A Logical Characterization of Well Branching Event Structures[1]

Madhavan Mukund and P. S. Thiagarajan

School of Mathematics
SPIC Science Foundation
92, G. N. Chetty Road
T. Nagar, Madras 600 017
India

Abstract

We develop a tense logic for reasoning about the occurrences of events in a subclass of prime event structures called well branching event structures. The well branching property ensures that two events being in conflict can always be traced back — via the causality relation — to two events being in minimal conflict. Two events are in minimal conflict if they are in conflict and their "unified" past is conflict-free. Thus the minimal conflict relation captures the branching points of the computations supported by the event structure. Our logical language has explicit modalities for talking about causality, conflict, concurrency and minimal conflict. We define the semantics of this logic using well branching event structures as Kripke frames. Our main result is a sound and complete axiomatization of the valid formulas over the chosen class of frames.

[1]This paper appears in *Theoretical Computer Science*, 96, 1, (1992) 35-72.

1 Introduction

Event structures have come to play a central role in the formal study of distributed systems. They clearly capture the intuition concerning the non-sequential and indeterminate behaviours of distributed systems. They have a rich mathematical structure [13, 18]. There are natural bridges to other theories of distributed systems such as net theory and the theory of trace languages [13, 14]. Event structures can be used to provide the non-interleaved denotational semantics of CCS-like languages as demonstrated by Winskel [18]. Hence there is a good deal of motivation for developing a logical framework to reason about the behaviours of distributed systems as represented by event structures.

The general problem of developing proof systems based on syntactic presentations of event structures (via, say, a CCS-like language) appears to be difficult. Here we attempt something more modest. Our aim will be to characterize a subclass of prime event structures called well branching event structures using a suitably chosen version of tense logic. The modalities that we introduce in our logic will permit us to explicitly talk about three basic features of distributed systems: causality, conflict and concurrency. The strong characterization result proved here should lead to a deeper understanding of the interplay between the behavioural aspects of distributed systems as captured by event structures and tense logic.

Prime event structures are in some sense the basic form of event structures. A prime event structure consists of a set of event occurrences partially ordered by a causality relation. In addition, the structure contains a binary conflict relation between the events. The conflict relation is required to be irreflexive, symmetric and inherited via the causality relation. The idea is that an event can occur in a computation only if all the events that lie in its past have occurred in the computation. No two events that are in the conflict relation can both occur in a computation. Two events that are neither causally related nor in conflict are said to be concurrent; whenever they both occur in a computation, they do so with no order over their occurrences.

In a well branching prime event structure two events being in conflict can always be traced back — via the causality relation — to two events being in minimal conflict. Two events are said to be in in minimal conflict if they are in conflict and no two events in their (unified) past are in conflict. In such an event structure, the branching points of a computation — where the system chooses between alternate courses of action — can be clearly identified in terms of a choice made between two events in minimal conflict. These ideas will become more transparent in the next section where we present a brief formal introduction to well branching prime event structures.

Prime event structures that model the behaviour of realizable distributed systems will certainly be well branching. In fact, such event structures will have at least two other attributes which we do not handle in this paper. Firstly, the events will be labelled with the elements of an action set. Secondly, these event structures will be finitary; every event will have at most a finite number of events lying in its past. The problems that arise when one tries to handle these two attributes will be discussed in the concluding section.

Turning now to our logical language, it is a tense (temporal) logic with the usual past and future modalities. In addition it has three unary modalities for expressing concurrency, conflict and minimal conflict. Actually, in the present work, where only well branching event structures are admitted, the conflict modality can be expressed in terms of the past, future and minimal conflict modalities. We have included the conflict modality as a first class object in our language mainly for technical reasons. More remarks regarding this point can be found in Section 3 which presents the language and develops a Kripke-style semantics for it with well branching prime event structures serving as the frames.

The conflict modality was introduced by Penczek [15]. He also introduced the logical means for talking about the maximal computations of an event structure called runs. In [16] he uses a modality called the immediate conflict modality. However this modality is not interpreted as the minimal conflict relation. Instead, it is given a rather weak interpretation. We will discuss in greater detail how our work relates to Penczek's work in the concluding section.

The main technical contribution of this paper is a sound and complete axiomatization of validity with respect to the chosen semantics. The axiomatization is presented in Section 4. The completeness proof, which is somewhat formidable, is spread over Section 5 and the Appendix. In the last section we discuss related work and issues.

Figure 1: An Event Structure

2 Event Structures

In this paper, we will deal only with prime event structures and hence we will simply call them event structures.

Definition 2.1 *An* event structure *is a triple* $ES = (E, <, \#)$ *where*

(i) E *is a set of* events *(or better,* event occurrences*).*

(ii) $< \subseteq E \times E$ *is an irreflexive and transitive* causality relation*.*

(iii) $\# \subseteq E \times E$ *is an irreflexive and symmetric* conflict relation*.*

(iv) $\#$ *is inherited via* $<$ *in the sense that* $e_1 \# e_2 < e_3$ *implies that* $e_1 \# e_3$ *for every* e_1, e_2, e_3 *in* E*.*

Usually the causality relation is required to be a partial ordering relation. We have made it a strict partial ordering relation because it fits in better with the completeness argument.

Let $ES = (E, <, \#)$ be an event structure. Then

$$\begin{array}{lcl} id & \stackrel{\text{def}}{=} & \{(e,e) \mid e \in E\} \\ > & \stackrel{\text{def}}{=} & \{(e,e') \mid (e',e) \in <\} \\ \leq & \stackrel{\text{def}}{=} & < \cup\, id \\ \geq & \stackrel{\text{def}}{=} & > \cup\, id \quad \text{and} \\ co & \stackrel{\text{def}}{=} & E \times E - (\leq \cup \geq \cup \#). \end{array}$$

The relation co is called the *concurrency* relation. Observe that the relations $\{<, >, \#, co, id\}$ partition $E \times E$.

It is necessary to define one more auxiliary relation. Let $ES = (E, <, \#)$ be an event structure and $e, e' \in E$. Then

$$e \#_\mu e' \stackrel{\text{def}}{=} e \# e' \textit{ and } \forall e_1, e_1' \in E : [\ e_1 \leq e \textit{ and } e_1' \leq e' \textit{ and } e_1 \# e_1' \textit{ implies } e_1 = e \textit{ and } e_1' = e'\ ].$$

$\#_\mu$ identifies the minimal elements (under $<$) of the $\#$ relation and is hence called the *minimal conflict* relation. The $\#$ relation identifies pairs of events which are inconsistent with each other and therefore cannot both occur during any run of the system. $\#_\mu$ identifies the branching points in the behaviour where choices are made between conflicting events. This "basic" conflict then propagates to causally related events and "generates" other conflicts.

In general, there may be events in $\#$ whose inconsistency cannot be traced back to a pair of events in $\#_\mu$ — a typical example consists of two infinite descending chains of events in $\#$ with each other. However, it is difficult to find useful examples of concurrent systems for modelling which one requires such event structures. We shall therefore restrict our attention to the class of well branching event structures.

Definition 2.2 *Let* $ES = (E, <, \#)$ *be an event structure.* ES *is* well branching *iff*

$$\forall e, e' \in E : e \# e' \textit{ implies } \exists e_1, e_1' \in E : e_1 \leq e \textit{ and } e_1' \leq e' \textit{ and } e_1 \#_\mu e_1'.$$

Assuming well branching, we can specify an event structure by displaying its $<$ and $\#_\mu$ relations. The $\#$ relation is then uniquely determined by part (iv) of Definition 2.1.

Figure 1 is an example of an event structure. The squiggly lines represent the $\#_\mu$ relation. The causality relation is shown in the form of the associated Hasse diagram. In this event structure, $e_1 \# e_6$ because $e_1 \#_\mu e_2 < e_6$. It is also easy to see that $e_6 \; co \; e_7$.

Notice that a well branching event structure may have infinite descending chains of events — and even pairs of infinite descending chains of events in $\#$ with each other. All that well branching ensures is that every pair of events in $\#$ is "guarded" below by a pair of events in $\#_\mu$. Thus, in Figure 2, every pair (e_i, f_j), $i, j \geq 0$, is in $\#$. However, each such element in $\#$ can be traced back to the pair $(e_k', f_k') \in \#_\mu$,

Figure 2: A Well Branching Event Structure that is not Well Founded

where $k = max(i, j)$, and so the event structure is well branching. Thus, well branching is a fairly weak restriction and does not, in particular, imply well foundedness with respect to $<$.

The states of an event structure are called configurations. A configuration identifies a set of events that have occurred "so far". An event can occur only if all the events in its past have occurred. Two events that are in conflict can never both occur in the same stretch of behaviour. Before formalizing these notions it will be convenient to adopt the following notation.

Let $ES = (E, <, \#)$ be an event structure and $X \subseteq E$. Then

$$\downarrow X = \{e' \mid \exists e \in X : e' \leq e\}$$

For the singleton $\{e\}$, we shall write $\downarrow e$ instead of $\downarrow\{e\}$.

Definition 2.3 *Let $ES = (E, <, \#)$ be an event structure and $c \subseteq E$. Then c is a* configuration *iff*

(i) $c = \downarrow c$ *(left-closed)*

(ii) $(c \times c) \cap \# = \emptyset$ *(conflict-free)*

For the event structure shown in Figure 1, $\{e_2, e_5, e_6\}$ is a configuration. $\{e_2, e_5, e_{10}\}$ is not a configuration because it is not left-closed and $\{e_3, e_7, e_8\}$ is not a configuration because it is not conflict-free.

Let $\mathcal{C}_{ES}$ denote the set of configurations of the event structure ES. In this paper, we will be concerned with only the *local configurations* of an event structure. The notion of a local configuration is based on a simple but crucial observation which lies at the heart of the theory of event structures [13].

Proposition 2.4 *Let $ES = (E, <, \#)$ be an event structure and $e \in E$. Then $\downarrow e$ is a configuration.*

Proof Follows easily from the definitions.

□

We now define $LC_{ES} = \{\downarrow e \mid e \in E\}$ to be the set of *local configurations* of the event structure $ES = (E, <, \#)$. We shall interpret the formulas of our logical language only at the local configurations of an event structure.

The configuration $\downarrow e$ corresponds to the state of the system when the event e has just occurred and thus represents the view of the system seen by observers participating in e. Suppose that $ES = (E, <, \#)$ models the behaviour of a system of communicating sequential processes. Then at least one process or *agent* will be involved in the occurrence of each event. If the agent j participates in e, then $\downarrow e$ will represent the local history of agent j upto the stage where e has occurred, together with the "latest" histories of all the agents that have communicated with the agent j upto the occurrence of e. This will be true of every agent that participates in e. In other words, $\downarrow e$ represents a "synchronized" set of local states of the agents that participate in the occurrence of e.

Another motivation for only considering local configurations comes from the work of Nielsen, Plotkin and Winskel. In [13], they show that the poset $(\mathcal{C}_{ES}, \subseteq)$ of configurations ordered under inclusion is prime algebraic and coherent. The fact that this poset is prime algebraic basically means that certain configurations in the poset are "special". These are called the *prime* elements of the poset. An arbitrary configuration is completely characterized by the prime elements that it dominates in the poset. It turns out that the prime elements of the poset $(\mathcal{C}_{ES}, \subseteq)$ are precisely the local configurations, so it makes sense to tie our assertions about the event structure to these configurations.

3 The Language and its Models

We fix $\mathcal{P} = \{p_1, p_2, \ldots\}$, a countably infinite set of atomic propositions. We let p, q with or without subscripts range over $\mathcal{P}$. The formulas of our language are then defined inductively as follows:

(i) Every member of $\mathcal{P}$ is a formula.

(ii) If α and β are formulas then so are $\neg\alpha, \alpha \vee \beta, \Box\alpha, \boxminus\alpha, \triangledown\alpha, \vartriangle\alpha$ and $\triangledown_\mu\alpha$.

$\Box$ and $\boxminus$ will denote the future and past modalities of tense logic respectively. $\triangledown$ will capture conflict, $\vartriangle$ will capture concurrency and $\triangledown_\mu$ will capture minimal conflict.

For the rest of the paper, Φ will denote the set of formulas of our language. We let α, β, γ and δ with or without subscripts range over Φ.

A *frame* is an ordered pair $Fr = (ES, LC_{ES})$ where $ES = (E, <, \#)$ is a well branching event structure and LC_{ES} is the set of local configurations of ES. The members of LC_{ES} will play the role of possible worlds in the Kripke-style semantics we are about to define.

A *model* is an ordered pair $M = (Fr, V)$ where $Fr = (ES, LC_{ES})$ is a frame and $V\colon \mathcal{P} \longrightarrow 2^{LC_{ES}}$ is a *valuation function.*

Let $M = (Fr, V)$ be a model with $Fr = (ES, LC_{ES})$, $ES = (E, <, \#)$, and $\downarrow e \in LC_{ES}$. Then the notion of a formula α being satisfied at the local configuration $\downarrow e$ in the model M is denoted as $M, \downarrow e \models \alpha$ and is defined inductively as follows:

$$
\begin{array}{ll}
M, \downarrow e \models p & \textit{iff } \downarrow e \in V(p), \textit{ for } p \in \mathcal{P}. \\
M, \downarrow e \models \neg\alpha & \textit{iff } M, \downarrow e \not\models \alpha. \\
M, \downarrow e \models \alpha \vee \beta & \textit{iff } M, \downarrow e \models \alpha \textit{ or } M, \downarrow e \models \beta. \\
M, \downarrow e \models \Box\alpha & \textit{iff } \forall e' \in E : e < e' \textit{ implies } M, \downarrow e' \models \alpha. \\
M, \downarrow e \models \boxminus\alpha & \textit{iff } \forall e' \in E : e' < e \textit{ implies } M, \downarrow e' \models \alpha. \\
M, \downarrow e \models \triangledown\alpha & \textit{iff } \forall e' \in E : e \# e' \textit{ implies } M, \downarrow e' \models \alpha. \\
M, \downarrow e \models \vartriangle\alpha & \textit{iff } \forall e' \in E : e \textit{ co } e' \textit{ implies } M, \downarrow e' \models \alpha. \\
M, \downarrow e \models \triangledown_\mu\alpha & \textit{iff } \forall e' \in E : e \#_\mu e' \textit{ implies } M, \downarrow e' \models \alpha.
\end{array}
$$

The formula α is *satisfiable* if there exists a model $M = ((ES, LC_{ES}), V)$ and a local configuration $\downarrow e \in LC_{ES}$ such that $M, \downarrow e \models \alpha$. α is *M-valid* — denoted $M \models \alpha$ — iff $M, \downarrow e \models \alpha$ for every $\downarrow e \in LC_{ES}$. α is *valid* — and this is denoted by $\models \alpha$ — iff α is M-valid for every model M.

We define the derived connectives of the Propositional Calculus (PC) such as $\wedge$, $\supset$ and $\equiv$ in terms of $\neg$ and $\vee$ in the usual way. In addition, we define

$$
\begin{array}{ccc}
\Diamond\alpha & \stackrel{\text{def}}{=} & \neg\Box\neg\alpha \\
\diamondminus\alpha & \stackrel{\text{def}}{=} & \neg\boxminus\neg\alpha \\
\bigtriangledown\alpha & \stackrel{\text{def}}{=} & \neg\triangledown\neg\alpha \\
\bigtriangleup\alpha & \stackrel{\text{def}}{=} & \neg\vartriangle\neg\alpha \\
\bigtriangledown_\mu\alpha & \stackrel{\text{def}}{=} & \neg\triangledown_\mu\neg\alpha
\end{array}
$$

It is easy to verify that $\triangledown\alpha \equiv \triangledown_\mu\alpha \vee \diamondminus\triangledown_\mu\alpha \vee \triangledown_\mu\Diamond\alpha \vee \diamondminus\triangledown_\mu\Diamond\alpha$ is a valid formula. Thus we could eliminate the modality $\triangledown$ from the definition of the language and instead make it a *derived* modality using $\triangledown_\mu$, $\Box$ and $\boxminus$. However, it will be very convenient to keep $\triangledown$ in the logic as a basic operator. We will instead throw in the above "definition" of $\triangledown$ in terms of $\triangledown_\mu$, $\Box$ and $\boxminus$ as an axiom of our logical system. In [9], it is shown that the language which contains the $\triangledown$ operator but not the $\triangledown_\mu$ operator is strictly less expressive for the class of models based on well branching event structures.

Also, notice that the usual reflexive forms of the future and past operators are easily expressible in our logic. $\alpha \wedge \Box\alpha$ defines a reflexive future operator while $\alpha \wedge \boxminus\alpha$ defines the corresponding reflexive past operator.

A number of interesting properties of distributed computations can be expressed in our language.

To illustrate, consider a finite set of sequential agents $A_1, A_2, \ldots, A_n$. Assume that $\{Q_1, Q_2, \ldots, Q_n\}$ is a collection of pair-wise disjoint finite non-empty sets. Q_i is the set of local states that agent A_i can assume. The agents communicate with each other asynchronously. The means for associating an event structure with such a system will be assumed to be available (see for example [10]). Let us call such event structures Communicating Sequential Agents (CSAs).

To reason about such systems in our logic, we specify certain "frame axioms", and focus on models based on CSAs over which these axioms are valid. We can then use these frame axioms in conjunction with the general proof system for event structures (which is presented in the next section) to derive facts about this class of models.

To begin with, suppose that for each i and each $q \in Q_i$ we have an atomic proposition, also called q for convenience, which denotes that the current local state is q. We first make the following assertions (where $Q = \bigcup\limits_i Q_i$).

$(i)\ \bigvee_{q \in Q} q$

$(ii)\ \bigwedge_{q,q' \in Q,\ q \neq q'} q \supset \neg q'$

These ensure that the local configurations of the models we consider can be mapped uniquely to the state space of the agents. In this framework, $at_i \stackrel{\text{def}}{=} \bigvee_{q \in Q_i} q$ will denote the fact that the current local configuration belongs to A_i.

We can capture the fact that the individual agents are *sequential* — i.e. they exhibit no concurrency in their local behaviours — by asserting

$(iii)\ \bigwedge_i (at_i \supset \triangle(\neg at_i))$

The next assertion expresses the important fact that (non-deterministic) choices in the behaviour are made locally by the individual agents.

$(iv)\ \bigwedge_i (at_i \supset \nabla_\mu at_i)$

To reason about the behaviour of the agents, it is convenient to have the means for talking about the computations of an event structure ES. Let $(\mathcal{C}_{ES}, \subseteq)$ be the poset of configurations of ES ordered under inclusion. The maximal elements of this poset correspond to the maximal (i.e. "non-extendable") computations of the event structure, which we shall call the *runs* of the event structure.

Following [15], we reserve a proposition ρ to designate a run. Let $ES = (E, <, \#)$ be an event structure. It is easy to verify that $r \subseteq E$ is a run iff

$$\forall e \in E : (e \in r \textit{ iff } \forall e' \in E : e \# e' \textit{ implies } e' \notin r).$$

Keeping this in mind, we can demand that the local configurations satisfying ρ constitute a run in any model we consider by asserting the following:

$(v)\ \rho \equiv \nabla \neg\rho$

Now suppose that the agents are running a common protocol to achieve a *stable property* — that is, once the property is satisfied in a run it remains true for the remainder of the run. For instance the agents may be running a protocol such as the snapshot algorithm [3], superimposed on the actual computation, to detect the termination of the main computation.

Let α denote the fact that an individual agent has "recognized", as a result of running the protocol, that the system has achieved a stable property – in this case, α would denote that the agent has detected that the main computation has terminated. If the protocol has worked correctly, α will be become true in any run within an agent only if the stable property has in fact been achieved, and so α itself must be "stable" within each agent. The formula $\bigwedge_j [at_j \wedge \rho \supset (\alpha \supset \Box((at_j \wedge \rho) \supset \alpha))]$ can be used to specify this.

In addition, to verify the protocol we have to check that in each run of the system, if A_i achieves a local state satisfying α then every other agent also eventually achieves a local state satisfying α. Notice that any pair of local states that occur in the same run must either be ordered or in *co* with each other. As a result, during a run A_i can "look across" to all the local states of some other agent A_j which belong to the same run using the modalities $\Box$, $\boxminus$ and $\triangle$. Hence, we can specify the correctness of the protocol by asserting:

$$\begin{aligned} at_i \wedge \rho \wedge \alpha \supset \bigwedge_{j \neq i} [\ & \Box((at_j \wedge \rho) \supset \Diamond(at_j \wedge \rho \wedge \alpha)) \wedge \\ & \boxminus((at_j \wedge \rho) \supset \Diamond(at_j \wedge \rho \wedge \alpha)) \wedge \\ & \triangle((at_j \wedge \rho) \supset \Diamond(at_j \wedge \rho \wedge \alpha))] \end{aligned}$$

On the other hand, suppose that the agents are meant to constitute a non-terminating system. We can specify that the system is free from "local" deadlock by asserting that every run is "perpetual" within each agent as follows:

$$\bigwedge_i ((at_i \wedge \rho) \supset \Diamond(at_i \wedge \rho))$$

We can also specify the weaker property that the system as a whole has no finite computations, though individual agents may terminate, as follows:

$$\bigwedge_i \{(at_i \wedge \rho) \supset \begin{array}[t]{l} [\Diamond(at_i \wedge \rho)) \vee \\ \bigvee_{j \neq i} (\begin{array}[t]{l} \Box((at_j \wedge \rho) \supset \Diamond(at_j \wedge \rho)) \wedge \\ \boxminus((at_j \wedge \rho) \supset \Diamond(at_j \wedge \rho)) \wedge \\ \triangle((at_j \wedge \rho) \supset \Diamond(at_j \wedge \rho)))]\} \end{array} \end{array}$$

This formula asserts that in case an agent is unable to make progress within a run, it must be able to look across to some other agent which is non-terminating, thereby ensuring that the run is infinite.

The intuitive notion of frame axioms used in the example can be formalized. Let $A = \{\alpha_1, \alpha_2, \ldots, \alpha_n\}$ be a finite set of formulas. Then, we will say that A *semantically entails* β — and denote this by $A \models \beta$ — iff for every model M, if $M \models \alpha_1 \wedge \alpha_2 \wedge \ldots \wedge \alpha_n$ then $M \models \beta$. Thus, in the example above, the intention is that the formulas we use for specifying properties about the class of models based on CSAs should be semantically entailed by the frame axioms (i) through (v) together with a suitably chosen set of assertions satisfied by a protocol. This will then ensure that the protocol (semantically) meets the specifications.

We can relate the notions of semantic entailment and validity. To do this, we introduce the modality $\mathcal{E}$ defined as follows:

$$\mathcal{E}\alpha \stackrel{\mathrm{def}}{=} \alpha \wedge \Box\alpha \wedge \boxminus\alpha \wedge \triangledown\alpha \wedge \triangle\alpha$$

$\mathcal{E}\alpha$ is to be read as "everywhere α". From the semantics it is obvious that if $\mathcal{E}\alpha$ is satisfied at a local configuration in a model then α is satisfied at every local configuration in the model. We then have the following result.

Theorem 3.1 $\{\alpha_1, \alpha_2, \ldots, \alpha_n\} \models \beta$ *iff* $\models \mathcal{E}(\alpha_1 \wedge \alpha_2 \wedge \ldots \wedge \alpha_n) \supset \beta$

Proof Follows easily from the definitions.

□

4 The Axiom System

Our axioms are a combination of standard modal logic axioms and tense logic axioms [1, 4], along with a few new axioms which reflect the restrictions imposed on the relations $<$, $\#$ and co in the definition of event structures. We first present the logical system in full, and then provide some explanatory remarks.

AXIOM SCHEMES

(A1) All the substitution instances of the tautologies of propositional logic.

(A2) (i) $\Box(\alpha \supset \beta) \supset (\Box\alpha \supset \Box\beta)$ (Deductive Closure)

(ii) $\boxminus(\alpha \supset \beta) \supset (\boxminus\alpha \supset \boxminus\beta)$

(iii) $\triangledown(\alpha \supset \beta) \supset (\triangledown\alpha \supset \triangledown\beta)$

(iv) $\triangle(\alpha \supset \beta) \supset (\triangle\alpha \supset \triangle\beta)$

(v) $\triangledown_\mu(\alpha \supset \beta) \supset (\triangledown_\mu\alpha \supset \triangledown_\mu\beta)$

(A3) (i) $\Box\alpha \supset \Box\Box\alpha$ (Transitivity of $<$)

(ii) $\boxminus\alpha \supset \boxminus\boxminus\alpha$

(A4) (i) $\alpha \supset \triangledown\triangledown\alpha$ (Symmetry of $\#$, $\#_\mu$ and co)

(ii) $\alpha \supset \triangledown_\mu\triangledown_\mu\alpha$

(iii) $\alpha \supset \triangle\triangle\alpha$

(A5) (i) $\alpha \supset \Box\diamondminus\alpha$ (Relating past and future)

(ii) $\alpha \supset \boxminus\Diamond\alpha$

(A6) $\bigtriangledown\alpha \supset \Box\bigtriangledown\alpha$ (Conflict inheritance)

(A7) $\bigtriangledown\alpha \equiv \bigtriangledown_\mu\alpha \vee \diamondminus\bigtriangledown_\mu\alpha \vee \bigtriangledown_\mu\Diamond\alpha \vee \diamondminus\bigtriangledown_\mu\Diamond\alpha$ (Well Branching)

(A8) (i) $\bigtriangledown_\mu\alpha \supset \boxminus(\Diamond\alpha \vee \bigtriangleup\alpha)$ (Conflict-free past)

(ii) $\bigtriangleup\alpha \supset \boxminus(\Diamond\alpha \vee \bigtriangleup\alpha)$

(A9) (i) $\Diamond\alpha \supset \Box(\alpha \vee \Diamond\alpha \vee \diamondminus\alpha \vee \bigtriangledown\alpha \vee \bigtriangleup\alpha)$ (Relating weak and strong modalities)

(ii) $\bigtriangledown\alpha \supset \dot{\bigtriangledown}(\alpha \vee \Diamond\alpha \vee \diamondminus\alpha \vee \bigtriangledown\alpha \vee \bigtriangleup\alpha)$

(iii) $\bigtriangleup\alpha \supset \dot{\bigtriangleup}(\alpha \vee \Diamond\alpha \vee \diamondminus\alpha \vee \bigtriangledown\alpha \vee \bigtriangleup\alpha)$

(iv) $\diamondminus\alpha \supset \boxminus(\alpha \vee \Diamond\alpha \vee \diamondminus\alpha \vee \bigtriangleup\alpha)$

(A10) $\bigtriangledown\alpha \supset \dot{\bigtriangleup}(\Diamond\alpha \vee \bigtriangledown\alpha \vee \bigtriangleup\alpha)$ (Relating $\dot{\bigtriangleup}$ and $\dot{\bigtriangledown}$)

(A11) $\bigtriangleup\alpha \supset \Box(\diamondminus\alpha \vee \bigtriangledown\alpha \vee \bigtriangleup\alpha)$ (Relating $\dot{\bigtriangleup}$ and $\Box$)

(A12) $\bigtriangledown\alpha \supset \boxminus(\Diamond\alpha \vee \bigtriangledown\alpha \vee \bigtriangleup\alpha)$ (Relating $\dot{\bigtriangledown}$ and $\boxminus$)

INFERENCE RULES

(MP) $\dfrac{\alpha \supset \beta,\ \alpha}{\beta}$

(TG) (i) $\frac{\alpha}{\Box\alpha}$ (ii) $\frac{\alpha}{\boxminus\alpha}$ (iii) $\frac{\alpha}{\dot{\bigtriangledown}\alpha}$ (iv) $\frac{\alpha}{\dot{\bigtriangleup}\alpha}$ (v) $\frac{\alpha}{\dot{\bigtriangledown}_\mu\alpha}$

(UNIQ) $\dfrac{\hat{p} \supset \alpha}{\alpha}$ where p is an atomic proposition not appearing in α and $\hat{p} \stackrel{\text{def}}{=} p \wedge \Box\neg p \wedge \boxminus\neg p \wedge \dot{\bigtriangleup}\neg p \wedge \dot{\bigtriangledown}\neg p$

Axioms A1 and A2 and inference rules MP and TG are standard. A3 and A4 are versions of the modal logic axioms T and B respectively which express the transitivity of $<$ and the symmetry of $\#$, $\#_\mu$ and *co*. A5 is the standard tense logic axiom relating the past and future modalities. A6 expresses the fact that conflict is inherited via $<$. A7 captures the fact that the event structure is well branching. A8(i) characterizes $\#_\mu$ as the minimal conflict relation while A8(ii) ensures that any two events related by *co* have consistent (i.e. conflict-free) pasts. The remaining axioms are necessary to capture the fact that the relations $<, <^{-1}, \#$ and *co* "cover" the event structure – i.e., any two distinct events are related by one of these relations.

The rule UNIQ is adapted from [2]. Notice that given a proposition p, the formula $\hat{p}$ can be true at at most one local configuration, by the definition of $\hat{p}$. Hence, we can label each local configuration $\downarrow e_i$ by a distinct formula $\hat{p}_i$. The rule UNIQ allows us to construct this labelling, which is crucial in demonstrating the completeness of the axiomatization.

A formula α is a *thesis* if it is derivable in our axiom system. We denote that α is a thesis by $\vdash \alpha$.

Theorem 4.1 (Soundness) *If* $\vdash \alpha$ *then* $\models \alpha$.

Proof It is routine to verify that the axioms A1 to A12 are valid and that the rules MP and TG preserve validity.

To show that UNIQ preserves validity requires a little bit of work. We have to show that if $\hat{p} \supset \alpha$ is valid, and p does not appear in α, then α is valid as well.

This is equivalent to showing that if $\neg\alpha$ is satisfiable then $\hat{p} \wedge \neg\alpha$ is satisfiable (i.e. $\not\models \alpha$ implies $\not\models \hat{p} \supset \alpha$.)

Suppose that $\neg\alpha$ is satisfiable. Then there exists a model $M = ((ES, LC_{ES}), V)$ and a local configuration $\downarrow e_0 \in LC_{ES}$ such that $M, \downarrow e_0 \models \neg\alpha$. For $\beta \in \Phi$, let $Voc(\beta)$ denote the set of propositions that appear in β — $Voc(\beta)$ is to be read as the vocabulary of β.

Define a new valuation function V' as follows:

$$\forall q \in \mathcal{P} : V'(q) = \begin{cases} V(q) & \text{if } q \in Voc(\neg\alpha) \\ \{\downarrow e_0\} & \text{otherwise} \end{cases}$$

Let $M' = ((ES, LC_{ES}), V')$. Let $\beta \in \Phi$ such that $Voc(\beta) \subseteq Voc(\neg\alpha)$. It is easy to verify the following, by induction on the structure of β.

$$\forall{\downarrow}e \in LC_{ES} : M', {\downarrow}e \models \beta \textit{ iff } M, {\downarrow}e \models \beta.$$

Thus we have $M', {\downarrow}e_0 \models \neg\alpha$. At the same time, since $p \notin Voc(\neg\alpha)$, $M', {\downarrow}e_0 \models \hat{p}$. So $M', {\downarrow}e_0 \models \hat{p} \wedge \neg\alpha$ and we are done.

□

We can now define the notion of a theory. Let $A = \{\alpha_1, \alpha_2, \ldots, \alpha_n\}$ be a *finite* set of formulas (here we mean concrete formulas and not schemes closed under substitution). We say that A derives β — which we denote as $A \vdash \beta$ — iff we can derive β in a finite number of steps using the formulas $\{\alpha_1, \alpha_2, \ldots, \alpha_n\}$ in conjunction with our axioms and inference rules.

To ensure that our axiomatization is sound for finite theories, we have to parametrize the rule (UNIQ) by the theory. For example, consider the following derivation, where $p, q \in \mathcal{P}$.

$\hat{p} \supset \hat{q} \vdash \hat{p} \supset \hat{q}$	(Assumption)
$\hat{p} \supset \hat{q} \vdash \hat{q}$	(UNIQ)

It is easy to verify that $\hat{p} \supset \hat{q} \not\models \hat{q}$, and so the rule (UNIQ) is not sound as it stands. Instead, for the theory A, we have to define the rule (UNIQ_A) as follows:

(UNIQ_A) $\quad \dfrac{A \vdash \hat{p} \supset \alpha}{A \vdash \alpha}$ $\quad$ where p is an atomic proposition appearing neither in α nor any of the formulas in A and $\hat{p} \stackrel{\text{def}}{=} p \wedge \Box\neg p \wedge \boxminus\neg p \wedge \triangle\neg p \wedge \triangledown\neg p$

The remaining axioms and inference rules do not change for finite theories. The axiomatization presented earlier then turns out to be a special case of the general axiomatization of finite theories, where the set of assumptions A is empty.

Recall that $\mathcal{E}\alpha$ stands for the formula $\alpha \wedge \Box\alpha \wedge \boxminus\alpha \wedge \triangledown\alpha \wedge \triangle\alpha$. We then have the following result.

Theorem 4.2 (Deduction) *$\{\alpha_1, \alpha_2, \ldots, \alpha_n\} \vdash \beta$ iff $\vdash \mathcal{E}(\alpha_1 \wedge \alpha_2 \wedge \ldots \wedge \alpha_n) \supset \beta$.*

Proof Set $\alpha = \alpha_1 \wedge \alpha_2 \wedge \ldots \wedge \alpha_n$. It suffices to show that $\alpha \vdash \beta$ iff $\vdash \mathcal{E}\alpha \supset \beta$.

Suppose $\vdash \mathcal{E}\alpha \supset \beta$. Then we have $\alpha \vdash \mathcal{E}\alpha \supset \beta$. Using the inference rule (TG), it is easy to see that the following is a derived inference rule:

$$\frac{A \vdash \gamma}{A \vdash \mathcal{E}\gamma}$$

Hence, since $\alpha \vdash \alpha$, we get $\alpha \vdash \mathcal{E}\alpha$. Using (MP), we can then derive $\alpha \vdash \beta$.

The proof in the other direction is by induction on the length of the derivation $\alpha \vdash \beta$. Let $A = \{\alpha_1, \alpha_2, \ldots, \alpha_n\}$. The base case, where β is an instance of one of the axioms or is one of the formulas $\alpha_i \in A$ is straightforward. For the induction step, the only interesting case is when the rule (UNIQ_A) is used to derive β. We then know that at some earlier step in the derivation we have $\alpha \vdash \hat{p} \supset \beta$, where p does not appear in α or β. By the induction hypothesis, we have $\vdash \mathcal{E}\alpha \supset (\hat{p} \supset \beta)$. Rewriting this we get $\vdash (\mathcal{E}\alpha \wedge \hat{p}) \supset \beta$ and hence $\vdash \hat{p} \supset (\mathcal{E}\alpha \supset \beta)$. We can then apply ($UNIQ_\emptyset$) to obtain $\vdash \mathcal{E}\alpha \supset \beta$.

□

From Theorems 3.1, 4.1 and 4.2 it follows that $A \vdash \beta$ implies $A \models \beta$. We will be able to prove the implication in the other direction when we establish the completeness of the axiomatization.

We shall need some theses and derived inference rules to prove the completeness of our axiom system. For convenience, we shall merely state these theses and rules here. The theses and rules are derived in detail in [12].

We begin with some standard theses and inference rules of the modal logic system **K** [4]. In what follows we let $\odot$ range over the set $\{\Box, \boxminus, \triangledown, \triangle, \triangledown_\mu\}$. $\bigcirc$ will abbreviate $\neg\odot\neg$. We have adopted the convention that in any formula (scheme) the same value is substituted for all appearances of $\odot$. Thus, for example, $\odot\alpha \supset \bigcirc\alpha$ would actually stand for $\Box\alpha \supset \Diamond\alpha$, $\boxminus\alpha \supset \diamondminus\alpha$, $\triangledown\alpha \supset \triangledown\alpha$, $\triangle\alpha \supset \triangle\alpha$ and $\triangledown_\mu\alpha \supset \triangledown_\mu\alpha$.

(TK1) $\quad \odot(\alpha \wedge \beta) \equiv \odot\alpha \wedge \odot\beta$

(TK2) $\quad \bigcirc(\alpha \wedge \beta) \supset \bigcirc\alpha \wedge \bigcirc\beta$

(TK3) $\bigodot\alpha \wedge \bigcirc\beta \supset \bigcirc(\alpha \wedge \beta)$

(DRK) (i) $\dfrac{\alpha \supset \beta}{\bigodot\alpha \supset \bigodot\beta}$ (ii) $\dfrac{\alpha \supset \beta}{\bigcirc\alpha \supset \bigcirc\beta}$

We shall also need the "dual" versions of some of the axioms.

(T6Dual) $\Diamond\alpha \supset \dot{\triangledown}\triangledown\alpha$

(T8Dual) $\diamondminus\alpha \supset \dot{\triangle}(\diamondminus\alpha \vee \triangle\alpha)$

(T10Dual) $\triangle\alpha \supset \dot{\triangledown}(\diamondminus\alpha \vee \triangledown\alpha \vee \triangle\alpha)$

(T11Dual) $\Diamond\alpha \supset \dot{\triangle}(\Diamond\alpha \vee \triangledown\alpha \vee \triangle\alpha)$

(T12Dual) $\diamondminus\alpha \supset \dot{\triangledown}(\diamondminus\alpha \vee \triangledown\alpha \vee \triangle\alpha)$

In addition, we need the following theses concerning $\hat{p}$.

(T$\hat{p}$) (i) $\Diamond\hat{p} \supset \neg p \wedge \boxminus\neg p \wedge \dot{\triangledown}\neg p \wedge \dot{\triangle}\neg p$

(ii) $\diamondminus\hat{p} \supset \neg p \wedge \Box\neg p \wedge \dot{\triangledown}\neg p \wedge \dot{\triangle}\neg p$

(iii) $\triangledown\hat{p} \supset \neg p \wedge \Box\neg p \wedge \boxminus\neg p \wedge \dot{\triangle}\neg p$

(iv) $\triangle\hat{p} \supset \neg p \wedge \Box\neg p \wedge \boxminus\neg p \wedge \dot{\triangledown}\neg p$

(v) $\triangledown(\hat{p} \wedge \alpha) \supset \dot{\triangledown}(\hat{p} \supset \alpha)$

Finally, we need a restricted type of substitution rule.

Let $f : \mathcal{P} \to \mathcal{P}$ be a function mapping atomic propositions to atomic propositions. We can extend f in an obvious way to $\hat{f} : \Phi \to \Phi$ through the syntax of our language, namely:

$\hat{f}(q) = f(q)$, for all $q \in \mathcal{P}$
$\hat{f}(\neg\alpha) = \neg\hat{f}(\alpha)$
$\hat{f}(\alpha \vee \beta) = \hat{f}(\alpha) \vee \hat{f}(\beta)$
$\hat{f}(\bigodot\alpha) = \bigodot(\hat{f}(\alpha))$, where $\bigodot \in \{\Box, \boxminus, \dot{\triangledown}, \dot{\triangle}, \dot{\triangledown}_\mu\}$

Since $\hat{f}$ is uniquely determined by f, henceforth we shall always write $f(\alpha)$ instead of $\hat{f}(\alpha)$ for all $f : \mathcal{P} \to \mathcal{P}$ and $\alpha \in \Phi$.

The derived inference rule that we want is

(SUB) $\dfrac{\alpha}{f(\alpha)}$ $(f : \mathcal{P} \to \mathcal{P})$

5 Completeness

We now wish to show that the axiomatization presented in the previous section is complete.

As usual, by a *consistent* formula, we shall mean a formula whose negation is not a thesis of our system. Our proof of completeness will establish that every consistent formula is satisfiable — in other words, we show that if $\nvdash \neg\alpha$ then $\not\models \neg\alpha$.

The finite set of formulas $\{\alpha_1, \alpha_2, \ldots, \alpha_m\}$ is consistent iff $\alpha_1 \wedge \alpha_2 \wedge \ldots \wedge \alpha_m$ is consistent. A set of formulas is consistent iff every finite subset is. A *Maximal Consistent Set (MCS)* is a consistent set of formulas which is not properly included in any other consistent set.

We shall assume the usual properties associated with MCSs. We shall also assume Lindenbaum's Lemma which says that any consistent set of formulas can be extended to an MCS.

In a model, the set of formulas satisfied at each local configuration is an MCS. Thus, given a frame, we can define a model by associating an MCS with each local configuration, instead of specifying a valuation function. Following [1] we call a function assigning MCSs to a frame a *chronicle*, and a frame together with a chronicle a *chronicle structure*.

Clearly not every chronicle structure corresponds to a model. We need to define coherency conditions which must be satisfied by the chronicle in order to yield a model. When these conditions are satisfied, the chronicle is said to be *perfect*. A perfect chronicle defines a model where the set of formulas satisfied by each local configuration is exactly equal to the MCS assigned to that local configuration by the chronicle.

To show that a consistent formula α_0 is satisfiable, we shall incrementally construct a chronicle structure which will eventually be perfect and yield a model for α_0.

Before defining chronicles and perfect chronicle structures, we need to define four binary "semantic" relations over MCSs as follows.

Definition 5.1 *Let A, B be maximal consistent sets.*

(i) $A \mathrel{\tilde{<}} B \stackrel{\text{def}}{=} \{\alpha \mid \Box\alpha \in A\} \subseteq B$.

(ii) $A \mathrel{\tilde{\#}} B \stackrel{\text{def}}{=} \{\alpha \mid \dot{\triangledown}\alpha \in A\} \subseteq B$.

(iii) $A \mathrel{\tilde{co}} B \stackrel{\text{def}}{=} \{\alpha \mid \dot{\triangle}\alpha \in A\} \subseteq B$.

(iv) $A \mathrel{\tilde{\#}_\mu} B \stackrel{\text{def}}{=} \{\alpha \mid \dot{\triangledown}_\mu\alpha \in A\} \subseteq B$.

The next result follows from standard arguments, using axiom A5 for (i) and axiom A4 for (ii) to (iv).

Proposition 5.2 *Let A, B be maximal consistent sets.*

$$
\begin{array}{llllll}
(i) & A \mathrel{\tilde{<}} B & \textit{iff} & \{\Diamond\beta \mid \beta \in B\} \subseteq A & \textit{iff} & \{\beta \mid \boxminus\beta \in B\} \subseteq A \\
 & & \textit{iff} & \{\Diamond\mkern-12mu-\,\alpha \mid \alpha \in A\} \subseteq B. & & \\
(ii) & A \mathrel{\tilde{\#}} B & \textit{iff} & \{\triangledown\beta \mid \beta \in B\} \subseteq A & \textit{iff} & \{\beta \mid \dot{\triangledown}\beta \in B\} \subseteq A \\
 & & \textit{iff} & \{\triangledown\alpha \mid \alpha \in A\} \subseteq B. & & \\
(iii) & A \mathrel{\tilde{co}} B & \textit{iff} & \{\triangle\beta \mid \beta \in B\} \subseteq A & \textit{iff} & \{\beta \mid \dot{\triangle}\beta \in B\} \subseteq A \\
 & & \textit{iff} & \{\triangle\alpha \mid \alpha \in A\} \subseteq B. & & \\
(iv) & A \mathrel{\tilde{\#}_\mu} B & \textit{iff} & \{\triangledown_\mu\beta \mid \beta \in B\} \subseteq A & \textit{iff} & \{\beta \mid \dot{\triangledown}_\mu\beta \in B\} \subseteq A \\
 & & \textit{iff} & \{\triangledown_\mu\alpha \mid \alpha \in A\} \subseteq B. & &
\end{array}
$$

We shall use $\tilde{>}$ to denote the relation $(\tilde{<})^{-1}$.

The semantic relations defined on MCSs are designed to behave like their counterparts in an event structure. We now verify that our definitions do have the desired effect.

Proposition 5.3

(i) $\tilde{<}$ *is transitive.*

(ii) $\tilde{\#}$, $\tilde{co}$ *and* $\tilde{\#}_\mu$ *are symmetric.*

(iii) $\tilde{\#}_\mu \subseteq \tilde{\#}$.

Proof (i) and (ii) follow at once from axioms A3 and A4. (iii) follows from A7.

□

We can now define the notion of a chronicle structure.

Definition 5.4 *A* chronicle structure *is a triple $CH = (ES, T, \mu)$ where*

(i) $ES = (E, <, \#)$ *is a well branching event structure with* $\#_\mu$ *as its minimal conflict relation.*

(ii) T *is a function, called a* chronicle *over* ES, *which assigns an MCS to each* $\downarrow e \in LC_{ES}$.

(iii) $\mu \subseteq \#_\mu$.

Abusing notation, we shall always talk of a chronicle T assigning an MCS to an event e rather than to the corresponding local configuration $\downarrow e$. Notice that we have augmented the chronicle structure with a parameter μ which singles out a special subset of $\#_\mu$. This will be needed because the relation $\#_\mu$ is not stable with respect to our incremental construction of chronicle structures. The exact role played by μ will become clearer when we describe the construction.

Next, we need the notion of coherence.

Definition 5.5 *Let $CH = (ES, T, \mu)$ be a chronicle structure, where $ES = (E, <, \#)$. CH is* coherent *iff it satisfies the following conditions.*

(i) $\forall e, e' \in E : e < e'$ *implies* $T(e) \tilde{<} T(e')$.

(ii) $\forall e, e' \in E : e \# e'$ *implies* $T(e) \tilde{\#} T(e')$.

(iii) $\forall e, e' \in E : e \ co \ e'$ *implies* $T(e) \ \tilde{co} \ T(e')$.

(iv) $\forall e, e' \in E : (e, e') \in \mu$ *implies* $T(e) \ \tilde{\#}_\mu \ T(e')$.

Coherence ensures that the chronicle is consistent with the underlying frame. Notice that for the $\#_\mu$ relation, the chronicle need only be coherent with respect to the special subset μ.

A coherent chronicle may still have "gaps". These gaps can be described in terms of "unfulfilled requirements", in the following sense.

Definition 5.6 *Let $CH = (ES, T, \mu)$ be a coherent chronicle structure, where $ES = (E, <, \#)$.*

(i) A live future requirement *is a pair $(e, \Diamond\alpha)$ such that $e \in E$ and $\Diamond\alpha \in T(e)$ and there does not exist e' such that $e < e'$ and $\alpha \in T(e')$.*

(ii) A live past requirement *is a pair $(e, \diamondminus\alpha)$ such that $e \in E$ and $\diamondminus\alpha \in T(e)$ and there does not exist e' such that $e' < e$ and $\alpha \in T(e')$.*

(iii) A live choice requirement *is a pair $(e, \bigtriangledown\alpha)$ such that $e \in E$ and $\bigtriangledown\alpha \in T(e)$ and there does not exist e' such that $e \# e'$ and $\alpha \in T(e')$.*

(iv) A live concurrent requirement *is a pair $(e, \triangle\alpha)$ such that $e \in E$ and $\triangle\alpha \in T(e)$ and there does not exist e' such that $e \ co \ e'$ and $\alpha \in T(e')$.*

(v) A live minimal choice requirement *is a pair $(e, \bigtriangledown_\mu\alpha)$ such that $e \in E$ and $\bigtriangledown_\mu\alpha \in T(e)$ and there does not exist e' such that $(e, e') \in \mu$ and $\alpha \in T(e')$.*

(vi) A live requirement *is a pair (e, β) such that (e, β) is a live future requirement or a live past requirement or a live choice requirement or a live concurrent requirement or a live minimal choice requirement.*

Once again, observe that the liveness of a minimal choice requirement is defined with respect to μ and not with respect to $\#_\mu$.

We can now state the criteria under which a chronicle structure is perfect.

Definition 5.7 *The chronicle structure $CH = (ES, T, \mu)$ is* perfect *iff it satisfies the following three conditions.*

(i) CH is coherent.

(ii) CH has no live requirements.

(iii) $\mu = \#_\mu$, where $\#_\mu$ is the minimal conflict relation of ES.

From the work of Burgess [1], it follows that demonstrating the satisfiability of a consistent formula reduces to the problem of constructing a certain type of perfect chronicle structure.

Lemma 5.8 *Let $CH = (ES, T, \mu)$ be a perfect chronicle structure, where $ES = (E, <, \#)$, and let α_0 be a formula such that for some $e_0 \in E$, $\alpha_0 \in T(e_0)$. Then α_0 is satisfiable.*

Proof Let $V_T : \mathcal{P} \to 2^{LC_{ES}}$ be given by:

$$\forall p \in \mathcal{P} : V_T(p) = \{\downarrow e \mid e \in E \ \text{and} \ p \in T(e)\}.$$

Consider the model $M_T = ((ES, LC_{ES}), V_T)$.

Claim $\forall \alpha \in \Phi : \forall e \in E : M_T, \downarrow e \models \alpha$ *iff* $\alpha \in T(e)$

Proof of Claim We omit the proof, which is a simple induction on the structure of α.

Now, since we know that $\alpha_0 \in T(e_0)$, we get $M_T, \downarrow e_0 \models \alpha_0$.

$\square$

For the rest of the section, we fix a consistent formula α_0. The aim is to construct a perfect chronicle structure $CH = (ES, T, \mu)$, with $ES = (E, <, \#)$, such that $\alpha_0 \in T(e_0)$ for some $e_0 \in E$.

We shall build up CH inductively. With this in mind, we fix a countably infinite set of events $\tilde{E} = \{e_0, e_1, e_2, \ldots\}$. Let $\{q_0, q_1, \ldots, q_x\}$ be the set of atomic propositions that appear in α_0. Fix an enumeration of $\mathcal{P}$ of the form $q_0, q_1, \ldots, q_x, p_0, p_1, \ldots$. Finally, fix an enumeration of $\tilde{E} \times \Phi$, where Φ is the set of formulas in our language.

For clarity of presentation, we shall omit a couple of lengthy proofs from this section. These proofs are provided in the Appendix. This section together with the Appendix provides a complete account of the inductive construction.

The general idea is to begin by extending α_0 to an MCS A_0. We can then define a chronicle structure $CH_0 = (ES_0, T_0, \emptyset)$, where $ES_0 = (\{e_0\}, \emptyset, \emptyset)$ and $T_0(e_0) = A_0$.

At stage i in our inductive construction, we shall eliminate one live requirement (e, β) from $CH_i = (ES_i, T_i, \mu_i)$, where $ES_i = (E_i, <_i, \#_i)$ and $e \in E_i$. In general, this will involve adding an event e_{i+1} to E_i. For example, if β is of the form $\triangle\beta_0$, we shall add an event e_{i+1} such that $e \; co \; e_{i+1}$ and extend T_i to T_{i+1} so that $\beta_0 \in T_{i+1}(e_{i+1})$ and $T_{i+1}(e) \; \tilde{co} \; T_{i+1}(e_{i+1})$.

When we add e_{i+1} to E_i, we also have to fix the relationship between e_{i+1} and all the other events in E_i. To then extend T_i to a coherent chronicle T_{i+1} on ES_{i+1}, we will need to selectively modify the MCSs assigned by T_i to the events in E_i. While doing this, we have to ensure that live requirements that have been killed at earlier steps remain killed. To achieve all this, we will need to "name" each event e_i that we add to our chronicle structure by the formula $\hat{p}_i$. Recall that for $p \in \mathcal{P}$, $\hat{p}$ abbreviates the formula $p \wedge \boxdot\neg p \wedge \boxminus\neg p \wedge \bigtriangledown\neg p \wedge \triangle\neg p$.

Keeping all this in mind, we shall assume that the chronicle structure $CH_i = (ES_i, T_i, \mu_i)$ constructed at stage i satisfies a set of inductive conditions $\mathcal{IC}$. $\mathcal{IC}$ captures the fact that E_i contains events from $\{e_0, e_1, \ldots, e_i\}$, each event being "named" by $\hat{p}_i$ — i.e. $\hat{p}_i \in T_i(e_i)$. $\mathcal{IC}$ also ensures that $\alpha_0 \in T_i(e_0)$ is maintained as an invariant condition, and that T_i is a coherent chronicle on ES_i. The detailed description of $\mathcal{IC}$ is given in the Appendix.

To ensure that CH_0 conforms to $\mathcal{IC}$, we have to be careful about extending α_0 to an MCS A_0. Since we want to have $T_0(e_0) = A_0$, we must ensure that $\hat{p}_0 \in A_0$.

Define the function $g_0 : \mathcal{P} \to \mathcal{P}$ as follows.

$$\forall p \in \mathcal{P} : g_0(p) = \begin{cases} p, & \textit{if } p \in \{q_0, q_1, \ldots, q_x\} \\ p_{i+1}, & \textit{if } p = p_i \textit{ for } i \geq 0 \end{cases}$$

In other words, replace p_0 by p_1, p_1 by p_2 ...but leave the atomic propositions $q_0, q_1, \ldots, q_x$ *untouched.* For $\alpha \in \Phi$, let α' abbreviate $g_0(\alpha)$.

Extend the consistent set $\{\alpha_0\}$ to a maximal consistent set A. Let $A' = g_0(A)$. Clearly A' is free of p_0.

Lemma 5.9 *A' is consistent and $\alpha_0 \in A'$. Moreover $A' \cup \{\hat{p}_0\}$ is consistent.*

Proof Suppose A' is not consistent. Then we must have $\vdash \neg\alpha'$ for some $\alpha' \in A'$. Let $h_0 : \mathcal{P} \to \mathcal{P}$ be given by:

$$\forall p \in \mathcal{P} : h_0(p) = \begin{cases} p, & \textit{if } p \in \{q_0, q_1, \ldots, q_x\} \cup \{p_0\} \\ p_{i-1}, & \textit{if } p = p_i \textit{ for } i \geq 1 \end{cases}$$

It is easy to check that $h_0(\alpha') = \alpha$. Then, by the derived inference rule SUB, $\vdash \neg\alpha$, which is a contradiction because $\alpha \in A$. Hence A' is consistent. Since $g_0(\alpha_0) = \alpha_0$, $\alpha_0 \in A'$.

Next, we must verify that $A' \cup \hat{p}_0$ is consistent. Suppose not. Then $\vdash \hat{p}_0 \supset \neg\alpha'$ for some $\alpha' \in A'$. By the choice of g_0, α' is free of p_0. Hence, by UNIQ, $\vdash \neg\alpha'$, which contradicts the consistency of A'.

$\square$

Now extend $A' \cup \{\hat{p}_0\}$ to an MCS A_0. Set $CH_0 = (ES_0, T_0, \emptyset)$, where $ES_0 = (\{e_0\}, \emptyset, \emptyset)$ and $T_0(e_0) = A_0$. We can verify that CH_0 satisfies the inductive conditions $\mathcal{IC}$.

We now proceed to construct CH_{i+1} from CH_i. The basic idea is to "kill" some live requirement (e, β) present in CH_i. While doing so, we do not wish to "disturb" the fact that $\alpha_0 \in T_i(e_0)$ and $\hat{p}_j \in T_i(e_j)$ for each $e_j \in E_i, 0 \le j \le i$. At the same time, in case the event e_{i+1} is added to E_i to kill (e, β), it should be named $\hat{p}_{i+1}$, for which we shall have to "shift up" propositions $p_{i+1}, p_{i+2} \ldots$ as we did for p_0 using g_0. Consequently, we shall demand two things of the live requirement (e, β) chosen to be "killed".

(i) At most the propositions $q_0, q_1, \ldots, q_x, p_0, p_1, \ldots, p_i$ may appear in β.

(ii) Among all the live requirements in CH_i which satisfy (i), (e, β) has the least index in the enumeration which we have fixed for $\tilde{E} \times \Phi$.

To be precise, for $j \ge 0$ let $\Sigma_j = \{q_0, q_1, \ldots, q_x\} \cup \{p_0, p_1, \ldots, p_j\}$. Let Φ_j be the least subset of Φ containing Σ_j such that if γ and δ are in Φ_j then so are $\neg\gamma, \gamma \vee \delta, \Box\gamma, \boxminus\gamma, \triangledown\gamma, \vartriangle\gamma$ and $\triangledown_\mu\gamma$. Clearly $\Phi = \bigcup_{j \ge 0} \Phi_j$.

Among all the live requirements of the form (e', β') in CH_i which satisfy the condition $\beta' \in \Phi_i$, let (e, β) be the live requirement with the least index (in the enumeration we have fixed for $\tilde{E} \times \Phi$). If there are no such live requirements, set $CH_{i+1} = CH_i$. Otherwise, assume that $e = e_j$, $j \in \{0, 1, \ldots, i\}$. We shall deal in detail with the case where β is of the form $\triangledown_\mu \beta_0$.

The live requirement $(e_j, \triangledown_\mu \beta_0)$ in CH_i can be killed in two different ways. In fact, this is the reason why we give a detailed treatment of this case. For all other cases (e, β) can be killed in only one way. This will become clearer as we proceed.

The first case corresponds to the situation where there already exists $e_k \in E_i$ such that $\beta_0 \in T_i(e_k)$ and $e_j \ \#_{\mu\,i}\ e_k$, but $(e_j, e_k) \notin \mu_i$. In other words, e_k was added at stage $k < i$ to kill some other live requirement, and (e_j, e_k) "accidentally" happens to belong to $\#_{\mu\,i}$. In general, we cannot assume that the event e_k kills the live requirement $(e_j, \triangledown_\mu \beta_0)$ because $\#_{\mu\,i}$ is not stable with respect to our inductive construction. For example, given $(e_j, e_k) \in \#_{\mu\,i}$, we may add e_{i+1} to E_i to kill a past requirement of the form $(e_k, \diamondminus\gamma)$, and set $e_{i+1} <_{i+1} e_k$ and $e_{i+1} \ \#_{i+1}\ e_j$ whereby $(e_j, e_k) \notin \#_{\mu\,i+1}$. As a result, if the liveness of minimal choice requirements were to be defined with respect to the relation $\#_{\mu\,i}$, requirements assumed to be killed at earlier stages in the construction may become live again at later stages.

This is precisely the reason why we identify a special subset of $\#_{\mu\,i}$ by μ_i. From the way we build up μ_i, it will turn out that if an element of $\#_{\mu\,i}$ is included in μ_i, then it will *always* remain in $\#_{\mu\,n}$ for all $n > i$. Thus, by killing a minimal choice requirement using an element in μ_i, we can ensure that the requirement never becomes live again.

On the other hand, it will turn out that $R_i \subseteq R_{i+1}$ for $R \in \{<, >, \#, co\}$. The inductive conditions $\mathcal{IC}$ will then ensure that for live requirements of the form $(e_j, \Diamond\beta_0)$, $(e_j, \diamondminus\beta_0)$, $(e_j, \triangledown\beta_0)$ and $(e_j, \vartriangle\beta_0)$, once the requirement has been killed at some stage k, the requirement never becomes live again at any later stage.

Suppose then that there exists $e_k \in E_i$ such that $e_j \ \#_{\mu\,i}\ e_k$ and $\beta_0 \in T_i(e_k)$ and, further, $T_i(e_j) \ \tilde{\#}_\mu\ T_i(e_k)$. Set $E_{i+1} = E_i$, $<_{i+1} = <_i$, $\#_{i+1} = \#_i$ and $T_{i+1} = T_i$. Finally, set $\mu_{i+1} = \mu_i \cup \{(e_j, e_k), (e_k, e_j)\}$.

It is easy to verify that CH_{i+1} is a well defined chronicle structure. We know that $ES_{i+1} = ES_i$ and $T_{i+1} = T_i$, so all we have to check is that $\mu_{i+1} \subseteq \#_{\mu\,i+1}$. Since $\mu_i \subseteq \#_{\mu\,i}$ by the inductive hypothesis and the new elements added to μ_i to obtain μ_{i+1} were assumed to be in $\#_{\mu\,i}$, we know that $\mu_{i+1} \subseteq \#_{\mu\,i}$. But clearly $\#_{\mu\,i} = \#_{\mu\,i+1}$ and so $\mu_{i+1} \subseteq \#_{\mu\,i+1}$ as well.

It is straightforward to check that CH_{i+1} satisfies the inductive conditions $\mathcal{IC}$. Clearly $(e_j, \triangledown_\mu\beta_0)$ is no longer a live requirement in CH_{i+1}. It is also easy to see that for all $\beta' \in \Phi_i$ and for all $e_k \in E_i$, if (e_k, β') is not a live requirement in CH_i then it is not a live requirement in CH_{i+1} either.

Now we consider the more difficult case. Assume that there does not exist $e_k \in E_i$ with $e_j \ \#_{\mu\,i}\ e_k$ such that $\beta_0 \in T_i(e_k)$ and $T_i(e_j) \ \tilde{\#}_\mu\ T_i(e_k)$.

To kill the live requirement $(e_j, \triangledown_\mu\beta_0)$ in this case, we will extend ES_i to ES_{i+1} so that $E_{i+1} = E_i \cup \{e_{i+1}\}$ and $R_{i+1} \cap (E_i \times E_i) = R_i$ for $R \in \{<, >, \#, co\}$. We will transform T_i to T_{i+1} in such a way that $T_{i+1}(e_k) \cap \Phi_i = T_i(e_k) \cap \Phi_i$ for all $e_k \in E_i$, $0 \le k \le i$. We will ensure that $(e_j, e_{i+1}) \in \mu_{i+1}$ and $\beta_0 \in T_{i+1}(e_{i+1})$ so that $(e_j, \triangledown_\mu\beta_0)$ is not a live requirement in CH_{i+1}. Naturally, in all this we must ensure that CH_{i+1} satisfies the inductive assumptions $\mathcal{IC}$.

The actual steps involved in the construction are as follows.

First we must "free" p_{i+1} so that e_{i+1} can be labelled with $\hat{p}_{i+1}$. To do this, define the function

$g_{i+1} : \mathcal{P} \rightarrow \mathcal{P}$ as follows:

$$\forall p \in \mathcal{P} : g_{i+1}(p) = \begin{cases} p, & \textit{if } p \in \{q_0, q_1, \ldots, q_x\} \cup \{p_0, p_1, \ldots, p_i\} \\ p_{k+1}, & \textit{if } p = p_k \textit{ for } k \geq i+1 \end{cases}$$

In other words, g_{i+1} replaces p_{i+1} by p_{i+2}, p_{i+2} by p_{i+3} ...but leaves $q_0, q_1, \ldots, q_x$ and $p_0, p_1, \ldots, p_i$ *untouched.* For $\alpha \in \Phi$, let α' abbreviate $g_{i+1}(\alpha)$.

For convenience, let A_k denote $T_i(e_k)$ and let A'_k denote $g_{i+1}(A_k)$ for $e_k \in E_i$, $0 \leq k \leq i$. By Lemma 5.9, A'_k is consistent for all such k.

The next step in the construction is to transform T_i to T_{i+1}. We begin by creating a consistent set of formulas containing β_0 and $\hat{p}_{i+1}$ which is in $\tilde{\#}_\mu$ with A'_j.

Lemma 5.10 $A'_{i+1} = \{\alpha' \mid \triangledown_\mu \alpha' \in A'_j\} \cup \{\beta_0\} \cup \{\hat{p}_{i+1}\}$ *is consistent.*

Proof Suppose not. Then $\vdash \hat{p}_{i+1} \supset \neg(\alpha' \wedge \beta_0)$ for some α' such that $\triangledown_\mu \alpha \in A_j$. The definition of g_{i+1} ensures that α' is free of p_{i+1}. Since we chose $(e_j, \triangledown_\mu \beta_0)$ as the live requirement to be killed at stage i, we know that $\triangledown_\mu \beta_0 \in \Phi_i$ and so $\beta_0 \in \Phi_i$ as well. But $p_{i+1} \notin \Sigma_i$, so β_0 is also free of p_{i+1}. Hence, by the rule UNIQ, we get $\vdash \neg(\alpha' \wedge \beta_0)$. By TG, we then get $\vdash \triangledown_\mu \neg(\alpha' \wedge \beta_0)$ and hence $\vdash \neg\triangledown_\mu(\alpha' \wedge \beta_0)$. Now, consider the function $h_{i+1} : \mathcal{P} \rightarrow \mathcal{P}$ given by:

$$\forall p \in \mathcal{P} : h_{i+1}(p) = \begin{cases} p, & \textit{if } p \in \{q_0, q_1, \ldots, q_x\} \cup \{p_0, p_1, \ldots, p_{i+1}\} \\ p_{k-1}, & \textit{if } p = p_k \textit{ for } k \geq i+2 \end{cases}$$

Since $h_{i+1}(\neg\triangledown_\mu(\alpha' \wedge \beta_0)) = \neg\triangledown_\mu(\alpha \wedge \beta_0)$, using the derived rule SUB we get $\vdash \neg\triangledown_\mu(\alpha \wedge \beta_0)$. But we began with $\triangledown_\mu \alpha \in A_j$. Since $\triangledown_\mu \beta_0 \in A_j$ as well and A_j is an MCS, by TK3 we get $\triangledown_\mu(\alpha \wedge \beta_0) \in A_j$. So $\triangledown_\mu(\alpha \wedge \beta_0)$ is consistent, which is a contradiction.

□

Extend A'_{i+1} to an MCS B_{i+1}. We now have to go back and extend A'_j to an MCS that is in $\tilde{\#}_\mu$ with B_{i+1}.

Lemma 5.11 $A'_j \cup \{\beta \mid \triangledown_\mu \beta \in B_{i+1}\}$ *is consistent.*

Proof Suppose not. Then $\vdash \alpha' \supset \neg\beta$ for some $\alpha \in A_j$ and $\triangledown_\mu \beta \in B_{i+1}$. But then, by axiom A4(ii), $\triangledown_\mu \triangledown_\mu \alpha \in A_j$ and hence $\triangledown_\mu \triangledown_\mu \alpha' \in A'_j$. As a result, $\triangledown_\mu \alpha' \in A'_{i+1} \subseteq B_{i+1}$. Now, applying DRK(ii) to $\vdash \alpha' \supset \neg\beta$, we get $\vdash \triangledown_\mu \alpha' \supset \triangledown_\mu \neg\beta$, which then implies that $\triangledown_\mu \neg\beta \in B_{i+1}$. This contradicts the assumption that $\triangledown_\mu \beta \in B_{i+1}$.

□

Now extend $A'_j \cup \{\beta \mid \triangledown_\mu \beta \in B_{i+1}\}$ to an MCS B_j. It is easy to verify that $B_j \;\tilde{\#}_\mu\; B_{i+1}$.

T_{i+1} will assign B_{i+1} to e_{i+1} and B_j to e_j. We must now extend T_{i+1} to the events in $E_i - \{e_j\}$. We shall deal with the concrete case where $e_k \in E_i - \{e_j\}$ and $e_k <_i e_j$. Let $A^*_k = \{\beta \mid \boxminus\beta \in B_j\}$.

Lemma 5.12 $A'_k \cup A^*_k$ *is consistent.*

Proof Suppose not. Then $\vdash \alpha' \supset \neg\beta$ for some $\alpha \in A_k$ and $\boxminus\beta \in B_j$. Since T_i was a coherent chronicle on ES_i, we must have had $A_k \tilde{<} A_j$. So, $\diamondminus\alpha \in A_j$ and hence $\diamondminus\alpha' \in A'_j \subseteq B_j$. Applying DRK(ii) to $\vdash \alpha' \supset \neg\beta$ we get $\vdash \diamondminus\alpha' \supset \diamondminus\neg\beta$, which implies that $\diamondminus\neg\beta \in B_j$. This contradicts the assumption $\boxminus\beta \in B_j$.

□

Extend $A'_k \cup A^*_k$ to an MCS B_k. It is straightforward to verify that $B_k \tilde{<} B_j$.

For the cases $e_k >_i e_j$, $e_k \#_i e_j$ and $e_k \; co_i \; e_j$, we set A^*_k to $\{\beta \mid \Box\beta \in B_j\}$, $\{\beta \mid \triangledown\beta \in B_j\}$ and $\{\beta \mid \triangle\beta \in B_j\}$ respectively. In each case, by an argument similar to Lemma 5.12, $A'_k \cup A^*_k$ is consistent and can be extended to an MCS B_k.

At this stage, for each $e_k \in E_{i+1}$, $0 \leq k \leq i+1$, we have managed to find an MCS B_k which shall use to define T_{i+1}. We can now define $CH_{i+1} = (ES_{i+1}, T_{i+1}, \mu_{i+1})$, where $ES_{i+1} = (E_{i+1}, <_{i+1}, \#_{i+1})$, as follows:

(i) $E_{i+1} = E_i \cup \{e_{i+1}\}$.

(ii) $<_{i+1} = <_i \;\cup\; \{(e_k, e_{i+1}) \mid e_k \in E_i, 0 \leq k \leq i, \text{ and } B_k \tilde{<} B_{i+1}\}$.
$\;\cup\; \{(e_{i+1}, e_k) \mid e_k \in E_i, 0 \leq k \leq i, \text{ and } B_{i+1} \tilde{<} B_k\}$

(iii) $>_{i+1} = (<_{i+1})^{-1}$.

(iv) $\#_{i+1} = \#_i \cup \{(e_k, e_{i+1}), (e_{i+1}, e_k) \mid e_k \in E_i, 0 \leq k \leq i, \text{ and } B_k \;\tilde{\#}\; B_{i+1}\}$.

(v) $co_{i+1} = (E_{i+1} \times E_{i+1}) - (<_{i+1} \cup >_{i+1} \cup \#_{i+1} \cup id_{i+1})$.

(vi) $\forall e_k \in E_{i+1} : T_{i+1}(e_k) = B_k$.

(vii) $\mu_{i+1} = \begin{cases} \mu_i \cup \{(e_j, e_{i+1}), (e_{i+1}, e_j)\} & \text{if the live requirement considered at stage } i \text{ was of the form } (e_j, \bigtriangledown_\mu \beta_0) \text{ (which we are assuming to be the case)} \\ \mu_i & \text{otherwise} \end{cases}$

We have to do a fair amount of work to prove the next result. The details are given in the appendix.

Lemma 5.13 *CH_{i+1} is a well defined chronicle structure which satisfies the inductive conditions $\mathcal{IC}$.*

Clearly, by our construction, the live requirement $(e_j, \bigtriangledown_\mu \beta_0)$ that we selected for killing in CH_i is not live in CH_{i+1}. To ensure that our inductive construction works properly, we finally have to show that all live requirements killed at earlier stages remain "dead" in CH_{i+1}.

Lemma 5.14 $\forall e \in E_i : \forall \alpha \in \Phi_i : (e, \alpha)$ *is not a live requirement in CH_i implies that it is not a live requirement in CH_{i+1}.*

Proof This follows from the following two observations.

(i) In extending CH_i to CH_{i+1} we preserve the semantic relationships between the MCSs assigned by T_i.

(ii) Formulas from Φ_i are left untouched in the MCSs assigned by T_i during the process of "shifting up" propositions to accommodate $\hat{p}_{i+1}$ in $T_{i+1}(e_{i+1})$.

□

The other four types of live requirements — $(e_j, \Diamond\beta_0), (e_j, \diamondminus\beta_0), (e_j, \bigtriangledown\beta_0), (e_j, \triangle\beta_0)$ — are also killed by adding an event e_{i+1} to E_i and extending ES_i to ES_{i+1} and T_i to T_{i+1} in a suitable manner. The procedure is almost identical to that described above for the case $(e_j, \bigtriangledown_\mu \beta_0)$. Virtually the only change to be made is in the step where B_{i+1} is constructed. For live requirements of the form $(e_j, \Diamond\beta_0), (e_j, \diamondminus\beta_0), (e_j, \bigtriangledown\beta_0)$ and $(e_j, \triangle\beta_0)$, B_{i+1} and B_j must be suitably defined so that $B_j \;\tilde{<}\; B_{i+1}, B_j \;\tilde{>}\; B_{i+1}, B_j \;\tilde{\#}\; B_{i+1}$ and $B_j \;\tilde{co}\; B_{i+1}$ respectively. The new event e_{i+1} can then be "hooked up" to the other events in E_i exactly as described above, yielding a chronicle structure CH_{i+1} in which the requirement considered at stage i is no longer live. In all four cases μ_i is left untouched — i.e $\mu_{i+1} = \mu_i$.

Thus, we can construct an infinite sequence of chronicle structures $CH_0, CH_1, \ldots$ so that each member of the sequence satisfies the inductive conditions $\mathcal{IC}$. Now define $CH = (ES, T, \mu)$ where:

- $ES = (E, <, \#)$, with $E = \bigcup_{i \geq 0} E_i, <= \bigcup_{i \geq 0} <_i$, and $\# = \bigcup_{i \geq 0} \#_i$.
- $T : E \rightarrow 2^\Phi$ is given by:
$$\forall e_i \in E : T(e_i) = \bigcup \{T_j(e_i) \cap \Phi_j \mid j \geq i\}.$$
- $\mu = \bigcup_{i \geq 0} \mu_i$.

We can then verify the following result, whose proof is given in the Appendix.

Lemma 5.15 *$CH = (ES, T, \mu)$ is a perfect chronicle structure in which $\alpha_0 \in T(e_0)$.*

From the preceding lemma and Lemma 5.8 it follows that every consistent formula α_0 is satisfiable. We have thus established our main result.

Theorem 5.16 (Completeness) *If $\models \alpha$ then $\vdash \alpha$.*

It is straightforward to extend this result to obtain completeness for finite theories.

Corollary 5.17 *Let $\{\alpha_1, \alpha_2, \ldots, \alpha_n\}$ be a finite set of formulas. If $\{\alpha_1, \alpha_2, \ldots, \alpha_n\} \models \beta$ then $\{\alpha_1, \alpha_2, \ldots, \alpha_n\} \vdash \beta$.*

Proof By Theorem 3.1, we know that $\{\alpha_1, \alpha_2, \ldots, \alpha_n\} \models \beta$ iff $\models \mathcal{E}(\alpha_1 \wedge \alpha_2 \wedge \ldots \wedge \alpha_n) \supset \beta$. Since our axiomatization is complete, we know that $\models \mathcal{E}(\alpha_1 \wedge \alpha_2 \wedge \ldots \wedge \alpha_n) \supset \beta$ implies $\vdash \mathcal{E}(\alpha_1 \wedge \alpha_2 \wedge \ldots \wedge \alpha_n) \supset \beta$. By Theorem 4.2, $\{\alpha_1, \alpha_2, \ldots, \alpha_n\} \vdash \beta$ iff $\vdash \mathcal{E}(\alpha_1 \wedge \alpha_2 \wedge \ldots \wedge \alpha_n) \supset \beta$ and we are done.

$\square$

6 Discussion

In this paper we have obtained a sound and complete axiomatization of the class of well branching prime event structures. We have achieved this using a temporal logic which appears to be rich in expressive power.

Our completeness proof leans heavily on the techniques developed in our earlier work [11]. However, the logic considered there did not have the $\bigtriangledown_\mu$ operator. Moreover the frames for the logic consisted of all prime event structures. As noted earlier, what we really need is a logic that characterizes finitary prime event structures. It is easy to check that every finitary prime event structure is also well branching. Thus the present work, in comparison to [11] represents advances on two fronts; the set of frames has been suitably shrunk (though not as much as we would have liked) and the logical language, due to the addition of the $\bigtriangledown_\mu$ operator, is provably more expressive [9]. The present completeness proof also requires some new ideas and is a lot more delicate.

Partial orders have also been considered as structures for temporal logics by — among others — Pinter and Wolper [17] and Katz and Peled [5]. In these studies assertions are tied directly to the global states of the system. Also, there are no modalities for directly expressing concurrency and conflict.

In [6, 7] and in the much improved [8] a subclass of event structures is used to model systems of communicating sequential agents. The logical language consists of future and past modalities indexed by the names of the agents. There are no modalities to express conflict and concurrency.

As mentioned earlier, Penczek [15, 16] has carried out closely related work. He was the first one to identify the $\bigtriangledown$ operator and use it to exploit the important notion of a run. In [15] a sound and complete axiomatization of (all) prime event structures is given. In our notation, the logical language consists of the $\Box$, $\boxminus$ and the $\bigtriangledown$ operators with neither the $\triangle$ operator nor the $\bigtriangledown_\mu$ operator. Moreover, a special proposition ρ is reserved in the language and is used to mark the runs of an event structure. A frame consists of an event structure together with a run. The valuation function is required to ensure that the special proposition ρ is assigned the run associated with the frame.

It is easy to check that we could, if we wished to, follow the same route using well branching event structures. As in [15] we would add $\rho \equiv \bigtriangledown\neg\rho$ as an axiom to our system. We would then get a sound and complete axiomatization (with respect to the extended notion of frames and models) in the presence of the additional operators $\triangle$ and $\bigtriangledown_\mu$. Consequently the resulting logic would be more powerful in comparison to [15] with the added advantage that non-well branching event structures would not arise in the process of producing models for consistent formulas.

We feel however that it is better to retain a "pure" notion of frames and embed the notion of a run into theories concerning specific applications as indicated in Section 3. This is so because this nice idea of Penczek to logically talk about runs can be extended in a number of useful ways. And these extensions, if brought into the "core" language, would lead to more and more elaborate notions of frames. It would then be difficult to distinguish the foundation from the superstructure, so to speak.

For instance one could reserve a special proposition τ and use it mark the sequential components of an event structure. By a sequential component we mean a maximal co-free subset of the events. The axiom $\tau \equiv \triangle\neg\tau$ will then capture this notion. One could reserve two more propositions λ and χ and use them to mark the maximal chains and anti-chains of an event structure. The respective characteristic axioms would be $\lambda \equiv \triangle\neg\lambda \wedge \bigtriangledown\neg\lambda$ and $\chi \equiv \Box\neg\chi \wedge \boxminus\neg\chi$. We could go further down this road but we will stop here. The point however is that the logic presented here appears to be well suited for capturing many interesting behavioural notions concerning event structures.

Turning now to [16], Penczek also introduces what he calls an immediate conflict operator which we will denote here as $\bigtriangledown_m$. However, all that is required of the corresponding frame relation $\#_m$ is that it should generate the conflict relation $\#$ via the causality relation in the obvious sense. As a result, for a fixed $\#$ very many subsets of $\#$ could play the role of $\#_m$. In particular, $\#_m$ could be $\#$ itself! Stated differently, all that is demanded semantically from $\triangledown_m$ is that the axiom A7 be sound. The problem of course is that A8(i) which is in some sense the characteristic axiom for the minimal conflict relation cannot be formulated in [16] due to the absence of the $\triangle$ operator.

It is worth pointing out in this connection that the term immediate conflict is suggested by net theory. With each event structure, one can associate a special kind of net called an occurrence net [18]. For occurrence nets the notion of two events being in immediate conflict is well-defined — they should share a pre-condition. However, the means for associating occurrence nets with event structures suggested in [18], and also in [13], would identify $\#_m$ with $\#$. Thus the notion of immediate conflict, applied to prime event structures, is not an interesting behavioural notion.

At present we do not know any thing about the decidability of our system. The finite model property does not hold and hence standard filtration techniques are not applicable. It seems difficult to identify a suitable notion of a pseudo model in which the causality relation will be just a preorder. Stated differently, what is lacking is a finite representation of event structures that are possibly infinite but "regular". The model checking problem is also at present not meaningful because we do not know what finite but "cyclic" event structures look like. Notice in this connection that the weaker system of [15] is shown by Penczek to be decidable.

As mentioned in Section 1, we lack the means for handling labelled event structures. One possibility would be to introduce next state and previous state operators and use reserved propositions to code up the label set. Once again due to the lack of a suitable notion of a pseudo model, standard techniques cannot be applied to obtain completeness.

Finally the problem of characterizing finitary event structures seems to be a difficult one. The standard well foundedness axiom for irreflexive frames $\diamondminus\alpha \supset \diamondminus(\alpha \wedge \boxminus\neg\alpha)$ will certainly be sound for finitary event structures. However, the past of an event will in general be a partially ordered set rather than a totally ordered set. Hence even well foundedness seems to be hard to capture let alone finitariness.

Acknowledgments Mogens Nielsen independently came up with the idea of replacing the conflict operator by the more useful minimal conflict operator. We thank him for useful discussions.

U. Goltz and W. Penczek found a mistake in an earlier version of this work. In the process of fixing this mistake, we were led to the well branching property.

We also thank the anonymous referees for their constructive comments and suggestions for improving the presentation.

References

[1] Burgess, J. P., Basic Tense Logic, In *The Handbook of Philosophical Logic, Vol. II*, eds. Gabbay, D. and Guenthner, F., D. Reidel Publishing Company, Dordrecht, The Netherlands (1984) 89-133.

[2] Burgess, J. P., Decidability for Branching Time, *Studia Logica*, **XXXIX**, 2/3 (1980) 203-218.

[3] Chandy, K. M. and Lamport, L., Distributed Snapshots: Determining Global States of Distributed Systems, *ACM Transactions on Computer Systems*, **3**, 1 (1985) 63-75.

[4] Hughes, G. E. and Creswell, M. J., *An Introduction to Modal Logic*, Methuen and Co., London (Reprinted in 1982).

[5] Katz, S. and Peled, D., An Efficient Verification Method for Parallel and Distributed Programs, *Springer Lecture Notes in Computer Science* **354** (1989) 489-507.

[6] Lodaya, K and Thiagarajan, P. S., A Modal Logic for a Subclass of Event Structures, *Springer Lecture Notes in Computer Science* **267** (1987) 290-303.

[7] Lodaya, K. and Thiagarajan, P. S., A Correction to "A Modal Logic for a Subclass of Event Structures", *Report* DAIMI-PB-275, Computer Science Department, Aarhus University, Aarhus, Denmark (1989).

[8] Lodaya, K., Ramanujam, R. and Thiagarajan, P. S., Temporal Logics for Communicating Sequential Agents, *Report* IMSc/89/15, The Institute of Mathematical Sciences, Madras, India (1989).

[9] Mukund, M., Expressiveness and Completeness of a Logic for Well Branching Prime Event Structures, *Report* TCS-90-1, School of Mathematics, SPIC Science Foundation, Madras, India (1990).

[10] Mukund, M. and Thiagarajan, P. S., A Petri Net Model of Asynchronously Communicating Sequential Processes, In *Perspectives in Theoretical Computer Science — Commemorative Volume for Gift Siromoney*, ed. Narasimhan, R., World Scientific, Singapore (1989) 165-198.

[11] Mukund, M. and Thiagarajan, P. S., An Axiomatization of Event Structures, *Springer Lecture Notes in Computer Science* **405** (1989) 143-160.

[12] Mukund, M. and Thiagarajan, P. S., An Axiomatization of Well Branching Prime Event Structures, *Report* TCS-90-2, School of Mathematics, SPIC Science Foundation, Madras, India (1990).

[13] Nielsen, M., Plotkin, G. and Winskel, G., Petri Nets, Event Structures and Domains: Part I, *Theoretical Computer Science*, **13**, 1 (1980) 86-108.

[14] Nielsen, M., Rozenberg, G. and Thiagarajan, P. S., Behavioural Notions for Elementary Net Systems, *Distributed Computing*, **4**, 1 (1990) 45-57.

[15] Penczek, W., A Temporal Logic For Event Structures, *Fundamenta Informaticae*, **XI**, (1988) 297-326.

[16] Penczek, W., A Temporal Logic for the Local Specification of Concurrent Systems, In *Information Processing 89*, ed. Ritter, G. X., Elsevier Science Publishers B.V. (1989) 857-862.

[17] Pinter, S. and Wolper, P., A Temporal Logic for Reasoning About Partially Ordered Computations, *Proc. 3^{rd} ACM Symposium on Principles of Distributed Computing*, Vancouver (1984) 28-37.

[18] Winskel, G., Event Structures, *Springer Lecture Notes in Computer Science* **255** (1987) 325-392.

A Appendix

Here we fill in the details of the completeness proof which are missing in Section 5. In this appendix, we describe in detail the inductive hypotheses $\mathcal{IC}$ which are used in our incremental construction of a perfect chronicle and then provide the proofs of Lemmas 5.13 and 5.15.

First, we need some preliminary results concerning the semantic relations defined on MCSs. We begin with a characterization of $\tilde{\#}_\mu$ in terms of $\tilde{\#}$.

Lemma A.1 *Suppose A and B are MCSs such that $A\ \tilde{\#}\ B$ and p is an atomic proposition such that $\hat{p} \in B$ and $\bigtriangledown_\mu \hat{p} \in A$. Then $A\ \tilde{\#}_\mu\ B$.*

Proof It suffices to show that $\{\bigtriangledown_\mu \beta \mid \beta \in B\} \subseteq A$. Suppose that $\beta \in B$ and $\bigtriangledown_\mu \beta \notin A$. Then, since A is an MCS, $\neg \bigtriangledown_\mu \beta \in A$ and hence $\bigtriangledown_\mu \neg \beta \in A$. We know that $\bigtriangledown_\mu \hat{p} \in A$, so in fact $\bigtriangledown_\mu \hat{p} \wedge \bigtriangledown_\mu \neg\beta \in A$. Hence, by TK3, $\bigtriangledown_\mu(\hat{p} \wedge \neg\beta) \in A$ and so, by axiom A7, $\bigtriangledown(\hat{p} \wedge \neg\beta) \in A$ as well. But, by T$\hat{p}$(v), $\bigtriangledown(\hat{p} \wedge \neg\beta) \supset \bigtriangledown(\hat{p} \supset \neg\beta)$ so we must have $\bigtriangledown(\hat{p} \supset \neg\beta) \in A$. Then, since $A\ \tilde{\#}\ B$, $(\hat{p} \supset \neg\beta) \in B$ by the definition of $\tilde{\#}$. But $(\hat{p} \supset \neg\beta) \in B$ and $\hat{p} \in B$ implies that $\neg\beta \in B$, which is a contradiction.

$\square$

Let $E = \{e_1, e_2, e_3\}$ be a set of three events. Suppose we choose relations $R_{12}, R_{13} \in \{<, >, \#, co\}$ and let $e_1\ R_{12}\ e_2$ and $e_1\ R_{13}\ e_3$. It is easy to verify that we can always complete the "triangle", choosing an appropriate $R_{23} \in \{<, >, \#, co\}$, such that setting $e_2\ R_{23}\ e_3$ yields a valid event structure $ES = (E, <, \#)$.

The next lemma establishes an analogous result for MCSs.

Lemma A.2 *Let A, B and C be* distinct *MCSs such that $A\ \tilde{R}_1\ B$ and $A\ \tilde{R}_2\ C$ for $\tilde{R}_1, \tilde{R}_2 \in \{\tilde{<}, \tilde{>}, \tilde{\#}, \tilde{co}\}$ as specified in Table 1. Then it must be the case that $B\ \tilde{R}_3\ C$, $\tilde{R}_3 \in \{\tilde{<}, \tilde{>}, \tilde{\#}, \tilde{co}\}$, for at least one of the options specified for $\tilde{R}_3$ in the corresponding row of Table 1.*

$\tilde{R}_1$	$\tilde{R}_2$	$\tilde{R}_3$
$\tilde{<}$	$\tilde{<}$	$\tilde{<}, \tilde{>}, \tilde{\#}, \tilde{co}$
$\tilde{<}$	$\tilde{>}$	$\tilde{>}$
$\tilde{<}$	$\tilde{\#}$	$\tilde{\#}$
$\tilde{<}$	$\tilde{co}$	$\tilde{>}, \tilde{\#}, \tilde{co}$
$\tilde{>}$	$\tilde{>}$	$\tilde{<}, \tilde{>}, \tilde{co}$
$\tilde{>}$	$\tilde{\#}$	$\tilde{<}, \tilde{\#}, \tilde{co}$
$\tilde{>}$	$\tilde{co}$	$\tilde{<}, \tilde{co}$
$\tilde{\#}$	$\tilde{\#}$	$\tilde{<}, \tilde{>}, \tilde{\#}, \tilde{co}$
$\tilde{\#}$	$\tilde{co}$	$\tilde{>}, \tilde{\#}, \tilde{co}$
$\tilde{co}$	$\tilde{co}$	$\tilde{<}, \tilde{>}, \tilde{\#}, \tilde{co}$

Table 1:

Proof Consider the case where $A\ \tilde{<}\ B$ and $A\ \tilde{<}\ C$. Since B and C are distinct, there exists $\beta \in B$ such that $\beta \notin C$. Now, suppose that $(B, C) \notin \{\tilde{>}, \tilde{<}, \tilde{\#}, \tilde{co}\}$. Then, by Proposition 5.2, there exist $\gamma_1, \gamma_2, \gamma_3, \gamma_4 \in B$ such that $\Diamond\gamma_1 \notin C$, $\diamondminus\gamma_2 \notin C$, $\bigtriangledown\gamma_3 \notin C$ and $\bigtriangleup\gamma_4 \notin C$.

Let $\delta \stackrel{\text{def}}{=} \beta \wedge \gamma_1 \wedge \gamma_2 \wedge \gamma_3 \wedge \gamma_4$. Since $A\ \tilde{<}\ B$ and $\delta \in B$, Proposition 5.2 guarantees that $\Diamond\delta \in A$. Hence, by axiom A9(i) we must have $\Box(\delta \vee \Diamond\delta \vee \diamondminus\delta \vee \bigtriangledown\delta \vee \bigtriangleup\delta) \in A$, since A is an MCS. But then, since $A\ \tilde{<}\ C$, we appeal to the definition of $\tilde{<}$ to conclude that $(\delta \vee \Diamond\delta \vee \diamondminus\delta \vee \bigtriangledown\delta \vee \bigtriangleup\delta) \in C$. Hence, $\delta \in C$ or $\Diamond\delta \in C$ or $\diamondminus\delta \in C$ or $\bigtriangledown\delta \in C$ or $\bigtriangleup\delta \in C$. If $\delta \in C$, then $\beta \in C$ as well. Alternatively, if $\Diamond\delta \in C$, then since $\delta \stackrel{\text{def}}{=} \beta \wedge \gamma_1 \wedge \gamma_2 \wedge \gamma_3 \wedge \gamma_4$ and $\vdash \Diamond(\beta \wedge \gamma_1 \wedge \gamma_2 \wedge \gamma_3 \wedge \gamma_4) \supset \Diamond\beta \wedge \Diamond\gamma_1 \wedge \Diamond\gamma_2 \wedge \Diamond\gamma_3 \wedge \Diamond\gamma_4$ (TK2), we must have $\Diamond\gamma_1 \in C$. Similarly, if

$\diamondsuit\delta \in C$ then $\diamondsuit\gamma_2 \in C$, if $\nabla\delta \in C$ then $\nabla\gamma_3 \in C$ and if $\triangle\delta \in C$ then $\triangle\gamma_4 \in C$. In any case, we obtain a contradiction.

The proofs of the other cases are similar. Instead of axiom A9(i), we would have to use axioms A5, A6, A9(ii) to A9(iv) and A10 and theses T8Dual, T11Dual and T12Dual.

The requirement that the MCSs be distinct is strictly necessary only for the cases which use axiom A9, but the more restrictive statement of the lemma will be sufficient for our purposes.

$\square$

Note that there are no entries for $\tilde{\#}_\mu$ in the table above. It turns out that it is sufficient to establish the following result in this connection.

Lemma A.3 *Let A, B and C be MCSs such that $A \;\tilde{\#}_\mu\; B$ and $A \;\tilde{>}\; C$. Then, either $B \;\tilde{>}\; C$ or $B \;\tilde{co}\; C$.*

Proof Similar to the proof of Lemma A.2, using axiom A8(i).

$\square$

The next lemma demonstrates why we have been naming MCSs using formulas of the form $\hat{p}$.

Lemma A.4 *Let A be an MCS such that $\hat{p} \in A$, $p \in \mathcal{P}$. Then, for any other MCS B, (A, B) can be in at most one of the semantic relations $\{\tilde{<}, \tilde{>}, \tilde{\#}, \tilde{co}\}$.*

Proof Consider the case where $A \;\tilde{\#}\; B$. By Proposition 5.2, we know that $\nabla\hat{p} \in B$. But $\nabla\hat{p} \supset \neg p \wedge \boxminus\neg p \wedge \boxminus\neg p \wedge \triangle\neg p$ by the thesis T$\hat{p}$(iii). By again appealing to Proposition 5.2, it follows that $(A, B) \notin \{\tilde{<}, \tilde{>}, \tilde{co}\}$. A similar argument can be used for the cases $A \;\tilde{<}\; B$, $A \;\tilde{>}\; B$ and $A \;\tilde{co}\; B$.

$\square$

With these preliminaries out of the way, we can now describe in detail the inductive hypotheses $\mathcal{IC}$. Recall that at stage i we assume that we have built up a chronicle structure $CH_i = (ES_i, T_i, \mu_i)$, where $ES_i = (E_i, <_i, \#_i)$. Henceforth let co_i and $\#_{\mu_i}$ denote the concurrency and minimal conflict relations of ES_i respectively. The inductive conditions satisfied by CH_i consist of five clauses.

(C1) $E_i \subseteq \{e_0, e_1, \ldots, e_i\}$.

(C2) For $0 \leq j \leq i$, if $e_j \in E_i$ then $\hat{p}_j \in T_i(e_j)$.

(C3) $\alpha_0 \in T_i(e_0)$.

(C4) $\forall e_j, e_k \in E_i, 0 \leq j, k \leq i : T_i(e_j) \;\tilde{R}\; T_i(e_k)$ *iff* $e_j \;R_i\; e_k$, where $R \in \{<, >, \#, co\}$.

(C5) $\forall e_j, e_k \in E_i, 0 \leq j, k \leq i : (e_j, e_k) \in \mu_i$ implies $T_i(e_j) \;\tilde{\#}_\mu\; T_i(e_k)$.

(C2) ensures that each event e_j present in E_i is assigned the unique name $\hat{p}_j$ by T_i. Note that the "$\Leftarrow$" direction of (C4) along with (C5) imply that CH_i is coherent. We have chosen to state the conditions separately because it turns out to be more convenient. Also, note that (C4) does *not* cover the case $e_j \;\#_{\mu_i}\; e_k$. It is thus possible that $e_j \;\#_{\mu_i}\; e_k$ whereas $(T_i(e_j), T_i(e_k)) \notin \tilde{\#}_\mu$.

Before beginning the proof of Lemma 5.13, we explicitly state and prove a result which has been mentioned in passing while describing the inductive construction.

Recall that we had abbreviated $T_i(e_k)$ by A_k for all $e_k \in E_i$. To kill the live requirement $(e_j, \nabla_\mu\beta_0)$, we constructed an MCS B_{i+1} containing β_0 and modified the MCS A_j to an MCS B_j such that $B_j \;\tilde{\#}_\mu\; B_{i+1}$. We then modified the MCSs A_k corresponding to events $e_k \in E_i - \{e_j\}$ to MCSs B_k, ensuring that the relationship between A_j and A_k was preserved between B_j and B_k.

We can now verify the following result.

Lemma A.5 *For all $e_k \in E_i$, $0 \leq k \leq i$, the following conditions hold:*

(i) $B_k \cap \Phi_i = T_i(e_k) \cap \Phi_i$ and $\hat{p}_k \in B_k$.

(ii) For $k \neq j$, if $A_k \;\tilde{R}\; A_j$ then $B_k \;\tilde{R}\; B_j$ for $\tilde{R} \in \{\tilde{<}, \tilde{>}, \tilde{\#}, \tilde{co}\}$.

Proof

(i) By the definition of g_{i+1}, for every $\alpha \in \Phi_i$ we have $g_{i+1}(\alpha) = \alpha$ and so $\alpha \in A'_k$ iff $\alpha \in A_k$. But then our construction ensures that $A'_k \subseteq B_k$ and hence $T_i(e_k) \cap \Phi_i \subseteq B_k \cap \Phi_i$.

To establish $B_k \cap \Phi_i \subseteq T_i(e_k) \cap \Phi_i$, suppose that there exists $\alpha \in B_k \cap \Phi_i$ such that $\alpha \in B_k - A_k$. Then $\alpha \notin A_k$ and consequently $\neg\alpha \in A_k$. But this would imply that $\neg\alpha \in A'_k \subseteq B_k$ as well, which contradicts the consistency of B_k.

Since $\hat{p}_k \in \Phi_i$ and $\hat{p}_k \in T_i(e_k)$ for all $e_k \in E_i$, it follows from the previous argument that $\hat{p}_k \in B_k$ as well.

(ii) Consider the case $A_k \tilde{<} A_j$. Then $\{\beta \mid \boxminus\beta \in B_j\} = A^*_k \subseteq B_k$ by construction and so $B_k \tilde{<} B_j$. (A^*_k was defined immediately preceding Lemma 5.12). Similar remarks apply to the other cases.

□

Now we can prove Lemma 5.13.

Lemma 5.13 *CH_{i+1} is a well defined chronicle structure which satisfies the inductive conditions $\mathcal{IC}$.*

Proof To verify that CH_{i+1} is a well defined chronicle structure, we have to check that ES_{i+1} is a well branching event structure and that $\mu_{i+1} \subseteq \#_{\mu\, i+1}$. It is then clear from the definition of T_{i+1} that T_{i+1} is a chronicle on ES_{i+1}.

Some of the inductive conditions can be verified before proving that CH_{i+1} is a chronicle structure.

Condition (C1) Assuming that $E_i \subseteq \{e_0, e_1, \ldots, e_i\}$, our construction guarantees that $E_{i+1} \subseteq \{e_0, e_1, \ldots, e_{i+1}\}$, so CH_{i+1} satisfies condition (C1).

Condition (C2) Lemma A.5(i) guarantees that we have not disturbed any of the formulas $\hat{p}_k$ in going from $T_i(e_k)$ to $T_{i+1}(e_k)$ for $e_k \in E_i$. We have also ensured that $\hat{p}_{i+1} \in B_{i+1} = T_{i+1}(e_{i+1})$, so CH_{i+1} satisfies (C2).

Condition (C3) We know that $\alpha_0 \in T_i(e_0)$ since CH_i satisfied condition (C3). Since $\alpha_0 \in \Phi_0 \subseteq \Phi_i$, by Lemma A.5(i), $\alpha_0 \in T_{i+1}(e_0)$ as well, so CH_{i+1} satisfies (C3).

Condition (C4) To verify that CH_{i+1} satisfies (C4) requires more effort. We have to show the following:

$$\forall e_k, e_m \in E_{i+1} : \forall R \in \{<, >, \#, co\} : e_k \neq e_m \textit{ implies } (e_k \ R_{i+1} \ e_m \textit{ iff } B_k \ \tilde{R} \ B_m)$$

There are two cases to consider — either $\{e_k, e_m\} \subseteq E_i$ or one of the two events is e_{i+1}.

Case 1: e_k and e_m both belong to E_i.

From the definition of ES_{i+1} it is clear that $<_i = <_{i+1} \cap (E_i \times E_i)$ and $\#_i = \#_{i+1} \cap (E_i \times E_i)$. From this it also follows at once that $co_i = co_{i+1} \cap (E_i \times E_i)$. Hence, we can conclude that $e_k \ R_{i+1} \ e_m$ iff $e_k \ R_i \ e_m$, for $R \in \{<, >, \#, co\}$. Since CH_i satisfies (C4) by the induction hypothesis, we know that $e_k \ R_i \ e_m$ iff $A_k \ \tilde{R} \ A_m$ for $R \in \{<, >, \#, co\}$. Thus, we only have to verify that $A_k \ \tilde{R} \ A_m$ iff $B_k \ \tilde{R} \ B_m$.

We first show that $A_k \ \tilde{R} \ A_m$ implies $B_k \ \tilde{R} \ B_m$, for $R \in \{<, >, \#, co\}$.

In case $e_k = e_j$ or $e_m = e_j$, the result follows from Lemma A.5(ii). Hence assume that $e_k \neq e_j$ and $e_m \neq e_j$.

Suppose that $A_k \ \tilde{R} \ A_m$ for some $R \in \{<, >, \#, co\}$. We shall deal with the case where $A_k \ \tilde{co} \ A_m$. The other cases can be handled similarly.

First, since CH_i satisfies (C4), we know that $e_k \ co_i \ e_m$. Since ES_i is an event structure, the pairs (e_k, e_j) and (e_m, e_j) must each belong to one of the relations $\{<_i, >_i, \#_i, co_i\}$. Thus, we have to consider all possible "triangles" involving these three events which fix $e_k \ co_i \ e_m$. Once again, we shall consider only one representative case.

Consider the case where $e_k \ co_i \ e_j$ and $e_m \ co_i \ e_j$. From the induction hypothesis (C4) and Lemma A.5(ii) it follows that $B_k \ \tilde{co} \ B_j$ and $B_m \ \tilde{co} \ B_j$. Note that $B_k \neq B_m$ because $\hat{p}_k \in B_k$ and $\hat{p}_k \notin B_m$. Hence, we can appeal to Lemma A.2 and conclude that (B_k, B_m) belongs to at least one relation from the set $\{\tilde{<}, \tilde{>}, \tilde{\#}, \tilde{co}\}$. However, by Lemma A.4 B_k and B_m cannot be related by more that one of these four relations. Hence B_k and B_m must be related by *exactly* one of these four relations. Since $\triangle\hat{p}_k \in A_m$ and $(\triangle\hat{p}_k)' = \triangle\hat{p}_k$, we know that $\triangle\hat{p}_k \in A'_m \subseteq B_m$ as well. Then by T$\hat{p}$(iv), we have $\boxminus\neg p_k \in B_m$ and

$\boxminus\neg p_k \in B_m$ and $\triangledown\neg p_k \in B_m$ and hence, by Proposition 5.2, $(B_k, B_m) \notin \{\tilde{<}, \tilde{>}, \tilde{\#}\}$. Thus, it must be the case that $B_k \ \tilde{co}\ B_m$.

Next, we must show that $B_k \ \tilde{R}\ B_m$ implies $A_k \ \tilde{R}\ A_m$ for $R \in \{<, >, \#, co\}$. Suppose that $(B_k, B_m) \in \tilde{R}$, for some $R \in \{<, >, \#, co\}$ and $(A_k, A_m) \notin \tilde{R}$. Since ES_i is an event structure, we know that $(e_k, e_m) \in R_i^*$ for some $R^* \in \{<, >, \#, co\}$. Then, since T_i satisfies condition (C4), it must be the case that $(A_k, A_m) \in \tilde{R}^*$ and so $R^* \neq R$. But, by the preceding argument, we must have $(B_k, B_m) \in \tilde{R}^*$ as well. This is a contradiction, because $\hat{p}_k \in B_k$ and thus, by Lemma A.4, (B_k, B_m) can belong to at most one semantic relation from $\{\tilde{<}, \tilde{>}, \tilde{\#}, \tilde{co}\}$.

Case 2: $e_k = e_{i+1}$ or $e_m = e_{i+1}$.

Assume without loss of generality that $e_m = e_{i+1}$. We first show that (B_k, B_{i+1}) is contained in *exactly* one of the semantic relations $\{\tilde{<}, \tilde{>}, \tilde{\#}, \tilde{co}\}$.

By Lemma A.4, (B_k, B_{i+1}) can belong to at most one of the four semantic relations because $\hat{p}_k \in B_k$. Hence it suffices to show that (B_k, B_{i+1}) is contained in at least one of the four relations.

If $e_k = e_j$, we know that $B_j \ \tilde{\#}_\mu\ B_{i+1}$, which implies that $B_j \ \tilde{\#}\ B_{i+1}$ since $\tilde{\#}_\mu \subseteq \tilde{\#}$.

On the other hand, suppose $e_k \in E_i - \{e_j\}$. Since ES_i is an event structure, we know that $(e_k, e_j) \in R_i$ for some $R \in \{<, >, \#, co\}$. Since ES_i satisfies (C4), $A_k \ \tilde{R}\ A_j$. So, by Lemma A.5(ii) we know that $B_k \ \tilde{R}\ B_j$. Since we also know that $B_j \ \tilde{\#}\ B_{i+1}$, we can appeal to Lemma A.2 to show that (B_k, B_{i+1}) must belong to at least one semantic relation.

Returning to the main proof, we know that (B_k, B_{i+1}) belongs to exactly one semantic relation from the set $\{\tilde{<}, \tilde{>}, \tilde{\#}, \tilde{co}\}$. By the definition of $<_{i+1}$ and $\#_{i+1}$, the result follows at once for the case where $(e_k, e_{i+1}) \in \{<_{i+1}, >_{i+1}, \#_{i+1}\}$. Therefore, consider the case where $e_k \ co_{i+1}\ e_{i+1}$. This implies that $(B_k, B_{i+1}) \notin \{\tilde{<}, \tilde{>}, \tilde{\#}\}$ and so it must be the case that $B_k \ \tilde{co}\ B_{i+1}$.

Having verified that (C4) holds, we can show that ES_{i+1} is well defined.

Claim: ES_{i+1} is a well branching event structure

It suffices to show that ES_{i+1} is an event structure. Since E_{i+1} is a finite set it will then follow that ES_{i+1} is in fact well branching.

We know that ES_i is an event structure and that $<_i = <_{i+1} \cap (E_i \times E_i)$ and $\#_i = \#_{i+1} \cap (E_i \times E_i)$. By the definition of $<_{i+1}$ and $\#_{i+1}$, it then follows that $<_{i+1}$ and $\#_{i+1}$ are irreflexive and that $\#_{i+1}$ is symmetric. Thus, all we have to establish is that $<_{i+1}$ is transitive and that $\#_{i+1}$ is inherited via $<_{i+1}$.

Now suppose that $e_k <_{i+1} e_l <_{i+1} e_m$. Since T_{i+1} satisfies (C4), we know that $B_k \ \tilde{<}\ B_l \ \tilde{<}\ B_m$. Hence, by Lemma A.2 it follows that $B_k \ \tilde{<}\ B_m$. Thus, by (C4), we have $e_k <_{i+1} e_m$ and so $<_{i+1}$ is transitive.

Next, suppose that $e_k \ \#_{i+1}\ e_l <_{i+1} e_m$. Once again, by (C4) we have $B_k \ \tilde{\#}\ B_l \ \tilde{<}\ B_m$ and thus $B_k \ \tilde{\#}\ B_m$ by Lemma A.2. So, by (C4), $e_k \ \#_{i+1}\ e_m$. Hence $\#_{i+1}$ is inherited via $<_{i+1}$. This establishes the claim.

Now all that remains to be shown is that CH_{i+1} satisfies (C5) and that $\mu_{i+1} \subseteq \#_{\mu\, i+1}$.

Condition(C5) We have to verify the following:

$$\forall e_k, e_m \in E_{i+1} : (e_k, e_m) \in \mu_{i+1} \ \textit{implies}\ T_{i+1}(e_k) \ \tilde{\#}_\mu\ T_{i+1}(e_m).$$

Suppose both e_k and e_m belong to E_i. Then $(e_k, e_m) \in \mu_i$. Since CH_i was a chronicle structure, in ES_i we must have had $e_k \ \#_{\mu\, i}\ e_m$. Hence $e_k \ \#_i\ e_m$ as well. So $e_k \ \#_{i+1}\ e_m$, since $\#_i = \#_{i+1} \cap (E_i \times E_i)$. Therefore, since T_{i+1} satisfies (C4), $B_k \ \tilde{\#}\ B_m$. We also know that $A_k \ \tilde{\#}_\mu\ A_m$ in CH_i by the inductive assumption (C5). Hence, $\triangledown_\mu \hat{p}_k \in A_m$. Since $(\triangledown_\mu \hat{p}_k)' = \triangledown_\mu \hat{p}_k$, $\triangledown_\mu \hat{p}_k \in A'_m \subseteq B_m$ as well. But, by Lemma A.1, from $B_k \ \tilde{\#}\ B_m$ and $\triangledown_\mu \hat{p}_k \in B_m$ we get $B_k \ \tilde{\#}_\mu\ B_m$.

On the other hand, suppose that $e_k = e_{i+1}$ or $e_m = e_{i+1}$. Without loss of generality, assume that $e_m = e_{i+1}$. Then $e_k = e_j$ and by the construction of B_j and B_{i+1} we are assured that $B_j \ \tilde{\#}_\mu\ B_{i+1}$.

Claim: $\mu_{i+1} \subseteq \#_{\mu\, i+1}$

Since T_{i+1} satisfies (C5), we know that $(e_k, e_m) \in \mu_{i+1}$ implies that $B_k \;\tilde{\#}_\mu\; B_m$ and so $B_k \;\tilde{\#}\; B_m$. Thus, since T_{i+1} satisfies (C4), $e_k \;\#_{i+1}\; e_m$. Now suppose that $e_l \in E_{i+1}$ such that $e_l <_{i+1} e_k$. Then, by (C4), we know that $B_l \tilde{<} B_k$. Hence, by Lemma A.3, either $B_l \tilde{<} B_m$ or $B_l \;\tilde{co}\; B_m$ which implies, again by (C4), $e_l <_{i+1} e_m$ or $e_l \; co_{i+1}\; e_m$. Symmetrically, for any $e_l <_{i+1} e_m$ we can argue that $e_l <_{i+1} e_k$ or $e_l \; co_{i+1}\; e_k$. Hence $e_k \;\#_{\mu\, i+1}\; e_m$.

□

Finally, we prove Lemma 5.15. Recall that our inductive construction yielded a chronicle structure $CH = (ES, T, \mu)$, where:

- $ES = (E, <, \#)$, with $E = \bigcup_{i\geq 0} E_i$, $<= \bigcup_{i\geq 0} <_i$, and $\# = \bigcup_{i\geq 0} \#_i$.
- $T : E \to 2^\Phi$ is given by:
$$\forall e_i \in E : T(e_i) = \bigcup \{T_j(e_i) \cap \Phi_j \mid j \geq i\}.$$
- $\mu = \bigcup_{i\geq 0} \mu_i$.

Lemma 5.15 *$CH = (ES, T, \mu)$ is a perfect chronicle structure in which $\alpha_0 \in T(e_0)$.*

Proof We establish this result by proving a series of claims. As usual, let $co = (E \times E) - (< \cup > \cup \;\#\; \cup\, id)$.

Claim 1 T assigns an MCS to each element of E.

Proof of Claim First, we show that for every $e \in E$, $T(e)$ is a consistent set. Consider $e_k \in E$. Suppose that $\alpha, \beta \in T(e_k)$. Let $\alpha \in \Phi_i$ and $\beta \in \Phi_j$. Let $n = MAX(i, j, k)$. Then, by Lemma A.5(i), it follows that $\alpha \wedge \beta \in T_n(e_k)$. Since $T_n(e_k)$ is an MCS, $\alpha \wedge \beta$ is consistent and so $T(e_k)$ must be consistent.

Next, we show that $T(e_k)$ is an MCS. Suppose that $T(e_k) \cup \{\beta\}$ is consistent and $\beta \notin T(e_k)$. Let $\beta \in \Phi_i$ and $n = MAX(i, k)$. Then, since $T_n(e_k)$ is an MCS, we must have $\beta \in T_n(e_k)$ or $\neg\beta \in T_n(e_k)$. If $\beta \in T_n(e_k)$, then $\beta \in T(e_k)$ by the definition of T, which is a contradiction. On the other hand, if $\neg\beta \in T_n(e_k)$ then $\neg\beta \in T(e_k)$, again by the definition of T, which implies that $T(e_k) \cup \{\beta\}$ is not consistent, which is again a contradiction.

Claim 2

(i) $\forall R \in \{<, >, \#, co\} : R_i = R \cap (E_i \times E_i)$.

(ii) $\forall e_j, e_k \in E : \forall R \in \{<, >, \#, co\} : e_j \; R \; e_k$ *iff* $e_j \; R_n \; e_k$ where $n = MAX(j, k)$.

(iii) ES is an event structure.

(iv) $\forall e_j \in E : \hat{p}_j \in T(e_j)$.

Proof of Claim (i) follows at once from the fact that at every stage i, in passing from CH_i to CH_{i+1} we ensured that $<_i = <_{i+1} \cap (E_i \times E_i)$ and $\#_i = \#_{i+1} \cap (E_i \times E_i)$, which also implies that $co_i = co_{i+1} \cap (E_i \times E_i)$. (ii) and (iii) follow immediately from (i). (iv) follows from Lemma A.5(i) and the definition of T.

Claim 3 $\forall e_j, e_k \in E : \forall R \in \{<, >, \#, co\} : e_j \; R \; e_k$ *iff* $T(e_j) \; \tilde{R} \; T(e_k)$.

Proof of Claim Let $e_j \; R \; e_k$ — say $e_j \; co \; e_k$. Consider $\triangle\alpha \in T(e_j)$. We have to verify that $\alpha \in T(e_k)$. Let $\triangle\alpha \in \Phi_i$ and let $n = MAX(i, j, k)$. Then $\triangle\alpha \in T_n(e_j)$. We know by Claim 2 that $e_j \; co_n \; e_k$ and, since T_n satisfied the inductive condition (C4), $T_n(e_j) \;\tilde{co}\; T_n(e_k)$. Hence $\alpha \in T_n(e_k)$ and so $\alpha \in T(e_k)$. Hence $T(e_j) \;\tilde{co}\; T(e_k)$. The other cases can be proved in a similar manner.

Next, suppose that $T(e_j) \; \tilde{R}_1 \; T(e_k)$ for some $R_1 \in \{<, >, \#, co\}$ and that $(e_j, e_k) \notin R_1$. Since ES is an event structure, (e_j, e_k) belongs to some relation $R_2 \in \{<, >, \#, co\}$ such that $R_1 \neq R_2$. By the first part of the argument, we must have $(T(e_j), T(e_k)) \in \tilde{R}_2$ as well. But, by Claim 2, $\hat{p_k} \in T(e_k)$. Hence, by Lemma A.4, $(T(e_j), T(e_k))$ belongs to at most one semantic relation and so we have a contradiction.

Henceforth, we let $\#_\mu$ denote the minimal conflict relation in ES.

Claim 4

(i) $\forall e_j, e_k \in E : (e_j, e_k) \in \mu$ implies $T(e_j) \tilde{\#}_\mu T(e_k)$.

(ii) $\mu \subseteq \#_\mu$.

Proof of Claim

(i) Let $e_j, e_k \in E$, such that $(e_j, e_k) \in \mu$. Then it must be the case that for some n, $(e_j, e_k) \in \mu_n$. But then, since CH_n satisfied (C5) we know that $T_n(e_j) \tilde{\#}_\mu T_n(e_k)$. From this, we can conclude that $T_n(e_j) \tilde{\#} T_n(e_k)$ and hence $e_j \#_n e_k$ since CH_n satisfied (C4). Hence $e_j \# e_k$ as well, and so, by Claim 3, $T(e_j) \tilde{\#} T(e_k)$. Also, since we must have had $\nabla_\mu \hat{p}_k \in T_n(e_j)$, we must have $\nabla_\mu \hat{p}_k \in T(e_j)$ as well. Hence, by Lemma A.1, it must be the case that $T(e_j) \tilde{\#}_\mu T(e_k)$.

(ii) By part (i), if $(e_j, e_k) \in \mu$ then $T(e_j) \tilde{\#}_\mu T(e_k)$. This implies that $T(e_j) \tilde{\#} T(e_k)$ and so we must have $e_j \# e_k$ by Claim 3. Suppose that $e_m < e_j$ in ES. Then, by Claim 3, we know that $T(e_m) \tilde{<} T(e_j)$ and hence, by Lemma A.3, $T(e_m) \tilde{<} T(e_k)$ or $T(e_m)\ \tilde{co}\ T(e_k)$. From this, using Claim 3 again, we can conclude that $e_m < e_k$ or $e_m\ co\ e_k$. Similarly, for any $e_l < e_k$, we can verify that $e_l < e_j$ or $e_l\ co\ e_j$. Thus $e_j \#_\mu e_k$.

Claim 5 T kills all requirements in CH.

Proof of Claim Suppose that $e_k \in E$ and $\beta \in T(e_k)$ where β is of the form $\Diamond\alpha$, $\diamondminus\alpha$, $\nabla\alpha$, $\triangle\alpha$ or $\nabla_\mu\alpha$. Then, we must show that there exists $e_j \in E$ such that $\alpha \in T(e_j)$ and $e_j > e_k$, $e_j < e_k$, $e_j \# e_k$, $e_j\ co\ e_k$ or $e_j\ \mu\ e_k$ respectively.

Once again, we establish this for only one concrete case. Suppose that $\beta = \nabla_\mu\alpha$. Let $\nabla_\mu\alpha \in \Phi_i$ and $m = MAX(i, k)$. Then $\nabla_\mu\alpha \in T_m(e_k)$. Let n be the index of $(e_k, \nabla_\mu\alpha)$ in the enumeration we have fixed for $\tilde{E} \times \Phi$. Then for some l, $m \leq l \leq m + (n-1)$, $(e_k, \nabla_\mu\alpha)$ must be the live requirement chosen to be killed in CH_l and hence in CH_{l+1} we have $e_j \in E_{l+1}$ such that $(e_j, e_k) \in \mu_{l+1}$ and $\alpha \in T_{l+1}(e_j)$. Clearly $\alpha \in \Phi_m$ and since $l + 1 > m$, we must have $\alpha \in T(e_j)$ and $(e_j, e_k) \in \mu$ as required.

Claim 6 $\mu = \#_\mu$.

Proof of Claim By Claim 4(ii), we know that $\mu \subseteq \#_\mu$. Hence it suffices to show that $\#_\mu \subseteq \mu$.

Suppose $e_j \#_\mu e_k$. Then $e_j \# e_k$ and so, by Claim 3, $T(e_j) \tilde{\#} T(e_k)$. We know that $\hat{p}_j \in T(e_j)$ and thus $\nabla\hat{p}_j \in T(e_k)$. From the axiom A7 we can conclude that $\{\nabla_\mu\hat{p}_j, \diamondminus\nabla_\mu\hat{p}_j, \nabla_\mu\Diamond\hat{p}_j, \diamondminus\nabla_\mu\Diamond\hat{p}_j\} \cap T(e_k) \neq \emptyset$.

Suppose that $\diamondminus\nabla_\mu\hat{p}_j \in T(e_k)$. By Claim 5, there exists e_m such that $e_m < e_k$ and $\nabla_\mu\hat{p}_j \in T(e_m)$. Once again, by Claim 5, there exists e_l such that $(e_m, e_l) \in \mu$ and $\hat{p}_j \in T(e_l)$. If e_l and e_j are distinct events, $T(e_l)$ and $T(e_j)$ must belong to some semantic relation from the set $\{\tilde{<}, \tilde{>}, \tilde{\#}, \tilde{co}\}$. However, for any $p \in \mathcal{P}$, from the definition of $\hat{p}$ it follows that no semantic relation can exist between two MCSs which both contain $\hat{p}$. Hence we must have $e_l = e_j$ and thus $(e_m, e_j) \in \mu$. Since $\mu \subseteq \#_\mu$, $e_m \#_\mu e_j$, which contradicts $e_j \#_\mu e_k$. In a similar fashion, we can show that $\nabla_\mu\Diamond\hat{p}_j$ and $\diamondminus\nabla_\mu\Diamond\hat{p}_j$ cannot be in $T(e_k)$.

Hence we must have $\nabla_\mu\hat{p}_j \in T(e_k)$. Again applying Claim 5 and the reasoning above, we must in fact have $(e_j, e_k) \in \mu$ and we are done.

Claim 7 ES is a well branching event structure.

Proof of Claim From Claim 2(iii) we already know that ES is an event structure, so all we have to verify is that ES is well branching.

Suppose $e_j, e_k \in E$ such that $e_j \# e_k$. Then, by Claim 3, $T(e_j) \tilde{\#} T(e_k)$ and, since $\hat{p}_j \in T(e_j)$, $\nabla\hat{p}_j \in T(e_k)$. Once again, from A7 we can conclude that $\{\nabla_\mu\hat{p}_j, \diamondminus\nabla_\mu\hat{p}_j, \nabla_\mu\Diamond\hat{p}_j, \diamondminus\nabla_\mu\Diamond\hat{p}_j\} \cap T(e_k) \neq \emptyset$.

If $\nabla_\mu \hat{p}_j \in T(e_k)$, then clearly $(e_j, e_k) \in \mu$ by previous arguments and hence $e_j \#_\mu e_k$ by Claim 4(ii). On the other hand, if $\diamondminus\nabla_\mu \hat{p}_j \in T(e_k)$, then we have shown in the proof of Claim 6 that there exists $e_m < e_k$ such that $e_m \#_\mu e_j$. Using a similar argument, we can show that if $\nabla_\mu \Diamond \hat{p}_j \in T(e_k)$, then there exists $e_l < e_j$ such that $e_l \#_\mu e_k$ and, finally, if $\diamondminus\nabla_\mu \Diamond \hat{p}_j \in T(e_k)$ then there exist $e_l < e_j$ and $e_m < e_k$ such that $e_l \#_\mu e_m$.

Hence, given $e_j \# e_k$ in ES, there must exist $e_l \leq e_j$ and $e_m \leq e_k$ such that $e_l \#_\mu e_m$ and so ES is well branching.

From Claims 1 through 4 and Claim 7, we can conclude that CH is a coherent chronicle structure. Claims 5 and 6 then show that CH is in fact perfect.

The fact that $\alpha_0 \in T(e_0)$ follows immediately from the fact that $\alpha_0 \in \Phi_0$, by definition, and that $\alpha_0 \in T_0(e_0)$ by construction.

$\square$