Causality between Nonatomic Poset Events in Distributed Computations

Ajay D. Kshemkalyani
ECECS Department, P.O. Box 210030, University of Cincinnati, Cincinnati, OH 45221-0030, USA
Email: ajayk@ececs.uc.edu

Abstract

Recently, a set of causality relations between distributed nonatomic events was proposed to provide a fine level of granularity in the specification of synchronization conditions between the events. This set of causality relations is complete in first-order predicate logic. In this paper, we examine a set of axioms on the proposed causality relations. The axioms provide a mechanism for reasoning with the set of relations and can be used to derive all possible implied relations from any valid predicate on the relations.

Keywords: *Atomicity, Causality, Distributed system, Synchronization, Time.*

1 Introduction

Motivation: Event abstraction in a computation (or system execution) deals with the grouping of elementary events in the computation into higher level nonatomic events [8, 11, 15, 18]. Distributed applications such as industrial process control, distributed debugging, navigation, planning, robotics, diagnostics, virtual reality, and coordination in mobile systems model such nonatomic events [9, 12, 13]. These applications deal with nonatomic events that are nonlinear, i.e., for each nonatomic event, at least two of its component atomic events occur concurrently at more than a single point in space [10, 15]. For these applications, the traditional causality relation [6, 14, 16, 19] defined between individual points in space-time does not suffice for the following reason. The interaction and synchronization conditions between two nonatomic events cannot be captured or specified at a fine level of granularity using various degrees of causality, as required for a sophisticated and realistic modeling of these applications. So a rich set of causality relations that allow the expression of various degrees of synchronization and causality to accurately represent and specify relationships between distributed nonatomic events was proposed [9, 12, 13]. The relations can then be composed to form global predicates involving several distributed nonatomic events. We propose a system of axioms to reason with the proposed relations.

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

Table 1. Relations in [10].

Model: We use the space-time model for a system execution. This model is a poset event structure model as in [10, 15, 19, 20]. Consider a poset $(E, \prec)$ where $\prec$ is an irreflexive partial order. Let $\mathcal{E}$ denote the power set of E and let $\mathcal{A}$ $(\neq \emptyset) \subseteq (\mathcal{E} - \emptyset)$. There is thus an implicit one-many mapping from $\mathcal{A}$ to E. Each element A of $\mathcal{A}$ is a nonempty subset of E, and is termed an *interval* or a *nonatomic event.* We will use the term "interval" interchangably with "event" when referring to nonatomic events. $(E, \prec)$ represents points in space-time which are the most primitive atomic events related by the causality relation. Elements of E are partitioned into local executions at a coordinate in the space dimensions. Each local execution E_i is a linearly ordered set of events in partition i. An event e in partition i is denoted e_i.

Previous Work: In the literature, relations between time durations and between instants have been studied in the context of time and interval algebras; several axiom systems have been proposed for these relations. Most previous work assumed that the nonatomic events were linearly ordered, e.g., [5, 6] – and confined the study of causality to relations between time durations or linear intervals. [6] includes an excellent review of related literature. The causality relations defined in the literature above also assumed that such linear nonatomic events occurred at a single point in space, implying the existence of a global time axis. But in a distributed system, there is no global time axis [1, 14, 16, 19]. The following literature deals with causality between nonatomic

1071-0485/97 $10.00 © 1997 IEEE

Authorized licensed use limited to: University of Illinois at Chicago Library. Downloaded on June 22,2020 at 22:34:15 UTC from IEEE Xplore. Restrictions apply.

relation of row header to column header	$R1$	$R2$	$R3$	$R4$
$R1$	$=$	$\sqsubseteq$	$\sqsubseteq$	$\sqsubseteq$
$R2$	$\sqsupseteq$	$=$	$\|\|$	$\sqsubseteq$
$R3$	$\sqsupseteq$	$\|\|$	$=$	$\sqsubseteq$
$R4$	$\sqsupseteq$	$\sqsupseteq$	$\sqsupseteq$	$=$

Table 2. Inclusion relationships between relations, from [10].

poset events in a distributed system execution and does not assume a global time axis.

Lamport defined system executions using two relations $\longrightarrow$ and $- - \rightarrow$ between primitive nonatomic elements and provided axioms A1 - A5 on these relations [15]. Informally, these relations are as follows. Let a nonatomic event be a set of atomic events. For two nonatomic events X and Y in $\mathcal{A}$, $X \longrightarrow Y$ iff every atomic event in X causally precedes every atomic event in Y. $X - - \rightarrow Y$ iff some atomic event in X causally precedes some atomic event in Y. The model and axioms in [15] were further examined in [1, 3].

Action refinement of posets is studied and surveyed along with a survey of related work in Petri nets in [8, 17, 18]. In these areas, there is no known work that addresses specific causality relations between nonatomic poset events.

It was shown earlier [10] that the two causality relations defined by Lamport are not sufficient to capture the essential temporal properties of system executions and specify synchronization and causality conditions between nonatomic events in distributed systems. In [10], we proposed a set of new causality relations between nonatomic events in a distributed system to capture a range of causality and synchronization specifications, without assuming a global time axis. These relations $R1$ - $R4$ and $R1'$ - $R4'$ from [10] are expressed in terms of the quantifiers over X and Y in Table 1.

Observe that all the relations in Table 1 are not independent relations. Table 2 gives the hierarchy and inclusion relationship of the causality relations $R1$ - $R4$. Each cell in the grid indicates the relationship of the row header to the column header. The notation for the inclusion relationship between causality relations on nonatomic events is as follows. The inclusion relation "is a subrelation of" is denoted '$\sqsubseteq$'. '$\sqsupseteq$' is the inverse of $\sqsubseteq$. '$=$' stands for equality between relations in addition to its standard usage as the equality in other contexts. For two causality relations r_1 and r_2, we define $r_1 \,||\, r_2$ to be $(r_1 \not\sqsubseteq r_2 \bigwedge r_2 \not\sqsubseteq r_1)$. The relations $\{ R1, R2, R3, R4 \}$ form a lattice hierarchy ordered by $\sqsubseteq$. Table 1 also defined relations $R1'$, $R2'$, $R3'$, and $R4'$, for which the order of quantifiers was reversed from the order in $R1$, $R2$, $R3$, and $R4$, respectively. Note that the relations $R2'$ and $R3'$ are different from relations $R2$ and $R3$, respectively, when applied to posets. However, for a linear interval, they are the same as $R2$ and $R3$, respectively. $R1'$ and $R4'$ are the same as $R1$ and $R4$, respectively.

The set of relations proposed in [10] formed an exhaustive set of causality relations to express all possible interactions between a pair of linear intervals and extended the incomplete hierarchy of relations in [15]. However, when the relations of [10] are applied to a pair of poset intervals, the hierarchy they form is incomplete. [9, 12, 13] formulated causality relations between a pair of nonatomic poset intervals by extending the results [9, 10] to nonatomic poset events. The relations form an "exhaustive" set of causality relations between nonatomic poset events using first-order predicate logic and fill in the existing partial hierarchy of causality relations between nonatomic poset events, formed by relations in [10, 15]. In this paper, we propose an axiom system on the causality relations, which extends the axiom systems of the relations in [10, 15]. The axioms provide a mechanism for reasoning with the set of relations and can be used to derive all possible implied relations from any valid predicate on the relations.

Organization: Section 2 reviews the fine-grained hierarchy of causality relations from [9, 12, 13]. Section 3 gives the axiom system on the relations. Section 4 concludes. The results of this paper are included in [9].

2 Relations between Nonatomic Poset Events

Let $\mathcal{A}$ be the set of all the sets that represent higher level groupings of the events of E, that are of interest to the particular application. An element of $\mathcal{A}$ is denoted A.

Definition 1 *An interval A is linear iff* $\forall x, y \in A, x \preceq y \bigvee y \preceq x$.

Definition 2 N_A*, the node set of interval A, is* $\{i | E_i \bigcap A \neq \emptyset\}$.

Our results apply to nonlinear, i.e., poset, intervals.

The relations in [10] are used to derive an exhaustive set of causality relations between nonatomic poset events, denoted $\mathcal{R}$. As an intermediate step, we propose definitions of certain proxies of a nonatomic event in Section 2.1.

2.1 Proxies of Nonatomic Poset Events

In the extensive literature on linear intervals and time durations, for example [5, 6, 7], an interval is identified by the instants of its beginning and end. The beginning and end instants of a linear interval are points in space-time which are atomic events in E. For a nonatomic poset interval, it is natural to identify counterparts for the beginning and end instants. These counterparts will serve as "proxy" events for the poset interval just as the events at the beginning and end of linear intervals such as time durations serve as proxies

Authorized licensed use limited to: University of Illinois at Chicago Library. Downloaded on June 22,2020 at 22:34:15 UTC from IEEE Xplore. Restrictions apply.

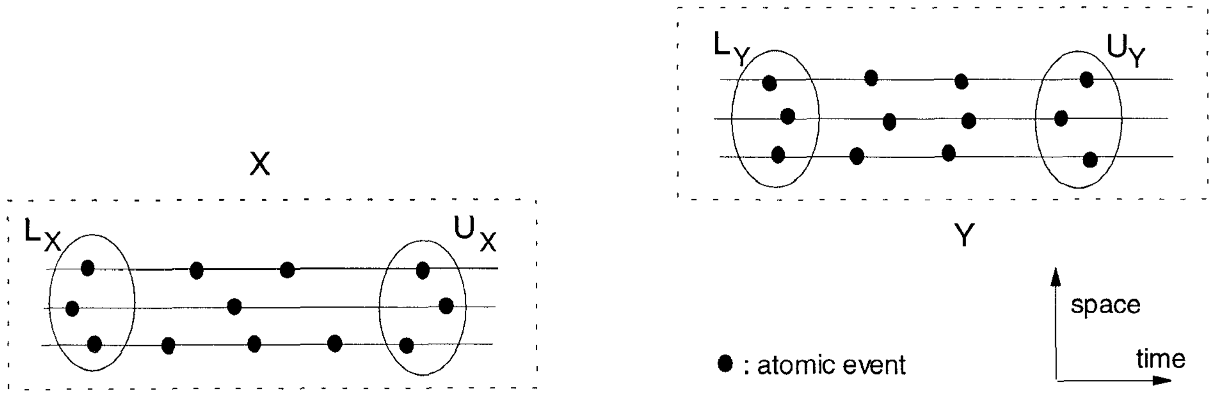

Figure 1. Poset events X and Y and their proxies.

Relation names: its quantifiers for $x \prec y$	$R1$, a (=$R1'$, a'): $\forall x \forall y$ (=$\forall y \forall x$)	$R2$, b: $\forall x \exists y$	$R2'$, b': $\exists y \forall x$	$R3$, c: $\exists x \forall y$	$R3'$, c': $\forall y \exists x$	$R4$, d (=$R4'$, d'): $\exists x \exists y$ (=$\exists y \exists x$)
$R1$, a (=$R1'$, a') : $\forall x \forall y$(= $\forall y \forall x$)	$=$	$\sqsubseteq$	$\sqsubseteq$	$\sqsubseteq$	$\sqsubseteq$	$\sqsubseteq$
$R2$, b: $\forall x \exists y$	$\sqsupseteq$	$=$	$\sqsupseteq$	$\parallel$	$\parallel$	$\sqsubseteq$
$R2'$, b': $\exists y \forall x$	$\sqsupseteq$	$\sqsubseteq$	$=$	$\parallel$	$\parallel$	$\sqsubseteq$
$R3$, c: $\exists x \forall y$	$\sqsupseteq$	$\parallel$	$\parallel$	$=$	$\sqsubseteq$	$\sqsubseteq$
$R3'$, c': $\forall y \exists x$	$\sqsupseteq$	$\parallel$	$\parallel$	$\sqsupseteq$	$=$	$\sqsubseteq$
$R4$, d (=$R4'$,d'): $\exists x \exists y$(= $\exists y \exists x$)	$\sqsupseteq$	$\sqsupseteq$	$\sqsupseteq$	$\sqsupseteq$	$\sqsupseteq$	$=$

Table 3. Full hierarchy of relations of Table 1 [10]. Relations $R1$, $R1'$, $R2$, $R2'$, $R3$, $R3'$, $R4$, $R4'$ of Table 1 are renamed a, a', b, b', c, c', d, d', respectively. Relations in the row and column headers are defined between X and Y.

for the linear interval. The proxies identify the durations on each node, in which the nonatomic event occurs.

We now define two proxies corresponding to the beginning and end of a nonatomic interval [9, 12, 13].

Definition 3 • $L_X = \{e_i \in X | \forall e_i' \in X, e_i \preceq e_i'\}$
• $U_X = \{e_i \in X | \forall e_i' \in X, e_i \succeq e_i'\}$

For any poset X, L_X and U_X are the sets of the minimal elements in X for each node and the set of the maximal elements in X for each node, respectively. L_X and U_X correspond to the beginning of the poset and the end of the poset, respectively, and can act as a *proxy* for poset X, depending on context and application. By Definition 3, each of L_X and U_X contains one event from each node in N_X.

An equally valid interpretation of the beginning and end of a poset are the sets of its minimal and maximal elements, respectively, as defined by the irreflexive partial order across the nodes. This gives an alternate definition of proxies.

Definition 4 • $L_X = \{e \in X | \forall e' \in X, e \not\succ e'\}$
• $U_X = \{e \in X | \forall e' \in X, e \not\prec e'\}$

L_X is the largest anti-chain containing the minimal elements of X. U_X is the largest anti-chain containing the maximal elements of X.

The causality relations between poset intervals are derived using proxies and depend on whether proxies are defined by Definition 3 or by Definition 4. Assume that any one of these definitions is consistently used, depending on context and application. Figure 1 depicts the proxies of X and Y and serves as a visual aid for the following discussion; recall that each poset X and Y represents a grouping of atomic events of interest to the application.

2.2 Deriving the Relations

The causality relations in [9, 12, 13] were defined using two aspects of specifying the relations. In the first aspect, a proxy needs to be chosen for X and Y; this can be done in 4 ways corresponding to relations $R1 - R4$ between linear intervals. These four relations form a lattice hierarchy ordered by "$\sqsubseteq$" ('is a subrelation of'). The second aspect of defining relations between nonatomic poset events involved defining relations between the elements of the proxies - there are 4 combinations of distinct quantifications $\exists$ and $\forall$ over the proxies of X and Y to express $r(X, Y)$, and for each combination, there are 2 permutations of the proxies of X and Y. The eight relations so formed correspond to $R1$, $R1'$, $R2$, $R2'$, $R3$, $R3'$, $R4$, $R4'$ of Table 1 and are renamed a, a', b, b', c, c', d, d', respectively, to avoid confusion with

Authorized licensed use limited to: University of Illinois at Chicago Library. Downloaded on June 22,2020 at 22:34:15 UTC from IEEE Xplore. Restrictions apply.

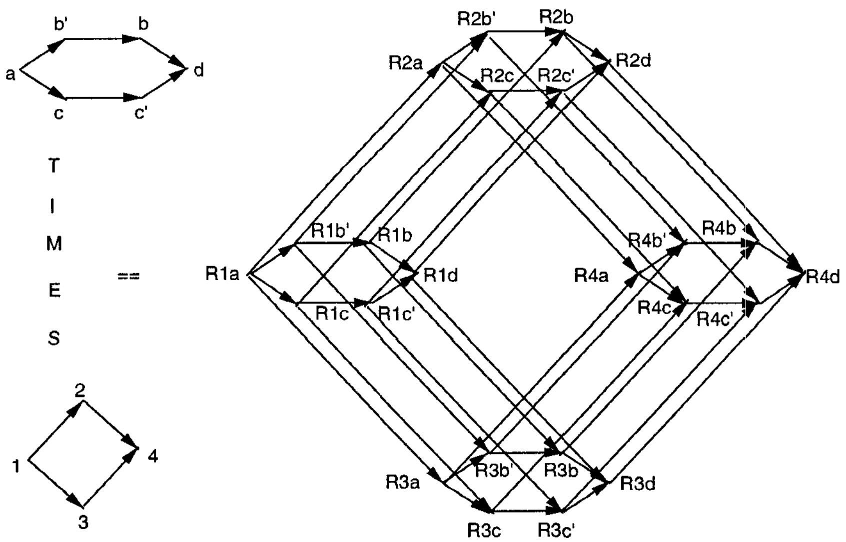

Figure 2. Hierarchy of relations in $\mathcal{R}$.

their original names used for choosing the proxies. a' and d' are the same as a and d, respectively; the six unique relations are ordered by $\sqsubseteq$, as shown in Table 3, and form a lattice hierarchy.

Each causality relation was formed by combining the two aspects of deriving causality relations described above. The relations $\{R1*, R2*, R3*, R4*\}$ between proxies for X and Y, and the relations $\{a, a', b, b', c, c', d, d'\}$ between the elements of the proxies, when multiplied give 32 relations over the domain $\mathcal{A} \times \mathcal{A}$ to express $r(X, Y)$. The resulting set of poset relations, denoted $\mathcal{R}$ and given in the second column of Table 4, forms a lattice hierarchy of 24 unique relations as shown in Figure 2. The set of relations is "complete" under first-order predicate logic and provides a fine-grained choice of causality relations.

2.3 Discussion

The set of relations [9, 12, 13] between nonatomic poset events is exhaustive using first-order predicate logic. The proposed relations form a lattice hierarchy. The strongest relation is $R1a$ and the weakest is $R4d$. The significance of a relation $R?\#(X,Y)$ is determined by examining ? for the choice of proxies of X and Y, and examining # for how these proxies are related. The proposed set of causality relations between nonatomic poset events is richer than the specific causality relations in the literature. The suite of two relations in [15], viz., $\longrightarrow$ and $--\rightarrow$, correspond to $R1a$ and $R4d$, respectively. The suite of relations in [10] and listed in Table 1 correspond to the new relations as follows: $R1{=}R1'$, $R2$, $R2'$, $R3$, $R3'$, $R4{=}R4'$ correspond to $R1a$, $R2b$, $R2b'$, $R3c$, $R3c'$, $R4d$, respectively. (This mapping is independent of whether the proxies used to derive $\mathcal{R}$ are defined by Definition 3 or 4.) The significance of the complete hierarchy of causality relations in first-order predicate logic is given in Section 4. Examples of applications that use the fine-grained relations are given in [13].

Note that by construction, $(\mathcal{R}, \sqsubseteq)$ is a partial order. For a given pair of posets X and Y, 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 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 nonatomic poset events. It follows from the result in [2], page 400, that there is a 1-1 correspondence between the set of all upward-closed subsets of a partial order and the set of anti-chains in the partial order. Therefore, an enumeration of the anti-chains in $(\mathcal{R}, \sqsubseteq)$ gives an enumeration of the upward-closed subsets of $(\mathcal{R}, \sqsubseteq)$ which correspond to all the combinations of the relations in $\mathcal{R}$ that can hold for a pair of nonatomic poset events. A recursive backtracking algorithm to enumerate the anti-chains of a poset is given in [4].

In the general case of defining causality between nonatomic events, causality between nonatomic events X and Y can be defined as "the composition of the causality relation between individual atomic events in unspecified subsets of X and Y." As applications become more sophisticated, they can use such causality relations.

3 Axiom System

The inclusion hierarchy of the relations in Table 4 is pictorally depicted in Figure 2. This hierarchy is captured by the following constraints (axioms) XP1-XP6. Let V_1 denote the set $\{1, 2, 3, 4\}$ and let V_2 denote the set $\{a, b, b', c, c', d\}$. Then the axioms are:

Authorized licensed use limited to: University of Illinois at Chicago Library. Downloaded on June 22,2020 at 22:34:15 UTC from IEEE Xplore. Restrictions apply.

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

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

Further axioms for the relations in Table 4 are derived from Tables 5, 6, 7 as follows. Table 5 is reproduced from [10] and represents the reflexivity, symmetry, and transitivity for the relations $R1$ - $R4$ defined in [10]. Table 6 is reproduced from [10] and gives the transitive axioms on the relations $R1$ - $R4$ defined in [10]. Table 7 indicates that if the proxies of X and Y in $r_1(X,Y)$ are related by the row header of the table, and if the proxies of Y and Z in $r_2(Y,Z)$ are related by the column header of the table, then the corresponding proxies of X and Z are related by the corresponding table entry; this entry is useful in deducing $r(X,Z)$. If $r_1(X,Y)$ and $r_2(Y,Z)$, then the transitive relation $r(X,Z)$ is determined by the algorithm *Trans_Poset_Axioms* using Tables 5, 6, 7 as follows.

Algorithm Trans_Poset_Axioms

1. Use the first two characters (*prefix*) of the identifier strings of $r_1(X,Y)$ and $r_2(Y,Z)$ as the inputs to Table 5 or 6. (From Table 5, $R4$ is not transitive. Hence, $R4(X,Y) \bigwedge R4(Y,Z) \Longrightarrow$ *true*.)
   $temp1$:= output of the appropriate table.
   /* $temp1$ gives the relation between X and Z if X, Y, Z were all linear intervals.*/
   If $temp1 = true$, then $r(X,Z) := true$; exit.
   /* no relation between X and Z can be inferred.*/
2. The row and column headers in Table 7 are the strings following the first two characters (*suffix*) of the identifier strings of the poset relations $\mathcal{R}$. Use the *suffixes* of $r_1(X,Y)$ and $r_2(Y,Z)$ as the row header and column header inputs, respectively, to Table 7.
   $temp2$:= output of Table 7.
   If $temp2 = true$, then $r(X,Z) := true$; exit.
   /* no relation between X and Z can be inferred.*/
3. Concatenate the values of $temp1$ and $temp2$ to get the value of $r(X,Z)$.

Example 1: If $R1c'(X,Y) \bigwedge R3b(Y,Z)$ then the algorithm yields $R1d(X,Z)$. In step 1, the inputs to Table 6 are R1 and R3, and the output $temp1$ is R1. In step 2, the inputs to Table 7 are c' and b, and its output $temp2$ is d. Step 3 concatenates $temp1$ and $temp2$ to yield $R1d$.
Example 2: If $R2a(X,Y) \bigwedge R1d(Y,Z)$ then the algorithm yields $R1b'(X,Z)$. In step 1, the inputs to Table 6 are R2 and R1, and the output $temp1$ is R1. In step 2, the inputs to Table 7 are a and d, and its output $temp2$ is b'. Step 3 concatenates $temp1$ and $temp2$ to yield $R1b'$.
Example 3: If $R3a(X,Y) \bigwedge R2b(Y,Z)$ then the algorithm yields $R4b'(X,Z)$. In step 1, the inputs to Table 6 are R3 and R2, and the output $temp1$ is R4. In step 2, the inputs to Table 7 are a and b, and its output $temp2$ is b'. Step 3 concatenates $temp1$ and $temp2$ to yield $R4b'$.
Example 4: If $R3b(X,Y) \bigwedge R2c'(Y,Z)$ then the algorithm yields $true$. In step 1, the inputs to Table 6 are R3 and R2, and the output $temp1$ is R4. In step 2, the inputs to Table 7 are b and c', and its output $temp2$ is $true$. Hence, no relation between X and Z can be inferred.

We specify the following axioms XP7-XP14 of the form $r_1(X,Y) \Longrightarrow r_2(Y,X)$ for the nonatomic poset events. For each relation $r_1(X,Y)$, we determine the strongest relation(s) $r_2(Y,X)$ that can be stated between Y and X in the hierarchy depicted in Figure 2 (Axioms XP1-XP6). Thus, given a relation between X and Y, the axioms give all possible relations between Y and X. The notation $\overline{R}$ indicates that the relation R is false. These axioms can be verified to be meaningful by examining each axiom with the aid of Figure 1 which shows X and Y in two-dimensional space-time.

XP7: $R1a(X,Y) \bigvee R1b(X,Y) \bigvee R1b'(X,Y) \bigvee R1c(X,Y) \bigvee R1c'(X,Y) \Longrightarrow \overline{R4d}(Y,X)$

XP8: $R1d(X,Y) \Longrightarrow \overline{R4b}(Y,X) \bigwedge \overline{R4c'}(Y,X)$

XP9: $R2a(X,Y) \bigvee R2b(X,Y) \bigvee R2b'(X,Y) \bigvee R2c(X,Y) \bigvee R2c'(X,Y) \Longrightarrow \overline{R2d}(Y,X)$

XP10: $R2d(X,Y) \Longrightarrow \overline{R2b}(Y,X) \bigwedge \overline{R2c'}(Y,X)$

XP11: $R3a(X,Y) \bigvee R3b(X,Y) \bigvee R3b'(X,Y) \bigvee R3c(X,Y) \bigvee R3c'(X,Y) \Longrightarrow \overline{R3d}(Y,X)$

XP12: $R3d(X,Y) \Longrightarrow \overline{R3b}(Y,X) \bigwedge \overline{R3c'}(Y,X)$

XP13: $R4a(X,Y) \bigvee R4b(X,Y) \bigvee R4b'(X,Y) \bigvee R4c(X,Y) \bigvee R4c'(X,Y) \Longrightarrow \overline{R1d}(Y,X)$

XP14: $R4d(X,Y) \Longrightarrow \overline{R1b}(Y,X) \bigwedge \overline{R1c'}(Y,X)$

In addition, we specify axiom XP15 that specifies the reflexivity and symmetry of the relations in $\mathcal{R}$.
XP15: The relations in $\mathcal{R}$ are not reflexive and are not symmetric.

$\mathcal{X}$ is the set of axioms XP1-XP6 (that specify hierarchy among relations), XP7-XP14 (that give all relations of the

Authorized licensed use limited to: University of Illinois at Chicago Library. Downloaded on June 22,2020 at 22:34:15 UTC from IEEE Xplore. Restrictions apply.

form $r_2(Y,X)$, given $r_1(X,Y)$), XP15 (that specifies reflexivity and symmetry), and the axioms that can be derived from algorithm *Trans_Poset_Axioms* to specify transitive relations. We do not attempt a completeness proof of this axiom system here. The axioms $\mathcal{X}$ provide a "sufficiently" rich framework to reason about poset intervals because:

- Axioms XP1-XP6, XP7-XP14 and XP15 give all enumerations of relations $r(X,Y)$ as well as relations $r(Y,X)$, implied by $R(X,Y)$, $\forall r \forall R \in \mathcal{R}$.
- Algorithm *Trans_Poset_Axioms* enumerates all relations $r(X,Z)$ implied by $r_1(X,Y) \bigwedge r_2(Y,Z)$, $\forall r$ $\forall r_1 \forall r_2 \in \mathcal{R}$.
- This set of axioms can be used to derive all possible implied relations from any given valid predicates on relations in $\mathcal{R}$.

Observe that depending on the choice of Definition 3 or 4 used for the proxy, there are two different sets of 32 relations $\mathcal{R}$, each of which satisfies the same set of axioms $\mathcal{X}$.

An application can specify global predicates using multiple relations from $\mathcal{R}$ between a pair of nonatomic poset events as well as between different pairs of nonatomic poset events. All the relations in $\mathcal{R}$ that hold between the involved nonatomic poset events can be inferred using the axiom system.

4 Conclusion

We examined a hierarchy of synchronization relations between nonatomic nonlinear events in a distributed system. The hierarchy of relations is complete using first-order predicate logic. We then presented an axiom system for reasoning with the proposed relations. This set of axioms can be used to derive all possible implications from any given valid predicates on the relations. The hierarchy of synchronization relations as well as the axiom system on the relations extend and complete both the hierarchy as well as the axiom system of Lamport [15], and the hierarchy and axiom system of [10], to nonatomic nonlinear events.

The results are useful for applications which use nonatomicity in reasoning and modeling and need a fine level of granularity of causality relations to specify synchronization relations and their composite global predicates. Each application can choose appropriate causality relations from the exhaustive fine-grained hierarchy to specify and capture causality and synchronization conditions between its nonatomic poset events at a fine level of granularity. The exhaustive classification gives an insight into the existing possibilities and can be used to select a number of primitive relations with good properties and clear intuitions. Examples of the use of the proposed relations by distributed real-time applications are given in [13]. The axiom system on the relations enables reasoning with different levels of causality relations between nonatomic poset events.

References

[1] U. Abraham, S. Ben-David, S. Moran, On the limitations of the global time assumption in distributed systems, *Proc. 5th Workshop on Distributed Algorithms*, LNCS 579, Springer-Verlag, 1-8, 1991.

[2] M. Aigner, *Combinatorial Theory,* Springer-Verlag, 1979.

[3] F. Anger, On Lamport's interprocessor communication model, *ACM TOPLAS*, 11(3), 404-417, 1989.

[4] M. Ball, J. S. Provan, Calculating bounds on reachability and connectedness in stochastic networks, *Networks,* 13, 253-278, 1983.

[5] *Linear Time, Branching Time, and Partial Orders in Logics and Models of Concurrency*, J. W. de Bakker, W. P. de Roever, G. Rozenberg (Eds.), LNCS 354, Springer-Verlag, 1989.

[6] J. van Benthem, *The Logic of Time*, Kluwer Academic Publishers, (1ed. 1983), 2ed. 1991.

[7] P. C. Fishburn, *Interval Orders and Interval Graphs: A Study of Partially Ordered Sets,* J. Wiley & Sons, 1985.

[8] W. Janssen, M. Poel, J. Zwiers, Action systems and action refinement in the development of parallel systems, In J.C.M. Baeten and J.F. Groote, (Eds.) *Concur'91*, LNCS 527, Springer-Verlag, 298-316, 1991.

[9] A. Kshemkalyani, Temporal interactions of intervals in distributed systems, *Technical Report TR-29.1933, IBM*, Sept. 1994.

[10] A. Kshemkalyani, Temporal interactions of intervals in distributed systems, *Journal of Computer and System Sciences*, 52(2), 287-298, April 1996. (Contains some parts of [9]).

[11] A. Kshemkalyani, Framework for viewing atomic events in distributed computations, *Theoretical Computer Science*, (in press). Abstract appears in *Proc. of Euro-Par'96*, LNCS 1123, Springer-Verlag, 496-505, Aug. 1996.

[12] A. Kshemkalyani, Relative timing constraints between complex events, *Proc. 8th IASTED Conference on Parallel and Distributed Computing Systems*, 324-326, Chicago, October 1996.

[13] A. Kshemkalyani, Synchronization for distributed real-time applications, *Proc. 5th IEEE Workshop on Parallel and Distributed Real-time Systems*, Geneva, April 1997.

[14] L. Lamport, Time, clocks, and the ordering of events in a distributed system, *C. ACM*, 558-565, 21(7), July 1978.

[15] L. Lamport, On interprocess communication, Part I: Basic formalism, Part II: Algorithms, *Distributed Computing*, 1:77-101, 1986.

[16] F. Mattern, Virtual time and global states of distributed systems, *Parallel and Distributed Algorithms*, North-Holland, 215-226, 1989.

[17] E.R. Olderog, *Nets, Terms, and Formulas*, Cambridge Tracts in Theoretical Computer Science, 1991.

[18] A. Rensink, *Models and Methods for Action Refinement,* Ph.D. thesis, University of Twente, Aug. 1993.

[19] R. Schwarz, F. Mattern, Detecting causal relationships in distributed computations: In search of the holy grail, *Distributed Computing*, 7:149-174, 1994.

[20] G. Winskel, Modeling concurrency with partial orders, *International Journal of Parallel Programming*, 15(1), 33-71, 1986.

Authorized licensed use limited to: University of Illinois at Chicago Library. Downloaded on June 22,2020 at 22:34:15 UTC from IEEE Xplore. Restrictions apply.

Relation name: its quantifiers for $x \prec y$	a (=a'): $\forall x \forall y$ (=$\forall y \forall x$)	b: $\forall x \exists y$	b': $\exists y \forall x$	c: $\exists x \forall y$	c': $\forall y \exists x$	d (=d'): $\exists x \exists y$ (=$\exists y \exists x$)
a (=a'): $\forall x \forall y (= \forall y \forall x)$	a ($\forall x \forall y$)	b' ($\exists y \forall x$)	b' ($\exists y \forall x$)	a ($\forall x \forall y$)	a ($\forall x \forall y$)	b' ($\exists y \forall x$)
b: $\forall x \exists y$	a ($\forall x \forall y$)	b ($\forall x \exists y$)	b' ($\exists y \forall x$)	$true$	$true$	$true$
b': $\exists y \forall x$	a ($\forall x \forall y$)	b' ($\exists y \forall x$)	b' ($\exists y \forall x$)	$true$	$true$	$true$
c: $\exists x \forall y$	c ($\exists x \forall y$)	d ($\exists x \exists y$)	d ($\exists x \exists y$)	c ($\exists x \forall y$)	c ($\exists x \forall y$)	d ($\exists x \exists y$)
c': $\forall y \exists x$	c ($\exists x \forall y$)	d ($\exists x \exists y$)	d ($\exists x \exists y$)	c ($\exists x \forall y$)	c'($\forall y \exists x$)	d ($\exists x \exists y$)
d (=d'): $\exists x \exists y (= \exists y \exists x)$	c ($\exists x \forall y$)	d ($\exists x \exists y$)	d ($\exists x \exists y$)	$true$	$true$	$true$

Table 7. Intermediate table to derive further axioms for poset relations $\mathcal{R}$. The relation names in the row and column headers are the suffixes of the poset relations $\mathcal{R}$ defined between X and Y.

Relation $r(X,Y)$	Relation definition specified by quantifiers for $x \prec y$, where $x \in X$, $y \in Y$
$R1a$ $R1a'$ (=$R1a$)	$\forall x \in U_X \forall y \in L_Y$ $\forall y \in L_Y \forall x \in U_X$
$R1b$	$\forall x \in U_X \exists y \in L_Y$
$R1b'$	$\exists y \in L_Y \forall x \in U_X$
$R1c$	$\exists x \in U_X \forall y \in L_Y$
$R1c'$	$\forall y \in L_Y \exists x \in U_X$
$R1d$ $R1d'$ (=$R1d$)	$\exists x \in U_X \exists y \in L_Y$ $\exists y \in L_Y \exists x \in U_X$
$R2a$ $R2a'$ (=$R2a$)	$\forall x \in U_X \forall y \in U_Y$ $\forall y \in U_Y \forall x \in U_X$
$R2b$	$\forall x \in U_X \exists y \in U_Y$
$R2b'$	$\exists y \in U_Y \forall x \in U_X$
$R2c$	$\exists x \in U_X \forall y \in U_Y$
$R2c'$	$\forall y \in U_Y \exists x \in U_X$
$R2d$ $R2d'$ (=$R2d$)	$\exists x \in U_X \exists y \in U_Y$ $\exists y \in U_Y \exists x \in U_X$
$R3a$ $R3a'$ (=$R3a$)	$\forall x \in L_X \forall y \in L_Y$ $\forall y \in L_Y \forall x \in L_X$
$R3b$	$\forall x \in L_X \exists y \in L_Y$
$R3b'$	$\exists y \in L_Y \forall x \in L_X$
$R3c$	$\exists x \in L_X \forall y \in L_Y$
$R3c'$	$\forall y \in L_Y \exists x \in L_X$
$R3d$ $R3d'$ (=$R3d$)	$\exists x \in L_X \exists y \in L_Y$ $\exists y \in L_Y \exists x \in L_X$
$R4a$ $R4a'$ (=$R4a$)	$\forall x \in L_X \forall y \in U_Y$ $\forall y \in U_Y \forall x \in L_X$
$R4b$	$\forall x \in L_X \exists y \in U_Y$
$R4b'$	$\exists y \in U_Y \forall x \in L_X$
$R4c$	$\exists x \in L_X \forall y \in U_Y$
$R4c'$	$\forall y \in U_Y \exists x \in L_X$
$R4d$ $R4d'$ (=$R4d$)	$\exists x \in L_X \exists y \in U_Y$ $\exists y \in U_Y \exists x \in L_X$

Table 4. Relations $r(X,Y)$ in $\mathcal{R}$ from [12, 13].

Relation	reflexive ?	symmetric ?	transitive ?
R1 [15]	no	no	yes
R2	no	no	yes
R3	no	no	yes
R4 [15]	no	no	no

Table 5. Reflexivity, symmetry and transitivity of R1, R2, R3, R4 from [10].

Axiom Label	$r_1(X,Y) \bigwedge r_2(Y,Z) \Longrightarrow r(X,Z)$
AL1	$R1(X,Y) \bigwedge R2(Y,Z) \Longrightarrow R2(X,Z)$
AL2	$R1(X,Y) \bigwedge R3(Y,Z) \Longrightarrow R1(X,Z)$
AL3	$R1(X,Y) \bigwedge R4(Y,Z) \Longrightarrow R2(X,Z)$
AL4	$R2(X,Y) \bigwedge R1(Y,Z) \Longrightarrow R1(X,Z)$
AL5	$R3(X,Y) \bigwedge R1(Y,Z) \Longrightarrow R3(X,Z)$
AL6	$R4(X,Y) \bigwedge R1(Y,Z) \Longrightarrow R3(X,Z)$
AL7	$R2(X,Y) \bigwedge R3(Y,Z) \Longrightarrow true$
AL8	$R2(X,Y) \bigwedge R4(Y,Z) \Longrightarrow true$
AL9	$R3(X,Y) \bigwedge R2(Y,Z) \Longrightarrow R4(X,Z)$
AL10	$R4(X,Y) \bigwedge R2(Y,Z) \Longrightarrow R4(X,Z)$
AL11	$R3(X,Y) \bigwedge R4(Y,Z) \Longrightarrow R4(X,Z)$
AL12	$R4(X,Y) \bigwedge R3(Y,Z) \Longrightarrow true$

Table 6. Axioms for causality relations R1, R2, R3, R4 from [10].

Authorized licensed use limited to: University of Illinois at Chicago Library. Downloaded on June 22,2020 at 22:34:15 UTC from IEEE Xplore. Restrictions apply.