17

An Axiomatization of Prior's Ockhamist Logic of Historical Necessity

MARK REYNOLDS

ABSTRACT. We present a complete axiomatization of Prior's Ockhamist branching time logic of historical necessity. This solves a long standing open problem.

1 Introduction

Logics of *historical necessity*, the most natural combination of tense with the modality of necessity and possibility, have a long and venerable history of development, especially in connection with arguments of human freedom versus determinism. Their variations and close relatives (such as the discrete branching time logics CTL[5], CTL*[8] and PCTL*[14, 18]), also have important recent applications including many in AI, agent reasoning and robotics [25, 1, 17], software engineering, reactive systems, and data communication protocols, [7, 21], and natural language understanding and philosophy[15].

These logics allow for reasoning about processes over time when the development is not necessarily determined. To achieve this, indeterminist time is usually modelled by a tree with each point having one linear past but branching towards the future. These future branches represent several alternative possibilities. The unique past represents the necessity of history: it can not be changed and there are no choices in that direction. Thus such logics are often termed logics of historical necessity.

Arthur Prior [16] recognized two main alternative approaches to historical necessity. The Peircean approach only allows statements of the form "p will occur" to be true at a point of indeterminist time if on all possible futures, p occurs at some time. In this paper we will concentrate on the other approach. This is named the Ockhamist after William of Ockham who is supposed to have held that it is meaningful to make statements whose truth is not yet known to us. The idea is that the truth of statements is evaluated with respect to a point on a particular whole linear history through time. The Ockhamist logic is the more expressive [20].

Advances in Modal Logic, Volume 4, 355–370.

Prior started to formalize these ideas in logics with the standard temporal G and H connectives in combination with a modality which quantifies over all histories through the point of evaluation. We might call this the *alethic* modality and we symbolize it by $\Box$. The semantics of Prior's Ockhamist (and Peircean) logic were provided in [3] where we also find the definition of a related "bundled" version of the logic. The bundled Ockhamist logic simply restricts evaluation and modal quantification to within a certain specified and fixed set or "bundle" of histories. Since then there has been much work in the area. Complete axiom systems have been presented for many branching time logics: CTL [6], Prior's Peircean logic [4] (or in [27] without a Gabbay-style IRR-rule), bundled CTL* [21], the bundled version of Prior's Ockhamist logic [26] (although an unpublished IRR-rule based axiomatization by Gabbay has been cited), bundled Ockhamist logic with "Until" and "Since" [28], CTL* [19], and PCTL* [18].

The decidability of many of these logics follows from the decidability of the second order monadic logic of trees [12]. Complexity results have been established for some of the variations.

This is an active area of research. Recent related work includes development of hybrid branching time logics [2], calculi [24] for first-order versions, decidability of fragments of the first-order versions [13], and decision procedures and axiomatizations for the versions allowing quantification over propositional variables[9].

Despite this effort and interest, technical difficulties have left the presentation of a sound and complete axiom system for Prior's most basic original Ockhamist logic of historical necessity as an open problem. In section 6 below we present a complete axiom system for the logic.

This conference paper also gives a brief sketchy overview of the long, complex and quite interesting completeness proof. The proof starts off in a traditional way but brings in new "banning" ideas from [19] and an unusual trans-countable construction. A full version of the proof (of over 100 pages) is in preparation.

Future work will include the extension of the technique to the language with "until" and "since" connectives, to combinations with epistemic logics, to the first-order version and to that including quantification of propositions. There is also need for the development of applications in reasoning about agents and about complex reactive systems. We will also try determining the complexity of decision procedures, and implementing efficient decision procedures.

2 The logic

A *tree* is a pair $(T, <)$ where T is a non-empty set, $<$ is a binary relation on T such that:

1. $<$ is irreflexive and transitive;
2. $<$ is past-linear, i.e. ,
   if $y < x$ and $z < x$ then either $y < z$, $z < y$ or $y = z$; and
3. $(T, <)$ is connected, i.e. ,
   for all $x, y \in T$, there is $z \in T$ with $z \leq x$ and $z \leq y$.

A *history* of $(T, <)$ is a maximal linearly $<$-ordered subset of $(T, <)$. Let $\mathbf{B}(T, <)$ be the set of all histories of $(T, <)$. A set of histories $B \subseteq \mathbf{B}(T, <)$ is called a *bundle* if for every $t \in T$ there is some $b \in B$ with $t \in b$. If B is a bundle on $(T, <)$ then $(T, <, B)$ is called a *bundled frame*. $(T, <, B)$ is *complete* iff $B = \mathbf{B}(T, <)$.

Fix a countable set $\mathcal{L}$ of atoms. *Bundled tree structures* $\mathcal{T} = (T, <, B, h)$ will have a bundled tree frame $(T, <, B)$ and a valuation h for the atoms i.e. for each atom $p \in \mathcal{L}$, $h(p) \subseteq T$. We have an *Ockhamist structure* iff $(T, <, B)$ is complete.

The set of formulas contains the atoms and recursively includes $\neg\alpha$, $\alpha \wedge \beta$, $G\alpha$, $H\alpha$ and $\Box\alpha$.

Formulas are evaluated at points on bundle histories in bundled tree structures:

$\mathcal{T}, \sigma, x \models p$	iff	$x \in h(p)$, for p atomic;
$\mathcal{T}, \sigma, x \models \neg\alpha$	iff	$\mathcal{T}, \sigma, x \not\models \alpha$;
$\mathcal{T}, \sigma, x \models \alpha \wedge \beta$	iff	both $\mathcal{T}, \sigma, x \models \alpha$ and $\mathcal{T}, \sigma, x \models \beta$;
$\mathcal{T}, \sigma, x \models G\alpha$	iff	for all $y > x$ in σ we have $\mathcal{T}, \sigma, y \models \alpha$;
$\mathcal{T}, \sigma, x \models H\alpha$	iff	for all $y < x$ in σ we have $\mathcal{T}, \sigma, y \models \alpha$;
$\mathcal{T}, \sigma, x \models \Box\alpha$	iff	for every history $\pi \in B$ containing x we have $\mathcal{T}, \pi, x \models \alpha$.

As well as the usual classical abbreviations, we have linear temporal ones $F\alpha \equiv \neg G \neg \alpha$, $P\alpha \equiv \neg H \neg \alpha$, $F_{\leq}\alpha \equiv \alpha \vee F\alpha$, $G_{\leq}\alpha \equiv \alpha \wedge G\alpha$, $P_{\leq}\alpha \equiv \alpha \vee P\alpha$ and $H_{\leq}\alpha \equiv \alpha \wedge H\alpha$. There is also the branching modal diamond, $\Diamond\alpha \equiv \neg\Box\neg\alpha$.

We say that α is *valid* in the bundled logic (and we write $\models_B \alpha$) iff for all bundled tree structures $\mathcal{T} = (T, <, B, h)$, for all histories $\sigma \in B$, for all points $x \in \sigma$, we have $\mathcal{T}, \sigma, x \models \alpha$.

Prior's Ockhamist logic is the logic of complete bundled structures. We write $(T, <, h), \sigma, x \models \alpha$ for $(T, <, \mathbf{B}(T, <), h), \sigma, x \models \alpha$. We say that α is *valid* in Prior's Ockhamist logic (and we write $\models_O \alpha$) iff for all complete bundled structures $\mathcal{T}$, for all histories σ in $\mathcal{T}$, for all points $x \in \sigma$, we have $\mathcal{T}, \sigma, x \models \alpha$.

It is worth noting that for us formulas are evaluated at points on histories but the valuation or assignment of atoms is determined only by the point. Thus the truth of an atom agrees at two histories through the same point. We can justify this prescription by assuming that atomic propositions do not contain any trace of futurity: an atom is either true now or false now regardless of how the future is to evolve. This assumption is arguable. Prior discusses the issue in [16] on page 124, and describes two different logics: one using two sorts of atomic propositions and the other (on page 126) being the one we look at, having only one sort of proposition with a single truth value at each point. Our approach is also taken in [22]. The alternative approach, with assignment depending on point and history is taken in [3] and [29] which paper also contains further interesting discussion on the question. In the completeness proof the non-futurity asssumption is relied on at various important places (see discussion later).

3 A system for bundled validity

An axiom system $\vdash_B$ for bundled validity will form part of our final axiom system. The inference rules are modus ponens, temporal and branching generalization, a Gabbay-style IRR-rule:

$$\frac{\alpha, \alpha \to \beta}{\beta} \quad \frac{\alpha}{G\alpha} \quad \frac{\alpha}{H\alpha} \quad \frac{\alpha}{\Box\alpha} \quad \frac{(p \wedge H\neg p) \to \alpha}{\alpha} \text{ if } p \text{ does not appear in } \alpha,$$

and a special rule (atomic non-futurity) saying that the truth of propositional atoms only depends on the present:

ANF $p \to \Box p$, for each atomic proposition p.

The axiom schemes include all substitution instances of propositional tautologies along with all instances of the following axioms.

Those for general linear temporal logic from the complete system given by [3]:

L1 $G(\alpha \to \beta) \to (G\alpha \to G\beta)$ and mirror image;
L2 $G\alpha \to GG\alpha$
L3 $\alpha \to GP\alpha$ and mirror image;
L4 $(F\alpha \wedge F\beta) \to (F(\alpha \wedge \beta) \vee F(\alpha \wedge F\beta) \vee F(\beta \wedge F\alpha))$
and mirror image.

We add axioms ensuring that the history modality $\Box$ behaves as in the modal logic S5:

BK $\Box(\alpha \to \beta) \to (\Box\alpha \to \Box\beta)$,
BT $\Box\alpha \to \alpha$,
BE $\Diamond\alpha \to \Box\Diamond\alpha$.

We need just one axiom to specify the interaction between modalities:

HN $P\alpha \to \Box P \Diamond \alpha$.

Finally, we require maximality of histories:
MB $G\bot \to \Box G\bot$.
Note that we do not use a substitution rule as it is not valid for $\models_B$. For example, $\Diamond p \to p$ is valid for each atomic proposition p (and can be derived easily) but $\Diamond \alpha \to \alpha$ is not generally valid: it is easy to construct a model of $\Diamond Fp \wedge \neg Fp$.

4 Bundled completeness and Kamp frames

By following the general idea of proofs in [11], [26], [10] and [29], we could show that

LEMMA 1 $\vdash_B$ *is sound and complete for* $\models_B$.

The proof is based on the traditional step by step construction of a perfect chronicle attributed to Kripke and Lemmon. A chronicle is a frame with points labelled by maximally consistent sets of formulas. The chronicle is perfect if the labels respect the accessibility relations between points in terms of both box and diamond formulas.

In the Kripke-Lemmon style axiomatic completeness proof for the linear temporal logic (see [3]), a chronicle is built on a linear frame. Each point is labelled with a maximal consistent set of formulas and we eventually ensure that the final chronicle is perfect. We can define a valuation on this frame by checking for the existence of an atom (rather than its negation) inside the label at each point. The perfection of the chronicle allows us to prove a "truth lemma" showing the equivalence of the truth of a formula at a point in the model to the appearance of that formula in the label at that point.

In the branching case there is an immediate problem with this basic strategy. Truth of formulas is not evaluated at points in a tree frame but instead at history-point pairs. Thus our chronicle will need to allow several different labels at each point of the tree. Fortunately, this problem has been neatly solved by the use of what is now called Kamp frames which were invented by Hans Kamp in a 1979 unpublished manuscript.

There are variations on this idea ([22], [23], [29]) but the general idea is that we make enough copies of each tree point and separate the various histories in the tree. We end up with a vaguely grid-like structure with the temporal $<$-relation increasing vertically and an equivalence relation $\equiv$ relating various nodes on a horizontal level. A whole $\equiv$-class of Kamp frame points correspond to just one point of a tree. Details can be found in [20]; see figures 1 and 2 for an example.

In our proof we use our own version of the idea:

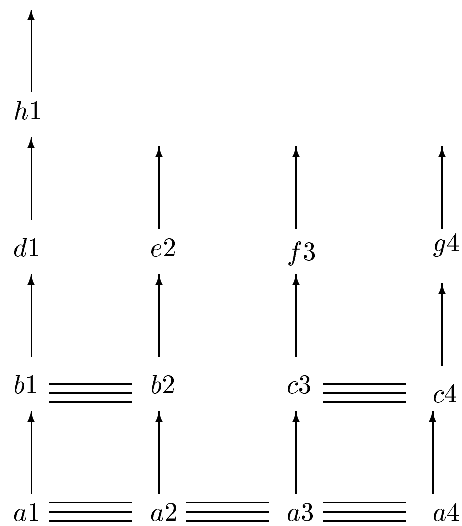

Figure 1. A Kamp frame

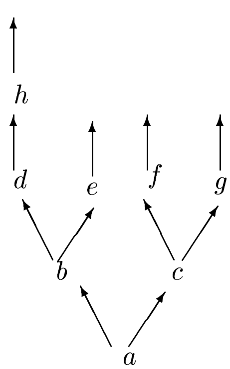

Figure 2. The corresponding tree

DEFINITION 2 A *Kamp* frame is $(T, <, \equiv)$ where:

KF1) $<$ is a transitive, irreflexive linear relation on T

KF2) $\equiv$ is an equivalence relation such that:
KF2.1) if $x \equiv y$ then we do not have $x < y$; and
KF2.2) if $x \equiv y$ and $u < x$ then there is $v < y$ such that $u \equiv v$;

KF3) the structure is connected, i.e.
for all $x, y \in T$ there is $u \leq x$ and $v \leq y$ such that $u \equiv v$; and

KF4) the columns are separated, i.e.
for each $x \equiv y$ with $x \neq y$ there exists $v > y$ such that there is no $u > x$ with $u \equiv v$.

A chronicle for the bundled logic is a Kamp frame with points labelled by maximally consistent sets of formulas. As we build the chronicle, starting from a one point one, we ensure that the labels of points respect the three box modalities, G, H and $\Box$. Thus, if $x < y$ and $G\alpha$ is in the label of x then α is in the label of y. Similarly, if $x \equiv y$ and $\Box\alpha$ is in the label of x then α is in the label of y.

Building a perfect chronicle requires us to reach, in the limit of the step by step construction, a structure in which the three modal diamonds F , P and $\Diamond$ are respected as well. Thus, for example, if $F\alpha$ is in the label of x then we require that there is a point $y > x$ with α in the label of y. If, at some finite stage, $F\alpha$ is in the label of x but there is no point $y > x$ with α in the label of y then we say that $(x, F\alpha)$ is a *defect*. One step in the construction consists in curing a defect by extending the chronicle slightly to add the needed y.

By enumerating the defects and being fair in curing them we can ensure that the limit is perfect. The Kamp frame corresponds directly with a bundled tree structure via a process of collapsing $\equiv$-classes (see [20] for details) and the truth lemma follows.

Note that the IRR rule is used in the construction to provide "names" for labels and hence to make it easy to match up labels around a square $x > y \equiv y' < x' \equiv x$. There is a complicated axiom in [26] which provides an alternative to using the IRR rule to axiomatize this logic.

5 Limit closure

The example $\gamma = \Box G(p \to \Diamond Fp) \to \Diamond G(p \to Fp)$ from [20] shows that $\models_B$ and $\models_O$ are distinct. A simple transfinite induction shows that $\models_O \gamma$. But now consider $T = \{(n, m) \in \mathbb{N} \times \mathbb{N} | 0 < n \leq m \text{ or } n = m = 0\}$. 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 the bundle to be $B = B(T, <) \setminus \{D\}$. Finally, define g so that $(n, m) \in g(p)$ iff $n = m$. We have $(T, <, B, g), (0, 0) \models \neg\gamma$. Thus $\not\models_B \gamma$.

There is a discussion of this difference in [20] in which it is related to a similar observation concerning CTL*. The difference in validity between bundled and complete versions of branching logics is connected with what is sometimes called the *limit closure* property of the complete structures: the complete bundle is closed under taking the (in our case transfinite) limit of a process of keeping on extending initial segments of histories towards the future.

It follows that the system for $\vdash_B$ is incomplete for $\models_O$. It is interesting to see where the completeness proof for $\vdash_B$ as above goes wrong for $\models_O$. In the limit of the step by step construction of a perfect Kamp frame chronicle we only construct a countable number of vertical lines (or columns) which correspond to histories when we factor out by $\equiv$. However, the resulting tree generally also has an uncountable number of other histories which do not correspond to columns. These may be called "emergent" histories and it is they which generally prevent the truth lemma from holding: they do not even have any labels constructed for them. More details of this problem can be found in [20].

Despite being incomplete the system $\vdash_B$ is, of course, sound for $\models_O$. There have been many attempts to "find the missing axiom". Here we are able to introduce a new axiom schema to allow us to derive the extra validities.

We solve the limit closure problem by the addition of what we call a limit closure schema. It is an infinite sequence of axioms: one for each $n > 0$. Suppose $\alpha_0, ..., \alpha_{n-1}$ are any formulas. Put $\alpha_n = \alpha_0$. The schema allows

LC $\quad \Box G_{\leq} \bigwedge_{i=0}^{n-1} (\Diamond\alpha_i \rightarrow \Diamond F \Diamond \alpha_{i+1}) \quad \rightarrow \quad \Diamond G_{\leq} \bigwedge_{i=0}^{n-1} (\Diamond\alpha_i \rightarrow F \Diamond \alpha_{i+1})$.

LEMMA 3 *LC is sound.*

Proof. The soundness proof is a straightforward transfinite induction. Basically, we cycle through the indicies gradually building an increasing sequence of points in the tree where the $\Diamond\alpha_i$ hold.

To proceed more formally suppose t is a point in a history b of a structure $(T, <, h)$ such that $(T, <, h), b, t \models \beta$ where $\beta = \Box G_{\leq} \bigwedge_{i=0}^{n-1} (\Diamond\alpha_i \rightarrow \Diamond F \Diamond \alpha_{i+1})$.

We observe the following. If there is some $s \geq t$, some history $b' \ni s$ and some $i = 0, 1, ..., n-1$, such that $(T, <, h), b', s \models \Diamond\alpha_i$ then there is a point $s' > s \geq t$ and a history $b'' \ni s'$ such that $(T, <, h), b'', s' \models \Diamond\alpha_{i+1}$.

Recalling that $\alpha_n = \alpha_0$, a short induction thus tells us that if there is some $s \geq t$, some history $b' \ni s$ and some $i = 0, 1, ..., n-1$, such that $(T, <, h), b', s \models \Diamond\alpha_i$ then there is a point $s' > s \geq t$ and a history $b'' \ni s'$ such that $(T, <, h), b'', s' \models \Diamond\alpha_0$.

If there is no $s \geq t$ with some history $b' \ni s$ and some $i = 0, 1, ..., n-1$, such that $(T, <, h), b', s \models \Diamond\alpha_i$ then clearly $(T, <, h), b, t \models G_{\leq} \bigwedge_{i=0}^{n-1}(\Diamond\alpha_i \rightarrow F\Diamond\alpha_{i+1})$ and we are done.

So now suppose that there is some $s \geq t$, some history $b' \ni s$ and some $i = 0, 1, ..., n-1$, such that $(T, <, h), b', s \models \Diamond\alpha_i$.

We find by transfinite induction (using the observation above) an ordinal Λ and a possibly transfinite sequence $\{t_{\lambda,i} | \lambda < \Lambda, i = 0, 1, ..., n-1\}$ of points and some histories $b_{\lambda,j}$ such that $t \leq t_{0,0}$ and for each $\lambda < \Lambda$, for each $i = 0, 1, 2, ..., n-1$, we have $t_{\lambda,i} \in b_{\lambda,i}$ and $(T, <, h), b, t_{\lambda,i} \models \Diamond\alpha_i$. We also ensure that if $\lambda < \mu < \Lambda$ then $t_{\lambda,i} < t_{\mu,j}$ and if $i < j < n$ then $t_{\lambda,i} < t_{\lambda,j}$.

The case of a successor ordinal is straightforward. Consider a limit ordinal λ such that we have defined $t_{\mu,i}$ for each $\mu < \lambda$ and each $i = 0, 1, ..., n-1$. If there is no history b', point $s \in b'$ and $i = 0, 1, ..., n-1$ such that each $t_{\mu,i} < s$ and $(T, <, h), b', s \models \Diamond\alpha_i$ then we let $\Lambda = \lambda$ and we are finished the construction. If there are such b', s and i then we know that there is also $t_{\lambda,0} > s$ and $b_{\lambda,0} \ni s$ such that $(T, <, h), b_{\lambda,0}, t_{\lambda,0} \models \Diamond\alpha_0$. This continues our construction.

Now suppose we have finished our construction at Λ: the construction can obviously not continue forever. Let

$$\pi = \{x \in T | \text{ there is some } \lambda, i \text{ such that } x \leq t_{\lambda,i}\}.$$

This is clearly linearly ordered by $<$. Thus there is a history $\sigma \supseteq \pi$.

We have $(T, <, h), \sigma, t \models G_{\leq} \bigwedge_{i=0}^{n-1}(\Diamond\alpha_i \rightarrow F\Diamond\alpha_{i+1})$.

As $t \leq s = t_{0,0}$ we have $t \in \pi \subseteq \sigma$ and so

$$(T, <, h), b, t \models \Diamond G_{\leq} \bigwedge_{i=0}^{n-1}(\Diamond\alpha_i \rightarrow F\Diamond\alpha_{i+1})$$

as required. ■

6 The axiom system for Ockamist validity

Let $\vdash_H$ be derivability in the axiom system with the schema LC added to $\vdash_B$.

THEOREM 4 $\vdash_H$ *is sound and complete for the logic* **HN** *of historical necessity.*

Soundness is the usual induction. To show completeness we have fixed a formula ϕ which is supposed to be $\vdash_H$-consistent, i.e. $\nvdash_H \neg\phi$. We are to show that it has a model. Most of the rest of the paper sketches an intricate construction of a model of ϕ: the full version is in preparation.

7 Overview of the completeness proof

The completeness proof proceeds in two stages. The first stage is based on the traditional step by step construction of a perfect Kamp frame chronicle described above. Stage two, which we return to briefly below, is a transfinite process of what we call grafting, copying the latter part of branches after "matching ends".

The solution by matching ends

The aim of stage one is to make a perfect Kamp frame chronicle in which the ends of emergent branches match initial segments of constructed branches in terms of the "colours" of labels (see below) which appear cofinally. In order to achieve the matching in stage one we use a banning technique along the lines of that used in [19]. To do this requires finite sets as labels, the IRR rule and the new LC rule.

Finite labels: hues and colours

During our construction we are really only interested in making a chronicle which is relatively perfect in respecting the finite set of subformulas of ϕ (and a few other formulas): this is enough to establish the truth lemma. This allows us to join parts of the chronicle together at points which have labels which agree on those formulas but maybe not on others. We say that the labels are of the same *hue* (defined below). Our banning procedure needs to do this and it also needs to control which hues recur along branches during the construction.

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, i.e. boolean combinations of those of the form $\Box\beta$.

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 histories in the limit of our construction.

The first of several closure sets we use is $\mathbf{cl}\phi = \{\psi, \neg\psi | \psi \leq \phi\}$ where $\psi \leq \phi$ denotes that ψ is a subformula of ϕ. Let

$$\mathbf{fcl}\phi = \{\psi, \neg\psi, \Diamond\psi, \neg\Diamond\psi, \Diamond\neg\psi, \neg\Diamond\neg\psi | \psi \leq \phi\}.$$

For $a \in \wp(\mathbf{fcl}\phi)$, define

$$\delta_a = \bigwedge_{\delta \in a} \delta \wedge \bigwedge_{\delta \in (\mathbf{fcl}(\phi) \setminus a)} \neg\delta.$$

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

$$\gamma_c = \bigwedge_{a \in c} \Diamond\delta_a \wedge \bigwedge_{a \in (\wp(\mathbf{fcl}\phi) \setminus c)} \neg\Diamond\delta_a.$$

Given a colour $c \in C$ 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

$$h(a,c) = \{\delta | \delta \in a\} \cup \{\neg\delta | \delta \in \mathbf{fcl}\phi \setminus a\} \cup \{\Diamond\delta_b | b \in c\} \cup \{\neg\Diamond\delta_b | b \in \wp(\mathbf{fcl}\phi) \setminus c\}.$$

The set $H(\phi)$ of all hues of ϕ is $H(\phi) = \{h(a,c) | a \in c \in C\}$. Each hue h is a hue of exactly one colour h^*.

Maximal consistent sets (MCSs) extend (consistent) hues: Let $\mathbf{hcl}\phi = \{\delta, \neg\delta | \delta \in \mathbf{fcl}\phi\} \cup \{\Diamond\delta_b, \neg\Diamond\delta_b | b \in \wp(\mathbf{fcl}\phi)\}$. Then for each MCS Δ, the set $h = \Delta \cap \mathbf{hcl}\phi$ is a hue. Furthermore, for each Δ this is the only hue which satisfies $h \subseteq \Delta$.

We define two useful relations on $H(\phi)$ as well as the traditional ones between MCSs. $\Gamma R_G \Delta$ iff for all α, if $G\alpha \in \Gamma$ then $\alpha \in \Delta$; and $\Gamma R_N \Delta$ iff for all α, if $\Box\alpha \in \Gamma$ then $\alpha \in \Delta$. Say that $hR_G h'$ iff there are MCSs Γ and Γ' such that $h \subseteq \Gamma$, $h' \subseteq \Gamma'$ and $\Gamma R_G \Gamma'$. Say that $hR_N h'$ iff there are MCSs Γ and Γ' such that $h \subseteq \Gamma$, $h' \subseteq \Gamma'$ and $\Gamma R_N \Gamma'$.

In our chronicle, points $x \equiv y$ have labels which extend hues of the same colour and points $x < y$ have respective labels Γ_x, Γ_y with $\Gamma_x R_G \Gamma_y$.

IRMCSs

Here we use the IRR rule for two reasons: matching squares $x > y \equiv y' < x' \equiv x$ as in the bundled case and as part of the banning procedure which we describe below. We just use a slight extension of the idea of IRR-theories as given in definition 6.2.3 of [10]: we allow $\Diamond$ as well as F and P. Points will be labelled with these what we call IRMCS sets which each have a unique name in that they contain some formula $p \wedge H\neg p$. Many useful results follow: eg, if $F\alpha$ is in IRMCS Γ then there is an IRMCS Δ with $\Gamma R_G \Delta \ni \alpha$.

The step by step

There is no space for a full definition of a chronicle but it is essentially a Kamp frame with each point labelled by two IRMCS and a ban and a banning query as defined shortly.

One of the IRMCS labels at each point is thought of as facing downwards, the other upwards. The two IRMCS at each point agree on hues, i.e. on formulas in $\mathbf{hcl}(\phi)$. In a traditional chronicle we would only have one label at a point but having two allows us to join together pieces of different (traditional) chronicles provided they agree enough at the joins on the finite set of formulas of interest to us. This is important to the banning process. At most points the two labels agree exactly.

At all stages all three box modalities are respected in labels. Each step in the construction is a cure of one of the three types defects corresponding to the three modal diamonds. Because of the joins in the chronicles where labels may only match huewise, we are mostly only concerned with the cure of defects involving formulas in $\mathbf{cl}\phi$ but there are a few complications which can not be discussed here.

Banning

The only time that banning comes in to play is when we cure an alethic defect, i.e. when $\Diamond\alpha$ is in the label of some point and we create a new vertical column of points containing a cure for the defect. The new point will be a joint and its downward facing label just any IRMCS cure for the defect will not necessarily agree with its upward facing label.

The upward facing label at this new point will have to take account of any new ban.

In order to make sure that the matching ends property results from this construction we use banning to prevent certain colours from occuring above a joint unless certain conditions are met. Ironically for historical necessity, banning is a way of gradually eliminating indeterminancy during the construction process. We want to only allow a colour to turn up as a label if it has to. This will mean, thanks to the LC rule, that the colours which turn up along an emergent history also turn up along a constructed history.

To do this we use three types of bans. Each of the three types has a corresponding formula of the language.

Suppose that we have a $\Diamond$ defect of $\Diamond\alpha$ in the label at x. We add a new column and a new point y such that $x \equiv y$ and α is in the label at y. Say that we choose the hue h of colour a for y. Now we consider making one of the three types of bans. We may decide whether to ban colour a from appearing again at any point up ($<$) and across ($\equiv$) from y. Suppose τ represents the current ban at x. We will ban a unless $\vdash (a \wedge \tau) \to \Diamond Fa$. If this query succeeds, i.e. if $\vdash (a \wedge \tau) \to \Diamond Fa$ then we do nothing: a may recur. Otherwise we add $\Box G \neg a$ to our bans and labels from y upwards. In this case colour a will not recur.

The second type of ban is to disallow a colour b from coming up again above y unless we then ban a from occurring again above that.

The third type of ban concerns a sequence $[a = e_0, e_1, e_2, ..., e_n]$ of colours and a disjoint set D of colours. We consider whether there must be a sequence of points $y \equiv y_0 < y_1 < ... < y_n$ such that y_i has colour e_i and there are no points between y_0 and y_n coloured from D. To capture the query and ban in this case needs to make use of the unique "names" that points have when we construct a model using the IRR rule. The query and ban formulas are lengthy.

At each joint, each time we cure an alethic defect, we schedule one of these ban queries in a fair way, and add the appropriate ban depending on the outcome of the query.

The limit chronicle as a tree

In the ω-limit of our construction we have a chronicle which corresponds to a bundled Ockhamist tree model $(S_0, <_0, B_0, g_0)$ of ϕ. It is relatively perfect in respecting formulas from $\mathbf{cl}(\phi)$.

This will not generally be able to be completed to an Ockhamist model of ϕ because of the problem of emergent branches mentioned above.

However, the banning procedure will have ensured that we have the important property of "matching ends" as follows. Suppose that η_0 is the labelling of $\{(b, s) | s \in b \in B_0\}$ by hues in the limit chronicle. We say that $\pi \subset S_0$ is an initial segment (of history σ) iff $\pi \subseteq \sigma$ and if $x < y \in \pi$ then $x \in \pi$. For each initial segment π of a history in $(S_0, <_0)$ we define $\mathrm{infcol}(\pi)$ to be the set of colours which appear cofinally as the colours of labels of points in π. We also let $\mathrm{ccreg}(\pi)$ be the set of points $x \in \pi$ such that if $x \leq y \in \pi$ then the colour of the label of y is in $\mathrm{infcol}(\pi)$. Then we have:

LEMMA 5 *Let σ be any history of $(S_0, <_0, B_0, \eta_0)$ which is not in B_0.*

Then for all $x \in \sigma$ there is an initial segment π of a history $\rho \in B_0$ such that:

- *M1)* $x \in \pi \cap \sigma$;
- *M2)* $\mathrm{ccreg}(\pi) \cap \mathrm{ccreg}(\sigma) \neq \emptyset$;
- *M3)* $\mathrm{infcol}(\pi) = \mathrm{infcol}(\sigma)$; *and*
- *M4)* $\pi \not\subseteq \sigma$; *and*
- *M5) either $\pi = \rho$ or colours not in* $\mathrm{infcol}(\pi)$ *appear arbitrarily soon after π in ρ.*

Grafting trees

A grafting tree is $(S, <, B, \eta)$ where $(S, <, B)$ is a bundled tree, $\eta : \{(b, s) \in B \times S | s \in b\} \to H(\phi)$ and η perfectly labels $(S, <, B)$ in the sense that both box and diamond formulas in hues are respected.

Starting with the grafting tree created by the limit chronicle, we proceed via a transfinite induction until the $\aleph_1$ stage. At each stage we extend the previous grafting trees to a new one.

The basic idea is that the grafting tree $\mathcal{S}_\kappa$ at stage κ only has non-bundled histories which each end just like some specific identified non-bundled history of $\mathcal{S}_0$. Consider a non-bundled history σ of $\mathcal{S}_\kappa$. Suppose that the end of it is just a copy of non-bundled history ρ of $\mathcal{S}_0$. By the matching ends property of $\mathcal{S}_0$ we can find am initial segment π of a bundled history of $\mathcal{S}_0$ which matches in terms of cofinal colours. To create $\mathcal{S}_{\kappa+1}$ we extend each such non-bundled σ as follows. Take each (bundled or not) history τ of $\mathcal{S}_0$ which extends π and copy all the points (with their labels) in $\tau \setminus \pi$ over to lie after σ. We create a copy of the branching sub-tree after π and put it after σ: this is what we call grafting. It can be shown that again we have a perfect grafting tree.

At limit ordinal stages we just take unions of the grafting trees at all earlier stages.

The truth lemma

During our transfinite grafting process we also record which points in $\mathcal{S}_0$ are copied to make the points in each $\mathcal{S}_\kappa$.

Certain technical arguments to do with the sets of cofinal colours on the matching ends of extended histories allow us to find, for each unbundled history σ in $\mathcal{S}_{\aleph_1}$ a gradually strictly increasing sequence of initial segments in $\mathcal{S}_0$ which are copied across to contribute to making σ. In fact there will be a set of ordinals, cofinal in $\aleph_1$, which enumerate a strictly increasing sequence of segments and hence points in $\mathcal{S}_0$.

As no countable set of ordinals can be cofinal in $\aleph_1$ this would contradict the countability of $\mathcal{S}_0$. The only conclusion is that there are no unbundled histories in $\mathcal{S}_{\aleph_1}$.

Thus we finally have a complete and perfect bundled grafting tree. Define a valuation g for the atoms using the hue labels $\eta_{\aleph_1}$ and the truth lemma for the resulting Ockhamist tree structure $\mathcal{S} = (S_{\aleph_1}, <_{\aleph_1}, g)$ follows directly.

8 Extensions

Follow on work could include investigating the version of the logic in which the truth values of atoms may depend on the history of evaluation, ie in which we do not make the assumption that atoms contain no trace of futurity. It is not clear whether the techniques in this paper will go through easily to this other logic. There is a particular problem with the copying and grafting technique in which we use the fact that points of the same colour sufficiently far into the cofinal closure region of branches agree on all formulas up to a given complexity. This would not hold if the colour of a point did not determine the truth of atoms.

Other useful related axiomatizations which are left as future work include the Ockhamist logic with until and since connectives, and Ockhamist logics over trees in which all histories have particular properties such as denseness or being the real numbers.

BIBLIOGRAPHY

[1] N. Belnap. Agents in branching time. In *Logic and Reality: Essays in Pure and applied Logic, in Memory of Arthur Prior.* OUP, 1996.

[2] P. Blackburn and V. Goranko. Hybrid ockhamist temporal logic. In C. Bettini and A. Montanari, editors, *Proc. of the Eighth International Symposium on Temporal Representation and Reasoning (TIME-01, Cividale del Friuli, Italy, June 14-16, 2001),*, pages 183–188. IEEE Computer Society Press, 2001.

[3] J. P. Burgess. Logic and time. *J. Symbolic Logic*, 44:566–582, 1979.

[4] J. P. Burgess. Decidability for branching time. *Studia Logica*, 39:203–218, 1980.

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

[6] E. Emerson and J. Halpern. Decision procedures and expressiveness in the temporal logic of branching time. *J. Comp and Sys. Sci*, 30(1):1–24, 1985.

[7] E. Emerson and J. Halpern. 'Sometimes' and 'not never' revisited: on branching versus linear time. *J. ACM*, 33, 1986.

[8] E. Emerson and A. Sistla. Deciding full branching time logic. *Information and Control*, 61:175 – 201, 1984.

[9] T. French. Decidability of quantified propositional branching time logics. In M. Stumptner, D. Corbett, and M. Brooks, editors, *AI 2001: Advances in Artificial Intelligence.* Springer, 2001.

[10] D. Gabbay, I. Hodkinson, and M. Reynolds. *Temporal Logic: Mathematical Foundations and Computational Aspects, Volume 1.* Oxford University Press, 1994.

[11] D. M. Gabbay. An irreflexivity lemma with applications to axiomatizations of conditions on tense frames. In U. Monnich, editor, *Aspects of Philosophical Logic*, pages 67–89. Reidel, Dordrecht, 1981.

[12] Y. Gurevich and S. Shelah. The decision problem for branching time logic. *J. of Symbolic Logic*, 50:668–681, 1985.

[13] I. Hodkinson, F. Wolter, and M. Zakharyashev. Decidable and undecidable fragments of first-order branching temporal logics. Proceedings of the 17th Annual IEEE Symposium on Logic in Computer Science, LICS 2002, pages 393-402.

[14] F. Laroussinie and Ph. Schnoebelen. Specification in CTL+past for verification in CTL. *Information and Computation*, 156:236–263, 2000.

[15] P. Øhrstrøm and P. F. V. Hasle. *Temporal Logic: From Ancient Ideas to Artificial Intelligence*, volume 57 of *Studies in Linguistic and Philosophy.* Dordrecht: Kluver Academic Publishers, 1996.

[16] A. Prior. *Past, Present and Future.* Oxford University Press, 1967.

[17] A. Rao and M. Georgeff. Decision procedures for BDI logics. *Journal of Logic and Computation*, 8(3):293–342, 1998.

[18] M. Reynolds. More past glories. In *Proceedings of Fifteenth Annual IEEE Symposium on Logic in Computer Science (LICS'2000), Santa Barbara, California, USA, June 26-28, 2000*, pages 229–240. IEEE, Los Alamitos, California, USA, 2000.

[19] M. Reynolds. An axiomatization of full computation tree logic. *J. of Symbolic Logic*, 66(3):1011–1057, 2001.

[20] M. Reynolds. Axioms for branching time. *J. Logic and Computation*, 2002. to appear.

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

[22] R. Thomason. Combinations of tense and modality. In D. Gabbay and F. Guenthner, editors, *Handbook of Philosophical Logic, Vol II: Extensions of Classical Logic*, pages 135–165. Reidel, Dordrecht, 1984.

[23] F. von Kutschera. $T \times W$ completeness. *Journal of Philosophical Logic*, 26(3):241–250, 1997.

[24] S. Woelfl. Combinations of Tense and Modality for Predicate Logic. *Journal of Philosophical Logic*, 28(4):371–398, 1999.

[25] M. Xu. Decidability of deliberative *stit* theories with multiple agents. In D. Gabbay and H. Ohlbach, editors, *Temporal Logic, Proceedings of ICTL '94*, number 827 in LNAI, pages 332–348. Springer-Verlag, 1994.

[26] A. Zanardo. A finite axiomatization of the set of strongly valid Ockamist formulas. *J. of Philosophical Logic*, 14:447–468, 1985.

[27] A. Zanardo. Axiomatization of 'Peircean' branching-time logic. *Studia Logica*, 49:183–195, 1990.

[28] A. Zanardo. A complete deductive system for since-until branching time logic. *J. Philosophical Logic*, 1991.

[29] A. Zanardo. Branching-time logic with quantification over branches: the point of view of modal logic. *J. Symbolic Logic*, 61:1–39, 1996.

Mark Reynolds
School of Information Technology
Murdoch University
Murdoch, Perth, Western Australia 6150
Australia
E-mail: m.reynolds@murdoch.edu.au