A Note on Modalities for Nonconjunctive Global Predicates

Ajay D. Kshemkalyani

Computer Science Department, Univ. of Illinois at Chicago
Chicago, IL 60607, USA
ajayk@cs.uic.edu

Abstract. Global predicate detection is an important problem in distributed executions. A conjunctive predicate is one in which each conjunct is defined over variables local to a single process. Polynomial space and time algorithms exist for detecting conjunctive predicates under the *Possibly* and *Definitely* modalities, as well as under a richer class of fine-grained modalities based on the temporal interaction of intervals. However, it is much more difficult to detect nonconjunctive predicates under the various modalities because the entire state lattice may need to be examined. We examine the feasibility of detecting nonconjunctive predicates under the fine-grained temporal modalities using the interval-based approach. We gain some insightful observations into how nonconjunctive predicates can be decomposed, and into the relationships among the intervals (at different processes) in which the local variables have values that can satisfy the nonconjunctive predicate.

1 Introduction

Predicate detection in a distributed system is important for various purposes such as debugging, monitoring, synchronization, and industrial process control [5, 6, 7, 8, 10, 12, 14]. Marzullo et al. defined two modalities under which predicates can hold for a distributed execution [5, 12].

- $Possibly(\phi)$: There exists a consistent observation of the execution such that ϕ holds in a global state of the observation.
- $Definitely(\phi)$: For every consistent observation of the execution, there exists a global state of it in which ϕ holds.

For any predicate ϕ, three orthogonal relational possibilities hold: (i) $Definitely(\phi)$, (ii) $\neg Definitely(\phi) \wedge Possibly(\phi)$, (iii) $\neg Possibly(\phi)$ [10]. The orthogonal set $\Re$ of 40 fine-grained temporal interactions between any pair of intervals [9] provides the basis for more expressive power than the *Possibly* and *Definitely* modalities for specifying any predicate. A mapping from the fine-grained interactions to the *Possibly* and *Definitely* modalities was given in [10].

A conjunctive predicate is of the form $\bigwedge_i \phi_i$, where ϕ_i is a predicate defined on variables local to process P_i, e.g., $x_i = 3 \wedge y_j > 20$, where x_i and y_j are local to P_i and P_j, resp.. Conjunctive predicates form an important class of

S. R. Das, S. K. Das (Eds.): IWDC 2003, LNCS 2918, pp. 11–25, 2003.

predicates. They have been studied in [2, 3, 6, 7, 8, 14]. The following result was shown in [10].

Theorem 1. *[10] For a conjunctive predicate $\phi = \wedge_{i \in N} \phi_i$, let ϕ_i denote the component of ϕ local to process P_i and let N denote the set of processes. The following results are implicitly qualified over a set of intervals, containing one interval from each process.*

- *$Definitely(\phi)$ holds if and only if $\bigwedge_{(\forall i \in N)(\forall j \in N)} Definitely(\phi_i \wedge \phi_j)$*
- *$\neg Definitely(\phi) \wedge Possibly(\phi)$ holds if and only if*
  - *$(\exists i \in N)(\exists j \in N) \neg Definitely(\phi_i \wedge \phi_j) \bigwedge (\bigwedge_{(\forall i \in N)(\forall j \in N)} Possibly(\phi_i \wedge \phi_j))$*
- *$\neg Possibly(\phi)$ holds if and only if $(\exists i \in N)(\exists j \in N) \neg Possibly(\phi_i \wedge \phi_j)$*

By Theorem 1, given a conjunctive predicate ϕ defined on any number of processes, $Definitely(\phi)$ and $Possibly(\phi)$ can be expressed in terms of the *Possibly* and *Definitely* modalities on predicates defined over all pairs of processes. Also, from the results in [10], the *Possibly* and *Definitely* modalities on any general predicates defined over a pair of processes have been mapped into the set $\Re$ of 40 orthogonal modalities between appropriately identified intervals on the two processes. Therefore, $Definitely(\phi)$ and $Possibly(\phi)$, where (conjunctive predicate) ϕ is defined over any number of processes, can be expressed in terms of the fine-grained orthogonal set $\Re$ of modalities over predicates defined over all pairs of processes.

Polynomial space and time algorithms exist for detecting conjunctive predicates under the *Possibly* and *Definitely* modalities [5, 12]. We have also designed polynomial complexity algorithms, to detect not just *Possibly* and *Definitely*, but also the exact fine-grained relation between each pair of processes when *Possibly* and *Definitely* are true [2]. Two factors make the problem of detecting conjunctive predicates, even under the harder fine-grained modalities, solvable with polynomial complexities. First, each process can locally determine the local *intervals* or *durations* in which the local predicate is true. Second, as a result of the first factor, each process can identify alternating intervals when the truth value of the local predicate alternates. *Nonconjunctive* predicates are of the form $(x_i + y_j = 5) \wedge (x_i + z_k = 10)$ and $x_i + y_j + z_k = 10$. It is much more difficult to detect nonconjunctive predicates because the above two factors do not hold and the entire state execution lattice may need to be examined, leading to exponential complexity [1, 5, 12, 13].

The use of the interval-based approach to specify and detect conjunctive predicates under the rich class of modalities $\Re$ [2] prompts us to examine the feasibility of a similar interval-based approach for nonconjunctive predicates.

The following can be seen from [10]. (1) For a pair of processes, the mapping from the fine-grained set of interactions $\Re$ to the *Possibly*/*Definitely* classification depends only on the intervals, and is independent of the "predicate type". (2) For more than 2 processes, a result similar to Theorem 1 can hold for nonconjunctive predicates *provided* the intervals can be first identified appropriately. The semantics of the intervals need to be defined and identified carefully for nonconjunctive predicates because the entire state lattice may need to be examined. In this paper, we examine how nonconjunctive predicates can be specified/ detected under various fine-grained modalities using the interval-based approach. We introduce the notion of a *composite interval* in an attempt to identify intervals at each process. We gain some insightful observations into how nonconjunctive predicates can be decomposed, and into the relationships among the intervals (at different processes) in which the local variables have values that can satisfy the nonconjunctive predicate under various modalities.

Section 2 gives the execution model. Section 3 states the objectives precisely and introduces the solution approach. Section 4 summarizes the results for conjunctive predicates. Section 5 presents the interval-based analysis for nonconjunctive predicates, and makes the main observations. Section 6 discusses the notion of minimal intervals for nonconjunctive predicates when detecting the $Definitely$ modality. Section 7 gives the conclusions.

2 System Model and Background

We assume an asynchronous distributed system in which $n = |N|$ processes communicate by reliable message passing. To model the system execution, let $\prec$ be an irreflexive partial ordering representing the causality relation on the event set E. E is partitioned into local executions at each process. Each E_i is a linearly ordered set of events executed by process P_i. An event e at P_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 [11].

A *cut* C is a subset of E such that if $e_i \in C$ then $(\forall e'_i) e'_i \prec e_i \implies e'_i \in C$. A *consistent cut* is a downward-closed subset of E in $(E, \prec)$ and denotes an execution prefix. The system state after the events in a cut is a global state; if the cut is consistent, the corresponding system state is a consistent global state. Each total ordering of $(E, \prec)$ is a *linear extension* that represents the global time ordering of events and globally observed states in some equivalent (isomorphic) execution. The global time interleaving of events is different in each such isomorphic execution, but all these executions have the same partial order. The *state lattice* of an execution represents all possible global states that can occur. There is a bijective mapping between the set of all paths in the state lattice and the set of all linear extensions of the execution, for a given execution. We assume that only consistent global states are included in the state lattice.

An interval at process P_i is identified by the (totally ordered) subset of adjacent events of E_i that occur in that interval. An interval of interest at a process is a duration in which the local predicate is true (for conjunctive predicates) or in which the local values may potentially satisfy the global predicate (for nonconjunctive predicates). Henceforth, unless otherwise specified, references to intervals will implicitly be to intervals of interest. For a nonconjunctive predicate, the intervals need to be identified carefully, based on appropriate semantics.

3 Objectives and Approach

In this paper, we examine how nonconjunctive predicates can be specified and detected under the various fine-grained modalities using the interval-based approach. We ask the following questions.

1. What semantics can be used to identify the intervals at each process when ϕ is a nonconjunctive predicate, in order to apply the interval-based approach [4, 10]? An example predicate is $x_i - y_j < 10$.
2. How can Theorem 1 be extended to nonconjunctive predicates? When ϕ is defined over more than two processes, can it be reexpressed in terms of predicates over pairs of processes such that Theorem 1 can then be used in some form? Example predicates are $(x_i + y_j = 5) \wedge (x_i + z_k = 10)$ and $x_i + y_j + z_k = 10$.

Knowing the inherent difficulty in dealing with nonconjunctive predicates, we do not expect startling new results. Rather, we hope to make some insightful observations into how nonconjunctive predicates can be decomposed, and into the relationships among the intervals (at different processes) where the local variables have values that can satisfy the nonconjunctive predicate. By answering these questions, we can specify/detect not just the $Possibly$ and $Definitely$ modalities but also the fine-grained modalities of $\Re$, on nonconjunctive predicates.

Based on Theorem 1 for conjunctive predicates, we reexpress the definition of $Possibly$ and $Definitely$ in terms of intervals when $n > 2$ processes.

Definition 1. *Let $\mathcal{I}$ be a set of intervals, containing one interval per process, such that during these intervals, the local predicates are true, (or more generally, the local variables using which global predicate ϕ is defined have values that may satisfy ϕ).*

- *$Possibly(\phi)$: (For some set $\mathcal{I}$ of intervals,) there exists a linear extension of $(E, \prec)$ such that for each pair of intervals X and Y in $\mathcal{I}$, the string $[min(X), max(X)]$ overlaps with the string $[min(Y), max(Y)]$.*
- *$Definitely(\phi)$: (For some set $\mathcal{I}$ of intervals,) for every linear extension of $(E, \prec)$, for each pair of intervals X and Y in $\mathcal{I}$, the string $[min(X), max(X)]$ overlaps with the string $[min(Y), max(Y)]$.*

This alternate definition cannot be applied to nonconjunctive predicates unless the semantics of the interval at each process is known, and the intervals can be identified somehow. To understand how the differences between conjunctive and nonconjunctive predicates affect the identification of intervals at each process, we first identify and discuss the salient features that lead to such differences.

Method of decomposing predicate. For conjunctive predicates, the global predicate can be simply decomposed as the conjunct of the local predicates. For nonconjunctive predicates, there are several choices for decomposing the global predicate. The most natural form is the disjunctive normal form

(DNF) where each disjunct is a conjunct over variables at all processes, and each disjunct can be satisfied by a set of "subintervals", one per process. Even a simple predicate defined over two processes, like $x_i + y_j = 10$ or $x_i + y_j > 10$ can lead to an infinite number of disjuncts. As one of our goals is to determine how a predicate over more than two processes can be reexpressed in terms of predicates over pairs of processes at a time, this reexpression needs to be done carefully.

Adjacency of local intervals. For conjunctive predicates, each local interval can be locally determined to be the maximum duration in which the local predicate is *true*, irrespective of values of the variables used to define the local predicate. Two such intervals can never be adjacent (otherwise they would form a larger interval). For nonconjunctive predicates, two adjacent intervals may be such that the values of the local variables can potentially satisfy global predicate ϕ. For the predicate $x_i + y_j = 10$, $x_i = 3 \wedge y_j = 7$ and $x_i = 4 \wedge y_j = 6$ correspond to two different disjuncts in the DNF expression. As such, the intervals in which x_i is 3 and in which x_i is 4 may be adjacent.

Composite intervals. A *composite interval* is defined as an interval containing multiple adjoining intervals which we term as *subintervals*. In each of the subintervals, the local variables using which the global predicate is defined may or may not satisfy that predicate under the specified modality, depending on which subintervals at other processes these subintervals overlap.
Composite intervals are not relevant to conjunctive predicates because in any interval in which the local predicate is true, the varying values of local variables do not matter. For nonconjunctive predicates, composite predicates are relevant because different sets of subintervals, each set containing one subinterval from each process, can cause predicate ϕ (or some disjunct(s) of ϕ when it is expressed in DNF) to be *true* in the desired modality.

We analyze the above features for the 8 combinations obtained by the choices: (i) conjunctive or nonconjunctive ϕ, (ii) *Possibly* or $Definitely$ modalities, and (iii) ϕ being defined on two or more than two processes. This gives a better insight into the use of the interval-based approach for detecting nonconjunctive predicates under the various modalities of $\Re$.

4 Conjunctive Predicates

There are 4 independent combinations to consider for conjunctive predicates.

- $Possibly(\phi)$, conjunctive predicate, $n = 2$. If $Possibly(\phi)$ holds, there is some linear extension in which some pair of intervals X and Y overlap.
- $Definitely(\phi)$, conjunctive predicate, $n = 2$. If $Definitely(\phi)$ holds, a common pair of intervals X and Y overlap for each linear extension. If X overlaps Y in only some linear extensions and a disjoint interval X' overlaps Y in only all the other linear extensions, then from the properties of linear extensions, there must exist a linear extension in which neither X nor X' overlaps Y [10].

- $Possibly(\phi)$, conjunctive predicate, $n > 2$. As shown in [10], there is some linear extension in which some set of intervals, containing one interval per process, overlap pairwise.
- $Definitely(\phi)$, conjunctive predicate, $n > 2$. As shown in [10], a common set of intervals, containing one interval per process, overlap pairwise for each linear extension.

For all these four cases, we have the following results using our approach.

1. **Decomposing global predicate**. In CNF, each conjunct is defined on variables local to a single process.
2. **Adjacency of local intervals**. Two intervals cannot be adjacent at a process.
3. **Composite intervals**. Not relevant because intervals cannot be adjacent at a process.

5 Nonconjunctive Predicates and Intervals

5.1 Nonconjunctive Predicates on Two Processes

$Possibly(\phi)$, nonconjunctive predicate, $n = 2$. For some linear extension, the local interval at one process overlaps the local interval at the other process and the values of the local variables in these intervals satisfy the predicate.

1. **Decomposing global predicate**. The predicate can be reexpressed in DNF. For the *Possibly* modality to be satisfied, it is sufficient if any one disjunct is satisfied in some linear extension.
2. **Adjacency of local intervals**. Intervals can be adjacent because when the predicate is expressed in DNF, two adjacent intervals at a process may satisfy two different disjuncts.
3. **Composite intervals**. Even though local intervals may be adjacent, it is not necessary to consider composite intervals because for the *Possibly* modality, any one interval (without subintervals) at a process may be used to satisfy the modality of the predicate.

Once the local intervals, one per process, are identified, the fine-grained modality for *Possibly* can be determined by using the tests and mappings from [10].

$Definitely(\phi)$, nonconjunctive predicate, $n = 2$. The predicate can be reexpressed in DNF. For the *Definitely* modality to be satisfied, it is sufficient if for every linear extension, some disjunct is satisfied. As it is sufficient that different disjuncts can be satisfied, it is necessary to consider composite intervals. **Examples 1(a,b):** Examples of composite intervals are given in Figure 1. Consider the predicate $Definitely(x = y)$. Two slightly differing executions are shown. The state lattices are labeled using event numbers at the two processes. In Figure 1(a), $Definitely(x = 1 \wedge y = 1)$ is *false*. Also, $Definitely(x = 2 \wedge y = 2)$

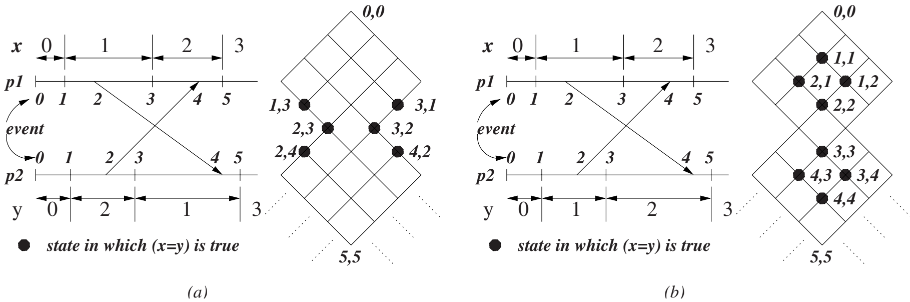

Fig. 1. Examples to show composite intervals for nonconjunctive predicates, for $n = 2$ processes

is *false*. However consider the following two composite intervals at P_1 and P_2, respectively. At P_1, the composite interval contains the subintervals when $x = 1$ and $x = 2$ in that order. At P_2, the composite interval contains the subintervals when $y = 2$ and $y = 1$ in that order. From the state lattice of this execution, observe that $Definitely(x = y)$ is *true* because exactly one of the following two states must hold in every execution: $x = 1 \wedge y = 1$, or $x = 2 \wedge y = 2$. The execution in Figure 1(b) differs only in that y first takes the value of 1 and then the value of 2. For this execution, the states in which $x = y$ are marked in the corresponding state lattice diagram. As can be observed from the lattice, $Definitely(x = y)$ is *false*. This example shows that the composite intervals need to be identified carefully after examining the state lattice.

Examples 2(a,b): Figure 2 shows two example executions with their corresponding state lattice diagrams. There are no messages exchanged in these executions. The state lattices are labeled so as to show only the values of the variables x and y. To detect $Definitely(x = y)$, observe that one has to seek recourse to examining the state lattice, and then determining the composite intervals. In Figure 2(a), $Definitely(x = y)$ is *true*. At P_1, the composite interval contains the subintervals when $x = 1, 2, 3, 4$. At P_2, the composite interval contains the subintervals when $y = 4, 3, 2, 1$. In Figure 2(b), $Definitely(x = y)$ is *false*.

These simple examples indicate that it is necessary to consider the state lattice to determine the composite intervals for the $Definitely$ modality.

1. **Decomposing global predicate**. The predicate can be reexpressed in DNF. For the $Definitely$ modality to be satisfied, it is sufficient if for every linear extension, some disjunct is satisfied.
2. **Adjacency of local intervals**. Intervals can be adjacent because when the predicate is expressed in DNF, two adjacent intervals at a process may satisfy two different disjuncts.

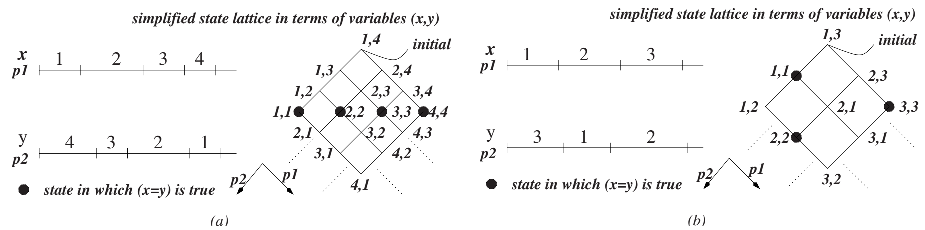

Fig. 2. Further examples for nonconjunctive predicates, for $n = 2$ processes

3. **Composite intervals**. As it is sufficient that different disjuncts can be satisfied in different linear extensions, it is necessary to consider composite intervals.

The identification of the semantics of the intervals for $Definitely(\phi)$ and the exact fine-grained relation(s) between the intervals will be dealt with in Section 6.

5.2 Nonconjunctive Predicates on More than Two Processes

The predicate can be reexpressed in DNF. For the $Definitely$ modality and $n = 2$, it was necessary to consider composite intervals, as also illustrated in Examples 1 and 2. For the $Definitely$ modality and $n > 2$, the same reasoning shows that because it is sufficient that different disjuncts can be satisfied in different linear extensions, therefore it is necessary to consider composite intervals. As the fine-grained modalities $\Re$ are expressed on processes pairwise, let us first try to adapt Definition 1 to nonconjunctive predicates. For $n > 2$, multiple pairs of processes need to be considered. In every linear extension, when the composite intervals pairwise overlap, only some combination(s) of the subintervals will actually overlap. For each pair of processes, there will exist multiple pairs of subintervals. However, if for each pair of processes, one of the pairs of subintervals overlap in each linear extension, that does not imply that there is a set of subintervals, one from each process, that will *collectively* overlap with each other in every linear extension. Thus for $Definitely(\phi)$, reexpressing ϕ as a conjunction of predicates on pairwise processes, the following *incorrect* Definition 2 would result.

Definition 2. *(Incorrect definition of $Definitely(\phi)$ when reexpressing ϕ as conjunction of predicates on pairwise processes:) (For some set $\mathcal{I}$ of composite intervals), for every linear extension, for each pair of composite intervals X and Y, one of the subinterval strings $[min(X'_i), max(X'_i)]$ overlaps with one of the subinterval strings $[min(Y'_j), max(Y'_j)]$.*

The above definition does not guarantee that if $[min(X'_i), max(X'_i)]$ overlaps with $[min(Y'_j), max(Y'_j)]$, and $[min(Y'_j), max(Y'_j)]$ overlaps with $[min(Z'_k),$-

$max(Z'_k)]$, then $[min(X'_i), max(X'_i)]$ overlaps with $[min(Z'_k), max(Z'_k)]$; instead, $[min(X''_i), max(X''_i)]$ may overlap with $[min(Z'_k), max(Z'_k)]$.

An argument analogous to the above shows that the same conclusion holds for $Possibly(\phi)$ if ϕ is reexpressed as a conjunction of predicates on pairwise processes, as in Definition 3.

Definition 3. *(Incorrect definition of $Possibly(\phi)$ when reexpressing ϕ as conjunction of predicates on pairwise processes:) (For some set $\mathcal{I}$ of composite intervals), for some linear extension, for each pair of composite intervals X and Y, one of the subinterval strings $[min(X'_i), max(X'_i)]$ overlaps with one of the subinterval strings $[min(Y'_j), max(Y'_j)]$.*

This unfortunately implies two negative results.

– The manner in which the global predicate is decomposed pairwise over processes must ensure that the same subinterval at any process is considered when determining overlaps with subintervals at other processes.
  Consider the predicate $x_i = y_j = z_k$. In DNF, this would be expressed as

$$(x_i = 1 \wedge y_j = 1 \wedge z_k = 1) \vee (x_i = 2 \wedge y_j = 2 \wedge z_k = 2) \vee \ldots\ldots.$$

  If this is reexpressed in terms of predicates over pairs of processes, as

$$((x_i = 1 \wedge y_j = 1) \wedge (x_i = 1 \wedge z_k = 1) \wedge (y_j = 1 \wedge z_k = 1)) \bigvee$$

$$((x_i = 2 \wedge y_j = 2) \wedge (x_i = 2 \wedge z_k = 2) \wedge (y_j = 2 \wedge z_k = 2)) \bigvee \ldots$$

  then care must be taken to consider each disjunct separately.
  Example 3(a): To detect ψ: $Definitely(x = y = z)$, if the predicate were reexpressed by splitting pairwise as per Definition 2, as

$$\psi' : Definitely(x = y) \wedge Definitely(y = z) \wedge Definitely(x = z),$$

  then ψ' would be *true* in the execution of Figure 3(a). However, ψ' is not equivalent to ψ, and in this example, ψ is *false*.
  Example 3(b): To detect ψ: $Possibly(x = y = z)$, if the predicate were reexpressed by splitting pairwise as per Definition 3, as

$$\psi' : Possibly(x = y) \wedge Possibly(y = z) \wedge Possibly(x = z),$$

  then ψ' would be *true* in the execution of Figure 3(b). $Possibly(x = y = z = 1)$ is *false*; $Possibly(x = y = z = 2)$ is also *false*. ψ' is not equivalent to ψ.
  Definition 3 implicitly assumed that composite intervals with their subintervals are considered. But for the *Possibly* modality, it is sufficient to ensure that *only one* subinterval from each process is considered, when checking for pairwise overlap (matching) between each pair of processes. However, to determine which subinterval from a process should be considered, one has to inevitably consider the global predicate (i.e., each disjunct of the global predicate) defined across all the processes, and hence the global state lattice. This suggests that the use of composite intervals with their subintervals is not useful for detecting *Possibly* modality.

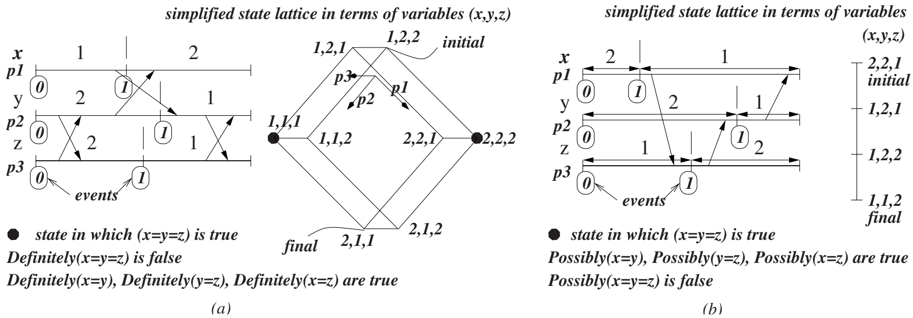

Fig. 3. Examples to show composite intervals, for $n = 3$ processes. Simplified state lattices that indicate only the values of the variables x, y, and z are shown for each execution. (a): $Definitely(x = y = z)$ is *false* although $Definitely(x = y)$, $Definitely(y = z)$, and $Definitely(x = z)$ are *true*. (b): $Possibly(x = y = z)$ is *false* although $Possibly(x = y)$, $Possibly(y = z)$, and $Possibly(x = z)$ are *true*

- For the $Definitely$ modality, one needs to inevitably examine the state lattice to determine whether multiple such sets of subintervals (one subinterval per process in each set) exist such that in each linear extension, there is mutual overlap between each pair of subintervals in at least one such set of subintervals. The subintervals at a process identify the composite interval at that process.
  Example 4: Figure 4 shows an example execution along with its state lattice diagram labeled using variable values. The predicate of interest here is $Definitely(x + y + z = 3)$. This is *true*, but can be determined only by examining the lattice and observing that each execution must necessarily pass through one of the 10 states marked in the lattice.

***Possibly*(ϕ), nonconjunctive predicate, $n > 2$.**

1. **Decomposing global predicate**. The predicate can be reexpressed in DNF. For the *Possibly* modality to be satisfied, it is sufficient if for *some* linear extension, *some* disjunct is satisfied. Consider each (instantiated) disjunct separately. If some disjunct is *true* in some global state, that is equivalent to there being pairwise overlap between the intervals, one per process, for which the local variable values in that state hold. $Possibly(\phi)$ is *true*, and the interval at each process that can satisfy that disjunct is identified by the duration in which the local variable values of that state persist.
   In this analysis, it is essential that ϕ must not be reexpressed as a conjunction of predicates on pairwise processes.
2. **Adjacency of local intervals**. Can be adjacent.

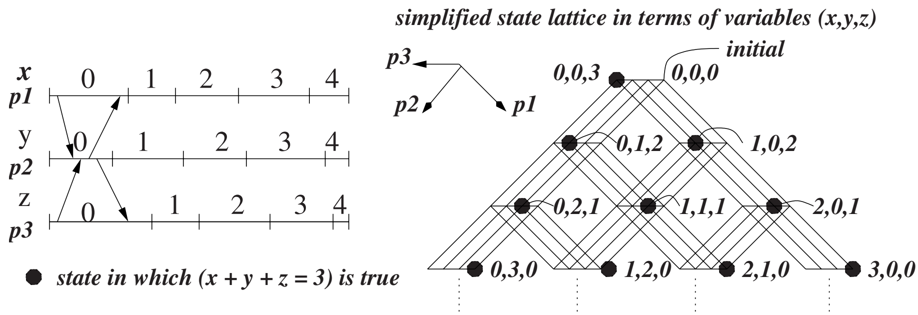

Fig. 4. Example to show composite intervals for nonconjunctive predicate $Definitely(x + y + z = 3)$, for $n = 3$ processes

3. **Composite intervals**. Not necessary and not useful.

Once the intervals (simple, not composite) are identified, Theorem 1 which was for conjunctive predicates can also be used along with the mapping from [10] to find the fine-grained modality for each pair of intervals corresponding to the disjunct that is satisfied.

$Definitely(\phi)$, nonconjunctive predicate, $n > 2$.

1. **Decomposing global predicate**. The predicate can be reexpressed in DNF. For the $Definitely$ modality to be satisfied, it is sufficient if for *every* linear extension, *some* disjunct is satisfied. This requires the state lattice to be examined.
2. **Adjacency of local intervals**. Can be adjacent.
3. **Composite intervals**. As in the case for $n = 2$, it is sufficient that different disjuncts can be satisfied in different linear extensions. Hence, it is necessary to consider composite intervals.

Section 6 shows how to identify the intervals and the fine-grained modalities between the intervals, for the $Definitely$ modality.

6 Minimal Intervals for $Def(\phi)$, for Nonconjunctive ϕ

For a conjunctive predicate, the interval at each process can span a duration in which local variables may take on multiple values, all of which satisfy the local predicate. More importantly, each interval can be determined locally. The following observation about the local intervals, one at each process, can be made. For each process P_i, let there be some contiguous range of events $D_i = [e_i^{x_i}, e_i^{x'_i}]$, such that the global predicate is *true* in at least one global state containing the local state after *each* event in D_i. Then the global predicate is necessarily *true* in *every* global state such that for each P_i, the local state is the state after some event in D_i. This is due to the conjunctive nature of the predicate. In terms of the state lattice, these states form a dense "convex" region. If $Definitely(\phi)$ is *true*, then every path in the lattice must pass through this region.

This observation can now be formally extended to nonconjunctive predicates.

Definition 4. (States in a sublattice in which ϕ is satisfied:)

1. *Let the set of states in the sublattice defined by the various D_i, one per process, be denoted $\prod_i D_i$.*
2. *Given the set of states $\prod_i D_i$ in a sublattice, the subset of states in which ϕ is* true *is denoted $S(\prod_i D_i)$. In general, $S(\prod_i D_i) \subseteq \prod_i D_i$ but for conjunctive predicates, $S(\prod_i D_i) = \prod_i D_i$.*

Thus, $S(\prod_i D_i)$ represents those states in the "convex" region in the lattice, where ϕ is *true*. Any equivalent execution must pass through at least one state in $S(\prod_i D_i)$ if $Definitely(\phi)$ is *true*.

Example 5(a): Consider the state lattice of an execution, shown in Figure 5(a). If the predicate ϕ is *true* in the states marked, and if ϕ is conjunctive, then ϕ is also necessarily *true* in all (consistent) states (v, w), where $1 \leq v \leq 7$ and $1 \leq w \leq 6$. Hence, $D_1 = [e_1^1, e_2^7]$ and $D_2 = [e_2^1, e_2^6]$. However, if ϕ is nonconjunctive, then the predicate can be *true* in only the 11 states $S(D_1 \times D_2)$ marked.

If $Definitely(\phi)$ holds for a nonconjunctive predicate, the identification of intervals at each process is useful for determining the fine-grained modality between each pair of processes. The sequence of events D_i identifies the interval at process P_i, provided that property `DEF-SUBLATTICE` is satisfied.

Property 1. (Property `DEF-SUBLATTICE(`$\prod_i D_i$`):`) Every equivalent execution must pass through at least one state in which ϕ is *true*, among the states in $\prod_i D_i$.

The pairwise orthogonal relation between the intervals can be specified by considering the set of intervals identified by each D_i, one per process P_i. Observe that these intervals can be refined further to get more specific information on the fine-grained modalities that can hold when $Definitely(\phi)$ is *true*.

Example 5(a) contd.: ϕ is *true* in only 11 of the 42 states of the sublattice $D_1 \times D_2$. Each of these 11 states corresponds to a potentially different disjunct.

For each possible disjunct of ϕ, one of which will necessarily become *true* in each equivalent execution through $\prod_i D_i$, the pairwise orthogonal relations can be determined by executing `DISJUNCT-FINE-GRAIN`($S(\prod_i D_i)$), shown in Figure 6.

The smaller the set of states $S(\prod_i D_i)$, the smaller is $\mathcal{O}^*$ likely to be, and the more precise the information about the fine-grained modalities. *Minimal* intervals are more useful because they can more accurately pinpoint the possible pairwise orthogonal relations.

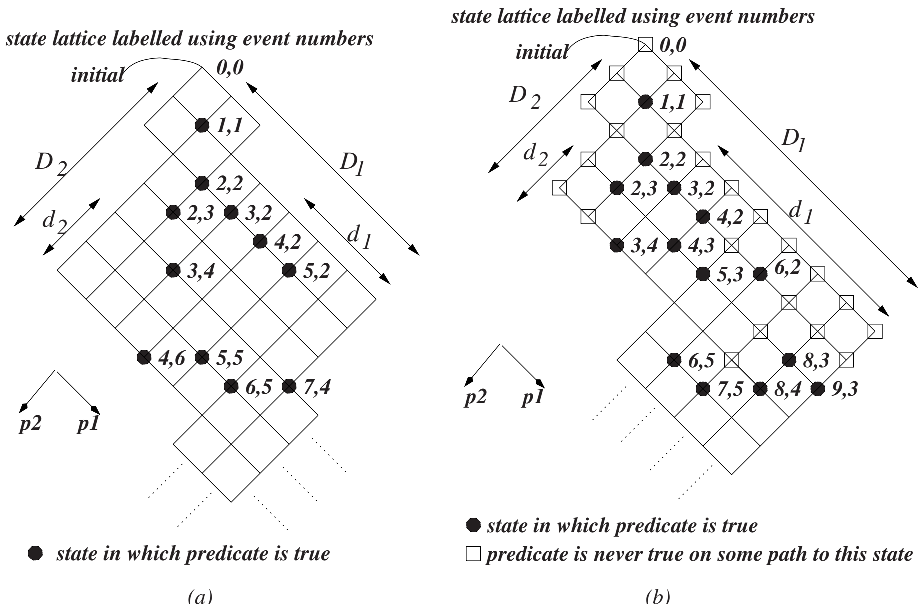

Fig. 5. Two examples to show minimal intervals for the *Definitely* modality for nonconjunctive predicates

Example 5(a) contd.: $D_1 = [e_1^1, e_1^7]$ and $D_2 = [e_2^1, e_2^6]$ are not *minimal*, in the sense that they have subintervals $d_1 = [e_1^4, e_1^7]$ and $d_2 = [e_2^4, e_2^6]$, resp., that satisfy the same property `DEF-SUBLATTICE`. Thus, every execution must pass through at least one state in which ϕ is *true* during the course of intervals d_1 and d_2.

Definition 5. (Minimal set of states in a sublattice, in which ϕ is satisfied:)

1. *Let $S_{min} = S(\prod_i d_i)$ represent the set of states in which ϕ is* true *in the minimal region satisfying* `DEF-SUBLATTICE`.
2. *Let S^*_{min} ($\subseteq S_{min}$) be such that S^*_{min} represents the minimal subset of states in which ϕ is* true *in the minimal region satisfying* `DEF-SUBLATTICE`.

Example 5(a) contd.: Every execution must pass through one state corresponding to $S_{min} = \{(e_1^3, e_2^4), (e_1^4, e_2^6), (e_1^5, e_2^5), (e_1^6, e_2^5), (e_1^7, e_2^4)\}$. There is a subset $S^*_{min} = \{(e_1^4, e_2^6), (e_1^5, e_2^5), (e_1^6, e_2^5), (e_1^7, e_2^4)\}$ of S_{min} such that every execution must pass through one state in S^*_{min}.
Example 5(b): In the example of Figure 5(b), ϕ is *true* in the 14 states shown. The example also shows each "reachable" state for which there is some path to that state such that the predicate is never *true* along that path.

DISJUNCT-FINE-GRAIN($S(\prod_i D_i)$)

1. For each global state ψ^k in $S(\prod_i D_i)$ (that satisfies predicate ϕ), do the following.
   (a) Identify I^k, the set of subintervals, one subinterval per process, corresponding to state ψ^k.
   (b) Using the tests and Theorem 1 [10], determine the orthogonal interaction $r^k_{i,j} \in \Re$, between each pair of subintervals in I^k. Denote the set of such interactions as $\mathcal{O}^k$.
2. As at least one ψ^k state must occur in any execution, the possible interaction types are given by $\bigvee_k \mathcal{O}^*$, where $\mathcal{O}^* = \{\mathcal{O}^k \mid \psi^k \in S(\prod_i D_i)\ \}$.

Fig. 6. Procedure DISJUNCT-FINE-GRAIN to identify possible pairwise fine-grained interaction types when $Definitely(\phi)$ holds

1. $D_1 = [e_1^1, e_1^9]$, $D_2 = [e_2^1, e_2^5]$, and $d_1 = [e_1^3, e_1^9]$, $d_2 = [e_2^3, e_2^5]$.
2. $S_{min}=\{(e_1^3, e_2^4), (e_1^4, e_2^3), (e_1^5, e_2^3), (e_1^6, e_2^5), (e_1^8, e_2^3), (e_1^7, e_2^5), (e_1^8, e_2^4), (e_1^9, e_2^3)\}$.
3. $S^*_{min} = \{\ (e_1^3, e_2^4), (e_1^4, e_2^3), (e_1^5, e_2^3), (e_1^6, e_2^5), (e_1^7, e_2^5), (e_1^8, e_2^4), (e_1^9, e_2^3)\ \}$.

The fine-grained modalities for $Definitely(\phi)$ are given by DISJUNCT-FINE--GRAIN($S(\prod_i d_i)$). Finally, we remark that it is possible to devise an algorithm MIN-DISJUNCT-FINE-GRAIN($S(\prod_i d_i)$) that identifies S^*_{min}. Such an algorithm would also be exponential in the number of states examined.

7 Conclusions

This paper examined the feasibility of using intervals to determine the fine-grained modality of nonconjunctive predicates. Although it is known that detecting nonconjunctive predicates (under the *Possibly*/*Definitely* modalities) involves examining an exponential number of states, nevertheless, this paper gave a better understanding of how the interval-based approach can be used to detect not just the *Possibly*/*Definitely* but also the fine-grained modalities. Three parameters were used for the analysis – how to decompose the global predicate, adjacency of local intervals, and the use of composite intervals. The analysis showed how the interval-based approach can be used to determine nonconjunctive predicates under fine-grained modalities. This included how to identify the intervals for nonconjunctive predicates, and how Theorem 1 [10] based on intervals can be adapted/extended to nonconjunctive predicates.

Acknowledgements

This material is based upon work supported by the National Science Foundation under Grant No. CCR-9875617.

References

[1] Alagar, S., Venkatesan, S.: Techniques to Tackle State Explosion in Global Predicate Detection, IEEE Trans. Software Engg., 27(8): 704-714, 2001. 12

[2] Chandra, P., Kshemkalyani, A. D.: Algorithms for Detecting Global Predicates under Fine-Grained Modalities, Proc. ASIAN 2003, LNCS, Springer, 2003 (to appear). 12

[3] Chandra, P., Kshemkalyani, A. D.: Distributed Algorithm to Detect Strong Conjunctive Predicates, Information Processing Letters, 87(5): 243-249, September 2003. 12

[4] Chandra, P., Kshemkalyani, A. D.: Detection of Orthogonal Interval Relations, Proc. High-Performance Computing Conference, 323-333, LNCS 2552, Springer, 2002. 14

[5] Cooper, R., Marzullo, K.: Consistent Detection of Global Predicates, Proc. ACM/ONR Workshop on Parallel & Distributed Debugging, 163-173, May 1991. 11, 12

[6] Garg, V. K., Waldecker, B.: Detection of Weak Unstable Predicates in Distributed Programs, IEEE Trans. Parallel & Distributed Systems, 5(3):299-307, Mar. 1994. 11, 12

[7] Garg, V. K., Waldecker, B.: Detection of Strong Unstable Predicates in Distributed Programs, IEEE Trans. Parallel & Distributed Systems, 7(12):1323-1333, Dec. 1996. 11, 12

[8] Hurfin, M., Mizuno, M., Raynal, M., Singhal, M.: Efficient Distributed Detection of Conjunctions of Local Predicates, IEEE Trans. Software Engg., 24(8): 664-677, 1998. 11, 12

[9] Kshemkalyani, A. D.: Temporal Interactions of Intervals in Distributed Systems, Journal of Computer and System Sciences, 52(2): 287-298, April 1996. 11

[10] Kshemkalyani, A. D.: A Fine-Grained Modality Classification for Global Predicates, IEEE Trans. Parallel & Distributed Systems, 14(8): 807-816, August 2003. 11, 12, 14, 15, 16, 21, 24

[11] Lamport, L.: Time, Clocks, and the Ordering of Events in a Distributed System, Communications of the ACM, 21(7): 558-565, July 1978. 13

[12] Marzullo, K., Neiger, G.: Detection of Global State Predicates, Proc. 5th Workshop on Distributed Algorithms, LNCS 579, Springer-Verlag, 254-272, October 1991. 11, 12

[13] Stoller, S., Schneider, F.: Faster Possibility Detection by Combining Two Approaches, Proc. 9th Workshop on Distributed Algorithms, 318-332, LNCS 972, Springer-Verlag, 1995. 12

[14] Venkatesan, S., Dathan, B.: Testing and Debugging Distributed Programs Using Global Predicates, IEEE Trans. Software Engg., 21(2): 163-177, Feb. 1995. 11, 12