (The proof of the dual is identical.) We only need to consider finite paths since an argument using König's Lemma and similar to the one in the proof of Lemma 4.3 can be used to show that any infinite path without an s-corresponding path must have a prefix without an s-corresponding path. We will use induction on the length of π to prove the slightly stronger result: If $|\pi| \leqslant n$, then there is an s-corresponding path π' such that, for some $v \in M$, $\text{last}(\pi) \models F(M, v)$ and $\text{last}(\pi') \models F(M, v)$.

Basis $(|\pi| = 1)$: In this case, $\pi = \langle s \rangle$. Let $B_1 = \langle s \rangle$ and $\pi' = B_1' = \langle s' \rangle$. By the outer inductive hypothesis, $s \models F(M, t)$ and $s' \models F(M, t)$ imply $s \, E_l \, s'$, so $B_1 \, E_l \, B_1'$. Therefore, the paths s-correspond. Since the last states of each path satisfy $F(M, t)$, the base case is true.

Induction: Assume the result for $|\pi| \leqslant n$. Suppose that $\pi = ss_1 s_2 \ldots s_n$, a path of length $n + 1$. Now, $ss_1 s_2 \ldots s_{n-1}$ is a path of length n, so, by the inner inductive hypothesis, there is an s-corresponding path π' such that $\text{last}(\pi') \models F(M, v)$ and $s_{n-1} \models F(M, v)$ for some $v \in M$. Let $B_1 B_2 \ldots B_m$ and $B_1' B_2' \ldots B_m'$ be the partitions that show that these paths s-correspond. There are three cases.

(1) $s_n \not\models \mathcal{C}(M, v)$. Since $s_{n-1} \models F(M, v)$, we can infer that $s_{n-1} \models \mathbf{A}[\mathcal{C}(M, v) \mathbf{W} \bigvee_i \mathcal{C}(M, w_i)]$, where the w_i are the extended successors of v. Since $s_{n-1} s_n$ is a path along which $\mathcal{C}(M, v) \mathbf{W} \bigvee_i \mathcal{C}(M, w_i)$ holds and since s_n does not satisfy $\mathcal{C}(M, v)$, we conclude that there must be an extended successor of v, x, such that $s_n \models \mathcal{C}(M, x)$. Since s_n is a successor of s_{n-1}, it must satisfy all of the **AG** formulas that s_{n-1} satisfies, so $s_n \models F(M, x)$.

From $\text{last}(\pi') \models F(M, v)$ we can infer that $\text{last}(\pi') \models \mathcal{C}(M, v) \wedge \mathbf{E}[\mathcal{C}(M, v) \mathbf{U} \mathcal{C}(M, x)]$. Therefore, there is a path $s_1' s_2' \ldots s_k'$ where $s_1' = \text{last}(\pi')$, $\forall i < k[s_i' \models \mathcal{C}(M, v)]$, and $s_k' \models \mathcal{C}(M, x)$. Now let

$$\pi = B_1 \ldots B_m \langle s_n \rangle \quad \text{and} \quad \pi' = B_1' \ldots B_{m-1}' \langle B_m', s_2' \ldots s_{k-1}' \rangle \langle s_k' \rangle.$$

Since s_n and s_k' both satisfy $F(M, x)$, the outer induction hypothesis gives $\langle s_n \rangle \, E_l \, \langle s_k' \rangle$. Similarly, since all the states in B_m, B_m', and $\langle s_2' \ldots s_{k-1}' \rangle$ satisfy $F(M, v)$, they are all E_l-related to each other. Therefore, π and π' s-correspond with $\text{last}(\pi) \models F(M, x)$ and $\text{last}(\pi') \models F(M, x)$.

(2) $s_n \models \mathcal{C}(M, v)$ and $v \models \mathbf{EG} \mathcal{C}(M, v)$. Since s_n must satisfy the same **AG** formulas as s_{n-1}, $s_n \models F(M, v)$. Now, $\text{last}(\pi') \models F(M, v)$, so $\text{last}(\pi') \models \mathbf{EG} \mathcal{C}(M, v)$. Therefore, $\text{last}(\pi')$ must have a successor s_1' which also satisfies $\mathcal{C}(M, v)$. Since this state must

also satisfy all of the **AG** formulas, $s_1' \models F(M, v)$. Therefore, by the outer induction hypothesis, $s_n E_l s_1'$. So if we let $B_{m+1} = \langle s_n \rangle$ and $B_{m+1}' = \langle s_1' \rangle$, the paths s-correspond.

(3) $s_n \models \mathscr{C}(M, v)$ and $v \not\models \mathbf{EG}\,\mathscr{C}(M, v)$. By the reasoning above, $s_n \models F(M, v)$, so $s_n E_l \mathrm{last}(B_m')$. Therefore, π s-corresponds to π' with the same partition except that s_n is added to B_m.

We must also show that the blocks of the partitions are finite. The only problem is case (3), in which we might add an infinite number of states to a block of π. In this case, each of the states added to B_m satisfy $F(M, v)$, so if we add an infinite number of states to this block $\mathrm{first}(B_m) \models \mathbf{EG}\,\mathscr{C}(M, v)$ must be true. But since $\mathrm{first}(B_m) \models F(M, v)$, $\mathrm{first}(B_m) \models \neg\mathbf{EG}\,\mathscr{C}(M, v)$, so we have a contradiction. Therefore, all of the blocks of the partition must be finite.

Therefore, the lemma is true. $\square$

Corollary 4.7. *Given two structures M and M' with initial states s_0 and s_0' respectively, $s_0 E s_0'$ if and only if, for all CTL* formulas f without the nexttime operator, $M, s_0 \models f \Leftrightarrow M', s_0' \models f$.*

Corollary 4.8. *Given two structures M and M' with initial states s_0 and s_0' respectively, if there is a formula of CTL* without the nexttime operator that is true in one and false in the other, then there is also a formula of CTL without the nexttime operator that is true in the one and false in the other.*

5. Algorithm for stuttering equivalence

In this section we show how to compute the relation for equivalence with respect to stuttering for states within a single Kripke Structure M. The method that we suggest is polynomial in the number of states of M. To determine equivalence between states in two different Kripke structures M_1 and M_2, we form a Kripke structure M_{12} that is the disjoint union of these structures and check equivalence between the corresponding states in the combined structure.

We construct a relation C on $S \times S$ that is identical to the relation E defined in Section 4. $C = \bigcap_n C_n$ where C_n is defined as follows:

- $C_0 = \{(s, s') \mid \mathscr{L}(s) = \mathscr{L}(s')\}$;
- In order to define C_{n+1} we must first define the set $\mathrm{NEXT}_{n+1}(s)$ of extended successors of s. We define this set in terms of the set $\mathrm{ST}_{n+1}(s)$ of *stuttering states* of s. $\mathrm{ST}_{n+1}(s) = \bigcup_k \mathrm{ST}_{n+1}^k(s)$, where

(i) $\mathrm{ST}_{n+1}^0(s) = \{s\}$,

(ii) $\mathrm{ST}_{n+1}^{k+1}(s) = \mathrm{ST}_{n+1}^k(s) \cup \{s' \mid s' \notin \mathrm{ST}_{n+1}^k(s) \wedge \exists s'' \in \mathrm{ST}_{n+1}^k(s)[s'' \to s'] \wedge s'\, C_n\, s\}$

$\mathrm{NEXT}_{n+1}(s) = \{s' \mid s' \notin \mathrm{ST}_{n+1}(s) \wedge \exists s'' \in \mathrm{ST}_{n+1}(s)[s'' \to s']\}$.

We will also use a predicate $\text{LOOP}_n(s)$ that is true iff there is a cycle containing only states in $\text{ST}_n(s)$.

Now we can define C_{n+1} as follows:

$$C_{n+1} = \{(s, s') | \text{LOOP}_{n+1}(s) = \text{LOOP}_{n+1}(s') \wedge s \, C_n \, s'$$

$$\wedge \, \forall s_1 \in \text{NEXT}_{n+1}(s) \exists s_1' \in \text{NEXT}_{n+1}(s')[s_1 \, C_n \, s_1']$$

$$\wedge \, \forall s_1' \in \text{NEXT}_{n+1}(s') \exists s_1 \in \text{NEXT}_{n+1}(s)[s_1 \, C_n \, s_1']\}.$$

The proof that the relation C constructed above is actually equal to the relation E defined in Section 4 is tedious but straightforward and will not be given in this paper. Since the inductive structures of the definitions of the two relations are different, it is necessary to split the proof into two parts: the first part shows that $C \subseteq E_i$ for every i; the second part shows that $E \subseteq C_i$ for every i. The intuition behind the proof is easy to understand. $\text{ST}_n(s)$ gives the set of states that are C_n-equivalent to s and can be reached from s along a path containing states which are all C_n-equivalent. Given a path π starting at s, the first block of that path in the definition of E is determined by the prefix of π in $\text{ST}_n(s)$.

Computing ST_n requires time $O(|S|^3)$. Computing C_{n+1} given C_n requires time $O(|S|^4)$ since at most $|S|^2$ pairs of states must be checked and each pair requires $O(|S|^2)$ time to check. The algorithm terminates as soon as $C_n = C_{n+1}$. Since at any previous step k, the number of equivalence classes of C_{k+1} is strictly greater than the number of equivalence classes of C_k and since C has at most $|S|$ equivalence classes, there are at most $|S|$ steps in the construction of C. It follows that the complexity of the entire algorithm is $O(|S|^5)$.

If we replace each equivalence class of C by a single state, this algorithm can also be used to minimize the number of states in the structure.

6. Conclusion

The results of our paper have a number of surprising implications. For example, if a specification of a finite-state concurrent program in CTL* is sufficiently detailed so that there is only one program (modulo one of our notions of equivalence) that meets the specification, then an equivalent specification could have been written in CTL instead. Another surprising consequence is that if a CTL* formula is not equivalent to any CTL formula, then it must have an infinite number of mutually inequivalent finite models. To see that this result is true, we first observe that since CTL* has the finite-model property, it must be the case that if two CTL* formulas have the same finite models, they must have the same infinite models as well. Otherwise, if f_1 had an infinite model M that was not a model of f_2, $f_1 \wedge \neg f_2$ would have an infinite model, but no finite models, contradicting the finite-model property of CTL* [7]. Therefore, we can characterize a CTL* formula by the set of *finite*

models in which it is satisfied. If a CTL* formula is satisfied by only a finite number of equivalence classes of finite models, then the formula is equivalent to the disjunction of the CTL formulas that characterize the individual equivalence classes.

There are a number of directions for further research. First, from our construction, it appears that the characteristic formula of a Kripke structure might be quite large. It would be nice to have a lower bound on the size of this formula in terms of the size of the Kripke structure. Also, we conjecture that the $O(|S|^5)$ algorithm in Section 5 can be improved significantly. Finally, it would be interesting to see which of our results carry over to Kripke structures with fairness constraints, i.e., Büchi automata.

References

[1] A. Arnold and P. Crubille, A linear algorithm to solve fixed-point equations on graphs, Tech. Rept. I-8632, Université de Bordeaux, November, 1986.

[2] M. Ben-Ari, A. Pnueli and Z. Manna, The temporal logic of branching time, *Acta Inform.* **20** (1983) 207–226.

[3] M. Browne, E. Clarke, D. Dill and B. Mishra, Automatic verification of sequential circuits using temporal logic, *IEEE Trans. Comput.* **35** (1986) 1035–1044.

[4] E.M. Clarke and E.A. Emerson, Design and synthesis of synchronization skeletons using branching time temporal logic, in: *Proc. Workshop on Logic of Programs*, Yorktown-Heights, NY, Lecture Notes in Computer Science **131** (Springer, Berlin, 1981) 52–71.

[5] E.M. Clarke, E.A. Emerson and A.P. Sistla, Automatic verification of finite-state concurrent systems using temporal logic specifications, *ACM Trans. Programming Languages and Systems* **8** (1986) 244–263.

[6] E.A. Emerson and J.Y. Halpern, "Sometimes" and "Not Never" revisited: on branching versus linear time temporal logic, in: *Proc. ACM Symp. on Principles of Programming Languages*, Austin, TX (1982) 127–140.

[7] E.A. Emerson and P. Sistla, Deciding full branching-time logic, in: *Proc. 16th Ann. ACM Symp. on Theory of Computing*, Washington, DC (1984) 19–24.

[8] S. Graf and J. Sifakis, From synchronization tree logic to acceptance model logic, in: *Logics of Programs*, Lecture Notes in Computer Science **193** (Springer, Berlin, 1985) 128–142.

[9] M. Hennessy and R. Milner, On observing nondeterminism and concurrency, in: *Proc. 7th ICALP*, Lecture Notes in Computer Science **85** (Springer, Berlin, 1980) 299–309.

[10] M. Hennessy and C. Stirling, The power of the future perfect in program logics, *Inform. and Control* **67** (1985) 23–52.

[11] G.E. Hughes and M.J. Creswell, *An Introduction to Modal Logic* (Methuen, London, 1977).

[12] L. Lamport, "Sometimes" is sometimes "Not Never", in: *Proc. 7th Ann. ACM Symp. on Principles of Programming Languages*, Las Vegas, NV (1980) 174–185.

[13] L. Lamport, What good is temporal logic?, in: *Proc. Internat. Federation for Information Processing* (1983) 657–668.

[14] O. Lichtenstein and A. Pnueli, Checking that finite state concurrent programs satisfy their linear specification, in: *Conf. Record of the 12th Ann. ACM Symp. on Principles of Programming Languages*, New Orleans, LA (1985) 97–107.

[15] R. Milner, *A Calculus of Communicating Systems*, Lecture Notes in Computer Science **92** (Springer, Berlin, 1979).

[16] B. Mishra and E. Clarke, Hierarchical verification of asynchronous circuits using temporal logic, *Theoret. Comput. Sci.* **38** (1985) 269–291.

[17] Z. Manna and P. Wolper, Synthesis of communicating processes from temporal logic specifications, *ACM Trans. on Programming Languages and Systems* **6** (1984) 68–93.

[18] A. Pnueli, Linear and branching structures in the semantics and logics of reactive systems, in: *Proc. 12th ICALP*, Lecture Notes in Computer Science **194** (Springer, Berlin, 1985) 15–32.

[19] J.P. Quielle and J. Sifakis, Specification and verification of concurrent systems in CESAR, in: *Proc. 5th Internat. Symp. on Programming* (1981) 337–350.

[20] A.P. Sistla and E.M. Clarke, Complexity of propositional linear temporal logics, *J. Assoc. Comput. Mach.* **32** (1985) 733–749.

[21] M.Y. Vardi and P. Wolper, An automata-theoretic approach to automatic program verification, in: *Proc. Logic in Computer Science*, Cambridge, MA (1986) 332–344.

[22] P. Wolper, Specification and synthesis of communicating processes using an extended temporal logic, in: *Proc. 9th Ann. ACM Symp. on Principles of Programming Languages*, Albuquerque, NM (1982) 20–33.