Orthogonal Relations for Reasoning about Abstract Events

Ajay Kshemkalyani[1] and Roshan Kamath[2]

[1] EECS Department, University of Illinois at Chicago, Chicago, IL 60607-7053, USA.
ajayk@eecs.uic.edu
[2] Motorola Inc. 1501 W. Shure Drive, IL27-1L20, Arlington Heights, IL 60004, USA.
Roshan.Kamath@motorola.com

Abstract. As systems become increasingly complex, event abstraction becomes an important issue in order to represent interactions and reason at the right level of abstraction. Abstract events are collections of more elementary events, that provide a view of the system execution at an appropriate level of granularity. Understanding how two abstract events relate to each other is a fundamental problem for knowledge representation and reasoning in a complex system. In this paper, we study how two abstract events in a distributed system are related to each other in terms of the more elementary causality relation. Specifically, we analyze the ways in which two abstract events can be related to each other *orthogonally*, that is, identify all the possible mutually independent relations by which two such events could be related to each other. Such an analysis is important because all possible relationships between two abstract events that can exist in the face of uncertain knowledge can be expressed in terms of the irreducible orthogonal relationships.

1 Introduction

As systems become increasingly complex, event abstraction becomes an important issue in order to represent interactions and reason at the right level of abstraction. Abstract events are collections of more elementary events, that provide a view of the system execution at an appropriate level of granularity. Understanding how two abstract events relate to each other is a fundamental problem for knowledge representation and reasoning in such a complex distributed system. This problem is of interest across philosophy, physics, artificial intelligence, computer science, and psychology [2].

Hamblin [10] and Allen [2] have shown that two linear time durations or *intervals* that are colocated can be related in one of 13 possible ways. These 13 relations form an *orthogonal* set of relations, i.e., the intervals must be related by one and only one of these relations, implying that the conjunction of any two relations is the empty relation. Orthogonal relations are important because they identify all possible mutually exclusive relations that can possibly hold between any given pair of intervals and because all possible relationships between two intervals that can exist in the face of uncertain knowledge can be expressed

S. Benferhat and P. Besnard (Eds.): ECSQARU 2001, LNAI 2143, pp. 726–735, 2001.

in terms of the irreducible orthogonal relationships. The set of 13 orthogonal relations between a pair of colocated linear intervals has been used extensively in the literature on artificial intelligence. For example, [8] developed a theory of temporal reasoning using semi-intervals which arise when there is uncertain and imprecise knowledge of intervals, using the 13 orthogonal relations of Allen. Examples of other uses of the 13 orthogonal relations between colocated linear intervals include [3,4,5,9,14,15,16].

The literature surveyed above considered the interactions and relative placement of time intervals, each of which can be viewed as a *linearly ordered* set of time instants. An additional assumption was that time was continuous, and hence the time intervals satisfy the density axiom (refer van Benthem [6] for the formal definitions and a detailed discussion of continuity and density).

Our objective is to study how two abstract events in a distributed system are related to each other in terms of the causality relation. The relativistic space-time model is an appropriate model of a distributed system execution for this study. We analyze the ways in which two abstract events can be related to each other orthogonally, that is, identify all the possible mutually independent relations by which two such events could be related to each other. The results of this paper differ from the work surveyed above in the following aspects. Each of the abstract events we consider is a partial order of more elementary events, unlike the time intervals which linearly order the component time instants. Additionally, the system model explicitly models individual events/actions/statement executions that occur at different processes in the execution of a complex distributed system, and hence models discrete events explicitly.

The work is motivated by the fact that in a distributed system, *abstract* events, wherein at least some of the component elementary events of the abstract event occur concurrently, are of great interest in simplifying the reasoning about distributed executions [12,13]. Henceforth, we also term such abstract events as *poset (partially ordered set) events*. Such poset events accurately model collaborative activity among multiple CPU subsystems in a distributed system, for various applications like navigation, planning, robotics, mobile computing, coordination among multiple participants in a virtual reality environment, and agent-based distributed cooperating programs. As a specific example, multiple autonomous robots need to cooperate to jointly solve a task such as to focus laser beams on a target so that the beams arrive at the target at a fixed moment. As another example, multiple roving mobile agents that can communicate only by message passing need to synchronize their actions in an adversarial environment. Causality between poset events has been studied in [12] wherein a spectrum of fine-grained causality relations between poset events was presented, along with an axiom system to reason with such relations. These relations provide a precise handle to express and represent a naturally occurring or enforce a desired fine-grained level of causality or synchronization among the cooperating agents. However, these relations are not orthogonal relations. In this paper, we present a methodology for deriving orthogonal relations between poset events. Section 2 gives the system model. Section 3 gives the main results. Section 4 concludes.

2 System Model and Preliminaries

A poset event structure model $(E, \prec)$, where $\prec$ is an irreflexive partial ordering representing the causality relation on the possibly infinite event set E, is used as the space-time model for a system execution, as in [12]. $(E, \prec)$ can follow either the discreteness or the density axioms [6]. E is partitioned into local executions at coordinates in the space dimensions. Each E_i is a linearly ordered set of events in partition i and corresponds to the execution of events by a distinct process i. An event e in partition i is denoted e_i. The causality relation on E is the transitive closure of the local ordering relation on each E_i and the ordering imposed by message send events and message receive events. In [11,12], poset events are defined as follows. Let $\mathcal{E}$ denote the power set of E. Let $\mathcal{A}$ $(\neq \emptyset)$ $\subseteq (\mathcal{E} - \emptyset)$. $\mathcal{A}$ is the set of all those sets that represent a higher level grouping of the events of E of interest to an application. Each element A of $\mathcal{A}$ is a subset of E, and is termed an *abstract event* or a *poset event*.

Table 1. The six basic relations, see [11,12].

Relation r	Expression for $r(X, Y)$
$R1$	$\forall x \in X \forall y \in Y, x \prec y$ $(= \forall y \in Y \forall x \in X, x \prec y)$
$R2$	$\forall x \in X \exists y \in Y, x \prec y$
$R2'$	$\exists y \in Y \forall x \in X, x \prec y$
$R3$	$\exists x \in X \forall y \in Y, x \prec y$
$R3'$	$\forall y \in Y \exists x \in X, x \prec y$
$R4$	$\exists x \in X \exists y \in Y, x \prec y$ $(= \exists y \in Y \exists x \in X, x \prec y)$

The causality relations between a pair of poset events were formulated in [12] using the notion of *proxies*. Each poset event X was defined to have two proxies – the set of its least elements L_X, and the set of its greatest elements U_X. These proxies were the equivalents of the beginning and end instants of the linearly ordered interval. Two alternate definitions of proxies were given:

- Definition 4 [12], viz., $L_X = \{e_i \in X \mid \forall e'_i \in X, e_i \preceq e'_i\}$ and $U_X = \{e_i \in X \mid \forall e'_i \in X, e_i \succeq e'_i\}$, and
- Definition 5 [12], viz., $L_X = \{e \in X \mid \forall e' \in X, e \not\succ e'\}$ and $U_X = \{e \in X \mid \forall e' \in X, e \not\prec e'\}$

Figure 1 depicts the proxies of X and shows the difference between the two definitions. In the figure, the time axis goes from left to right, and the lines with arrows denote the messages that impose causality across different processes (points in space). Depending on the problem domain, an application chooses and consistently uses one definition of proxy. For example, for events in a distributed sensor/robot system, where the various sensors/robots cooperate to perform loosely synchronized actions, the former definition is more suitable to represent the start and end of interactions. When different mobile agents invoke services offered by other agents/servers in a nested Remote Procedure Call (RPC) form, the latter definition is more suitable to represent the start and end of interactions.

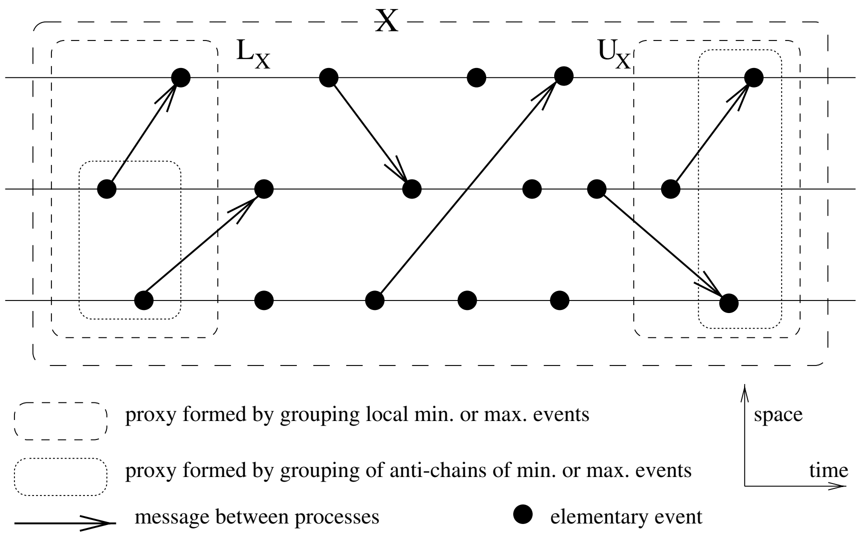

Fig. 1. Poset event X and its proxies L_X and U_X. The proxies defined by Definition 4 are shown by the closely spaced dashed lines. The proxies defined by Definition 5 are shown by dotted lines.

The causality relations in [12] were defined using the following two aspects of specifying the relations, based on the concept of proxies. (i) As there is a choice of two proxies of X and a choice of two proxies of Y, there are four combinations between the proxies. (ii) The six causality relations in Table 1 can be specified for each combination, thus yielding 24 relations between X and Y. The set of these causality relations is denoted $\mathcal{R}$. The following nomenclature was adopted to name the relations in $\mathcal{R}$. Relation $R?\#(X,Y)$ was such that $R?$ was a value from $\{R1, R2, R3, R4\}$ and indicated the choice of proxies of X and Y, whereas # indicated how the chosen proxies were related to each other, and took a value from { a, b, b', c, c', d }, where $R1$, $R2$, $R2'$, $R3$, $R3'$, $R4$ were renamed a, b, b', c, c', d, respectively, to avoid confusion with the previous usage of the relations $R1$ - $R4$. The set of relations $\mathcal{R}$ between poset events was complete using first-order predicate logic and only the $\prec$ relation between elementary events. The relation algebra given in [12] can be viewed as a power algebra [7].

In this paper, the label $\mathcal{R}$ is used to denote the set of the above relations when the discussion is common to the relations defined using either definition of proxies, viz., Definition 4 or 5 [12]. If the distinction matters, the notations $\mathcal{R}^{\prec_i}$ and $\mathcal{R}^{\prec}$ are used to denote the sets of relations that result when Definition 4 and 5 of proxies, respectively, are used. Intuitively, $\mathcal{R}^{\prec_i}$ indicates the set of relations resulting when the proxies are defined using the $\prec$ relation on each E_i, and $\mathcal{R}^{\prec}$

indicates the set of relations resulting when the proxies are defined using the $\prec$ relation on E. Each of $\mathcal{R}^{\prec}$ and $\mathcal{R}^{\prec_i}$ forms a hierarchy of *dependent* relations as shown in Figure 2. The relative hierarchy among relations in $\mathcal{R}^{\prec}$ and relations in $\mathcal{R}^{\prec_i}$ is given in [12].

A set of axioms to reason with the relations in $\mathcal{R}^{\prec}$ was given in [12]. The set of axioms was complete in the sense that (i) given any $R(X,Y)$, the axioms gave all enumerations of valid relations $r(X,Y)$ and $r'(Y,X)$, for $r, r', R \in \mathcal{R}^{\prec}$, and (ii) given $r_1(X,Y) \wedge r_2(Y,Z)$, the axioms gave all relations $r(X,Z)$ (and from (i), all $r'(Z,X)$), for $r, r', r_1, r_2 \in \mathcal{R}^{\prec}$. Hence, the axioms could be used to derive all possible implications from any given predicates on relations in $\mathcal{R}^{\prec}$.

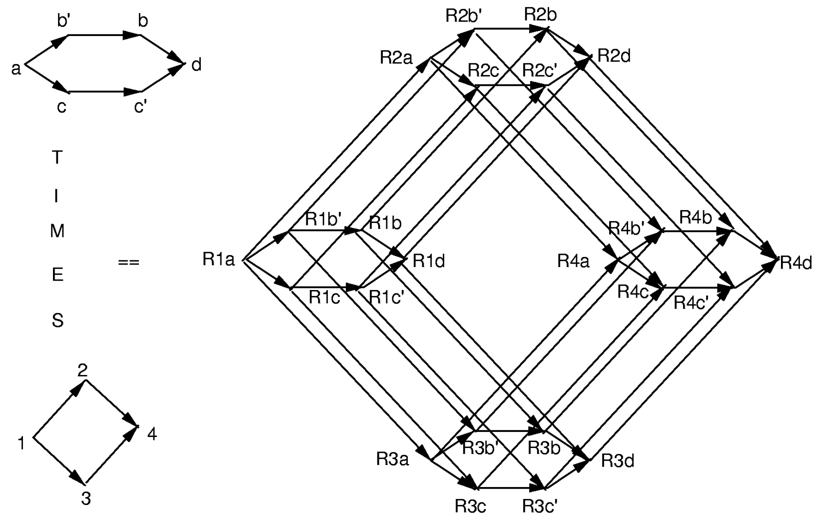

Fig. 2. Hierarchy of causality relations, ordered by "is a subrelation of" [12]. An edge from $r1$ to $r2$ indicates that $r1$ is a subrelation of $r2$.

In the next section, we give a methodology to enumerate the set of orthogonal relations for $\mathcal{R}$. The results of implementing this methodology for $\mathcal{R}^{\prec}$ using the axioms of [12] are then given. In this paper, we also modify the axiom system to make it applicable to $\mathcal{R}^{\prec_i}$. We then apply the above methodology to enumerate the set of orthogonal relations for $\mathcal{R}^{\prec_i}$ and give the results.

3 Orthogonal Relations

We now propose a method to derive and enumerate the orthogonal relations between any pair of poset events, using the set of dependent relations $\mathcal{R}$. We also present the numerical results of enumerating the orthogonal relations for $\mathcal{R}^{\prec}$ and $\mathcal{R}^{\prec_i}$ based on the appropriate axiom system. Specifically, for $\mathcal{R}^{\prec}$, we use axioms XP1-XP14 given in [12]. For $\mathcal{R}^{\prec_i}$, we use axioms XP1-XP6 and eight new axioms XP7$^{\prec_i}$-XP14$^{\prec_i}$. The results of the two enumerations were obtained by implementing the methodology in XSB Prolog.

The algorithm proposed here has the following two steps to create a (complete and mutually independent) set of orthogonal relations from the set of dependent relations $\mathcal{R}$.

1. Identify all possible combinations of relations $r(X, Y) \in \mathcal{R}$ that can hold simultaneously for a given X and Y.
2. For each of the identified combinations of relations $r(X, Y)$, identify all combinations of $r(Y, X)$ that can simultaneously hold for the same X and Y.

3.1 Step 1: All Possible Relations $r(X, Y)$

As a first step, we identify all the combinations of relations $r(X, Y)$, for $r \in \mathcal{R}$, that hold between poset events X and Y. Note that by construction, $(\mathcal{R}, \sqsubseteq)$, where $\sqsubseteq$ is the relation "is a subrelation of", is a lattice as illustrated in Figure 2. For a given pair of posets X and Y, it may be the case that a combination of the relations in $\mathcal{R}$ may hold. Specifically, if $R(X, Y)$ holds, then $\forall R' \mid R \sqsubseteq R'$, $R'(X, Y)$ holds. Thus, if $R(X, Y)$ holds, then for each R' in the upward-closed subset[1] of $\mathcal{R}$, $R'(X, Y)$ holds. In the partial order $(\mathcal{R}, \sqsubseteq)$, all upward-closed subsets of $\mathcal{R}$ correspond exactly to the combinations of relations in $\mathcal{R}$ that can hold concurrently for a given pair of poset events. It follows from the result on page 400 [1] that there is a 1-1 correspondence between the set of all upward-closed subsets of a partial order and the set of antichains[2] in the partial order. Therefore, an enumeration of the antichains in $(\mathcal{R}, \sqsubseteq)$ gives an enumeration of the upward-closed subsets of $(\mathcal{R}, \sqsubseteq)$, which corresponds to all the combinations of the relations in $\mathcal{R}$ that can hold for a pair of poset events. Let $\mathcal{RAC}$ be the set of all such antichains. A member of $\mathcal{RAC}$, denoted $rac(X, Y)$, is an antichain of $\mathcal{R}$ and can be expressed as the conjunction of the members of the antichain, each of which is a member of $\mathcal{R}$, i.e., $rac(X, Y)$ can be viewed as $\bigwedge_{r \in rac(X,Y)} r(X, Y)$. The number of antichains in $\mathcal{RAC}$ was computed by the implementation of axioms XP1-XP6 (given below), to be as follows. There are 1, 24, 147, 350, 341, 168, 44, 2, and 0 antichains of size 0 through 8, respectively, giving a total of 1077 antichains. The antichain of size 0 denotes the empty-set upward-closed subset of $\mathcal{R}$, equivalent to $\overline{R4d}(X, Y)$, where $\overline{R4d}(X, Y)$ denotes that $R4d(X, Y)$ is false. Observe from Figure 2 that the size of the largest antichain is 7.

The axioms XP1 - XP6 from [12] are reproduced here. The relation $||(r_1, r_2)$ stands for $\not\sqsubseteq (r_1, r_2) \wedge \not\sqsubseteq (r_2, r_1)$. V_1 denotes the set $\{1, 2, 3, 4\}$ and V_2 denotes the set $\{a, b, b', c, c', d\}$.

XP1. $R1? \sqsubseteq R2? \sqsubseteq R4?$, where ? is instantiated from V_2
XP2. $R1? \sqsubseteq R3? \sqsubseteq R4?$, where ? is instantiated from V_2
XP3. $R2?||R3\#$, where ? and # are separately instantiated from V_2

[1] A set $\Re \subseteq \mathcal{R}$ is upward-closed iff for every $r, r' \in \mathcal{R}$, $(r \in \Re \bigwedge r \sqsubseteq r') \implies r' \in \Re$.
[2] A set $\Re$ is an anti-chain iff for every r and r' in $\Re$, $r \not\sqsubseteq r' \wedge r' \not\sqsubseteq r$.

XP4. $R?a \sqsubseteq R?b' \sqsubseteq R?b \sqsubseteq R?d$, where ? is instantiated from V_1
XP5. $R?a \sqsubseteq R?c \sqsubseteq R?c' \sqsubseteq R?d$, where ? is instantiated from V_1
XP6. $R?b||R?c'$, $R?b'||R?c'$, $R?b||R?c$, $R?b'||R?c$, where ? is instantiated from V_1

3.2 Step 2: Relations $r(Y, X)$, Given That Certain $r(X, Y)$ Hold

The computed combinations of relations in $\mathcal{R}$, viz., antichains in $(\mathcal{R}, \sqsubseteq)$, are useful to determine all the orthogonal relations that can exist between any two poset events. For each of the $|\mathcal{RAC}|$ antichains that hold between X and Y, there are potentially $|\mathcal{RAC}|$ antichains that hold between Y and X, thus leading to a potential $|\mathcal{RAC}|^2$ orthogonal relations between X and Y. Several of these relations will be illegal because they contradict the relations $r(X, Y)$. The objective is to determine exactly all the orthogonal relations that are admissible by the axiom system. For each $rac1(X, Y)$, where $rac1 \in \mathcal{RAC}$, determine which $rac2(Y, X)$ can hold, where $rac2 \in \mathcal{RAC}$, using the axiom system which allows the derivation of all $r'(Y, X)$ from any $r(X, Y)$, where r, $r' \in \mathcal{R}$. Then each conjunction of an antichain $rac1(X, Y)$ and a compatible antichain $rac2(Y, X)$ is orthogonal from every other such conjunction; denote this set of conjunctions as $\mathcal{RO}$, which then represents all the possible orthogonal relations between two posets, based on the $\prec$ relation among elementary events.

Let us denote the sets of orthogonal relations obtained for relations in $\mathcal{R}^{\prec}$ and $\mathcal{R}^{\prec_i}$ by $\mathcal{RO}^{\prec}$ and $\mathcal{RO}^{\prec_i}$, respectively.

Table 2. Number of orthogonal relations in $\mathcal{RO}^{\prec}$, classified based on size of antichains.

Size/Number of $rac(X,Y)$ antichains	Number of antichains $rac(Y,X)$ of size $s = 0 \ldots 7$								$\sum_{s=0}^{7} col_s$
	$s=0$	$s=1$	$s=2$	$s=3$	$s=4$	$s=5$	$s=6$	$s=7$	
0 / 1	1	24	147	350	341	168	44	2	1077
1 / 24	24	261	898	1285	822	264	34	1	3589
2 / 147	147	898	1911	1683	642	130	4	0	5415
3 / 350	350	1285	1683	937	180	8	0	0	4443
4 / 341	341	822	642	180	18	0	0	0	2003
5 / 168	168	264	130	8	0	0	0	0	570
6 / 44	44	34	4	0	0	0	0	0	82
7 / 2	2	1	0	0	0	0	0	0	3

Relations $\mathcal{RO}^{\prec}$. Axioms XP7-XP14 along with XP1-XP6 were used to determine all the orthogonal relations $\mathcal{RO}^{\prec}$, counted in Table 2. Axioms XP7-XP14 are reproduced below with labels XP7$^{\prec}$ - XP14$^{\prec}$, respectively.

XP7$^{\prec}$. $R1a(X, Y) \bigvee R1b(X, Y) \bigvee R1b'(X, Y) \bigvee R1c(X, Y) \bigvee R1c'(X, Y) \Longrightarrow \overline{R4d}(Y, X)$.
XP8$^{\prec}$. $R1d(X, Y) \Longrightarrow \overline{R4b}(Y, X) \bigwedge \overline{R4c'}(Y, X)$.
XP9$^{\prec}$. $R2a(X, Y) \bigvee R2b(X, Y) \bigvee R2b'(X, Y) \bigvee R2c(X, Y) \bigvee R2c'(X, Y) \Longrightarrow \overline{R2d}(Y, X)$.

XP$10^{\prec}$**.** $R2d(X,Y) \Longrightarrow \overline{R2b}(Y,X) \bigwedge \overline{R2c'}(Y,X)$.
XP$11^{\prec}$**.** $R3a(X,Y) \bigvee R3b(X,Y) \bigvee R3b'(X,Y) \bigvee R3c(X,Y) \bigvee R3c'(X,Y) \Longrightarrow \overline{R3d}(Y,X)$.
XP$12^{\prec}$**.** $R3d(X,Y) \Longrightarrow \overline{R3b}(Y,X) \bigwedge \overline{R3c'}(Y,X)$.
XP$13^{\prec}$**.** $R4a(X,Y) \bigvee R4b(X,Y) \bigvee R4b'(X,Y) \bigvee R4c(X,Y) \bigvee R4c'(X,Y) \Longrightarrow \overline{R1d}(Y,X)$.
XP$14^{\prec}$**.** $R4d(X,Y) \Longrightarrow \overline{R1b}(Y,X) \bigwedge \overline{R1c'}(Y,X)$.

Table 2 consists of three parts, separated by vertical double-lines. The first part categorizes the $|\mathcal{RAC}(X,Y)|$ antichains of Figure 2, based on size which ranges from 0 to 7. Each row i, $i \in [0 \ldots 7]$, in the entire table is used to compute the orthogonal relations in which antichains $rac(X,Y)$ have size i. Consider any row i. For each antichain $rac(X,Y)$ of size i, the number of the corresponding legal (as per XP$7^{\prec}$–XP$14^{\prec}$) antichains $rac(Y,X)$ of size s, $s \in [0,\ldots,7]$, are added to column s in the second part of the table. The entry in row i in the last part of the table sums up the row entires of columns $s = 0$ through $s = 7$ of that row, and gives the total number of orthogonal relations in which antichains $rac(X,Y)$ have size i. The sum of the last column is 17,185 = $|\mathcal{RO}^{\prec}|$.

Note that $\mathcal{RAC}$ needs to consider all the antichains in $\mathcal{R}$, not just the maximal antichains, because even a subset of a maximal antichain identifies a different upward-closed subset of $\mathcal{R}$ than does the maximal antichain, indicating a different set of relations that hold. Also note that for any $rac1(X,Y)$, all relations in the upward-closed subset of $\mathcal{R}$ hold and those not in the upward-closed subset do not hold. Thus, for any $rac1(X,Y)$, there is a bit-vector of size 24 where each bit corresponds to a relation in $\mathcal{R}$, such that there is a "1" for each relation in the upward-closed subset of $rac1(X,Y)$ and a "0" for each relation not in the upward-closed subset of $rac1(X,Y)$. Analogously, for any $rac2(Y,X)$ that is compatible with $rac1(X,Y)$ as per the axioms, there is a bit-vector of size 24 where each bit corresponds to a relation in $\mathcal{R}$, such that there is a "1" for each relation in the upward-closed subset of $rac2(Y,X)$ and a "0" for each relation not in the upward-closed subset of $rac2(Y,X)$. Each orthogonal relation can thus be represented by a 48-bit vector.

Example: For the $rac1(X,Y)$ antichain $R2b(X,Y) \bigwedge R2c(X,Y) \bigwedge R3a(X,Y)$ of size three, the axioms XP$7^{\prec}$-XP$14^{\prec}$ give $\overline{R2d}(Y,X) \bigwedge \overline{R3d}(Y,X)$. The only possible antichains $rac2(Y,X)$ can be from the set of relations $\{ R4{*}(Y,X) \}$ – this gives 11 possible antichains $rac2(Y,X)$, counting the antichain of size 0, that are compatible with $rac1(X,Y)$. Each of these 11 combinations of $rac2(Y,X)$ with $rac1(X,Y)$ yields a unique 48-bit vector.

Relations $\mathcal{RO}^{\prec_i}$. Observe that the axioms XP7-XP14 given in [12] are applicable only to relations in $\mathcal{R}^{\prec}$ which use Definition 5 of proxies [12], and not to relations in $\mathcal{R}^{\prec_i}$ which use Definition 4 of proxies [12]. If proxies are defined by Definition 4 and not Definition 5, then the axioms XP7-XP14 need to be replaced by the following axioms XP$7^{\prec_i}$–XP$14^{\prec_i}$ to obtain all the orthogonal relations $\mathcal{RO}^{\prec_i}$.

XP$7^{\prec_i}$**.** $R1a(X,Y) \Longrightarrow \overline{R4d}(Y,X)$;
$R1b(X,Y) \bigvee R1b'(X,Y) \Longrightarrow \overline{R4b}(Y,X)$;
$R1c(X,Y) \bigvee R1c'(X,Y) \Longrightarrow \overline{R4c'}(Y,X)$.
XP$8^{\prec_i}$**.** $R1d(X,Y) \Longrightarrow \overline{R4a}(Y,X)$.
XP$9^{\prec_i}$**.** $R2a(X,Y) \Longrightarrow \overline{R2d}(Y,X)$;
$R2b(X,Y) \bigvee R2b'(X,Y) \Longrightarrow \overline{R2b}(Y,X)$;
$R2c(X,Y) \bigvee R2c'(X,Y) \Longrightarrow \overline{R2c'}(Y,X)$.
XP$10^{\prec_i}$**.** $R2d(X,Y) \Longrightarrow \overline{R2a}(Y,X)$.
XP$11^{\prec_i}$**.** $R3a(X,Y) \Longrightarrow \overline{R3d}(Y,X)$;
$R3b(X,Y) \bigvee R3b'(X,Y) \Longrightarrow \overline{R3b}(Y,X)$;
$R3c(X,Y) \bigvee R3c'(X,Y) \Longrightarrow \overline{R3c'}(Y,X)$.
XP$12^{\prec_i}$**.** $R3d(X,Y) \Longrightarrow \overline{R3a}(Y,X)$.
XP$13^{\prec_i}$**.** $R4a(X,Y) \Longrightarrow \overline{R1d}(Y,X)$;
$R4b(X,Y) \bigvee R4b'(X,Y) \Longrightarrow \overline{R1b}(Y,X)$;
$R4c(X,Y) \bigvee R4c'(X,Y) \Longrightarrow \overline{R1c'}(Y,X)$.
XP$14^{\prec_i}$**.** $R4d(X,Y) \Longrightarrow \overline{R1a}(Y,X)$.

Axioms XP1-XP6 and XP$7^{\prec_i}$-XP$14^{\prec_i}$ are used to derive the orthogonal relations $\mathcal{RO}^{\prec_i}$, instead of axioms XP1-XP6 and XP$7^{\prec}$-XP$14^{\prec}$ that were used to obtain $\mathcal{RO}^{\prec}$. Results analogous to those in Table 2 for $\mathcal{RO}^{\prec}$ are obtained for $\mathcal{RO}^{\prec_i}$ and shown in Table 3. The sum of the last column is 123,474 = $|\mathcal{RO}^{\prec_i}|$.

Table 3. Number of orthogonal relations in $\mathcal{RO}^{\prec_i}$, classified based on size of antichains.

Size/Number of $rac(X,Y)$ antichains	Number of antichains $rac(Y,X)$ of size $s = 0 \ldots 7$								$\sum_{s=0}^{7} col_s$
	$s=0$	$s=1$	$s=2$	$s=3$	$s=4$	$s=5$	$s=6$	$s=7$	
0 / 1	1	24	147	350	341	168	44	2	1077
1 / 24	24	405	1926	3695	3084	1326	293	11	10764
2 / 147	147	1926	7097	11493	7963	2768	527	18	31939
3 / 350	350	3695	11493	16469	9406	2654	469	16	44552
4 / 341	341	3084	7963	9406	4158	802	132	4	25890
5 / 168	168	1326	2768	2654	802	18	0	0	7736
6 / 44	44	293	527	469	132	0	0	0	1465
7 / 2	2	11	18	16	4	0	0	0	51

4 Conclusions

Orthogonal relations between events provide an understanding of all possible mutually exclusive relations that can hold between the events when complete and precise knowledge is available. These form the basis of relation algebras, and allow the derivation of relations to represent knowledge when imprecise and incomplete information is available. Abstract events, each of which is a partially ordered collection of elementary events, are important when reasoning and representing actions in complex distributed systems. We derived orthogonal relations $\mathcal{RO}$ between abstract events using the space-time model for a distributed system execution. Relations in $\mathcal{RO}$ are analogous to the 13 orthogonal relations between linear intervals at a point in space [2]. Relations in $\mathcal{RO}$ are also analogous to the following sets of orthogonal relations based on the elementary causality relation: (i) the three orthogonal relations between two points in space-time ($a \prec b$, $b \prec a$, $a \not\prec b \wedge b \not\prec a$), (ii) the six orthogonal relations between a linear interval and a point in space-time [11], (iii) the 29 orthogonal relations between two linear intervals in space-time using the dense model of time [11], and (iv) the 40 orthogonal relations between two linear intervals in space-time using the nondense model of time [11]. We expect that as distributed agent-based programs and applications become more common, specific uses for these orthogonal relations between abstract events will emerge, similar to the uses of the 13 orthogonal relations between colocated linear intervals.

Acknowledgements. This work was supported by the U.S. National Science Foundation grants CCR-9875617 and EIA-9871345.

References

1. M. Aigner, *Combinatorial Theory,* Springer-Verlag, 1979.
2. J. Allen, Maintaining knowledge about temporal intervals, *Communications of the ACM*, 26(11):832-843, 1983.
3. J. Allen, Towards a general theory of action and time, *Artificial Intelligence*, 23:123-154, 1984.
4. P. van Beek, Reasoning about qualitative temporal information, *Artificial Intelligence*, 58:297-326, 1992.
5. P. van Beek, R. Cohen, Exact and approximate reasoning about temporal relations, *Computational Intelligence*, 6:132-144, 1990.
6. J. van Benthem, *The Logic of Time*, Kluwer Academic Publishers, (1ed. 1983), 2ed. 1991.
7. C. Brink, Power structures, *Algebra Universalis,* Vol. 30, 177-216, 1993.
8. C. Freksa, Temporal reasoning based on semi-intervals, *Artificial Intelligence*, 54: 199-227, 1992.
9. A. Gerevini, L. Schubert, Efficient algorithms for qualitative reasoning about time, *Artificial Intelligence*, 74: 207-248, 1995.
10. C. L. Hamblin, Instants and intervals, In *The Study of Time*, pp. 324-332, Springer-Verlag, 1972.
11. A. Kshemkalyani, Temporal interactions of intervals in distributed systems, *Journal of Computer and System Sciences*, 52(2), 287-298, April 1996.
12. A. Kshemkalyani, Reasoning about causality between distributed nonatomic events, *Artificial Intelligence*, 92(2): 301-315, May 1997.
13. L. Lamport, On interprocess communication, Part I: Basic formalism, *Distributed Computing*, 1:77-85, 1986.
14. B. Nebel, H. -J. Buerckert, Reasoning about temporal relations: A maximal tractable subclass of Allen's interval algebra, *Journal of the ACM*, 42(1): 43-66, Jan. 1995.
15. R. Rodriguez, F. Anger, K. Ford, Temporal reasoning: a relativistic model, *International Journal of Intelligent Systems*, 11: 237-254, 1991.
16. P. Terenziani, P. Torasso, Time, action-types, and causation, *Computational Intelligence*, 11(3):529-552, 1995.