Università di Pisa

Dipartimento di Informatica

Technical Report : TR-99-07

Modalities for Asynchronous Distributed Systems

L. Semini, C. Montangero

March 31, 1999

ADDR: Corso Italia 40,56125 Pisa,Italy. TEL: +39-50-887111. FAX: +39-50-887226

Modalities for Asynchronous Distributed Systems

Carlo Montangero[1] and Laura Semini[2]

[1] Dip. di Informatica, Università di Pisa, monta@di.unipi.it
[2] Dip. di Sistemi e Informatica, Università di Firenze, semini@dsi.unifi.it

Abstract. The purpose of this work is to establish some foundational ground for the logics that can be used to reason on distributed systems communicating asyncronously via message passing, e.g. distributed workflow systems. These systems rule out synchronized local clocks, and reasoning on global time or state: message passing is the only possible mechanism for communication, and therefore for causality across components. We introduce *adtl*, a logic with modalities for time and locality. We show that its axiom system characterizes the class of structures that reflect the nature of the causal relation when communication is asyncronous. Correspondingly, the axioms are such that any attempt to formulate in the logic properties about the unspeakable entities, like a global state, results in a contradiction.

1 Introduction

The design of quality software for distributed systems is becoming more and more critical, due to the current impact of software on every technical accomplishment, and the fact that networks pervade any current application. The problem has two facets: the complexity of the systems under development, and the need of continuous update to keep the pace with moving requirements. We are convinced that a design process based on formal refinements, when centered on the system architecture and applied in the early phases of development, would mitigate the problems related to both complexity and change. To this end, we have been working on a refinement calculus for distributed systems that integrates in a natural way local refinements (i.e. within a single component) and coordination templates [12, 9, 10].

The essential element of the approach is the logic which is used to specify the components of the distributed systems and their interactions. Indeed, the approach can be effective only if it comes with simple, workable concepts and notations, well matched to the underlying paradigm. We are interested in distributed work–flow systems, a class of systems that are becoming more and more needed to solve problems in domains like geografically distributed manufacturing, electronic commerce, round–the–clock world–wide software development, etc. They essentally consist of loosely coupled asynchronous long running services.

In the context of this refinement–centered design approach, in [10] we introduced Oikos_adtl, an extention of Unity [3] to deal with components and events. As in Unity, in Oikos_adtl the emphasys is on high level operators and application oriented theorems , to be used in the everyday activity of formal design. Complementarily, the purpose of this work is to establish some foundational ground: we provide a characterization of the class of frames which have a structure that reflects the essential properties of the causal relation in distributed systems communicating asyncronously via message passing. In this paper we will refer to this class of systems as *asynchronous systems*.

Before delving into the technicalities, in the rest of this introduction we discuss informally the structure of the frames, and its impact on the logic presented here, especially with respect to the kind of properties of the computations that can be meaningfully expressed.

Asynchronous systems are characterized in [1] by the following properties: there are no bounds on the relative speeds of processes as well as on message delays. Therefore, they rule out synchronized local clocks, and reasoning on global time: message passing is the only possible mechanism for communication. This restricts the ways in which events in the computation can be ordered. In an asynchronous system, two events are constrained to occur in a certain order only if the occurrence of the first may affect the outcome of the second, that is only if there is a flow of information from the first to the second. This in turn implies that, if the events occur in different processes, they must be linked by the exchange of a message. This establishes an asymmetry between the processes involved in the communication, and is the obvious root of the differencies with distributed systems based on synchronous communication. These differences are discussed extensively in the section devoted to related works.

As it is fairly common in the specification of distributed systems, we use a modal logic, with modalities for time and locality. Oikos_adtl has been designed to support an essential facet of any design method: the composition of specifications. It fosters the expression of the global properties of a system in terms of the local properties and the interactions of the components. However, the global properties that may hold in a system are constrained by the underlying causal structure: any attempt to express properties that entail the syncronization of remote components must end up in a contradiction.

In our setting, the causal relation can be conveniently depicted by diagrams, reminiscent of Lamport's space–time diagrams [6]. They show the evolution of the states of the components, as the result of internal transitions or message passing. Fig 1 shows an example: w_j^i is the j^{th} state of component i. Plain arrows denote single steps, while dotted arrows denote paths composed of an undefined number of steps.

Fig. 1. A space–time diagram

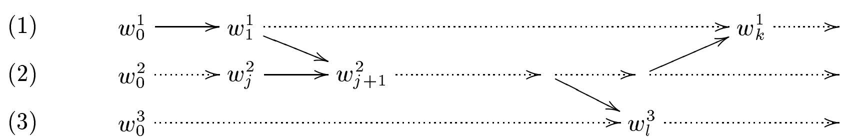

In particular, the diagram shows that w_{j+1} in component 2 is in causal relation with w_k in component 1 and with w_l in 3 as well, without entailing a causal relation of w_{j+1}^2 with a sort of global state made of w_k^1 and w_l^3.

Indeed, in asynchronous systems, no set of local states can be seen as a global one. Consequently, no logical formula can be true, which requires some properties of states in different components to hold at the same time instant. To satisfy this constraint, though maintaining the possibility of reasoning on global properties of a system, we define the logic *adtl*: formulae are interpreted over states of space–time diagrams, and the conjunction of properties of remote states is always false.

The main result of the paper is the definition of an axiom system for *adtl*, which characterizes the space–time diagrams structure.

The organization of the paper is the following. Sections 2 and 3 present the syntax and the semantics of *adtl*, respectively. Section 4 introduces the axioms. Their correspondence to the properties of the underlying frame is shown in Section 5. Before some concluding remarks, Section 6 discusses related work.

2 The logic *adtl*

The syntax of *adtl* formulae is given over a denumerable set $Prop$ of propositional symbols, and a finite, nonempty set I of component names. We consider a fixed set $I = \{1, \ldots, n\}$, and we let p, q, $\ldots$ range over $Prop$, and i, j, $\ldots$ range over I.

Definition 1. *An **adtl** **formula** F is defined by:*

$$F ::= p \mid \neg F \mid F \wedge F \mid \bigcirc_{\Box} F \mid \Box F \mid [i]F$$

Modality $\bigcirc_{\Box}$ means "in all states reachable in one step", and $\Box$ means "in all reachable states", as usual. The new modality $[i]$ is read "if in i then". Informally, $[i]F$ holds in a state w if either F holds in w, or if w is not a state of component i. Formal semantics is provided with Def. 4.

Some dual modalities can be introduced,

$$\Diamond F \overset{def}{\equiv} \neg\Box\neg F \qquad \bigcirc_{\Diamond} F \overset{def}{\equiv} \neg\bigcirc_{\Box}\neg F \qquad \langle i\rangle F \overset{def}{\equiv} \neg[i]\neg F$$

They read "in some reachable state", "in some state reachable in one step", and "in i and", respectively.

Example 1.

$\Box[\langle 1\rangle p \rightarrow \Diamond\langle 2\rangle q]$: this is a typical meaningful formula. Its intended meaning is that whenever p holds in 1, q will hold in 2 (and there will be a communication in between, from 1 to 2).

$\Diamond[\langle 1\rangle p \wedge \langle 2\rangle q]$: conversely, this is a typical unsatisfiable formula. Formulae are interpreted in local states, and no state can belong to more than a component.

$\Box[i]F$ says that F is an invariant of component i.

3 Semantics

In our setting, the interpretation domain for a formula is a space–time diagram shaping like the one in Fig. 1, composed of n noetherian denumerable chains, each representing the evolution of a single component, and an n-partite oriented graph, representing the asynchronous communications among the components.

Next, we define the frames for asynchronous systems, and list the properties satisfied by the accessibility relations. In Sect. 4, we will provide the axioms that characterize these properties.

Definition 2. *A* **frame** $\mathcal{F}$ *is a tuple* $\langle W, R_{\bigcirc}, R_1, \ldots, R_n\rangle$*, where* W *is a set of worlds;* $R_{\bigcirc} \subseteq W \times W$ *is an accessibility relation called* "next state"*; for each* $i = 1, \ldots, n$*,* $R_i \subseteq W \times W$ *is an accessibility relation called* "in i".

Worlds in W represent the states of the components. The "in i" relations have enough structure to partition W. Indeed, we require:

R_i is a partial identity: $(w, w') \in R_i \ \rightarrow \ w = w'$. The only accessible state from w according to R_i, if any, is w itself.

Disjunction of the R_i's: $R_i \cap R_j = \emptyset$ for any pair of $i \neq j$. Each world in W is a state of at most one component.

Reflexivity of $\bigcup_{i \in I} R_i$: $\forall w.\ (w, w) \in \bigcup_{i \in I} R_i$. Each world in W is a state of some component.

We will write $w \in i$ (w belongs to component i) as a shorthand for $(w, w) \in R_i$.

The accessibility relation "next state" is constrained to be linear in each component. Besides, access to other components is also linear: there is a unique state which is the recipient of a communication from a component to another one. That is, for any two pairs (w, w') and (w, w'') in $R_{\bigcirc}$, either $w' = w''$ or w' and w'' belong to distinguished components. Formally, the required properties of "next state" are:

Distributed linearity of $R_{\bigcirc}$: for all i, w, w', and w'', $(w, w') \in R_{\bigcirc}$, $(w, w'') \in R_{\bigcirc}$, $w', w'' \in i$ imply $w' = w''$.

Stepping of $\bigcirc_{\Box}$ through $\Box$: the accessibility relation used to give semantics to $\Box$ is the transitive and reflexive closure $R^*_{\bigcirc}$ of $R_{\bigcirc}$.

It is useful to distinguish the local next states, i.e. those due to local transitions, from the remote ones, i.e. those due to the passing of a message from one component to another. For each i in I, we define the *next i–local state accessibility* relation $R_{\bigcirc_i}$, as the set

$$\{(w,\ w') \mid (w,\ w') \in R_{\bigcirc} \text{ and } w,\ w' \in i\}$$

The *next i–local state accessibility* is examplified in Fig. 1 by the pair of states $(w^i_j,\ w^i_{j+1})$. We require that $R_{\bigcirc_i}$ is total, i.e. each world has a successor in the same component:

Totality of $R_{\bigcirc_i}$: for any $w \in i$ there exists a w' such that $(w, w') \in R_{\bigcirc_i}$.

We now define the models for *adtl* formulae. As usual, a model is a frame enriched with a valuation function which defines the truth of propositional formulae on the worlds of the frame [5].

Definition 3. *A* **model** *on a frame* $\mathcal{F} = \langle W, R_\bigcirc, R_1, \ldots, R_n\rangle$ *is a tuple* $\mathcal{M} = \langle \mathcal{F}, V\rangle$ *where* V *is a mapping (valuation):*

$$V : W \rightarrow 2^{Prop}$$

A formula F is true in a model $\mathcal{M}$ if and only if F is true in all the initial worlds of $\mathcal{M}$. The truth of a formula in a world of $\mathcal{M}$ is defined as follows.

Definition 4. *Given a model* $\mathcal{M} = \langle W, R_\bigcirc, R_1, \ldots, R_n, V\rangle$, *and a world* $w \in W$*:*

$$\langle \mathcal{M}, w\rangle \models p \quad \text{iff} \quad p \in V(w)$$

$$\langle \mathcal{M}, w\rangle \models \neg F \quad \text{iff} \quad \langle \mathcal{M}, w\rangle \not\models F$$

$$\langle \mathcal{M}, w\rangle \models F \wedge F' \quad \text{iff} \quad \langle \mathcal{M}, w\rangle \models F \text{ and } \langle \mathcal{M}, w\rangle \models F'$$

$$\langle \mathcal{M}, w\rangle \models \bigcirc_\square F \quad \text{iff} \quad \forall w' : wR_\bigcirc w'.\langle \mathcal{M}, w'\rangle \models F$$

$$\langle \mathcal{M}, w\rangle \models \square F \quad \text{iff} \quad \forall w' : wR^*_\bigcirc w'.\langle \mathcal{M}, w'\rangle \models F$$

$$\langle \mathcal{M}, w\rangle \models [i]F \quad \text{iff} \quad \forall w' : wR_i w'.\langle \mathcal{M}, w'\rangle \models F$$

Example 2. As an example, the structure (each world is labelled with the propositions holding in it):

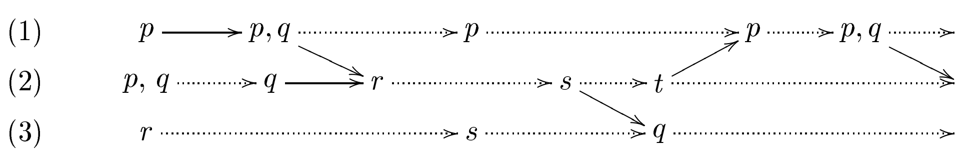

is a model of the formulae:

$\square\,[1]\,p$: p holds in all states of component 1 ($\square\langle 1\rangle p$ is false instead);

$[1]\,p \wedge [3]\,r$: p holds in the initial state of component 1, and r holds in the initial state of component 3;

$\square\,[\langle 2\rangle\, q \;\rightarrow\; (\Diamond\,\langle 1\rangle\, p \;\wedge\; \Diamond\,\langle 3\rangle\, q)]$: Any state in 2 satisfying q is eventually followed by a state in 1 satisfying p, and by a state in 3 satisfying q;

$\square\,(\langle 1\rangle\, q \;\rightarrow\; \bigcirc_\Diamond\,\langle 2\rangle\, true)$: A state of component 1 satisfying q is immediately followed by a state in 2, that is, every time q holds in 1, there is a communication from 1 to 2;

4 Axiom System

In this section we present the axiom system for *adtl*. F, F', F'' are *adtl* formulae built on the set $I = \{1, \ldots, n\}$, and $i, j \in I$.

Axioms of the propositional calculus:

P1 $F \to (F' \to F)$

P2 $(F \to (F' \to F'')) \to ((F \to F') \to F \to F''))$

P3 $((\neg F \to \neg F') \to ((\neg F \to F') \to F))$

Axiom **K**, for all modalities:

K$\Box$ $\Box(F \to F') \;\to\; (\Box F \to \Box F')$

K$\bigcirc_{\Box}$ $\bigcirc_{\Box}(F \to F') \;\to\; (\bigcirc_{\Box} F \to \bigcirc_{\Box} F')$

K$[i]$ $[i](F \to F') \;\to\; ([i]F \to [i]F')$

The following axioms are those normally used to define a temporal logic. Axioms **t1** and **t2** model the stepping of $\bigcirc_{\Box}$ through $\Box$. Axioms **4** and **T** define the transitive and reflexive properties of time, respectively, and axiom **D** states totality: a next state always exists.

t1 $\Box(F \to \bigcirc_{\Box} F) \to (F \to \Box F)$

t2 $\Box F \to \bigcirc_{\Box}\Box F$

4 $\Box F \;\to\; \Box\Box F$

T $\Box F \;\to\; F$

D $\bigcirc_{\Box} F \to \bigcirc_{\Diamond} F$

The last set of axioms deal with the new modality $[i]$. Correspondence with properties of the frames is discussed in Sect. 5.

a[i] $F \;\to\; [i]F$

b[i] $[1]F \wedge \ldots \wedge [n]F \;\to\; F$

c[i] $[i][j]F$ with $i \neq j$.

d[i] $\bigcirc_{\Diamond}\langle i\rangle F \;\to\; \bigcirc_{\Box}[i]F$

e[i] $\langle i\rangle\bigcirc_{\Box} F \;\to\; \bigcirc_{\Diamond}\langle i\rangle F$

Finally, we have the inference rules:

$$\textbf{MP} \quad \frac{F \quad F \to F'}{F'}$$

$$\textbf{Nec} \quad \frac{F}{\Box F}$$

Example 3. The following formulae are consequences of the axioms:

1) $\langle i \rangle \langle i \rangle F \;\leftrightarrow\; \langle i \rangle F$
2) $[i]F \;\leftrightarrow\; [i][i]F$
3) $F \;\rightarrow\; [i]\langle j \rangle F$
4) $\langle i \rangle F \wedge \langle j \rangle F' \;\leftrightarrow\; \langle i \rangle \langle j \rangle (F \wedge F')$
5) $[i](F \vee F') \;\leftrightarrow\; [i]F \vee [i]F'$

Proposition 1. *(Soundness) If F can be derived, then F is valid.*

Proof. The soundness of the axiom system is not difficult to show. We prove, as an example, the correctness of **c**[i].

We need to prove that $\langle \mathcal{M}, w \rangle \models [i][j]F$ holds, for any model $\mathcal{M}$, and any world w. By definition, $\langle \mathcal{M}, w \rangle \models [i][j]F$ if and only if $\forall w' : (w, w') \in R_i. \langle \mathcal{M}, w' \rangle \models [j]F$.

We recall that, according to R_i, w can only access itself, and distinguish two cases: $(w, w) \notin R_i$ and $(w, w) \in R_i$. In the first case, the thesis is immediately true. In the second case we are reduced to prove that $\langle \mathcal{M}, w \rangle \models [j]F$. This is true since R_i and R_j are disjoint. □

The correspondence results of the next section are a first step towards proving completeness.

5 Correspondence

We show that the axioms of the previous section correspond to the requirements on the accessibility relations listed in Sect. 3, i.e. they *characterize* the frames we are interested in.

A formula F is said to be *valid on a frame* $\mathcal{F}$ if and only if it is valid in all the models $\mathcal{M}$ built on $\mathcal{F}$. A formula F *characterizes* a class $\mathcal{C}$ of frames if and only if $\mathcal{C} = \{\mathcal{F} \mid F \text{ is valid in } \mathcal{F}\}$.

It is well known that axioms **P1**, **P2**, **P3**, **K**$\Box$, **K**$\bigcirc_\Box$, **t1**, **t2**, **4**, **T**, and **D** characterize the frames with a total next state relation ($R_\bigcirc$), stepping through a reflexive and transitive relation ($R^*_{\bigcirc}$). Correspondence results for the modalities $[i]$ are listed below. Some of these results are not new (for instance the statement of Prop. 2 can be found in [2]), however the axioms for $[i]$ do not hold in the modal logics normally used in computer science. We think that presenting these proofs might provide useful insights into the new modality.

Proposition 2. *The formula $F \rightarrow [i]F$ characterizes the class of frames* $\mathcal{C} = \{\langle W, C, R_1, \ldots, R_n\rangle \mid (w, w') \in R_i \;\rightarrow\; w = w'\}$

Proof. ($\Rightarrow$: if $\mathcal{F} \in \mathcal{C}$, then $F \rightarrow [i]F$ is valid in $\mathcal{F}$). By contradiction: suppose $F \rightarrow [i]F$ is not valid. This means that there exists a model $\mathcal{M} = \langle \mathcal{F}, V\rangle$, and a world w in W with $\langle \mathcal{M}, w\rangle \models F$ and $\langle \mathcal{M}, w\rangle \not\models [i]F$, i.e., there exists w' such that $(w, w') \in R_i$ and $\langle \mathcal{M}, w'\rangle \not\models F$. But, since $(w, w') \in R_i$ implies $w = w'$, we must have $\langle \mathcal{M}, w\rangle \not\models F$, a contradiction.

($\Leftarrow$: if the formula $F \rightarrow [i]F$ is valid in a frame $\mathcal{F}$, then $\mathcal{F} \in \mathcal{C}$). By contradiction: let $\mathcal{F}$ be a frame, $F \rightarrow [i]F$ be valid in it, but $\mathcal{F} \notin \mathcal{C}$. This means that there exist a pair of worlds $(w, w') \in R_i$ with $w \neq w'$. Since $F \rightarrow [i]F$ is valid in $\mathcal{F}$ for all F, in particular $p \rightarrow [i]p$ is valid in $\mathcal{F}$. Now, consider the model $\mathcal{M} = \langle \mathcal{F}, V\rangle$, where $V(w) = p$ and $V(w') = \emptyset$. We have, $\langle \mathcal{M}, w\rangle \models p$ and $\langle \mathcal{M}, w\rangle \not\models [i]p$, a contradiction.

□

Proposition 3. *The formula* $[1]F \wedge \ldots \wedge [n]F \;\rightarrow\; F$ *characterizes the class of frames* $\mathcal{C} = \{\langle W, C, R_1, \ldots, R_n\rangle \mid \bigcup_i R_i$ is reflexive$\}$

Proof. We consider $n = 2$. The generalization to the case of arbitrary n is straight.

($\Rightarrow$: if $\mathcal{F} \in \mathcal{C}$, then $[1]F \wedge [2]F \;\rightarrow\; F$ is valid in $\mathcal{F}$). By contradiction: suppose $[1]F \wedge [2]F \;\rightarrow\; F$ not to be valid. This means that there exists a model $\mathcal{M} = \langle \mathcal{F}, V\rangle$, and a world w in W with $\langle \mathcal{M}, w\rangle \models [1]F \wedge [2]F$ (i.e., $\langle \mathcal{M}, w\rangle \models [1]F$ and $\langle \mathcal{M}, w\rangle \models [2]F$), and $\langle \mathcal{M}, w\rangle \not\models F$. Since $R_1 \cup R_2$ is reflexive, $(w, w) \in R_1$ or $(w, w) \in R_2$. In both cases the conjunction of $\langle \mathcal{M}, w\rangle \models [1]F$ and $\langle \mathcal{M}, w\rangle \models [2]F$ imply $\langle \mathcal{M}, w\rangle \models F$, which is a contradiction.

($\Leftarrow$: if $[1]F \wedge [2]F \;\rightarrow\; F$ is valid in a frame $\mathcal{F}$, then $\mathcal{F} \in \mathcal{C}$). By contradiction, with $F = p$: let $[1]p \wedge [2]p \;\rightarrow\; p$ be valid in a frame $\mathcal{F} \notin \mathcal{C}$, i.e. in a frame including a world w such that $(w, w) \notin R_1 \cup R_2$. We take a model $\mathcal{M} = \langle \mathcal{F}, V\rangle$, where $V(w') = p$ for all w' with $(w, w') \in R_1$ or $(w, w') \in R_2$ and $V(w'') = \emptyset$ for all other worlds. Clearly, $\langle \mathcal{M}, w\rangle \models [1]p$ and $\langle \mathcal{M}, w\rangle \models [2]p$, consequently $\langle \mathcal{M}, w\rangle \models p$, a contradiction.

□

Proposition 4. *The formula* $[i][j]F$ *characterizes the class of frames* $\mathcal{C} = \{\langle W, C, R_1, \ldots, R_n\rangle \mid R_i \cap R_j = \emptyset\}$*, for* $i \neq j$.

The proof follows the usual schema and it is skipped here, and the same holds for the proofs of the following propositions. Prop. 5 states that **d**[i] characterizes the distributed linearity of $R_\bigcirc$, Prop. 6 states that **e**[i] corresponds to the totality of $R_{\bigcirc_i}$.

Proposition 5. *The formula* $\bigcirc_\Diamond\langle i\rangle F \rightarrow \bigcirc_\Box[i]F$ *characterizes the class of frames* $\mathcal{C} = \{\langle W, C, R_\bigcirc, R_1, \ldots, R_n\rangle \mid (w, w'), (w, w'') \in R_\bigcirc, w', w'' \in i$*, imply* $w' = w''\}$

Proposition 6. *The formula* $\langle i\rangle\bigcirc_\Box F \rightarrow \bigcirc_\Diamond\langle i\rangle F$ *characterizes the class of frames* $\mathcal{C} = \{\langle W, C, R_\bigcirc, R_1, \ldots, R_n\rangle \mid w \in C_i$ *implies* $\exists w' \in i$ *such that* $(w, w') \in R_\bigcirc\}$

Finally, we observe that relation R_i is symmetric and transitive. Indeed, we proved that formulae $F \rightarrow [i]\langle i\rangle F$, and $[i]F \rightarrow [i][i]$, characterizing symmetry and transitivity, respectively, are a consequence of the axioms of Sect. 4.

6 Related Work

Various extensions of temporal logic to deal with distributed systems have been defined in the literature. We present and discuss the proposals which are more closely related to ours.

One of the first logics defined to reason on distributed systems is TTL [8]. In this logic, for each local state of the system, a *visibility* function specifies which information of remote components is accessible. The visibility function is defined on the basis of a relation among states which is symmetric in the case of states belonging to distinguished components.

A trace based extension of linear time temporal logic, called *TrPTL*, has been defined in [13] (see also [14]). The logic has been designed to be interpreted over infinite traces, i.e., labelled partial orders of actions, which respect some dependence relations associated to the alphabet of actions.

In [7], a temporal logic, StepTL, is defined and interpreted over multistep transition systems. These are a well known extension of transition systems, allowing to describe as concurrent the steps of computation that can actually be executed in parallel. A multistep transition system thus contains transitions of the form $s\,A\,s'$, where A is a set of actions, instead of a single one.

Three distinguished logics are presented in [11] to describe systems composed of sets of communicating *agents*. The logics differentiate on the amount of information each agent can have on the other agents running on the system, but share a common setting: agents communicate via common actions. The models for these logics are runs of networks of synchronizing automata. The logics $\mathbf{D}_0$ and $\mathbf{D}_1$ presented in [4] are based on a similar approach.

Anyway, in all these proposals, components communicate via some form of synchronization, and logic formulae are interpreted on models shaping:

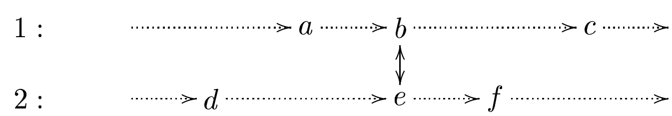

In these models, and consequently in any logic defined over them, it is not possible to express the asymmetric nature of causality as required when modelling the behaviour of agents communicating asynchronously by message passing. Indeed, in the previous model we can both assert that a causes f and that d causes c.

Moreover, the model above satisfies a formula like $\Diamond (\langle 1 \rangle b \ \wedge \ \langle 2 \rangle e)$, implicitly talking about global states. Indeed, if we interpret a logic in these models, the only way to prevent formulae on compound states from being true in the same time instant is to restrict the formulae syntax in some (often unnatural) way.

Conclusions

We have introduced *adtl*, a logic to reason on asynchronous distributed systems communicating via message passing. Models for *adtl* are based on space–time diagrams, which reflect the asymmetric nature of the causal relation in this case. Further, models built on space–time diagrams falsify any formula predicating about global states. Finally, we have introduced a sound axiom system for *adtl* and proved its correspondence to the properties of the considered frames.

An immediate extension of the work presented in this paper is to prove completeness and to explore decidability in *adtl*. Since Oikos_adtl can be defined in terms of *adtl*, decidability results would find a natural application in the refinement calculus we are developing for asynchronous systems.

References

1. Ozalp Babaoglu and Keith Marzullo. Consistent global states of distributed systems: Fundamental concepts and mechanisms. In S. Mullender, editor, *Distributed Systems*, pages 55–96. Addison-Wesley, 1993.
2. J. Van Benthem. Correspondence theory. In D. Gabbay and F. Guenther, editors, *Handbook of Philosophical Logic*, volume II, chapter II.4, pages 167–248. Reidel Publishing Company, 1984.
3. K.M. Chandy and J. Misra. *Parallel Program Design: A Foundation*. Addison-Wesley, Reading Mass., 1988.
4. H.-D. Ehrich, C. Caleiro, A. Sernadas, and G. Denker. Logics for specifying concurrent information systems. In J. Chomicki and G. Saake, editors, *Logic for Databases and Information Syustems*, pages 167–198. Kluver Academic Publishers, 1998.
5. S.A. Kripke. A completeness theorem in modal logic. *The Journal of Symbolic logic*, 24:1–14, 1959.
6. Leslie Lamport. Time, clocks, and the ordering of events in a distributed system. *Communications of the ACM*, 21(7):558–565, July 1978.
7. K. Lodaya, R. Parikh, R. Ramanujam, and P.S. Thiagarajan. A Logical Study of Distributed Transition Systems. *Information and Computation*, 119(1):91–118, 1995.
8. A. Masini and A. Maggiolo-Schettini. TTL: A formalism to describe local and global properties of distributed systems. *Informatique théorique et Applications/Theoretical Informatics and Applications*, 26(2), 1992.
9. C. Montangero and L. Semini. Refining by Architectural Styles or Architecting by Refinements. In L. Vidal, A. Finkelstein, G. Spanoudakis, and A.L. Wolf, editors, *2nd International Software Architecture Workshop, Proceedings of the SIGSOFT '96 Workshops, Part 1*, pages 76–79, San Francisco, CA, Oct 1996. ACM Press.
10. C. Montangero and L. Semini. Composing specifications for coordination. In P. Ciancarini and A. Wolf, editors, *Proc. 3nd Int. Conf. on Coordination Models and Languages*, volume to appear of *Lecture Notes in Computer Science*, Amsterdam, April 1999. Springer-Verlag.
11. R. Ramanujam. Locally linear time temporal logic. In *Proc.* 11^{th} *IEEE Symp. on Logic In Computer Science*, pages 118–127. IEEE Computer Society, 1996.
12. L. Semini and C. Montangero. A Refinement Calculus for Tuple Spaces. To appear in Science of Computer Programming.
13. P. S. Thiagarajan. A trace based extension of linear time temporal logic. In *Proceedings, Ninth Annual IEEE Symposium on Logic in Computer Science*, pages 438–447, Paris, Jul 1994. IEEE Computer Society Press.
14. P. S. Thiagarajan and J. G. Henriksen. Distributed versions of linear time temporal logic: A trace perspective. *Lecture Notes in Computer Science*, 1491, 1998.