

TEMPORAL LOGICS FOR COMMUNICATING SEQUENTIAL AGENTS: I

KAMAL LODAYA and R. RAMANUJAM

The Institute of Mathematical Sciences, Madras 600 113, India

P. S. THIAGARAJAN

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

Received March 1990

First revision June 1991

Second revision October 1991

Communicated by R. Parikh

ABSTRACT

We introduce a class of distributed systems called *Communicating Sequential Agents (CSAs)*. Sound and complete axiomatizations are provided for various subclasses using a family of indexed temporal logics. Some of the important features of these logics are:

- Both the formulas and the structures for the logics reflect the fact that a system is composed out of a number of participating sequential agents.
- Formulas of the logics are interpreted only at local states.
- An agent makes a definite assertion about another agent only if it has received — directly or indirectly — some communication from that agent supporting that assertion.

Keywords: Logics of programs, temporal logic, distributed systems, knowledge.

0. Introduction

In this paper, we introduce a class of distributed systems called *Communicating Sequential Agents (CSAs)* and propose a family of indexed temporal logics to reason about them.

During the eighties, modal logics have been extensively used to reason about distributed systems. Temporal and tense logics allow clean and often decidable theories in which specifications seem to be easily written [1,2]. Model checkers can be constructed to verify specifications [3,4]. However, temporal logics have been interpreted over sequences (*linear time* [2]) and trees (*branching time* [3,5]) rather than the more general partial orders which characterize the semantics of distributed systems.

Much of the work on temporal logics is based on the idea of global states. Usually, this presumes that a global observer exists and specifications are assertions

of behaviour as seen by this observer. However, often an individual agent in a distributed system has no access to the global state of the system due to spatial separation and autonomy of the agents. Knowledge of the state of another agent is based strictly on the messages received from that agent.

Consequently we restrict our interest to distributed systems and temporal logics with the features outlined below. The distributed systems we consider shall support:

- (a) The underlying structures of the models are partial orders.
- (b) Both the formulas and the structures for the logics reflect the fact that a system is composed of a number of participating agents.
- (c) Agents are characterized as being sequential, and all choices can be traced back to them.

The logics we propose shall have the following features:

- (d) Formulas of the logics are interpreted only at local states.
- (e) An agent makes a definite assertion about another agent only if it has received — directly or indirectly — some communication from the agent supporting that assertion.

Our study covers various such classes of distributed systems and we provide sound and complete axiomatizations for them using a family of such logics.

The distributed systems we consider can be viewed as special kinds of prime event structures arising out of the work of Nielsen, Plotkin and Winskel [6] and later developed in a variety of ways by Winskel [7]. One subclass of systems we identify (called n -ACSAs) model distributed programs composed out of a finite number of sequential programs that communicate with each other by message passing. Such systems have been repeatedly considered in the literature (e.g. [8]). The semantics of CCS and related languages [9,10] can also be the systems we study here; this follows from the work of Winskel [7]. An appendix relates our models to such event structure models [7,11].

Temporal logics for partially-ordered structures have been proposed by Pinter and Wolper [12] and by Katz and Peled [13]. In these studies the local state of the agents is not the focus of attention as is the case here. Logics of knowledge [14,15,16] certainly emphasize the notion of agents, but they also assume states of knowledge as seen by a global observer. Reif and Sistla [17] consider local state but use a spatial *modality* to refer to other processors.

A logic for n -ACSAs was presented in [11] with a more elaborate axiomatization. This paper has a simpler completeness proof for this class. We use new proof strategies to obtain completeness for the various other classes considered here.

In Sec. 1, we describe systems of n Asynchronously Communicating Sequential Agents (n -ACSAs), a formal model for the intuitive notion of n sequential processes communicating with each other by sending and receiving messages.

Section 2 describes our logical language, gives a formal semantics for it in our class of models and demonstrates how a specification may be written in it.

The following two sections, Secs. 3 and 4, are devoted to presenting an axiomatization for our models and proving it sound and complete. In Sec. 5, we show how the restriction to a bounded finite number of agents can be dispensed with.

An important subclass of models is that in which all events have finite causes. This *finitariness* also yields discrete, well-founded models. Axiomatizing this subclass requires a different proof idea and is undertaken in Sec. 6.

Section 7 expands our models to incorporate those allowing synchronous (“handshake”) communication [9,10]. It turns out that the same logical language is adequate to describe these systems.

If we drop condition (e) above, we can have a temporal logic which evaluates formulas at local states but can make assertions about other agents without necessarily receiving any information from them. In such a situation, we can dispense with indexed modalities and use basic tense logic with some type propositions. In the forthcoming Part II of this paper, we consider various such alternatives within tense logic for axiomatizing the class of CSAs.

We conclude with a discussion.

1. Frames

In this section, we introduce systems of communicating sequential agents, for which we will design and study logics in the subsequent sections.

Our model of a distributed system consists of a finite set of agents that communicate with each other. An agent is simply a set of events together with a “tree-like” ordering relation over their occurrences. The idea is that an agent represents the “unfolded” behaviour of a sequential nondeterministic process.

Let (X, \leq) be a poset and $x \in X$. Then

$$\downarrow x \triangleq \{y \in X \mid y \leq x\} .$$

Definition 1.1. An agent is a pair (E, \leq) , where

- E is a set of event occurrences and
- $\leq \subseteq (E \times E)$ is a partial order called the **causality relation** such that $\forall e \in E$. $\downarrow e$ is totally ordered by \leq . □

The restriction imposed on the causality relation in the definition of an agent can be formulated in a different fashion:

$$\forall e_1, e_2, e_3 \in E. (e_1 \leq e_3 \text{ and } e_2 \leq e_3) \Rightarrow (e_1 \leq e_2 \text{ or } e_2 \leq e_1) .$$

This is referred to as **backward linearity** of the agent.

An agent is then a poset (E, \leq) in which $\downarrow e$ is a totally ordered subset of E for every $e \in E$. Suppose that $e_1 \leq e_2$ in the agent (E, \leq) . Then this will be taken to mean that in any computation that this agent participates in, e_2 can occur only if e_1 has already occurred in that computation. When neither $e_1 \leq e_2$ nor $e_2 \leq e_1$ holds, we interpret this as a *choice* between the occurrences of e_1 and e_2 in the behaviour

of the agent. Consequently in no computation that the agent participates in can both e_1 and e_2 occur. The motivation for imposing backward linearity on agents should now be clear: we do not wish an event occurrence to causally depend upon conflicting event occurrences.

$\downarrow e$ in an agent can be thought of as the *state* of that agent when the event e has “just” occurred, the state containing information about all events that have occurred. A computation of the agent is then a chain in the agent, corresponding to the standard idea of a computation as a sequence of states.

We now proceed to consider systems of such sequential agents that *asynchronously* communicate with each other. Informally, a system has finitely many sequential processes, which communicate with each other by sending messages asynchronously. The computation of each process proceeds sequentially; any waiting is caused only when the process requires a message from another. We can think of each event as being a send, receive or internal event. Since our model refers only to event occurrences, we can think of multiple occurrences of the same message as being distinguished using some scheme like affixing sequence numbers. Studies of distributed systems typically consider such models [8,18].

Definition 1.2. A system of n **Asynchronously Communicating Sequential Agents** (abbreviated n -ACSA) is a tuple $(E_1, \dots, E_n; \leq)$ where

- (i) $E_1 \cap E_j = \emptyset$, for $i \neq j \in \{1, \dots, n\}$,
- (ii) $\leq \subseteq E \times E$ is a partial order, where $E = \bigcup_j E_j$, and
- (iii) $\forall e \in E. \forall i: 1 \leq i \leq n. \downarrow e \cap E_i$ is totally ordered by \leq . □

When $e_1 \leq e_2$, $e_1 \in E_i$, $e_2 \in E_j$ and $i \neq j$, we have the occurrence of a j -event causally dependent on the occurrence of an i -event, and we think of this as a behaviour where agent j receives information about agent i . This information could be a message from agent i (the receipt of a message can never precede its sending), or a chain of indirect messages between e_1 and e_2 .

We shall use the symbol \upharpoonright for restriction. Let \leq_i denote the restriction of \leq to the agent i , $\leq \upharpoonright (E_i \times E_i)$. Note that the definition above says more than the statement that for each i , (E_i, \leq_i) is a sequential agent. Of course, this is implied by Def. 1.2, but the converse is not the case. For example, in Fig. 1, we have a system where (E_1, \leq_1) and (E_2, \leq_2) are sequential agents, but $(E_1, E_2; \leq)$ is not a 2-ACSA, as it violates condition (iii) of the definition.

Figure 2 gives an example of a 2-ACSA, where each agent chooses symmetrically between internal action and sending a message to the other agent. In each agent, if the internal action is chosen, the next event that may occur is receiving a message from the other agent. If both agents choose to send or both agents choose to perform an internal action, the system can deadlock. Figure 3 gives an example of a 2-ACSA consisting of a producer-consumer system. The producer chooses to produce an item or to stop, while the consumer receives the produced items.




Fig. 1. Two sequential agents.




Fig. 2. A 2-ACSA




Fig. 3. A Producer-Consumer System.

We had earlier defined the state of an agent (E, \leq) to be any set $\downarrow e$, where $e \in E$. In systems of communicating agents, $\downarrow e$ can be regarded as a *local state* of agent j , where $e \in E_j$. It includes the local history of agent j as well as the “latest” local histories of all other agents from which j has had a communication upto this state. Thus $\downarrow e$ represents agent j ’s view of the system state, which is in general partial. For an n -ACSA S , we let \mathcal{L}_S denote the set of local states of the system,

$$\mathcal{L}_S \triangleq \left\{ \downarrow e \mid e \in \bigcup_j E_j \right\}.$$

If e_1 and e_2 are events of distinct agents in an n -ACSA and they are unordered by \leq , we cannot read this to mean that e_1 and e_2 are concurrent. If an agent chooses between e'_1 and e'_2 , and we have $e'_1 \leq e_1$, $e'_2 \leq e_2$, then in no computation can e_1

and e_2 both occur. Thus, local choice gets “inherited” in computations. However, it is easy to discover which events are concurrent.

Definition 1.3. Let $(E_1, \dots, E_n; \leq)$ be an n -ACSA and let $E = \bigcup_j E_j$. If $e_1, e_2 \in E$, and neither $e_1 \leq e_2$ nor $e_2 \leq e_1$ holds, then we say e_1 and e_2 are **concurrent** iff for all $i \in \{1, \dots, n\}$, $(\downarrow e_1 \cup \downarrow e_2) \cap E_i$ is totally ordered by \leq . \square

Clearly, we can see that no two events of the same agent can be concurrent. Further, unordered events of distinct agents are not concurrent only if they causally depend upon unordered events of some agent in the system. Hence a conflict can always be traced back to a choice made by some agent.

For convenience, our definition of frames is parametrized by the number of agents. In a later section, we will generalize the definition to include systems with finite but arbitrarily many agents. Further, we will also relax our assumption above that communication between agents is purely asynchronous; we will consider events shared by many agents, representing synchronization among them. For the time being, we restrict our attention to n -ACSAs and a logic to reason about them.

n -ACSAs are closely related to *event structures* [6,7]. In fact, there is a bijection between the class of n -ACSAs and that of n -agent event structures introduced in [11]. In the Appendix, we formally establish this relationship.

We let S, S', \dots (with or without subscripts) range over the class of n -ACSAs. Often we will write $S = (E; \leq)$ to mean $S = (E_1, \dots, E_n; \leq)$, where an implicit partitioning of E into E_1, \dots, E_n will be assumed and write $\downarrow e$ with respect to \leq , without specifically mentioning the partial order relation. We will also use $\uparrow e$ to denote $\{e' \in E \mid e \leq e'\}$.

2. The Language and its Models

For this section and the next one we fix an $n \in \mathbb{N}$ and let i, j, k range over $\{1, 2, \dots, n\}$.

We fix, in addition, a countable set of atomic propositions $P = \{p_1, p_2, \dots\}$ and let p, q range over P . We also fix a set consisting of n atomic *type* propositions, $T_n = \{\tau_1, \dots, \tau_n\}$ such that $P \cap T_n = \emptyset$ and set $P' = P \cup T_n$. The type propositions will be used to identify particular agents.

The formulas of our language are then given by:

Definition 2.1.

- (i) Every member of P' is a formula.
- (ii) If α and β are formulas, then so are $\sim \alpha$, $\alpha \vee \beta$, $\exists_i \alpha$ and $\Box_i \alpha$, for $1 \leq i \leq n$. \square

Let Φ_n be the set of all formulas. We let $\alpha, \beta, \gamma, \delta$ (with or without subscripts) range over Φ_n . The following derived modalities and logical connectives will

prove useful:

- (i) $\alpha \wedge \beta \triangleq \sim (\sim \alpha \vee \sim \beta)$
- (ii) $\alpha \oplus \beta \triangleq (\sim \alpha \wedge \beta) \vee (\sim \beta \wedge \alpha)$
- (iii) $\alpha \supset \beta \triangleq \sim \alpha \vee \beta$
- (iv) $\alpha \equiv \beta \triangleq (\alpha \supset \beta) \wedge (\beta \supset \alpha)$
- (v) $\diamond_i \alpha \triangleq \sim \Box_i \sim \alpha$
- (vi) $\Box_i \alpha \triangleq \sim \square_i \sim \alpha$

A **frame** is an n -ACSA $S = (E_1, E_2, \dots, E_n; \leq)$. A model is an ordered pair $M = (S, V)$, where

- (i) $S = (E_1, E_2, \dots, E_n; \leq)$ is a frame and
- (ii) $V : \mathcal{L}_S \rightarrow 2^{P'}$ is the **valuation function** defined on the local states of S satisfying:

$$\forall e \in E. \tau_i \in V(\downarrow e) \text{ iff } e \in E_i .$$

Let $M = (S, V)$ be a model, where $S = (E; \leq)$, and let $e \in E$. For a formula α , the notation of α being **true** at a local state $\downarrow e$ in M is denoted as $M, \downarrow e \vDash \alpha$ and is defined inductively as follows:

Definition 2.2.

- (i) $M, \downarrow e \vDash \alpha$ iff $\alpha \in V(\downarrow e)$, for $\alpha \in P'$.
- (ii) $M, \downarrow e \vDash \sim \alpha$ iff $M, \downarrow e \not\vDash \alpha$.
- (iii) $M, \downarrow e \vDash \alpha \vee \beta$ iff $M, \downarrow e \vDash \alpha$ or $M, \downarrow e \vDash \beta$.
- (iv) $M, \downarrow e \vDash \Box_i \alpha$ iff for all $e' \in \downarrow e \cap E_i$, $M, \downarrow e' \vDash \alpha$.
- (v) $M, \downarrow e \vDash \square_i \alpha$ iff
 - (Case 1: $e \in E_i$ — for all $e' \in \uparrow e \cap E_i$, $M, \downarrow e' \vDash \alpha$).
 - (Case 2: $e \notin E_i$ — for some $e' \in \downarrow e \cap E_i$, $M, \downarrow e' \vDash \square_i \alpha$). □

The first three clauses of the definition are standard and require no explanation. The meaning of $\Box_i \alpha$ asserted at a local state $\downarrow e$, where $e \in E_i$, can be expressed as:

“at the state when e has just occurred, as far as agent j knows, α has always been true in agent i .”

This works like the standard tense logic past operator.

On the other hand, the semantics of $\square_i \alpha$ is nonstandard — when it is asserted at a local state $\downarrow e$, where $e \in E_j$, it must be read as:

“at the state when e has just occurred, agent j knows that α will be true henceforth in agent i .”

When $i = j$, $\square_i \alpha$ works in the same way as the tense logic future operator. Note that the two cases in clause (v) could be combined into a single one:

$$M, \downarrow e \vDash \square_i \alpha \text{ iff for some } e' \in \downarrow e \cap E_i, \\ \text{for all } e'' \in \uparrow e' \cap E_i, M, \downarrow e'' \vDash \alpha.$$

Note that $\Box_i\alpha$ is strong and $\Box_i\alpha$ is weak in the following sense: when $\Box_i\alpha$ is asserted in an agent other than i , it refers to knowledge about another agent. Hence the semantics ensures that the agent i has communicated at some stage in the past that α would hold henceforth in its future. Thus if there has been no communication from agent i at all at some local state for agent j , then agent j can assert $\Box_i\beta$ for arbitrary β at that state, but not $\Box_i\beta$.

The semantics of the derived dual modalities can be understood in a similar manner. It can be easily seen that

$M, \downarrow e \models \Diamond_i\alpha$ iff

(Case 1: $e \in E_i$;) for some $e' \in \uparrow e \cap E_i$, $M, \downarrow e' \models \alpha$.

(Case 2: $e \notin E_i$;) for all $e' \in \downarrow e \cap E_i$, $M, \downarrow e' \models \Diamond_i\alpha$.

Thus while in agent i itself $\Diamond_i\alpha$ asserts that α can hold in future at some time, in another agent $j \neq i$, it merely says:

“as far as agent j knows at this local state, α can hold eventually in agent i ”.

Therefore, if agent i has never communicated with agent j , then $\Diamond_i\alpha$ will hold in agent j at that state for any α .

The semantics of $\Diamond_i\alpha$ is standard:

$M, \downarrow e \models \Diamond_i\alpha$ iff for some $e' \in \downarrow e \cap E_i$, and $M, \downarrow e' \models \alpha$.

The intended meaning is (assume $e \in E_j$):

“agent j knows at $\downarrow e$ that α was true in the past in agent i .”

Definition 2.3.

- (i) α is **satisfiable** (α has a model) iff there exists a model $M = ((E; \leq), V)$ such that $M, \downarrow e \models \alpha$ for some $e \in E$.
- (ii) For a model $M = ((E; \leq), V)$, $M \models \alpha$ iff for all $e \in E$, $M, \downarrow e \models \alpha$.
- (iii) α is **valid** (denoted $\models \alpha$) iff $M \models \alpha$ for every model M .

Since $\downarrow e$ is uniquely determined by e in a frame, we will also write $M, e \models \alpha$ for $M, \downarrow e \models \alpha$. When we talk of an event e satisfying a formula α , we mean the local state $\downarrow e$ satisfies α .

We now give an example to illustrate the use of our logical language. Consider a distributed database managed by n processes which communicate with each other by exchanging messages. A protocol is needed whereby the processes can *commit* to a distributed transaction. When each committed process knows that all the others have also committed it can go ahead and perform the distributed transaction. For this, the following requirement must be met.

If any process commits to the transaction then it knows that all processes in the system can eventually commit.

Such *distributed transaction commit protocols* commonly arise in the design of distributed systems [12].

We now specify the protocol requirement in our logical language. Let $\{C_1, \dots, C_n\}$ be a set of propositions, where C_j is read to mean “process j has committed

to the transaction". The formula

$$\bigwedge_i \left(\tau_i \wedge C_i \supset \diamond_i \left(\bigwedge_j \diamond_j C_j \right) \right)$$

expresses the requirement above.

A two-stage implementation of this protocol may use two local variables in each process P_i :

- (i) a variable L_i in which process P_i records whether it can participate in the transaction or not, and
- (ii) a variable, which we also call C_i , to record the commitment of the process to the transaction.

The implementation can perhaps run as follows:

Process P_i :

- (i) As soon as a local decision L_i is made, broadcast L_i to all other processes.
- (ii) When L_j is heard from all j , set C_i to TRUE.
- (iii) As soon as C_i is set, broadcast it to all other processes.
- (iv) When C_j is heard from all j , perform transaction.
- (v) Acknowledge all incoming messages.

All processes follow the same protocol in a symmetric manner. This is, of course, a naive protocol. However, our aim here is to merely illustrate the use of our logical language. Let us again, by abuse of notation, use $\{L_1, \dots, L_n\}$ to denote another set of propositions. Consider now the following formulas:

$$(1) \bigwedge_i \left(\tau_i \supset \left(C_i \equiv \bigwedge_j \diamond_j L_j \right) \right)$$

$$(2) \bigwedge_i \left(\tau_i \wedge C_i \supset \diamond_i \bigwedge_j \diamond_j \diamond_i C_i \right).$$

The first of the two formulas says that a process sets C_i to be TRUE only when it has heard L_j from all other processes P_j . The second formula asserts that if any process P_i sets C_i , then in its future there is a state when it has heard an acknowledgement from all other processes for a broadcast from P_i that C_i has been set. Note that here an agent has to assert something about the state of other agents and this can be done using messages from them.

It is easy to verify that the formulas (1) and (2) together imply the requirement above. In fact, in the next section, we use an axiom system and show how we can logically deduce the requirement from (1) and (2). This verifies that the simple protocol above meets its specification.

Note that the protocol above works for only one transaction. When a protocol is needed for several transactions, we can index the transaction by sequence numbers and modify the specification above appropriately. In a deterministic system, the formula expresses a strong requirement. In a nondeterministic system, we can only assert the existence of a future "committing" execution.

3. The Axiom System $\mathcal{A}(n\text{-ACSA})$

Most of our axioms are indexed versions of Burgess's axioms for tense logic [19]. The new axioms that we introduce reflect the way knowledge about other agents depends on communication from them.

Axioms

- (A0) All substitutional instances of the tautologies of propositional logic.
- (A1) (a) $\Box_i(\alpha \supset \beta) \supset (\Box_i\alpha \supset \Box_i\beta)$ (Deductive closure)
 (b) $\Box_i(\alpha \supset \beta) \supset (\Box_i\alpha \supset \Box_i\beta)$
- (A2) (a) $\tau_i \supset (\Box_i\alpha \supset \alpha)$ (Local reflexivity)
 (b) $\tau_i \supset (\Box_i\alpha \supset \alpha)$
- (A3) $\Diamond_i\Diamond_j\alpha \supset \Diamond_j\alpha$ (Transitivity)
- (A4) $\Diamond_i\alpha \wedge \Diamond_i\beta \supset \Diamond_i(\alpha \wedge \Diamond_i\beta) \vee \Diamond_i(\beta \wedge \Diamond_i\alpha)$ (Backward linearity)
- (A5) (a) $\Diamond_i\alpha \supset \Box_i\Diamond_i\alpha$ (Relating past and future)
 (b) $\Diamond_i\alpha \supset \Box_i\Diamond_i\alpha$
- (A6) $\Box_i\alpha \supset \Diamond_i\Box_i\alpha$ (Communication)
- (A7) (a) $\tau_i \equiv \bigwedge_{j \neq i} \sim \tau_j$ (Type axioms)
 (b) $\Box_i\tau_i$
 (c) $\tau_i \supset \Box_i\tau_i$

Inference Rules

- (MP)
$$\frac{\alpha, \alpha \supset \beta}{\beta}$$
- (TG \Box_i)
$$\frac{\alpha}{\Box_i\alpha}$$
- (TG \Diamond_i)
$$\frac{\alpha}{\tau_i \supset \Box_i\alpha}$$

Firstly we note that reflexivity (Axiom A2) holds only within agents, and hence $\alpha \wedge \Box_i \sim \alpha$ may well be consistent. The formula $\Box_i \sim \alpha$, when it is asserted by an agent $j \neq i$, talks only about events in agent i as viewed by j , and hence a j -event may satisfy α as well as $\Box_i \sim \alpha$. Similarly, $\alpha \wedge \Box_i \sim \alpha$ can also be consistent.

(A3) asserts transitivity across agents. As we shall see later, in the case of $\Diamond_i\alpha$, we only have transitivity within agents. (A4) states that individual agents are tree-like.

Note that the standard form of (A5.a), namely, $\alpha \supset \Box_i\Diamond_i\alpha$ is not sound in our logical system. An event e in E_j , where $j \neq i$, may satisfy α , but may have no communication from i at all to support $\Box_i\Diamond_i\alpha$. Hence the axiom refers only to i -events or other events where i has communicated. A similar remark holds in the case of (A5.b).

The communication axiom (A6) says that a strong assertion $\Box_i\alpha$ must be supported by communication from agent i to that effect. The type axiom (A7.a) captures the fact that each event belongs to exactly one agent. The other type axioms reflect the nature of our indexed modalities $\Box_i\alpha$ and $\Diamond_i\alpha$ being assertions about

agent i . Note that $\Box_i \tau_i$ is not in general valid, since an agent j , in the absence of any communication from i ($\neq j$), can well assert that $\Diamond_i \sim \tau_i$ holds!

The inference rules (MP) and (TG \Box_i) are standard, whereas the rule (TG \Box_i) again reflects the nonstandard nature of our future modality.

A formula α which can be derived using the axioms and the inference rules will be called a **thesis**. We will use $\vdash \alpha$ to denote the fact that α is a thesis in the system $\mathcal{A}(n\text{-ACSA})$.

We say a formula is **consistent** if its negation is not a thesis in our system. The finite set of formulas $\{\alpha_1, \dots, \alpha_m\}$ is consistent if and only if the formula $\alpha_1 \wedge \dots \wedge \alpha_m$ is consistent. A set of formulas is consistent if and only if every finite subset is consistent.

Theorem 3.1. (Soundness)

If $\vdash \alpha$ then $\models \alpha$.

Proof. The soundness of Axioms (A1) and (A5) and the inference rules (MP) and (TG \Box_i) are standard arguments in tense logic. Axioms (A2), (A3), (A4) and (A7) can be easily seen to be sound from the definition of frames.

Now consider (A6). Let $M = ((E; \leq), V)$ be a model and assume for some $e \in E$ that $M, e \models \Box_i \alpha$. If $e \in E_i$, since $e \leq e$, by semantics of \Diamond_i , $M, e \models \Diamond_i \Box_i \alpha$, as required. Otherwise, by semantics of $\Box_i \alpha$, there exists $e' \in E_i$ such that $e' \leq e$ and $M, e' \models \Box_i \alpha$. Clearly, by semantics of \Diamond_i , $M, e \models \Diamond_i \Box_i \alpha$.

To see that (TG \Box_i) preserves validity, assume that α is valid, and that $M, e \models \tau_i$. Then $e \in E_i$. If $e' \in E_i$ such that $e \leq e'$, then since $\models \alpha$, $M, e' \models \alpha$ as well. Thus, by the semantics of $\Box_i \alpha$, $M, e \models \Box_i \alpha$. \square

Proposition 3.2.

- (i) If $\Diamond_i \alpha$ is consistent, then so is α .
- (ii) If $\tau_i \wedge \Diamond_i \alpha$ is consistent, then so is α .

Proof.

(i) and (ii) are, respectively, the contrapositive versions of the inference rules (TG \Box_i) and (TG \Box_i). \square

We now state some useful theses and derived inference rules.

Theses

- | | |
|------|--|
| (T1) | (a) $\Box_i(\alpha \wedge \beta) \equiv (\Box_i \alpha \wedge \Box_i \beta)$
(b) $\Box_i(\alpha \wedge \beta) \equiv (\Box_i \alpha \wedge \Box_i \beta)$ |
| (T2) | (a) $\Diamond_i(\alpha \wedge \beta) \supset (\Diamond_i \alpha \wedge \Diamond_i \beta)$
(b) $\Diamond_i(\alpha \wedge \beta) \supset (\Diamond_i \alpha \wedge \Diamond_i \beta)$ |
| (T3) | (a) $\Box_i \alpha \wedge \Diamond_i \beta \supset \Diamond_i(\alpha \wedge \beta)$
(b) $\Box_i \alpha \wedge \Diamond_i \beta \supset \Diamond_i(\alpha \wedge \beta)$ |

- (T4) (a) $\diamond_i \alpha \supset \diamond_i(\tau_i \wedge \alpha)$
 (b) $\tau_i \wedge \diamond_i \alpha \supset \diamond_i(\tau_i \wedge \alpha)$
- (T5) $\Box_i \alpha \supset \Box_i \Box_i \alpha$
- (T6) $\diamond_i(\alpha \wedge \Box_i \beta_1) \wedge \dots \wedge \diamond_i(\alpha \wedge \Box_i \beta_k)$
 $\supset \diamond_i(\alpha \wedge \Box_i \beta_1 \wedge \dots \wedge \Box_i \beta_k) \quad (k > 0)$

Derived Inference Rules

$$(DR0) \quad \frac{\alpha \equiv \beta}{\gamma[\alpha/\beta] \equiv \gamma}$$

where $\gamma[\alpha/\beta]$ is the formula obtained by replacing α by β uniformly throughout γ .

$$(DR1) \quad \frac{\alpha \supset \beta}{\Box_i \alpha \supset \Box_i \beta}$$

$$(DR2) \quad \frac{\alpha \supset \beta}{\Box_i \alpha \supset \Box_i \beta}$$

$$(DR3) \quad \frac{\alpha \supset \beta}{\diamond_i \alpha \supset \diamond_i \beta}$$

$$(DR4) \quad \frac{\alpha \supset \beta}{\tau_i \supset (\diamond_i \alpha \supset \diamond_i \beta)}$$

The derivations of (T1) through (T4) and of the derived inference rules are easily obtained from [19,20]. Here we just derive (T5) and (T6).

- (T5) $\Box_i \alpha \supset \Box_i \Box_i \alpha$
- (1) $\Box_i \alpha \supset \diamond_i \Box_i \alpha$ (A6)
 - (2) $\diamond_i \Box_i \alpha \supset \Box_i \diamond_i \Box_i \alpha$ (A5.a, subst)
 - (3) $\diamond_i \Box_i \alpha \supset \Box_i \alpha$ (A5.b, contrapos.)
 - (4) $\Box_i \diamond_i \Box_i \alpha \supset \Box_i \Box_i \alpha$ (3, DR2)
 - (5) $\Box_i \alpha \supset \Box_i \Box_i \alpha$ (1, 2, 4)

$$(T6) \quad \diamond_i(\alpha \wedge \Box_i \beta_1) \wedge \dots \wedge \diamond_i(\alpha \wedge \Box_i \beta_k) \supset$$

$$\diamond_i(\alpha \wedge \Box_i \beta_1) \wedge \dots \wedge \Box_i \beta_k \quad (k > 0)$$

The derivation is by induction on k . The base case, when $k = 1$, is trivial to see. Below, let $\hat{\beta}$ abbreviate the formula $\beta_1 \wedge \dots \wedge \beta_{k-1}$.

- (1) $\diamond_i(\alpha \wedge \Box_i \beta_i) \wedge \dots \wedge \diamond_i(\alpha \wedge \Box_i \beta_{k-1}) \supset \diamond_i(\alpha \wedge \Box_i \hat{\beta})$
 (IH, T1.a, DR3)
- (2) $\diamond_i(\alpha \wedge \Box_i \beta_1) \wedge \dots \wedge \diamond_i(\alpha \wedge \Box_i \beta_k)$
 $\supset \diamond_i(\alpha \wedge \Box_i \hat{\beta}) \wedge \diamond_i(\alpha \wedge \Box_i \beta_k)$ (1, PC)

$$\begin{aligned}
 (3) \quad & \diamond_i(\alpha \wedge \Box_i \hat{\beta}) \wedge \diamond_i(\alpha \wedge \Box_i \beta_k) \\
 & \supset \diamond_i(\alpha \wedge \Box_i \hat{\beta} \wedge \diamond_i(\alpha \wedge \Box_i \beta_k)) \\
 & \quad \vee \diamond_i(\alpha \wedge \Box_i \beta_k \wedge \diamond_i(\alpha \wedge \Box_i \hat{\beta})) \quad (A4) \\
 (4) \quad & \diamond_i(\alpha \wedge \Box_i \hat{\beta} \wedge \diamond_i(\alpha \wedge \Box_i \beta_k)) \\
 & \supset \diamond_i(\alpha \wedge \diamond_i(\alpha \wedge \Box_i \hat{\beta} \wedge \Box_i \beta_k)) \quad (A3, T3.a, DR3) \\
 (5) \quad & \diamond_i(\alpha \wedge \Box_i \hat{\beta} \wedge \diamond_i(\alpha \wedge \Box_i \beta_k)) \\
 & \supset \diamond_i \alpha \wedge \diamond_i \diamond_i(\alpha \wedge \Box_i \hat{\beta} \wedge \Box_i \beta_k) \quad (4, T2.a) \\
 (6) \quad & \diamond_i(\alpha \wedge \Box_i \hat{\beta} \wedge \diamond_i(\alpha \wedge \Box_i \beta_k)) \\
 & \supset \diamond_i(\alpha \wedge \Box_i \hat{\beta} \wedge \Box_i \beta_k) \quad (5, A3, PC) \\
 (7) \quad & \diamond_i(\alpha \wedge \Box_i \beta_k \wedge \diamond_i(\alpha \wedge \Box_i \hat{\beta})) \\
 & \supset \diamond_i(\alpha \wedge \Box_i \hat{\beta} \wedge \Box_i \beta_k) \quad (6, \text{subst}) \\
 (8) \quad & \diamond_i(\alpha \wedge \Box_i \beta_1) \wedge \dots \wedge \diamond_i(\alpha \wedge \Box_i \beta_k) \\
 & \supset \diamond_i(\alpha \wedge \Box_i \beta_1 \wedge \dots \wedge \Box_i \beta_k) \quad (2, 3, 6, 7, T1.a, DR3)
 \end{aligned}$$

We now present a proof of the protocol given in Sec. 2. We need to show that the formula

$$(R) \quad \bigwedge_i (\tau_i \wedge C_i \supset \diamond_i \bigwedge_j \diamond_j C_j)$$

is implied by the formulas (1) and (2) below:

$$\begin{aligned}
 (1) \quad & \bigwedge_i (\tau_i \supset (C_i \equiv \bigwedge_j \diamond_j L_j)) \\
 (2) \quad & \bigwedge_i (\tau_i \wedge C_i \supset \diamond_i \bigwedge_j \diamond_j \diamond_i C_i)
 \end{aligned}$$

The proof goes as follows. Each line of the derivation uses DR3 and DR4 in addition to the theses cited.

$$\begin{aligned}
 (3) \quad & \bigwedge_i (\tau_i \wedge C_i \supset \diamond_i \bigwedge_j \diamond_j \diamond_i(\tau_i \wedge C_i)) \quad (2, T4.a) \\
 (4) \quad & \bigwedge_i (\tau_i \wedge C_i \supset \diamond_i \bigwedge_j \diamond_j \diamond_i \bigwedge_k \diamond_k L_k) \quad (3, 1) \\
 (5) \quad & \bigwedge_i (\tau_i \wedge C_i \supset \diamond_i \bigwedge_j \diamond_j \bigwedge_k \diamond_k L_k) \quad (4, A3) \\
 (6) \quad & \bigwedge_i (\tau_i \wedge C_i \supset \diamond_i \bigwedge_j \diamond_j (\tau_j \wedge \bigwedge_k \diamond_k L_k)) \quad (5, T4.a) \\
 (7) \quad & \bigwedge_i (\tau_i \wedge C_i \supset \diamond_i \bigwedge_j \diamond_j C_j) \quad (6, 1)
 \end{aligned}$$

4. Completeness of $\mathcal{A}(n\text{-ACSA})$

We now proceed to demonstrate that our axiom system is complete. The proof follows Burgess [19] in style. Completeness is proved by showing that every consistent formula is satisfiable.

By an MCS (Maximal Consistent Set), we mean a consistent set of formulas which is not properly included in any other consistent set. The next two results are standard.

Proposition 4.1. Any consistent set of formulas can be extended to an MCS. \square

Proposition 4.2. Let A be an MCS.

- (i) $\sim \alpha \in A$ iff $\alpha \notin A$.
- (ii) $\alpha \vee \beta \in A$ iff $\alpha \in A$ or $\beta \in A$.
- (iii) If $\vdash \alpha$ then $\alpha \in A$.
- (iv) If $\alpha \in A$ and $\vdash \alpha \supset \beta$, then $\beta \in A$. \square

We will be using these two propositions throughout without explicitly referring to them.

Proposition 4.3. Let A be an MCS. For some $i \in \{1, \dots, n\}$, $\tau_i \in A$ and for all $i \neq j$, $\tau_j \notin A$.

Proof. Follows from Axiom (A7.a). \square

Definition 4.4. Let A and B be MCSs, where $\tau_i \in A$. Then

$$A \preceq B \triangleq \{\diamond_i \alpha \mid \alpha \in A\} \subseteq B. \quad \square$$

Proposition 4.5. Let A and B be MCSs, where $\tau_i \in A$. Then

$$A \preceq B \text{ iff } \{\alpha \mid \exists i; \alpha \in B\} \subseteq A.$$

Proof.

(\Rightarrow):

Suppose $\exists i; \alpha \in B$ and $\alpha \notin A$. Then $\sim \alpha \in A$ and since $\tau_i \in A$ and $A \preceq B$, $\diamond_i \sim \alpha \in B$, which is a contradiction.

(\Leftarrow):

Similar to the previous case. \square

Proposition 4.6. \preceq is reflexive and transitive.

Proof. Let A be an MCS and let $\tau_i \in A$. If $\alpha \in A$, then by axiom (A2.a), $\diamond_i \alpha \in A$. Thus \preceq is clearly reflexive.

To show transitivity, assume MCSs A , B and C , where $\tau_i \in A$, $\tau_j \in B$, $A \preceq B$ and $B \preceq C$. If $\alpha \in A$, then $\diamond_i \alpha \in B$ and hence $\diamond_j \diamond_i \alpha \in C$. By Axiom (A3), $\diamond_i \alpha \in C$. Thus $A \preceq C$, as required. \square

\preceq is not only a preorder, but is also “backward-connected” within agents. The following proposition will be useful later.

Proposition 4.7. Let A , B and C be MCSs, where $\tau_i \in A \cap B$, $A \preceq C$ and $B \preceq C$. Then either $A \preceq B$ or $B \preceq A$.

Proof. Suppose that we have neither $A \preceq B$ nor $B \preceq A$. Then there exist formulas α and β such that $\alpha \wedge \sim \diamond_i \beta \in A$ and $\beta \wedge \sim \diamond_i \alpha \in B$. Since $A \preceq C$ and $B \preceq C$,

we have $\diamond_i(\alpha \wedge \Box_i \sim \beta) \wedge \diamond_i(\beta \wedge \Box_i \sim \alpha) \in C$. Hence by Axiom (A4), $\diamond_i(\alpha \wedge \Box_i \sim \beta \wedge \diamond_i(\beta \wedge \Box_i \sim \alpha))$ or $\diamond_i(\beta \wedge \Box_i \sim \alpha \wedge \diamond_i(\alpha \wedge \Box_i \sim \beta))$ is in C . Without loss of generality, assume that the former is in C . By Proposition 3.2(i), $\alpha \wedge \Box_i \sim \beta \wedge \diamond_i(\beta \wedge \Box_i \sim \alpha)$ is consistent. By (T2.a, T3.a), we get $\diamond_i(\beta \wedge \sim \beta)$ is consistent. Again using Proposition 3.2(i), $\beta \wedge \sim \beta$ is consistent, which is absurd. \square

Lemma 4.8. Let A be an MCS and let $\diamond_i \alpha \in A$. Then there exists an MCS B such that $B \preceq A$ and $\{\tau_i, \alpha\} \subseteq B$.

Proof. Consider the set $\Sigma \triangleq \{\beta \mid \Box_i \beta \in A\} \cup \{\tau_i, \alpha\}$. It suffices to show that Σ is consistent, because in that case we can extend Σ to an MCS B , and by Proposition 4.5, $B \preceq A$.

Let $\Sigma' \triangleq \{\tau_i, \alpha, \beta_1, \dots, \beta_k\}$ be an arbitrary finite subset of Σ . Since $\Box_i \beta_1, \dots, \Box_i \beta_k \in A$, by Thesis (T1.a), $\Box_i(\beta_1 \wedge \dots \wedge \beta_k) \in A$. Since $\diamond_i \alpha \in A$, by (T4.a), $\diamond_i(\tau_i \wedge \alpha) \in A$. Hence, by (T3.a), $\diamond_i(\tau_i \wedge \alpha \wedge \beta_1 \wedge \dots \wedge \beta_k) \in A$ and must be consistent. But then, by Proposition 3.2(i), Σ' and consequently Σ must be consistent. \square

Lemma 4.9. Let A be an MCS and let $\tau_i \wedge \diamond_i \alpha \in A$. Then there exists an MCS B such that $A \preceq B$ and $\{\tau_i, \alpha\} \subseteq B$.

Proof. Consider the set $\Sigma \triangleq \{\diamond_i \beta \mid \beta \in A\} \cup \{\tau_i, \alpha\}$. It suffices to show that Σ is consistent because for any MCS B that contains Σ we will have, by the definition of \preceq , that $A \preceq B$.

Let $\Sigma' \triangleq \{\tau_i, \alpha, \diamond_i \beta_1, \dots, \diamond_i \beta_k\}$ be an arbitrary finite subset of Σ . Since τ_i as well as $\beta_1, \dots, \beta_k \in A$, applying (A2.a) and (A5.a), $\Box_i \diamond_i(\beta_1 \wedge \dots \wedge \beta_k) \in A$. By (T2.a) and (DR2), $\Box_i(\diamond_i \beta_1 \wedge \dots \wedge \diamond_i \beta_k) \in A$. Since τ_i and $\diamond_i \alpha \in A$, by (T4.b), $\diamond_i(\tau_i \wedge \alpha) \in A$. Hence, using (T.3b), the formula $\tau_i \wedge \diamond_i(\tau_i \wedge \alpha \wedge \diamond_i \beta_1 \wedge \dots \wedge \diamond_i \beta_k) \in A$ and must be consistent. But then, by Proposition 3.2(ii), Σ' and consequently Σ must be consistent. \square

Definition 4.10. Let $S = (E; \leq)$ be a frame. Then

(i) A **chronicle** on S is a function T which assigns an MCS to each $e \in E$ such that for $e \in E$, $\tau_i \in T(e)$ iff $e \in E_i$.

Let T be a chronicle on the frame S . Then

(ii) T is **coherent** iff it satisfies for $e, e' \in E$, $e \leq e' \Rightarrow T(e) \preceq T(e')$.

(iii) T is **prophetic** iff whenever $e \in E_i$ and $\diamond_i \alpha \in T(e)$, there exists $e' \in E_i$ such that $e \leq e'$ and $\alpha \in T(e')$.

(iv) T is **historic** iff whenever $e \in E$ and $\diamond_i \alpha \in T(e)$, there exists $e' \in E_i$ such that $e' \leq e$ and $\alpha \in T(e')$.

(v) T is **perfect** iff it is coherent, historic and prophetic. \square

Definition 4.11. Let $S = (E; \leq)$ be a frame and T a chronicle on it. The **valuation induced by T** , denoted V_T , is given by:

$$\forall e \in E, V_T(\downarrow e) \triangleq T(e) \cap P' .$$

We use M_T to denote (S, V_T) . □

Lemma 4.12. Let T be a perfect chronicle on a frame $S = (E; \leq)$. Then for any $e \in E$ and formula α , $\alpha \in T(e)$ iff $M_T, e \vDash \alpha$.

Proof. The proof proceeds by induction on the structure of α .

The cases when $\alpha \in P'$ or α is of the form $\sim \beta$ or $\beta_1 \vee \beta_2$ are routine. Now assume that α is of the form $\Box_i \beta$.

(\Rightarrow):

Suppose $\Box_i \beta \in T(e)$, for some $e \in E$. To show $M_T, e \vDash \Box_i \beta$, consider $e' \in E_i$ such that $e' \leq e$. We have $\tau_i \in T(e')$. By coherence of T , $T(e') \preceq T(e)$, and by Proposition 4.5, $\beta \in T(e')$. Then by induction hypothesis, $M_T, e' \vDash \beta$, as required.

(\Leftarrow):

Suppose that $\Box_i \beta \notin T(e)$. Then $\Diamond_i \sim \beta \in T(e)$, and since T is historic, there exists $e' \in E_i$ such that $e' \leq e$ and $\sim \beta \in T(e')$. Hence $\beta \notin T(e')$. By the induction hypothesis, $M_T, e' \not\vDash \beta$. Thus, $M_T, e \not\vDash \Box_i \beta$.

Next assume that α is of the form $\Box_i \beta$. We first consider the case when $e \in E_i$. Then $\tau_i \in T(e)$.

(\Rightarrow):

Suppose $\Box_i \beta \in T(e)$. To show $M_T, e \vDash \Box_i \beta$, consider $e' \in E_i$ such that $e \leq e'$. Therefore $\tau_i \in T(e')$. By coherence of T , $T(e) \preceq T(e')$. Hence, $\Diamond_i \Box_i \beta \in T(e')$. By the dual of Axiom (A5.b), $\Box_i \beta \in T(e')$. Since $\tau_i \in T(e')$, Axiom (A2.b) gives $\beta \in T(e')$. By the induction hypothesis, $M, e' \vDash \beta$.

(\Leftarrow):

Suppose that $\Box_i \beta \notin T(e)$. Then $\Diamond_i \sim \beta \in T(e)$, and as T is prophetic, there exists $e' \in E_i$ such that $e \leq e'$ and $\sim \beta \in T(e')$. Hence $\beta \notin T(e')$. By induction hypothesis, $M_T, e' \not\vDash \beta$. Thus, $M_T, e \not\vDash \Box_i \beta$.

Now consider the case when $e \in E_j$, ($j \neq i$).

(\Rightarrow):

Suppose $\Box_i \beta \in T(e)$. By Axiom (A6), $\Diamond_i \Box_i \beta \in T(e)$. Since T is historic, there exists $e' \in E_i$ such that $e' \leq e$ and $\Box_i \beta \in T(e')$. From the previous case, we know that $M_T, e' \vDash \Box_i \beta$ and hence, by the semantics of $\Box_i \beta$, we have $M_T, e \vDash \Box_i \beta$.

(\Leftarrow):

Suppose that $\Box_i \beta \notin T(e)$. Then $\Diamond_i \sim \beta \in T(e)$. By Axiom (A5.b), $\Box_i \Diamond_i \sim \beta \in T(e)$. Now consider any $e' \in E_i$ such that $e' \leq e$. $\tau_i \in T(e')$ and by coherence of T , $T(e') \preceq T(e)$, and by Proposition 4.5, $\Diamond_i \sim \beta \in T(e')$. From the previous case, we have $M_T, e' \vDash \Diamond_i \sim \beta$. Hence $M_T, e \not\vDash \Box_i \beta$. □

Hence in order to show that a formula α is satisfiable, it suffices to construct a frame $S = (E; \leq)$ and a perfect chronicle T over S such that $\alpha \in T(e)$ for some $e \in E$. This will be our proof strategy. Firstly we show that a coherent but imperfect chronicle can be “improved” in some sense. For this we will find it useful to consider a specific kind of chronicle.

Definition 4.13. Let $S = (E; \leq)$ be a frame and T a coherent chronicle over S . T is said to be **strict** iff for all $e, e' \in E$, whenever $T(e) \preceq T(e')$ or $T(e') \preceq T(e)$, we have $e \leq e'$ or $e' \leq e$. \square

Definition 4.14. Let $S = (E; \leq)$ be a frame and T a coherent chronicle over S .

(i) A **live historic requirement** is a pair $(e, \diamond_i \alpha)$ such that $e \in E$, $\diamond_i \alpha \in T(e)$ and there does not exist $e' \in E_i$ such that $e' \leq e$ and $\alpha \in T(e')$.

(ii) A **live prophetic requirement** is a pair $(e, \diamond_i \alpha)$ such that $e \in E_i$, $\diamond_i \alpha \in T(e)$ and there does not exist $e' \in E_i$ such that $e \leq e'$ and $\alpha \in T(e')$.

(iii) A **live requirement** is either a live historic requirement or a live prophetic requirement. \square

Lemma 4.15. Let $S = (E; \leq)$ be a frame, $\hat{e} \notin E$ and T a *strict* and coherent chronicle on S . If $(e, \diamond_i \alpha)$ is a live *historic* requirement for T in S , then there exists a frame $S' = (E'; \leq')$ and a strict and coherent chronicle T' over S' such that

- (i) $E' = E \cup \{\hat{e}\}$.
- (ii) $\leq = \leq' \upharpoonright (E \times E)$.
- (iii) $T = T' \upharpoonright E$.
- (iv) $(e, \diamond_i \alpha)$ is *not* a live requirement for T' in S' .

Proof. Let $A = T(e)$. By Lemma 4.8, there exists an MCS B such that $B \preceq A$ and $\{\tau_i, \alpha\} \subseteq B$. Define

$$\begin{aligned} E'_j &\triangleq \begin{cases} E_j \cup \{\hat{e}\} & \text{if } i = j \text{ and} \\ E_j & \text{otherwise.} \end{cases} \\ T'(e') &\triangleq \begin{cases} B & \text{if } e' = \hat{e}, \text{ and} \\ T(e') & \text{otherwise.} \end{cases} \\ \leq' &\triangleq \leq \cup \{(e_1, \hat{e}) \mid T(e_1) \preceq B \text{ and not } B \preceq T(e_1), e_1 \in E\} \\ &\quad \cup \{(\hat{e}, e_2) \mid B \preceq T(e_2), e_2 \in E\} \\ &\quad \cup \{(\hat{e}, \hat{e})\}. \end{aligned}$$

Clearly, conditions (i) through (iii) are satisfied. The definition of \leq' ensures that $\hat{e} \leq' e$, since $B \preceq A$, by choice of B . Thus, if S' is a frame and T' is a strict and coherent chronicle over it, $(e, \diamond_i \alpha)$ cannot be a live requirement for T' in S' . To show that S' is a frame, we have to prove that $E'_i \cap E'_j = \emptyset$, for $i \neq j$, and that \leq' is a partial order which is backward linear within agents.

Since $\hat{e} \notin E$, and $E'_j = E_j$ for $j \neq i$, we also have $E'_i \cap E'_j = \emptyset$, for $i \neq j$. Reflexivity of \leq' follows from the definition, since \leq is reflexive.

To see that \leq' is antisymmetric, suppose that $e_1 \leq' e_2$ and $e_2 \leq' e_1$. Since $\leq = \leq' \upharpoonright (E \times E)$ is antisymmetric, if e_1, e_2 are both in E , then $e_1 = e_2$. If $e_1 \notin E$ and $e_2 \notin E$, then $e_1 = e_2 = \hat{e}$. So consider the case when $e_1 \notin E$ and $e_2 \in E$. That is, $e_1 = \hat{e}$. Since $\hat{e} \leq' e_2$, we have $B \preceq T(e_2)$. But $e_2 \leq' \hat{e}$, hence $T(e_2) \preceq B$ and not $B \preceq T(e_2)$, which is a contradiction. Hence $e_2 \notin E$. So $e_2 = \hat{e}$, as required.

To show transitivity of \leq' , assume that $e_1 \leq' e_2 \leq' e_3$. As we have that $\leq = \leq' \upharpoonright (E \times E)$, it suffices to consider the case when one of e_1, e_2, e_3 equals \hat{e} .

($e_1 = \hat{e}$):

We have $B \preceq T(e_2) \preceq T(e_3)$. By transitivity of \preceq (Proposition 4.6), $B \preceq T(e_3)$, and by definition, $\hat{e} \leq' e_3$.

($e_2 = \hat{e}$):

We have $T(e_1) \preceq B \preceq T(e_3)$, and not $B \preceq T(e_1)$. Clearly we cannot have $T(e_3) \preceq T(e_1)$ either. (For, otherwise, we get $B \preceq T(e_3) \preceq T(e_1)$, and by transitivity of \preceq , we get $B \preceq T(e_1)$.) Thus we get $T(e_1) \preceq T(e_3)$ and not $T(e_3) \preceq T(e_1)$. By coherence and strictness of T , we get $e_1 \leq e_3$. Hence $e_1 \leq' e_3$ as well.

($e_3 = \hat{e}$):

We have $T(e_1) \preceq T(e_2) \preceq B$, and not $B \preceq T(e_2)$. Thus we get $T(e_1) \preceq B$ and not $B \preceq T(e_1)$. Hence $e_1 \leq' \hat{e}$.

Thus, \leq' is a partial order on $E' \times E'$.

From the definitions of T' and \leq' , it is easy to see that for any $e \in E'$, $e \in E'_j$ iff $\tau_j \in T'(e)$ and that for any $e_1, e_2 \in E'$,

(*) $e_1 \leq' e_2$ implies $T'(e_1) \preceq T'(e_2)$, and

(**) if $T'(e_1) \preceq T'(e_2)$ or $T'(e_2) \preceq T'(e_1)$ then $e_1 \leq' e_2$ or $e_2 \leq' e_1$.

Now, to show that \leq' is backward linear within agents, let e_1, e_2 and $e_3 \in E'$ such that for some j , $\{e_1, e_2\} \subseteq E'_j$, $e_1 \leq' e_3$ and $e_2 \leq' e_3$. By (*) above, we have $T'(e_1) \preceq T'(e_3)$ and $T'(e_2) \preceq T'(e_3)$. By Proposition 4.7, either $T'(e_1) \preceq T'(e_2)$ or $T'(e_2) \preceq T'(e_1)$. By (**) above, either $e_1 \leq' e_2$ or $e_2 \leq' e_1$, as required.

Thus $S' = (E'; \leq')$ is a frame and T' is a coherent and strict chronicle over it. \square

Lemma 4.16. Let $S = (E; \leq)$ be a frame, $\hat{e} \notin E$ and T a strict and coherent chronicle on S . If $(e, \Diamond_i \alpha)$ is a live prophetic requirement for T in S , then there exists a frame $S' = (E'; \leq')$ and a strict and coherent chronicle T' over S' such that

(i) $E' = E \cup \{\hat{e}\}$.

(ii) $\leq = \leq' \upharpoonright (E \times E)$.

(iii) $T = T' \upharpoonright E$.

(iv) $(e, \Diamond_i \alpha)$ is not a live requirement for T' in S' .

Proof. Let $A = T(e)$. $\tau_i \wedge \Diamond_i \alpha \in T(e)$, and by Lemma 4.9, there exists an MCS B such that $A \preceq B$ and $\{\tau_i, \alpha\} \subseteq B$. Define

$$\begin{aligned} E'_j &\triangleq \begin{cases} E_j \cup \{\hat{e}\} & \text{if } i = j \text{ and} \\ E_j & \text{otherwise} . \end{cases} \\ T'(e') &\triangleq \begin{cases} B & \text{if } e' = \hat{e}, \text{ and} \\ T(e') & \text{otherwise} . \end{cases} \\ \leq' &\triangleq \leq \cup \{(e_1, \hat{e}) \mid T(e_1) \preceq B, e_1 \in E\} \\ &\quad \cup \{(\hat{e}, e_2) \mid B \preceq T(e_2) \text{ and not } T(e_2) \preceq B, e_2 \in E\} \\ &\quad \cup \{(\hat{e}, \hat{e})\} . \end{aligned}$$

The proof that S' is a frame and that T' is a strict and coherent chronicle over it is similar to the one in the proof of Lemma 4.15. Since $T(e) \preceq B$ by choice of B , by definition of \preceq' , $e \preceq' \hat{e}$ and hence $\diamond_i \alpha$ is not a live requirement for T' in S' . \square

Theorem 4.17. (Completeness)

If $\vDash \alpha$ then $\vdash \alpha$.

Proof. We show that every consistent formula is satisfiable. Let $\hat{E} = \{e_0, e_1, e_2, \dots\}$ be a countably infinite set. Fix an enumeration of $\hat{E} \times \Phi_n$ (Φ_n is the set of all formulas).

Now, let α be a consistent formula. Pick an MCS A containing α . We now define, for all $k \geq 0$, S^k and T^k .

$$S^0 \triangleq (\{e_0\}, \{(e_0, e_0)\}) \text{ and } T^0(e_0) \triangleq A.$$

Clearly, T^0 is strict and coherent over S^0 .

Inductively assume that $S^k = (E^k; \leq^k)$ and T^k have been defined, where $E^k = \{e_0, e_1, \dots, e_k\}$ and T^k is strict and coherent over S^k . Suppose there are no live requirements for T^k in S^k . Then set $S^{k+1} = S^k$ and $T^{k+1} = T^k$. Otherwise, among all the live requirements for T^k in S^k , choose the least one in the enumeration of $\hat{E} \times \Phi_n$, say (e, β) . By Lemma 4.15 and Lemma 4.16, we can extend S^k and T^k to a frame S^{k+1} and a chronicle T^{k+1} over it with $S^{k+1} = (E^{k+1}; \leq^{k+1})$ such that

- (i) $E^{k+1} = E^k \cup \{e_{k+1}\}$,
- (ii) $\leq^{k+1} \upharpoonright (E^k \times E^k)$ is \leq^k ,
- (iii) $T^{k+1} \upharpoonright E^k = T^k$,
- (iv) T^{k+1} is strict and coherent over S^{k+1} and
- (v) (e, β) is not a live requirement for T^{k+1} in S^{k+1} .

Finally set $S = (E; \leq)$, where $E = \bigcup_k E^k$, $\leq = \bigcup_k \leq^k$, and define a chronicle T over S by:

$$\text{for } e \in E, T(e) \triangleq T^k(e), \text{ where } e \in E^k.$$

It can be easily checked that T is a perfect chronicle over S . By Lemma 4.12, $M_T, e_0 \vDash \alpha$, where $M_T = (S, V_T)$, V_T being the valuation induced by T . Thus, α is indeed satisfiable. \square

5. Unboundedly Many Agents

So far we have a parametrized collection of classes of frames, where the parameter is n , the number of asynchronously communicating agents. Correspondingly, we have a collection of logical languages Φ_n . For each $n \in \mathbb{N}$, the required sound and complete axiomatization is obtained by varying the Axiom (A7.a).

We now wish to provide a uniform way of handling these classes of frames, languages and axiomatizations. Stated differently, we now wish to extend our study to handle systems of finite but unboundedly many agents. For the frames the generalization we have in mind is easy to achieve. We postulate a countably finite set \mathcal{N} of *names* (of agents). In order to tie up smoothly with the results of the

previous sections we will in fact set $\mathcal{N} = \mathbb{N}$, the set of *positive integers*. We let i, j, k range over \mathbb{N} . Through the rest of the paper we also let U range over the set of *non-empty finite subsets* of \mathbb{N} .

Definition 5.1. A system of **asynchronously Communicating Sequential Agents (ACSA)** is a triple (E, \leq, η) , where

- (i) (E, \leq) is a poset,
- (ii) $\eta: E \rightarrow \mathbb{N}$ is a (**naming**) function such that:
for all $e \in E$, for all $j \in \mathbb{N}$,
 $\downarrow e \cap \eta^{-1}(j)$ is totally ordered by \leq . □

In what follows we will usually write E_j instead of $\eta^{-1}(j)$ for $j \in \mathbb{N}$. Note that an ACSA (E, \leq, η) with the range of η restricted to $\{1, 2, \dots, n\}$ is in fact an n -ACSA. The notion of local state will remain as before. We will continue to interpret formulas at the local states of ACSAs.

The required logical language is obtained by postulating a countably infinite set of *type* propositions $T \triangleq \{\tau_i \mid i \in \mathbb{N}\}$ which is required to be disjoint from P , the countably infinite set of *atomic* propositions. The set of formulas is then inductively given by:

- (i) Every member of $P \cup T$ is a formula.
- (ii) If α and β are formulas then so are $\sim \alpha$, $\alpha \vee \beta$, $\Box_i \alpha$ and $\Box_i \alpha$ where $i \in \mathbb{N}$.

We thus have a language $\Phi_{\mathbb{N}}$ with an infinite set of modalities. However, since formulas are finite objects, each formula can refer to only a finite number of agents. Further, it is easy to see that $\Phi_{\mathbb{N}} = \bigcup_{n \in \mathbb{N}} \Phi_n$.

A **frame** is an ACSA $S = (E, \leq, \eta)$. A **model** is a pair $M = (S, V)$ where $S = (E, \leq, \eta)$ is a frame and $V: E \rightarrow 2^{P \cup T}$ is a **valuation function** such that

$$\text{for every } e \in E \text{ and every } i \in \mathbb{N}, \quad \tau_i \in V(e) \text{ iff } e \in E_i.$$

Let $M = (S, V)$ be a model with $S = (E, \leq, \eta)$ and $e \in E$ and α a formula. Then $M, e \models \alpha$ denotes the notion of α being satisfied at the local state $\downarrow e$ in the model M and is defined exactly as in Sec. 2. We will once again use $\models \alpha$ to denote that α is a valid formula in this new set-up. The derived modalities are defined as before.

As for an axiomatization, we retain the axiom schemes (A1) through (A6) and the inference rules (MP), (TG \Box_i) and (TG \Box_i) from Sec. 3. The two latter parts of (A7) will also be carried over unchanged. However, (A7.a) will no longer be sound. An equivalent version in the present set-up would require an infinite disjunction. To get around this problem we first capture the fact that each event can belong to *at most* one agent by replacing (A7.a) with an infinite set of *axioms*: $\tau_i \supset \sim \tau_j$, for every pair $i, j \in \mathbb{N}, i \neq j$.

We will “weakly” capture the fact that each event must belong to *at least one* agent by introducing a new inference rule. In doing so we will use $\text{TYP}(\alpha)$ to denote the set of agents referred to by the formula α . This notation can be inductively

defined as follows:

$$\begin{aligned}
 \text{TYP}(p) &= \emptyset, \text{ for } p \in P. \\
 \text{TYP}(\tau_i) &= \{i\}. \\
 \text{TYP}(\sim \alpha) &= \text{TYP}(\alpha). \\
 \text{TYP}(\alpha \vee \beta) &= \text{TYP}(\alpha) \cup \text{TYP}(\beta). \\
 \text{TYP}(\exists_i \alpha) &= \text{TYP}(\Box_i \alpha) = \text{TYP}(\alpha) \cup \{i\}.
 \end{aligned}$$

We can now state the inference rule (TE) which guarantees that each event belongs to at least one agent. The rule is similar to those used in [21,22].

$$(\text{TE}) \quad \frac{\bigvee_{i \in U} \tau_i \supset \alpha}{\alpha} \quad (\text{where } U \text{ is such that } \text{TYP}(\alpha) \subsetneq U)$$

Observe that U is a non-empty finite subset of \mathbb{N} and hence the disjunction is allowed. Moreover, $\text{TYP}(\alpha)$ is a *proper* subset of U . The proposed axiomatization $\mathcal{A}(\text{ACSA})$ is given below.

Axioms

- (A0) All substitutional instances of the tautologies of propositional logic.
- (A1) (a) $\exists_i(\alpha \supset \beta) \supset (\exists_i \alpha \supset \exists_i \beta)$ (Deductive closure)
 (b) $\Box_i(\alpha \supset \beta) \supset (\Box_i \alpha \supset \Box_i \beta)$
- (A2) (a) $\tau_i \supset (\exists_i \alpha \supset \alpha)$ (Local reflexivity)
 (b) $\tau_i \supset (\Box_i \alpha \supset \alpha)$
- (A3) $\diamond_i \diamond_j \alpha \supset \diamond_j \alpha$ (Transitivity)
- (A4) $\diamond_i \alpha \wedge \diamond_i \beta \supset \diamond_i(\alpha \wedge \diamond_i \beta) \vee \diamond_i(\beta \wedge \diamond_i \alpha)$ (Backward linearity)
- (A5) (a) $\diamond_i \alpha \supset \Box_i \diamond_i \alpha$ (Relating past and future)
 (b) $\diamond_i \alpha \supset \exists_i \diamond_i \alpha$
- (A6) $\Box_i \alpha \supset \diamond_i \Box_i \alpha$ (Communication)
- (A7) (a) $\exists_i \tau_i$ (Type axioms)
 (b) $\tau_i \supset \Box_i \tau_i$
- (A8) $\tau_i \supset \sim \tau_j \quad (i \neq j)$ (Disjoint agents)

Inference Rules

$$(\text{MP}) \quad \frac{\alpha, \alpha \supset \beta}{\beta}$$

$$(\text{TG}\exists_i) \quad \frac{\alpha}{\exists_i \alpha}$$

$$(\text{TG}\Box_i) \quad \frac{\alpha}{\tau_i \supset \Box_i \alpha}$$

$$(\text{TE}) \quad \frac{\bigvee_{i \in U} \tau_i \supset \alpha}{\alpha} \quad (\text{where } U \text{ is such that } \text{TYP}(\alpha) \subsetneq U)$$

As before we shall use $\vdash \alpha$ to denote the fact that α is a thesis in this new logical system. The notion of a consistent formula is defined as usual.

Theorem 5.2. (Soundness)

If $\vdash \alpha$ then $\models \alpha$.

Proof. It is easy to check the soundness of the axiom schemes, axioms and the first three inference rules. What requires an argument is the soundness of the new inference rule (TE).

Let α be a formula and U a proper superset of $\text{TYP}(\alpha)$, such that the formula $\bigvee_{i \in U} \tau_i \supset \alpha$ is valid. To prove that α is also valid, we must show that for any model $M = (S, V)$, where $S = (E, \leq, \eta)$, and for any $e_0 \in E$, $M, e_0 \vDash \alpha$.

Suppose that $\eta(e_0) = j \in U$. Then $M, e_0 \vDash \tau_j$ and hence $M, e_0 \vDash \bigvee_{i \in U} \tau_i$. But then $\bigvee_{i \in U} \tau_i \supset \alpha$ is a valid formula and hence is satisfied at e_0 in the model M . Consequently $M, e_0 \vDash \alpha$ as well.

Suppose that $\eta(e_0) = j \notin U$. Let $i \in U - \text{TYP}(\alpha)$. The existence of such an i is assured by the fact that $\text{TYP}(\alpha)$ is a proper subset of U . Now define $S' = (E, \leq, \eta')$ where η' is given by

$$\forall e \in E. \eta'(e) \triangleq \begin{cases} i, & \text{if } \eta(e) = j. \\ j, & \text{if } \eta(e) = i. \\ \eta(e), & \text{otherwise.} \end{cases}$$

It is easy to check that S' is a frame. Now define the model $M' = (S', V')$ where V' is given by

$$\forall e \in E. V'(e) \triangleq \begin{cases} (V(e) - \{\tau_j\}) \cup \{\tau_i\}, & \text{if } \eta(e) = j. \\ (V(e) - \{\tau_i\}) \cup \{\tau_j\}, & \text{if } \eta(e) = i. \\ V(e), & \text{otherwise.} \end{cases}$$

It is easy to check that M' is indeed a model.

Claim. For every formula δ such that $\{i, j\} \cap \text{TYP}(\delta) = \emptyset$ and every $e \in E$, $M, e \vDash \delta$ iff $M', e \vDash \delta$.

Proof of Claim. By induction on the structure of δ .

Suppose that $\delta \in P$. Then since V and V' assign identical subsets of P to identical events, the claim follows.

Suppose that $\delta \in T$. Let $\delta = \tau_k$. Then $k \notin \{i, j\}$. But this implies that $\eta(e) = k$ iff $\eta'(e) = k$ and once again the claim follows.

If δ is of the form $\sim \gamma$ or of the form $\delta_1 \vee \delta_2$ then standard arguments lead to the claim.

If δ is of the form $\exists_k \gamma$ or $\Box_k \gamma$, $k \in \text{TYP}(\delta)$ and hence $k \notin \{i, j\}$. Therefore, for all $e' \in E$, $\eta(e') = k$ iff $\eta'(e') = k$. Then it is easy to verify that $\forall e \in E$, $M, e \vDash \delta$ iff $M', e \vDash \delta$, using the induction hypothesis. This establishes the claim. \square

Returning to our soundness proof, by the definition of η' , $\eta'(e_0) = i \in U$. Then $M', e_0 \vDash \tau_i$ and hence $M', e_0 \vDash \bigvee_{i \in U} \tau_i$. But then $\bigvee_{i \in U} \tau_i \supset \alpha$ is a valid formula and hence is satisfied at e_0 in the model M' . Consequently, $M', e_0 \vDash \alpha$. Since $\{i, j\} \cap$

$\text{TYP}(\alpha) = \emptyset$, by the claim above, $M, e_0 \models \alpha$ as well, and thus the inference rule (TE) is sound. \square

Proposition 5.3. If α is consistent then there exists $i \in \mathbb{N}$ such that $\alpha \wedge \tau_i$ is also consistent.

Proof. Suppose not. Then for all $i \in \mathbb{N}$, we have $\vdash \alpha \supset \sim \tau_i$. Let $U = \text{TYP}(\alpha) \cup \{j\}$, where $j \notin \text{TYP}(\alpha)$. We have $\vdash \alpha \supset \bigwedge_{i \in U} \sim \tau_i$. Thus, $\vdash \bigvee_{i \in U} \tau_i \supset \sim \alpha$, and using (TE), we get $\vdash \sim \alpha$, which contradicts the consistency of α . \square

It is fairly easy to prove the completeness of our axiomatization using the results of the previous section. A direct approach to produce a model for a consistent formula will however yield, in general, a model consisting of an infinite set of agents even though the consistent formula we started out with could only talk about a finite number of agents. To eliminate this “slack” we will relativize the notions concerning chronicles w.r.t. the *closure* of a formula. This idea will turn out to be essential in the next section where we axiomatize finitary ACSAs.

Definition 5.4. Let α be a formula.

(i) $\text{CL}'(\alpha)$ is the least set of formulas containing α that satisfies the following conditions:

- (a) If $i \in \text{TYP}(\alpha)$ then $\tau_i \in \text{CL}'(\alpha)$.
- (b) If $\sim \beta \in \text{CL}'(\alpha)$ then $\beta \in \text{CL}'(\alpha)$.
- (c) If $\beta_1 \vee \beta_2 \in \text{CL}'(\alpha)$ then $\beta_1, \beta_2 \in \text{CL}'(\alpha)$.
- (d) If $\exists_i \beta \in \text{CL}'(\alpha)$ then $\beta \in \text{CL}'(\alpha)$.
- (e) If $\Box_i \beta \in \text{CL}'(\alpha)$ then $\beta, \diamond_i \Box_i \beta \in \text{CL}'(\alpha)$.

(ii) The **closure** of α (denoted $\text{CL}(\alpha)$) is given by:

$$\text{CL}(\alpha) \triangleq \text{CL}'(\alpha) \cup \{\sim \beta \mid \beta \in \text{CL}'(\alpha)\} . \quad \square$$

It is easy to check that $\text{CL}(\alpha)$ is a finite set for every α . (In fact there exists a constant c such that the cardinality of $\text{CL}(\alpha)$ is at most c times the length of the formula α).

Let A be an MCS. We will say that the MCS A is **good** if there exists *at least one* $i \in \mathbb{N}$ such that $\tau_i \in A$. Throughout the completeness proof, we shall restrict our attention to good MCSs; in fact, we will usually say “an MCS” to mean “a good MCS”. It is easy to see that the \preceq relation defined in Sec. 4 by:

$$A \preceq B \text{ iff } \{\diamond_i \alpha \mid \alpha \in A\} \subseteq B$$

is meaningful on good MCSs. Further, the results Proposition 4.1 through Lemma 4.9 go through smoothly for good MCSs. The proof of Proposition 4.3, however, is different now, as it appeals to Axiom (A8). A chronicle over a frame (E, \leq, η) is a function T which assigns a good MCS to each e in E . As before T is said to be

coherent iff $e \leq e'$ implies that $T(e) \preceq T(e')$ for all $e, e' \in E$. However, the notions of historic and prophetic chronicles are now relativized, as given below.

Definition 5.5. Let L be a set of formulas and T a chronicle over the frame (E, \leq, η) .

- (i) T is **L -historic** iff whenever $e \in E$ and $\diamond_i \beta \in T(e) \cap L$ then there exists $e' \in E_i$ such that $e' \leq e$ and $\beta \in T(e')$.
- (ii) T is **L -prophetic** iff whenever $e \in E_i$ and $\diamond_i \beta \in T(e) \cap L$ then there exists $e' \in E_i$ such that $e \leq e'$ and $\beta \in T(e')$.
- (iii) T is **L -perfect** iff T is coherent, L -historic and L -prophetic. \square

Definition 5.6. Let $S = (E, \leq, \eta)$ be a frame, T a chronicle for it and L a set of formulas. The **valuation induced by T for L** , denoted V_T^L , is given by:

$$\text{for } e \in E, V_T^L(e) \triangleq T(e) \cap (P \cup T) \cap L.$$

We use M_T^L to denote (S, V_T^L) . \square

Lemma 5.7. Let T be a $\text{CL}(\alpha)$ -perfect chronicle over the frame (E, \leq, η) . Then for every $e \in E$ and every $\beta \in \text{CL}(\alpha)$, $\beta \in T(e)$ iff $M_T^{\text{CL}(\alpha)}, e \vDash \beta$.

Proof. Similar to the proof of Lemma 4.12; the only necessary observation is that when $\Box_i \beta \in \text{CL}(\alpha)$, $\diamond_i \Box_i \beta \in \text{CL}(\alpha)$ as well. \square

Corresponding to L -perfect chronicles, we now have “ L -imperfections” in a chronicle, which can be judged with the help of live L -requirements.

Definition 5.8. Let T be a chronicle over the frame $S = (E, \leq, \eta)$ and L a set of formulas.

- (i) A **live L -historic requirement** is a pair $(e, \diamond_i \beta)$ such that $e \in E$ and $\diamond_i \beta \in T(e) \cap L$ and there does *not* exist $e' \in E_i$ such that $e' \leq e$ and $\beta \in T(e')$.
- (ii) A **live L -prophetic requirement** is a pair $(e, \diamond_i \beta)$ such that $e \in E_i$ and $\diamond_i \beta \in T(e) \cap L$ and there does *not* exist $e' \in E_i$ such that $e \leq e'$ and $\beta \in T(e')$.
- (iii) A **live L -requirement** is either a live L -historic requirement or a live L -prophetic requirement. \square

We can once again verify that the Lemmas 4.15 and 4.16 go through when we use good MCSs in the context of $\text{CL}(\alpha)$ -requirements, for any formula α .

Theorem 5.9. (Completeness)

If $\vDash \alpha$ then $\vdash \alpha$.

Proof. We will show that every consistent formula is satisfiable. To this end, let α be a consistent formula. Then according to Proposition 5.3, there exists an i in \mathbb{N} such that $\alpha \wedge \tau_i$ is consistent. Set $\hat{\alpha} = \alpha \wedge \tau_i$. Fix a countably infinite set of events \hat{E} and fix an enumeration e_0, e_1, \dots , of \hat{E} . Also fix an enumeration of $\hat{E} \times \text{CL}(\hat{\alpha})$.

The proof proceeds exactly as in the proof of Theorem 4.17; we build up a sequence of frames and chronicles S^k and T^k as before. The only changes to be

noted are:

(a) $T^0(e_0)$ is set to A , where A is any maximal consistent set containing $\hat{\alpha}$ (and hence A is good),

(b) given the frame S^k and a chronicle T^k over it, we choose that live $\text{CL}(\hat{\alpha})$ -requirement which has the least index in the enumeration of $\hat{E} \times \text{CL}(\hat{\alpha})$.

Again, as before, the countable componentwise union yields the desired frame S and it can be easily checked that the chronicle T over S defined by

$$\forall e \in E. T(e) \triangleq T^k(e) \text{ where } e \in E^k$$

is a $\text{CL}(\hat{\alpha})$ -perfect chronicle over S with $\alpha \in T(e_0)$. Hence by Lemma 5.7, we have $M_T^{\text{CL}(\hat{\alpha})}, e_0 \models \alpha$. \square

We noted earlier that $\bigcup_n \Phi_n = \Phi_{\mathbb{N}}$. What is more interesting is that such an equation holds also for the *satisfiable formulas* of these languages. Let $\text{SAT}(n\text{-ACSA})$ and $\text{SAT}(\text{ACSA})$ be as given below:

$$\text{SAT}(n\text{-ACSA}) \triangleq \{\alpha \in \Phi_n \mid \alpha \text{ is satisfiable in a model based on an } n\text{-ACSA}\}.$$

$$\text{SAT}(\text{ACSA}) \triangleq \{\alpha \in \Phi_{\mathbb{N}} \mid \alpha \text{ is satisfiable in a model based on an ACSA}\}.$$

Then the method of constructing models for consistent formulas in $\Phi_{\mathbb{N}}$ leads at once to the following result.

Corollary 5.10. $\bigcup_n \text{SAT}(n\text{-ACSA}) = \text{SAT}(\text{ACSA})$. \square

While we have used the rule (TE) to ensure that we always work with good MCSs, it is not necessary for completeness. To capture the effect of the rule in the model construction process, one has to relativize the notion of a chronicle so that the condition ($\tau_i \in T(e)$ iff $e \in E_i$) needs to be maintained only for $\tau_i \in \text{CL}(\alpha)$. Completeness proofs of this kind will be found in Part II.

6. Finitary Frames

In this section, we consider the problem of axiomatizing finitary n -ACSAs (and later, finitary ACSAs). A finitary n -ACSA is one in which every event has a finite past. The motivation for studying such frames is clear: if n -ACSAs were to be used to give the semantics of n asynchronously communicating sequential programs, we would expect the semantics to yield only finitary structures. Moreover, finitariness implies well-foundedness and hence induction principles based on well-foundedness become available for proving properties of n -ACSAs [2].

Definition 6.1. A finitary n -ACSA is an n -ACSA $(E_1, \dots, E_n; \leq)$, such that $\downarrow e$ is a finite set for every $e \in E = \bigcup_i E_i$.

We will continue to work with the language Φ_n defined in Sec. 2. By a frame we will mean a finitary n -ACSA and by a model, we will mean a pair (S, V) where S is a finitary frame and V is a valuation as before. The semantics of the various modalities will continue to remain the same.

The axiom system $\mathcal{A}(\text{fin-}n\text{-ACSA})$ is defined to be $\mathcal{A}(n\text{-ACSA})$ enriched by the additional axiom schemes:

$$(A9) \quad \begin{array}{ll} \text{(a)} \ \diamond_i \alpha \supset \diamond_i (\alpha \wedge \Box_i (\sim \alpha \supset \Box_i \sim \alpha)) & \text{(Well-founded agents)} \\ \text{(b)} \ \diamond_i \alpha \supset \diamond_i (\alpha \wedge \Box_j \Box_i \sim \alpha) \quad (j \neq i) & \text{(Well-founded communications)} \end{array}$$

The first of the two axioms rules out infinite descending chains of i -events and the second rules out infinite descending chains of communication. Interestingly, both axioms are indexed versions of well-known axioms for well-founded structures in modal logic [23]: (A9.a) is a weak form of the *Grz* axiom for reflexive structures, while (A9.b) derives from the *W* axiom for irreflexive structures. Note that \Box_i is a reflexive operator and $\Box_j \Box_i$ (when $j \neq i$) is an irreflexive operator in our system.

Proposition 6.2. (A9) is sound.

Proof. (Soundness of (A9.a)).

Assume some model $M = (S, V)$, where $S = (E; \leq)$ and $e \in E$ such that $M, e \not\models (A9.a)$.

We have $M, e \models \diamond_i \alpha \wedge \Box_i (\alpha \supset \diamond_i (\sim \alpha \wedge \diamond_i \alpha))$. Thus, for some $e_1 \leq e$, $e_1 \in E_i$, $M, e_1 \models \alpha \wedge \Box_i (\alpha \supset \diamond_i (\sim \alpha \wedge \diamond_i \alpha))$, and hence $M, e_1 \models \diamond_i (\sim \alpha \wedge \diamond_i \alpha)$. Now we can find $e_2 \in E_i$ such that $e_2 \leq e_1$ and $M, e_2 \models \sim \alpha \wedge \diamond_i \alpha \wedge \Box_i (\alpha \supset \diamond_i (\sim \alpha \wedge \diamond_i \alpha))$, where $e_2 \leq e_1$. Since $M, e_1 \models \alpha$ and $M, e_2 \models \sim \alpha$, $e_2 \neq e_1$ and thus $e_2 < e_1$. We proceed now to find an e_3 such that $e_3 \leq e_2$ and $M, e_3 \models \alpha \wedge \Box_i (\alpha \supset \diamond_i (\sim \alpha \wedge \diamond_i \alpha))$. Clearly, $e_3 < e_2$. We have returned to the situation as in e_1 and hence can repeat the construction to get the desired infinite descending chain.

(Soundness of (A9.b)).

Assume some model $M = (S, V)$, where $S = (E; \leq)$ and $e \in E$ such that $M, e \not\models (A9.b)$.

We have $M, e \models \diamond_i \alpha \wedge \Box_i (\alpha \supset \diamond_j \diamond_i \alpha)$ with $i \neq j$. Thus, for some $e_1 \leq e$, we have $e_1 \in E_i$ and $M, e_1 \models \alpha \wedge \Box_i (\alpha \supset \diamond_j \diamond_i \alpha)$, and hence $M, e_1 \models \diamond_j \diamond_i \alpha$. Now we can find $e_2 \in E_j$ such that $e_2 \leq e_1$ and $M, e_2 \models \diamond_i \alpha \wedge \Box_i (\alpha \supset \diamond_j \diamond_i \alpha)$. Since $e_1 \in E_i$ and $e_2 \in E_j$, $e_2 \neq e_1$. We proceed now to find an e_3 in E_i such that $e_3 \leq e_2$ and $M, e_3 \models \alpha \wedge \Box_i (\alpha \supset \diamond_j \diamond_i \alpha)$. Clearly, $e_3 \leq e_2$. We have returned to the situation as in e_1 and hence can repeat the construction to get the desired infinite descending chain. \square

This is to be understood relative to the larger axiom system we now have. $\vdash \alpha$ denotes that α is a thesis in $\mathcal{A}(\text{fin-}n\text{-ACSA})$. The only additional thesis required to prove completeness is:

$$(T7) \quad \diamond_i \alpha \supset \diamond_i \left(\alpha \wedge \bigwedge_{j \neq i} \Box_j \Box_i \sim \alpha \wedge \Box_i (\sim \alpha \supset \Box_i \sim \alpha) \right)$$

Derivation.

$$(1) \quad \bigwedge_{j \neq i} \Box_j \Box_i \alpha \supset \Box_i \bigwedge_{j \neq i} \Box_j \Box_i \alpha \quad (A3)$$

$$(2) \quad \diamond_i \alpha \supset \diamond_i \left(\alpha \wedge \bigwedge_{j \neq i} \Box_j \Box_i \sim \alpha \right) \quad (\text{A9.b, similar to T6})$$

$$(3) \quad \diamond_i \alpha \subset \diamond_i \left(\alpha \wedge \Box_i \bigwedge_{j \neq i} \Box_j \Box_i \sim \alpha \right) \quad (1, 2, \text{DR3})$$

$$(4) \quad \diamond_i \alpha \subset \diamond_i (\alpha \wedge \Box_i \Box_i (\sim \alpha \supset \Box_i \sim \alpha)) \quad (\text{A9.a, A3, DR3})$$

$$(5) \quad \diamond_i \alpha \supset \diamond_i \left(\alpha \wedge \Box_i \left(\bigwedge_{j \neq i} \Box_j \Box_i \sim \alpha \wedge \Box_i (\sim \alpha \supset \Box_i \sim \alpha) \right) \right) \quad (3, 4, \text{T6, T1.a})$$

$$(6) \quad \diamond_i \alpha \supset \diamond_i \left(\alpha \wedge \tau_i \wedge \Box_i \left(\bigwedge_{j \neq i} \Box_j \Box_i \sim \alpha \wedge \Box_i (\sim \alpha \supset \Box_i \sim \alpha) \right) \right) \quad (5, \text{T4.a})$$

$$(7) \quad \diamond_i \alpha \supset \diamond_i \left(\alpha \wedge \bigwedge_{j \neq i} \Box_j \Box_i \sim \alpha \wedge \Box_i (\sim \alpha \supset \Box_i \sim \alpha) \right) \quad (6, \text{A2.a, DR3})$$

In the rest of this section, by α_i , we shall denote the formula $\alpha \wedge \bigwedge_{j \neq i} \Box_j \Box_i \sim \alpha \wedge \Box_i (\sim \alpha \supset \Box_i \sim \alpha)$. We have the thesis $\vdash \diamond_i \alpha \equiv \diamond_i \alpha_i$.

We now proceed to show that $\mathcal{A}(\text{fin-}n\text{-ACSA})$ is complete for finitary n -ACSAs. The proof will proceed along the same lines as in Sec. 4.

It should firstly be noted that the model constructed in the proof of Theorem 4.17 *may not* be finitary. It turns out that a good deal of work must be put in, using (A9.a) and (A9.b), in order to ensure that finitariness is maintained during the model construction process. Below we present a proof, where we ensure that once the past of an event is determined, it will not be changed ever again in the process of killing a requirement.

Hereafter we shall say MCS to mean a maximal consistent set of formulas in the enlarged system. We assume the same ordering \leq between MCSs. Propositions 4.5 to 4.9 go through without any change. However, we shall find occasion to use a stronger form of Lemma 4.8. This is due to the fact that we now need to work with a strengthened notion of historic chronicles.

Let $S = (E; \leq)$ be a frame and T a chronicle over S . Let $e, e' \in E$ and $\diamond_i \alpha \in T(e)$. Then e' **kills** the requirement $(e, \diamond_i \alpha)$ in S for T if $e' \leq e$, $\{\tau_i, \alpha\} \subseteq T(e')$ and for any MCS C containing $\{\tau_i, \alpha\}$ and satisfying $C \preceq T(e)$ it is the case that $T(e') \preceq C$. $(e, \diamond_i \alpha)$ is said to be a **live historic requirement** for the chronicle T in the frame $S = (E; \leq)$ if $\diamond_i \alpha \in T(e)$ and there does not exist an $e' \in E$ which kills $(e, \diamond_i \alpha)$ in S for T . T is said to be **historic** iff there is no live historic requirement in S for T .

Thus in some sense $T(e')$ is the “virtually earliest” MCS that contains $\{\tau_i, \alpha\}$ and is “earlier than” $T(e)$.

The notion of a prophetic chronicle remains unchanged. Recall the relativized notions of historic and prophetic chronicles w.r.t a set L of formulas introduced in the previous section. L -historic chronicles are to be understood now in terms of the stronger notion of a historic chronicle introduced above. The definition of an

L -prophetic chronicle remains unchanged. Naturally, an L -perfect chronicle is one which is coherent, L -historic in the new sense and L -prophetic.

The relativized notion of a live L -historic requirement, where L is a set of formulas, is defined accordingly. The definition of a live prophetic requirement is the same as the one used in Sec. 4. The relativized notion of an L -prophetic requirement is defined in the obvious way.

Lemma 6.3. Let A be an MCS and let $\diamond_i \alpha \in A$. Then there exists an MCS B such that $B \preceq A$, $\{\tau_i, \alpha\} \subseteq B$ and for all MCSs $C \preceq A$, if $\{\tau_i, \alpha\} \subseteq C$, $B \preceq C$.

Proof. Consider the set

$$\Sigma \triangleq \{\gamma \mid \diamond_i(\alpha \wedge \Box_i \gamma) \in A\} \cup \{\tau_i, \alpha\}.$$

Suppose Σ is consistent. Then we can extend Σ to an MCS B . To show $B \preceq A$, let $\beta \in B$. If $\diamond_i \beta \notin A$, $\Box_i \sim \beta \in A$, and hence $\Box_i \Box_i \sim \beta \in A$. Since $\diamond_i \alpha \in A$, we get $\diamond_i(\alpha \wedge \Box_i \sim \beta) \in A$, and so $\sim \beta \in \Sigma \subseteq B$, a contradiction. Thus $B \preceq A$.

To show consistency of Σ , consider an arbitrary finite subset of Σ , say, $\Sigma' \triangleq \{\gamma_1, \dots, \gamma_k, \tau_i, \alpha\}$. It suffices to prove that Σ' is consistent. $\diamond_i(\alpha \wedge \Box_i \gamma_1), \dots, \diamond_i(\alpha \wedge \Box_i \gamma_k)$ is in A by definition of Σ . $\diamond_i(\alpha \wedge \Box_i \gamma_1 \wedge \dots \wedge \Box_i \gamma_k) \in A$ by thesis (T6). Hence $\diamond_i(\tau_i \wedge \alpha \wedge \Box_i \gamma_1 \wedge \dots \wedge \Box_i \gamma_k) \in A$ by thesis (T4.a). We get $\diamond_i(\tau_i \wedge \alpha \wedge \gamma_1 \wedge \dots \wedge \gamma_k) \in A$ by axiom (A2.a) and (DR3) and hence $\tau_i \wedge \alpha \wedge \gamma_1 \wedge \dots \wedge \gamma_k$ is consistent. That is, Σ' is consistent, as required.

Now consider an MCS C such that $C \preceq A$ and $\{\tau_i, \alpha\} \subseteq C$. We have to show that $B \preceq C$. Suppose not. Then there must be a formula $\delta \in B$ such that $\Box_i \sim \delta \in C$. Thus $\alpha \wedge \Box_i \sim \delta \in C$, and since $C \preceq A$, $\diamond_i(\alpha \wedge \Box_i \sim \delta) \in A$. By construction, $\sim \delta \in \Sigma \subseteq B$, which is a contradiction. Hence B is the required MCS. \square

In the previous section we define the closure of a formula in $\Phi_{\mathbb{N}}$. Here we must do the same for formulas in Φ_n . Let $\alpha \in \Phi_n$. Then $CL'(\alpha)$ is the least subset of Φ_n containing α which satisfies the following conditions:

- $\{\tau_1, \tau_2, \dots, \tau_n\} \subseteq CL'(\alpha)$.
- If $\sim \beta \in CL'(\alpha)$ then $\beta \in CL'(\alpha)$.
- If $\beta_1 \vee \beta_2 \in CL'(\alpha)$ then $\beta_1, \beta_2 \in CL'(\alpha)$.
- If $\Box_i \beta \in CL'(\alpha)$ then $\beta \in CL'(\alpha)$.
- If $\diamond_i \beta \in CL'(\alpha)$ then $\beta, \diamond_i \Box_i \beta \in CL'(\alpha)$.

The closure of α , denoted $CL(\alpha)$, is then given by:

$$CL(\alpha) \triangleq CL'(\alpha) \cup \{\sim \beta \mid \beta \in CL'(\alpha)\}.$$

We again have a $CL(\alpha)$ -perfect chronicle inducing a model.

Lemma 6.4. Let γ be a formula and let $L = CL(\gamma)$. If T is an L -perfect chronicle on a frame $S = (E; \preceq)$, then for any $e \in E$ and for any formula $\alpha \in L$, $\alpha \in T(e)$ iff $M_T^L, e \models \alpha$.

Proof. As for Lemma 5.7. \square

To show that a consistent formula is satisfiable, we fix a consistent formula α_0 and construct a $\text{CL}(\alpha_0)$ -perfect chronicle on a frame S . While doing so, every time we extend the “current” frame by adding an event e we also add a *finite* set of “past” events of e , which is fixed once and for all. The bulk of the proof will be devoted to constructing this fixed past of a new event. First let us observe that a frame can always be “historically” improved.

Lemma 6.5. Let $S = (E; \leq)$ be a frame, $\hat{e} \notin E$ and T a strict and coherent chronicle on S . If $(e, \diamond_i \alpha)$ is a live *historic* requirement for T in S , then there exists a frame $S' = (E'; \leq')$ and a strict and coherent chronicle T' over S' such that

- (i) $E' = E \cup \{\hat{e}\}$.
- (ii) $\leq = \leq' \upharpoonright (E \times E)$.
- (iii) $T = T' \upharpoonright E$.
- (iv) \hat{e} kills the requirement $(e, \diamond_i \alpha)$ in S' for T' .

Proof. Similar to the proof of Lemma 4.15, except that instead of appealing to Lemma 4.8 to find MCS B , we use Lemma 6.3 to get the virtually “earliest” B . \square

Now we can turn to the task of constructing the fixed finite past of a new event as dictated by the historical requirement generated by a *finite* set of formulas. Finiteness will be achieved by ensuring that live requirements, once killed, do not come up again. We crucially use the thesis (T7) for this; recall that α_i denotes the formula $\alpha \wedge \bigwedge_{j \neq i} \Box_j \Box_i \sim \alpha \wedge \Box_i (\sim \alpha \supset \Box_i \sim \alpha)$.

Proposition 6.6. For any MCSs A and B , if $A \preceq B$, $\alpha \in A$ and $\alpha_i \in B$, then $\alpha_i \in A$.

Proof. Let $\tau_k \in A$. By axioms (A3) and thesis (T1.a), the formula

$$\Box_k \left(\bigwedge_{j \neq i} \Box_j \Box_i \sim \alpha \wedge \Box_i (\sim \alpha \supset \Box_i \sim \alpha) \right) \in B.$$

Therefore, since $A \preceq B$, the formula $\bigwedge_{j \neq i} \Box_j \Box_i \sim \alpha \wedge \Box_i (\sim \alpha \supset \Box_i \sim \alpha) \in A$. But $\alpha \in A$ as well, and thus $\alpha_i \in A$. \square

Proposition 6.7. Let T be a coherent chronicle over $S = (E; \leq)$. Then for any e and e' in E , if e' kills the requirement $(e, \diamond_i \alpha)$ for T in S , then $\alpha_i \in T(e')$.

Proof. Suppose e_2 kills $(e_1, \diamond_i \alpha)$. $\diamond_i \alpha \in T(e_1)$ and by (T7), $\diamond_i \alpha_i \in T(e_1)$ as well. By Lemma 6.3, there exists $C \preceq T(e_1)$, such that $\{\tau_i, \alpha_i\} \subseteq C$. But then $\{\tau_i, \alpha\} \subseteq C$ also and since e_2 kills $(e_1, \diamond_i \alpha)$, $T(e_2) \preceq C$. But $\alpha \in T(e_2)$ and $\alpha_i \in C$; by Proposition 6.6, $\alpha_i \in T(e_2)$. \square

Proposition 6.8. Let $S = (E; \leq)$ be a frame and T a coherent and strict chronicle over S . Let $e_1, e_2 \in E$ such that $e_2 \leq e_1$ and $(e_1, \diamond_i \alpha)$ is not a live requirement for T in S . Then $(e_2, \diamond_i \alpha)$ is also not a live requirement for T in S .

Proof. Assume the hypothesis. If $\diamond_i \alpha \notin T(e_2)$, we are done. So assume that $\diamond_i \alpha \in T(e_2)$. Since $e_2 \leq e_1$ and T is coherent, we get $\diamond_i \alpha \in T(e_1)$. But $(e_1, \diamond_i \alpha)$

is not a live requirement for T in S , so there exists $e_3 \in E$ which kills $(e_1, \diamond_i \alpha)$. Since $\diamond_i \alpha \in T(e_2)$, by Lemma 6.3, there exists an MCS B such that $B \preceq T(e_2)$, $\{\tau_i, \alpha\} \subseteq B$. But $T(e_2) \preceq T(e_1)$, hence $B \preceq T(e_1)$. But e_3 kills $(e_1, \diamond_i \alpha)$; therefore $T(e_3) \preceq B$. By transitivity, we get $T(e_3) \preceq T(e_2)$.

By strictness of T , either $e_3 \leq e_2$ or $e_2 \leq e_3$. In the former case, it is easy to show that e_3 kills $(e_2, \diamond_i \alpha)$, since if an MCS C has α and $C \preceq T(e_2)$, then $C \preceq T(e_1)$ as well, so $T(e_3) \preceq C$. Assume the latter case. Since e_3 kills $(e_1, \diamond_i \alpha)$, by Proposition 6.7, $\alpha_i \in e_3$. Now $e_2 \leq e_3$, so by coherence, $T(e_2) \preceq T(e_3)$. If $e_2 \in E_j$, $j \neq i$, $\Box_i \sim \alpha \in T(e_2)$, contradicting $\diamond_i \alpha \in T(e_2)$. Thus $e_2 \in E_i$. Then the formula $(\sim \alpha \supset \Box_i \sim \alpha)$ is contained in $T(e_2)$, therefore to avoid a contradiction α must be in $T(e_2)$. We now show that e_2 itself kills $(e_2, \diamond_i \alpha)$.

We already have that $\{\tau_i, \alpha\} \subseteq T(e_2)$. Let D be an MCS such that $D \preceq T(e_2)$ and $\{\tau_i, \alpha\} \subseteq D$. Since $T(e_2) \preceq T(e_1)$, we get $D \preceq T(e_1)$. But e_3 kills $(e_1, \diamond_i \alpha)$, so $T(e_3) \preceq D$, and by transitivity, $T(e_2) \preceq D$. Thus $(e_2, \diamond_i \alpha)$ is not a live requirement for T in S . \square

Lemma 6.9. Let A be an MCS and L a finite set of formulas. Then there exists a *finite* frame $S = (E; \leq)$ and an L -historic, strict and coherent chronicle T over S such that for some $e_0 \in E$, $T(e_0) = A$ and such that

(i) for all $e \in E$, $e \leq e_0$.

(ii) for all $e \in E - \{e_0\}$, there exists a formula $\diamond_i \alpha \in L \cap A$ such that $\{\tau_i, \alpha\} \subseteq T(e)$.

Proof. We construct S inductively below. Let \hat{E} be a countably infinite set. To begin the construction pick $e_0 \in \hat{E}$ and let $S^0 \triangleq (\{e_0\}, \{(e_0, e_0)\})$ and define T^0 by setting $T^0(e_0) \triangleq A$. T^0 satisfies conditions (i) and (ii) above, and T^0 is a coherent and strict chronicle over S^0 .

Assume inductively that $S^k = (E^k, \leq^k)$ and T^k are defined, where T^k is a coherent and strict chronicle over S^k satisfying conditions (i) and (ii) above. Let $e \in E^k$ such that $(e, \diamond_i \alpha)$ is a live L -historic requirement for T^k in S^k . By Lemma 6.5, there exists a frame $S' = (E'; \leq')$ such that $E' = E^k \cup \{\hat{e}\}$, for some $\hat{e} \in E' - E^k$, $\leq \subseteq \leq'$ and T' is coherent and strict over S' . To show that S' and T' satisfy conditions (i) and (ii), let $e' \in E'$. If $e' \in E^k$, by induction hypothesis, $e' \leq e_0$. Otherwise, $e' = \hat{e}$. Since $\hat{e} \leq' e$ and $e \leq e_0$, we have $\hat{e} \leq' e_0$. Consider $e' \in E' - \{e_0\}$. If $e' \in E^k$, we again get (ii) by induction hypothesis. Otherwise $e' = \hat{e}$. \hat{e} kills $(e, \diamond_i \alpha)$. Since $\diamond_i \alpha \in T^k(e) \cap L$, $e \leq e_0$ and T^k is coherent, $\diamond_i \alpha \in T^k(e_0) \cap L = T^0(e_0) \cap L$. Thus $\diamond_i \alpha \in A \cap L$ and $\{\tau_i, \alpha\} \subseteq T'(\hat{e})$. Set $S^{k+1} = (E'; \leq')$, $T^{k+1} = T'$.

If there is no such live historic requirement for T^k in S^k , then set $S^{k+1} = S^k$ and $T^{k+1} = T^k$.

Define $S = (E; \leq)$, where $E = \bigcup_k E^k$ and $\leq = \bigcup_k \leq^k$. Define the chronicle T on S by $T(e) \triangleq T^k(e)$, for $e \in E^k$. Clearly, T is a coherent, strict and L -historic chronicle over S , as required.

It only remains to be shown that S is a *finite* frame. To prove this consider a tree structure defined on E as follows:

e_0 is the root of the tree, and

$e_1 \succ e_2$ iff for some k , $(e_1, \diamond_i \alpha)$ was the chosen live L -requirement at stage k and $E^{k+1} = E^k \cup \{e_2\}$. (In this case, we write $e_1 \succ^k e_2$).

The tree structure here refers only to the construction of S using S^k and not to the \leq ordering on the n -ACSA which is in general not a tree.

Since L is a finite set, this tree is finitely branching. Further, it is easy to see that the tree covers E . Hence, to prove finiteness of E , it is sufficient (by König's Lemma) to prove that every path in this tree is finite. For this, we need to show that along every path in the tree, after a finite depth, there are no more live requirements. To this end, define

$L(e, k) \triangleq \{\diamond_i \alpha \in L \cap T^k(e) \mid (e, \diamond_i \alpha) \text{ is a live requirement for } T^k \text{ in } S^k\}$.

Clearly, for all e , for all $k \geq 0$, $L(e, k)$ is finite and bounded by the size of L .

Claim 1. For all e , $L(e, k) \supseteq L(e, k+1)$, $k \geq 0$.

Claim 2. Let $e_1 \succ^k e_2$. $L(e_1, k) \supsetneq L(e_2, k+1)$.

The first claim trivially follows from the observation that $\leq^k \subseteq \leq^{k+1}$. For Claim 2, observe that T^{k+1} is coherent and strict and hence by Proposition 6.8, we have $L(e_1, k+1) \supseteq L(e_2, k+1)$. Further $L(e_1, k) \supseteq L(e_1, k+1)$ from Claim 1 and by construction, $L(e_1, k) \neq L(e_1, k+1)$.

Now let $e_1 \succ e_2 \geq \dots$ be a path in the tree. Then there exist k_1, k_2, \dots such that $e_1 \succ^{k_1} e_2 \succ^{k_2} \dots$, where we have $\dots > k_2 > k_1$. By claims above, $L(e_1, k_1) \supsetneq L(e_2, k_2) \supsetneq \dots$, forming a strictly descending chain, which cannot be infinite, since L is finite. Thus every path in the tree is finite and hence S is finite, and the lemma is proved. \square

Given a historic chronicle, to kill a prophetic requirement, we add as many events as necessary to get a historic chronicle, in such a way that the past of events in the given frame is unaltered; for every new event added, a fixed finite past as dictated by the historic requirements generated by a suitable subset of $\text{CL}(\alpha_0)$ is attached.

Lemma 6.10. Let $S = (E; \leq)$ be a finite frame and T a coherent and L -historic chronicle over S , where L is a finite set of formulas. If $(e_0, \diamond_i \alpha)$ is a live *prophetic* requirement for T in S , then there exists a finite frame $S' = (E'; \leq')$ and an L -historic and coherent chronicle T' over S' such that

- (i) $E' = E \cup E''$, where E'' is finite and disjoint from E .
- (ii) $\leq = \leq' \upharpoonright (E \times E)$.
- (iii) $T = T' \upharpoonright E$.
- (iv) $(e_0, \diamond_i \alpha)$ is *not* a live requirement for T' in S' .
- (v) for all $e \in E$, $\{e' \mid e' \leq e\} = \{e' \mid e' \leq' e\}$.

Proof. Let $e_0 \in E$ and let $(e_0, \diamond_i \alpha)$ be a live prophetic requirement for T in the frame S . Let $A = T(e_0)$. $\tau_i \wedge \diamond_i \alpha \in T(e_0)$, and by Lemma 4.8, there exists an MCS B such that $A \preceq B$ and $\{\tau_i, \alpha\} \subseteq B$.

Let $L' \triangleq (B \cap L) - A$. L' is a finite set and by Lemma 6.9, there exists a finite frame $S_B = (E_B; \leq_B)$ and a coherent and L' -historic chronicle T_B over S_B with an $\hat{e} \in E_B$ such that $T_B(\hat{e}) = B$ and

(a) for all $e \in E_B$, $e \leq \hat{e}$.

(b) for all $e \in E_B - \{\hat{e}\}$, there exists a formula $\diamond_i \alpha \in L'$ such that $\{\tau_i, \alpha\} \subseteq T_B(e)$.

Without loss of generality, assume $E \cap E_B = \emptyset$. Define $S' = (E'; \leq')$ and T' as given below:

$$E' \triangleq E \cup E_B ;$$

$$\leq' \triangleq \leq \cup \leq_B \cup \{(e_1, e_2) \mid e_1 \leq e_0, e_2 \in E_B \text{ and } T(e_1) \leq T_B(e_2)\};$$

$$T'(e) \triangleq T(e), \text{ for } e \in E, \text{ and } T_B(e), \text{ for } e \in E_B .$$

Note that \hat{e} kills the requirement $(e_0, \diamond_i \alpha)$ in S' for T' provided we show that S' is a frame and T' is a coherent and L -historic chronicle over S' .

Further, from the definition of \leq' , it is clear that we cannot have $e_1 \leq' e_2$, where $e_1 \in E_B$ and $e_2 \in E$. Hence for all $e \in E$, $e' \leq e$ iff $e' \leq' e$. Also $\leq = \leq' \upharpoonright (E \times E)$.

Clearly, from the coherence of T and T_B and the definition of \leq' , we have $e_1 \leq' e_2 \Rightarrow T'(e_1) \preceq T'(e_2)$. Though we have not yet established that S' is a frame, we refer to this property as coherence of T' .

Further, since T and T' are chronicles and $E \cap E_B = \emptyset$, we have $e \in E'_j$ iff $\tau_j \in T'(e)$, for all $e \in E'$.

To check reflexivity of \leq' , consider $e \in E'$. Then either $e \in E$, or $e \in E_B$. In either case, we have $e \leq e$ or $e \leq_B e$ and hence $e \leq' e$.

To see that \leq' is antisymmetric, suppose that $e'_1 \leq e_2$ and $e_2 \leq' e_1$. Since $\leq = \leq' \upharpoonright (E \times E)$, if $\{e_1, e_2\} \subseteq E$, $e_1 = e_2$, since \leq is antisymmetric. If $e_1, e_2 \notin E$, then $\{e_1, e_2\} \subseteq E_B$ and hence $e_1 = e_2$ since \leq_B is antisymmetric. So consider the case when $e_1 \notin E$ and $e_2 \in E$. But then $e_1 \leq' e_2$ is impossible, by definition of \leq' . Similarly, the case when $e_2 \notin E$ and $e_1 \in E$ is also impossible.

To check transitivity of \leq' , consider $e_1 \leq' e_2 \leq' e_3$. If $\{e_1, e_2, e_3\} \subseteq E$, $e_1 \leq' e_3$, by transitivity of \leq , since $\leq = \leq' \upharpoonright (E \times E)$. Hence at least one of e_1, e_2, e_3 is in E_B . If $e_1 \in E_B$, then by definition of \leq' , $\{e_1, e_2, e_3\} \subseteq E_B$, and since $\leq_B = \leq' \upharpoonright (E_B \times E_B)$, $e_1 \leq' e_3$, by transitivity of \leq_B . If $e_2 \in E_B$, we have $e_1 \leq' e_2 \leq_B e_3$, hence $e_1 \leq e_0$ and by coherence of T and T_B , $T(e_1) \preceq T_B(e_2) \preceq T_B(e_3)$. By transitivity of \preceq , $T(e_1) \preceq T_B(e_3)$, and hence by definition of \leq' , $e_1 \leq' e_3$. If $e_3 \in E_B$, we have $e_1 \leq e_2 \leq' e_3$, and hence by coherence of T and T_B , $T(e_1) \preceq T(e_2) \preceq T_B(e_3)$. Again by transitivity of \preceq , we have $T(e_1) \preceq T_B(e_3)$, and by definition $e_1 \leq' e_3$, as required.

Note that if $e_1 \in E$, $e_1 \leq e_0$, $e_2 \neq \hat{e}$ and $e_2 \in E_B$, we cannot have $T_B(e_2) \preceq T(e_1)$. This is because, by condition (b), there exists $\diamond_j \beta \in L'$ such that $\{\tau_j, \beta\} \subseteq T_B(e_2)$. If $T_B(e_2) \preceq T(e_1)$, we get $\diamond_j \beta \in T(e_1)$. If $\tau_k \in T(e_1)$, since $T(e_1) \preceq T(e_0)$, we get $\diamond_k \diamond_j \beta \in T(e_0)$ and hence by transitivity, we get $\diamond_j \beta \in T(e_0) = A$, contradicting the fact that $\diamond_j \beta \in L' = (B \cap L) - A$.

Now, to check backward linearity of \leq' within agents, assume events e_1 and e_2 in E'_j for some j and $e_3 \in E'$ such that $e_1 \leq' e_3$ and $e_2 \leq' e_3$. If $e_3 \in E$, then $\{e_1, e_2\} \subseteq E$, and by backward linearity of \leq , either $e_1 \leq' e_2$ or $e_2 \leq' e_1$. Hence let $e_3 \in E_B$. If both e_1 and e_2 are in E_B , we are done since \leq_B is backward linear. If $\{e_1, e_2\} \subseteq E$, by definition of \leq' , $e_1 \leq e_0$ and $e_2 \leq e_0$; again, by backward linearity of \leq , either $e_1 \leq e_2$ or $e_2 \leq e_1$. The only remaining case is when one of e_1 and e_2 is in E and the other is in E_B . Let $e_1 \in E$ and $\{e_2, e_3\} \subseteq E_B$. If $e_2 = \hat{e}$, then we must have $e_3 = \hat{e}$ as well, so $e_1 \leq e_2$. Therefore, let $e_2 \neq \hat{e}$. We have $T(e_1) \preceq T(e_0) \preceq T_B(\hat{e})$ and $T_B(e_2) \preceq T_B(\hat{e})$. Hence by Proposition 4.7, either $T(e_1) \preceq T_B(e_2)$ or $T_B(e_2) \preceq T(e_1)$. By the remark in the previous paragraph, the latter case is impossible. But the former case ensures that $e_1 \leq' e_2$, as required.

Thus, $S' = (E'; \leq')$ is a finite frame and T' is a coherent chronicle over it. It only remains to show that T' is L -historic.

Let $e \in E'$. If $e \in E$ and $\diamond_j \beta \in T(e) \cap L$, since T is L -historic, we can find required $e' \in E$ such that $e' \leq e$. Now let $e \in E_B$ and let $\diamond_j \beta \in T_B(e) \cap L$. Since $T_B(e) \preceq T_B(\hat{e})$, we have $\diamond_j \beta \in B \cap L$. If $\diamond_j \beta \notin A$, then $\diamond_j \beta \in L'$ and since T_B is L' -historic, we get required e' in E_B .

Otherwise, let $\diamond_j \beta \in A$. Since T is L -historic, there exists $e_1 \leq e_0$ such that e_1 kills the requirement $(e_0, \diamond_j \beta)$ for T in S . Now $\diamond_j \beta \in T_B(e)$, and hence there exists an MCS D such that $D \preceq T_B(e)$ and $\{\tau_j, \beta\} \subseteq D$. We have $T(e_1) \preceq T_B(\hat{e})$ as well as $D \preceq T_B(e)$. Hence either $D \preceq T(e_1)$ or $T(e_1) \preceq D$. In the former case, $D \preceq T(e_0)$ as well; but e_1 kills $(e_0, \diamond_j \beta)$, so we have $T(e_1) \preceq D$. Thus in either case we have $T(e_1) \preceq D$. By transitivity, $T(e_1) \preceq T_B(e)$ and by construction, $e_1 \leq' e$. Further for arbitrary $D \preceq T_B(e)$ such that $\{\tau_i, \alpha\} \subseteq D$, we have already shown that $T(e_1) \preceq D$. Thus e_1 kills $(e, \diamond_j \beta)$ as well. Hence, T' is L -historic.

This completes the proof of the lemma. \square

Theorem 6.11. (Completeness)

If $\models \alpha$ then $\vdash \alpha$.

Proof. We show that every consistent formula is satisfiable. Let $\hat{E} = \{e_0, e_1, e_2, \dots\}$ be a countably infinite set. Fix an enumeration of $\hat{E} \times \Phi_n$, where Φ_n is the set of all formulas in the language.

Now, let α be a consistent formula. Fix $L = \text{CL}(\alpha)$. Pick an MCS A containing α . We now define, for all $k \geq 0$, S^k and T^k . By Lemma 6.9, there exists a finite frame S^0 and a coherent and L -historic chronicle T^0 over it.

Inductively assume that $S^k = (E^k; \leq^k)$ and T^k have been defined, where $E^k = \{e_0, e_1, \dots, e_k\}$ and T^k is L -historic and coherent over S^k . Suppose there are no live requirements for T^k in S^k . Then set $S^{k+1} = S^k$ and $T^{k+1} = T^k$. Otherwise, among all the live requirements for T^k in S^k , choose the least one in the enumeration of $\hat{E} \times \Phi_n$, say (e, β) . β is a live prophetic requirement. By Lemma 6.10, we can extend S^k and T^k to a frame $S^{k+1} = (E^{k+1}; \leq^{k+1})$ and a chronicle T^{k+1} over it such that

- (i) $E^{k+1} = E^k \cup \{e_{k+1}, \dots, e_{k+m}\}$, for some $m > 0$,

- (ii) $T^{k+1} \upharpoonright E^k = T^k$,
- (iii) T^{k+1} is L -historic and coherent over S^{k+1} and
- (iv) (e, β) is not a live requirement for T^{k+1} in S^{k+1} .

Finally set $S = (E; \leq)$, where $E = \bigcup_k E^k$, $\leq = \bigcup_k \leq^k$, and define a chronicle T over S by:

$$\text{for } e \in E, T(e) \triangleq T^k(e), \text{ where } e \in E^k .$$

It can be easily checked that T is an L -perfect chronicle over S . Further, for any $e \in E^k$, $\{e' \mid e' \leq^k e\}$ is finite and equals the set $\{e' \mid e' \leq^{k+1} e\}$. Thus every event has a finite past in S . Now, by Lemma 6.4, $M_T^L, e_0 \models \alpha$, where $M_T^L = (S, V_T^L)$, V_T^L being the valuation induced by T and L . Thus, α is indeed satisfiable. \square

We have so far considered only finitary n -ACSAs. Finitary ACSAs can be axiomatized in a similar manner. The logical language is then $\Phi_{\mathbb{N}}$ as defined in the last section. Define

$$\mathcal{A}(\text{fin-ACSA}) \triangleq \mathcal{A}(\text{ACSA}) + (\text{A9}) .$$

(Note that (A9.b) now stands for infinitely many axioms, one for each pair of distinct i and j in \mathbb{N} .)

Theorem 6.12. $\mathcal{A}(\text{fin-ACSA})$ is sound and complete for the class of finitary ACSAs.

Proof. It is easy to check that the axiom system is sound for finitary ACSAs. For completeness, let α_0 be a consistent formula. Using (TE) we can find a τ_i , for some $i \in \mathbb{N}$, such that $\alpha_0 \wedge \tau_i$ is consistent. Let $L = \text{CL}(\alpha_0 \wedge \tau_i)$. Define n to be the maximum of $\{j \mid \tau_j \in L\}$. Clearly, $L \subseteq \Phi_n$. The proof that α_0 is satisfiable in a finitary n -ACSA proceeds exactly as given earlier in this section, except that we now restrict ourselves to good MCSs; recall that a good MCS is one which has at least one τ_j . Axiom (A8) ensures that a good MCS has exactly one type proposition. We can check that all the results 6.3 through 6.10 hold when we confine our attention to good MCSs. \square

As before, let $\text{SAT}(\text{fin-}n\text{-ACSA})$ denote the set of formulas satisfiable in models based on finitary n -ACSAs and $\text{SAT}(\text{fin-ACSA})$ denote the set of those satisfiable in models based on finitary ACSAs. Once again, we have the equation between satisfiable formulas:

$$\text{Corollary 6.13. } \text{SAT}(\text{fin-ACSA}) = \bigcup_n \text{SAT}(\text{fin-}n\text{-ACSA}). \quad \square$$

7. Shared Events

We have so far maintained the condition that each event belongs to exactly one agent. We would like to relax this requirement and allow events to be *shared* by agents.

The motivation for this comes about when we try to represent joint actions by agents. A typical example of this is a “handshake” in systems with synchronous

communication [9,10], i.e. where all communicating agents have to wait until they can synchronize and perform a joint action. In our producer-consumer example of Sec. 1 (Fig. 3), for instance, we might require the producer and the consumer to interact synchronously with a buffer. We would represent this by allowing events shared by the producer and buffer and by the consumer and buffer respectively. A system may allow both synchronous and asynchronous communication; for instance, the consumer may asynchronously request items from the producer, but will have to wait for them to be delivered via the buffer.

Definition 7.1. A system of n Communicating Sequential Agents (n -CSA) is a tuple $(E_1, E_2, \dots, E_n; \leq)$ such that:

- (i) $(E; \leq)$ is a poset, where $E = \bigcup_j E_j$, and
- (ii) for all $e \in E$, for $1 \leq i \leq n$,

$$\downarrow e \cap E_i \text{ is totally ordered by } \leq . \quad \square$$

Note that we have just dropped the requirement that E_1, E_2, \dots, E_n be disjoint from Definition 1.2.

As before, we shall let S, S', \dots range over n -CSAs and often write $S = (E; \leq)$ assuming an implicit definition of the agents E_1, \dots, E_n such that $E = \bigcup_i E_i$.

Note that every n -ACSA is also an n -CSA; thus CSAs can model both synchronous and asynchronous communication.

The axiom system $\mathcal{A}(n\text{-CSA})$ is defined by retaining the axioms and inference rules of $\mathcal{A}(n\text{-ACSA})$, *except* that axiom (A7.a) is now relaxed to (A7'.a) as shown:

- (A7') (a) $\tau_1 \vee \dots \vee \tau_n$ (n -agents)
 (b) $\boxplus_i \tau_i$
 (c) $\tau_i \supset \boxminus_i \tau_i$

Consistency and theses are now defined with respect to the new axiom system. Note that in the new axiom system, formulas like $\tau_1 \wedge \tau_2$ are consistent; this formula specifies an event shared by agents 1 and 2. Similarly, $\boxplus_i (\tau_j \wedge \alpha)$ (where $j \neq i$) specifies a synchronization event in the past where α held.

The soundness of axiom (A7'.a) is trivial. The completeness proof closely follows that in Sec. 4. In the rest of this section, we point out some of the finer points to be considered.

Firstly, we relax Proposition 4.3 to assert that an MCS A can have more than one type proposition:

Proposition 7.2. Let A be an MCS. For some $i \in \{1, \dots, n\}$, $\tau_i \in A$.

Proof. Follows from axiom (A7'.a). □

The definition of the semantic ordering is rewritten as:

Definition 7.3. Let A and B be MCSs. Then

$$A \leq' B \triangleq \{\boxplus_i \alpha \mid \tau_i \wedge \alpha \in A\} \subseteq B . \quad \square$$

Since more than one type proposition may belong to A , this seems to be a stronger ordering. The following observation shows that it is not.

Proposition 7.4. Let A, B be MCSs and let $\tau_i, \tau_j \in A$. Then

$$\{\diamond_i \alpha \mid \alpha \in A\} \subseteq B \text{ iff } \{\diamond_j \alpha \mid \alpha \in A\} \subseteq B .$$

Proof. Assume that the left-hand side is true and that $\beta \in A$. Since $\tau_j \in A$, by Axiom (A2.a), $\diamond_j \beta \in A$. By assumption, $\diamond_i \diamond_j \beta \in B$. By Axiom (A3), $\diamond_j \beta \in B$. The other case is symmetric. \square

Clearly, \preceq' remains reflexive, transitive and backwards connected within an agent. The next results to be proved are the new versions of Lemma 4.8 and 4.9.

Lemma 7.5. Let A be an MCS and let $\diamond_i \alpha \in A$. Then there exists an MCS B such that $B \preceq' A$ and $\{\tau_i, \alpha\} \subseteq B$.

Proof. The proof of Lemma 4.8 showed that, given A as above, there is an MCS B where $\{\tau_i, \alpha\} \subseteq B$ and $\{\diamond_i \beta \mid \beta \in B\} \subseteq A$. But by Proposition 7.4, this is enough to show that $B \preceq' A$. \square

In an analogous manner, we can prove:

Lemma 7.6. Let A be an MCS and let $\tau_i \wedge \diamond_i \alpha \in A$. Then there exists an MCS B such that $A \preceq' B$ and $\{\tau_i, \alpha\} \subseteq B$. \square

The definitions relating to chronicles and live requirements and the chronicle construction remain unchanged. A model can be extracted from a perfect chronicle as before. The earlier completeness proof can therefore be used. Consider, for example, Lemma 4.15, which “kills” a live historic requirement. The required MCS is obtained by invoking Lemma 4.8, in place of which we now appeal to Lemma 7.5. For the construction, always add a new event to all the agents it should belong to. Thus the initial frame S^0 is defined by letting e_0 belong to all agents j such that $\tau_j \in T(e_0)$. Similarly, the structure E' is defined in terms of E as:

$$E'_j \triangleq \begin{cases} E_j \cup \{\hat{e}\} & \text{if } j \in T(\hat{e}), \text{ and} \\ E_j & \text{otherwise .} \end{cases}$$

This ensures that the requirements of a chronicle are satisfied. Proceeding thus, we can prove:

Theorem 7.7. The axiom system $\mathcal{A}(n\text{-CSA})$ is sound and complete for the class of n -CSAs. \square

Finitary frames

We now consider *finitary* n -CSAs, i.e. those satisfying the condition that every event has a finite past.

The axiom system $\mathcal{A}(\text{fin-}n\text{-CSA})$ is given by adding the following scheme to $\mathcal{A}(n\text{-CSA})$:

- (A9') (a) $\diamond_i \alpha \supset \diamond_i (\alpha \wedge \Box_i (\sim \alpha \supset \Box_i \sim \alpha))$ (Well-founded agents)
 (b) $\diamond_i \alpha \supset \diamond_i (\alpha \wedge \Box_j (\sim \tau_i \supset \Box_i \sim \alpha))$ (Well-founded communication)

Earlier, we used Axiom (A9) to ensure finitariness:

- (A9) (a) $\diamond_i \alpha \supset \diamond_i (\alpha \wedge \Box_i (\sim \alpha \supset \Box_i \sim \alpha))$
 (b) $\diamond_i \alpha \supset \diamond_i (\alpha \wedge \Box_j \Box_i \sim \alpha)$ ($j \neq i$)

(A9.b) is unsound when shared events are allowed. For instance, if the event e is shared by agents 1 and 2, it must satisfy the formula $\diamond_1 \tau_1 \wedge \Box_1 (\tau_1 \supset \diamond_2 \diamond_1 \tau_1)$, contradicting the validity of (A9.b). As we explained in Sec. 6, (A9.b) ruled out infinitely descending chains of communication events. (A9'.b) rules out infinite descending chains of *asynchronous* communication events. (A9.a) itself is strong enough to rule out an infinite chain of shared events in which agent i participates, so it is retained unchanged.

To show soundness of (A9'.b), suppose not. Then we have a model $M = (S, V)$, where $S = (E; \leq)$ is an n -CSA. Then there exists $e_1 \in E$ such that $M, e_1 \models \diamond_i \alpha \wedge \Box_i (\alpha \supset \diamond_j (\sim \tau_i \wedge \diamond_i \alpha))$. Hence there is an $e_2 \in E_i$ such that $e_2 \leq e_1$ and $M, e_2 \models \alpha \wedge \diamond_j (\sim \tau_i \wedge \diamond_i \alpha)$. So there is an $e_3 \in E_j$ such that $e_3 \leq e_2$ and $M, e_3 \models \sim \tau_i \wedge \diamond_i \alpha$. Since $e_2 \in E_i$ and $e_3 \notin E_i$, $e_2 \neq e_3$. Next we find an $e_4 \in E_i$ such that $M, e_4 \models \alpha \wedge \diamond_j (\sim \tau_i \wedge \diamond_i \alpha)$. Note that $e_3 \neq e_4$, and hence the argument from e_1 can be repeated again to yield an infinite descending chain of events, contradicting the fact that e_1 had a finite past.

The completeness proof follows the one given in Sec. 5 in the same manner as the completeness proof for n -CSAs given earlier followed that of Sec. 4. Firstly the required thesis now is:

$$(T7') \quad \diamond_i \alpha \supset \diamond_i \left(\alpha \wedge \bigwedge_j \Box_j ((\tau_i \supset \sim \alpha) \supset \Box_i \sim \alpha) \right).$$

The derivation of (T7') is easy and follows the same lines as that of (T7). Given Propositions 7.4 to 7.6, we can verify that the results 6.3 through 6.5 hold when we use the semantic ordering \preceq' on MCSs. Now let $\alpha_i \triangleq \alpha \wedge \bigwedge_j \Box_j ((\tau_i \supset \sim \alpha) \supset \Box_i \sim \alpha)$. With this notation, Propositions 6.6 and 6.7 go through easily for MCSs ordered by \preceq' . We now prove the analog of Proposition 6.8.

Proposition 7.8. Let $S = (E; \leq)$ be a finitary n -CSA and T a coherent and strict chronicle over S . Let $e_1, e_2 \in E$ such that $e_1 \leq e_2$ and $(e_1, \diamond_i \alpha)$ is not a live requirement for T in S . Then $(e_1, \diamond_i \alpha)$ is also not a live requirement for T in S .

Proof. Assume the hypothesis and let $e_2 \in E$ such that e_3 kills $(e_i, \diamond_i \alpha)$. If $\diamond_i \alpha \notin T(e_2)$, we are done. Otherwise, there exists an MCS C such that $C \preceq' T(e_2)$, $\{\tau_i, \alpha\} \subseteq C$. By coherence of T , we have $T(e_2) \preceq' T(e_1)$ and hence $C \preceq' T(e_1)$. But e_3 kills $(e_1, \diamond_i \alpha)$; therefore, $T(e_3) \preceq' C$. By transitivity, we get $T(e_3) \preceq' T(e_2)$.

By strictness of T , either $e_3 \leq e_2$ or $e_2 \leq e_3$. In the former case, it is easy to show that e_3 kills $(e_2, \diamond_i \alpha)$, since if an MCS C has α and $C \preceq' T(e_2)$, then $C \preceq' T(e_1)$ as well, so $T(e_3) \preceq' C$. Assume the latter case. Since e_3 kills $(e_1, \diamond_i \alpha)$, by the analog of Proposition 6.7, $\alpha_i \in e_3$. Now $e_2 \leq e_3$, so $T(e_2) \preceq' T(e_3)$. Hence

$M, e_2 \vDash (\tau_i \supset \sim \alpha) \supset \Box_i \sim \alpha$. If $e_2 \notin E_i$, $M, e_2 \vDash \sim \tau_i$ and hence $M, e_2 \vDash \Box_i \sim \alpha$, contradicting the fact that $\Diamond_i \alpha \in T(e_2)$. Hence $e_2 \in E_i$. Again, if $M, e_2 \vDash \sim \alpha$, then we get a contradiction, and thus $M, e_2 \vDash \alpha$. That is, $\alpha \in T(e_2)$. Now, if D is an MCS such that $D \preceq' T(e_2)$ and $\{\tau_i, \alpha\} \subseteq D$, since $T(e_2) \preceq' T(e_1)$, we get $D \preceq T(e_1)$. But e_3 kills $(e_1, \Diamond_i \alpha)$, so $T(e_3) \preceq' D$, and by transitivity, $T(e_2) \preceq' D$. Thus e_2 kills $(e_2, \Diamond_i \alpha)$. \square

Note that the completeness proof in Sec. 6 uses (A9) and (T7) only in proving results 6.6 to 6.8. Since these results follow for n -CSAs, we can easily verify that the results 6.9 through 6.11 hold in the new set-up. The definitions of $\text{CL}(\alpha)$ and requirements are the same and the proof construction yields a finitary n -CSA.

Note that (A9'.b) is sound for the frames in Sec. 6 and the above proof could be used to provide completeness there as well, but we preferred to use a simpler axiom for n -ACSAs.

CSAs

We can also generalize n -CSAs to allow for systems with unboundedly many agents. This is done in exactly the same way as in Sec. 5.

Definition 7.9. A system of **Communicating Sequential Agents (CSA)** is a triple (E, \leq, η) , where

- (i) (E, \leq) is a poset,
- (ii) $\eta: E \rightarrow 2^{\mathbb{N}}$ is a (**naming**) function such that: for all $e \in E$, for all $j \in \mathbb{N}$,
 - (a) $\eta(e)$ is a nonempty finite subset of \mathbb{N} ,
 - (b) $\downarrow e \cap E_j$ is totally ordered by \leq for every $j \in \mathbb{N}$, where $E_j \triangleq \{e \in E \mid j \in \eta(e)\}$. \square

The axiom system $\mathcal{A}(\text{CSA})$ is formed by dropping (A8) from $\mathcal{A}(\text{ACSA})$. It is easy to check that the new axiom system is sound over CSAs. Completeness necessitates some minor changes.

The definition of a chronicle has to be weakened.

Definition 7.10. An η -chronicle on a CSA $S = (E, \leq, \eta)$ is a function T which assigns an MCS to each $e \in E$ such that

$$\text{for every } e \in \bigcup_{e \in E} \eta(e), \tau_i \in T(e) \text{ iff } i \in \eta(e). \quad \square$$

The definitions of coherent, L -historic, L -perfect η -chronicles and the corresponding L -requirements are as in Sec. 5. An η -chronicle induces a model as before.

Definition 7.11. Let $S = (E, \leq, \eta)$ be a frame, T an η -chronicle on it and L a set of formulas. The valuation induced by T for L , denoted V_T^L , is given by:

$$\text{for } e \in E, V_T^L(e) \triangleq T(e) \cap (P \cup T) \cap L.$$

We use M_T^L to denote (S, V_T^L) . \square

Lemma 7.12. Let T be a $\text{CL}(\alpha)$ -perfect η -chronicle over the frame (E, \leq, η) . Then for every $e \in E$ and every $\beta \in \text{CL}(\alpha)$, $\beta \in T(e)$ iff $M_T^{\text{CL}(\alpha)}, e \models \beta$. \square

For construction of a relativized perfect chronicle, we need to confine our attention to good MCSs as in Sec. 5. Observe that results 7.4 to 7.6 go through for good MCSs. While killing a live requirement (L -historic or L -prophetic, with L a finite set of formulas), the new event \hat{e} is added to finitely many agents as follows:

$$\eta'(\hat{e}) = \{j \mid \tau_j \in T'(\hat{e}) \cap L\} .$$

This ensures that T' is an η' -chronicle and the lemmas for killing requirements can be proved.

Now, given a consistent formula α_0 , use the rule (TE) to find τ_i such that $\alpha_0 \wedge \tau_i$ is consistent. Let $L = \text{CL}(\alpha_0 \wedge \tau_i)$ and A be an MCS containing $\alpha_0 \wedge \tau_i$. A is good. In the initial step of the construction, we set $E^0 = \{e_0\}$, $\leq^0 = \{(e_0, e_0)\}$, $T^0(e_0) = A$ and $\eta^0(e_0) = \{i \mid \tau_i \in A \cap L\}$. T^0 is an η^0 -chronicle over the frame (E^0, \leq^0, η^0) . By the remarks above, we can proceed as before and construct inductively an L -perfect η -chronicle T over a frame S . Then using Lemma 7.12 we get:

Theorem 7.13. $\mathcal{A}(\text{CSA})$ is sound and complete for the class of CSAs. \square

Let $\text{SAT}(\text{CSA})$ denote the set of formulas in $\Phi_{\mathbb{N}}$ satisfiable over CSAs and let $\text{SAT}(n\text{-CSA})$ denote the set of formulas in Φ_n satisfiable over n -CSA, for $n \in \mathbb{N}$. Note that the proof above builds an n -CSA where n is the maximum of $\{j \mid \tau_j \in L\}$. Hence we again have an equation on satisfiable formulas.

Corollary 7.14.

$$\text{SAT}(\text{CSA}) = \bigcup_n \text{SAT}(n\text{-CSA}) . \quad \square$$

For the subclass of finitary CSAs, we define the axiom system $\mathcal{A}(\text{fin-CSA}) \triangleq \mathcal{A}(\text{CSA}) + (A9')$. Soundness of the axiom system is routine and the strategy for completeness consists in constructing a finitary CSA as a model for a consistent formula. This can be done exactly as we outlined earlier in this section, but with η -chronicles. All those results hold for good MCSs.

We conclude the section with the relevant equation on satisfiable formulas, where we use the obvious notation.

Corollary 7.15.

$$\text{SAT}(\text{fin-CSA}) = \bigcup_n \text{SAT}(\text{fin-}n\text{-CSA}) . \quad \square$$

8. Discussion

We have introduced in this paper several classes of distributed systems, which can be compactly represented by Fig. 4, where the arrows represent inclusions.

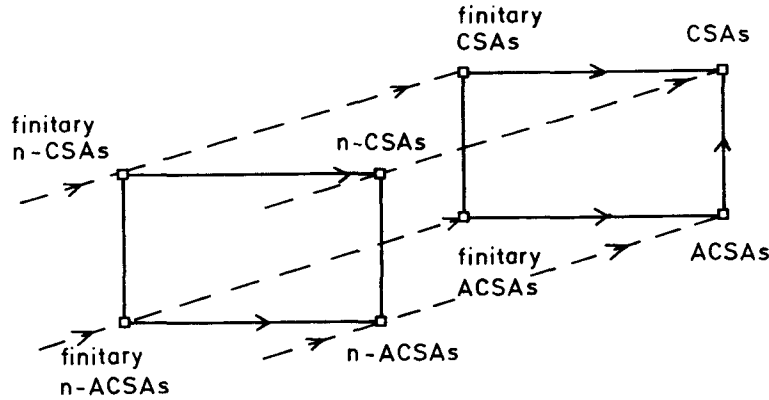


Fig. 4. Hierarchy of CSAs.

The figure above represents a (strict) hierarchy of CSAs. The hierarchy actually extends through 1-CSAs, 2-CSAs, ... to CSAs, and similarly for the other classes. Let \mathcal{C} range over the classes shown above.

The logical languages we introduced were Φ_n for the bounded (n -)agent systems, and their union $\Phi_{\mathbb{N}}$ for the unbounded agent systems. For each class of CSAs \mathcal{C} we provided sound and complete axiomatizations $\mathcal{A}(\mathcal{C})$ in the appropriate Φ language. For \mathcal{C} ranging over CSAs, ACSAs, finitary CSAs and finitary ACSAs, we showed that $\text{SAT}(\mathcal{C})$, the satisfiable formulas in the unbounded class, are exactly the union of $\text{SAT}(n\text{-}\mathcal{C})$, the satisfiable formulas in the n -bounded class.

While we have used indexed modalities here, Parikh [24] has pointed out that one can eliminate them by using the usual tense modalities (\diamond and \heartsuit) along with the type propositions. On the other hand we can also consider indexed tense logic *without* type propositions. In Part II of this paper, we consider these and other alternatives to the logic studied here. There we present languages of varying expressive power in all of which the class of CSAs can be axiomatized.

Other logical languages are also of interest. As far as tense logic goes, indexed versions of next-state, now, until and other operators as well as the path operators of branching-time logic [5] can be considered, but it is not clear how they should be interpreted. More interestingly, explicit operators can be used to express conflict and concurrency [22,25,26]. The full class of event structures has been axiomatized in this fashion [22].

One major technical question which remains is that of decidability. We conjecture that all the logics we have considered are decidable.

As mentioned in the Introduction, temporal logics are extensively used to specify properties of distributed systems and model checkers have been built for verification purposes. We intend to explore the convenience of specifying protocols in indexed tense logics like the one studied in this paper and study the complexity of model checking such logics.

Many subclasses of CSAs are of interest. *Sequential* systems, where no two events are concurrent, and systems of *deterministic* agents can be axiomatized by simple extensions of the logics given here. A more interesting subclass is one where the agents are allowed to be nondeterministic, but communication follows a deterministic behaviour. Another restriction which would be expected of systems is that of *finite branching*.

Partial order computational models for which modal logics have been used include nets [27] and event structures [22,25]. We show in the Appendix that CSAs can be obtained by enriching event structures with the notion of agents. General partial orders have been considered by Pinter and Wolper [12] and Katz and Peled [13].

It is straightforward, though tedious, to take a concurrent programming language such as TCSP [9] or CCS [10] and give semantics to its programs in terms of CSAs. One can then use our logics to reason about the behaviour of these programs. However, it would be far more satisfactory to have a proof system which reasons compositionally at the level of the program syntax. That seems to be somewhat elusive at present.

Acknowledgement

We thank Madhavan Mukund and Rohit Parikh for discussions and comments.

References

- [1] S. S. Owicki and L. Lamport, "Proving liveness properties of concurrent programs", *ACM Trans. Prog. Lang. Syst.* **4**, 3 (1982) 455–495.
- [2] A. Pnueli, "Applications of temporal logic to the specification and verification of reactive systems: a survey of current trends", *Lecture Notes in Computer Science* **224** (Springer-Verlag, 1986) pp. 510–584.
- [3] E. M. Clarke, E. A. Emerson and A. P. Sistla, "Automatic verification of finite-state concurrent programs using temporal logic specifications", *ACM Trans. Prog. Lang. Syst.* **8**, 2 (1986) 244–263.
- [4] E. A. Emerson and C.-L. Lei, "Modalities for model checking: branching time logic strikes back", *Sci. Comput. Program.* **8** (1987) 275–306.
- [5] E. A. Emerson and J. Y. Halpern, "Decision procedures and expressiveness in the temporal logic of branching time", *J. Comput. System Sci.* **30**, 1 (1986) 1–24.
- [6] M. Nielsen, G. Plotkin and G. Winskel, "Petri nets, event structures and domains, Part I", *Theoret. Comput. Sci.* **13**, 1 (1980) 86–108.
- [7] G. Winskel, "Event structures", *Lecture Notes in Computer Science* **255** (Springer-Verlag, 1987), pp. 325–392.
- [8] K. M. Chandy and L. Lamport, "Distributed snapshots: determining global states of distributed systems", *ACM Trans. Comput. Syst.* **3**, 1 (1985) 63–75.
- [9] C. A. R. Hoare, *Communicating Sequential Processes* (Prentice-Hall, 1984).
- [10] R. Milner, *Communication and Concurrency* (Prentice-Hall, 1989).
- [11] K. Lodaya and P. S. Thiagarajan, "A modal logic for a subclass of event structures", *Lecture Notes in Computer Science* **267** (Springer-Verlag, 1987) pp. 290–303; "A correction to 'A modal logic for a subclass of event structures'", Report DAIMI-PB-275, Computer Science Department, Aarhus University (Aarhus, 1989).

- [12] S. Pinter and P. Wolper, “A temporal logic for reasoning about partially ordered computations”, *Proc. 3rd ACM Conf. Principles of Distributed Computing*, Vancouver, 1984, pp. 28–37.
- [13] S. Katz and D. Peled, “An efficient verification method for parallel and distributed programs”, *Lecture Notes in Computer Science* **354** (Springer-Verlag, 1989) pp. 489–507.
- [14] J. Y. Halpern and Y. Moses, “Knowledge and common knowledge in a distributed environment”, *Proc. 3rd ACM Conf. Principles of Distributed Computing*, Vancouver, 1984, pp. 50–61.
- [15] R. Parikh and R. Ramanujam, “Distributed processes and the logic of knowledge”, *Lecture Notes in Computer Science* **193** (Springer-Verlag, 1985) pp. 223–229.
- [16] J. Y. Halpern and R. Fagin, “Modelling knowledge and action in distributed systems”, *Distrib. Comput.* **3**, 4 (1989) 159–177.
- [17] J. Reif and A. P. Sistla, “A multiprocess network logic with temporal and spatial modalities”, *J. Comput. Syst. Sci.* **30**, 1 (1985).
- [18] K. M. Chandy and J. Misra, “How processes learn”, *Distrib. Comput.* **1**, 2 (1986) 40–52.
- [19] J. P. Burgess, “Basic tense logic”, in *Handbook of Philosophical Logic*, Vol. II, eds. D. Gabbay and F. Guentner (D. Reidel, 1984) pp. 89–133.
- [20] G. E. Hughes and M. J. Cresswell, *An Introduction to Modal Logic* (Methuen, reprinted in 1982).
- [21] J. P. Burgess, “Decidability for branching time”, *Studia Logica* **XXXIX**, 2/3 (1980) 203–218.
- [22] M. Mukund and P. S. Thiagarajan, “An axiomatization of event structures”, *Lecture Notes in Computer Science* **405** (Springer-Verlag, 1989) pp. 143–160.
- [23] R. Bull and K. Segerberg, “Basic modal logic”, in *Handbook of Philosophical Logic*, Vol. II, eds. D. Gabbay and F. Guentner (D. Reidel, 1984) pp. 1–88.
- [24] R. Parikh, Personal communication (1987).
- [25] W. Penczek, “A temporal logic for event structures”, *Fundamenta Informaticae* **XI** (1988) 297–326.
- [26] S. Christensen, “A logical characterization of linear n -agent event structures”, Master’s thesis, Computer Science Department, Aarhus University (1990).
- [27] W. Reisig, “Towards a temporal logic for causality and choice in distributed systems”, *Lecture Notes in Computer Science* **354** (Springer-Verlag, 1989) pp. 603–627.

Appendix

We show that n -ACSAs can be mapped to the class of n -agent event structures introduced in [11] and vice versa. We also show that CSAs can be mapped to event structures [7].

Definition A.1. [7]

An *event structure* is a triple $(E, \leq, \#)$ where (E, \leq) is a poset and $\#$ is an irreflexive symmetric relation such that

- (I) $\forall e_1, e_2, e_3 \in E. e_1 \# e_2$ and $e_2 \leq e_3$ implies $e_1 \# e_3$. □

Definition A.2. [11]

An event structure $(E, \leq, \#)$ is *sequential* if

- (S) $\forall e_1, e_2 \in E. e_1 \leq e_2$ or $e_2 \leq e_1$ or $e_1 \# e_2$. □

Definition A.3. [11]

Let $n \in \mathbb{N}$, $E = E_1 \cup \dots \cup E_n$. $(E_1, \dots, E_n; \leq; \#)$ is an n -agent event structure if $(E, \leq, \#)$ is an event structure and

- (i) $E_i \cap E_j = \emptyset$, for $i \neq j$.
- (ii) $(E_i, \leq_i, \#_i)$ is a sequential event structure for every i , where $\leq_i = \leq \upharpoonright (E_i \times E_i)$ and $\#_i = \# \upharpoonright (E_i \times E_i)$.
- (iii) $\# = \{(e_1, e_2) \mid \exists i. \exists (e'_1, e'_2) \in \#_i. e'_1 \leq e_1, e'_2 \leq e_2\}$. □

Proposition A.4.

Let $(E_1, \dots, E_n; \leq; \#)$ be an n -agent event structure. Then $(E_1, \dots, E_n; \leq)$ is an n -ACSA.

Proof.

We only need to verify the backwards linearity condition. So let $e_1, e_2 \in E_i$, $e \in E$, $e_1 \leq e$ and $e_2 \leq e$. By (S), $e_1 \leq e_2$ or $e_2 \leq e_1$ or $e_1 \# e_2$. If $e_1 \# e_2$, then by the (I) condition on the underlying event structure, $e \# e$, which contradicts the irreflexivity of $\#$. Hence $e_1 \leq e_2$ or $e_2 \leq e_1$, as desired. □

Proposition A.5.

Let $(E_1, \dots, E_n; \leq)$ be an n -ACSA.

Define $\#_i \triangleq (E_i \times E_i) - (\leq \cup \geq)_i$, and
 $\# \triangleq \{(e_1, e_2) \mid \exists i. \exists (e'_1, e'_2) \in \#_i. e'_1 \leq e_1, e'_2 \leq e_2\}$.

then $(E_1, E_2, \dots, E_n; \leq; \#)$ is an n -agent event structure.

Proof. Let $E = E_1 \cup \dots \cup E_n$. We only need to demonstrate that $(E, \leq, \#)$ is an event structure. (E, \leq) is a poset and $\#$ is easily seen to be irreflexive and symmetric. It remains to establish condition (I); so suppose $e_1 \# e_2$ and $e_2 \leq e_3$. Then for some i , $\exists (e'_1, e'_2) \in \#_i$ such that $e'_1 \leq e_1$, $e'_2 \leq e_2$. But then $e'_2 \leq e_3$ and hence $e_1 \# e_3$. □

Corollary A.6. Let (E, \leq, η) be a CSA and $E_i = \eta^{-1}(i)$, for all i , and $\#$ as defined in Proposition A.5. Then $(E, \leq, \#)$ is an event structure.

Using Propositions A.5 and A.4, a bijection can be constructed between the class of n -ACSAs and the class of n -agent event structures. Corollary A.6 allows construction of a map from the class of ACSAs to that of event structures. In the reverse direction, given an event structure $(E, \leq, \#)$, it is possible to have many different naming functions η such that (E, \leq, η) is a CSA. For instance, given E a countable set, we can let each event be an agent by itself. Hence ACSAs have more information built into them than event structures.