THE JOURNAL OF SYMBOLIC LOGIC
Volume 66, Number 3, Sept. 2001

AN AXIOMATIZATION OF FULL COMPUTATION TREE LOGIC

M. REYNOLDS

Abstract. We give a sound and complete axiomatization for the full computation tree logic, CTL*, of R-generable models. This solves a long standing open problem in branching time temporal logic.

§1. Introduction. CTL*, which is occasionally called full computation tree logic, was first described in [Emerson and Sistla, 1984] and [Emerson and Halpern, 1986]. By using a slightly unusual semantics based on paths through Kripke (or transition) structures, CTL* is able to extend, in expressiveness, both the computation tree logic, CTL, of [Clarke and Emerson, 1981], a simple branching logic, and the standard propositional linear temporal logic, PLTL of [Pnueli, 1977].

The language of CTL*, which is a propositional temporal language, is built recursively from the atomic propositions using the next X and until U operators of PLTL, and the existential path switching operator E of CTL as well as classical connectives. This language is appropriate for describing any situation with paths as countable sequences of states, the propositions having a truth evaluation at states and the possibility of at least some states lying on more than one path. We will be interested in the logic obtained by restricting attention to Kripke structures with states, a total accessibility relation between them and the set of all paths which arise by moving from state to state along the accessibility relation (which is usually called R). This standard semantics for CTL* is thus called the semantics over R-generable models.

The main uses of CTL* in computer science are for developing and checking the correctness of complex reactive systems. See [Emerson, 1990] for a survey. CTL* is also used widely as a framework for comparing other languages more appropriate for specific reasoning tasks of this type. See the description in [Emerson, 1996]. These include the purely linear and purely branching sub-languages as well as languages such as [Bernholtz and Grumberg, 1994] which allow a limited amount of interplay between these two aspects.

Validity of formulas of CTL* is known to be decidable. This was proved in [Emerson and Sistla, 1984] using an automata-theoretic approach. Essentially one

Received November 18, 1998; revised March 10, 1999.

The work was partially supported by EPSRC grants GR/K54946 and GR/L82441. Thanks to Alberto Zanardo, Colin Stirling, the Logic and Computation Group at Manchester Metropolitan University and the Algebraic Logic Group at the University of London for helpful discussions. Especial thanks to the anonymous referee for very careful reading, many useful suggestions and some serious rewrites.

0022-4812/01/6603-0002/$5.70

finds an equivalent Rabin tree automaton to the negation of the formula, i.e. one accepting exactly the models of the negated formula, and then tests the automaton for emptiness. To get efficient procedures, one is able to exploit a normal form result to allow most of the work to be done by procedures for finding linear automata equivalents to PLTL formulas. Making use of the specific form of such linear automata allowed [Emerson and Jutla, 1988] to describe a decision procedure of deterministic double exponential time complexity in the length of the formula. This agrees with the lower bound found in [Vardi and Stockmeyer, 1985].

Even though a decision procedure exists, and so we know that CTL* is certainly recursively axiomatisable, there are still reasons why an explicit and simple axiomatization might be useful and interesting. The immediate CTL* uses of a sound and complete axiomatization include providing intuitive axioms for manual proofs, for automated proof assistance, and for providing a means of showing completeness of other theorem proving methods (such as tableau or resolution based approaches). As pointed out in [Emerson and Sistla, 1984], [Emerson, 1990], [Stirling, 1992] and [Kaivola, 1996], up until now, providing such an axiomatization has remained an open problem.

There are some related results in computer science. In [Stirling, 1992], an axiomatization is given for a logic ∀LTFC which uses the language of CTL* but uses a more general semantics. In this logic we may restrict the use of the path quantifier to a given subset of all of the paths through a Kripke structure. The only requirements are that the chosen set of paths is closed under taking suffixes (i.e. it is suffix closed) and is closed under putting together a finite prefix of one path with the suffix of any other path such that the prefix ends at the same state as the suffix begins (i.e. the set is fusion closed). This logic can also be defined over trees of height ω provided we again restrict the path quantifier to a certain set, or *bundle*, of paths (or branches). Thus, in this paper we call this the logic of bundled ω-trees.

If we restrict this logic by requiring the bundle to contain all the branches of the tree, i.e. we consider the *completely* bundled ω-trees, then the resulting logic is standard CTL* again. This restriction is effectively the same as putting an extra closure condition on the set of paths which define the semantics of ∀LTFC. The extra restriction is that the set of paths must be limit closed, i.e. if a sequence of prefixes of paths from the set is strictly increasing then the path which is the limit of the sequence is also in the set.

These equivalent restrictions— requiring completeness of the bundle or limit closure of the path set— result in some extra validities and so, to get an axiomatization for CTL* we need to add some extra axioms or rules to Stirling's axiomatization for ∀LTFC (i.e. for bundled ω-trees).

There are other logics in which the extra limit closure condition causes problems for axiomatization. In computer science the problem has been solved in the case of the less expressive CTL (in [Emerson and Halpern, 1982] and) in [Stirling, 1992] by the addition of a simple induction axiom (seen later as example 1 in section 16 here) and recently in the case of the much more expressive νCTL* in [Kaivola, 1996] by the addition of a limit closure rule.

In philosophical considerations of branching time, such as in the logics of historical necessity of [Burgess, 1980], [Zanardo, 1996], the move from bundled trees

to complete trees is also a long standing open problem as far as providing a complete axiom system is concerned. The language is usually Prior's F and P along with a branch-switching modality $\Diamond$, and the models are much more general trees but the problem is very similar. A complete axiom system for the bundled version has been given in [Zanardo, 1985]. The problem with completeness is surveyed in [Zanardo, 1996].

The general problem of definability in many of these branching time temporal logics is considered in [Dam, 1992] and [Zanardo, Barcellan, and Reynolds, 1999]. Here we consider the question of finding formulas whose validity on a frame across all possible valuations for the atoms is equivalent to the completeness of the frame. In the latter paper it is shown that a simple formula called δ, which looks very much like a possible limit closure axiom, defines completeness in the class of all bundled ω-trees. In contrast, as shown in the former paper, there is no such formula which defines limit closure within the class of all suffix and fusion closed path frames. This just shows that (frame) definability is certainly not closely related to axiomatization questions.

In this paper we are able to provide a sound and complete explicit (Hilbert-style) axiom system for the standard CTL* logic, i.e. over R-generable models.

The axiomatization itself provides some interest. There is a new axiom and a new rule. The new axiom is, as is probably widely expected, clearly an inductive path construction axiom. I call this the limit closure (LC) axiom. It has similarity with the limit closure axiom for CTL, the limit closure rule of [Kaivola, 1996] and the formula δ of [Zanardo, Barcellan, and Reynolds, 1999], which defines limit closure in the class of all bundled ω-trees.

The new rule is novel. I call it the Auxiliary Atoms rule (AA). It has some affinity with Gabbay's IRR rule [Gabbay, 1981] (for characterizing irreflexivity of time in general temporal logics) and generalizations (see, for example, [Gabbay, Hodkinson, and Reynolds, 1994]). Like the IRR rule, AA allows the introduction of fresh atoms to a proof and allows the interpretation of the fresh atoms to have particular properties. Also like the IRR rule, which by design is only generally useful when the flow of time is irreflexive, the AA rule is generally only useful when the Kripke (transition) frame is a tree. See the proof of lemma 6 for details. Unlike the IRR rule which involves one fresh atom, the AA rule allows the simultaneous expansion of the language by a possibly large (but finite) number of fresh atoms. The assumption we can make about their interpretations is also quite complex, although, particularly for automata users, it should be easily able to be applied and understood intuitively. Most of the work in stating the conditions of application of the rule is done in the side condition and so this is a little complicated. To some extent, then, like other IRR-style rules, the AA rule might be regarded as an infinite set of rules. However, the side condition is clearly an easily decidable syntactic test and so the AA rule is no worse than the common substitution rule of many logics (—a rule which, incidentally, is not valid for CTL*).

The axiomatic completeness proof used here is also interesting. Of course, as CTL* is not compact we can only manage a weak completeness result, i.e. showing that any given consistent formula is satisfiable in some Kripke structure. We use a basic step-by-step filtration style construction but we let a deterministic Rabin

linear automaton loose in the background and we impose an elaborate banning mechanism as we go along.

The basic filtration-like idea is to build a model from sets of formulas maximally consistent in some finite closure set determined by the given formula. Certain equivalence classes of these sets form the states of the model. The way of defining the accessibility relation between states (and hence defining paths) has some similarities with the scheduling ideas of the bundled completeness proof in [Stirling, 1992] but is more akin to the usual step-by-step construction in historical necessity proofs such as [Zanardo, 1985]. Because these latter proofs, and ours, generally use the same set of formulas at many different places in the construction we rather think of the set of formulas as a label on some abstract point objects.

The reason that such proofs have trouble with the limit closure condition on paths is that, in the limit, the step-by-step construction produces many more paths than were ever chosen explicitly to start being constructed at any finite stage. In a filtration-based completeness proof we want each of the points of the limit structure to model exactly the formulas in its label. However, there is no guarantee that one of these non-explicit paths will make a formula $A\psi$ true at its initial point even if $A\psi$ is in the label of that point. The solution adopted here is to consider a particular linear automaton A which can be set loose on each path— whether explicitly constructed or not. We make sure that the automaton accepts exactly paths which mess up our construction in this way. We also make sure that no paths at all are accepted by A.

So how do we make sure that no paths are accepted? The answer is as follows. We use the AA rule to introduce fresh atoms corresponding to each state of the automaton A and we interpret the atoms according to the state that A would be in if it traveled to that point of the construction. A is deterministic so there is only one such state. The Rabin acceptance condition is defined by a set of pairs (U_i, V_i) of sets of its states and requires that there is some such pair such that some state in U_i comes up infinitely often along the path and no state in V_i does. To ensure non-acceptance along each path of our construction, we impose a banning mechanism which ensures that atoms from U_i do not come up very often in labels of points after it seems that atoms from V_i have stopped being in the labels.

A linear automaton is also used in an axiomatic completeness proof in [Kesten and Pnueli, 1995]. Here, the linear time logic with quantification of propositions is axiomatized and the proof uses a back and forth technique involving automata. Of course, having quantification of propositions available in the logic makes AA style rules unnecessary.

Another use of a linear automaton is made in [Kaivola, 1996] in which a sound and complete axiomatization is given for the extended CTL* logic, νCTL*, which allows operators of the linear time mu-calculus (which is more expressive than PLTL) as well as the path operator A. The models are R-generable and so Kaivola has the same problem as we have in ensuring limit closure of the constructed model. The solution in this case involves the combination of a deterministic Rabin automaton (of an even more restricted form) and the transformation of formulas to a corresponding normal form. The axiom system, however, needs no new rule apart from a new limit closure rule to succeed, because the fixed-point operators in νCTL* allow any number of extra propositions with carefully defined interpretations to be brought to bear on the proof.

We can also mention that an automaton is used to deal with infinite paths in a tableau refutation in [Walukiewicz, 1995] in the proof of completeness for an axiom system for the modal μ-calculus. Modal μ-calculus, like CTL*, is a logic for reasoning about transition structures but it uses fixpoint operators and is more expressive. There is no issue of limit closure, however, as infinite paths are only definable in the calculus in terms of limits of successor relations. Essentially there is limit closure by definition.

In this paper we first review the linear temporal logic PLTL, its axiom system and the result about Rabin automata which we need. After defining CTL* in section 3, we look at some of the variant semantics and relate the corresponding logics. In section 5, we recall the axiom system for bundled ω-trees and then, in the next two sections, introduce our new axiom LC and rule AA. Most of the remainder of the paper consists of a fairly detailed presentation of the completeness proof of the full system.

After a few examples of the axiom system in action, we mention a few possible extensions to this work and suggestions for how the results or techniques might have other applications. One important question is whether the AA rule is really necessary. Another is whether it, along with the use of an automaton, could have benefits in finding complete axiomatizations of other similar logics.

1.1. General notation. A *sequence* will usually mean an ω-sequence $\langle s_0, s_1, \dots \rangle$ (unless otherwise specified). If $\sigma = \langle s_0, s_1, s_2, \dots \rangle$ is a sequence (of anything) then we will refer to s_i as σ_i and the *suffix* sequence $\langle s_i, s_{i+1}, s_{i+2}, \dots \rangle$ as $\sigma_{\geq i}$. The set of all ω-sequences of objects from a set S will be denoted by ${}^{\omega}S$.

The set of all finite sequences of objects from a set S will be denoted by ${}^{<\omega}S$.

The set of all subsets of a set S will be denoted $\wp(S)$.

If $A \subseteq C$ and $\sigma \in {}^{\omega}(\wp(C))$ then the restriction $\sigma|_A$ of σ to A is the sequence $\langle \sigma_0 \cap A, \sigma_1 \cap A, \dots \rangle$ from ${}^{\omega}(\wp(A))$.

§2. Propositional linear temporal logic. Much of the work in the completeness proof is done by reasoning along branches of trees so we review the linear temporal logic PLTL.

We fix a countable set $\mathscr{L}$ of atomic propositions. Formulas are evaluated in *ω-structures* in the signature $\mathscr{L}$. An ω-structure $\sigma = \langle \sigma_0, \sigma_1, \dots \rangle$ is a countable sequence of subsets of $\mathscr{L}$ where $p \in \sigma_i$ represents the atom p being true at time i in the structure.

The formulas of PLTL are built from ***true*** and the atomic propositions in $\mathscr{L}$ recursively using classical connectives $\neg$ and $\wedge$ as well as the temporal connectives X and U: if α and β are formulas then so are $X\alpha$ and $\alpha U \beta$.

Truth of formulas is evaluated at ω-structures. We write $\sigma \models \alpha$ iff the formula α is true of the sequence σ. This is defined formally recursively by:

$\sigma \models \mathbf{true}$

$\sigma \models p$ iff $p \in \sigma_0$, any $p \in \mathscr{L}$

$\sigma \models \neg\alpha$ iff $\sigma \not\models \alpha$

$\sigma \models \alpha \wedge \beta$ iff $\sigma \models \alpha$ and $\sigma \models \beta$

$\sigma \models X\alpha$ iff $\sigma_{\geq 1} \models \alpha$

$\sigma \models \alpha U \beta$ iff there is some $i \geq 0$ such that $\sigma_{\geq i} \models \beta$ and for each j, if $0 \leq j < i$ then $\sigma_{\geq j} \models \alpha$

If $\sigma \models \alpha$ then we say that σ is a model of α. If α has a model then we say that α is satisfiable. If $\sigma \models \alpha$ for any ω-structure σ then we say that α is valid in PLTL and we write $\models_L \alpha$.

As well as the usual abbreviations $\vee$, $\rightarrow$ and $\leftrightarrow$, we have $F\alpha \equiv \textit{true}\, U\, \alpha$ and $G\alpha \equiv \neg F \neg \alpha$.

Complete axiom systems exists in various versions in the literature. The one we present here will be incorporated into our branching-time system. The rules are modus ponens and temporal generalization,

$$\frac{\alpha, \alpha \rightarrow \beta}{\beta} \qquad \frac{\alpha}{G\alpha}.$$

The axioms are all substitution instances of the following:

C0 any propositional tautology
C1 $F \neg\neg \alpha \leftrightarrow F\alpha$
C2 $G(\alpha \rightarrow \beta) \rightarrow (G\alpha \rightarrow G\beta)$
C3 $G\alpha \rightarrow (\alpha \wedge X\alpha \wedge X(G\alpha))$
C4 $X \neg \alpha \leftrightarrow \neg X\alpha$
C5 $X(\alpha \rightarrow \beta) \rightarrow (X\alpha \rightarrow X\beta)$
C6 $G(\alpha \rightarrow X\alpha) \rightarrow (\alpha \rightarrow G\alpha)$
C7 $(\alpha\, U\, \beta) \leftrightarrow (\beta \vee (\alpha \wedge X(\alpha\, U\, \beta)))$
C8 $(\alpha\, U\, \beta) \rightarrow F\beta$

We define derivability $\vdash_L$ in the usual way and say that a formula α is consistent iff we do not have $\vdash_L \alpha \rightarrow \textit{false}$.

Note that the original system in [Gabbay, Pnueli, Shelah, and Stavi, 1980] used axioms rather than substitution instances of axiom schemas and so also included the substitution rule. We will avoid the substitution rule in this paper as it is not valid for CTL*.

In [Gabbay, Pnueli, Shelah, and Stavi, 1980], it was proved that

THEOREM 1. *The system above is sound and complete for* PLTL: *i.e.* $\vdash_L \alpha$ *iff* $\models_L \alpha$.

Soundness is established by the usual induction on the length of proofs. The vaguely filtration-based completeness proof, with a novel fair scheduling idea, is interesting and some of its elements appear in our branching-time completeness proof.

2.1. Automata. Suppose that P is a finite set of propositional atoms. A *deterministic (Rabin) automaton* recognizing ω-sequences from $\wp(P)$ is a 4-tuple $A = (Q, s_0, \rho, L)$ where

- Q is a finite non-empty set called the set of *states*,
- $s_0 \in Q$ is the *initial state*,
- $\rho : Q \times \wp(P) \rightarrow Q$ is the *transition function* and
- the finite set $L \subseteq (\wp(Q) \times \wp(Q))$ is the *set of accepting pairs*.

A *run* of A on an ω-structure σ in the signature P is a sequence of states $s_0, s_1, s_2, \ldots$ from Q such that for each $i < \omega$, $\rho(s_i, \sigma_i) = s_{i+1}$.

We say that the Rabin automaton $A = (Q, s_0, \rho, L)$ accepts σ iff there is some run $s_0, s_1, s_2, \ldots$ on σ and some pair $(U, V) \in L$ such that no state in V is visited infinitely often but there is some state in U visited infinitely often. It is clear that we may assume that for each pair (U, V) we have $U \cap V = \emptyset$: just replace each original

U by $U' = U \setminus V$. Note that a deterministic automaton will have a unique run on any given structure.

We will be wanting to translate a temporal formula into an equivalent automaton: i.e. one that accepts exactly the models of the formula. Various well-known results including those in [McNaughton, 1966], and [Safra, 1988] give us:

THEOREM 2. *For any* PLTL *formula* α *using atoms from the finite set* P *there is a deterministic Rabin automaton* $(Q, s_0, \rho, \{(U_1, V_1), \ldots, (U_k, V_k)\})$ *which recognizes* ω*-sequences of elements of* $\wp(P)$ *and accepts exactly the models of* α.

§3. CTL*. The language of CTL* is used to describe several different types of structures and so there are really several different logics here. We will be mostly interested in the logic of R-generable sets of paths on transition structures, which we will call Kripke structures. In most papers it is this logic which is referred to as CTL*: this is the standard CTL* logic. In the next section we briefly look at some other semantics for the language as we need to use some of them in the axiomatic completeness proof.

We fix a countable set $\mathscr{L}$ of atomic propositions.

DEFINITION 1. For us a *Kripke frame* is a pair (S, R) where:

S is the non-empty set of *states*

R is a total binary relation $\subseteq S \times S$
(i.e. for every $s \in S$, there is some $t \in S$ such that $(s, t) \in R$)

Note that usually in modal logic a Kripke frame's accessibility relation R is not necessarily assumed to be total.

Formulas are evaluated in *(Kripke) structures*:

DEFINITION 2. A structure is a triple $M = (S, R, g)$ where:

(S, R) is a Kripke frame

g : $S \to \wp(\mathscr{L})$ is a labelling of the states with sets of atoms

Such structures are often called *transition* structures.

A *fullpath* in M (or in (S, R)) is an infinite sequence $s_0, s_1, s_2, \ldots$ of states of M such that for each i, $(s_i, s_{i+1}) \in R$. For the fullpath $b = s_0, s_1, s_2, \ldots$, and any $i \geq 0$, we write b_i for the state s_i and $b_{\geq i}$ for the fullpath $s_i, s_{i+1}, s_{i+2}, \ldots$.

The formulas of CTL* are built from ***true*** and the atomic propositions in $\mathscr{L}$ recursively using classical connectives $\neg$ and $\wedge$ as well as the temporal connectives X, U and E: if α and β are formulas then so are $X\alpha$, $\alpha U \beta$ and $E\alpha$. As well as the linear abbreviations, $\vee$, $\to$, $\leftrightarrow$, F and G, we have $A\alpha \equiv \neg E \neg \alpha$.

We shall write $\psi \leq \phi$ if ψ is a subformula of ϕ. If S is a finite set of formulas, define $\bigwedge S = \alpha_1 \wedge \cdots \wedge \alpha_n$, after enumerating $S = \{\alpha_1, \ldots, \alpha_n\}$ in some particular order.

Truth of formulas is evaluated at fullpaths in structures. We write $M, b \models \alpha$ iff the formula α is true of the fullpath b in the structure $M = (S, R, g)$. This is defined formally recursively by:

$M, b \models \textbf{\textit{true}}$
$M, b \models p$ iff $p \in g(b_0)$, any $p \in \mathscr{L}$
$M, b \models \neg\alpha$ iff $M, b \not\models \alpha$
$M, b \models \alpha \wedge \beta$ iff $M, b \models \alpha$ and $M, b \models \beta$
$M, b \models X\alpha$ iff $M, b_{\geq 1} \models \alpha$
$M, b \models \alpha\, U\, \beta$ iff there is some $i \geq 0$ such that $M, b_{\geq i} \models \beta$ and for each j, if $0 \leq j < i$ then $M, b_{\geq j} \models \alpha$
$M, b \models E\alpha$ iff there is some fullpath b' such that $b_0 = b'_0$ and $M, b' \models \alpha$

We say that α is valid in CTL* iff for all Kripke structures M, for all fullpaths b in M, we have $M, b \models \alpha$. Let us write $\models_C \alpha$ in that case. The C in $\models_C$ can stand for CTL* or, as we shall see, for complete.

We say that α is satisfiable in CTL* iff for some Kripke structure M and for some fullpath b in M, we have $M, b \models \alpha$. Clearly α is satisfiable in a Kripke structure iff $\not\models_C \neg\alpha$.

Some presentations of CTL* proceed via the definition of a certain subset of the formulas which only depend, for their truth, on an evaluation point rather than fullpath. We can make some use of these formulas. Call a formula a *state* formula if it is a boolean combination of atoms and formulas of the form $E\beta$. It is easy to show that

LEMMA 1. *If α is a state formula and b and b' are fullpaths with $b_0 = b'_0$ then $M, b \models \alpha$ iff $M, b' \models \alpha$.*

In the case of α being a state formula we can thus write $M, x \models \alpha$ to mean that for some, or equivalently all, fullpaths b with $b_0 = x$, we have $M, b \models \alpha$.

§4. Other semantics. There are other semantics for the formulas of CTL*. Several are worth introducing as they are used in the proof or cast light on the issues. In the end we will see that there really are only two distinct notions of validity of interest to us here.

4.1. Path structures. One of the most general semantics is on what I will call path frames. A pair (S, Π) is a *path frame* if S is a set (of states) and Π is any set of paths, i.e. of ω-sequences from S. Three closure properties are often assumed of the set Π of paths. We will assume two: Suffix Closure (SC), i.e. that Π is closed under taking suffixes of paths; and Fusion Closure (FC), i.e. that the beginning of one path (in Π) can be joined at a common state to the tail of another path (in Π) and the result is in Π. A *path structure* is (S, Π, g) where (S, Π) is a path frame and $g: S \to \wp(\mathscr{L})$ is a labelling. We will say that the structure is an $SC + FC$ path structure iff the set of paths Π is suffix and fusion closed.

We can give the formulas of CTL* a new semantics by defining truth of formulas on paths from $SC + FC$ path structures in the obvious way: temporal connectives are evaluated along paths while E allows switching to the same state on another path containing the current state. We define $SC + FC$ path validity by saying $\models_B \alpha$ iff $(S, \Pi, g), \pi \models \alpha$ for all $SC + FC$ path structures (S, Π, g) and all paths $\pi \in \Pi$. We will see later that the B in $\models_B$ stands for bundle.

4.2. ***R*-generable validity.** Various computing concerns to do with applications for CTL* (including the desire to reason explicitly about fairness constraints) motivate us to restrict our attention further to certain classes of path frames: in particular, we can require the other closure property of interest, Limit Closure (LC), i.e. that if for any n, a path π agrees up to its nth state with a path in Π then π itself is in Π.

The conjunction of all three closure properties is interesting because, as shown by Emerson in [Emerson, 1983], this is equivalent to Π being the set of all fullpaths in some Kripke frame (S, R). In the case that Π is the set of all paths, we say that it is R-generable. We thus can talk of R-generable validity of CTL* formulas: i.e. α is R-generable valid iff, for all path frames (S, Π) in which Π is R-generable, for all labellings g, for all $\pi \in \Pi$, we have $(S, \Pi, g), \pi \models \alpha$.

4.3. Trees. It has been observed, for example in [Emerson and Sistla, 1984], that it is sometimes useful to consider satisfiability of CTL* formulas in special tree-like structures. Indeed, we find it so in this paper. Let us define an ω-tree (frame) to be a pair $(T, <)$ where $<$ is transitive and irreflexive, for each $t \in T$, the past $\{s \in T | s < t\}$ of t is linearly ordered by $<$, there is a $<$-smallest element and each maximal linearly $<$-ordered subset of T is order-isomorphic to the natural numbers. In an ω-tree each point t will have a non-empty set N_t of immediate successors, and the future $\{s | t < s\}$ is the disjoint union of N_t and the futures of each of the elements of N_t. A *branch* of an ω-tree frame is an ω-sequence $\langle t_0, t_1, \dots \rangle$ such that each t_{i+1} is an immediate successor of t_i.

4.4. Bundled tree validity. A set B of branches on an ω-tree frame $(T, <)$ is a *bundle* iff every point $t \in T$ lies on at least one branch in B and the set B is suffix closed and also closed under superbranches, i.e. if b is a branch of $(T, <)$ and for some n, $b_{\geq n} \in B$ then $b \in B$. Say that $(T, <, B)$ is a *bundled* ω-tree frame. We give the formulas of CTL* the bundled semantics on *bundled* ω-tree structures $(T, <, B, g)$. Truth is defined recursively at branches $\pi \in B$ in a straightforward way with the clauses for atoms using the labelling g at the initial point of the branch and the temporal connectives directed along the branch. The clause for E is as follows:

$$(T, <, B, g), \pi \models E\alpha \quad \text{iff} \quad \text{there is some } \pi' \in B \text{ such that } \pi_0 = \pi'_0 \text{ and } (T, <, B, g), \pi' \models \alpha$$

We thus have a notion of *bundled* ω-tree validity.

4.5. Complete tree validity. Let $B(T, <)$ be the set of all branches of the ω-tree $(T, <)$. In case that the bundle B is just $B(T, <)$ then we talk of complete ω-tree frames and complete ω-tree validity, or sometimes just ω-tree validity. Write $(T, <, g), \pi \models \alpha$ for $(T, <, B(T, <), g), \pi \models \alpha$.

4.6. Ockhamist frames. In our completeness proof we will make use of yet another semantics.

DEFINITION 3. A (floored) *Ockhamist* frame (of countable height) is $(T, <, \equiv)$ where:

1) T is the set of points;
2) $<$ is transitive, anti-symmetric, irreflexive order satisfying
$\forall xyz(x < y \wedge x < z \rightarrow (y < z \vee y = z \vee z < y))$ and

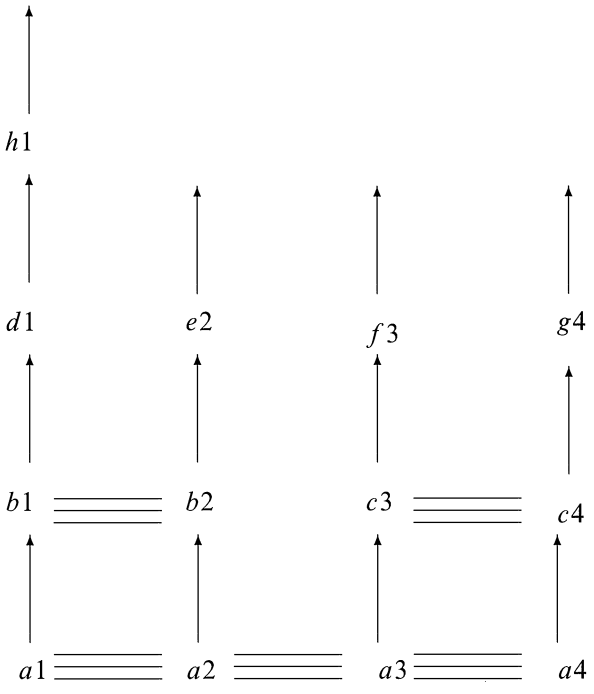

FIGURE 1. An Ockhamist frame

$\forall xyz(y < x \wedge z < x \rightarrow (y < z \vee y = z \vee z < y))$
(that is, $<$ is a strict linear order which may have parallel lines but may not have branching);

3) for each $x \in T$, $\{y | y < x\}$ is finite;
4) $\equiv$ is an equivalence relation such that:
   if $x \equiv y$ then we do not have $x < y$
   if $x \equiv y$ and $u < x$ then there is $v < y$ such that $u \equiv v$; and
5) there is an element $0 \in T$, such that for each $t \in T$,
   there is $t' \in T$ such that $0 \equiv t'$ and either $t' < t$ or $t' = t$
   ($0/\equiv$ is known as the *floor*).

Our use of the word *Ockhamist* here derives from its use in a more general (not necessarily countable) setting in [Zanardo, 1996]. Such Ockhamist frames are closely related to the Kamp frames seen in [Thomason, 1984]. In this paper we will later find it useful to think of the points in T as being arranged in an imperfect two-dimensional grid with $<$ increasing vertically and $\equiv$ relating some of the points on each horizontal level. Figure 1 portrays a very simple example. As we will see below, when we consider the semantics of CTL* formulas on Ockhamist frames, a state in a Kripke or path based structure corresponds to a whole $\equiv$-class of Ockhamist points and a path corresponds to a vertical line of points.

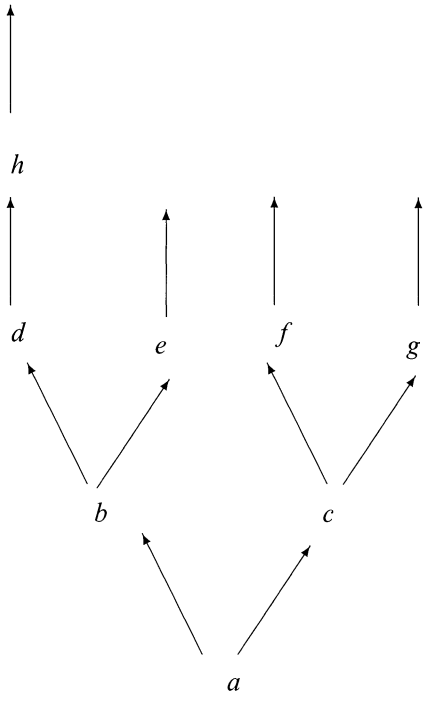

FIGURE 2. The corresponding tree

Despite using an Ockhamist frame in our completeness proof construction we do not actually use any notion of Ockhamist validity or even of truth of CTL* formulas on Ockhamist frames. However, both notions can be developed and it is worth doing so to aid with understanding the proof. First we need to impose extra restrictions on the frames and labellings of points.

The vertical lines of points need to be isomorphic to the whole natural numbers. Say that the Ockhamist frame $(T, <, \equiv)$ is an $(\mathbb{N} \times W)$-frame iff

(1) there is some set W such that $T = \mathbb{N} \times W$ and
(2) the order $<$ is defined by $(n, u) < (m, v)$ iff $n < m$ and $u = v$.

We also need to require that labels of all points in any $\equiv$-class agree on the atoms. Say that the structure $(T, <, \equiv, g)$ is an $(\mathbb{N} \times W)$-*structure* iff $(T, <, \equiv)$ is an $(\mathbb{N} \times W)$-frame and for all $n \in \mathbb{N}$, for all $u, v \in W$, if $(n, u) \equiv (n, v)$ then $g(n, u) = g(n, v)$.

Truth in Ockhamist structures is a more traditional modal logic concept as formulas are evaluated at points (rather than at states on paths). We define truth by having the temporal connectives operate vertically upwards and E allow a horizontal move within an $\equiv$-class. For example,

$$\begin{array}{lll} (T, <, \equiv, g), (n, u) \models X\alpha & \text{iff} & (T, <, \equiv, g), (n+1, u) \models \alpha, \text{ and} \\ (T, <, \equiv, g), (n, u) \models E\alpha & \text{iff} & \text{there is some } v \in W \text{ with } (n, u) \equiv (n, v) \text{ and} \\ & & (T, <, \equiv, g), (n, v) \models \alpha. \end{array}$$

Corresponding to each $(\mathbb{N} \times W)$-structure is a bundled ω-tree structure. We simply take the nodes of the tree to be the equivalence classes, define successors using $<$ and let the bundle contain each branch of the form $\boldsymbol{up}(n, w) = \{(m, w)/\equiv | n \leq m\}$ for each $(n, w) \in T$. Figure 2 shows the tree corresponding to the frame in figure 1; the bundle contains all suffixes of the four maximal branches shown. In particular note that a point in an Ockhamist frame corresponds to the combined notion of a state on a path. This correspondence preserves truth of CTL* formulas in the obvious way.

With such an Ockhamist semantics we can define a notion of what might be called $(\mathbb{N} \times W)$-*Ockhamist validity* on $(\mathbb{N} \times W)$-structures. In lemma 2 below we show (using the correspondence) that this notion of validity is just bundled validity.

However, to find an Ockhamist version of the notion of validity identified in lemma 3 below requires a complicated extra restriction of the frame in a similar vein to the limit closure property. We omit this as we do not use either of these definitions of validity anyway.

4.7. Equivalence results. There are two equivalence results which are useful for relating some of the semantic approaches to CTL*. The proofs are straightforward.

LEMMA 2. *The following are equivalent*:

a) $\models_B$ *validity i.e. SC + FC path frame validity*;
b) *bundled* ω*-tree validity*;
c) $(\mathbb{N} \times W)$-*Ockhamist validity.*

The following equivalence result is of vital importance to the proof of soundness of our axiom system for $\models_C$. It follows from results in [Emerson and Sistla, 1984].

LEMMA 3. *The following are equivalent*:

a) $\models_C$ *validity i.e. Kripke validity*
b) *suffix, fusion and limit closed path frame validity*
c) *R-generable path frame validity*
d) *complete* ω*-tree validity*

Note that the equivalence between a) and d) follows from a particularly close relationship between Kripke frames and ω-trees: a tree is just a Kripke frame with the accessibility relation defined in terms of immediate successors while a tree can be produced by unraveling the fullpaths on a given Kripke frame.

4.8. An inequivalence result. We now show that $\models_B$ and $\models_C$ are distinct notions of validity.

Note that for complete ω-trees we have $\models_C \gamma$ where $\gamma = A\,G(p \rightarrow E\,X\,p) \rightarrow (p \rightarrow E\,G\,p)$.

However, this is not valid on bundled trees. To see this, in a slightly indirect way, consider the $(\mathbb{N} \times W)$-frame shown in figure 3 in which $W = \mathbb{N}$ and $(n, u) \equiv (m, v)$ iff either $(n, u) = (m, v)$ or $n = m$ and both $n \leq u$ and $n \leq v$. Suppose that $p \in g(n, u)$ iff $n \leq u$. It is clear that we have $(T, <, \equiv, g), (0, 0) \models A\,G(p \rightarrow E\,X\,p) \wedge p \wedge A\,F\,\neg p$ and so $\not\models_B \gamma$.

It is very interesting to also consider the bundled ω-tree corresponding to this Ockhamist structure. Take $T = \{(n, m) \in \mathbb{N} \times \mathbb{N} | n \leq m\}$. Put $(a, b) < (c, d)$ iff$(a = c$ and $b < d)$ or $a = b < c$. Let $D = \{(n, n) | n \in \mathbb{N}\} \in B(T, <)$ and define

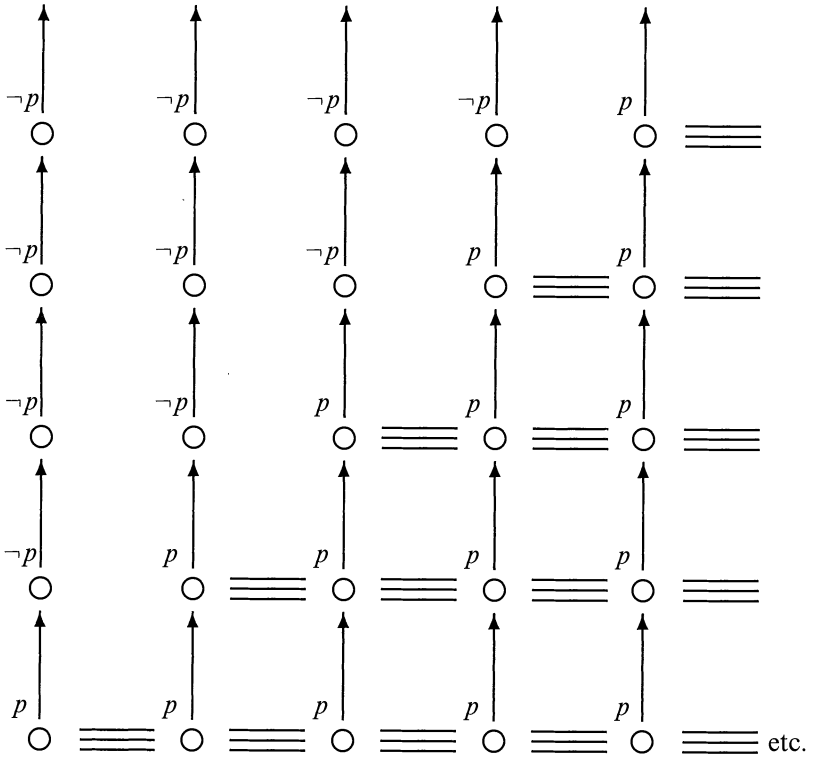

FIGURE 3. The Ockhamist Counter-example

the bundle to be $B = B(T, <) \setminus \{D\}$. Finally define g so that $(n, m) \in g(p)$ iff $n = m$. See figure 4. Of course we have $(T, <, B, g), (0, 0) \models \neg\gamma$ as well.

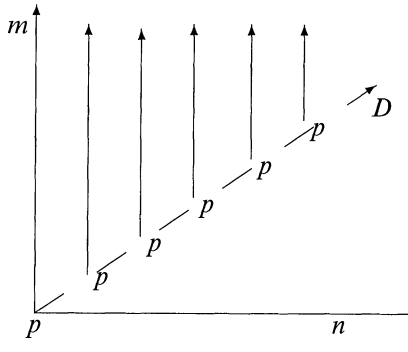

FIGURE 4. The Counter-example, tree-style

4.9. Emergent branches. As well as establishing the inequivalence of $\models_B$ and $\models_C$, this example also gives a good illustration of the idea of an emergent branch appearing from an Ockhamist construction. In particular note that the branch

D in the ω-tree does not correspond to any set $\boldsymbol{up}(n, w)$ in the Ockhamist frame. It is not in the bundle defined in the correspondence. In general converting an Ockhamist frame, even with a countable set W, to an ω-tree can produce a possibly uncountable number of such emergent branches. This fact will cause us a major difficulty in the completeness proof: we can make sure all the constructed branches satisfy a desired property at all finite stages of the construction of an Ockhamist frame but then discover a large number of extra branches emerge when we take the limit. We will need extra machinery to make sure that these emergent branches satisfy the desired property.

In terms of Kripke structures the fact(problem) of emergent branches is seen here by viewing the ω-tree $(T, <)$ as a Kripke structure (T, R, g) where $(n, m)R(n, m+1)$ and $(n, n)R(n+1, n+1)$. This is just the Kripke structure we get by defining R in terms of immediate $<$-successors as we need to do in lemma 3. It is clear that we have a fullpath $\langle (0,0), (1,1), (2,2), \ldots \rangle$ which is not in the bundle which is defined by the Ockhamist frame. The fullpath is emergent.

§5. The axiom system for bundled trees. Axioms for bundled validity will form part of our final axiom system. Consider a system for bundled validity. The inference rules are modus ponens and temporal and path generalization:

$$\frac{\alpha, \alpha \to \beta}{\beta} \qquad \frac{\alpha}{G\alpha} \qquad \frac{\alpha}{A\,\alpha}.$$

The axiom schemes include the usual ones for PLTL (as seen earlier): plus axioms ensuring that the path modality A behaves as in the modal logic S5

C9 $A(\alpha \to \beta) \to (A\,\alpha \to A\,\beta)$
C10 $A\,\alpha \to A\,A\,\alpha$
C11 $A\,\alpha \to \alpha$
C12 $\alpha \to A\,E\,\alpha$
C13 $A \neg\alpha \leftrightarrow \neg E\,\alpha$
plus propositional atoms only depend on states
C14 $p \to A\,p$, for each atomic proposition p
plus some interaction between modalities
C15 $A\,X\,\alpha \to X\,A\,\alpha$

We define derivability $\vdash_B$ using this system in the usual way.

Note that we do not use a substitution rule as it is not valid for $\models_B$. For example $\models_B (E\,p \to p)$ is valid for each atomic proposition p (and can be derived easily) but $\models_B (E\,\alpha \to \alpha)$ is not generally valid.

We could show that

LEMMA 4 ([Stirling, 1992]). $\vdash_B$ *is sound and* (*weakly*) *complete for* $\models_B$.

Soundness is the usual straightforward induction on the lengths of proofs. Completeness (for an equivalent but slightly different set of axioms) is shown in [Stirling, 1992] by building a SC+FC path structure. In what is actually quite a similar technique, despite appearances, one could also use a step by step method via Ockhamist frames to produce a bundled ω-tree structure. This style of proof is seen in proofs by Burgess, Zanardo and von Kutschera for logics of historical necessity. In fact, this method forms the basis for our $\models_C$ completeness proof later in the paper.

By removing most of the machinery associated with the automaton and the banning procedure, we would be left with a completeness proof for $\vdash_B$.

Note that the example γ from the previous section shows that the system for $\vdash_B$ is incomplete for $\models_C$. Of course, this system is sound for $\models_C$. In the next two sections we introduce a new axiom and new rules of inference to allow us to derive the extra validities.

§6. The limit closure axiom. Intuitively, the new limit closure (LC) axiom schema captures a particular form of limit closure. Suppose that some state formula $E\alpha$, say, is always (at all points on all branches) at the start of a finite path of $E\beta$ states leading up to a new $E\alpha$ state. Also suppose that $E\alpha$ holds at some node. Then it is clear that from that node extends a branch on which $(E\beta)\,U(E\alpha)$ always holds.

LC: $\quad A\,G(E\alpha \to EX((E\beta)\,U(E\alpha))) \to (E\alpha \to EG((E\beta)\,U(E\alpha)))$

LEMMA 5. *The LC axiom schema is sound for* $\models_C$.

PROOF. Suppose that $(T, <, V)$ is a (complete) ω-tree structure with a node $t_0 \in T$ and

$$(T, <, V), t_0 \models A\,G(E\alpha \to EX((E\beta)\,U(E\alpha))) \wedge E\alpha.$$

We are to show that $(T, <, V), t_0 \models EG((E\beta)\,U(E\alpha))$.

We will find a sequence of nodes $t_0 < t_1 < \ldots$ from T such that

- $(T, <, V), t_k \models E\alpha$ and
- for each s such that $t_k < s < t_{k+1}$, we have $(T, <, V), s \models E\beta$.

We already have t_0 chosen. Suppose for the induction we have chosen an appropriate t_k with $(T, <, V), t_k \models E\alpha$. By assumption

$$(T, <, V), t_k \models EX((E\beta)\,U(E\alpha)).$$

This gives us some $t_{k+1} > t_k$ along some branch starting at t_k, with $(T, <, V), t_{k+1} \models E\alpha$ and $(T, <, V), s \models E\beta$ for each s such that $t_k < s < t_{k+1}$.

Let π be the branch of $(T, <)$ which starts at t_0 and includes $t_0 < t_1 < \ldots$.

By construction we have

$$(T, <, V), \pi \models G((E\beta)\,U(E\alpha)) :$$

for all $n < \omega$, either $\pi_n = t_0$ when we put $k = 0$ or there is some $k > 0$ such that $t_{k-1} < \pi_n \leq t_k$ and so we have $(T, <, V), t_k \models E\alpha$ and at any s with $\pi_n \leq s < t_k$ we have $(T, <, V), s \models E\beta$. It follows immediately that $(T, <, V), t_0 \models EG((E\beta)\,U(E\alpha))$.

So the LC axiom is valid on ω-tree structures. By lemma 3, it is valid for $\models_C$. □

§7. The auxiliary atoms rule. The Auxiliary Atoms (AA) rule allows the use of a certain arrangement of fresh atoms in a proof. Suppose that L and Q are disjoint sets of atoms. Suppose that ψ, which only uses atoms from L, is the formula which we are interested in deriving. The rule will involve another formula θ, using atoms from $L \cup Q$. The formula θ describes the arrangement of the fresh atoms, ie those in Q, in terms of those in L. The AA rule will allow $\vdash \psi$ to be derived from $\vdash (\theta \to \psi)$ under certain side-conditions.

As we will see in example derivations below in section 16, the AA rule can be used in conjunction with the LC schema to derive a formula stating the existence of a fullpath satisfying a certain desired property. The LC schema allows us to deduce the existence of the infinite fullpath from the existence of a finite cycle of states. The fresh atoms can be used as markers placed along the cycle for various reasons. For example they can be used to record that various states are visited in the cycle. They can also be used to distinguish several visits to one state within the cycle, indicating that at one visit a certain successor state should be chosen and on another visit a different successor state should be chosen. For the purposes of theorem-proving, deriving ψ say, the AA rule will be used as follows. First show that ψ is true when the fresh atoms are arranged according to θ, i.e. show $\vdash (\theta \rightarrow \psi)$ and then use the AA rule to immediate effect.

Of course, the rule is also "used" in our completeness proof. We use the fresh atoms to record the states of a certain linear time automaton as it trundles along the branches of a tree. We eventually want to check whether the branches are accepted or not to see whether they satisfy a certain temporal formula. In the completeness proof, we will have assumed that a formula $\neg\psi$, say, is consistent and we will be wanting to build a model of $\neg\psi$. Using the AA rule (in a contrapositive way) we will immediately know that $\theta \wedge \neg\psi$ is also consistent. The specific formula θ which we will use here will describe an arrangement of the atoms to record the running of our automaton. So $\neg\psi$ being consistent implies a description of a model of $\neg\psi$ with the automaton's state recorded in fresh atoms is also consistent.

Note that there is no necessary connection between the AA rule and automata. (In fact, the acceptance criteria of automata has no counterpart in the AA rule, and so we should really only be discussing finite state machines here). The AA rule is just a rule which allows fresh atoms to be introduced into a proof in accordance with a fairly prescriptive arrangement. However, in our completeness proof we are going to use the rule to allow our construction to make use of an automaton. Further, as we will see in some example derivations later, uses of the AA rule in theorem-proving may well be motivated by consideration of certain automata. Finally, we will see that an understanding of automata will help with an intuitive understanding of the rule.

Obviously the AA rule must have a substantial side-condition: it can not always be sound to derive $\vdash \psi$ from $\vdash (\theta \rightarrow \psi)$. In terms of consistency/satisfiability, we must state sufficient conditions on θ under which θ can be added as a conjunct to any formula $\neg\psi$ with satisfiability preserved.

In fact, a very general class of such formulas θ exists. It is easy to see that we just need to require that, given any ω-tree structure in the language of L, and any point, we can choose the valuations of the atoms in Q so that the state formula θ holds at that point of the expanded structure. We might call such formulas $L + Q$-expandable. However, this is a semantic condition on θ and we instead find a syntactic condition which is sufficient to ensure that θ is in this class.

The actual syntactic side-condition we use for our purposes is slightly complicated. It says that θ is a formula which prescribes exactly which atom from Q to make true at each point in any ω-tree structure in the language of L by working up from the root along each branch. We require that θ makes use of a finite pairwise-inconsistent set $\{\alpha_1, \ldots, \alpha_n\}$ of state formulas in the language of L and a function

ρ to determine which atom to make true at the next point along a branch. We will say that such formula θ is functionally $L + Q$-expandable.

DEFINITION 4. Say that $\{\alpha_1, \ldots, \alpha_n\}$ is a set of formulas using atoms only from L such that

- each α_i is a state formula, i.e. a boolean combination of atoms and formulas of the form $E\beta$,
- for each $i \neq j$, $(\alpha_i \wedge \alpha_j) \rightarrow \textbf{\textit{false}}$ is a substitution instance of a classical tautology and
- $\bigvee_{i=1}^{n} \alpha_i$ is a substitution instance of a classical tautology.

Suppose that Q is finite (and enumerated in some way), there is $b_0 \in Q$ and a function $\rho : (Q \times \{1, \ldots, n\}) \rightarrow Q$ such that

$$\theta = b_0 \wedge A\,G \bigwedge_{b \in Q} \bigwedge_{i=1}^{n} ((b \wedge \alpha_i) \rightarrow A\,X\,\rho(b, i)) \wedge A\,G \bigwedge_{b \neq b' \in Q} \neg(b \wedge b').$$

Then we say that θ is *functionally $L + Q$-expandable*.

We state the AA rule using this concept as follows:

$$\frac{\vdash \theta \rightarrow \psi}{\vdash \psi}$$

AA: provided there are disjoint sets L and Q of atoms such that:

- ψ only uses atoms from L
- θ is functionally $L + Q$-expandable.

LEMMA 6. *The AA rule is sound for* $\models_C$.

PROOF. We prove this by contraposition: suppose that $\neg\psi$ is satisfiable, i.e. has a complete ω-tree model in the language of L. Say that $\neg\psi$ is true of the branch π.

But now a simple recursion on the length of finite prefixes of branches starting at π_0, defines an interpretation of the Q atoms to make θ true of π.

We have a complete ω-tree model of $\theta \wedge \neg\psi$ and so, via lemma 3, have our result. □

§8. The theorem. Now we have all the axioms and rules of our proof system for standard CTL* validity.

DEFINITION 5. Let $\vdash_C$ be derivability in the system which results from adding the limit closure axiom schema and the auxiliary atoms rule to the system $\vdash_B$ for bundled trees.

THEOREM 3. $\vdash_C$ *is sound and (weakly) complete for standard CTL* validity,* $\models_C$.

PROOF. Soundness follows from the usual induction on the lengths of proofs. We can show validity of the axioms including LC in lemma 5 and, from lemma 6, we know that the rules preserve validity. Most of the rest of the paper contains the completeness proof. We will show that any given consistent formula has a Kripke model.

Note that, as with PLTL and $\models_B$, the logic is not compact and we can thus only manage a weak completeness result. □

§9. Completeness preliminaries. In what follows we have fixed a CTL* formula ϕ which is supposed to be $\vdash_C$-consistent, i.e. $\nvdash_C \neg\phi$. We are to show that it has a Kripke model. We use $\vdash$ for $\vdash_C$ throughout the proof. Let L be the set of atoms appearing in ϕ.

In this section we outline the quite intricate completeness proof and motivate the need to use three languages: the CTL* language of ϕ with atoms from L, a PLTL language and an extended CTL* language with some fresh atoms.

We will use a step-by-step construction to build an Ockhamist frame with each point t labelled by a set of formulas $\lambda(t)$. After taking the limit $(T, <, \equiv)$ of this construction we will factor out by $\equiv$ and be left with a Kripke structure (S, R) with R defined in terms of $\equiv$-classes of $<$-successor points. By making sure that λ does not vary on atoms within $\equiv$-classes we will be able to use it to define a labelling g on S and we will be able to show, via a truth lemma relating labels to true formulas, that (S, R, g) is a model of ϕ as required.

The first of many complications is that we deal with only a finite set of formulas of interest. This is because we need to be fair in our scheduling of labels as we make choices in our construction. We do this for exactly the same reason that it is done in the proof of completeness for the PLTL system: we need to satisfy eventualities such as $\alpha\, U\, \beta$. There are subtleties here and it is important for our proof to be familiar with the (PLTL) idea of working with finite labels instead of the usual infinite maximal consistent sets of formulas. In particular, note that as we proceed in the construction we can make choices about which label to place on the new successor to an existing labelled point.

So, for some purposes including proving the truth lemma, we restrict our attention to the finite set of subformulas of ϕ and their negations. The simplest of the several closure sets of formulas used in the proof is thus

$$\boldsymbol{cl}\,\phi = \{\psi, \neg\psi | \psi \leq \phi\}$$

recalling that $\psi \leq \phi$ denotes that ψ is a subformula of ϕ. This set is not quite closed under taking negations of its formulas. However, in the proof we can use ψ as the negation of $\neg\psi$.

A straightforward step-by-step construction using such labels and fair scheduling will eventually produce an Ockhamist model of ϕ. The maximal sets of $<$-related points will each give us a branch. Unfortunately, this construction will not in general give us a Kripke model but only a bundled model. This is because, in taking the limit of the construction, a possibly uncountable number of new branches emerge in addition to the ones which were constructed as maximal sets of $<$-related points. (See the example in section 4.9 above.)

The problem with emergent branches in our construction can be briefly described as follows. When we add a new labelled point, most importantly a successor to an existing point, to our finite Ockhamist frame, we need to decide on its label. Obviously for our truth lemma to work there are strict requirements on this new label. One of the most important is that if $\neg E\,\alpha$ is in the label then we also need $\neg\alpha$ in the label. This will eventually ensure that the fullpath built from the equivalence

class of that point and the points $<$-above does not make α true. However, it may be that an emergent fullpath beginning at this class does make α true. In that case we will not after all have $\neg E\alpha$ holding here as required by the truth lemma.

We might call emergent fullpaths *bad* if they mess up the truth lemma in this way. Our task, then, is to prevent the construction of bad emergent fullpaths.

The solution we employ is as follows. We first note that we can tell which formulas will hold along a fullpath (even an emergent one) by looking at the ω-sequence of sets of formulas of the form $E\beta$ appearing in the labels of points along the fullpath. Due to the fact that labels of points in $\equiv$-classes agree on such formulas we do not even have be careful of which representative we look at. So now to ensure that no fullpaths are bad we just need to ensure the truth, along each fullpath, of a certain PLTL formula which has atoms corresponding to the $E\beta$ formulas. Hence the next few definitions.

Define $\boldsymbol{ecl}\,\phi = \{E\psi, E\neg\psi | \psi \leq \phi\}$ and for each $c \subseteq \boldsymbol{ecl}\,\phi$ let

$$\zeta_c = \bigwedge_{E\psi \in c} E\psi \wedge \bigwedge_{E\psi \in (\boldsymbol{ecl}\,\phi \setminus c)} \neg E\psi.$$

If α is a subformula or the negation of a subformula of a formula in $\boldsymbol{ecl}\,\phi$, then we define a formula $\overline{\alpha}$ of the linear language with atoms from $\overline{\boldsymbol{ecl}\,\phi} = \{\overline{E\beta} | E\beta \in \boldsymbol{ecl}(\phi)\}$. Simply put $\overline{\boldsymbol{true}} = \boldsymbol{true}$, $\overline{p} = \overline{E\,p}$ for atoms $p \in L$, $\overline{\neg\alpha} = \neg\overline{\alpha}$, $\overline{\alpha \wedge \beta} = \overline{\alpha} \wedge \overline{\beta}$, $\overline{\alpha\, U\, \beta} = \overline{\alpha}\, U\, \overline{\beta}$, $\overline{X\alpha} = X\overline{\alpha}$ and $\overline{E\beta}$ is just the atom $\overline{E\beta}$ itself. For any $c \subseteq \boldsymbol{ecl}\,\phi$, let $\overline{c} = \{\overline{\beta} | \beta \in c\}$.

Now we are able to introduce the PLTL formula which must be satisfied by every fullpath in the limit model. Define

$$\chi_0 = G \bigwedge_{\gamma \in \boldsymbol{cl}(\phi)} (\overline{\gamma} \rightarrow \overline{E\gamma}),$$

a PLTL formula in the language with atoms from $\overline{\boldsymbol{ecl}(\phi)}$.

For any given fullpath (emergent or otherwise) we can interpret the atom $\overline{E\beta}$ in $\overline{\boldsymbol{ecl}(\phi)}$ in accordance with the appearance of $E\beta$ in the label of a point along the fullpath. We must make sure that on no fullpath the ω-sequence so defined is a model of $\neg\chi_0$. Such a fullpath would be a bad one.

The task is to ensure this by making the right choices (of labels) during the finite construction. We need some way of constructing many interconnected PLTL models (of χ_0) in a step-by-step way. To help us do this we now bring in the promised automaton to tell us how well we are approaching the goal (of a model of χ_0) on each fullpath. The simple form of the Rabin acceptance criteria is the key to guiding our construction. Essentially we can tell how we are going by looking at the history of the construction so far and we can thus assess our progress on emergent and constructed fullpaths alike.

So by theorem 2, find a deterministic Rabin automaton $A = (Q, s_0, \rho, \{(U_1, V_1), \ldots, (U_K, V_K)\})$ recognizing ω-sequences of subsets of $\overline{\boldsymbol{ecl}(\phi)}$ and accepting exactly the models of $\neg\chi_0$. Thus $\rho : Q \times (\wp(\overline{\boldsymbol{ecl}\,\phi})) \rightarrow Q$ and each $U_i, V_i \subseteq Q$. We can choose Q so that $Q \cap \overline{\boldsymbol{ecl}(\phi)} = \emptyset$ and, as mentioned in section 2, we can choose the (U_i, V_i) so that each $U_i \cap V_i = \emptyset$.

Given the automaton our task can now be restated as the requirement to proceed so that in the limit no fullpaths will be accepted by A. So somehow we have to make sure that as A runs along any fullpath (and interprets the labels appropriately) there is no $i = 1, \ldots, K$ such that a state in U_i comes up infinitely often while no state in V_i does.

The next step is to realize that in achieving this goal during our construction we do need to look ahead a little to make sure that we do not force ourselves to end up with a bad path by making a bad choice (of label) at a finite stage. In order to make sure that we can consistently continue after making a particular choice we need to bring the whole automata machinery (and a banning mechanism which we will meet later) inside the language, as it were. That is, we need to be able to reason about the progress of the automaton within the language of the labels. This is a crucial observation and motivates the need for the AA rule. In particular, we need to bring in fresh atoms to record the current state of A (if it started at the floor of our frame and worked up to any given point).

So now we introduce our third language: a CTL* language with atoms from L and some new ones as well. We use the symbols of Q as the new atoms and define the following branching formula using atoms from $Q \cup L$:

$$\theta_A = s_0 \wedge A\,G \bigwedge_{s \in Q, c \subseteq \boldsymbol{ecl}\,\phi} ((s \wedge \zeta_c) \rightarrow A\,X\,\rho(s, \overline{c})) \wedge A\,G \bigwedge_{s \neq r \in Q} \neg(r \wedge s).$$

This formula, which we want to hold at the root, will ensure that the atoms in Q are interpreted exactly according to the state which A would be in at that point if it trundled along reading the PLTL version of the labels as we desire it to. Note that for each $c \subseteq \boldsymbol{ecl}(\phi)$, the truth of the CTL* state formula ζ_c (or its appearance in a label) is exactly what we mean by saying that the automaton is reading the current set of $\overline{\boldsymbol{ecl}(\phi)}$ atoms as being $\overline{c}$.

In order to allow us to reason ahead about whether we are in danger of creating a bad fullpath we also define $\nu_i = \bigvee_{s \in V_i} s$, and $\mu_i = \bigvee_{s \in U_i} s$ and let $\chi_1 = \bigvee_{i=1}^{K} (F\,G \neg \nu_i \wedge G\,F\,\mu_i)$. Thus, given the described setting, χ_1 holds at a fullpath b iff A accepts b iff b is bad.

Now we have a language which allows us to specify a model of ϕ in which we have no fullpaths accepted by A. We simply use the formula $\phi^+ = \theta_A \wedge \phi \wedge \neg E\,\chi_1$ in the CTL* language with the fresh atoms.

Fortunately, the AA rule and a little PLTL reasoning allows us to conclude that ϕ^+ is consistent.

LEMMA 7. *If ϕ is consistent then so is $\phi^+ = \theta_A \wedge \phi \wedge \neg E\,\chi_1$.*

PROOF. The proof proceeds via three claims.

CLAIM 1. In PLTL we have the following:

$$\models_L (\chi_1 \wedge s_0 \wedge G \bigwedge_{s \in Q, c \subseteq \boldsymbol{ecl}\,\phi} (s \wedge \overline{\zeta_c} \rightarrow X\,\rho(s, \overline{c})) \wedge G \bigwedge_{r \neq s \in Q} \neg(r \wedge s)) \rightarrow \neg\chi_0.$$

PROOF. Suppose that σ is an ω-sequence of subsets of $\overline{\boldsymbol{ecl}\,\phi} \cup Q$ such that

$$\sigma \models \chi_1 \wedge s_0 \wedge G \bigwedge_{s \in Q, c \subseteq \boldsymbol{ecl}\,\phi} (s \wedge \overline{\zeta_c} \rightarrow X\,\rho(s, \overline{c})) \wedge G \bigwedge_{r \neq s \in Q} \neg(r \wedge s).$$

Say that the run of A on $\sigma|_{\overline{\boldsymbol{ecl}\,\phi}}$ is $\langle s_0, s_1, \ldots\rangle$ so that $s_{i+1} = \rho(s_i, \sigma_i|_{\overline{\boldsymbol{ecl}\,\phi}})$.

I claim that for all $i < \omega$, for all $s \in Q$, we have $s \in \sigma_i$ iff $s = s_i$. We prove this by induction on i. After we prove the converse direction, the forward direction follows from the fact that $\sigma \models G \bigwedge_{r \neq s \in Q} \neg(r \wedge s)$. The case of $i = 0$ follows as $\sigma \models s_0$. Now assume that the hypothesis holds for $i \geq 0$: we show the converse direction for $i + 1$. Thus we are to show that $s_{i+1} \in \sigma_{i+1}$.

We know that $\sigma_{\geq i} \models \bigwedge_{s \in Q, c \subseteq \boldsymbol{ecl}\,\phi}(s \wedge \overline{\zeta_c} \rightarrow X\,\rho(s, \overline{c}))$ and $s_i \in \sigma_i$. Thus $\sigma_{\geq i} \models s_i \wedge \overline{\zeta_c} \rightarrow X\,\rho(s_i, \overline{c})$ where $c = \{\beta \in \boldsymbol{ecl}\,\phi | \overline{\beta} \in \sigma_i\}$ so that $\overline{c} = \sigma_i|_{\overline{\boldsymbol{ecl}\,\phi}}$.

Now for each $\psi \in \boldsymbol{ecl}\,\phi$, if $E\,\psi \in c$ then $\overline{E\,\psi} \in \sigma_i$ so $\sigma_{\geq i} \models \overline{E\,\psi}$. Similarly, if $E\,\psi \in \boldsymbol{ecl}\,\phi \setminus c$ then $\sigma_{\geq i} \models \neg\overline{E\,\psi}$. Thus $\sigma_{\geq i} \models \overline{\zeta_c}$.

Since $\sigma_{\geq i} \models s_i$ we can conclude that $\sigma_{\geq i} \models X\,\rho(s_i, \overline{c})$. But $\rho(s_i, \overline{c}) = \rho(s_i, \sigma_i|_{\overline{\boldsymbol{ecl}\,\phi}})$ $= s_{i+1}$ so $\sigma_{\geq i+1} \models s_{i+1}$. We conclude that $s_{i+1} \in \sigma_{i+1}$ as required.

Thus we have shown that for all $i < \omega$, for all $s \in Q$, we have $s \in \sigma_i$ iff $s = s_i$.

Since $\sigma \models \chi_1$ there is some $j = 1, \ldots, K$ such that, no atom in V_j comes up infinitely often in the sequence σ of sets of atoms but some atom from U_j does. By the foregoing, in the run $\langle s_0, s_1, s_2, \ldots\rangle$ no state in V_j is visited infinitely often but some state in U_j is. Thus A accepts $\sigma|_{\overline{\boldsymbol{ecl}\,\phi}}$. We deduce that $\sigma \models \neg\chi_0$ as required. □

CLAIM 2. $\vdash \theta_A \rightarrow \neg E\,\chi_1$.

PROOF. C9, C11, C14 and some simple PLTL reasoning imply that

$$\vdash \theta_A \wedge E\,\chi_1 \rightarrow E(\chi_1 \wedge s_0 \wedge G \bigwedge_{s \in Q, c \subseteq \boldsymbol{ecl}\,\phi} (s \wedge \zeta_c \rightarrow X\,\rho(s, \overline{c})) \wedge G \bigwedge_{r \neq s \in Q} \neg(r \wedge s)).$$

From claim 1, we know that the PLTL axiom system can be used to show

$$\vdash_L (\chi_1 \wedge s_0 \wedge G \bigwedge_{s \in Q, c \subseteq \boldsymbol{ecl}\,\phi} (s \wedge \overline{\zeta_c} \rightarrow X\,\rho(s, \overline{c})) \wedge G \bigwedge_{r \neq s \in Q} \neg(r \wedge s)) \rightarrow (\neg G \bigwedge_{\gamma \in \boldsymbol{cl}\,\phi} (\overline{A\,\gamma} \rightarrow \overline{\gamma})).$$

Following the same proof in our axiom system using $E\,\gamma$ substituted for each atom $\overline{E\,\gamma}$ gives us

$$\vdash (\chi_1 \wedge s_0 \wedge G \bigwedge_{s \in Q, c \subseteq \boldsymbol{ecl}\,\phi} (s \wedge \zeta_c \rightarrow X\,\rho(s, \overline{c})) \wedge G \bigwedge_{r \neq s \in Q} \neg(r \wedge s)) \rightarrow (\neg G \bigwedge_{\gamma \in \boldsymbol{cl}\,\phi} (A\,\gamma \rightarrow \gamma)).$$

But C12 can be used to show that for all $\gamma \in \boldsymbol{cl}\,\phi$, $\vdash A\,\gamma \rightarrow \gamma$. So

$$\vdash G \bigwedge_{\gamma \in \boldsymbol{cl}\,\phi} (A\,\gamma \rightarrow \gamma).$$

Putting this altogether we conclude that $\theta_A \wedge E\,\chi_1$ is inconsistent as required. □

CLAIM 3. If $\vdash \theta_A \wedge \phi \rightarrow E\,\chi_1$ then $\vdash \neg\phi$.

PROOF. Propositional reasoning from claim 2 and the assumption implies that $\vdash \theta_A \rightarrow \neg\phi$.

Now we can show that θ_A is functionally $L + Q$-expandable. Say that $\wp(\boldsymbol{ecl}\,\phi) = \{c_1, \ldots, c_N\}$. We use the set $\{\zeta_{c_1}, \ldots, \zeta_{c_N}\}$ of formulas and a function $\rho' : Q \times \{1, \ldots, N\} \to Q$ given by $\rho'(s, i) = \rho(s, \overline{c_i})$. Since $\neg\phi$ uses atoms only from L, we may apply the AA rule with premise $\vdash \theta_A \to \neg\phi$.

This gives us $\vdash \neg\phi$ as required. □

Our lemma follows immediately. □

Before we continue with the completeness proof, let us look ahead to the truth lemma to see how the three languages and the automaton come together. We will define a Kripke frame (S, R) from equivalences classes in the limit $(T, <, \equiv)$ of our Ockhamist construction and on it we will define a $(L \cup Q)$-labelling g from the labels $\lambda(t)$ on the points t in the construction. A point t of the Ockhamist limit will determine the state $[t]$ which is an equivalence class of $\equiv$-related points and it will also determine a fullpath $\boldsymbol{up}(t)$ in (S, R) being the sequence of equivalence classes of points $<$-above t.

The truth lemma will establish, by induction on the construction of formulas in $\boldsymbol{cl}(\phi)$, that truth of a formula on $\boldsymbol{up}(t)$ in (S, R, g) agrees exactly with membership of the formula in $\lambda(t)$.

There are two difficult parts of the truth lemma. One is to show that if a formula of the form $\alpha\, U\, \beta$ is in $\lambda(t)$ then β eventually makes it into the label of a point above t. This involves fair scheduling argument (just as in the PLTL proof) but is somewhat complicated by our banning mechanism. More about this later.

The other difficult case is, as we have foreshadowed, showing that if $E\,\alpha$ holds of $\boldsymbol{up}(t)$ in (S, R, g) then $E\,\alpha$ is in $\lambda(t)$. So we assume that there is a fullpath b starting at $[t]$ in (S, R) on which α holds. The emergent paths cause the trouble here because the fullpath b might be emergent and thus not of the form $\boldsymbol{up}(t')$ for any t'. So we need to show $E\,\alpha$ is in $\lambda(t)$ but we do not necessary have immediate access to any label $\lambda(t')$ containing α. Note that if we were just attempting a bundled completeness proof then we would define the bundle to contain only the $\boldsymbol{up}(t')$ fullpaths and at this stage of that proof we could use the bundled semantics to give us the required t'.

Instead, we need some other guarantee that we have $E\,\alpha$ in $\lambda(t)$. This is where the automaton A and the idea of bad paths comes in. We can define an ω-sequence σ in the PLTL language of $\overline{\boldsymbol{ecl}(\phi)}$ by saying that $\overline{E\,\beta}$ is in σ_i iff $E\,\beta$ is in any (equivalently all) $\lambda(t')$ for t' in the $\equiv$-class b_i. We know that $(S, R, g), b \models \alpha$. Our truth lemma inductive hypothesis (on subformulas of α) will, in lemma 25, be able to be used to show that $\sigma \models \overline{\alpha}$. This is partly just the observation that the truth of CTL* formulas along fullpaths is determined by the linear arrangement of truth of $E\,\beta$ formulas along the fullpath. Of course, we also need the inductive hypothesis in the truth lemma to relate truth of these formulas to the contents of labels (which is how σ is defined).

From the fact that $\sigma \models \overline{\alpha}$ and that we were able to ensure that each fullpath, and in particular b is not bad, it follows that $\sigma \models \overline{E\,\alpha}$. This is because we have seen that badness can be defined in terms of being a model of $\neg\chi_0$. Note that the actual argument at this stage is slightly complicated by the need to start certain fullpaths at the floor of the structure to reason correctly about the behaviour of the automaton.

It follows immediately, by definition of σ, that $\lambda(t)$ contains $E\,\alpha$ and we are done.

§10. The standard basis of the construction. In this section we consider the standard modal logic aspects of our step-by-step construction of a labelled Ockhamist frame. We will be building a model of $\phi^+ = \theta_A \wedge \phi \wedge \neg E\chi_1$ under the assumption that it is consistent– which we have just determined happens when ϕ is consistent. The construction will be the usual procedure of laying out maximal consistent subsets of some closure set of formulas and putting accessibility relations between them. The unusual aspects will be a fair approach to scheduling the use of the sets and a banning mechanism for preventing the use of some of them.

At each stage of the step-by-step construction we will have part of an Ockhamist structure, a vaguely grid-like two-dimensional structure with points labelled by what we will call hues. Construction proceeds by adding a successor to a point. We might say that we cure a defect.

In picturing the construction we will, as usual, think of the irreflexive transitive order $<$ as increasing vertically and the equivalence relation $\equiv$ as being horizontal.

The hues which label points will indicate which formulas from a small finite closure set are to hold at that point. This closure set $\boldsymbol{fcl}\,\phi^+$ contains all the subformulas of ϕ^+ and is closed under one application of both negation and E. Recall that in an Ockhamist frame the equivalence relation $\equiv$ relates points which collapse to a state under the usual branching time semantics. Thus we need to ensure that equivalent points' labels agree on state formulas.

The hue of a point will also indicate which other hues will hold at other points which are $\equiv$-related to that point. A set of hues sported by the members of a whole $\equiv$-class of points will be called a colour. All equivalent points will share a colour and so it is not surprising that corresponding to a colour will be a state formula γ_c: the colour is a property of the whole equivalence class. This fact will ensure that we will be able to talk about the sequence of colours (but not hues) along the emergent fullpaths in the limit of our construction.

Let us formalize hues and colours.

10.1. Colours and hues. Let

$$\boldsymbol{fcl}\,\phi^+ = \{\psi, \neg\psi, E\psi, \neg E\psi, E\neg\psi, \neg E\neg\psi | \psi \leq \phi^+\}.$$

For $a \in \wp(\boldsymbol{fcl}\,\phi^+)$, define

$$\delta_a = \bigwedge_{\delta\in a} \delta \wedge \bigwedge_{\delta\in(\boldsymbol{fcl}(\phi^+)\setminus a)} \neg\delta$$

(with the conjunction taken in some fixed order).

Let $\mathrm{C} = \wp(\wp(\boldsymbol{fcl}\,\phi^+))$: the set of ϕ^+-*colours*. For $c \in \mathrm{C}$ define

$$\gamma_c = \bigwedge_{a\in c} E\delta_a \wedge \bigwedge_{a\in(\wp(\boldsymbol{fcl}\,\phi^+)\setminus c)} \neg E\delta_a.$$

Given a colour $c \in \mathrm{C}$ the various $a \in c$ determine what we will call the *hues* of c. So c has at most $|c|$ different hues. The hue $h(a, c)$ of c corresponding to $a \in c$ is given by

$$h(a,c) = \{\delta | \delta \in a\} \cup \{\neg\delta | \delta \in \boldsymbol{fcl}\,\phi^+ \setminus a\} \cup \{E\delta_b | b \in c\} \cup \{\neg E\delta_b | b \in \wp(\boldsymbol{fcl}\,\phi^+) \setminus c\}.$$

The set $H(\phi^+)$ of all hues of ϕ^+ is

$$H(\phi^+) = \{h(a,c) | a \in c \in \mathrm{C}\}.$$

LEMMA 8. *Each hue is a hue of exactly one colour.*

PROOF. We claim that $c = \{b \in \wp(\boldsymbol{fcl}\,\phi^+) | E\delta_b \in h(a,c)\}$; from this the lemma is immediate. The inclusion $\subseteq$ follows by definition. For the other inclusion, consider $b \in \wp(\boldsymbol{fcl}\,\phi^+)$ with $E\delta_b \in h(a,c)$. By definition of $h(a,c)$ we could either have $E\delta_b \in a$ or $b \in c$. In the latter case we are done but the former case can be eliminated by a consideration of lengths of formulas as follows. We show that in general, for $a \in \wp(\boldsymbol{fcl}\,\phi^+)$, we do not have $E\delta_a \in \boldsymbol{fcl}\,\phi^+$.

Suppose that the length $|\phi^+|$ of ϕ^+ is n. Clearly each formula in $\boldsymbol{fcl}\,\phi^+$ has length at most $n+3$.

But $\phi^+, \neg\phi^+, E\phi^+, \neg E\phi^+, E\neg\phi^+, \neg E\neg\phi^+$ are six distinct formulas in $\boldsymbol{fcl}\,\phi^+$ and for each of them, either the formula or its negation is a separate conjunct of $E\delta_a$. This makes $|E\delta_a|$ at least $6n+10$ and gives us our result. □

If h is a hue of a colour then we denote the colour by h^*. From the proof of the lemma we see that $h^* = \{b \in \wp(\boldsymbol{fcl}\,\phi^+) | E\delta_b \in h\}$ and that h is a hue of h^*.

As usual, we will say that a set Γ of formulas is ($\vdash_C$-)inconsistent iff there is some $\alpha_1, \ldots, \alpha_n \in \Gamma$ such that $\vdash (\bigwedge_{i=1}^n \alpha_i) \rightarrow \boldsymbol{false}$. Otherwise, we will say the set Γ is ($\vdash_C$-)consistent. By the way, due to the lack of compactness of the logic, a consistent set of formulas is not necessarily satisfiable. We will sometimes need to relate hues and colours to the maximally $\vdash_C$-consistent sets of formulas: called MCSs. As usual we define the relations R_X and R_A on MCSs: $\Gamma R_X \Delta$ iff for all α, if $X\alpha \in \Gamma$ then $\alpha \in \Delta$; and $\Gamma R_A \Delta$ iff for all α, if $A\alpha \in \Gamma$ then $\alpha \in \Delta$.

Often we will use the fact that if Γ is an MCS then $\Delta = \{\alpha | X\alpha \in \Gamma\}$ is also one: to prove this just uses axioms from PLTL.

Also, the usual Lindenbaum technique gives us:

LEMMA 9. *If Σ is a consistent set of formulas then there is an MCS $\Gamma \supseteq \Sigma$.*

Let

$$\boldsymbol{hcl}\,\phi^+ = \{\delta, \neg\delta | \delta \in \boldsymbol{fcl}\,\phi^+\} \cup \{E\delta_b, \neg E\delta_b | b \in \wp(\boldsymbol{fcl}\,\phi^+)\}.$$

Then for each MCS Δ, the set $h = \Delta \cap \boldsymbol{hcl}\,\phi^+$ is a hue. Furthermore, for each Δ this is the only hue which satisfies $h \subseteq \Delta$.

We say that a hue is consistent if it is consistent as a set of formulas. We say that a colour c is consistent iff γ_c is a consistent formula. Note that

LEMMA 10. *a consistent hue h is maximally consistent in $\boldsymbol{hcl}\,\phi^+$, i.e. for all $\alpha \in \boldsymbol{hcl}\,\phi^+$, either $\alpha \in h$, $\neg\alpha \in h$ or $\alpha = \neg\beta$ and $\beta \in h$.*

We define two useful relations on $H(\phi^+)$. Say that hR_Xh' iff there are MCSs Γ and Γ' such that $h \subseteq \Gamma$, $h' \subseteq \Gamma'$ and $\Gamma R_X \Gamma'$. Say that hR_Ah' iff there are MCSs Γ and Γ' such that $h \subseteq \Gamma$, $h' \subseteq \Gamma'$ and $\Gamma R_A \Gamma'$.

We say that a hue is in V_i (or U_i) iff it contains an atom $s \in Q$ which is in V_i (or U_i respectively). We say that a colour is in V_i (or U_i) iff it has a hue which is in V_i (or U_i respectively).

10.2. Some lemmas. Consideration of the constituents of hues, some simple uses of C9-C14, and extensions to MCSs give us:

LEMMA 11. 1) *Consistent hues of the same colour agree on atoms and on formulas of the form* $E\beta$.

2) *If* h *is a consistent hue then* h^* *is consistent (i.e.* γ_{h^*} *is) and all hues of* h^* *are consistent.*

The following "down and across" lemma will allow us to fill in hues in the past of each of the members of the equivalence class of a new point.

LEMMA 12. *If* $h_{11}R_X h_{21}$ *then each hue* h_{22} *of* h_{21}^* *has a (consistent) predecessor hue* h_{12} *of* h_{11}^*, *i.e.* $h_{12}R_X h_{22}$.

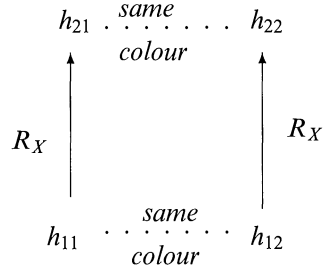

PROOF. As $h_{11}R_X h_{21}$ we have MCSs $\Gamma \supseteq h_{11}$ and $\Delta \supseteq h_{21}$ with $\Gamma R_X \Delta$.

If h_{22} is a hue of h_{21}^* then there is $a \in h_{21}^*$ with

$$h_{22} = \{\delta | \delta \in a\} \cup \{\neg\delta | \delta \in \mathbf{fcl}\,\phi^+ \setminus a\} \cup \{E\delta_b | b \in h_{21}^*\} \cup \{\neg E\delta_b | b \in \wp(\mathbf{fcl}\,\phi^+) \setminus h_{21}^*\}.$$

Also we have $E(\delta_a \wedge \gamma_{h_{21}^*}) \in \Delta$. Thus $XE(\delta_a \wedge \gamma_{h_{21}^*}) \in \Gamma$. By the contrapositive of C15, $EX(\delta_a \wedge \gamma_{h_{21}^*})$ is also in Γ.

By C9-C13, $\{X(\delta_a \wedge \gamma_{h_{21}^*})\} \cup \{\alpha | A\alpha \in \Gamma\}$ is consistent and can be extended to an MCS Ξ say. Let $h_{12} = \Xi \cap \mathbf{hcl}\,\phi^+$: this is a consistent hue.

We can show that h_{12} is a hue of h_{11}^*. We just need to show that $E\delta_b \in h_{11}^*$ iff $E\delta_b \in h_{12}$. Assuming $E\delta_b \in h_{11}^*$, we have $E\delta_b \in \Gamma$ and so $AE\delta_b \in \Gamma$ (C12 and C10). Thus $E\delta_b \in \Xi$ and in h_{12} as required. Similarly if $\neg E\delta_b \in h_{11}^*$.

Finally we show $h_{12}R_X h_{22}$ as required. We just need to show that if $\alpha \in h_{22}$ then $X\alpha \in \Xi$. In the case of α being in $\mathbf{fcl}\,\phi^+$ then $\vdash X\delta_a \rightarrow X\alpha$ and so $X\alpha \in \Xi$. Otherwise, α is $E\delta_b$ or $\neg E\delta_b$ for some $b \in \wp(\mathbf{fcl}\,\phi^+)$. But then we have $\vdash \gamma_{h_{21}^*} \rightarrow \alpha$ and $X\gamma_{h_{21}^*} \in \Xi$ to give us our result. □

§11. The idea of banning. We have seen that the basis of the construction is that for building a bundled model of ϕ^+. However, not just any bundled model of ϕ^+ will do because of the "no bad fullpaths" requirement.

We must ensure that when all the branches/fullpaths of the structure are included in the model– in particular the possibly uncountable number of "emergent" ones that were never constructed explicitly at any finite stage– then the truth of formulas along constructed branches is unaffected. We have seen that it is sufficient to ensure that along all branches, even emergent ones, we have χ_0 holding on the sequence of state formulas in the labels of the points, i.e. reading $E\gamma$ in a label as the atom $\overline{E\gamma}$ holding. Thus we require the automaton A to not accept any branch and thus, as χ_1 captures its acceptance criteria, we require χ_1 to be false along every branch. This means that we want to avoid constructing a fullpath on which some colour in

some U_i occurs infinitely often while colours in the corresponding V_i occur only finitely often. The banning mechanism will be used to provide an effective finite mechanism for bringing about this infinite property.

The basic idea of banning is as follows. We label each point with a list $[(c_1, p_1), \ldots, (c_n, p_n)]$ of banned colour-index pairs. The aim is to not allow any c_i to occur again as the colour of a hue above the point unless some colour in V_{p_i} occurs first. So, if we see a colour c in some U_p occurring on a branch in the construction but no colour in V_p has recently occurred (on that branch) then we will want to place (c, p) in the banned list. The banned list is inherited by successors with specific changes in certain circumstances.

The same banned list will apply to all points within an equivalence class. This ensures that we will be able to talk about banned lists (like colours) in the context of emergent fullpaths.

There are many subtleties. In particular note that we put colour-index pairs in the banned list and not just colours. One may wonder why we just don't keep a list of banned colours and ban all colours in U_i if we see that no colour in V_i has come up for long enough. The reason is that branches bifurcate during the construction and ideas of waiting long enough cease to be coherent in that situation: after bifurcation, a long unseen colour may turn up. For example, consider constructing a branch on which $G(q \wedge E\, X\, \neg q)$ holds. After a million steps in which we have seen the atom q being put in the label of every point of this branch at every stage, we are still going to want to make a new, parallel, branch (vertical line of points) of $\equiv$-related points (with the same coloured labels up until then) and extend it by adding a successor point (not $\equiv$-related to any point on the original branch) with a hue not containing q. The possibility of bifurcation requires that we only place colours conditionally in the banned list. If (c, p) is in the list then a colour in V_p might well eventually turn up above here (but not necessarily vertically above) in a new related branch and so we will want to unban c.

There is another conditionality which we capture by recording the banned pairs in a list rather than just a set. Suppose that we ban $(c_1, 1)$ at some stage. Because of the strict relationships between hues and successor hues, this might have the effect of subsequently stopping all colours in V_2 and V_3 from occurring along this branch. Thus, later on we may ban $c_2 \in U_2$ and $c_3 \in U_3$. Now, after a bifurcation, colours in V_1 may suddenly start showing up. However, the continued banning of c_2 may prevent colours in V_3 appearing and vice versa. It might turn out that if we allowed c_2 and c_3 to appear again then colours in V_2 and V_3 can show up too. In some sense the undesirability of c_2 and c_3 was actually dependent on the banning of c_1. We do not want to continue banning c_2 and c_3 in this situation: banning too much will frustrate our efforts to make sure that eventualities are satisfied along branches. The solution we use is to record bannings in a certain order and remove all later listed pairs when we unban a pair. There is plenty of time for colours to get banned again.

The actual order of listing of pairs (as specified in a clause K4B below) is not exactly the order of being put in the list. It turns out to be easier to reason with an equally effective order based on the order of V_ps ceasing to appear.

Another important subtlety is that we must plan ahead a little to prevent being forced, by a limited choice of successors hues, to choose a hue of a colour which

contravenes the banning condition. This planning ahead effectively limits our choice of successor hues as we avoid early errors.

To explain this idea and help with notation we introduce a formula τ_L which captures the requirement in terms of the current banned list L. Recall that the formula ν_i identifies when an atom in V_i holds while μ_i identifies U_i. If $L = [(c_1, p_1), \ldots, (c_n, p_n)]$ is a finite sequence from $\mathrm{C} \times \{1, \ldots, K\}$ then we put

$$\tau_L = \chi_1 \vee X((\neg \nu_{p_1}) \, U \, \gamma_{c_1}) \vee \cdots \vee X((\neg \nu_{p_1} \wedge \cdots \wedge \neg \nu_{p_n}) \, U \, \gamma_{c_n}).$$

If L is empty then $\tau_L = \chi_1$.

The purpose of τ_L is to describe a path which is either bad or leads through a finite sequence of colours to a future place where the banning requirement is contravened. In order to prevent us making a bad choice of hue, we make sure (in K5A below) that

$$\forall \text{ colour now} \rightarrow E \, \tau_L$$

holds at each point. In our construction, we choose the successor hue which satisfies this restriction but, to ensure fairness, has been used least recently.

Consider a banned list L as above. The first disjunct of τ_L says that there is a bad fullpath starting here. The second disjunct (if there is one) says that the path goes on to avoid V_{p_1} and ends up at c_{p_1}: a clear contravention. Later disjuncts reflect the conditionality of the ordering of the list. Consider the third disjunct. It says that the path proceeds to a point coloured with c_{p_2} without any colour in either V_{p_1} or V_{p_2} showing up on the way. Again this is a contravention as (c_2, p_2) would have remained banned until then.

11.1. Lulls. It will be important in our construction to prevent too much frenetic banning and unbanning behaviour. This is because we want to use the PLTL completeness proof idea of fair scheduling to make sure that eventualities of the form $\alpha \, U \, \beta$ are satisfied by β occurring in a label at a point above (if not in the current label itself). The fair scheduling idea requires us to try all possible hues and it is easy to see how banning could get in the way: a hue could infinitely often be allowed to appear at a later point but constantly be prevented by the time we actually get there.

Thus, after some colour-index pair becomes unbanned then we want to try to wait for quite a while to see if some colour in some V_i shows up before newly banning a colour in U_i. A new banning should only be instituted if we have waited long enough since the last time that some colour was unbanned. Hence, we here introduce the idea of a lull in unbanning.

Note that there is no need to wait after a new banning before making a subsequent new banning: bannings only further restrict the possibility of colours in V_i coming up.

We will see later that the appropriate "waiting time", during a lull in unbanning, is $e_H = 2^{|H(\phi^+)|}$.

Let x be a point in our construction which has a successor x^+. Suppose that the list of banned pairs at x is $\boldsymbol{bl}(x)$ and the list of banned pairs at x^+ is $\boldsymbol{bl}(x^+)$. We say that there has been an *unbanning* at x^+ iff $\boldsymbol{bl}(x)$ is not a prefix of $\boldsymbol{bl}(x^+)$ (written $\boldsymbol{bl}(x) \not\leq \boldsymbol{bl}(x^+)$).

Suppose that there has been an unbanning at x^+, so $\boldsymbol{bl}(x) \not\leq \boldsymbol{bl}(x^+)$. Say that $[(c_1, p_1), \ldots, (c_n, p_n)]$ is the longest common prefix of $\boldsymbol{bl}(x)$ and $\boldsymbol{bl}(x^+)$,

$$\boldsymbol{bl}(x) = [(c_1, p_1), \ldots, (c_n, p_n), (c_{n+1}, p_{n+1}), \ldots, (c_m, p_m)]$$

and

$$\boldsymbol{bl}(x^+) = [(c_1, p_1), \ldots, (c_n, p_n), (c'_{n+1}, p'_{n+1}), \ldots, (c'_{m'}, p'_{m'})].$$

Then we say that each (c_i, p_i), for $i \in \{n+1, \ldots, m\}$, has been *unbanned* and each (c'_i, p'_i), for $i \in \{n+1, \ldots, m'\}$, has been *newly banned* at x^+. It is clear that if there is an unbanning at x^+ then something is unbanned at x^+.

If $x_n < x_{n+1} < \cdots < x_{n+e_H}$ are points in our construction such that each x_{i+1} is the (immediate) successor of x_i and there has not been an unbanning at any $x_i (n \leq i < n + e_H)$ then we say that there has been a *lull* in unbanning between x_n and x_{n+e_H}.

§12. **Chronicles.** At each stage during the construction and in the limit, we will have a chronicle.

A *chronicle* is $(T, <, \equiv, \lambda, \boldsymbol{bl}, X_0)$ where:

$(T, <, \equiv)$ is a floored Ockhamist frame;
$\lambda : T \to H(\phi^+)$;
$\boldsymbol{bl} : T \to ({}^{<\omega}(\mathrm{C} \times \{1, \ldots, K\}))$ (where K, recall, is the number of accepting pairs of A);
$X_0 \in T$;

such that for all $x, y \in T$:

K0 The floor is $X_0/\equiv$ and ϕ^+ is in $\lambda(X_0)$.

K1 $\lambda(x)$ is consistent.

K2 If x has an immediate $<$-successor in $(T, <)$ —call it x^+: it is unique— then $\lambda(x) R_X \lambda(x^+)$.

K3A If $x \equiv y$ then $\lambda(x)$ and $\lambda(y)$ are hues of the same colour.

K3B If h is a hue of $\lambda^*(x)$ then there is $y \in T$ such that $x \equiv y$ and $h = \lambda(y)$.

K4A If (c, p) occurs in $\boldsymbol{bl}(x)$ then:

1) c has appeared at or below x (i.e. $c = \lambda^*(y)$ for some $y \leq x$); and
2) there is a lull in unbanning reaching its e_Hth step below x such that no colour in V_p appears between the start of the lull and x;

K4B If (d, q) precedes (c, p) in $\boldsymbol{bl}(x)$ then: 1) the most recent appearance of V_q (if any) below x is below (or equal to) the most recent appearance of V_p; and 2) if there is no appearance of V_p before x then there is no appearance of V_q before x.

K4C If $x \equiv y$ then $\boldsymbol{bl}(x) = \boldsymbol{bl}(y)$.

K5A $\forall \gamma_{\lambda^*(x)} \to E\, \tau_{\boldsymbol{bl}(x)}$.

K5B If x^+ exists and (c, p) occurs in $\boldsymbol{bl}(x)$ then $\lambda^*(x^+) \neq c$.

K5C No (c, p) occurs more than once in $\boldsymbol{bl}(x)$.

K5D If x^+ exists and (c, p) is newly banned at x^+ then $\lambda^*(x^+) = c$.

Most of this definition is motivated by the standard modal logic ideas of step-by-step model construction. The clauses K4A-K5A have been motivated in section 11 above. K5B says that a currently banned colour does not appear in a label, K5C ensures banned lists stay manageable and K5D says that only the current colour is allowed to enter the banned list at a point.

Throughout our construction we will have a finite chronicle, i.e. T is finite. Only in the limit will T be infinite but even then each point will only have a finite past in accordance with the requirements of an Ockhamist frame.

A straightforward induction using K0 (in particular $\theta_A \in \lambda(X_0)$), K1, K2, K3A and lemma 11 gives us the following useful result about the distribution within the labels of the atoms from Q, the atoms of the automaton's state.

LEMMA 13. *For each x in a chronicle we have the following*:

a) $A\,G \bigwedge_{s\in Q, c\subseteq \mathbf{ecl}\,\phi}((s \wedge \zeta_c) \rightarrow A\,X\,\rho(s,\overline{c})) \in \lambda(x)$ *and for all $s \in Q$ and for all $c \subseteq \mathbf{ecl}\,\phi$, if each of the conjuncts of ζ_c is in $\lambda(x)$, $s \in \lambda(x)$ and x^+ exists then $\rho(s,\overline{c}) \in \lambda(x^+)$.*

b) $A\,G \bigwedge_{s\neq r\in Q} \neg(r \wedge s) \in \lambda(x)$ *and there is only one s from Q in $\lambda(x)$.*

This will allow us to deduce that the atoms do reflect the states of A as it reads the $\overline{\mathbf{ecl}(\phi)}$ versions of the labels along branches. See lemma 26 below for a formal statement of this when we need it.

12.1. Pioneers. We will identify some of the elements of T as *pioneers* of their $\equiv$-classes and some others as the *siblings* of pioneers. Note that some elements of T may be neither pioneers nor siblings of pioneers. We need to know which points are pioneers because the fair hue scheduling technique will only apply to them.

We will ensure:

P1: Every point in T is below (or equal to) a pioneer or sibling of a pioneer.

P2: If x is a pioneer or sibling of a pioneer and x^+ exists then x^+ is a pioneer.

P3: If x is a pioneer and x^+ exists then

$$\forall\,(\bigwedge \lambda(x) \wedge X(\bigwedge \lambda(x^+))) \rightarrow E\,\tau_{\mathbf{bl}(x)}.$$

P4: If x is a pioneer and x^+ exists then $\lambda(x^+)$ is a fair choice amongst the possible hues which satisfy P3 in the sense that there are no other such hues which have come up less recently as $\lambda(y)$ in the past $\{y | y < x\}$ of x.

P3 is a stronger property than K5A but for the same purpose: we need to make sure that when we choose a successor hue then the fact of the combined pair of hues being successors does not force us to go on contravene the banning restrictions.

12.2. The start.

LEMMA 14. *There is a finite chronicle satisfying P1-P4.*

PROOF. Let Γ_0 be any MCS extending ϕ^+ and let $h = \Gamma_0 \cap \mathbf{hcl}\,\phi^+$. Say that h^* has n hues, $h_0, \ldots, h_{n-1}$ including $h = h_0$. We know that these will be consistent hues. Choose n objects $X_0, \ldots, X_{n-1}$.

Our construction starts with $(T_0, <_0, \equiv_0, \lambda_0, \mathbf{bl}_0, X_0)$ where $T_0 = \{X_0, \ldots, X_{n-1}\}$, $<_0$, and $\mathbf{bl}_0(X_0), \ldots, \mathbf{bl}_0(X_{n-1})$ are empty, $\equiv_0 = \{(X_i, X_j) | i, j < n\}$ and $\lambda_0(X_i) = h_i$. It is a simple matter to check that the conditions hold. K5A holds as $\neg E\,\chi_1$ is in each $\lambda(X_i)$. We have $\neg E\,\chi_1 \in \mathbf{fcl}\,\phi^+ \subseteq \mathbf{hcl}\,\phi^+$. We also have $\neg E\,\chi_1 \in \Gamma_0$ as it is a

conjunct of ϕ^+. Thus $\neg E\chi_1 \in h_0$. By lemma 11, $\neg E\chi_1$ is in each h_i. We will say that X_0 is the pioneer of the $\equiv_0$-class and the other X_i are its siblings. □

§13. Curing a defect. In this section we show that we can add a successor to any point in a chronicle.

LEMMA 15. *Suppose* $(T, <, \equiv, \lambda, \boldsymbol{bl}, X_0)$ *is a finite chronicle satisfying* $P1$-$P4$. *Say that* $x \in T$ *and there is no* $y \in T$ *with* $x < y$. *We can define a new chronicle* $(T', <', \equiv', \lambda', \boldsymbol{bl}', X_0)$ *to slightly extend* $(T, <, \equiv, \lambda, \boldsymbol{bl}, X_0)$ *with the addition of a new element to be* x^+ *and probably some other new elements as well.* $(T', <', \equiv', \lambda', \boldsymbol{bl}', X_0)$ *is also a finite chronicle satisfying* $P1$-$P4$.

The rest of this section is devoted to proving this.

If it wasn't for the banning mechanism then the set of possible hues we use as $\lambda'(x^+)$ is

$$\{h \in H(\phi^+) | \lambda(x) R_X h\}.$$

In order to respect banning we have to be a little more selective. For a hue h and banned list B we will define the set $Pn(h, B)$ of hues to contain just those which we will allow after a point with hue h and banned list B:

DEFINITION 6. Suppose that $h \in H(\phi^+)$ and $B \in ({}^{<\omega}(\mathrm{C} \times \{1, \ldots, K\}))$. We define $Pn(h, B)$ to be the set of $h' \in H(\phi^+)$ such that

$$\nvdash (\bigwedge h \wedge X(\bigwedge h')) \rightarrow E\tau_B.$$

We now show that $Pn(\lambda(x), \boldsymbol{bl}(x))$ is non-empty (and some other useful things).

LEMMA 16. *Suppose* $h \in H(\phi^+)$ *and* $B \in ({}^{<\omega}(\mathrm{C} \times \{1, \ldots, K\}))$. *Then*

1. *if* h *and* B *satisfy K5A, i.e.* $\nvdash \gamma_{h^*} \rightarrow E\tau_B$, *then* $Pn(h, B)$ *is non-empty*;
2. *all the hues in* $Pn(h, B)$ *are consistent*; *and*
3. *if* $h' \in Pn(h, B)$ *then* $hR_X h'$.

PROOF. (1.) Suppose for contradiction that it is. This means that for all $h' \in H(\phi^+)$, we have $\vdash (\bigwedge h \wedge X \bigwedge h') \rightarrow E\tau_B$. Since $\vdash \bigvee_{h' \in H(\phi^+)} \bigwedge h'$, as it is just a substitution instance of a propositional tautology, we have, using a few PLTL axioms, that $\vdash \bigwedge h \rightarrow \bigvee_{h' \in H(\phi^+)} (\bigwedge h \wedge X \bigwedge h')$. Putting these facts together gives us $\vdash \bigwedge h \rightarrow E\tau_B$. Now, suppose that $h = h(a, c)$ so that $a \in c = h^*$. As γ_{h^*} is just a conjunction of formulas of the form $E\beta$ and their negations, and $E\delta_a$ is one of the conjuncts, we can use the S5 axioms to show that $\vdash \gamma_{h^*} \rightarrow E(\delta_a \wedge \gamma_{h^*})$. But clearly, $\vdash (\delta_a \wedge \gamma_{h^*}) \rightarrow \bigwedge h$. This gives us $\vdash \gamma_{h^*} \rightarrow E\tau_B$ contradicting K5A.

(2.) follows as

$$\nvdash (\bigwedge \lambda(x) \wedge X \bigwedge h') \rightarrow E\tau_B$$

for all $h' \in Pn(\lambda(x), \boldsymbol{bl}(x))$.

(3.) Assume $h' \in Pn(h, B)$. Then $\bigwedge h \wedge X \bigwedge h'$ is consistent and we can extend it to an MCS Γ say. Let $\Delta = \{\alpha | X\alpha \in \Gamma\}$. This is also an MCS and, as $h' \subset \Delta$, it is clear that $hR_X h'$ as required. □

Say that all the elements of T which are below x are exactly

$$x_1 < x_2 < \cdots < x_r = x.$$

This gives us a sequence of hues $h_i = \lambda(x_i)$. Choose a possible hue h which we can put at x^+, so that h has come up least recently as an h_i. That is, for each $h \in Pn(\lambda(x), \boldsymbol{bl}(x))$, put $l(h) = \max(\{-1\} \cup \{i | h = h_i\})$. Choose one of the $h \in Pn(\lambda(x), \boldsymbol{bl}(x))$ with the smallest $l(h)$. Say it is h_+.

Say that the hues of h_+^* are exactly $h_+^0, \ldots, h_+^{n-1}$ with $h_+^0 = h_+$.

We first choose new objects $z_0, \ldots, z_{n-1} \notin T$. Also, for each predecessor $x_j (j = 1, \ldots, r)$ of x (including $x = x_r$ itself) and each $i = 1, \ldots, n-1$, choose a new object z_{ij}.

Let $T' = T \cup \{z_0, \ldots, z_{n-1}\} \cup \{z_{ij} | i = 1, \ldots, n-1, j = 1, \ldots, r\}$. Extend $<$ to $<'$ as follows:

$$
\begin{aligned}
<' = \; & < \\
& \cup \{(x_j, z_0) | j = 1, \ldots, r\} \\
& \cup \{(z_{ij}, z_i) | i = 1, \ldots, n-1, j = 1, \ldots, r\} \\
& \cup \{(z_{ij}, z_{ik}) | i = 1, \ldots, n-1, j = 1, \ldots, r, k = j+1, \ldots, r\}.
\end{aligned}
$$

Thus z_0 is the successor of x in $(T', <', \equiv')$. Extend $\equiv$ to $\equiv'$ by adding

$$
\{(z_i, z_q)\} \cup \{(x_j, z_{ij})\} \cup \{(z_{ij}, x_j)\} \cup \{(z_{ij}, z_{kj})\} \cup \{(y, z_{ij}), (z_{ij}, y) | y \equiv x_j\}.
$$

In Figure 5, we can see the new successor z_0 added above $x = x_r$ and the grid of $n-1$ parallel vertical towers, each of height $r+1$, added along side the old chronicle to preserve K3B and the property of being an Ockhamist frame (amongst other things).

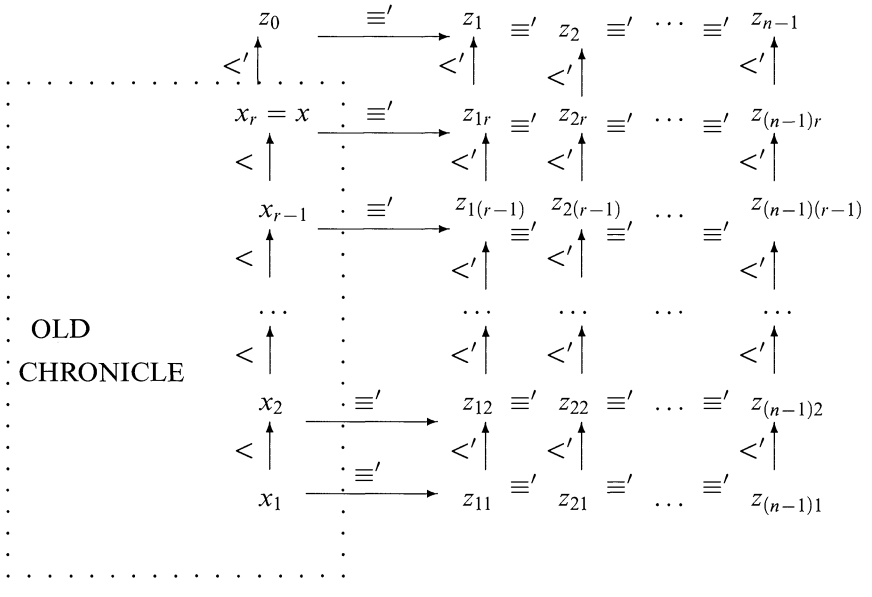

Figure 5. The new elements

We call $z_0 = x^+$ the pioneer of its $\equiv'$-class and $z_1, \ldots, z_{n-1}$ its siblings. Other new points are neither pioneers nor siblings. Old points inherit their previous classifications.

Let $\lambda'(z_i) = h^i_+$. We fill in the hues of their predecessors from the top down using "the down and across lemma", lemma 12. That is we choose appropriate hues in the order

$$\lambda'(z_{1r}), \lambda'(z_{2r}), \ldots, \lambda'(z_{(n-1)r}), \lambda'(z_{1(r-1)}), \lambda'(z_{2(r-1)}), \ldots, \lambda'(z_{(n-1)1}).$$

13.1. Bannings. We also need to decide on what to ban at x^+ (and its siblings).

Recall that for any y, if $\boldsymbol{bl}(y)$ is not a prefix of $\boldsymbol{bl}(y^+)$ then we say that there is an *unbanning* at y^+.

First, what to unban. Suppose that $\boldsymbol{bl}(x) = [(c_1, p_1), \ldots, (c_n, p_n)]$. If $\lambda'^*(x^+) \notin V_{p_i}$ for all $i \leq n$ then put $e = n$; otherwise let e be such that $\lambda'^*(x^+) \in V_{p_{e+1}}$ but $\lambda'^*(x^+) \notin V_{p_i}$ for all $i \leq e$. Then we are only going to keep (at most) $[(c_1, p_1), \ldots, (c_e, p_e)]$ in the new banned list. Recall that any pairs from (c_{e+2}, p_{e+2}) onwards are unbanned because their banning is deemed to be dependent on the banning of (c_{e+1}, p_{e+1}) which is now rescinded.

Now let us consider whether we want to ban anything new. The general idea is to newly ban $(\lambda'^*(x^+), p)$ at x^+ (and its siblings) if and only if $\lambda'^*(x^+) \in U_p$ and there are points $y < z < x^+$ such that

- the period (y, z) contained no unbannings and was over e_H steps long (–we will see later that this is long enough for the process of cycling through hues to have covered all hues it will ever get to–) and
- no colour in V_p appeared between y and x^+.

The order of listing these new banned pairs (if there are more than one) is determined by which V_p came up least recently in the past: the least recent ones go into the list first. Say that the order so determined for the new pairs is

$$[(\lambda'^*(x^+), q_1), \ldots, (\lambda'^*(x^+), q_l)].$$

Recall that the purpose of this ordering as stipulated in K4B is to reflect a conditionality on the banning.

All the new banned pairs are added after the ones surviving from $\boldsymbol{bl}(x)$. However, some of the new banned pairs might be more important to put earlier in the new banned list than some surviving (c_i, p_i) pairs. We just need look at q_1 versus $p_1, \ldots, p_e$. Suppose that in the past of x, the latest appearance of V_{q_1}, if there was one, was before the latest appearance of $V_{p_f} (f \leq e)$ but after (or equal to) the latest appearance of V_{p_i} for each $i < f$. Alternatively suppose that there has been no appearance of V_{q_1} or V_{p_i} for each $i < f$ but that V_{p_f} has appeared. In either of those cases we throw away $(c_f, p_f), \ldots, (c_e, p_e)$ as well. We might say that $(c_f, p_f), \ldots, (c_e, p_e)$ are *usurped* by $(\lambda'^*(x^+), q_1), \ldots, (\lambda'^*(x^+), q_l)$. Note that this counts as an unbanning.

If, in the past of x, the latest appearance of V_{q_1} was not strictly before the latest appearance of V_{q_e} then we do not throw away any extra pairs. Put $f = e + 1$.

The new banned list is

$$\boldsymbol{bl}'(x^+) = [(c_1, p_1), \ldots, (c_{f-1}, p_{f-1}), (\lambda'^*(x^+), q_1), \ldots, (\lambda'^*(x^+), q_l)].$$

The banned list at all the siblings of x^+ is the same as that at x^+.

Note that we will see soon (in lemma 18) that no successor hue can be a hue of a currently banned colour: thus the colours we newly ban at a particular point are not already in the banned list.

To preserve K4C we put $\boldsymbol{bl}'(z_{ij}) = \boldsymbol{bl}(x_j)$ for each i, j.

13.2. A new chronicle? We check that $(T', <', \equiv', \lambda', \boldsymbol{bl}', X_0)$ is a new chronicle.

It is clear that $(T', <', \equiv')$ is a finite floored Ockhamist frame satisfying P1 and P2. P3 and P4 hold by construction.

K0 continues to hold and K1 holds as the new hues we have used as labels were all chosen to be consistent.

K2 follows from lemma 16(part 3) and our down-and-across construction.

K3A–K4C and K5D all hold by construction.

It remains to check that K5A, K5B and K5C continue to hold. First K5A. In this lemma we establish a sufficient condition for K5A to be preserved. It is a not particularly strong condition, allowing some flexibility in what we unban and newly ban, but it is clearly satisfied by our actual banning procedure. The lemma is crucial (and is where the new LC axiom is used) so we go into some detail in the proof.

LEMMA 17. *Suppose that* $h, h' \in H(\phi^+)$ *and* $B, B' \in {}^{<\omega}(\mathrm{C} \times \{1, \ldots, K\})$. *Suppose that* $h' \in Pn(h, B)$. *Suppose that* $B = [(c_1, p_1), \ldots, (c_n, p_n)]$, $m \leq n$ *and*

$$B' = [(c_1, p_1), \ldots, (c_m, p_m), (h'^*, q_1), \ldots, (h'^*, q_l)],$$

with h'^* *not in* V_{p_i} *for all* $i = 1, \ldots, m$ *and* h'^* *in* U_{q_i} *for all* $i = 1, .., l$. *Then* $\not\vdash \gamma_{h'^*} \rightarrow E\,\tau_{B'}$.

PROOF. Suppose for contradiction that $\vdash \gamma_{h'^*} \rightarrow E\,\tau_{B'}$, i.e. that choosing the hue h' will inevitably lead to the new banning condition B' being contravened.

First assume $l > 0$: we return to the case of $l = 0$ below.

The case of $l > 1$ follows from the argument for $l = 1$ as, for each j,

$$\vdash X((\bigwedge_{i=1}^{m} \neg v_{p_i} \wedge \bigwedge_{i=1}^{j} \neg v_{q_i})\, U\, \gamma_{h'^*}) \rightarrow X((\bigwedge_{i=1}^{m} \neg v_{p_i} \wedge \neg v_{q_1})\, U\, \gamma_{h'^*})$$

showing that $\vdash \tau_{B'_l} \rightarrow \tau_{B'_1}$ (with obvious definitions).

So we now assume that $l = 1$.

If $m = 0$ then put $\theta = \boldsymbol{false}$ and otherwise put

$$\theta = X((\neg v_{p_1})\, U\, \gamma_{c_1}) \vee X((\neg v_{p_1} \wedge \neg v_{p_2})\, U\, \gamma_{c_2}) \vee$$
$$\cdots \vee X((\neg v_{p_1} \wedge \neg v_{p_2} \wedge \cdots \wedge \neg v_{p_m})\, U\, \gamma_{c_m})$$

and

$$\theta' = X((\neg v_{p_1} \wedge \neg v_{p_2} \wedge \cdots \wedge \neg v_{p_m} \wedge \neg v_{q_1})\, U\, \gamma_{h'^*})$$

so that $\tau_{B'} = \chi_1 \vee \theta \vee \theta'$ and our assumption is $\vdash \gamma_{h'^*} \rightarrow E(\chi_1 \vee \theta \vee \theta')$.

Some basic uses of C9-C13 gives us

$$\vdash (\gamma_{h'^*} \wedge \neg E\,\chi_1 \wedge \neg E\,\theta) \rightarrow E\,\theta'.$$

Let $\alpha = \gamma_{h'^*} \wedge \neg E\,\chi_1 \wedge \neg E\,\theta$, and $\beta = \neg v_{p_1} \wedge \cdots \wedge \neg v_{p_m} \wedge \neg v_{q_1}$ so that C11 gives us $\vdash E\,\alpha \rightarrow E\,\theta'$.

Now α will hold at any point with hue h' which is also the successor of a point with banned list B (or even just the m-long prefix of B). We have thus just established that

an $\equiv$-class of points containing one satisfying α must begin a finite path witnessing θ'. The next few steps of the proof will show that such a path will not pass through $V_{p_1}, \ldots, V_{p_m}, V_{q_1}$ (i.e., β will constantly hold) and will end up at another point satisfying α.

From $\vdash E\alpha \rightarrow EX((E\beta)\, U\, \gamma_{h'^*})$, C0-C13 allows us to deduce that $\vdash E\alpha \rightarrow EX((E\beta)\, U(\gamma_{h'^*} \wedge (\neg E\chi_1 \vee E\chi_1)))$. Further uses, distributing connectives over disjunction gives us $\vdash E\alpha \rightarrow (EX((E\beta)\, U(\gamma_{h'^*} \wedge \neg E\chi_1)) \vee EX((E\beta)\, U(\gamma_{h'^*} \wedge E\chi_1)))$.

As $\vdash F E\chi_1 \rightarrow E\chi_1$ and $\vdash E\alpha \rightarrow \neg E\chi_1$, it follows that

$$\vdash E\alpha \rightarrow EX((E\beta)\, U(\gamma_{h'^*} \wedge \neg E\chi_1)).$$

Now $\vdash EX((E\beta)\, U(\gamma_{h'^*} \wedge \neg E\chi_1 \wedge E\theta)) \rightarrow E\theta$ and so we can similarly conclude that

$$\vdash E\alpha \rightarrow EX((E\beta)\, U(E\alpha)).$$

Intuitively, any point with hue h' which is also the successor of a point with banned list B must begin a finite path avoiding $V_{p_1}, \ldots, V_{p_m}, V_{q_1}$ and ending at another such point. What follows is a use of the LC axiom to allow us to deduce that such a point begins an infinite sequence of such finite paths joined end to end.

By generalization we have

$$\vdash A\, G(E\alpha \rightarrow EX((E\beta)\, U(E\alpha)))$$

and so modus ponens and the LC axiom give us

$$\vdash E\alpha \rightarrow EG((E\beta)\, U(E\alpha)).$$

The infinite path is clearly bad: a short proof using the PLTL system gives us

$$\vdash G((E\beta)\, U(E\alpha)) \rightarrow \chi_1$$

as h'^* is in U_{q_1} and not in V_{q_1}. This needs the assumption that $V_{q_1} \cap U_{q_1} = \emptyset$.

Putting these together we get

$$\vdash E\alpha \rightarrow E\chi_1$$

which, by using C9-C13, gives us

$$\vdash (\gamma_{h'^*} \wedge \neg E\chi_1 \wedge \neg E\theta) \rightarrow E\chi_1.$$

Thus, from $\vdash \bigwedge h' \rightarrow \gamma_{h'^*}$, it follows that the formula $(\bigwedge h' \wedge \neg E\chi_1 \wedge \neg E\theta)$ is inconsistent.

The following claim gives us our contradiction to $h' \in Pn(h, B)$: it is not acceptable to choose the hue h' at a successor to a point with banned list B. The proof above has established that the hue h' necessarily begins a fullpath which is either bad or satisfies θ.

The claim is that $\vdash (\bigwedge h \wedge X \bigwedge h') \rightarrow E\tau_B$, ie $h' \notin Pn(h, B)$.

Proof. From above we have, $\vdash \bigwedge h' \rightarrow E(\chi_1 \vee \theta)$.

Also, as h'^* is not in any V_{p_j}, $\vdash \bigwedge h' \rightarrow \bigwedge_{j=1}^{m} \neg v_{p_j}$ and so $\vdash \bigwedge h' \rightarrow E(\chi_1 \vee \bigvee_{j=1}^{m}((\bigwedge_{i=1}^{j} \neg v_{p_i})\, U\, \gamma_{c_j}))$.

By C14, $\vdash X E\zeta \to E X\zeta$ (for any ζ), so

$$\vdash X \bigwedge h' \to E(X\chi_1 \vee \bigvee_{j=1}^{m} X((\bigwedge_{i=1}^{j} \neg v_{p_i})\, U\, \gamma_{c_j})).$$

And $\vdash X\chi_1 \to \chi_1$ so $\vdash (\bigwedge h \wedge X \bigwedge h') \to E(\chi_1 \vee \theta)$. It is clear that $\vdash E(\chi_1 \vee \theta) \to E\tau_B$ and the claim is proved. □

The case of $l = 0$ is a simplified version of the above: it follows almost immediately from the initial assumption that $\bigwedge h' \wedge \neg E\chi_1 \wedge \neg E\theta$ is inconsistent. We then just use the claim above to conclude the argument. □

Now we must show that K5B continues to hold. First consider x which we have just given a successor.

LEMMA 18. *If (c, p) occurs in $\boldsymbol{bl}(x)$ then we do not have $\lambda'^*(x^+) = c$.*

PROOF. Assume otherwise and suppose that $\boldsymbol{bl}(x) = [(c_1, p_1), \ldots, (c_n, p_n)]$. As $\lambda'(x^+) \in Pn(\lambda, \boldsymbol{bl}(x))$, we have

$$\not\vdash \bigwedge \lambda(x) \wedge X \bigwedge \lambda'(x^+) \to E(\chi_1 \wedge (\neg v_{p_1})\, U\, \gamma_{c_1} \wedge \cdots \wedge (\neg v_{p_1} \wedge \cdots \wedge \neg v_{p_n})\, U\, \gamma_{c_n})$$

and $\lambda'^*(x^+) = c_j$ for some j. Now we can not have $\neg v_{p_1}, \ldots, \neg v_{p_j}$ all in $\lambda(x)$ because then

$$\vdash \bigwedge \lambda(x) \wedge X \bigwedge \lambda'(x^+) \to ((\neg v_{p_1} \wedge \cdots \wedge \neg v_{p_j})\, U\, \gamma_{c_j}).$$

Thus there is some $i \leq j$ with $\neg v_{p_i} \notin \lambda(x)$, i.e. $\lambda^*(x)$ is in V_{p_i}. Thus (c_i, p_i) would have been unbanned at x if it was in the banned list of the predecessor of x. The only way it could be in $\boldsymbol{bl}(x)$ is if it was put in at x. Thus c_i is $\lambda^*(x)$ and c_i in U_{p_i}. Given lemma 13, this contradicts our assumption about the Rabin automaton that each $U_{p_i} \cap V_{p_i}$ is empty. □

K5B continues to hold for any old points. It remains to check it for the z_{ij} as the z_j do not have successors. However, $\boldsymbol{bl}'(z_{ij}) = \boldsymbol{bl}(x_j)$, by definition, and $\lambda'^*(z_{ij}) = \lambda^*(x_j)$, by the construction using the down and across lemma, so we just need to call on K5B for x_i to give us our result.

From lemma 18 we can see that K5C also continues to hold. This is because the pairs (c, p) which are added into the banned list $\boldsymbol{bl}'(x^+)$ all have $c = \lambda'^*(x^+)$. By lemma 18, such pairs do not appear amongst the pairs which are inherited from $\boldsymbol{bl}(x)$. We also do not add repeats amongst the new pairs.

§14. The limit. If we begin with the simple starting chronicle described above (lemma 14) and continue to cure defects, i.e. add successor points, using lemma 15 in such a way that every point eventually gets a successor, then we clearly end up with a chronicle satisfying P1 and P2.

The limit chronicle $(T, <, \equiv, \lambda, \boldsymbol{bl}, X_0)$ has a few extra properties:

L1 every point has an immediate successor;
L2 and every point has a pioneer above it.

14.1. A Kripke structure. Now we use the standard correspondence between a $(\mathbb{N} \times W)$-frame and a bundled ω-tree (as described above in section 4). The tree can be seen as a Kripke structure and this will be our model.

Let $S = T/\equiv$. Define $R \subseteq S \times S$ by $(s, s') \in R$ iff there is some $t \in s$ with $t^+ \in s'$. This is total. Define $g : S \to \wp(L \cup Q)$ by $p \in g(s)$ iff there is some $t \in s$ with $p \in \lambda(t)$. Thus (S, R, g) is a Kripke-structure.

For each $t \in T$ we shall write $[t]$ for the $\equiv$-class $t/\equiv \in S$. By K3A and K4C, we can define $\lambda^*([t]) = \lambda^*(t)$ and $\boldsymbol{bl}([t]) = \boldsymbol{bl}(t)$ and these are well-defined.

For each $t \in T$ define a sequence $t_0, t_1, \ldots$ by $t_0 = t$ and $t_{i+1} = t_i^+$. This gives us a fullpath $\boldsymbol{up}(t) = \langle [t_0], [t_1], [t_2], \ldots \rangle$ in (S, R).

It is important to realise that the $\boldsymbol{up}(t)$ are generally not the only fullpaths through the Kripke structure (S, R). As we have seen in the example in section 4, there maybe other "emergent" fullpaths. Recall that a fullpath through (S, R) is any sequence $s_0, s_1, \ldots$ from S such that each $(s_i, s_{i+1}) \in R$. It may not be possible to find some $t \in s_0$ such that each t_i is in s_i. A lot of the work we have done and we still have to do will be to make sure that we can handle these emergent fullpaths along with the constructed $\boldsymbol{up}(t)$ ones in the truth lemma.

Note that as $(T, <, \equiv)$ is floored, with $[X_0]$ as the floor, for every $s \in S$ there is a fullpath starting at $[X_0]$ and passing through s.

14.2. Universal non-acceptance. In this subsection we show that no fullpaths, not even emergent ones, are bad: none are accepted by the automaton A. We show that along any fullpath in (S, R) we do not have some $i = 1, \ldots, K$ such that there are an infinite number of labels along the fullpath containing an atom from U_i but only a finite number containing an atom from V_i. Recall that this is a crucial observation to allow the truth lemma to work.

We will talk of colour-index pairs becoming newly banned and unbanned along a fullpath and of lulls (in unbanning) along a fullpath. These terms are defined in the analogous way to the respective terms for points in the chronicles in terms of the banned list $\boldsymbol{bl}(s_i)$ at each $\equiv$-class s_i along a fullpath. Equivalently we could define them by inheriting the terms from any $<$-related $\equiv$-class representatives.

First a helpful lemma implying that there are infinitely many lulls along any fullpath.

LEMMA 19. *Along any fullpath $\langle s_0, s_1, \ldots \rangle$ in (S, R) there are infinitely many different pairs of indices (j, k) such that $k \geq j + e_H$ and there is no unbanning at any $t \in s_i$ with $j \leq i < k$; i.e there are infinitely many lulls.*

PROOF. Assume otherwise for contradiction, i.e there are only finitely many lulls along the fullpath. Thus there must be some pair (c, p) which is newly banned and unbanned infinitely often. To be newly banned infinitely often, c must come up as $\lambda^*(s_i)$ infinitely often (K5D). There are two cases.

A colour in V_p may come up infinitely often along the fullpath. But then after the end of the last lull, and after a subsequent appearance of a colour in V_p, (c, p) can not be newly banned (K4A). So this case does not occur.

Thus we may assume that colours in V_p only come up finitely often along the fullpath. Consider the non-empty set G of such pairs (c', p') with $c' \in U_{p'}$ appearing infinitely often, colours in $V_{p'}$ appearing only finitely often and the pair (c', p') being newly banned and unbanned infinitely often. Of all the pairs in G,

find one (c', p') of those such that the very last appearance of a colour in $V_{p'}$ along the fullpath is earliest.

Now find an index i after the end of the last lull along $\sigma = \langle s_0, s_1, \ldots \rangle$ and after all the colours not in $\inf(\sigma) = \{c \in \mathrm{C} \,|\, c = \lambda^*(s_j) \text{ for infinitely many } j\}$ have stopped coming up forever. Find a later index after all the colours in $\inf(\sigma)$ have come up at least once subsequently. Also, wait until from then on, the only pairs being added or taken away from the banned lists are those which will do so infinitely often. Find the next index at which (c', p') is newly banned.

We can show that (c', p') can never subsequently be unbanned and this will be our contradiction. If it is unbanned then it is not because a colour in $V_{p'}$ has come up. Thus, there must be a pair (d, q) before (or equal to) (c', p') in the banned list which either gets directly unbanned (by V_q coming up) or gets usurped by an "older" pair.

Because (d, q) is before (or equal to) (c', p') we know that the most recent appearance of V_q was before or equal to the most recent appearance of $V_{p'}$ (K4B). This was long enough ago for us to deduce that V_q only comes up finitely often: no colour in V_q has come up since i and yet all members in $\inf(\sigma)$ have. Thus the unbanning is not because V_q came up.

This leaves the possibility that a colour $e \in U_r$ comes up such that the most recent appearance of V_r was strictly before the most recent appearance of (d, q) and so strictly before the most recent appearance of $V_{p'}$. It is clear that (e, r) is in G: e has come up since i, no colour in V_r has and yet the colours coming up since i are precisely those that come up infinitely often. But then we have a contradiction to the choice of (c', p'). Such an e is thus not a possibility.

This completes the proof. □

So now let us state our main claim. Recall that we say that a colour is in V_i (or U_i) iff it has a hue containing an atom which is in V_i (or U_i respectively).

LEMMA 20. *Along any fullpath $\langle s_0, s_1, \ldots \rangle$ in (S, R) we do not have some $i = 1, \ldots, K$ such that there are an infinite number of js with $\lambda^*(s_j)$ in U_i but only a finite number of js with $\lambda^*(s_j)$ in V_i.*

PROOF. Suppose for contradiction that there were such a fullpath $\pi = \langle s_0, s_1, \ldots \rangle$ in (S, R) and such an $p = 1, \ldots, K$. Say that the colour $c \in U_p$ comes up infinitely often (as $\lambda^*(s_j)$) but no colour in V_p does.

Consider the non-empty set G of such pairs (c', p') with $c' \in U_{p'}$ appearing infinitely often, and all colours in $V_{p'}$ appearing only finitely often. Of all the pairs in G, choose one (c', p') of those such that the very last appearance of any colour in $V_{p'}$ along π is earliest. The rest of the proof is similar to that of the previous lemma.

Find an index i after all the colours not in $\inf(\pi) = \{c \in \mathrm{C} \,|\, c = \lambda^*(s_j)$ for infinitely many $j\}$, (so including all those in $V_{p'}$) have stopped coming up forever in π. Find a later index after all the colours in $\inf(\pi)$, have come up subsequently. Using the last lemma we know that there are infinitely many lulls along π so we can find an even greater index after the end of a subsequent lull. Find the next index i' at which c' comes up. Clearly (c', p') will be put in the banned list here.

We claim that it is never subsequently unbanned. If it is unbanned then there must be a pair (d, q) before (or equal to) (c', p') in the banned list which either gets directly unbanned (by V_q coming up) or gets usurped by an "older" pair.

Because (d, q) is before (or equal to) (c', p') we know that the most recent appearance of V_q was before or equal to the most recent appearance of $V_{p'}$. This was long enough ago for us to deduce that V_q only comes up finitely often. Thus the unbanning is not because V_q came up.

This leaves the possibility that a colour $e \in U_r$ comes up such that the most recent appearance of V_r was strictly before the most recent appearance of (d, q) and so strictly before the most recent appearance of $V_{p'}$. As in the proof of the previous lemma, it is clear that (e, r) is in G. But then we have a contradiction to the choice of (c', p'). Such an e is thus not a possibility.

So we can conclude that (c', p') is never subsequently unbanned. By K5B, this is a contradiction to our assumption that c' comes up infinitely often. □

14.3. Hue scheduling. Assuming that there is no unbanning going on, how long does it take to try all the possible hues which are ever going to come up? Here we will show that $e_H = 2^{|H(\phi^+)|}$ steps is long enough.

Our result follows from the following graph theoretic result.

Suppose that we have a finite set H and a sequence of relations $E = \langle E_0, E_1, \dots \rangle$. Denote $\{h' \in H | Ehh'\}$ by $E(h)$. Call (H, E) an *evaporating graph* iff P for all $h \in H$, $E_{i+1}(h) \subseteq E_i(h)$.

For a finite sequence $\sigma = \langle h_0, h_1, \dots, h_{n-1} \rangle$ let $|\sigma| = n$ and for σ an ω-sequence let $|\sigma| = \omega$.

DEFINITION 7. We say that a finite or ω sequence $h_0, h_1, h_2, \dots$ from H is a *fair path* through the evaporating graph (H, E) iff each $h_{i+1} \in E_i(h_i)$ and each h_{i+1} is either new (not seen so far as h_j for $j \leq i$) or all the elements of $E_i(h_i)$ have appeared before and h_{i+1} is the one which appeared least recently.

LEMMA 21. *Suppose that (H, E) is an evaporating graph and $\pi = \langle x_0, x_1, \dots \rangle$ is a fair path through it. Then*

(i) *$x_{2^{|H|}}$ already occurs in $\{x_i | 0 \leq i < min(2^{|H|}, |\pi|)\}$, and*
(ii) *$ran(\pi) = \{x_i | i < |\pi|\} \subseteq \{x_i | 0 \leq i < min(2^{|H|}, |\pi|)\}$.*

PROOF. (A more reader friendly version courtesy of the referee.) We prove (i) and (ii) by induction on $|H|$. Note that (ii) easily follows from (i) by a subinduction.

To prove (i) let $k = 2^{|H|-1}$, so $2^{|H|} = 2k$. Suppose for contradiction that the claim does not hold: that is, all the elements $x_0, \dots, x_{2k-1}$ are distinct from x_{2k}. It is easy to see that now any subsequence $\langle x_i, x_{i+1}, \dots, x_j \rangle$ with $0 \leq i$ and $j < 2k$ must be a fair path through (H', E') where $H' = H \setminus \{x_i\}$ and each $E'_j(h) = E_j(h) \cap H'$. Hence, from IH(i) on the path $\langle x_{k-1}, \dots, x_{2k-1} \rangle$ it follows that x_{2k-1} occurs earlier on the list, say, as x_a with $k - 1 \leq a \leq 2k - 2$. Now if x_{2k} occurs somewhere in the sequence $\langle x_0, \dots, x_{a+1} \rangle$ we are finished, so suppose otherwise; in particular, this means that $x_{2k} \neq x_{a+1}$. Now consider the sequence $\langle x_0, \dots, x_{a+1} \rangle$ which is a fair path through (H', E'), so by IH(ii) x_{a+1} has already appeared in the sequence $\langle x_0, \dots, x_{k-1} \rangle$. But then it is unfair not to take x_{2k} as x_{a+1}. □

Let x_0 be a point in our limit construction. Suppose that its successors are $x_{n+1} = x_n^+$. Let $h_i = \lambda(x_i)$ for each i. Suppose that the banned list $\boldsymbol{bl}(x_i)$ at x_i is L_i.

We show that, if there is no unbanning between x_0 and x_N — i.e. $L_0 \leq L_1 \leq \cdots \leq L_N$ (prefix inclusion)— and $e_H \leq N$ then there is $i \in \{0, 1, \ldots, e_H\}$ such that $h_i = h_N$.

LEMMA 22. *Suppose that* $\sigma = \langle h_0, h_1, \ldots \rangle$ *is a finite or infinite sequence of hues from* $H(\phi^+)$ *and* $\langle L_0, L_1, \ldots \rangle$ *is a sequence of elements of* $\mathrm{C} \times \{1, \ldots, K\}$ *such that*

1. *each* $h_{i+1} \in Pn(h_i, L_i)$
2. *each* $L_i \leq L_{i+1}$
3. *each* h_{i+1} *is either new or is the least recent member of* $Pn(h_i, L_i)$ *to show up.*

Then $\{h_i | i < |\sigma|\} \subseteq \{h_i | i < e_H\}$.

PROOF. For all i, for all h let $E_i(h) = Pn(h, L_i)$. As $L_i \leq L_{i+1}$, $\vdash \tau_{L_i} \rightarrow \tau_{L_{i+1}}$ and so $E_{i+1}(h) \subseteq E_i(h)$. Thus $(H(\phi^+), E)$ is an evaporating graph.

Clearly $\langle h_0, h_1, \ldots \rangle$ is a fair path through it and so lemma 21 gives us the result. □

§15. Truth lemma. K0 and the truth lemma below gives us our result as (S, R, g), $\boldsymbol{up}(X_0) \models \phi$.

For each $\alpha \in \boldsymbol{cl}\,\phi$, for each $t \in T$,

$$\alpha \in \lambda(t) \text{ iff } (S, R, g), \boldsymbol{up}(t) \models \alpha$$

PROOF. By induction on the construction of α.

p, ($\Rightarrow$): As $\boldsymbol{up}(t)_0$ is just the $\equiv$-class of t, $p \in g(\boldsymbol{up}(t)_0)$.

p, ($\Leftarrow$): So there is some $t' \equiv t$ such that $p \in \lambda(t')$. By K1, K3A and lemma 11, $p \in \lambda(t)$.

$\boldsymbol{true}, \neg\alpha, \alpha \wedge \beta$, ($\Leftrightarrow$): Use the fact that $\lambda(t)$ is maximally consistent in $\boldsymbol{hcl}\,\phi^+$ and the inductive hypothesis.

$X\alpha$, ($\Rightarrow$): By K2 and L1, $\alpha \in \lambda(t^+)$. By the inductive hypothesis, $(S, R, g), \boldsymbol{up}(t^+) \models \alpha$. By definition of $\boldsymbol{up}$ and the semantics of X, $(S, R, g), \boldsymbol{up}(t) \models X\alpha$ as required.

$X\alpha$, ($\Leftarrow$): So $(S, R, g), \boldsymbol{up}(t^+) \models \alpha$. By the inductive hypothesis, $\alpha \in \lambda(t^+)$. By K2, K1 and lemma 10, $X\alpha \in \lambda(t)$ as required.

$\alpha\, U\, \beta$, ($\Rightarrow$): Say that $\alpha\, U\, \beta \in \lambda(t)$. Recall that we have let $t_0 = t$ and defined each t_{i+1} as the immediate $<$-successor of each t_i. We are going to use the fair scheduling idea to show that β must turn up in some $\lambda(t_i)$ with α in all the labels in between.

If $\beta \in \lambda(t)$ then, by IH, $(S, R, g), \boldsymbol{up}(t) \models \beta$ and we have our result.

Otherwise, as $\lambda(t)$ is consistent we have $\alpha \in \lambda(t)$ by axiom C7. Also by K2, $\lambda(t_0) R_X \lambda(t_1)$. Thus there are MCSs Γ and Δ such that $\lambda(t_0) \subseteq \Gamma$, $(\Gamma, \Delta) \in R_X$ and $\lambda(t_1) \subseteq \Delta$. By C7 we must have $\beta \vee (\alpha \wedge X(\alpha\, U\, \beta)) \in \Gamma$. As we have assumed that $\beta \notin \Gamma$, we have $X(\alpha\, U\, \beta)) \in \Gamma$. By definition of R_X, $\alpha\, U\, \beta \in \Delta$. Thus $\alpha\, U\, \beta \in \lambda(t_1)$.

Continuing in this way we either find some $n > 0$ such that $\beta \in \lambda(t_n)$ and for all j, if $0 \leq j < n$ then $\alpha \in \lambda(t_j)$ (which means we are done) or $\alpha, \alpha\, U\, \beta$ and $\neg\beta$ are

in each $\lambda(t_i)$ for $i \geq 0$. We must rule out this latter case. Assume for contradiction that it happens.

If we follow $\boldsymbol{up}(t)$ up high enough then all the points in $\boldsymbol{up}(t)$ above there will be pioneers of their $\equiv$-classes (L2 and P2). Call this the pioneering region of $\boldsymbol{up}(t)$. Eventually, in the pioneering region of $\boldsymbol{up}(t)$ there is a lull from unbanning: this follows from the infinite lulls lemma 19. Choose any one lull totally contained in the pioneering region. We work within the pioneering region because we need to use properties of fair scheduling of hues and in other areas of the structure hues are not necessarily chosen fairly, they may be chosen by the down and across construction.

Say that the banned list at the end of the lull (i.e. after e_H steps) is B. Note that because there is no unbanning during the lull, the banning restrictions become stricter: colours which are allowed by B, so to speak, are allowed all through the lull. The general idea now is to consider a hypothetical construction which continues, after the end of the lull, fairly choosing a sequence of hues under the assumption that the banning list remains fixed as B. We use the fair scheduling idea to conclude that any hue which can ever come up above here must have come up during the lull. Because $\alpha \, U \, \beta$ is a consequence of the hues here (so β should come up in the future) but so is $\neg\beta$ (so β does not come up) we will derive our contradiction.

Say that the lull starts at t_n and so extends to t_{n+e_H}. For $i = n, \ldots, n + e_H$, put $B_i = \boldsymbol{bl}(t_i)$ and $h_i = \lambda(t_i)$. Let $B = B_{n+e_H}$, the banned list at the end of the lull. We have assumed that $\alpha, \alpha \, U \, \beta$ and $\neg\beta$ are in each hue h_i with $n \leq i \leq n + e_H$. We can also conclude (from K4A part 2) that for any such hue h, for any pair (c, p) occurring in B, h is not in V_p.

To derive a contradiction we need to consider the set of hues which would come up infinitely often if we continued forever fair scheduling of hues respecting a fixed banning condition B. So now we see how to recursively choose h_i for each $i \geq n+e_H$ such that $\nvdash \gamma_{h_i^*} \rightarrow E\,\tau_B$. The inductive hypothesis (ih') that we can do so holds for $i = n+e_H$. Assume (ih') true for $i \geq n+e_H$. By lemma 16, $Pn(h_i, B)$ is non-empty. Select $h_{i+1} \in Pn(h_i, B)$ fairly in terms of the previous $h_j (j = n, \ldots, i)$. We also know that h_{i+1} is not in any V_p for p mentioned in B: lemma 22 implies that h_{i+1} has come up during the lull in the actual construction and we have just seen that it is impossible that such a p is mentioned in B. By lemma 17, the inductive hypothesis (ih') holds for $i + 1$.

We put $I = \{h \in H\phi^+ | h = h_i \text{ for infinitely many } i\}$.

Note that by the hue scheduling lemma 22, all the hues in I came up during the lull in the actual construction, i.e. as h_i for some i with $n \leq i \leq n + e_H$. Thus we have

V1) for all $h \in I$, $\alpha, \alpha \, U \, \beta$ and $\neg\beta$ are in h;
V2) for all $h \in I$, h is not in V_p for any p mentioned in B;
V3) for all $h \in I$, $\nvdash \bigwedge h \rightarrow E\,\tau_B$ (from ih' and the fact that $\vdash \gamma_{h^*} \rightarrow E \bigwedge h$).

Let $\theta = \bigvee_{h \in I} \bigwedge h$: so θ says that a point satisfies one of the hues in I. Fair scheduling allows us to conclude that θ must be preserved from point to successor (unless we contravene the B restriction):

LEMMA 23. $\vdash \theta \rightarrow (X\,\theta \vee E\,\tau_B)$.

PROOF. Suppose not, i.e. there is an MCS Γ extending $\theta \wedge \neg E\,\tau_B \wedge \neg X\,\theta$. Let Δ be the MCS containing $\{\delta | X\,\delta \in \Gamma\}$. Let $h_1 = \Gamma \cap \boldsymbol{hcl}\,\phi^+$ and $h_2 = \Delta \cap \boldsymbol{hcl}\,\phi^+$. It

is clear that $h_1 \in I$ but $h_2 \notin I$. We are done if we show that $h_2 \in Pn(h_1, B)$: if h_1 comes up infinitely often then so should h_2 and this would imply that h_2 should be in I.

For contradiction suppose that $\vdash \bigwedge h_1 \wedge X \bigwedge h_2 \rightarrow E\tau_B$. But then Γ would be inconsistent. □

The rest of this argument is a fairly straightforward proof theoretic version of the idea that it is contradictory to have θ holding forever when it implies β is both eventually true and never true. The only complication is the constant assumption associated with the banning condition.

By V2, we have $\vdash \theta \rightarrow \bigwedge_p$ mentioned in $_B \neg v_p$ so that some simple uses of C0-C13 gives us $\vdash (\theta \wedge X E\tau_B) \rightarrow E\tau_B$.

Combining this with lemma 23 gives us

$$\vdash (\theta \wedge \neg E\tau_B) \rightarrow X(\theta \wedge \neg E\tau_B).$$

By generalization and C6, $\vdash \theta \wedge \neg E\tau_B \rightarrow G\theta$.

We have noted in V1 that $\alpha, \alpha U \beta$ and $\neg\beta$ are all in each $h \in I$. Thus each $\vdash \bigwedge h \rightarrow (\alpha \wedge (\alpha U \beta) \wedge \neg\beta)$. Thus $\vdash \theta \rightarrow (\alpha \wedge (\alpha U \beta) \wedge \neg\beta)$. Thus $\vdash (\theta \wedge \neg E\tau_B) \rightarrow ((\alpha U \beta) \wedge G\neg\beta)$. By C8, $\vdash \theta \rightarrow E\tau_B$. Now choose any $h \in I$. Thus $\vdash \bigwedge h \rightarrow \theta$ and so $\vdash \bigwedge h \rightarrow E\tau_B$ contradicting V3.

So we are done.

$\alpha U \beta$, ($\Leftarrow$): Suppose that $(S, R, g), \boldsymbol{up}(t) \models \alpha U \beta$. Say $\boldsymbol{up}(t) = (s_0, s_1, s_2, \dots)$ which are the classes of $t = t_0 < t_1 < t_2 < \dots$ respectively.

So there is $i \geq 0$ such that $(S, R, g), (s_i, s_{i+1}, s_{i+2}, \dots) \models \beta$ and for any j, if $0 \leq j < i$ then $(S, R, g), (s_j, s_{j+1}, s_{j+2}, \dots) \models \alpha$.

By IH, $\beta \in \lambda(t_i)$. If $i = 0$ then we are done. Otherwise, if $0 \leq j < i$ then $\alpha \in \lambda(t_j)$ and without loss of generality we can assume $\neg\beta \in \lambda(t_j)$. Also $\alpha \in \lambda(t_0)$. Suppose for contradiction that $\neg(\alpha U \beta) \in \lambda(t_0)$. By C7, $\neg(\alpha U \beta) \in \lambda(t_1)$. Continuing in this way we show that $\neg(\alpha U \beta) \in \lambda(t_i)$ and $\beta \in \lambda(t_i)$ contrary to C7.

$E\alpha$, ($\Rightarrow$): So $E\alpha \in \lambda(t)$. By the lemma below, there is a hue h' of $\lambda^*(t)$ with $\alpha \in h'$. Say that t'' is the pioneer of the $\equiv$-class of t. So $\lambda^*(t) = \lambda^*(t'')$ and t'' has a sibling, t' say, with $\lambda(t') = h'$. So $t' \equiv t$ and $\alpha \in \lambda(t')$. By IH, $(S, R, g), \boldsymbol{up}(t') \models \alpha$. But $\boldsymbol{up}(t')_0 = \boldsymbol{up}(t)_0$ is the $\equiv$-class of both t and t'. So $(S, R, g), \boldsymbol{up}(t) \models E\alpha$.

LEMMA 24. *If $E\alpha \in \boldsymbol{cl}\,\phi$ and $E\alpha \in h \in H(\phi^+)$ and h is consistent then h^* has a hue h' such that $\alpha \in h'$.*

PROOF. Just extend h to an MCS Γ. Let $\Sigma = \{\alpha\} \cup \{\beta \mid A\beta \in \Gamma\}$. This is consistent by C9-12. Extend Σ to an MCS Δ. Let $h' = \Delta \cap \boldsymbol{hcl}\,\phi^+$. To show $h'^* = h^*$, just consider any $E\delta_b$ for $b \in \wp(\boldsymbol{fcl}\,\phi^+)$: it is clearly in Δ iff it is in Γ. □

$E\alpha$, ($\Leftarrow$): Say $(S, R, g), \boldsymbol{up}(t) \models E\alpha$. Let b be an R-path through S starting at $b_0 = [X_0]$ and going through $b_N = [t]$ such that $(S, R, g), b_{\geq N} \models \alpha$. We are to show that $E\alpha$ is in $\lambda(t)$.

As has been mentioned, emergent paths cause the trouble here because b might be emergent and thus not of the form $\boldsymbol{up}(t')$. If b was $\boldsymbol{up}(t')$ then we could use the inductive hypothesis to immediate effect. If we were attempting a completeness proof for a bundled logic then we could have defined the bundle to contain only paths of the form $\boldsymbol{up}(t')$ and so we would thus be able to conclude the proof.

Instead, we need some other guarantee that we have $E\alpha$ in $\lambda(t)$ and it is here that the automaton, the fresh atoms and the banning procedure are brought together.

Recall that concept of bad fullpaths plays an important role here. These are fullpaths which are accepted by A. First we clarify how a linear automaton recognizing sequences of subsets of $\overline{\boldsymbol{ecl}\,\phi}$ can have a run along a fullpath through the Kripke structure. Define an ω-sequence σ of subsets of $\overline{\boldsymbol{ecl}\,\phi}$ via $\overline{E\delta} \in \sigma_i$ iff there is some $t'_i \in b_i$ such that $E\delta \in \lambda(t'_i)$. We intend A to read the $E\delta$ formulas in the labels along b: by lemma 11, this is well-defined.

We know that $(S, R, g), b_{\geq N} \models \alpha$. Our truth lemma inductive hypothesis is now able to be used to show that $\sigma_{\geq N} \models \overline{\alpha}$. This is partly just the observation that the truth of CTL* formulas along fullpaths is determined by the linear arrangement of truth of $E\beta$ subformulas along the fullpath. We need the inductive hypothesis in the truth lemma to relate truth of these formulas to the contents of labels (which is how σ is defined). That $\sigma_{\geq N} \models \alpha$ follows from:

LEMMA 25. *For all linear temporal combinations β of the subformulas of α, for all i,*

$$\sigma_{\geq i} \models \overline{\beta} \text{ iff } (S, R, g), b_{\geq i} \models \beta.$$

PROOF. By induction on the construction of β. The cases of ***true***, atoms, negation, conjunction, X and U are immediate. This leaves the case of $E\beta$ being a subformula of α.

If $\sigma_{\geq i} \models \overline{E\beta}$ then $E\beta \in \lambda(t')$ for some $t' \in b_i$. By the overall inductive hypothesis, $(S, R, g), \boldsymbol{up}(t') \models E\beta$ and so $(S, R, g), b_{\geq i} \models E\beta$.

Conversely, $(S, R, g), b_{\geq i} \models E\beta$ implies that $(S, R, g), \boldsymbol{up}(t') \models E\beta$ for any $t' \in b_i$; now apply the overall inductive hypothesis (on α). (Note that the inductive hypothesis applies only to paths of the form $\boldsymbol{up}(t)$.) So we have $E\beta \in \lambda(t')$ and by definition, $\sigma_{\geq i} \models \overline{E\beta}$, as required. □

Next, in preparation for recognizing bad paths, we show that the distribution of Q atoms along the path represents the run of the automaton A along b. Consider the run $\langle q_0, q_1, \ldots \rangle$ of A on σ.

LEMMA 26. *For all $i < \omega$, for all $s \in Q$, we have $s \in g(b_i)$ iff $s = q_i$.*

PROOF. We prove this by induction on i. After we prove the converse direction, the forward direction follows from lemma 13 and lemma 11. The case of $i = 0$ follows from K0. Now assume that the hypothesis holds for $i \geq 0$: we show the converse direction for $i + 1$. Thus we are to show that $q_{i+1} \in g(b_{i+1})$.

Let $c = \{\beta \in \boldsymbol{ecl}\,\phi | \overline{\beta} \in \sigma_i\}$. Now choose some $t_i \in b_i$ with $t_i^+ \in b_{i+1}$. By the inductive hypothesis and lemma 11, we will have $q_i \in \lambda(t_i)$. I claim that each conjunct of ζ_c is in $\lambda(t_i)$: lemma 13 will then tell us that, as $q_i \in \lambda(t_i)$, $q_{i+1} = \rho(q_i, \overline{c}) \in \lambda(t_i^+)$. This gives us $q_{i+1} \in g(b_{i+1})$ as required.

To prove this claim we need consider conjuncts of two forms: $E\psi \in c$ and $\neg E\psi$ for $E\psi \in \boldsymbol{ecl}\,\phi \setminus c$.

If $E\psi \in c$ then $\overline{E\psi} \in \sigma_i$. Thus there is some $t' \in b_i$ such that $E\psi \in \lambda(t')$. By lemma 11, $E\psi \in \lambda(t)$.

If $E\psi \in \boldsymbol{ecl}\,\phi \setminus c$ then $\overline{E\psi} \notin \sigma_i$. Thus there is no $t' \in b_i$ with $E\psi \in \lambda(t')$ and, in particular, $E\psi \notin \lambda(t)$. As $\boldsymbol{ecl}\,\phi \subseteq \boldsymbol{fcl}\,\phi^+$, and, $\lambda(t)$ is maximally consistent in $\boldsymbol{fcl}\,\phi^+$, we have $\neg E\psi \in \lambda(t)$ as required. □

From the facts that $\sigma_{\geq N} \models \overline{\alpha}$ and that we were able to ensure that each fullpath, and in particular b is not bad, it follows that $\sigma_{\geq N} \models \overline{E\,\alpha}$:

LEMMA 27. $\sigma_{\geq N} \models \overline{E\,\alpha}$.

PROOF. Suppose not for contradiction. So $\sigma \models \neg\chi_0$ and A accepts σ. Thus there is some $i = 1, \ldots, K$ such that the run $\langle q_0, q_1, \ldots \rangle$ of A on σ contains only a finite number of occurrences of states in V_i but an infinite number of occurences of states in U_i.

Thus, by lemmas 11 and 26, along b, atoms in V_i are in the labels only finitely often while atoms in U_i are in the labels infinitely often. By universal non-acceptance (lemma 20) this can not happen. □

By definition of σ and lemma 11, $E\,\alpha \in \lambda(t)$ and we are done. □

§16. Examples.

1. The following validity is one that can be used as an axiom to capture limit closure in the CTL language:

$$A\,G(p \to E\,X\,p) \to (p \to E\,G\,p).$$

To derive it in our system use the LC axiom

$$A\,G(E\,p \to E\,X((E\,\mathit{false})\,U(E\,p))) \to (E\,p \to E\,G((E\,\mathit{false})\,U(E\,p)))$$

after showing $\vdash p \leftrightarrow E\,p$ and $\vdash p \leftrightarrow ((E\,\mathit{false})\,U((E\,p))$.

2. This example comes from [Stirling, 1992]:

$$\vdash A\,G(A\,\alpha \to E\,X\,F\,A\,\alpha) \to (A\,\alpha \to E\,G\,F\,A\,\alpha).$$

To derive it, simply apply the LC axiom

$$A\,G(E\,A\,\alpha \to E\,X((E\,\mathit{true})\,U(E\,A\,\alpha))) \to (E\,A\,\alpha \to E\,G((E\,\mathit{true})\,U(E\,A\,\alpha)))$$

after showing $\vdash E\,A\,\alpha \leftrightarrow A\,\alpha$.

3. We prove the following using both AA and LC:

$$\vdash (A\,G(p \to E\,X\,r) \wedge A\,G(r \to E\,X\,p)) \to (p \to E\,G(F\,p \wedge F\,r)).$$

If every p place has an r successor and every r place has a p successor then there is a fullpath on which both p and r hold infinitely often. The idea behind the proof we present involves the setting up of an automaton with three states b, c and e. We establish the existence of a path on which it alternates between state b at a place where p holds and state c at a place where r holds. The state e will indicate that we have gone astray.

To use the AA rule we define $L = \{p, r\}$ and $Q = \{b, c, e\}$. The specified set of formulas using atoms only from L contains $\alpha_1 = \neg p \wedge \neg r$, $\alpha_2 = \neg p \wedge r$, $\alpha_3 = p \wedge \neg r$ and $\alpha_4 = p \wedge r$. The function $\rho : (Q \times \{1, 2, 3, 4\}) \to Q$ (capturing the transition table of the automaton) is given by:

ρ	1	2	3	4
b	e	e	c	c
c	e	b	e	b
e	e	e	e	e

The functionally $L + Q$-expandable formula θ is $b \wedge \theta_1 \wedge \theta_2$ where

$$\theta_1 = A\,G \bigwedge_Q \bigwedge_i (q \wedge \alpha_i \rightarrow A\,X\,\rho(q, i))$$

and

$$\theta_2 = A\,G \bigwedge_{q \neq q'} \neg(q \wedge q').$$

Let $\gamma = A\,G(p \rightarrow E\,X\,r) \wedge A\,G(r \rightarrow E\,X\,p)$ and $\alpha = \gamma \wedge \neg e \wedge \theta_1 \wedge (p \vee r)$.

In the full derivation there are long stretches of reasoning within $\vdash_B$ (bundled derivability) including PLTL reasoning. This can be quite subtle (eg., using the induction C6) but we will omit such steps. For our purposes, we just concentrate on the couple of steps where LC and AA are used.

We have $\vdash_B E\,\alpha \rightarrow E\,X\,E\,\alpha$ from which more bundled reasoning gives us

$$\vdash_B A\,G(E\,\alpha \rightarrow E\,X((E\,\textbf{\textit{false}})\,U(E\,\alpha))).$$

The LC axiom (and modus ponens) allows us to conclude

$$\vdash E\,\alpha \rightarrow E\,G((E\,\textbf{\textit{false}})\,U(E\,\alpha))).$$

from which we easily derive $\vdash E\,\alpha \rightarrow E\,G\,E\,\alpha$.

As $\vdash_B (\gamma \wedge \theta \wedge p) \rightarrow E\,\alpha$ and $\vdash_B E\,\alpha \rightarrow \neg e$ we conclude that

$$\vdash \gamma \wedge \theta \wedge p \rightarrow E\,G\neg e \tag{1}$$

i.e. under the given assumptions, there is a fullpath on which the automaton never visits e.

It is interesting to note here that, although the acceptance criteria of the automaton do not play a part in the AA rule, we are bringing them in here: the automaton playing its role in guiding our intuitions in this derivation, accepts exactly paths which never visit e. In this derivation we are looking for a path which is accepted by the automaton.

Some PLTL reasoning (guided by intuitions about the transition function) gives us

$$\vdash_B (G\neg e \wedge \theta) \rightarrow (G(p \rightarrow X\,r) \wedge G(r \rightarrow X\,p))$$

and

$$\vdash_B (p \wedge G(p \rightarrow X\,r) \wedge G(r \rightarrow X\,p)) \rightarrow G(F\,p \wedge F\,r).$$

Bringing (1) in as well gives us

$$\vdash \theta \rightarrow (\gamma \wedge p \rightarrow E\,G(F\,p \wedge F\,r))$$

and so the AA rule gives us

$$\vdash \gamma \wedge p \rightarrow E\,G(F\,p \wedge F\,r)$$

as required.

Note that it is very hard to see how this could be derived without the AA rule. Using LC we could certainly derive that from $\gamma \wedge p$ it follows that there is a fullpath on which p holds infinitely often and on which either p or r holds at each state. However, to deduce that both p and r hold infinitely often seems beyond the capabilities of LC. In order to do away with AA one may propose other forms

of limit closure axioms which can cope with two events being required to recur. However, an infinite set of axioms will probably be needed to cope with any number n of such recurring events.

4. Here is another example using both the AA and LC rules:

let $\gamma = A\,G(p \to E\,X((\neg q)\,U(\neg q \wedge X\,p))) \wedge A\,G(p \to E\,X(q\,U(q \wedge X\,p)))$;
let $\delta = G\,F(\neg q \wedge X\,p) \wedge G\,F(q \wedge X\,p)) \wedge G(((\neg q)\,U\,p) \vee (q\,U\,p)))$;
and we will show $\vdash \gamma \to (p \to E\,\delta)$.

Use $\alpha_1 = \neg p \wedge \neg q$, $\alpha_2 = \neg p \wedge q$, $\alpha_3 = p \wedge \neg q$ and $\alpha_4 = p \wedge q$ along with $Q = \{b_0, \ldots, b_{10}\}$, ρ given by:

ρ	1 $\neg p \wedge \neg q$	2 $\neg p \wedge q$	3 $p \wedge \neg q$	4 $p \wedge q$
b_0	b_3	b_4	b_2	b_1
b_1	b_3	b_4	b_5	b_5
b_2	b_3	b_4	b_2	b_1
b_3	b_3	b_{10}	b_2	b_1
b_4	b_{10}	b_4	b_5	b_5
b_5	b_9	b_8	b_6	b_7
b_6	b_9	b_8	b_0	b_0
b_7	b_9	b_8	b_6	b_7
b_8	b_{10}	b_8	b_6	b_7
b_9	b_9	b_{10}	b_0	b_0
b_{10}	b_{10}	b_{10}	b_{10}	b_{10}

and the corresponding expandable formula $\theta = b_0 \wedge \theta_1 \wedge \theta_2$ analogous to that in the last example.

The automaton which guides our intuition is defined by the transition table above with acceptance determined by b_0 coming up infinitely often. The reader can check that any path accepted by this automaton is a model of δ. To see this note that the initial state b_0 is only reached again after b_5 has been visited and $\neg q$ immediately followed by p holds (i.e. $\neg q \wedge X\,p$ has been true). The state b_5 will only be reached after b_0 when $q \wedge X\,p$ has just held. The states $b_1 - b_4$ record the progress between b_0 and b_5: for example, b_1 indicates that $p \wedge q$ has just been seen while b_4 is encountered during a sequence of $q \wedge \neg p$ states. The states $b_6 - b_9$ are for similar purposes in recording the progress between b_5 and b_0. If $((\neg q)\,U\,p) \vee (q\,U\,p)$ is ever violated then the automaton ends up in the sink state b_{10} (and so does not accept the structure).

Define $\alpha = \gamma \wedge b_0 \wedge \theta_1 \wedge \theta_2$. We can show that when α holds then it holds again at some later state: $\vdash_B A\,G(E\,\alpha \to E\,X((E\,\textbf{\textit{true}})\,U(E\,\alpha))$. The LC axiom allows us to deduce that then there is a fullpath on which α holds infinitely often: $\vdash E\,\alpha \to E\,G((E\,\textbf{\textit{true}})\,U(E\,\alpha))$.

Now $E\,X\,E\,\alpha$ holds at the root: $\vdash_B \gamma \wedge p \wedge \theta \to E\,X\,E\,\alpha$. Any path of recurring $E\,\alpha$ is accepted by the automaton and so satisfies δ: $\vdash_B G((E\,\textbf{\textit{true}})\,U(E\,\alpha)) \to \delta$. Putting these together we deduce that there is a fullpath satisfying δ: $\vdash \theta \to (\gamma \to (p \to E\,\delta))$.

The AA rule finishes the derivation.

Again, this example would be a good test of any alternative axiom systems. It is not good enough here just to establish that there is a fullpath on which either $E(q\ U(q \wedge X\ p))$ or $E((\neg q)\ U(\neg q \wedge X\ p))$ holds at every state and that both these formulas hold infinitely often. Instead, it is required that either $q\ U(q \wedge X\ p)$ or $(\neg q)\ U(\neg q \wedge X\ p)$ holds along the chosen fullpath at every state (and each holds infinitely often).

§17. Conclusion. We have been able to give a simple axiomatization of validity in standard CTL*.

Interesting aspects of the system include a very intuitive limit closure axiom and a slightly complicated, but possibly more generally useful, rule for the systematic introduction of fresh atoms into a proof.

Interesting aspects of the completeness proof include the use of a linear Rabin automaton and the use of a new banning mechanism working alongside a strict variant of the usual scheduling mechanism in the vaguely filtration-based construction.

The result and proof suggest several avenues for future work. The most important question regarding this axiomatization is whether the auxiliary atoms rule is really needed. Even if the rule is not needed, and especially if it is, there is the possibility of using it to good effect in other similar logics such as other branching time logics, including those from philosophical logic, or in more general areas of modal and temporal logic reasoning. Interesting examples are the CTL* logic with past operators from [Zanardo and Carmo, 1993] as well as the long-unaxiomatized logic of historical necessity.

REFERENCES

[Bernholtz and Grumberg, 1994] O. Bernholtz and O. Grumberg, *Buy one, get one free!!!*, ***Temporal logic, Proceedings of ICTL'94*** (D. Gabbay and H. Ohlbach, editors), LNAI, no. 827, Springer-Verlag, 1994, pp. 210–224.

[Burgess, 19➡ J. P. Burgess, *Decidability for branching time*, ***Studia Logica***, vol. 39 (1980), pp. 203–218.

[Clarke and Emerson, 1981] E. Clarke and E. Emerson, *Synthesis of synchronization skeletons for branching time temporal logic*, ***Proc. IBM Workshop on Logic of Programs, Yorktown Heights, NY*** (Berlin), Springer, 1981, pp. 52–71.

[Dam, 1992] M. Dam, *R-generability, and definability in branching time logics*, ***Information Processing Letters***, vol. 41 (1992), pp. 281–287.

[Emerson, 1983] E. Emerson, *Alternative semantics for temporal logics*, ***Theoretical Computer Science***, vol. 26 (1983).

[Emerson, 1996] ———, *Automated temporal reasoning for reactive systems*, ***Logics for concurrency*** (F. Moller and G. Birtwistle, editors), Springer Verlag, 1996, pp. 41–101.

[Emerson and Halpern, 1982] E. Emerson and J. Halpern, *Decision procedures and expressiveness in the temporal logic of branching time*, ***Proc. 14th ACM Symp. on Theory of Computing***, 1982.

[Emerson and Halpern, 1986] ———, *'Sometimes' and 'not never' revisited: on branching versus linear time*, ***Journal of the ACM***, vol. 33 (1986).

[Emerson and Jutla, 1988] E. Emerson and C. Jutla, *Complexity of tree automata and modal logics of programs*, ***29th IEEE Foundations of Computer Science, Proceedings***, IEEE, 1988.

[Emerson and Sistla, 1984] E. Emerson and A. Sistla, *Deciding full branching time logic*, ***Information and Control***, vol. 61 (1984), pp. 175–201.

[Emerson, 1990] E. A. Emerson, *Temporal and modal logic*, ***Handbook of theoretical computer science*** (J. van Leeuwen, editor), vol. B, Elsevier, Amsterdam, 1990.

[Gabbay, Hodkinson, and Reynolds, 1994] D. GABBAY, I. HODKINSON, and M. REYNOLDS, ***Temporal logic: Mathematical foundations and computational aspects***, vol. 1, Oxford University Press, 1994.

[Gabbay, 1981] D. M. GABBAY, *An irreflexivity lemma with applications to axiomatizations of conditions on tense frames*, ***Aspects of philosophical logic*** (U. Monnich, editor), Reidel, Dordrecht, 1981, pp. 67–89.

[Gabbay, Pnueli, Shelah, and Stavi, 1980] D. M. GABBAY, A. PNUELI, S. SHELAH, and J. STAVI, *On the temporal analysis of fairness*, ***7th ACM Symposium on Principles of Programming Languages, Las Vegas***, 1980, pp. 163–173.

[Kaivola, 1996] R. KAIVOLA, *Axiomatising extended computation tree logic, in trees in algebra a programming*, ***CAAP'96, 21st International Colloquium, Proceedings***, vol. 1059, Springer, 1996, pp. 87–101.

[Kesten and Pnueli, 1995] YONIT KESTEN and AMIR PNUELI, *A complete proof systems for QPTL*, ***Proceedings, Tenth Annual IEEE Symposium on Logic in Computer Science*** (San Diego, California), IEEE Computer Society Press, 26-29 June 1995, pp. 2–12.

[McNaughton, 1966] R. MCNAUGHTON, *Testing and generating infinite sequences by finite automata*, ***Information and Control***, vol. 9 (1966), pp. 521–530.

[Pnueli, 1977] A. PNUELI, *The temporal logic of programs*, ***Proceedings of the Eighteenth Symposium on Foundations of Computer Science*** (Providence, RI), 1977, pp. 46–57.

[Safra, 1988] S. SAFRA, *On the complexity of ω-automata*, ***Proceedings of 29th IEEE Symposium on the Foundations of Computer Science***, 1988.

[Stirling, 1992] C. STIRLING, *Modal and temporal logics*, ***Handbook of Logic in Computer Science, Volume 2*** (S. Abramsky, D. Gabbay, and T. Maibaum, editors), OUP, 1992, pp. 477–563.

[Thomason, 1984] R. THOMASON, *Combinations of tense and modality*, ***Handbook of philosophical logic, Vol II: Extensions of classical logic*** (D. Gabbay and F. Guenthner, editors), Reidel, Dordrecht, 1984, pp. 135–165.

[Vardi and Stockmeyer, 1985] M. VARDI and L. STOCKMEYER, *Improved upper and lower bounds for modal logics of programs*, ***17th ACM Symp. on Theory of Computing, Proceedings***, ACM, 1985, pp. 240–251.

[Walukiewicz, 1995] I. WALUKIEWICZ, *A complete deductive system for the μ-calculus*, ***BRICS Research Report RS-95-6***, Department of Computer Science, University of Aarhus, Denmark, 1995.

[Zanardo, 1985] A. ZANARDO, *A finite axiomatization of the set of strongly valid Ockamist formulas*, ***Journal of Philosophical Logic***, vol. 14 (1985), pp. 447–468.

[Zanardo, 1996] ———, *Branching-time logic with quantification over branches: the point of view of modal logic*, this JOURNAL, vol. 61 (1996), pp. 1–39.

[Zanardo, Barcellan, and Reynolds, 1999] A. ZANARDO, B. BARCELLAN, and M. REYNOLDS, *Non-definability of the class of complete bundled trees*, ***Logic Journal of the IGPL***, vol. 7 (1999), no. 1, pp. 125–136.

[Zanardo and Carmo, 1993] ALBERTO ZANARDO and JOSÉ CARMO, *Ockhamist computational logic: Past-sensitive necessitation in CTL*, ***Journal of Logic and Computation***, vol. 3 (1993), no. 3, pp. 249–268.

SCHOOL OF INFORMATION TECHNOLOGY
MURDOCH UNIVERSITY
SOUTH STREET
PERTH, WESTERN AUSTRALIA 6150
E-mail: m.reynolds@murdoch.edu.au