Predicate Control: Synchronization in Distributed Computations with Look-ahead *

Ashis Tarafdar
Dept. of Computer Sciences
University of Texas at Austin
Austin, TX 78712
(ashis@cs.utexas.edu)

Vijay K. Garg
Dept. of Electrical and Computer Engg.
University of Texas at Austin
Austin, TX 78712
(garg@ece.utexas.edu)

- **Proposed Running Head:** Predicate Control
- **Address for Manuscript Correspondence:**
Ashis Tarafdar, 111 North St., Newtonville, MA 02460
Email: ashis@akamai.com, ashis@cs.utexas.edu
Phone: (617) 613 2564 (Day), (617) 527 3369 (Evening)
Fax: (617) 556 6810, (617) 556 4269

*supported in part by the NSF ECS-9414780, CCR-9520540, a General Motors Fellowship, Texas Education Board ARP-320 and an IBM grant

Abstract

The predicate control problem involves synchronizing a distributed computation to maintain a given global predicate. In contrast with many popular distributed synchronization problems such as mutual exclusion, readers writers, and dining philosophers, predicate control assumes a look-ahead, so that the computation is an off-line rather than an on-line input. Predicate control is targeted towards applications such as rollback recovery, debugging, and optimistic computing, in which such computation look-ahead is natural.

We define predicate control formally and show that, in its full generality, the problem is NP-Complete. We find efficient solutions for some important classes of predicates including "disjunctive predicates", "mutual exclusion predicates", and "readers writers predicates". For each class of predicates, we determine the necessary and sufficient conditions for solving predicate control and describe an efficient algorithm for determining a synchronization strategy. In the case of "independent mutual exclusion predicates", we determine that predicate control is NP-Complete and describe an efficient algorithm that finds a solution under certain constraints.

Keywords: Predicate Control, Synchronization, Distributed Computations, Mutual Exclusion, Global Predicates, Rollback recovery, Distributed Debugging, Controlled Re-execution

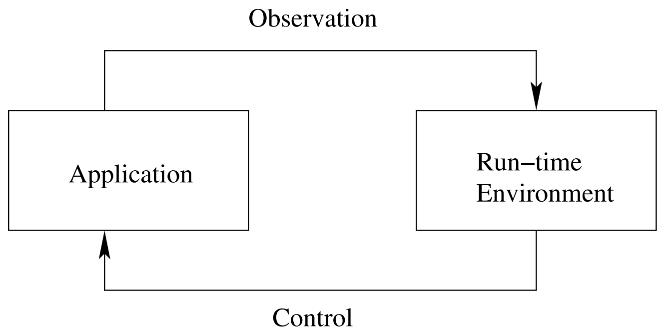

Figure 1: Observation and Control

1 Introduction

The interactions between an application and its run-time environment can be divided into two categories: *observation* and *control*. Examples of observation include monitoring load, detecting failures, checking variable values for debugging, and logging usage characteristics. Examples of control include dynamically balancing load, recovering from failures, resetting variable values during debugging, and terminating processes which violate usage restrictions. Such a duality between observation and control has also been noted in other areas, such as systems analysis.

In distributed applications, the problems of observation and control are made more difficult. In the case of distributed observation, this is because no one process has a global view of the system. Therefore, a global view must be correctly deduced from the local observations of each process. Distributed observation has been formalized as the widely-studied *predicate detection problem* [7, 9, 12]. In the case of distributed control, the difficulty lies in ensuring global conditions through appropriate local control actions. While there are many forms of control, our focus is on an important form of control – synchronization – involving control of the relative timing among processes. In this paper we define and study the *predicate control problem* that formalizes this form of distributed control.

In contrast to existing distributed synchronization problems [3] such as mutual exclusion, readers writers, and dining philosophers, predicate control assumes a look-ahead, so that the computation is known *a priori*. In other words, the computation is an off-line rather than an on-line input, and so, the predicate control problem is an off-line variant of on-line synchronization problems. Such a form of look-ahead is natural for certain applications, such as rollback recovery, debugging, and optimistic computing, in which the computation look-ahead is determined from a previous run. An important advantage of solving the off-line rather than the on-line synchronization problem is that, given a pre-knowledge of the computation, it is possible to solve the off-line problem in a general context under which the on-line problem would be impossible to solve. For example, the on-line mutual exclusion problem [29] is solvable under the assumption that critical sections do not block, whereas, as we will show, off-line mutual exclusion can be solved without such an assumption.

Informally, the predicate control problem is to determine how to add synchronizations to a distributed computation so that it maintains a global predicate. As an example, consider the computation shown in Figure 2(a) with three processes, P1, P2, and P3. The processes have synchronized by sending messages to one another. Suppose the stated global predicate is the mutual exclusion predicate, so that no two processes are to be in critical sections (labeled CS1, CS2, CS3, and CS4) at the same time. Clearly, the given computation does not maintain mutual exclusion at all times. Figure 2(b) shows the same computation with added synchronizations that ensure that mutual exclusion is maintained at all times. We call such a computation a "controlling computation". The main difficulty in determining such a controlling computation lies in adding the synchronizations in such a manner as to maintain the given property without causing deadlocks with the existing synchronizations.

Overview

We first define the model and state the predicate control problem formally in Section 2. Informally, modeling a computation as a partially ordered set of events, the predicate control problem is to determine how to add edges to a computation (that is, make the partial order of events stricter) so that it maintains a global predicate.

In its full generality, the predicate control problem deals with *any* predicates

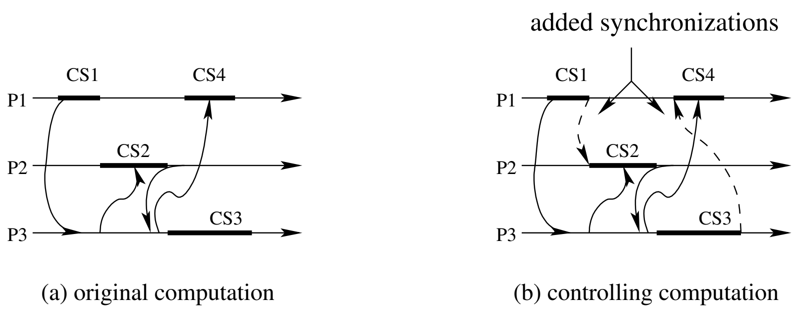

Figure 2: The Predicate Control Problem

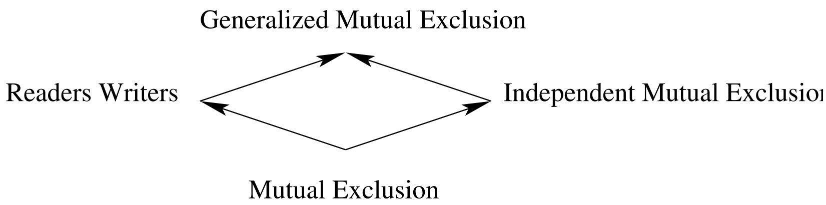

Figure 3: Variants of Mutual Exclusion Predicates

and it would, therefore, be expected that it is hard to solve. We establish in Section 3 that it is NP-Complete. However, as we will see, there are useful predicates for which the problem can be solved efficiently.

The first class of predicates that we study in Section 4.1 is the class of "disjunctive predicates" that are specified as a disjunction of local predicates. Intuitively, these predicates can be used to express a global condition in which at least one local condition has occurred, or, in other words, in which a bad combination of local conditions has not occurred. For example: "at least one server is available". First, we show that a necessary condition to solve the problem for disjunctive predicates is the absence of a clique of n false intervals in the interval graph. Next, we design an algorithm for solving the problem when the necessary condition is satisfied, thus establishing that the condition is also sufficient. The algorithm is of complexity $O(np)$, where n is the number of processes and p is the number of false intervals in the interval graph.

The next class of predicates is that of "mutual exclusion predicates" which state that no two processes are in critical sections at the same time. Mutual exclusion predicates are particularly useful in software fault tolerance since they correspond to races, a very common type of synchronization faults. In Section 4.2, our two results show that the necessary and sufficient conditions for solving predicate control for mutual exclusion predicates is the absence of a non-trivial cycle of critical sections in the interval graph. We also design an algorithm of complexity $O(np)$, where p is the number of critical sections in the computation.

We can generalize the mutual exclusion predicates to "readers writers predicates" specifying that only "write" critical sections must be exclusive, while "read" critical sections need not be exclusive. Another generalization is the "independent mutual exclusion predicates" where critical sections have "types" associated with them, such that no two critical sections of the same type can execute simultaneously. Finally, "generalized mutual exclusion predicates" allow read and write critical sections *and* multiple types of critical sections. Figure 3 illustrates the relative generality of the four problems.

In Section 4.3, we find that the necessary and sufficient conditions for solving predicate control for readers writers predicates is the absence of a non-trivial cycle of critical sections with at least one write critical section. For independent mutual exclusion, however, we discover that the problem is NP-Complete in general. We show this in Section 4.4, and also show that a sufficient condition for solving the problem is the absence of a non-trivial cycle of critical sections with two critical sections of the same type. However, in general, this condition is not necessary. The results for generalized mutual exclusion follow along similar lines. We do not describe individual algorithms for readers writers and independent mutual exclusion predicates. Instead, in Section 4.5, we describe a general $O(np)$ algorithm that solves predicate control for general mutual exclusion predicates (under the sufficient conditions). This algorithm can be applied to solving readers writers predicates (in general) and independent mutual exclusion predicates (under the sufficient conditions).

The predicate control problem is targeted for applications in which look-ahead is natural. In Section 5 we discuss three such applications: software fault tolerance, debugging, and optimistic computing. Finally, we give an outline of related research in Section 6, and we make some concluding remarks in Section 7. In this paper, we present the shorter proofs in full and, for the longer proofs, we present a proof outline together with a proof intuition where necessary. The complete proofs may be found in the appendix.

2 Model and Problem Statement

The predicate control problem is based on a model of distributed computations and concepts related to distributed computations. Our model is similar to the *happened before* model [20], described in more detail in [2, 11].

2.1 Computations

Since we are going to be concerned mainly with distributed computations, we drop the qualifier *distributed* and call them simply *computations*.

- $\langle$**computation**$\rangle$: A *computation* is a tuple $\langle E_1, E_2, \cdots E_n, \rightarrow\rangle$ where the E_i's are disjoint finite sets of "events" and $\rightarrow$ (precedes) is an irreflexive partial order on $E = \bigcup_i E_i$ such that, for each i, $\rightarrow_i$ ($\rightarrow$ restricted to E_i) is a total order. We abbreviate $\langle E_1, E_2, \cdots E_n, \rightarrow\rangle$ to $\langle E, \rightarrow\rangle$ with the understanding that n is to be always used to represent the size of the partition of E into $E_1, E_2, \cdots E_n$.

Informally, each E_i represents a *process*, and the partial order $\rightarrow$ represents a partial ordering of events such that $e \rightarrow f$ means that the event e *may* have directly or indirectly caused the event f. We say e "causally precedes" f. Note that, in general, the $\rightarrow$ relation is merely an approximation of causality. An example of approximating causality to obtain a $\rightarrow$ relation on events is given by the *happened before* model [20], in which $\rightarrow$ is defined as the smallest relation satisfying the following: (1) If e and f are events in the same process and e comes before f according to the local process clock, then $e \rightarrow f$, (2) If e is the send event of a message and f is the receive event for the same message by another process, then $e \rightarrow f$, and (3) If $e \rightarrow f$ and $f \rightarrow g$, then $e \rightarrow g$.

A special class of computations are runs, defined as:

- $\langle$**run**$\rangle$: A computation $\langle E, \rightarrow\rangle$ is a *run* if $\rightarrow$ is a total order.

A run is an even weaker approximation of causality than locally ordered computations. A run corresponds to an interleaving of all the events in the distributed computation. Finally, we define some notation concerning computations:

- $\langle \underline{\rightarrow} \rangle$: $(e \underline{\rightarrow} f) \equiv (e = f) \vee (e \rightarrow f)$
  In a similar way, we represent a corresponding reflexive relation for any given relation, e.g.: $(e \preceq f) \equiv (e = f) \vee (e \prec f)$
- $\langle e.proc \rangle$: $(e.proc = i) \equiv (e \in E_i)$

Informally, $e.proc$ is the identifier for the process containing e.

2.2 Cuts

- $\langle$**cut**$\rangle$: A *cut* is a subset of E containing at most one event from each E_i.

A cut corresponds to the intuitive notion of a global state. Sometimes a cut has been defined to correspond to the notion of a history (all events preceding a global state). In such definitions, the *frontier of a cut* corresponds to our *cut*. Since, we will be dealing more frequently with the notion of a global state than a history, we formalize the notion of history in terms of global state rather than the other way around.

If a cut intersects a process at an event e, then the local state on that process corresponding to the cut is the one reached by executing event e. If the cut does not intersect a process, then the local state on that process corresponding to the cut is the initial state on that process. To represent an initial state explicitly, we augment our model with initialization events:

- $\langle \perp_i \rangle$: Corresponding to each E_i we define a special event $\perp_i$ ($\perp_i \notin E$).
- $\langle \perp \rangle$: Let $\perp = \bigcup_i \{\perp_i\}$.

Each $\perp_i$ corresponds to a special "dummy" event that initializes the state of process E_i. Note that $\perp_i$ is not a real event. Introducing the dummy events, $\perp$, allows a one-to-one correspondence between the informal notion of local states corresponding to a process E_i and the formal notion of $E_i \cup \{\perp_i\}$, in which an event corresponds to the local state that follows it. We next define a local ordering between events and dummy events:

- $\langle \prec_i \rangle$: For each i, let $\prec_i$ be the smallest relation on $E_i \cup \{\perp_i\}$ such that:
  $\forall e \in E_i : \perp_i \prec_i e$, and
  $\forall e, f \in E_i : e \rightarrow_i f \Rightarrow e \prec_i f$.
- $\langle \prec \rangle$: Let $\prec = \bigcup_i \prec_i$.

Thus, the $\prec$ relation augments the $\rightarrow_i$ relation by ordering the initialization event $\perp_i$ before all events in E_i. Using the initialization events $\perp$, we can now define, for any cut, a unique event/initialization event corresponding to each process as:

- $\langle C[i] \rangle$: For a cut C, define $C[i] \in E \cup \perp$ such that:
  $C[i] = e$, if $C \cap E_i = \{e\}$, and
  $C[i] = \perp_i$, if $C \cap E_i = \emptyset$.

Finally, we define a relation $\leq$ on cuts as follows:

- $\langle \leq \rangle$: For two cuts C_1 and C_2 define a relation $\leq$ as:
  $C_1 \leq C_2 \quad \equiv \quad \forall i : C_1[i] \preceq C_2[i]$

2.3 Predicates

Predicates represent the informal notion of properties on local or global states:

- $\langle$**local predicate**$\rangle$: A *local predicate* of an E_i in a computation is a (polynomial-time computable) boolean function on $E_i \cup \{\perp_i\}$.
- $\langle$**global predicate**$\rangle$: A *global predicate* of a computation $\langle E, \rightarrow \rangle$ is a (polynomial-time computable) boolean function on the set of cuts of the computation.

Local predicates are used to define properties local to a process, such as: whether the process is in a critical section, whether the process has a token, or whether, for a process variable x, $x < 4$. Global predicates define properties across multiple processes, such as: whether two processes are in critical sections, whether at least one process has a token, or whether, for variables x and y on different processes, $(x < 4) \wedge (y < 3)$. We use *predicate* instead of global predicate, when the meaning is clear from the context.

2.4 Consistent Cuts

All cuts in a computation do not correspond to global states that may have happened in the computation. To see why this is true, consider the set of events which causally precede any event in a given cut, which we call the *causal past* of the cut:

- $\langle G.past \rangle$ **:** The *causal past* of a set of events $G \subseteq E$ in a computation $\langle E, \rightarrow \rangle$ is $G.past = \{e \in E : \exists f \in G : e \underline{\rightarrow} f\}$

Suppose there is an event e in the causal past of a cut C and suppose e has not yet occurred in the global state corresponding to C. Since a cause must always precede the effect, the cut C does not represent a global state that may have happened. Cuts which represent global states that may have happened are called *consistent* global states. The reason we say that a consistent global state *may* have happened is that its occurrence depends on the particular interleaving of the partially ordered set of events that occurred in real time.

- $\langle$**consistent cut**$\rangle$ **:** A cut C is *consistent* if $\forall e \in C.past : C[e.proc] \not\prec e$.

Note that, in the case of a locally ordered computation, $C[e.proc] \not\prec e$ is the same as $e \preceq C[e.proc]$.

We now prove two useful results about consistent cuts. The first one allows us to extend a single event to a consistent cut that contains it. This corresponds to the intuition that the local state following the event must have occurred at some time. Therefore, there must be some consistent cut (global state that may have happened) that contains it.

Lemma 1 *(Extensibility) Let $\mathcal{C}$ be the set of consistent cuts of a computation $\langle E, \rightarrow \rangle$, then:*

$$\forall e \in E : \exists C \in \mathcal{C} : e \in C$$

Proof: Let e be an event in E. Consider a maximal (w.r.t. $\leq$) cut C contained in $\{e\}.past$. Since C is maximal, it must contain e. Consider any event f in $C.past$. Since C is contained in $\{e\}.past$, f is also in $\{e\}.past$. Suppose $C[f.proc] \prec f$. Then since f is in $\{e\}.past$, C is not maximal and we have a contradiction. Therefore, $C[f.proc] \not\prec f$. So, C is consistent. □

The next result tells us that a cut that is consistent in a computation is also consistent in any computation that is less strict (that is, the partial ordering of events in the latter is less strict than the first).

Lemma 2 *If $\langle E, \rightarrow' \rangle$ and $\langle E, \rightarrow \rangle$ are computations such that $\rightarrow \subseteq \rightarrow'$, then any cut that is consistent in $\langle E, \rightarrow' \rangle$ is also consistent in $\langle E, \rightarrow \rangle$.*

Proof: Let C be a consistent cut in $\langle E, \rightarrow' \rangle$. Let e be any event in the causal past of C with respect to $\rightarrow$. Since $\rightarrow \subseteq \rightarrow'$, e is also in the causal past of C with respect to $\rightarrow'$. Therefore, $C[e.proc] \not\prec e$. So, C is consistent in $\langle E, \rightarrow \rangle$. □

In case of locally ordered computations, every event has a unique next event and a unique previous event:

- $\langle e.next \rangle$ **:** For $e \in E_i \cup \{\perp_i\}$, $e.next$ denotes the immediate successor to e in the total order defined by the $\prec_i$ relation. If no successor exists, $e.next = null$, where $null$ does not belong to $E \cup \perp$.

- $\langle e.prev \rangle$ **:** For $e \in E_i \cup \{\perp_i\}$, $e.prev$ denotes the immediate predecessor to e in the total order defined by the $\prec_i$ relation. If no predecessor exists, $e.prev = null$, where $null$ does not belong to $E \cup \perp$.

Note that $null$ (not an event) should not be confused with $\perp_i$ (the initialization event). In fact, $\perp_i.prev = null$.

The following result shows us an alternative way of proving that a cut is consistent.

Lemma 3 *In a computation $\langle E, \rightarrow \rangle$:*

a cut C is consistent $\equiv \quad \forall i, j : C[i].next, C[j] \in E : C[i].next \not\rightarrow C[j]$

Proof: (1) Suppose C is consistent. We show by contradiction that:

$\forall i, j : C[i].next, C[j] \in E : C[i].next \not\rightarrow C[j]$.

Suppose there is an i and j such that $C[i].next \rightarrow C[j]$. Then $C[i].next \in C.past$. But $C[i] \prec C[i].next$ contradicting the consistency of C.

(2) Suppose C is not consistent. We prove by contradiction that:

$\exists i, j : C[i].next, C[j] \in E : C[i].next \rightarrow C[j]$.

Since C is not consistent, there is an event, say e, such that $e \in C.past$ and $C[e.proc] \prec e$. Therefore, $C[e.proc].next \preceq e$. By transitivity, $C[e.proc].next \in C.past$. Therefore, $\exists j : C[j] \in E : C[e.proc].next \rightarrow C[j]$. $\Box$

The next is a well-known result about the set of consistent cuts in a computation [23]. It is easy to verify that:

Lemma 4 *The set of consistent cuts of a computation forms a lattice under* $\leq$.

It can further be verified that the lattice is a distributive one. This result allows us to define the consistent cut that is the greatest lower bound or least upper bound of a non-empty set of consistent cuts. Consider the set A of consistent cuts containing an event e. Lemma 1 tells us that A is non-empty. By Lemma 4, there is a consistent cut that is the greatest lower bound of A. Further, it is easy to check that this consistent cut contains e. Therefore, we can define the following:

- $\langle lcc(e) \rangle$ **:** For each event e in a computation, let $lcc(e)$ denote the *least consistent cut* that contains e.

Informally, $lcc(e)$ represents the earliest global state for which e has occurred. It is verifiable that $lcc(e)$ is the maximum cut contained in $\{e\}.past$ (similar to the proof of Lemma 1). Therefore, we have:

Lemma 5 $lcc(e).past = \{e\}.past$.

This result shows us that $lcc(e)$, which corresponds to the earliest global state for which e has occurred, is formed by the "frontier" of the causal past of e (that is, the set of last events per process in the causal past of e).

There can be only one global state that occurs immediately before $lcc(e)$ and it is represented by the following cut:

- $\langle lccprev(e) \rangle$ **:** For each event e in a computation, let $lcc\text{-}prev(e)$ denote the cut formed from $lcc(e)$ by deleting e, and then by adding $e.prev$ if $e.prev \notin \perp$.

To show that $lcc\text{-}prev(e)$ is indeed the global state that occurs immediately before $lcc(e)$, we must show that it is consistent.

Lemma 6 *For any event e in a computation, $lcc\text{-}prev(e)$ is consistent.*

Proof: We prove by contradiction. Let f be any event in $lcc\text{-}prev(e).past$ such that $lcc\text{-}prev(e)[f.proc] \prec f$. Clearly, by the definition of $lcc\text{-}prev(e)$ we have $f.proc \neq e.proc$. Therefore, $lcc\text{-}prev(e)[f.proc] = lcc(e)[f.proc]$ (by the defn. of $lcc\text{-}prev(e)$). Further, since f is in $lcc\text{-}prev(e).past$, it is also in $lcc(e).past$. This contradicts the consistency of $lcc(e)$. $\Box$

2.5 Runs

In this section we state some definitions and results specific to runs. In the case of a run, we can make a stronger assertion about the set of consistent cuts than the fact that it forms a lattice (Lemma 4):

Lemma 7 *The set of consistent cuts of a run is totally ordered under* $\leq$.

Proof: We prove by contradiction. Suppose that C_1 and C_2 are two consistent cuts in a computation $\langle E, \rightarrow \rangle$ such that $C_1 \not\leq C_2$ and $C_2 \not\leq C_1$. Therefore, there must exist i and j such that $C_1[i] \prec C_2[i]$ and $C_2[j] \prec C_1[j]$. Since $\rightarrow$ is a total order, without loss of generality, let $C_2[i] \rightarrow C_1[j]$. Therefore, $C_2[i]$ is in $C_1.past$ contradicting the consistency of C_1. $\Box$

This corresponds to the intuition that since the events in a run are totally ordered, the global states must also occur in sequence. Further, for any global state that occurs, there is a unique event that can lead up to it. This unique event is defined as:

- $\langle ge(C) \rangle$ **:** For each consistent cut C of a run $\langle E, \rightarrow \rangle$, such that $C \neq \emptyset$, let the greatest event in C with respect to $\rightarrow$ be denoted by $ge(C)$.

The next result shows us that there is a one-to-one correspondence between the consistent cuts (excluding $\emptyset$) and the events in a run.

Lemma 8 *In a given run, ge and lcc are inverses of each other.*

Proof: Consider a computation, $\langle E, \rightarrow \rangle$. It follows directly from Lemma 5 that $\forall e \in E : ge(lcc(e)) = e$. Therefore, it remains to prove that $lcc(ge(C)) = C$, for a consistent cut $C \neq \emptyset$. We prove by contradiction. Suppose $lcc(ge(C)) \neq C$, then there is an i such that $lcc(ge(C))[i] \neq C[i]$. By the definition of lcc, we must have: $lcc(ge(C))[i] \prec C[i]$. By the definition of ge, $C[i] \rightarrow ge(C)$ and so, $C[i] \in lcc(ge(C)).past$. This contradicts the consistency of $lcc(ge(C))$. $\Box$

2.6 Interval Graphs

Let $\langle E, \rightarrow \rangle$ be a locally ordered computation and let $\alpha_1, \alpha_2, \cdots \alpha_n$ be a set of local predicates. Each local predicate α_i defines a partition of the sequence of events $\langle E_i \cup \{\perp_i\}, \prec_i \rangle$ into "intervals" in which α_i is alternately true and false. We define this formally as:

- $\langle$**interval**$\rangle$: An interval I is a non-empty subset of an $E_i \cup \{\bot_i\}$ corresponding to a maximal subsequence in the sequence of events in $\langle E_i \cup \{\bot_i\}, \prec_i\rangle$, such that all events in I have the same value for α_i.

We next introduce a few notations and definitions related to intervals.

- $\langle \mathcal{I}_i \rangle$: Let $\mathcal{I}_i$ denote the set of intervals of $E_i \cup \{\bot_i\}$.
- $\langle \mathcal{I} \rangle$: Let $\mathcal{I} = \bigcup_i \mathcal{I}_i$.
- $\langle I.proc \rangle$: $(I.proc = i) \equiv (I \subseteq (E_i \cup \{\bot_i\}))$
- $\langle I.first \rangle$: For an interval I, $I.first$ denotes the minimum event in I with respect to $\prec_{I.proc}$.
- $\langle I.last \rangle$: For an interval I, $I.last$ denotes the maximum event in I with respect to $\prec_{I.proc}$.
- $\langle \prec_i$ **(for intervals)**$\rangle$: For any i, let the $\prec_i$ relation on events apply to intervals as well, such that $I_1 \prec_i I_2 \equiv I_1.first \prec_i I_2.first$.
- $\langle I.next \rangle$: $I.next$ denotes the immediate successor interval of interval I with respect to $\prec_{I.proc}$, or $null$ if none exists ($null$ is distinct from all intervals).
- $\langle I.prev \rangle$: $I.prev$ denotes the immediate predecessor interval of interval I with respect to $\prec_{I.proc}$, or $null$ if none exists ($null$ is distinct from all intervals).
- $\langle$**local predicate (on intervals)**$\rangle$: A local predicate α_i also applies to an interval $I \subseteq E_i \cup \{\bot_i\}$ such that $\alpha_i(I) \equiv \alpha_i(I.first)$.
- $\langle$**true interval**$\rangle$: An interval I is called a *true interval* if $\alpha_{I.proc}(I)$.
- $\langle$**false interval**$\rangle$: An interval I is called a *false interval* if $\neg\alpha_{I.proc}(I)$.

The following relation on the set of intervals, $\mathcal{I}$, is defined so that $I_1 \mapsto I_2$ represents the intuitive notion that: "I_1 must enter before I_2 can leave".

- $\langle \mapsto \rangle$: $\mapsto$ is a relation on intervals defined as :

$$I_1 \mapsto I_2 \equiv \begin{cases} I_1.first \to I_2.next.first & \text{if } I_1.prev \neq null \text{ and } I_2.next \neq null \\ true & \text{if } I_1.prev = null \text{ or } I_2.next = null \end{cases}$$

Note that, the conditions $I_1.prev \neq null$ and $I_2.next \neq null$ imply that $I_1.first \notin \perp$ (so $\rightarrow$ is defined) and $I_2.next.first$ is defined. If $I_1.prev = null$ or $I_2.next = null$, then we define $I_1 \mapsto I_2$ to be true. Since the execution starts in the first interval of every process and ends in the last interval of every process, this corresponds to the intuition that I_1 must enter before I_2 can leave. Further, note that the relation $\mapsto$ is reflexive, corresponding to the intuition that an interval must enter before it can leave. Unlike many of the relations we deal with, it is to be noted that the relation $\mapsto$ is not a partial order and may have cycles.

The set of intervals $\mathcal{I}$ together with the relation $\mapsto$ forms a graph, called an interval graph.

- $\langle$**interval graph**$\rangle$: $\langle \mathcal{I}, \mapsto \rangle$ is called the *interval graph* of computation $\langle E, \rightarrow \rangle$ under the set of local predicates $\alpha_1, \alpha_2, \cdots \alpha_n$.

The interval graph represents a higher granularity view of a computation, in which the intervals may be viewed as "large events". However, there is one important difference — an interval graph may be cyclic, while a computation is partially ordered.

2.7 Problem Statement

In order to state the problem, we first require the concept of a "controlling computation". Informally, given a predicate and a computation, a controlling computation is a stricter computation for which all consistent cuts satisfy the predicate. Formally, it is defined as:

- $\langle$**controlling computation**$\rangle$: Given a computation $\langle E, \rightarrow \rangle$ and a global predicate ϕ, a computation $\langle E, \rightarrow^c \rangle$ is called a *controlling computation* of ϕ in $\langle E, \rightarrow \rangle$, if: (1) $\rightarrow \subseteq \rightarrow^c$, and
(2) for all consistent cuts C in $\langle E, \rightarrow^c \rangle$: $\phi(C)$

Given this definition, the predicate control problem is:

The Predicate Control Problem: Given a computation $\langle E, \rightarrow \rangle$ and a global predicate ϕ, is there a controlling computation of ϕ in $\langle E, \rightarrow \rangle$?

The search problem corresponding to the predicate control problem is to find such a controlling computation, if one exists.

3 Predicate Control is NP-Complete

We prove that predicate control is NP-Complete in two steps. We first define the predicate control problem for runs and show that it is equivalent to the predicate control problem. We next show that the predicate control problem for runs is NP-Complete.

The Predicate Control Problem for Runs: Given a computation $\langle E, \rightarrow \rangle$ and a global predicate ϕ, is there a controlling *run* of ϕ in $\langle E, \rightarrow \rangle$?

Lemma 9 *The Predicate Control Problem is equivalent to the Predicate Control Problem for Runs.*

Proof: Let $\langle E, \rightarrow \rangle$ be a computation and ϕ be a global predicate. We have to show that a controlling computation of ϕ in $\langle E, \rightarrow \rangle$ exists if and only if a controlling run of ϕ in $\langle E, \rightarrow \rangle$ exists.

Since a run is a special case of a computation, the "if" proposition follows.

For the "only if" proposition, suppose a controlling computation $\langle E, \rightarrow^c \rangle$ of ϕ in $\langle E, \rightarrow \rangle$ exists. Let $<^c$ be any linearization of $\rightarrow^c$ (a *linearization* of a partial order is a totally ordered superset). If C is a consistent cut of $\langle E, <^c \rangle$, then C is also a consistent cut of $\langle E, \rightarrow^c \rangle$ (Lemma 2). Since, $\langle E, \rightarrow^c \rangle$ is a controlling computation, we have $\phi(C)$. Therefore, $\langle E, <^c \rangle$ is also a controlling computation of ϕ in $\langle E, \rightarrow \rangle$. $\Box$

Lemma 10 *The Predicate Control Problem for Runs is NP-Complete.*

Proof Outline: Since all the cuts of a run can be enumerated in polynomial time, the problem is in NP. A transformation from SAT proves the NP-Hardness. $\Box$

Theorem 1 *The Predicate Control Problem is NP-Complete.*

Proof: Follows directly from Lemmas 9 and 10. $\Box$

4 Solving Predicate Control

4.1 Disjunctive Predicates

- $\langle$**disjunctive predicates**$\rangle$ **:** Given n local predicates $\alpha_1, \alpha_2, \cdots, \alpha_n$, the *disjunctive predicate* ϕ_{disj} is defined as:

$$\phi_{disj}(C) \equiv \bigvee_{i \in \{1, \cdots, n\}} \alpha_i(C[i])$$

Some examples of disjunctive predicates are:

(1) At least one server is available:	$avail_1 \ \vee \ avail_2 \ \vee \ \dots \ avail_n$
(2) x must happen before y:	$after\ x \ \vee \ before\ y$
(3) At least one philosopher is thinking:	$think_1 \ \vee \ think_2 \ \vee \ \dots \ think_n$
(4) $(n-1)$-mutual exclusion:	$\neg cs_1 \ \vee \ \neg cs_2 \ \vee \ \dots \ \neg cs_n$

Note how we can even achieve the fine-grained control necessary to cause a specific event to happen before another as in property (3). This was done using local predicates to check if the event has happened yet. $(n-1)$-mutual exclusion is a special case of k-mutual exclusion, which states that at most k processes can be in the critical sections at the same time.

We next determine a necessary condition for solving the predicate control problem for disjunctive predicates.

Theorem 2 *Let $\langle \mathcal{I}, \mapsto \rangle$ be the interval graph of a computation $\langle E, \rightarrow \rangle$ under local predicates $\alpha_1, \alpha_2, \cdots, \alpha_n$. If $\langle \mathcal{I}, \mapsto \rangle$ contains a clique of n false intervals, then there is no controlling computation of ϕ_{disj} in $\langle E, \rightarrow \rangle$.*

Proof Intuition and Outline:
• **Intuition:** $I \mapsto I'$ means "I must enter before I' can leave". Therefore, before any interval in the clique can leave, all the other intervals must enter. Thus, there must be a time when each process is within the corresponding interval of the clique.
• **Outline:** Given the existence of a clique, clearly the n intervals must be on distinct processes. Any cut that intersects each interval of the clique does not satisfy ϕ_{disj}. We show that at least one such cut is consistent in any stricter computation. Therefore, no controlling computation exists. □

Algorithm Description

In order to show that the necessary condition is also sufficient, we design an algorithm that finds a controlling computation whenever the condition is true. The algorithm is shown in Figure 4.

The algorithm takes as input the n sequences of intervals for each of the processes. The output contains a sequence of added edges. The central idea is that this sequence of added edges links true intervals into a continuous "chain" from the start of the computation to the end. Any cut must either intersect this chain in a true interval, in which case it satisfies the disjunctive predicate, or in an edge, in which case the cut is made inconsistent by the added edge.

Types:

$event$: ($proc$: int; v: vector clock)
$interval$: ($proc$: int; α: boolean; $first$: event; $last$: event)

Input:

$\mathcal{I}_1, \mathcal{I}_2, \cdots, \mathcal{I}_n$: list of $interval$, initially non-empty

Output:

$\mathcal{O}$: list of ($event$, $event$), initially $null$

Vars:

$I_1, I_2, \cdots, I_n$: $interval$, initially $\forall i: \; I_i = \mathcal{I}_i.head$
$valid\_pairs$: set of ($interval$, $interval$)
$chain$: list of $interval$
$anchor, crossed$: $interval$
$prev, curr$: $interval$

Notation:

$$N_i = \begin{cases} \min X \in \mathcal{I}_i : \; I_i \preceq_i X \;\wedge\; X.\alpha = false, & \text{if exists} \\ null, & \text{otherwise} \end{cases}$$

$$select(Z) = \begin{cases} \text{arbitrary element of set } Z, & \text{if } Z \neq \emptyset \\ null, & \text{otherwise} \end{cases}$$

Procedure:

while ($\forall i: \; N_i \neq null$) do
    $valid\_pairs := \{(I_i, N_j) | \; (I_i.\alpha = true) \;\wedge\; (N_i \not\mapsto N_j)\}$
    if ($valid\_pairs = \emptyset$)
        exit("no controlling computation exists")
    $(anchor, crossed) := select(valid\_pairs)$
    if ($anchor.prev = null$) then
        $chain := null$
    $chain.add\_head(anchor)$
    for ($i \in \{1, \cdots n\}$: $i \neq crossed.proc$) do
        while ($I_i.next \neq null \;\wedge\; I_i.next \mapsto crossed$)
            $I_i := I_i.next$
    $I_{crossed.proc} := crossed.next$
$anchor := select(\{I_i : \; i \in \{1, \cdots n\} \;\; \wedge N_i = null\})$
if ($anchor.prev = null$) then
    $chain := null$
$chain.add\_head(anchor)$
$prev := chain.delete\_head()$
while ($chain \neq null$) do
    $curr := chain.delete\_head()$
    $\mathcal{O}.add\_head(\; (prev.first, curr.next.first) \;)$
    $prev := curr$

Figure 4: Algorithm for Predicate Control of the Disjunctive Predicate

The purpose of each iteration of the main loop in lines L12-L23 is to add one true interval, *anchor*, to the *chain* of true intervals. The intervals $I_1, \cdots, I_n$ form a "frontier" of intervals that start at the beginning of the computation and continue until one of them reaches the end. In each iteration, a *valid_pair*, consisting of a true interval, *anchor*, and a false interval, *crossed*, is selected at line L13, such that *anchor* does not have to leave before *crossed* leaves. If no such *valid_pair*, we can (and will) show that the necessary condition must be true so that an n-sized clique of false intervals exists in the interval graph. Next, the *anchor* is added to the *chain* of true intervals at line L19. Finally, the frontier advances such that *crossed* is crossed and all intervals which must enter as a result of this are entered (lines L20 - L23). After the loop terminates, the output is constructed from the chain of true intervals by connecting them in reverse order.

The time complexity of the algorithm is $O(np)$ where p is the number of false-intervals in the computation The naive implementation of the algorithm would be $O(n^2p)$ because the outer while loop iterates $O(np)$ times and calculating the set $ValidPairs()$ can take $O(n^2)$ time to check every pair of processes. However, an optimized implementation avoids redundant comparisons in computing the set $ValidPairs()$. Since, in this approach, each new false-intervals has to be compared with $n-1$ existing false-intervals, the time complexity is $O(np)$.

We have assumed, of course, that the algorithm is correct. We now establish this. In doing so, we also prove that the necessary condition – the existence of an n-sized clique of false intervals – is also a sufficient one for solving predicate control for disjunctive predicates.

We first show that the algorithm is well-specified so that it is indeed an algorithm – all the terms used are well-defined and it terminates.

Lemma 11 *The algorithm in Figure 4 is well-specified.*

Proof Outline: It can be verified that each term in the algorithm (especially the term $curr.next.first$ at line L31) is well-defined. The algorithm terminates since the main loop in the algorithm crosses at least one interval in each iteration. □

Next, we prove two useful invariants maintained by the algorithm.

Lemma 12 *The algorithm in Figure 4 maintains the following invariants on $I_1, I_2 \cdots, I_n$ (except in lines L20-L23 while I_i's are being updated):*

INV1: $\forall i, j$: *if* $I_i.next \neq null$ *and* $I_j.prev \neq null$ *then* $I_i.next \not\mapsto I_j.prev$

INV2: $\forall i$: *if* $I_i.prev \neq null$ *then* $I_i.\alpha = true \ \vee \ (\exists j : \ I_i \mapsto I_j.prev \ \wedge \ I_j.\alpha = true)$

Proof Intuition and Outline:

• **Intuition:** Intuitively, INV1 means that for any pair I_i and I_j, it is not the case that I_i must leave before I_j enters. The key observation that leads to INV1 is that, in each iteration, L20-L23 advances $I_1, I_2, \cdots, I_n$ only across the intervals that must leave before $crossed.next$ enters.

Intuitively, INV2 means that any I_i is either a true interval or there is another true interval I_j such that I_i must enter before I_j can enter. The key observation that leads to INV2 is that, in each iteration, L20-L23 advances $I_1, I_2, \cdots, I_N$ such that $I_{crossed.proc}$ becomes a true interval and all the other I_j's that are entered, must enter before $I_{crossed.proc}$ can enter.

• **Outline:** Each invariant is independently proved by structural induction treating L20-L23 as atomic. The base cases are trivial.

To prove the inductive case for INV1, consider two possibilities: I_j is either updated or remains constant. If I_j is updated, then it can be verified that L20-L23 ensures that INV1 is true. If I_j remains constant, then INV1 is true by the inductive hypothesis.

To prove the inductive case for INV2, consider two possibilities: I_i is either updated or remains constant. If I_i is updated, then it is easy to check that L20-L23 ensures that INV2 is true. If I_i remains constant, then the inductive hypothesis implies that, previously, there was an interval I_j so that INV2 is maintained. However, I_j may have been advanced. If so, then L20-L23 ensures that selecting $I_{crossed}$ as the corresponding interval satisfies INV2. $\Box$

We next show that if the algorithm exits abnormally at line L15, failing to produce a controlling computation, then no controlling computation exists for the problem instance.

Lemma 13 *If the algorithm in Figure 4 exits at line* L15, *then no controlling computation exists.*

Proof: If $\forall i, j: \; N_i \mapsto N_j$ then we have a clique of n false intervals and the result follows from Theorem 2. Therefore, let i and j be such that $N_i \not\mapsto N_j$. Since the set $valid\_pairs = \emptyset$ (line L14), we must have $I_i.\alpha = false$ (from line L13). –[A]

Therefore, by the definition of N_i, we have: $I_i = N_i$. So, since $N_i \not\mapsto N_j$, we have: $I_i.prev \neq null$ –[B]

Applying INV2 and using [A] and [B] let k be such that: $I_i \mapsto I_k.prev \; \wedge \; I_k.\alpha = true$. Since $N_i \not\mapsto N_j$ and $N_i = I_i$ and $I_i \mapsto I_k.prev$ we have: $N_k \not\mapsto N_j$. This contradicts the fact that $valid\_pairs = \emptyset$ (line L14). $\Box$

Finally, we show that the output does form a controlling computation.

Lemma 14 *If the algorithm in Figure 4 terminates normally, then* $\langle E, \rightarrow^c \rangle$ *is a controlling computation of* ϕ_{disj} *in* $\langle E, \rightarrow \rangle$*, where* $\rightarrow^c$ *is the transitive closure of the union of* $\rightarrow$ *with the set of edges in* $\mathcal{O}$*.*

Proof Intuition and Outline:
• **Intuition:** For two anchors A and A' in successive iterations, INV1 ensures that it is not true that A must leave before A' can enter. So, it is possible to ensure that A' enters before A leaves, by adding an edge from $A'.first$ to $A.next.first$. Each added edge ensures that the following anchor must enter before the previous anchor leaves. Therefore, at any given time during a run of the computation, at least one process must be within an anchor.
• **Outline:** We prove that the output edges define a valid computation by showing that the existence of a cycle would contradict INV1. The computation is proved to be a controlling computation by showing that any consistent cut must intersect one of the anchors, and therefore satisfies ϕ_{disj}. □

Our final theorem in this section states the sufficient condition for solving disjunctive predicate control as demonstrated by the correctness of our algorithm.

Theorem 3 *Let* $\langle \mathcal{I}, \mapsto \rangle$ *be the interval graph of a computation* $\langle E, \rightarrow \rangle$ *under local predicates* $\alpha_1, \alpha_2, \cdots, \alpha_n$*. If* $\langle \mathcal{I}, \mapsto \rangle$ *does not contain a clique of* n *false intervals, then there is a controlling computation of* ϕ_{disj} *in* $\langle E, \rightarrow \rangle$*.*

Proof: By Lemmas 11, 13, and 14, the algorithm in Figure 4 determines a controlling computation if $\langle \mathcal{I}, \mapsto \rangle$ does not contain a clique of n false intervals. □

4.2 Mutual Exclusion Predicates

- $\langle$**mutual exclusion predicates**$\rangle$ **:** Let $critical_1, critical_2, \cdots, critical_n$ be n local predicates and let $critical(e) \equiv critical_{e.proc}(e)$. The *mutual exclusion predicate* ϕ_{mutex} is defined as:

$$\phi_{mutex}(C) \equiv \forall \, distinct \; i, j : \neg \, (\, critical(C[i]) \, \wedge \, critical(C[j]) \,)$$

- $\langle$**critical section**$\rangle$ **:** A *critical section* (*non-critical section*) denotes a true (false) interval in $\langle E, \rightarrow \rangle$ with respect to $critical_1, critical_2, \cdots, critical_n$.

Mutual exclusion is one of the most common forms of synchronization in distributed applications. Since a "data race" is a violation of mutual exclusion, determining a controlling computation of mutual exclusion in a computation is equivalent to synchronizing a computation to eliminate all data races.

The next two theorems demonstrate the necessary and sufficient conditions for solving predicate control for the mutual exclusion predicates.

Theorem 4 *Let $\langle \mathcal{I}, \mapsto \rangle$ be the interval graph of a computation $\langle E, \rightarrow \rangle$ under local predicates $critical_1, critical_2, \cdots, critical_n$. If $\langle \mathcal{I}, \mapsto \rangle$ contains a non-trivial cycle of critical sections, then there is no controlling computation of ϕ_{mutex} in $\langle E, \rightarrow \rangle$.*

Proof Intuition and Outline:
• **Intuition:** In the cycle, consider the first critical section CS_1 to leave during a run of the computation. Let CS_2 be its predecessor in the cycle, so that CS_2 must enter before CS_1 can leave. Therefore, at some point in the run processes are executing within both CS_1 and CS_2.
• **Outline:** We prove by contradiction. Suppose there is a non-trivial cycle of critical sections and a controlling computation. Consider any run of the controlling computation. We demonstrate that there is a consistent cut in the run that does not satisfy ϕ_{mutex}. The consistent cut is either: (1) $lcc - prev(e)$, where e is the minimum event such that $e = CS.next.first$, for some critical section CS in the cycle, if such an e exists, or (2) the maximum cut in the run, if no such e exists. □

Theorem 5 *Let $\langle \mathcal{I}, \mapsto \rangle$ be the interval graph of a computation $\langle E, \rightarrow \rangle$ under local predicates $critical_1, critical_2, \cdots, critical_n$. If $\langle \mathcal{I}, \mapsto \rangle$ does not contain a non-trivial cycle of critical sections, then there is a controlling computation of ϕ_{mutex} in $\langle E, \rightarrow \rangle$.*

Proof Outline: Since there are no non-trivial cycles, we topologically sort the critical section graph. An edge is added from each critical section to the successor in topologically sorted order. The absence of non-trivial cycles in the critical section graph ensures that these added edges define a valid computation. Further, we show that, in the resulting computation, any cut that intersects two critical sections is not consistent. □

Algorithm Description

The proof of the above theorem provides a simple algorithm for finding a controlling computation based on topologically sorting the interval graph of critical sections. We provide a more efficient algorithm in Figure 5, making use of the fact that the critical sections in a process are totally ordered. The key idea used in the algorithm

Types:

event:	(*proc*: int; v: vector clock)
interval:	(*proc*: int; *critical*: boolean; *first*: event; *last*: event)

Input:

$\mathcal{I}_1, \mathcal{I}_2, \cdots, \mathcal{I}_n$:	list of *interval*, initially non-empty

Output:

$\mathcal{O}$:	list of (*event*, *event*), initially *null*

Vars:

$I_1, I_2, \cdots, I_n$:	*interval*, initially $\forall i : \; I_i = \mathcal{I}_i.head$
valid_cs:	set of *interval*
chain:	list of *interval*
crossed:	*interval*
prev, *curr*:	*interval*

Notation:

$$CS_i = \begin{cases} \min X \in \mathcal{I}_i : \; I_i \preceq_i X \;\wedge\; X.critical = true, & \text{if exists} \\ null, & \text{otherwise} \end{cases}$$

$$CS = \{ \; CS_i \mid \; (1 \leq i \leq n) \;\wedge\; (CS_i \neq null) \; \}$$

$$select(Z) = \begin{cases} \text{arbitrary element of set } Z, & \text{if } Z \neq \emptyset \\ null, & \text{otherwise} \end{cases}$$

Procedure:

while ($CS \neq \emptyset$) do
    $valid\_cs := \{ \; c \in CS \mid \; \forall c' \in CS : \; (c' \neq c) \Rightarrow (c' \not\rightarrow c) \; \}$
    if ($valid\_cs = \emptyset$)
        exit("no controlling computation exists")
    $crossed := select(valid\_cs)$
    $chain.add\_head(crossed)$
    $I_{crossed.proc} := crossed.next$
if ($chain \neq null$)
    $prev := chain.delete\_head()$
while ($chain \neq null$) do
    $curr := chain.delete\_head()$
    $\mathcal{O}.add\_head(\; (curr.next.first, prev.first) \;)$
    $prev := curr$

Figure 5: Algorithm for Predicate Control of the Mutual Exclusion Predicate

is to maintain a frontier of critical sections $(CS_1, \cdots, CS_n)$ that advances from the start of the computation to the end. Instead of finding a minimal critical section of the whole interval graph, we merely find the minimal critical section in the current frontier. It is guaranteed to be a minimal critical section of the remaining critical sections in the interval graph at that point. Therefore, this procedure achieves a topological sort.

The main while loop of the algorithm executes p times in the worst case, where p is the number of critical sections in the computation. Each iteration takes $O(n^2)$, since it must compute the *valid_cs*. Thus, a simple implementation of the algorithm will have a time complexity of $O(n^2 p)$. However, a better implementation of the algorithm would amortize the cost of computing *valid_cs* over multiple iterations of the loop. Each iteration would compare each of the critical sections that has newly reached the head of its list with the existing critical sections. Therefore, each of the p critical section reaches the head of the list just once, when it is compared with $n-1$ critical sections. The time complexity of the algorithm with this improved implementation is, therefore, $O(np)$. Note that a naive algorithm based directly on the constructive proof of the sufficient condition in Theorem 5 would take $O(p^2)$. We have reduced the complexity significantly by using the fact that the critical sections in a process are totally ordered.

We next prove the correctness of the algorithm.

Lemma 15 *The algorithm in Figure 5 is well-specified.*

Proof Outline: It can be shown that each term in the algorithm (especially $curr.next.first$ at line L23) is well-defined. The algorithm terminates since the main loop in the algorithm crosses at least one critical section in each iteration. □

Lemma 16 *If the algorithm in Figure 5 exits at line L16, then no controlling computation exists.*

Proof: At the point of exit at line 16 $valid\_cs = \emptyset$. Therefore, the graph $\langle CS, \mapsto \rangle$ has no minimal element. Since $CS \neq \emptyset$ (by the terminating condition at line L13), there must be a non-trivial cycle in $\langle CS, \mapsto \rangle$. Therefore, by Theorem 4, no controlling computation exists. □

Lemma 17 *If the algorithm in Figure 5 terminates normally, then $\langle E, \rightarrow^c \rangle$ is a controlling computation of ϕ_{mutex} in $\langle E, \rightarrow \rangle$, where $\rightarrow^c$ is the transitive closure of the union of $\rightarrow$ with the set of edges in $\mathcal{O}$.*

Proof Outline: We show that, at the termination of loop L13-L19, *chain* is a topological sort of the graph $\langle \mathcal{C}, \mapsto \rangle$. Then, as in the proof of Theorem 5, the computation defined by the added edges is a controlling computation. □

Theorem 6 *The algorithm in Figure 5 solves the predicate control problem for the mutual exclusion predicate.*

Proof: This follows directly from Lemmas 15, 16, and 17. □

4.3 Readers Writers Predicates

Let $critical_1, critical_2, \cdots, critical_n$ be n local predicates and let $critical(e) \equiv critical_{e.proc}(e)$ and let a *critical section* (*non-critical section*) denote a true (false) interval in $\langle E, \rightarrow \rangle$ with respect to $critical_1, critical_2, \cdots, critical_n$.

- $\langle$**read-critical/write-critical**$\rangle$ **:** A critical section is either *read-critical* or *write-critical* .
- $\langle write\_critical(I) \rangle$ **:** We define $write\_critical(I)$ to be true for a write-critical section I and false for all other intervals. We also say $write\_critical(e)$ for all events e in a write-critical section.
- $\langle read\_critical(I) \rangle$ **:** We define $read\_critical(I)$ to be true for a read-critical section I and false for all other intervals. We also say $read\_critical(e)$ for all events e in a read-critical section.
- $\langle$**readers writers predicate**$\rangle$ **:** The *readers writers predicate* ϕ_{rw} is defined as:

$$\phi_{rw}(C) \equiv \forall \; distinct \; i, j : \neg \; (\; critical(C[i]) \; \wedge \; write\_critical(C[j]) \;)$$

The readers writers predicate is a generalized form of mutual exclusion allowing critical sections to be read or write-critical. Two critical sections cannot enter at the same time if one of them is write-critical. The next two theorems establish the necessary and sufficient conditions for solving the predicate control problem for readers writers predicates. Since the algorithm that will be presented in Section 4.5 is applied to readers writers predicates without significant simplification, we do not present a specialized algorithm for readers writers predicates.

Theorem 7 *Let $\langle \mathcal{I}, \mapsto \rangle$ be the interval graph of a computation $\langle E, \rightarrow \rangle$ under local predicates $critical_1, critical_2, \cdots, critical_n$. If $\langle \mathcal{I}, \mapsto \rangle$ contains a non-trivial cycle of critical sections containing at least one write-critical section, then there is no controlling computation of ϕ_{rw} in $\langle E, \rightarrow \rangle$.*

Proof Intuition and Outline:
• **Intuition:** In the cycle, consider any *write* critical section W. Consider the point when W leaves. There are two possibilities: either (Case 1) at least one critical section has not yet left, or (Case 2) all the critical sections have left. (Case 1:) The predecessor of W in the cycle must have entered. If it has left, then its predecessor must have entered. Repeating this argument, we determine at least one critical section which must have entered but not left. Therefore, the readers writers predicate is violated. (Case 2:) Consider the point when the successor of W left. At that time, W must have entered, thereby violating the readers writers predicate.
• **Outline:** We prove by contradiction. Suppose there is a non-trivial cycle of critical sections with at least one write critical section and a controlling computation exists. Consider any run of the controlling computation. We demonstrate that there is a consistent cut in the run that does not satisfy ϕ_{rw}.

The consistent cut is determined as follows. Consider any write critical section W and its successor CS in the cycle. The consistent cut is one of: (1) $lcc-prev(W.next.first)$, (2) $lcc-prev(CS.next.first)$, or (3) the maximum cut in the run. At least one of these cuts is well-defined and does not satisfy ϕ_{rw}. □

Theorem 8 *Let $\langle \mathcal{I}, \mapsto \rangle$ be the interval graph of a computation $\langle E, \rightarrow \rangle$ under local predicates $critical_1, critical_2, \cdots, critical_n$. If $\langle \mathcal{I}, \mapsto \rangle$ does not contain a non-trivial cycle of critical sections containing at least one write-critical section, then there is a controlling computation of ϕ_{rw} in $\langle E, \rightarrow \rangle$.*

Proof Outline: We topologically sort the strongly connected components (scc's) of the critical section graph. An edge is added from each critical section in one scc to each critical section in the successor scc in topologically sorted order. We show that the added edges define a valid computation. Further, we show that, in the resulting computation, any cut that intersects a write-critical section and another critical section is not consistent. □

4.4 Independent Mutual Exclusion Predicates

Let $critical_1, critical_2, \cdots, critical_n$ be n local predicates and let $critical(e) \equiv critical_{e.proc}(e)$ and let a *critical section* (*non-critical section*) denote a true (false) interval in $\langle E, \rightarrow \rangle$ with respect to $critical_1, critical_2, \cdots, critical_n$.

- $\langle$**k-critical**$\rangle$ **:** A critical section is k-critical for some $k \in \{1, \cdots m\}$.

- $\langle k\_critical(I) \rangle$ **:** We define $k\_critical(I)$ to be true for a k-critical section I and false for all other intervals. We also say $k\_critical(e)$ for all events e in a k-critical section.

- $\langle$**independent mutual exclusion predicate**$\rangle$ **:** The *independent mutual exclusion predicate* ϕ_{ind} is defined as:

  $$\phi_{ind}(C) \equiv \forall \, distinct \; i, j : \forall \, k : \neg \, (\, k\_critical(C[i]) \wedge k\_critical(C[j]) \,)$$

The independent mutual exclusion predicate is a generalized form of mutual exclusion allowing critical sections to have types in $k \in \{1, \cdots m\}$. Two critical sections cannot enter at the same time if they are of the same type. The next result shows us that the problem becomes hard for this generalization. However, we will determine sufficient conditions for solving it that allow us to solve the problem efficiently under certain conditions.

Theorem 9 *The predicate control problem for the independent mutual exclusion predicate is NP-Complete.*

Proof Outline: The problem is in NP since the general predicate control problem is in NP. The NP-Hardness is established by a transformation from 3SAT. □

Although the problem is NP-Complete, the next result states a sufficient condition under which it can be solved. The condition is the absence of cycles containing two critical sections of the same type. Under these conditions, we can construct an efficient algorithm. Since the algorithm that will be presented in Section 4.5 is applied to independent mutual exclusion predicates without significant simplification, we do not present a specialized algorithm in this section.

Theorem 10 *Let $\langle \mathcal{I}, \mapsto \rangle$ be the interval graph of a computation $\langle E, \rightarrow \rangle$ under local predicates $critical_1, critical_2, \cdots, critical_n$. If $\langle \mathcal{I}, \mapsto \rangle$ does not contain a non-trivial cycle of critical sections containing two k-critical section for some k, then there is a controlling computation of ϕ_{rw} in $\langle E, \rightarrow \rangle$.*

Proof: The proof is along similar lines to the proof of Theorem 8. □

4.5 Generalized Mutual Exclusion Predicates

Using the definitions of the previous two sections:

- $\langle$**generalized mutual exclusion predicate**$\rangle$: The *generalized mutual exclusion predicate* ϕ_{gen_mutex} is defined as:
$\phi_{gen\_mutex}(C) \equiv \forall\ distinct\ i,j :$
$$\neg\ (\ critical(C[i])\ \wedge\ write\_critical(C[j])\ )\ \wedge$$
$$\forall\ k :\ \neg\ (\ k\_critical(C[i])\ \wedge\ k\_critical(C[j])\ )$$

Generalized mutual exclusion predicates allow critical sections to have types and be read/write-critical. Clearly the predicate control problem is NP-Complete for generalized mutual exclusion predicates. Further, a similar sufficient condition can be proved combining the sufficient conditions for readers writers and independent mutual exclusion predicates.

Algorithm Description

Based on the proof of the sufficient conditions, we can design a simple algorithm based on determining the strongly connected components in the critical section graph and then topologically sorting them. Instead, we present a more efficient algorithm in Figure 6.

In order to understand how the algorithm operates, we require the concept of a "general interval". A general interval is a sequence of intervals in a process that belong to the same strongly connected component of the interval graph. For the purposes of the algorithm, it is convenient to treat such sequences of intervals as a single general interval. We now define general intervals and define a few notations.

- $\langle$**general interval**$\rangle$: Given an interval graph $\langle \mathcal{I}, \mapsto \rangle$, a *general interval* is a contiguous sequence of intervals in an $\langle \mathcal{I}_i, \prec_i \rangle$ subgraph.

- $\langle g.first/g.last \rangle$: For a general interval g, let $g.first$ ($g.last$) represent the first and last intervals in g.

- $\langle g.set \rangle$: Let $g.set$ represent the set of intervals in the sequence g.

- $\langle\mapsto$ **(for general intervals)**$\rangle$: Let $\mathcal{G}$ denote the set of general intervals in a computation. We define the (overloaded) relation $\mapsto$ for general intervals as:

$$\forall g, g' \in \mathcal{G} :\ g \mapsto g'\ \equiv\ g.first \mapsto g'.last$$

Types:

event: ($proc$: int; v: vector clock) (L1)
gen_interval: ($proc$: int; $critical$: boolean; $first$: event; $last$: event) (L2)
str_conn_comp: set of *gen_interval* (L3)

Input:

$\mathcal{I}_1, \mathcal{I}_2, \cdots, \mathcal{I}_n$: list of *gen_interval*, initially non-empty (L4)

Output:

$\mathcal{O}$: list of (*event*, *event*), initially *null* (L5)

Vars:

$I_1, I_2, \cdots, I_n$: *gen_interval*, initially $\forall i: \ I_i = \mathcal{I}_i.head$ (L6)
scc_set: set of *str_conn_comp* (L7)
valid_scc: set of *str_conn_comp* (L8)
chain: list of *str_conn_comp* (L9)
candidate: *str_conn_comp* (L10)
prev_scc, *curr_scc*: *str_conn_comp* (L11)
prev, *curr*: *gen_interval* (L12)

Notation:

$CS_i = \begin{cases} \min X \in \mathcal{I}_i : \ I_i \preceq_i X \ \wedge \ X.critical = true, & \text{if exists} \\ null, & \text{otherwise} \end{cases}$ (L13)
$CS = \{ \ CS_i \mid \ (1 \leq i \leq n) \ \wedge \ (CS_i \neq null) \ \}$ (L14)
$get\_scc(CS)$ = set of strongly connected components in $\langle CS, \mapsto \rangle$ (L15)
$select(Z) = \begin{cases} \text{arbitrary element of set } Z, & \text{if } Z \neq \emptyset \\ null, & \text{otherwise} \end{cases}$ (L16)
$not\_valid(X)$ = X has a non-trivial cycle of critical sections with either one write critical section or two k-critical sections (L17)
$merge(c, c') = (c.last := c'.last; \ c.next := c'.next)$ (L18)

Procedure:

while ($CS \neq \emptyset$) do (L19)
  $scc\_set := get\_scc(CS)$ (L20)
  $valid\_scc := \{ \ s \in scc\_set \mid \ \forall s' \in scc\_set : \ (s' \neq s) \Rightarrow (s' \not\mapsto s) \ \}$ (L21)
  $candidate := select(valid\_scc)$ (L22)
  $mergeable := \{ \ c.next.next \mid \ c \in candidate \ \wedge \ c.next.next \neq null \ \wedge \ \exists c' \in candidate : \ c.next.next \mapsto c' \}$ (L23)
  if ($mergeable = \emptyset$) (L24)
    if ($not\_valid(candidate)$) (L25)
      exit("cannot find controlling computation") (L26)
    $chain.add\_head(candidate)$ (L27)
    for ($c \in candidate$) do (L28)
      $I_{c.proc} := c.next$ (L29)
  else (L30)
    for ($c \in mergeable$) do (L31)
      $merge(CS_{c.proc}, c)$ (L32)
if ($chain \neq null$) (L33)
  $prev\_scc := chain.delete\_head()$ (L34)
while ($chain \neq null$) do (L35)
  $curr\_scc := chain.delete\_head()$ (L36)
    for ($curr \in curr\_scc$, $prev \in prev\_scc$) do (L37)
  $\mathcal{O}.add\_head(\ (curr.next.first, prev.first)\ )$ (L38)
  $prev\_scc := curr\_scc$ (L39)

Figure 6: Algorithm for Generalized Mutual Exclusion Predicate Control

- $\langle \alpha(g) \rangle$: We say that a general interval g is *true* under a local predicate α if $\alpha(g.first)$. In particular, if g is *true* under a local predicate *critical*, we call g a *general critical section*.

- $\langle \hookrightarrow \rangle$: Let $\mathcal{G}_1 \subseteq \mathcal{G}$ be a set of general intervals. Let $\mathcal{S}_{\mathcal{G}_1}$, the set of strongly connected components in the graph $\langle \mathcal{G}_1, \mapsto \rangle$. We define a relation $\hookrightarrow$ on $\mathcal{S}_{\mathcal{G}_1}$ as follows:
$$\forall s, s' \in \mathcal{S}_{\mathcal{G}_1} : \; s \hookrightarrow s' \;\equiv\; \exists g \in s, g' \in s' : \; g \mapsto g'$$
Clearly, $\langle \mathcal{S}_{\mathcal{G}_1}, \hookrightarrow \rangle$ has no cycles.

The algorithm maintains a frontier of general critical sections that advances from the beginning of the computation to the end. In each iteration, the algorithm finds the strongly connected components (scc's) of the general critical sections in the frontier. Then, it picks a minimal strongly connected component, *candidate*, from among them (line L22). However, the *candidate* is not necessarily a minimal scc of the entire critical section graph. In fact, it need not even be an scc of the entire graph. To determine if it is, we find the *mergeable* set of critical sections that immediately follow the general critical sections and belong to the same scc (line L23). If *mergeable* is not empty, the critical sections in *mergeable* are merged with the general critical sections in *candidate* to give larger general critical sections (line L32) and the procedure is repeated. If *mergeable* is empty, then it can be shown that *candidate* is a minimal scc of the graph. Therefore, we check that it meets the sufficient conditions of validity (line L25), and then append it to the *chain* (line L27).

Finally, after the main loop terminates, the scc's in the *chain* are connected using added edges which define the controlling computation (lines L33-L39). Note how the use of general critical sections allows us to reduce the number of edges that need to connect two consecutive scc's as compared to the simple algorithm that would be defined by the proof of Theorem 8.

The main while loop of the algorithm executes p times in the worst case, where p is the number of critical sections in the computation. Each iteration takes $O(n^2)$, since it must compute the scc's. Thus, a simple implementation of the algorithm will have a time complexity of $O(n^2 p)$. However, a better implementation of the algorithm would amortize the cost of computing scc's over multiple iterations of the loop. Each iteration would compare each of the critical sections that have newly reached the heads of the lists with the existing critical sections. Therefore, each of the p critical section reaches the head of the list just once, when it is compared with $n-1$ critical sections to determine. The time complexity of the algorithm with this improved implementation is, therefore, $O(np)$. Note that a naive algorithm based directly on the constructive proof of the sufficient condition in Theorem 8 would take $O(p^2)$. We have reduced the complexity significantly by using the fact that the critical sections in a process are totally ordered.

Finally, we prove the correctness of the algorithm.

Lemma 18 *The algorithm in Figure 6 is well-specified.*

Proof Outline: It can be verified that each term in the algorithm (especially the term $curr.next.first$ at line L37) is well-defined. The algorithm terminates since, in each iteration of the main loop, either one critical section is crossed or two distinct critical sections are merged into one. □

Lemma 19 *If the algorithm in Figure 6 exits at line* L26*, then there is a non-trivial cycle of critical sections with:*

(1) at least one write-critical, or

(2) two k-critical sections for some $k \in \{1, \cdots, m\}$.

Proof: Follows directly from the condition at line L25. □

Lemma 20 *If the algorithm in Figure 6 terminates normally, then* $\langle E, \rightarrow^c \rangle$ *is a controlling computation of* ϕ_{gen_mutex} *in* $\langle E, \rightarrow \rangle$*, where* $\rightarrow^c$ *is the transitive closure of the union of* $\rightarrow$ *with the set of edges in* $\mathcal{O}$.

Proof Outline: Consider the graph of regular critical sections $\langle \mathcal{C}, \mapsto \rangle$ and the corresponding graph of scc's $\langle \mathcal{S}, \hookrightarrow \rangle$. We first prove that any two regular critical sections within the same general critical section belong to the same scc of $\mathcal{S}$. This follows from the way in which two critical sections are merged into one in L32.

Next, we show that when a *candidate* scc of general critical sections is added to *chain* at L27, then it corresponds to a valid scc of regular critical sections. Further, we show that it also corresponds to a minimal scc of regular critical sections in the remaining critical section graph.

Therefore, at the termination of loop L19-L32, *chain* corresponds to a topological sort of the graph $\langle \mathcal{S}, \hookrightarrow \rangle$. Then, as in the proof of Theorem 8, the computation defined by the added edges is a controlling computation. □

Theorem 11 *The algorithm in Figure 6 solves the predicate control problem for the generalized mutual exclusion predicate if the input computation has no non-trivial* *cycles of critical sections containing two k-critical sections for some $k \in \{1, \cdots, m\}$.*

Proof: By Lemma 18, the algorithm always terminates, and by Lemma 20, if it terminates normally then it outputs a controlling computation. By Lemma 19, if it terminates at line L26, then there is a non-trivial cycle with:

(1) at least one write-critical, or

(2) two k-critical sections for some $k \in \{1, \cdots, m\}$.

In case (1), Theorem 7 indicates that no controlling computation exists. Therefore, it is only in case (2) that the algorithm fails to find a controlling computation that does exist, (which is to be expected since the problem is NP-Complete by Theorem 3). □

5 Applications

In this section, we describe three examples of application domains in which predicate control has potential benefits.

5.1 Software Fault-tolerance

Distributed programs often encounter failures, such as races, owing to the presence of synchronization faults (bugs). Rollback recovery [16, 41, 6] is an important technique that is used to tolerate such faults. It involves rolling back the program to a previous state and re-executing, in the hope that the failure does not recur. The disadvantage of this approach is that it cannot guarantee that the re-execution will itself be failure-free.

By tracing relevant state information during the first execution, the look-ahead required by predicate control can be achieved. It is then possible to solve the predicate control problem for the traced computation and for the predicate expressing that the failure does not recur. For example, a mutual exclusion predicate expresses that a data race failure does not occur. If a controlling computation exists, it is used to synchronize the computation appropriately during re-execution to guarantee that a failure does not recur. If no controlling computation exists, then a failure must always occur during any re-execution of the computation.

An implementation of a rollback recovery system based on predicate control for tolerating data races in distributed MPI applications has been described in [36]. Experiments showed that the tracing costs involved less than 1% time overhead and a space cost of less than 3 MB/hour per process. It was also experimentally validated that re-execution based on predicate control greatly improved the ability to recover from failures as compared to simple re-execution and re-execution with locking.

5.2 Distributed Debugging

Software faults (bugs) in distributed programs are difficult to detect and eliminate since they are often difficult to reproduce. An example of such software faults are synchronization faults such as races which depend on the relative timing among processes. In order to facilitate the localization of software faults, trace-replay methods have been developed in many different environments: message-passing parallel programs [27], shared-memory parallel programs [26], distributed shared memory programs [31], and multi-threaded programs [8]. These methods involve tracing the non-deterministic events during execution and replaying them during re-execution.

Applying existing trace-replay methods leads to a cycle of replay and passive observation to determine the cause of the anomaly. A more active process would involve the user specifying a constraint that is imposed on the replayed computation. A common constraint that would be required is mutual exclusion of various functions. The user could then replay the computation with the added constraint to observe whether the anomaly occurs under mutual exclusion. Since tracing is already carried out for the purposes of replay, look-ahead is provided naturally. Predicate control can, therefore, be applied to determine a controlling computation that can be replayed to present an observation under the specified constraints.

More details about the application of predicate control to distributed debugging may be found in [37] and [24]. Note that in debugging it is more acceptable to incur tracing costs than in software fault-tolerance.

5.3 Optimistic Computing

Optimistic (or speculative) computing involves making an optimistic assumption that a certain condition will not occur. This allows the computation to proceed without delaying for a check to be carried out. In the eventuality that the condition does occur, then the computation is rolled back and an alternative path is followed. Optimistic computing typically finds applications in shared-memory multiprocessor environments [33, 35, 1] and distributed shared memory environments [14, 18, 17, 4, 25]. The optimistic assumption is used to express conditions such as: a certain branch is taken [33, 35], and a certain data access is consistent with a specified memory consistency model [14, 18, 17, 4, 1, 25].

In optimistic computing the advantage of optimism is in improving performance. Therefore, it is important that the optimism does not fail very often. Furthermore, it is important that when the optimism does fail and the program rolled back, the resulting re-execution is safe. Predicate control allows an optimistic assumption to be specified as a general predicate, such as mutual exclusion. This eliminates the need for expensive synchronization operations such as locking or barrier synchronizations. If the optimism fails, then predicate control can be applied since traced information can provide the necessary computation look-ahead. This ensures that the re-execution is safe. The predicate control approach differs from existing approaches based on memory consistency models in that it allows atomicity to be specified over operations on multiple objects.

There is a similarity between optimistic computing and software fault-tolerance since in both cases a failure is detected and the program is rolled back. However, in the case of optimistic computing, the failure is intentionally allowed to occur, whereas in software fault-tolerance, the failure is unexpected. This implies that optimistic computing can integrate the observation and control mechanisms into the application development process. In the case of software fault tolerance, the mechanisms must be isolated in a transparent way.

6 Related Work

The results in this paper were first published in [37] and [38]. We are aware of two previous studies of controlling distributed systems to maintain classes of global predicates. One study [21] allows global properties within the class of conditional elementary restrictions [21]. Unlike our model of a distributed system, their model uses an off-line specification of pair-wise mutually exclusive states and does not use causality. [30] and [39] study the on-line maintenance of a class of global predicates based on ensuring that a sum or sum-of-product expression on local variables does not exceed a threshold. In contrast to these approaches, our focus is on general global boolean predicates, disjunctive predicates, and mutual exclusion predicates.

Since our results in predicate control were published, some advances in solving predicate control have been made in [24]. The authors present an equivalent formulation of the predicate control problem in terms of "admissible sequences". Based on this, they provide an algorithm for disjunctive predicates that has the same time complexity as the algorithm presented in this paper, but has the added advantage of generating the minimum synchronizations. They also solve the predicate control problem for the class of "region predicates" that can express such conditions as deadlock and termination.

Relation to Predicate Detection

Since observation and control are duals of each other, it is not surprising that there is a relation between predicate detection, a formalization of observation, and predicate control, a formalization of control. To understand the relationship, consider a computation in which *definitely* : $\neg\phi$ is detected, where the modality *definitely* for a predicate [9] implies that the predicate is true in *any* possible execution (run) corresponding to the computation. Clearly, there can be no controlling computation of ϕ in the computation since any execution corresponding to the controlling computation is also an execution corresponding to the original computation. Therefore, a controlling computation of a predicate ϕ in a given computation exists if and only if *definitely* : $\neg\phi$ is not detected in the computation. As evidence of this, the necessary and sufficient conditions for detecting *definitely* : ϕ for conjunctive predicates were demonstrated in [13], and are similar to our results in predicate control for disjunctive predicates. An important difference between the two problems is that, though the predicate detection problem for *definitely* : $\neg\phi$ can be used to determine whether a controlling computation of ϕ exists, it is not concerned with actually finding the controlling computation. Another interesting conclusion is that all of our predicate control algorithms can be used to solve a corresponding predicate detection problem.

Relation to On-line Synchronization

The predicate control problem is an off-line synchronization problem since it assumes that the computation is known *a priori*. In contrast, existing synchronization problems [3] such as mutual exclusion, readers writers, and dining philosophers are on-line and require synchronization without such pre-knowledge. The predicate control problem for each of the predicates we have studied has a corresponding on-line synchronization problem, in which the computation is an on-line rather than an off-line input. Predicate control for mutual exclusion predicates corresponds to the widely-studied distributed mutual exclusion problem [29, 3, 40, 32]. Predicate control for readers writers predicates corresponds to the readers writers problem [10, 19, 3]. Predicate control for independent mutual exclusion predicates does not correspond to a new on-line problem. Instead, multiple distributed mutual exclusion protocols operate independently assuming that deadlocks are prevented using a programming discipline such as two-phase locking. Predicate control for disjunctive predicates corresponds to a special case of the k-mutual exclusion problem [28, 34, 22, 15, 5] (that allows up to k processes inside a critical section at the same time) in which $k = (n - 1)$.

Each of the on-line problems above make assumptions on the programming discipline such as non-blocking critical sections and two-phase locking. In a general context, without these assumptions, it can be shown that the on-line synchronization problems are impossible to solve [37], and therefore, it is impossible to maintain the predicate while avoiding deadlocks. Our solutions to the predicate control problems do not make such assumptions. Instead, they make use of the pre-knowledge of the computation to maintain the predicate while avoiding deadlocks. Thus, off-line synchronization is, inherently, a more powerful technique than on-line synchronization.

Another distinction between the on-line and off-line problems is that the predicate control problems operate on a single computation of the distributed program. The on-line synchronization problems, on the other hand, operate at the level of the programming language and must, therefore, correctly synchronize for all possible computations of the program.

7 Conclusions

The interaction between an application and its run-time environment has conventionally been perceived as a passive one consisting primarily of observation. The dual function of control has existed only in simple forms such as application termination. In this paper we have aimed at demonstrating a more active interaction in which the control function of the environment is enhanced. As a first step in this direction, we have studied how one form of control – synchronization – can be achieved. Predicate control formalizes the problem of synchronizing a distributed computation with look-ahead (that is, the computation is an off-line input). The look-ahead allows for improved control with respect to existing synchronization techniques. The predicate control problem is, therefore, targeted to applications like distributed debugging and rollback recovery for which such a computation look-ahead occurs naturally. We have demonstrated efficient algorithms for solving predicate control for some important classes of predicates. This is encouraging, especially since the corresponding on-line synchronization problems are impossible to solve without constraining assumptions.

An interesting future direction would be to generalize predicate control to include other forms of control. Some examples of other forms of control are: controlling the order of message delivery, controlling the allocation and deallocation of memory, and controlling the physical location of the application in a distributed system.

In the study of the predicate control problem, there are some interesting open questions. We have solved predicate control for mutual exclusion predicates and for disjunctive predicates. These are two extreme subcases of the more general *k-mutual exclusion* predicates. The k-mutual exclusion predicates specify that at most k of the processes can enter critical sections together. Thus, we have solved the cases of 1-mutual exclusion and $(n-1)$-mutual exclusion. We do know of a necessary condition for solving the predicate control problem for k-mutual exclusion. This corresponds to a set of critical sections in which each critical section has at least k incoming $\mapsto$ edges from the other critical sections. Note that from this, we can derive the necessary and sufficient conditions for both 1-mutual exclusion and $(n-1)$-mutual exclusion. However, it still remains an open question whether the k-mutual exclusion problem can be efficiently solved.

It would also be useful to solve the predicate control problem in the context of *relational predicates* [11]. An example of such a predicate is $(x_1 + x_2 + \cdots + x_n > k)$, where x_i's are integer variables in different processes and k is a constant. This allows us to express conditions such as "the usage of a limited shared resource is within acceptable bounds".

References

[1] S. V. Adve, V. S. Pai, and P. Ranganathan. Recent advances in memory consistency models for hardware shared memory systems. *Proceedings of the IEEE*, 87(3):445–455, March 1999.

[2] O. Babaoglu and K. Marzullo. Consistent global states of distributed systems: fundamental concepts and mechanisms. In S. Mullender, editor, *Distributed Systems*, chapter 4. Addison-Wesley, 1993.

[3] M. Ben-Ari. *Principles of concurrent and distributed programming.* Prentice Hall, 1990.

[4] R. Biachini, R. Pinto, and C. Amorim. Data prefetching for software DSMs. In *Proceedings of the International Conference on Supercomputing*, pages 385–392, Melbourne, Australia, July 1998. ACM.

[5] S. Bulgannawar and N. H. Vaidya. A distributed K-mutual exclusion algorithm. In *Proceedings of the 15th International Conference on Distributed Computing Systems*, pages 153–160. IEEE, 1995.

[6] S. Chandra and P. M. Chen. How fail-stop are faulty programs? In *Proc. of the Symposium on Fault-tolerant Computing (FTCS)*, Munich, Germany, June 1998. IEEE.

[7] K. M. Chandy and L. Lamport. Distributed snapshots: Determining global states of distributed systems. *ACM Transactions on Computer Systems*, 3(1):63 – 75, February 1985.

[8] J. D. Choi and H. Srinivasan. Deterministic replay of Java multithreaded applications. In *2nd SIGMETRICS Symp. on Parallel and Distr. Tools*, pages 48 – 59, Aug. 1998.

[9] R. Cooper and K. Marzullo. Consistent detection of global predicates. In *Proceedings of the ACM/ONR Workshop on Parallel and Distributed Debugging*, pages 163 – 173, Santa Cruz, California, 1991.

[10] P. J. Courtois, F. Heymans, and D. L. Parnas. Concurrent control with "readers" and "writers". *Communications of the ACM*, 14(10):667–668, October 1971.

[11] V. K. Garg. *Principles of Distributed Systems.* Kluwer Academic Publishers, 1996.

[12] V. K. Garg and B. Waldecker. Detection of weak unstable predicates in distributed programs. *IEEE Transactions on Parallel and Distributed Systems*, 5(3):299 – 307, March 1994.

[13] V. K. Garg and B. Waldecker. Detection of strong unstable predicates in distributed programs. *IEEE Transactions on Parallel and Distributed Systems*, 7(12):1323 – 1333, December 1996.

[14] Gudjon Hermannsson and Larry D. Wittie. Optimistic synchronization in distributed shared memory. In *Proceedings of the 14th International Conference on Distributed Computing Systems*, pages 345–354. IEEE, 1994.

[15] S.-T. Huang, J.-R. Jiang, and Y.-C. Kuo. k-coteries for fault-tolerant k entries to a critical section. In *Proceedings of the 13th International Conference on Distributed Computing Systems*, pages 74 – 81. IEEE, 1993.

[16] Y. Huang and C. Kintala. Software implemented fault tolerance: technologies and experience. In *Proc. IEEE Fault-Tolerant Comp. Symp.*, pages 138 – 144, June 1993.

[17] M. Karlsson and P. Stenstrom. Effectiveness of dynamic prefetching in multiple-writer distributed virtual shared memory systems. *Journal of Parallel and Distributed Computing*, 43(7):79–93, July 1997.

[18] P. Keleher, A. L. Cox, S. Dwarkadas, and W. Zwaenepoel. An evaluation of software-based release consistent protocols. *Journal of Parallel and Distributed Computing*, 29:126–141, September 1995.

[19] L. Lamport. Concurrent reading and writing. *Communications of the ACM*, 20(11):806–811, November 1977.

[20] L. Lamport. Time, clocks, and the ordering of events in a distributed system. *Communications of the ACM*, 21(7):558 – 565, July 1978.

[21] A. Maggiolo-Schettini, H. Wedde, and J. Winkowski. Modeling a solution for a control problem in distributed systems by restrictions. *Theoretical Computer Science*, 13(1):61 – 83, January 1981.

[22] K. Makki, P. Banta, K. Been, N. Pissinou, and E. Park. A token based distributed k mutual exclusion algorithm. In *Proceedings of the Symposium on Parallel and Distributed Processing*, pages 408 – 411. IEEE, December 1992.

[23] F. Mattern. Virtual time and global states of distributed systems. In *Parallel and Distributed Algorithms: Proc. of the International Workshop on Parallel and Distributed Algorithms*, pages 215 – 226. Elsevier Science Publishers B. V. (North Holland), 1989.

[24] N. Mittal and V. K. Garg. Debugging distributed programs using controlled re-execution. In *Proc. of the 19th ACM Symposium on Principles of Distributed Computing*, pages 239–248, Portland, USA, July 2000.

[25] F. Mueller. Adaptive DSM-runtime behavior via speculative data distribution. In J. Rolim et al., editor, *Parallel and Distributed Computing, Proceedings of the IPPS/SPDP99 Workshops, Workshop on Run-Time Systems for Parallel Programming*, LNCS 1586, pages 553–567. Springer-Verlag, April 1999.

[26] R. H. B. Netzer. Optimal tracing and replay for debugging shared-memory parallel programs. In *Proc. of ACM/ONR Workshop on Parallel and Distributed*

Debugging, pages 1 – 11, San Diego, USA, May 1993. Also available as ACM SIGPLAN Notices Vol. 28, No. 12.

[27] R. H. B. Netzer and B. P. Miller. Optimal tracing and replay for debugging message-passing parallel programs. In *Supercomputing '92*, pages 502 – 511, November 1992.

[28] K. Raymond. A distributed algorithm for multiple entries to a critical section. *Information Processing Letters*, 30:189–193, February 1989.

[29] M. Raynal. *Algorithms for Mutual Exclusion.* MIT Press, 1986.

[30] M. Raynal. *Distributed Algorithms and Protocols.* John Wiley and Sons Ltd., 1988.

[31] M. Ronnse and W. Zwaenepoel. Execution replay for TreadMarks. In *Proc. of the 5th EUROMICRO Workshop on Parallel and Distributed Processing (PDP'97)*, pages 343–350, January 1997.

[32] M. Singhal. A taxonomy of distributed mutual exclusion. *Journal of Parallel and Distributed Computing*, 18:94 – 101, 1993.

[33] J. E. Smith. A study of branch prediction strategies. In *Proc. of the 8th International Symposium on Computer Architecture*, pages 135–148, May 1981.

[34] P. K. Srimani and R. L. Reddy. Another distributed algorithm for multiple entries to a critical section. *Information Processing Letters*, 41:51–57, January 1992.

[35] J. Stark and Y. Patt. Variable length path branch prediction. In *Architectural Support for Programming Languages and Operating Systems*, pages 170–179, 1998.

[36] A. Tarafdar. *Software Fault Tolerance in Distributed Systems using Controlled Re-execution.* PhD thesis, University of Texas at Austin, 2000.

[37] A. Tarafdar and V. K. Garg. Predicate control for active debugging of distributed programs. In *Proc. of the 9th Symposium on Parallel and Distributed Processing*, pages 763–769, Orlando, USA, April 1998. IEEE.

[38] A. Tarafdar and V. K. Garg. Software fault tolerance of concurrent programs using controlled re-execution. In *Proc. of the 13th International Symposium on*

Distributed Computing, pages 210–224, Bratislava, Slovak Republic, September 1999.

[39] A. I. Tomlinson and V. K. Garg. Maintaining global assertions on distributed sytems. In *Computer Systems and Education*, pages 257 – 272. Tata McGraw-Hill Publishing Company Limited, 1994.

[40] Martin G. Velazquez. A survey of distributed mutual exclusion algorithms. Technical Report 93-116, Computer Science Dept., Colorado State University, 1993.

[41] Y. M. Wang, Y. Huang, W. K. Fuchs, C. Kintala, and G. Suri. Progressive retry for software failure recovery in message-passing applications. *IEEE Trans. on Computers*, 46(10):1137–1141, October 1997.

A Proofs

Lemma 4 *The set of consistent cuts of a computation forms a lattice under* $\leq$.

Proof: Let C_1 and C_2 be any two consistent cuts. Let C be a consistent cut defined such that for each i, $C[i]$ is the minimum of $C_1[i]$ and $C_2[i]$. Clearly, this ensures that for any lower bound C' of C_1 and C_2, $C' \leq C$. Therefore, to show that C is the greatest lower bound of C_1 and C_2, it remains to show that C is consistent.

Suppose C is not consistent. Using Lemma 3, there must exist i and j such that $C[i].next \rightarrow C[j]$. Without loss of generality, suppose $C[i] = C_1[i]$. Clearly, $C[j] \neq C_1[j]$ since, otherwise, the consistency of C_1 is contradicted. Therefore, $C[j] = C_2[j]$. By the definition of C, $C_2[j] \preceq C_1[j]$. Therefore, by transitivity, $C_1[i].next \rightarrow C_1[j]$. This contradicts the consistency of C_1. This shows that C is consistent and, therefore, C is the greatest lower bound of C_1 and C_2.

In a similar way, we can demonstrate the existence of the least upper bound of C_1 and C_2. $\square$

Lemma 10 *The Predicate Control Problem for Runs is NP-Complete.*

Proof: First, we show that the problem is in NP. A non-deterministic algorithm need only guess a run $\langle E, <^c \rangle$ such that $<^c$ is a linearization of $\rightarrow$. To check that a run is a controlling computation, we first enumerate all the consistent cuts in the run. By Lemma 8, each consistent cut is the *lcc* of some event, and, therefore, we can achieve the enumeration by finding the *lcc* of each event in the run. Using Lemma 5 and the fact that the causal past of a cut is unique, it can be easily shown that the *lcc* of an event can be constructed in polynomial time. Further, by definition, it can be checked in polynomial time that the given predicate satisfies each *lcc*.

To show that the problem is NP-Hard, we transform SAT. Let $U = \{u_1, \cdots u_m\}$ be a set of boolean variables, and $D = \{d_1, \cdots d_k\}$ be a set of clauses forming an instance of SAT. Let ψ be the boolean function over the set of truth assignments such that $\psi(t) \equiv t$ satisfies D. (ψ can be represented as a boolean formula in conjunctive normal form).

For each variable $u_i \in U$, construct a singleton set of events $E_i = \{e_i\}$. Construct one more set $E_{m+1} = \{e_{m+1}, e'_{m+1}\}$. Let $E = \bigcup_{i \in \{1, \cdots m+1\}} E_i$. Let relation $\rightarrow$ be defined so that $e_{m+1} \rightarrow e'_{m+1}$ and no other events are related. Clearly $\rightarrow$ is a partial order on E and $\rightarrow_i$ is a total order. Therefore, $\langle E, \rightarrow \rangle$ is a (locally ordered) computation.

For each i, let α_i be a local predicate on $E_i \cup \{\perp_i\}$ as follows. Let $\alpha_{m+1}(\perp_{m+1}) = \mathit{true}$, $\alpha_{m+1}(e_{m+1}) = \mathit{false}$, and $\alpha_{m+1}(e'_{m+1}) = \mathit{true}$. For each $i \leq m$ let $\alpha_i(\perp_i) = \mathit{false}$ and $\alpha_i(e_i) = \mathit{true}$. Note that for each cut C, there is a truth assignment t_C such that $t_C(u_i) \equiv \alpha_i(C[i])$. Define a global predicate $\phi = \alpha_{m+1}(C[m+1]) \vee \psi(t_C)$.

Thus, we have a polynomial construction from an instance (U, D) of SAT to an instance $(\langle E, \rightarrow \rangle, \phi)$ of the Predicate Control Problem for Runs. We next show that this is indeed a transformation.

Part 1: Suppose that $\langle E, <^c \rangle$ is a controlling run of ϕ in $\langle E, \rightarrow \rangle$. Let C be any consistent cut containing e (by Lemma 1). Since $\phi(C) = \mathit{true}$ and $\alpha_{m+1}(C[m+1]) = \mathit{false}$, we must have $\psi(t_C)$. Therefore, t_C is a truth assignment that satisfies the set of clauses D.

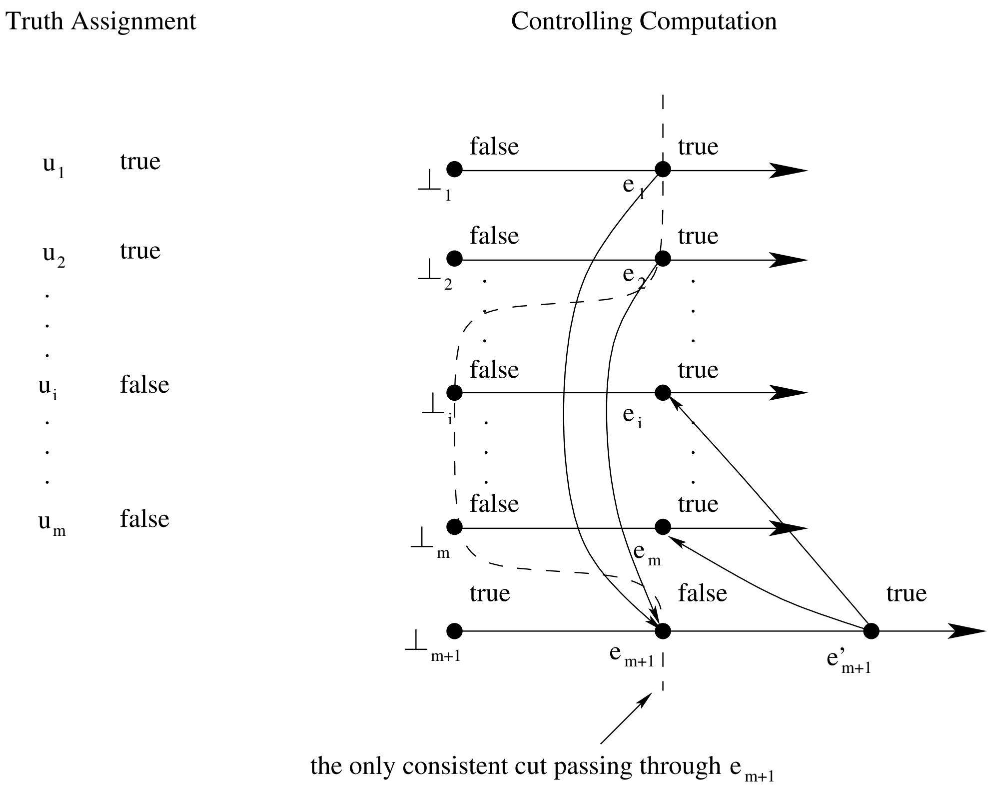

Figure 7: Proof: The Predicate Control Problem for Runs is NP-Complete

Part 2: Suppose that t is a truth assignment of variables in U that satisfies D. Define $\rightsquigarrow^c$ such that $\rightarrow \subset \rightsquigarrow^c$ and $\rightsquigarrow^c$ contains the following edges in addition to edges in $\rightarrow$: (1) for each $i \leq m$ such that $t(u_i) = \mathit{true}$, the edge $e_i \rightsquigarrow^c e_{m+1}$, and (2) for each $i \leq m$ such that $t(u_i) = \mathit{false}$, the edge $e'_{m+1} \rightsquigarrow^c e_i$. Let $\rightarrow^c$ be the

transitive closure of $\rightsquigarrow^c$. Since each E_i except E_{m+1} does not have both incoming and outgoing $\rightsquigarrow^c$ edges, there can be no cycles in $\rightsquigarrow^c$, and, therefore, $\rightarrow^c$ is a partial order and $\langle E, \rightarrow^c \rangle$ is a computation. Let $<^c$ be any linearization of $\rightarrow^c$.
Claim: The run $\langle E, <^c \rangle$ is a controlling run of ϕ in $\langle E, \rightarrow \rangle$.
Proof: Let C be a consistent cut of the run $\langle E, <^c \rangle$. We show that $\phi(C) = \mathit{true}$. If $\alpha_{m+1}(C[m+1]) = \mathit{true}$, we directly have $\phi(C) = \mathit{true}$. Therefore, we consider the remaining possibility: $\alpha_{m+1}(C[m+1]) = \mathit{false}$. By the definition of α_{m+1}, we have: $C[m+1] = e_{m+1}$. We next show that $t_C = t$. Consider any $i \leq m$. We consider two cases depending on the value of $t(u_i)$:

Case 1: $t(u_i) = \mathit{true}$:
Therefore, $e_i <^c e_{m+1}$ (by the definition of $\rightsquigarrow^c$ and $\rightsquigarrow^c \subset <^c$). Since C is consistent and $C[m+1] = e_{m+1}$, we must have $C[i] = e_i$. By the definition of α_i, we have $\alpha_i(C[i]) = \mathit{true}$. Therefore, $t_C(u_i) = \mathit{true}$.

Case 2: $t(u_i) = \mathit{false}$:
Therefore, $e'_{m+1} <^c e_i$ (by the definition of $\rightsquigarrow^c$ and $\rightsquigarrow^c \subset <^c$). Since C is consistent and $C[m+1] = e_{m+1}$, we must have $C[i] = \perp_i$. By the definition of α_i, we have $\alpha_i(C[i]) = \mathit{false}$ and, therefore, $t_C(u_i) = \mathit{false}$.

This proves that $t_C = t$. Since t satisfies the set of clauses, D, we have $\psi(t_C) = \mathit{true}$. Therefore, $\phi(C) = \mathit{true}$. $\langle$ *End: Proof of Claim* $\rangle$ $\Box$

Theorem 2 *Let $\langle \mathcal{I}, \mapsto \rangle$ be the interval graph of a computation $\langle E, \rightarrow \rangle$ under local predicates $\alpha_1, \alpha_2, \cdots, \alpha_n$. If $\langle \mathcal{I}, \mapsto \rangle$ contains a clique of n false intervals, then there is no controlling computation of ϕ_{disj} in $\langle E, \rightarrow \rangle$.*

Proof: Let $\langle \mathcal{I}, \mapsto \rangle$ have a clique X of n false intervals. Since for any two distinct false intervals I and I' in X, $I \mapsto I'$ and $I' \mapsto I$, we must have $I.\mathit{proc} \neq I'.\mathit{proc}$. Therefore, each of the false intervals in X is on a distinct E_i. Let $I_1, I_2, \cdots I_n$ be the false intervals of X, such that $I_i.\mathit{proc} = i$). We prove by contradiction that there is no controlling computation of ϕ_{disj} in $\langle E, \rightarrow \rangle$.

Suppose that $\langle E, \rightarrow^c \rangle$ is a controlling computation of ϕ_{disj} in $\langle E, \rightarrow \rangle$. Let $<^c$ be any linearization of $\rightarrow$. Therefore, by Lemma 2, $\langle E, <^c \rangle$ is also a controlling computation. We next show that there is a consistent cut C in the run $\langle E, <^c \rangle$ such that $\neg \phi_{disj}(C)$, thus, contradicting the fact that $\langle E, <^c \rangle$ is a controlling computation.

Consider two cases:

Case 1: $\forall i : I_i.next = null$:

Let C be the maximum cut in the run $\langle E, \to \rangle$. It is easy to verify that C is consistent (as in any computation). Since each I_i is the last interval and is also a false interval, $\neg(\phi_{disj}(C))$.

Case 2: $\exists i : I_i.next \neq null$:

Let e be the minimum event in $\langle E, <^c \rangle$ such that $e = I_i.next.first$ for some i. Let $C' = lcc(e)$. Consider any $j \in \{1, \cdots n\}$ such that $j \neq i$. Consider two cases:

Case a: $I_j.prev = null$: In this case $I_j.first = \bot_j$. Therefore, $I_j.first \preceq_j C'[j]$.

Case b: $I_j.prev \neq null$: Since $I_j \mapsto I_i$, we have $I_j.first \to e$. Further, since $\to \subseteq <^c$, we have $I_j.first <^c e$. Therefore, by the consistency of C', we have $I_j.first \preceq_j C'[j]$.

Therefore, in both cases: $I_j.first \preceq_j C'[j]$ –[A]

Consider two cases:

Case a: $I_j.next = null$: Clearly, $C'[j] \preceq_j I_j.last$.

Case b: $I_j.next \neq null$: Since $e <^c I_j.next.first$ (by the definition of e) and since $e = ge(C')$ (by Lemma 8), we have $C'[j] <^c e <^c I_j.next.first$. Therefore, $C'[j] \preceq_j I_j.last$.

Therefore, in both cases: $C'[j] \preceq I_j.last$ –[B]

By [A] and [B], we have $C'[j] \in I_j$ and therefore $\neg\alpha_j(C'[j])$. Let $C = lcc\text{-}prev(e)$. Since, $\forall j : j \neq i : C[j] = C'[j]$, we have: $\neg\alpha_j(C[j])$. Further, since $e = I_i.next.first$, $C[i] = I_i.last$, and therefore, $\neg\alpha_i(C[i])$. So, $\neg\phi_{disj}(C)$. Also, by Lemma 6, C is consistent.

□

Lemma 11 *The algorithm in Figure 4 is well-specified.*

Proof: We first check that all terms in the algorithm are well-defined. The second condition, $I_i.next \mapsto crossed$, in L21 is well-defined assuming that the first condition, $I_i.next \neq null$, is checked prior to it. The only remaining non-trivial check is for the term $curr.next.first$ at line L31. To show that it is well-defined, first note that, for all the intervals *anchor* added to *chain* at line L19, $anchor.next \neq null$ since $anchor.\alpha = true$ (by lines L13 and L16) and $N_{anchor.proc} \neq null$ (by terminating condition L12). Therefore, immediately after line L27, for all intervals X in *chain*, we have: $(X = chain.head) \ \vee \ (X.next \neq null)$. Lastly, note that the interval *chain.head* immediately after line L27 is never assigned to *curr* (by L28).

We next check termination. The while loop in lines L21-L22 terminates since I_i advances once per iteration. The while loop in lines L12-L23 terminates since $I_{crossed.proc}$ advances by at least one interval in each iteration. The loop in lines L29-L32 terminates since *chain* reduces by one interval in each iteration. □

Lemma 12 *The algorithm in Figure 4 maintains the following invariants on* $I_1, I_2 \cdots, I_n$ *(except in lines L20-L23 while* I_i*'s are being updated):*

INV1: $\forall i, j$: *if* $I_i.next \neq null$ *and* $I_j.prev \neq null$ *then* $I_i.next \not\mapsto I_j.prev$

INV2: $\forall i$: *if* $I_i.prev \neq null$ *then* $I_i.\alpha = true \ \vee \ (\exists j : \ I_i \mapsto I_j.prev \ \wedge \ I_j.\alpha = true)$

Proof: (structural induction treating L20-L23 as atomic)

Base Case: Initially, $\forall j : I_j = \mathcal{I}_j.head$ and, therefore, $\forall j : I_j.prev = null$ and INV1 and INV2 follow.

Inductive Case: The only statements that update any I_i are L20-L23. Therefore, we assume that INV1 and INV2 are true before an occurrence of L20-L23. We denote the values of the variables I_i at this point by I'_i. We must prove that INV1 and INV2 are true after the occurrence of L20-L23. We denote the values of the variables I_i at this point by I_i

To prove that INV1 holds, let $I_i.next \neq null$ and $I_j.prev \neq null$ and consider two cases:

Case 1: $I_j \neq I'_j$**:** Therefore, the variable I_j was updated at either L22 or L23.

- **Case a:** I_j **updated at** L23**:** Therefore, $I_j.prev = crossed$. Also, by the terminating condition of while loop L21 and since $I_i.next \neq null$, $I_i.next \not\mapsto crossed$. Therefore, $I_i.next \not\mapsto I_j.prev$.
- **Case b:** I_j **updated at** L22**:** Therefore, $I_j \mapsto crossed$. Since $I_j.prev \neq null$ and $crossed.next \neq null$ (since $anchor \not\mapsto crossed$ - line L13, L16), we have: $I_j.first \rightarrow crossed.next.first$. –[A]

  Also, by the terminating condition of loop L21 and since $I_i.next \neq null$, $I_i.next \not\mapsto crossed$. –[B]

From [A] and [B], we have $I_i.next.first \not\to I_j.first$. Therefore, $I_i.next \not\mapsto I_j.prev$.

Case 2: $I_j = I'_j$**:** We know that $I'_i.next \neq null$ since $I_i.next \neq null$ and $I'_i \preceq I_i$ (since variable I_i is always advanced). We also know that $I'_j.prev \neq null$ since $I'_j = I_j$. Therefore, by the inductive hypothesis, we have: $I'_i.next \not\mapsto I'_j.prev$. Since $I_j = I'_j$, we further have: $I'_i.next \not\mapsto I_j.prev$. Therefore, $I'_i.next.first \not\to I_j.first$. Since $I'_i \preceq I_i$, we have $I'_i.next.first \underline{\to} I_i.next.first$. Therefore, by transitivity, $I_i.next.first \not\to I_j.first$. It follows that $I_i.next \not\mapsto I_j.prev$.

We next prove that INV2 holds after lines L20-L23. Let $I_i.prev \neq null$. We have two cases:

Case 1: $I_i = I'_i$**:** By the inductive hypothesis, let j be such that $I_i \mapsto I'_j.prev \;\wedge\; I'_j.\alpha = true$. If $I_j = I'_j$, then INV2 clearly holds. Suppose $I_j \neq I'_j$. Therefore, either I_j was updated at L22 or at L23.

Case a: I_j **updated at** L23**:** Therefore, $I_j = crossed.next$. Since $crossed.\alpha = false$, we have $I_j.\alpha = true$. Further, since $I_i \mapsto I'_j.prev$ and $I'_j \preceq_j crossed$, we have: $I_i \mapsto I_j.prev$. Therefore, INV2 holds.

Case b: I_j **updated at** L22**:** From the terminating condition of the loop at L21, we have: $I'_j.next \mapsto crossed$ and so, $I'_j.next.first \to crossed.next.first$. –[A]

Since $I_i \mapsto I'_j.prev$ and $I_i.prev \neq null$, we have: $I_i.first \to I'_j.first$. –[B]

Using [A] and [B] and transitivity, we have: $I_i.first \to crossed.next.first$ and so, $I_i \mapsto crossed$. Further, $crossed.next.\alpha = true$ and so INV2 follows (instantiating $crossed.proc$ for j).

Case 2: $I_i \neq I'_i$**:** In this case, either I_i was updated in L23 or L22. If at L23, then $I_i.\alpha = true$ and INV2 follows. If at L22, then by the loop condition at L21, we have: $I_i \mapsto crossed$. Therefore, INV2 follows (instantiating $crossed.proc$ for j).

□

Lemma 14 *If the algorithm in Figure 4 terminates normally, then* $\langle E, \to^c\rangle$ *is a controlling computation of* ϕ_{disj} *in* $\langle E, \to\rangle$*, where* $\to^c$ *is the transitive closure of the union of* $\to$ *with the set of edges in* $\mathcal{O}$*.*

Proof: First, we show that $\langle E, \rightarrow^c \rangle$ is a computation. Just before loop L29-32, we have $\forall I \in chain : I.prev \neq null$ (by lines L17-L18 and and L25-L26). Therefore, in line L31, $prev.first \in E$ (i.e. $prev.first \notin \bot$). Clearly, the same holds for $curr.next.first$. Therefore, $\mathcal{O}$ determines a set of edges on E. Let $\rightsquigarrow^c$ represent this set of edges.

We show by contradiction that $\rightarrow^c$ is an irreflexive partial order. Since $\rightarrow^c$ is transitive by definition, we assume that it is reflexive. Therefore, $\rightarrow \cup \rightsquigarrow^c$ has a cycle, say $\mathcal{C}$. Among the edges in $\mathcal{C}$ let $X.first \rightsquigarrow^c Y.next.first$ be the edge such that X is the latest to be added to *chain*. Let $U.first \rightsquigarrow^c V.next.first$, be the next $\rightsquigarrow^c$ edge in the cycle $\mathcal{C}$ (note that $U = X$ if there is only one $\rightsquigarrow^c$ edge in $\mathcal{C}$). Consider the algorithm step in which X is the *anchor* and added to *chain*. Since Y was the *anchor* immediately before X, and since the *anchor* is not affected by the update to I_i's in lines L20-L23, we have: $I_{Y.proc} = Y$. –[A]

Since U cannot be added after X (by the definition of X), we have: $U \preceq I_{U.proc}$. –[B]

Since $U.first \rightsquigarrow^c V.next.first$ is the next $\rightsquigarrow^c$ edge in the cycle $\mathcal{C}$ following the edge $X.first \rightsquigarrow^c Y.next.first$, we have: $Y.next.first \rightarrow U.first$. (Note: $Y.next.first \neq U.first$ since both Y and U are true intervals). Therefore, using [A] and [B], we have: $I_{Y.proc}.next.first \rightarrow I_{U.proc}.first$ –[C]

Since $\rightsquigarrow^c$ is a set of edges on E, we know that $U.prev \neq null$. Therefore, using [B], we have: $I_{U.proc}.prev \neq null$. Therefore, applying INV1: $I_{Y.proc}.next \not\rightarrow I_{U.proc}.prev$ and, so: $I_{Y.proc}.next.first \not\rightarrow I_{U.proc}.first$, which contradicts [C]. Therefore, $\rightarrow^c$ is an irreflexive partial order and $\langle E, \rightarrow^c \rangle$ is a computation.

It remains to prove that $\langle E, \rightarrow^c \rangle$ is a controlling computation. Let C be a consistent cut in $\langle E, \rightarrow^c \rangle$. Let I be the last interval to be added to *chain* such that: $I.first \preceq C[I.proc]$ (such an I must exist, since the initial conditions together with lines L17-L18 ensure that for first interval X, $X.first \in \bot$). Consider two cases:

Case 1: ***I is the last interval added to*** *chain:* Since I was added in line L27, $I.next = null$ (by L24). Therefore, $C[I.proc] \preceq I.last$. Together with the definition of I, this implies that $C[I.proc] \in I$. Therefore $\alpha_{I.proc}(C[I.proc])$ and so, $\phi_{disj}(C)$.

Case 2: ***I is not the last interval added to*** *chain:* Let I' be the next interval to be added to *chain*. Therefore, $I'.first \rightsquigarrow^c I.next.first$ (by L31). –[A]

By the definition of I, we have: $C[I'.proc] \prec I'.first$. Therefore, since C is consistent in $\langle E, \rightarrow^c \rangle$, $I'.first \notin C.past$. In particular, $I'.first \not\rightarrow^c C[I.proc]$.

From [A], we must have: $C[I.proc] \prec I.next.first$. Together with the definition of I, we have $C[I.proc] \in I$. Therefore $\alpha_{I.proc}(C[I.proc])$ and so, $\phi_{disj}(C)$.

□

Theorem 4 *Let $\langle \mathcal{I}, \mapsto \rangle$ be the interval graph of a computation $\langle E, \rightarrow \rangle$ under local predicates $critical_1, critical_2, \cdots, critical_n$. If $\langle \mathcal{I}, \mapsto \rangle$ contains a non-trivial cycle of critical sections, then there is no controlling computation of ϕ_{mutex} in $\langle E, \rightarrow \rangle$.*

Proof: We prove the result by contradiction. Let X be a non-trivial cycle of critical sections in $\langle \mathcal{I}, \mapsto \rangle$ and let $\langle E, \rightarrow^c \rangle$ be a controlling computation of ϕ_{mutex} in $\langle E, \rightarrow \rangle$. Let $<^c$ be a linearization of $\rightarrow^c$ forming a controlling run $\langle E, <^c \rangle$ (by Lemma 2).

We next show that there is a consistent cut C in the run $\langle E, <^c \rangle$ such that $\neg\phi_{mutex}(C)$, thus, contradicting the fact that $\langle E, <^c \rangle$ is a controlling computation. Consider two cases:

Case 1: $\forall CS \in X : CS.next = null$ **:**
Let C be the maximum cut in the run $\langle E, \rightarrow \rangle$. It is easy to verify that C is consistent (as in any computation). Since X has at least two critical sections, $\neg\phi_{mutex}(C)$.

Case 2: $\exists CS \in X : CS.next \neq null$ **:**
Let e be the minimum event in $\langle E, <^c \rangle$ such that $e = CS.next.first$ for some $CS \in X$. Let CS' be the critical section preceding CS in the cycle X. Therefore, $CS' \mapsto CS$. Let $C' = lcc(e)$. Consider two cases:

Case a: $CS'.prev = null$: Since $CS'.first \in \perp$, we have $CS'.first \preceq C'[CS'.proc]$.

Case b: $CS'.prev \neq null$: Since $CS' \mapsto CS$, we have $CS'.first \rightarrow e$ and so, $CS'.first <^c e$. Therefore, since C' is consistent, $CS'.first \preceq C'[CS'.proc]$

In both cases, $CS'.first \preceq C'[CS'.proc]$. –[A]

Consider two cases:

Case a: $CS'.next = null$: Clearly, $C'[CS'.proc] \preceq CS'.last$.

Case b: $CS'.next \neq null$: Since $e <^c CS'.next.first$ (by the definition of e) and since $e = ge(C')$ (by Lemma 8), we have $C'[CS'.proc] <^c e <^c CS'.next.first$. Therefore, $C'[CS'.proc] \preceq CS'.last$.

In both cases, $C'[CS'.proc] \preceq CS'.last$ –[B]

By [A] and [B], we have $C'[CS'.proc] \in CS'$ and therefore $critical(C'[CS'.proc])$. Let $C = lcc\text{-}prev(e)$. Since, $\forall j : j \neq e.proc : C[j] = C'[j]$, we have: $critical(C[CS'.proc])$. Further, since $e = CS.next.first$, $C[CS.proc] = CS.last$, and therefore, $critical(C[CS.proc])$. So, $\neg\phi_{mutex}(C)$. Also, by Lemma 6, C is consistent.

□

Theorem 5 *Let $\langle \mathcal{I}, \mapsto \rangle$ be the interval graph of a computation $\langle E, \rightarrow \rangle$ under local predicates $critical_1, critical_2, \cdots, critical_n$. If $\langle \mathcal{I}, \mapsto \rangle$ does not contain a non-trivial cycle of critical sections, then there is a controlling computation of ϕ_{mutex} in $\langle E, \rightarrow \rangle$.*

Proof: Since there are no non-trivial cycles of critical sections in $\langle \mathcal{I}, \mapsto \rangle$, we can topologically sort the critical sections to form a sequence: $CS_1, CS_2, \cdots CS_m$ such that $CS_i \mapsto CS_j \Rightarrow i \leq j$. –[A]

Define a relation $\rightsquigarrow^c$ on events in E as follows. For each $i \in \{1, \cdots, m-1\}$, if $CS_i.next \neq null$ and $CS_{i+1}.prev \neq null$, then $CS_i.next.first \rightsquigarrow^c CS_{i+1}.first$. Further, no other events are related by $\rightsquigarrow^c$.

Define $\rightarrow^c$ to be the transitive closure of $\rightarrow \cup \rightsquigarrow^c$. We first show that $\langle E, \rightarrow^c \rangle$ is a computation and next show that it is a controlling computation.

If $\langle E, \rightarrow^c \rangle$ is not a computation, then $\rightarrow^c$ is not an irreflexive partial order. Since $\rightarrow^c$ is transitive by definition, it must be reflexive. Therefore, $\rightarrow \cup \rightarrow^c$ must have a cycle, say X. Since $\langle E, \rightarrow \rangle$ is a computation, there can be no cycle that involves no $\rightsquigarrow^c$ edges. Therefore, let $CS_i.next.first \rightsquigarrow^c CS_{i+1}.first$ be the $\rightsquigarrow^c$ edge in X such that i has the maximum value. Let $CS_j.next.first \rightsquigarrow^c CS_{j+1}.first$ be the next $\rightsquigarrow^c$ edge in the cycle X (note that it is possible that $i = j$). Therefore, $CS_{i+1}.first \underline{\rightarrow} CS_j.next.first$. Since by definition two critical sections are not contiguous, the equality cannot hold. Therefore, $CS_{i+1}.first \rightarrow CS_j.next.first$ and so, $CS_{i+1} \mapsto CS_j$. By [A], we have: $i + 1 \leq j$ and so $i < j$. This contradicts the choice of i as maximum. Therefore, $\langle E, \rightarrow^c \rangle$ is a computation.

Before showing that $\langle E, \rightarrow^c \rangle$ is a controlling computation, we first prove a claim:

Claim 1: For all critical sections CS_i and CS_j such that $i < j$, if $CS_i.next \neq null$ and $CS_j.prev \neq null$ then $CS_i.next.first \rightarrow^c CS_j.first$.

Proof: (by induction on $j - i$)
Base Case: $j - i = 1$

Directly from the definition of $\leadsto^c$.

Inductive Case: $j - i > 1$

Consider two critical sections CS_i and CS_j such that $i < j$ and $CS_i.next \neq null$ and $CS_j.prev \neq null$. We have two cases:

Case 1: $CS_{j-1}.prev = null$:
Therefore, by definition, $CS_{j-1} \mapsto CS_i$. By [A], we have $j - 1 \leq i$. Together with $i < j$, this implies $j = i+1$, which violates the inductive case assumption.

Case 2: $CS_{j-1}.prev \neq null$:
By the inductive hypothesis, we have: $CS_i.next.first \rightarrow^c CS_{j-1}.first$. –[IH]
Consider two sub-cases:

Case a: $CS_{j-1}.next = null$:
Therefore, by definition, $CS_j \mapsto CS_{j-1}$. By [A], we have $j \leq j-1$ giving a contradiction.

Case b: $CS_{j-1}.next \neq null$:
Since by the choice of j, $CS_j.prev \neq null$, we have by the definition of $\leadsto^c$: $CS_{j-1}.next.first \leadsto^c CS_j.first$. Together with [IH], by transitivity, this gives: $CS_i.next.first \rightarrow^c CS_j.first$.

□ (Proof of Claim)

We now show that $\langle E, \rightarrow^c \rangle$ is a controlling computation. Suppose it is not. Then, let C be a consistent cut in $\langle E, \rightarrow^c \rangle$ such that $\neg\phi_{mutex}(C)$. Therefore, we have distinct i and j such that $critical(C[i])$ and $critical(C[j])$.

Let $C[i] \in CS_k$ and $C[j] \in CS_l$. Without loss of generality, $k < l$. (Since $i \neq j$, $k \neq l$).

Consider two cases:

Case 1: $CS_k.next = null \ \vee \ CS_l.prev = null$:
Therefore, by definition, $CS_l \mapsto CS_k$. By [A], $l \leq k$, contradicting $k < l$.

Case 2: $CS_k.next \neq null \ \wedge \ CS_l.prev \neq null$:
Therefore, by Claim 1, we have: $CS_k.next.first \rightarrow^c CS_l.first$. Since $C[i] \in CS_k$, we have: $C[i].next \preceq CS_k.next.first$ and since $C[j] \in CS_j$, we have: $CS_l.first \preceq C[j]$. Therefore, by transitivity, $C[i].next \rightarrow^c C[j]$, violating the consistency of C.

□

Lemma 15 *The algorithm in Figure 5 is well-specified.*

Proof: To prove that the algorithm is well-specified, the only non-trivial check is for the term $curr.next.first$ at line L23. To show that $curr.next \neq null$ at this point, we prove that, after the termination of loop L13-L19, for all critical sections c in $chain$: $(c.next = null) \Rightarrow (c = chain.head)$. Let c be a critical section in $chain$ such that $c.next = null$. Consider the point when c is about to be added to $chain$ in line L18. Since $c \in valid\_cs$, we have $\forall c' \in valid\_cs : (c' \neq c) \Rightarrow (c' \not\mapsto c)$ However, since $c.next = null$, we have $\forall c' \in \mathcal{C} : c' \mapsto c$. Therefore, c is the only element of $valid\_cs$ at this point. It follows that after line L19, $CS = \emptyset$ and the loop terminates. Therefore, after the termination of loop L13-L19, $c = chain.head$.

The algorithm terminates since the loop L13-L19 advances one critical section in each iteration and the loop L20-L24 decreases $chain$ by one critical section in each iteration. □

Lemma 17 *If the algorithm in Figure 5 terminates normally, then $\langle E, \rightarrow^c \rangle$ is a controlling computation of ϕ_{mutex} in $\langle E, \rightarrow \rangle$, where $\rightarrow^c$ is the transitive closure of the union of $\rightarrow$ with the set of edges in $\mathcal{O}$.*

Proof: We show that, at the termination of loop L13-L19, $chain$ is a topological sort of the graph $\langle \mathcal{C}, \mapsto \rangle$. Then, since lines L20-L24 construct the same $\rightsquigarrow^c$ relation as in the proof of Theorem 5, $\langle E, \rightarrow^c \rangle$ is a controlling computation of ϕ_{mutex} in $\langle E, \rightarrow \rangle$.

Let $chain.set$ denote the set of intervals in $chain$. First, we note that the algorithm maintains the following invariant:

INV1: $\forall c \in \mathcal{C} : c \prec CS_{c.proc} \Rightarrow c \in chain.set$

Next, consider the point just before L17, when $crossed$ is about to be added to $chain$. Let c be any critical section in $\mathcal{C} - chain.set$ such that $c \neq crossed$. By INV1, we have $CS_{c.proc} \preceq c$. If $c \mapsto crossed$ then $CS_{c.proc} \mapsto crossed$, contradicting $crossed \in valid\_cs$. Therefore $c \not\mapsto crossed$. Therefore, $crossed$ is minimal in $\langle C - chain.set, \mapsto \rangle$. Therefore, $chain$ is a topological sort of $\langle C, \mapsto \rangle$. □

Theorem 7 *Let $\langle \mathcal{I}, \mapsto \rangle$ be the interval graph of a computation $\langle E, \rightarrow \rangle$ under local predicates $critical_1, critical_2, \cdots, critical_n$. If $\langle \mathcal{I}, \mapsto \rangle$ contains a non-trivial cycle* *of critical sections containing at least one write-critical section, then there is no controlling computation of ϕ_{rw} in $\langle E, \rightarrow \rangle$.*

Proof:

We prove the result by contradiction. In interval graph $\langle \mathcal{I}, \mapsto \rangle$, let X be a non-trivial cycle of critical sections containing a write-critical section W. Let $\langle E, \rightarrow^c \rangle$ be a controlling computation of ϕ_{mutex} in $\langle E, \rightarrow \rangle$. Let $<^c$ be a linearization of $\rightarrow^c$ forming a controlling run $\langle E, <^c \rangle$ (by 2).

We next show that there is a consistent cut C in the run $\langle E, <^c \rangle$ such that $\neg\phi_{rw}(C)$ contradicting the fact that $\langle E, <^c \rangle$ is a controlling computation. Consider two cases:

Case 1: $W.next = null$ **:**

Let CS be the critical section following W in X. Since X is non-trivial, $CS \neq W$. Consider two cases:

Case a: $CS.next = null$**:** Let C be the maximum cut in the run $\langle E, \rightarrow \rangle$. It is easy to verify that C is consistent (as in any computation). Since $W.next = null$ and $CS.next = null$, $\neg\phi_{rw}(C)$.

Case b: $CS.next \neq null$**:** Let $e = CS.next.first$ and $C' = lcc(e)$. Consider two cases:

$W.prev = null$: Since $W.first \in \perp$, we have $W.first \preceq C'[W.proc]$.

$W.prev \neq null$: Since $W \mapsto CS$, we have $W.first \rightarrow e$ and so, $W.first <^c e$. Therefore, since C' is consistent, $W.first \preceq C'[W.proc]$

In both cases, $W.first \preceq C'[W.proc]$ –[A]

Since $W.next = null$, we have: $C'[W.proc] \preceq W.last$ –[B]

By [A] and [B], we have $C'[W.proc] \in W$ and therefore: $write\_critical(C'[W.proc])$. Let $C = lcc\text{-}prev(e)$. Since, $\forall j : j \neq e.proc : C[j] = C'[j]$, we have: $write\_critical(C[W.proc])$. Further, since $e = CS.next.first$, $C[CS.proc] = CS.last$, and therefore, $critical(C[CS.proc])$. So, $\neg\phi_{rw}(C)$. Also, by Lemma 6, C is consistent.

Case 2: $W.next \neq null$ **:**

Let $C' = lcc\text{-}prev(W.next.first)$. If there is a CS in X distinct from W such that: $CS.first \preceq C'[CS.proc] \preceq CS.last$, then since C' is consistent and $\neg\phi_{rw}(C)$, we are done. Therefore, assume that for all critical sections CS in X distinct from W: $(C'[CS.proc] \prec CS.first) \vee (CS.last \prec C'[CS.proc])$. –[D]

Let CS_1 and CS_2 be any two consecutive critical sections in X. Suppose that $CS_2.first \preceq C'[CS_2.proc]$. –[E]

Therefore, applying [D], we have: $CS_2.last \prec C'[CS_2.proc]$. –[F]

Consider two cases:

Case a: $CS_1.prev = null$**:** In this case, since $CS_1.first \in \perp$ we have: $CS_1.first \preceq C'[CS_1.proc]$.

Case b: $CS_1.prev \neq null$**:** Since $CS_1 \mapsto CS_2$ and since $CS_2.next \neq null$ (from [F]), we have: $CS_1.first \rightarrow CS_2.next.first$. This, together with [F] and transitivity gives us: $CS_1.first \in C'.past$. Since C' is consistent, $CS_1.first \preceq C'[CS_1.proc]$.

Therefore, in both cases: $CS_1.first \preceq C'[CS_1.proc]$. Combining this with our supposition [E], we have: $\forall CS_1, CS_2 \in X : CS_2.first \preceq C'[CS_2.proc] \Rightarrow CS_1.first \preceq C'[CS_1.proc]$. Therefore, since $W.first \preceq C'[W.proc]$ (by the definition of C'), we have: $\forall CS \in X : CS.first \preceq C'[CS.proc]$. Combining this with [D] we have: $\forall CS \in X : CS = W \lor CS.next.first \preceq C'[CS.proc]$. –[G]

Let CS be the critical section following W in the cycle X. Let $C = lcc\text{-}prev(CS.next.first)$. From [G], $CS.next.first \preceq C'[CS.proc]$, and so: $C[CS.proc] \prec C'[CS.proc]$. Therefore, by Lemma 7 we have: $C < C'$. So, we have: $C[W.proc] \preceq C'[W.proc]$. Using the definition of C': $C[W.proc] \preceq W.last$. –[H]

Consider three cases:

Case a: $W.prev = null$**:** Therefore, $W.first \in \perp$, and so: $W.first \preceq C[W.proc]$.

Case b: $W.proc = CS.proc$**:** In this case, $W \prec CS$ and, therefore, $W.first \preceq C[w.proc]$.

Case c: $W.prev \neq null \land W.proc \neq CS.proc$ **:** Since $W \mapsto CS$ and $CS.next \neq null$ (because of [G]), we have: $W.first \rightarrow CS.next.first$. Therefore, $W.first \in lcc(CS.next.first).past$. By the consistency of $lcc(CS.next.first)$, we have: $W.first \preceq lcc(CS.next.first)[W.proc]$. Since $W.proc \neq CS.proc$, by the definition of *lcc-prev*: $lcc\text{-}prev(CS.next.first)[W.proc] = lcc(CS.next.first)[W.proc]$. Therefore, $W.first \preceq C[W.proc]$.

In all cases: $W.first \preceq C[W.proc]$. This, together with [H], gives us: $C[W.proc] \in W$. Therefore, $write\_critical(C[W.proc])$, and by the definition of C, we have

$critical(C[CS.proc])$. Therefore, $\neg\phi_{rw}(C)$. Further C is consistent by definition.

□

Theorem 8 *Let $\langle \mathcal{I}, \mapsto \rangle$ be the interval graph of a computation $\langle E, \rightarrow \rangle$ under local predicates $critical_1, critical_2, \cdots, critical_n$. If $\langle \mathcal{I}, \mapsto \rangle$ does not contain a non-trivial cycle of critical sections containing at least one write-critical section, then there is a controlling computation of ϕ_{rw} in $\langle E, \rightarrow \rangle$.*

Proof: Let $\mathcal{C}$ be the set of critical sections in $\mathcal{I}$. Let $\mapsto_{\mathcal{C}}$ denote the relation $\mapsto$ restricted to the set $\mathcal{C}$. Let $\mathcal{S}$ be the set of strongly connected components of the graph $\langle \mathcal{C}, \mapsto_{\mathcal{C}} \rangle$. Let $\hookrightarrow$ be a relation defined on strongly connected components as:

$$SCC \hookrightarrow SCC' \equiv \exists CS \in SCC, CS' \in SCC' : CS \mapsto CS'$$

Clearly, $\hookrightarrow$ is a partial order on $\mathcal{S}$. Therefore, we can topologically sort the graph $\langle \mathcal{S}, \hookrightarrow \rangle$ to give a sequence of strongly connected components:

$$SCC_1, SCC_2, \cdots, SCC_l \text{ such that } SCC_i \hookrightarrow SCC_j \Rightarrow i \le j \quad \text{–[A]}$$

Define a relation $\leadsto^c$ on events in E as the set of tuples:

$$\{ (CS_i.next.first, CS_j.first) \mid CS_i.next \neq null \wedge CS_j.prev \neq null \wedge \exists k \in \{1, \cdots, l-1\} : CS_i \in SCC_k \wedge CS_j \in SCC_{k+1} \}$$

Define $\rightarrow^c$ to be the transitive closure of $\rightarrow \cup \leadsto^c$.

We first show by contradiction that $\langle E, \rightarrow^c \rangle$ is a computation. If it is not, then there must be a cycle, say X, in the graph $\langle E, \rightarrow \cup \leadsto^c \rangle$. Since $\langle E, \rightarrow \rangle$ is a computation, X must have at least one $\leadsto^c$ edge. Consider the set of $\leadsto^c$ edges in X, and let $CS_i.next.first \leadsto^c CS_j.first$ be the edge such that CS_i belongs to the maximum strongly connected component (among such edges) in the topological sorted order (i.e. $CS_i \in SCC_k$ and k has the maximum value). Let $CS_i \in SCC_k$, and so, $CS_j \in SCC_{k+1}$. Let $CS_a.next.first \leadsto^c CS_b.first$ be the next $\leadsto^c$ edge in the cycle X (note that it is possible that $a = i$). Therefore, $CS_j.first \underline{\rightarrow} CS_a.next.first$. Since, by definition, two critical sections cannot be contiguous, the equality cannot hold. Therefore, $CS_j.first \rightarrow CS_a.next.first$, and so, by definition, $CS_j \mapsto CS_a$. So, if $CS_a \in SCC_p$, we have: $SCC_{k+1} \hookrightarrow SCC_p$. Therefore, by [A], $k+1 \le p$ and so $k < p$. This contradicts the choice of k as maximum. Therefore, $\langle E, \rightarrow^c \rangle$ is a computation.

Before showing that $\langle E, \rightarrow^c \rangle$ is a controlling computation, we first prove a claim:

Claim 1: For all critical sections $CS \in SCC_p$ and $CS' \in SCC_q$ such that $p < q$, if

$CS.next \neq null$ and $CS'.prev \neq null$, then: $CS.next.first \;\rightarrow^c\; CS'.first$.
Proof: (by induction on $q - p$)
Base Case: (q - p = 1): Directly from the definition of $\rightsquigarrow^c$.
Inductive Case: (q - p ¿ 1): Let $CS \in SCC_p$ and $CS' \in SCC_q$ such that $p < q$, if $CS.next \neq null$ and $CS'.prev \neq null$. Let CS'' be any critical section in SCC_{q-1}. We have two cases:

Case 1: ($CS''.prev = null$):

Therefore, by definition: $CS'' \mapsto CS$, and so $SCC_{q-1} \;\hookrightarrow\; SCC_p$. Using [A], we have: $(q-1) \leq p$. Together with $p < q$, this implies that: $q = p + 1$ which violates the inductive case assumption.

Case 2: ($CS''.prev \neq null$):

Subcase 2.1: ($CS''.next = null$)

Therefore, by definition, $CS' \;\mapsto\; CS''$, and so $SCC_q \;\hookrightarrow\; SCC_{q-1}$. Using [A], it follows that $q \leq (q-1)$. Thus, we have a contradiction.

Subcase 2.2: ($CS''.next \neq null$)

Therefore, $CS''.next.first \rightsquigarrow^c CS'.first$. Further, by the inductive hypothesis, $CS.next.first \;\rightarrow^c\; CS''.first$. If follows, by transitivity, that: $CS.next.first \;\rightarrow^c\; CS''.first$.

$\square$ (Proof of Claim 1)

We now show that $\langle E, \rightarrow^c \rangle$ is a controlling computation. Suppose it is not. Then, let C be a consistent cut of $\langle E, \rightarrow^c \rangle$ such that $\neg\phi_{rw}(C)$. Therefore, we have distinct i and j such that *write_critical*($C[i]$) and *critical*($C[j]$). Let $C[i] \in CS$ and $C[j] \in CS'$. Further let, $CS \in SCC_p$ and $CS \in SCC_q$. Without loss of generality, let $p \leq q$. Consider two cases:

Case 1: $p = q$:

In this case, both CS and CS' are in the same strongly connected component of $\langle \mathcal{C}, \mapsto_{\mathcal{C}} \rangle$. Further CS is a write-critical section and CS' is a distinct critical section. Therefore, there is a non-trivial cycle of critical sections in $\langle \mathcal{I}, \mapsto \rangle$ containing a write-critical section, thus, contradicting the assumptions.

Case 2: $p < q$:

Subcase 2.1: *$CS.next = null \;\vee\; CS'.prev = null$:*

Therefore, by definition, $CS' \;\mapsto\; CS$, and so, $SCC_q \;\hookrightarrow\; SCC_p$. Using [A], we have: $q \leq p$ which is a contradiction.

Subcase 2.2: $CS.next \neq null \ \wedge \ CS'.prev \neq null$:

Therefore, by Claim 1, we have: $CS.next.first \ \rightarrow^c \ CS'.first$. Since $C[i] \in CS$, we have: $C[i].next \preceq CS.next.first$ and since $C[j] \in CS'$, we have: $CS'.first \preceq C[j]$. Therefore, by transitivity, $C[i].next \rightarrow^c C[j]$ violating the consistency of C.

□

Theorem 9 *The predicate control problem for the independent mutual exclusion predicate is NP-Complete.*

Proof:

Let PC-IND denote the predicate control problem for the independent mutual exclusion predicate. PC-IND is in NP for the same reasons that the general Predicate Control problem is in NP. We prove that PC-IND is NP-Hard by transforming 3SAT.

Let $U = \{u_1, u_2, \cdots u_w\}$ be a set of variables, and $C = \{c_1, c_2, \cdots c_x\}$ be a set of clauses, forming an instance of 3SAT. Let l_{i1}, l_{i2}, l_{i3} be the three literals in clause c_i.

To construct an instance of PC-IND from this instance of 3SAT, we first construct a set of events E as follows:

- For each clause c_i define 6 processes with event sets $E_{i1}, E_{i2}, E_{i3}, \cdots, E_{i6}$. Each E_{ip} has two events denoted by e_{ip1} and e_{ip2}. Let $E_i = E_{i1} \cup E_{i2} \cup \cdots \cup E_{i6}$.
- For each variable u_r define 2 processes with event sets E'_{r1}, E'_{r2}. Each E'_{rq} has two events denoted by e'_{rq1} and e'_{rq2}. Let $E'_r = E'_{r1} \cup E'_{r2}$.

Let $E = \bigcup_i E_i \ \cup \ \bigcup_r E'_r$ and $n = |E| = 2w + 6x$.

Let $\rightsquigarrow$ be a relation on E defined as follows:

- Within each E_i: $e_{i22} \rightsquigarrow e_{i31}$, $e_{i42} \rightsquigarrow e_{i51}$, $e_{i62} \rightsquigarrow e_{i11}$
- For each literal $l_{ip} \in c_i$ such that $l_{ip} = u_r$ for some variable u_r: $e_{i(2p)1} \rightsquigarrow e'_{r22}$, $e'_{r11} \rightsquigarrow e_{i(2p-1)1}$, $e_{i(2p-1)1} \rightsquigarrow e'_{r12}$, $e'_{r21} \rightsquigarrow e_{i(2p)1}$
- For each literal $l_{ip} \in c_i$ such that $l_{ip} = \overline{u_r}$ for some variable u_r: $e_{i(2p)1} \rightsquigarrow e'_{r12}$, $e'_{r21} \rightsquigarrow e_{i(2p-1)1}$, $e_{e(2p-1)1} \rightsquigarrow e'_{r22}$, $e'_{r11} \rightsquigarrow e_{i(2p)1}$

Let $\prec$ be a relation on E such that the local events in an E_{ip} or an E_{rq} are totally ordered by $\prec$ so that: $e_{ip1} \ \prec \ e_{ip2}$ and $e'_{rq1} \prec e'_{rq2}$.

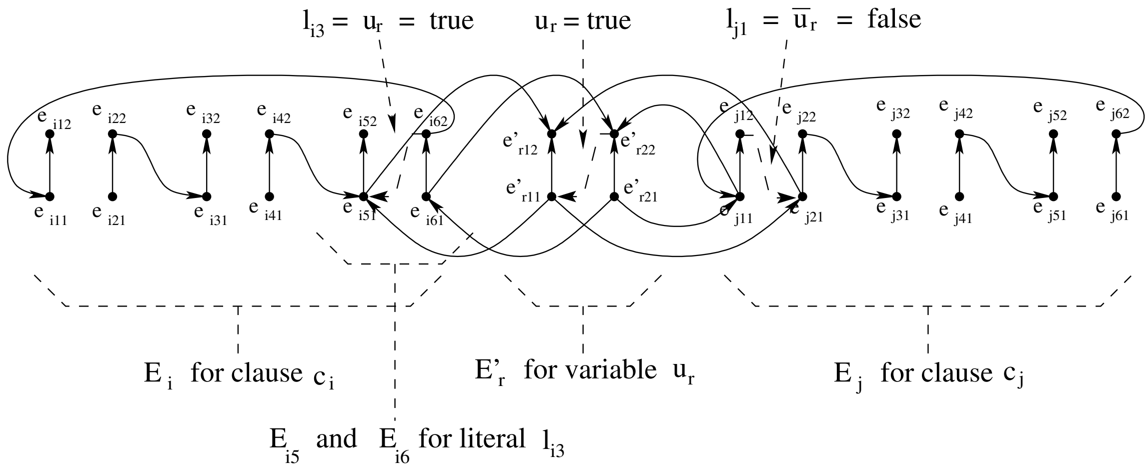

Figure 8: Proof: PC-IND is NP-Complete

Let $\rightarrow$ be the transitive closure of $\rightsquigarrow \cup \prec$. To show that $\rightarrow$ is irreflexive, we prove that there are no cycles in the graph $\langle E, \rightsquigarrow \cup \prec \rangle$. Any event in an E'_{rq} has either incoming or outgoing edges but not both. Therefore, no event in an E'_{rq} can be involved in a cycle. Since any distinct E_i and E_j have no edges between them, the only remaining possibility is a cycle within an E_i. Considering only edges within E_i, any E_{ip} has either incoming edges to its events or outgoing edges from its events but not both. Further, there are no cycles within an E_{ip}. Therefore, there is no cycle within an E_i. Thus, we have shown that $\rightarrow$ is irreflexive. Clearly, $\rightarrow$ is also transitive and so, $\langle E, \rightarrow \rangle$ is a computation.

Define local predicates $critical_1, critical_2, \cdots, critical_n$, such that each process consists of exactly one critical section containing only its first event. So, for each E_{ip}, there is exactly one critical section, $CS_{ip} = \{e_{ip1}\}$ and for each E'_{rq}, there is exactly one critical section $CS'_{rq} = \{e'_{rq1}\}$.

The critical sections in an E_i are $i\_critical$ and the critical sections in an E'_r are $(x+r)\_critical$. Therefore, there are $m = (x+w)$ independent types of critical sections. Define global predicate ϕ_{ind} based on these critical section definitions.

Thus, we have a polynomial construction from an instance $(U,\ C)$ of 3SAT to an instance $(\langle E, \rightarrow \rangle,\ \phi_{ind})$ of PC-IND. We next show that the construction is indeed a transformation.

Part 1: Let t be a truth assignment on variables in U such that t satisfies the set of clauses in C. Define a relation $\rightsquigarrow^{c1}$ on E as follows:

- For each variable u_r:

if $t(u_r) = \mathit{true}$ then $e'_{r22} \leadsto^{c1} e'_{r11}$
if $t(u_r) = \mathit{false}$ then $e'_{r12} \leadsto^{c1} e'_{r21}$

- For each literal l_{ip} in a clause c_i:
  if l_{ip} is *true* under t then $e_{i(2p)2} \leadsto^{c1} e_{i(2p-1)1}$
  if l_{ip} is *false* under t then $e_{i(2p-1)2} \leadsto^{c1} e_{i(2p)1}$

For any i, let $\mathcal{C}_i$ denote the set of 6 critical sections in E_i. Define $\mapsto_i^{c1}$ relation on $\mathcal{C}_i$ as:

$$\forall j, k \in \{1, \cdots, 6\} : \ CS_{ij} \ \mapsto_i^{c1} \ CS_{ik} \ \equiv \ e_{ij2} \leadsto^{c1} e_{ik1} \ \vee \ e_{ij2} \leadsto e_{ik1}$$

Depending on whether each l_{ip} is *true* or *false* under t, we have 8 cases (by the definition of $\leadsto^{c1}$). Since t satisfies clause c_i, at least one of l_{i1}, l_{i2}, l_{i3} must be *true* under t, and one of the cases is not possible. It is easy to verify that for each of the remaining 7 cases, there are no cycles in the graph $\langle \mathcal{C}_i, \mapsto_i^{c1} \rangle$.

Therefore, we can topologically sort $\langle \mathcal{C}_i, \mapsto_i^{c1} \rangle$. Consider a topological sort in which all the maximal critical sections are selected last. Therefore, if p_i is a permutation of $\{1, \cdots, 6\}$ representing the topological sort, then:

TOP1: $\forall j, k : \ CS_{ij} \ \mapsto_i^{c1} \ CS_{ik} \ \Rightarrow \ p_i(j) < p_i(k)$, and
TOP2: $\forall j : \ CS_{ij}$ is maximal in $\langle \mathcal{C}_i, \mapsto_i^{c1} \rangle \ \Rightarrow$
$(\ \forall k : \ p_i(k) > p_i(j) \ \Rightarrow \ CS_{ik}$ is maximal in $\langle \mathcal{C}_i, \mapsto_i^{c1} \rangle \ )$

Define a relation $\leadsto^{c2}$ on E that orders critical sections in $\mathcal{C}_i$ in the topologically sorted order as follows:

- In each E_i:
  $\forall j, k : \ (\ p_i(j) = p_i(k) - 1 \ \wedge \ CS_j \not\mapsto_i^{c1} CS_k \ ) \ \Rightarrow \ e_{ij2} \leadsto^{c2} e_{ik1}$

Let $\rightarrow^c$ be the transitive closure of $\rightarrow \ \cup \leadsto^{c1} \ \cup \ \leadsto^{c2}$. We next prove that $\rightarrow^c$ is irreflexive. This is equivalent to proving the following:
Claim 1: There is no cycle in $\prec \ \cup \leadsto \ \cup \ \leadsto^{c1} \ \cup \ \leadsto^{c2}$.
Proof: We consider three cases for a cycle and prove that each case is not possible.

Case 1: A cycle within an E'_r: The cycle cannot be within a single E'_{rq} since there is only one edge ($\prec$) between the two events. Since t assigns either *true* or *false* to variable u_r, there is only one edge ($\leadsto^{c1}$) between E'_{r1} and E'_{r2}. Therefore, there can be no cycles within an E'_r.

Case 2: A cycle within an E_i: The cycle cannot be within a single E_{ij} since there is only one edge ($\prec$) between the two events. We prove by contradiction that there cannot be a cycle intersecting multiple E_{ij}'s. Suppose there is such a cycle X. Among all the intersected E_{ij}'s, consider E_{ik} such that $p_i(k) = max\{p_i(j) : E_{ij} \cap X \neq \emptyset\}$ (i.e. CS_{ik} is the last critical section in the topological sort of $\mapsto_i^{c1}$ among all intersected $CSij$'s). Let the cycle X intersect E_{il} immediately after E_{ik}. Therefore, there must be an edge ($\leadsto$, $\leadsto^{c1}$, or $\leadsto^{c2}$) from e_{ik2} to e_{il1}. Therefore, either it is a $\leadsto^{c2}$ edge or $CS_{ik} \mapsto_i^{c1} CS_{il}$. So, using the definition of $\leadsto^{c2}$ or TOP1, we have $p_i(k) < p_i(l)$. This contradicts the choice of k.

Case 3: All other cycles: For any distinct E_i and E_j there are no edges between them. Therefore, we have only to consider cycles that intersect at least one E'_r. Let X be such a cycle intersecting E'_r. We consider the case when $t(u_r) = true$. The case $t(u_r) = false$ follows along similar lines. By the definition of $\leadsto^{c1}$, we have $e'_{r22} \leadsto^{c1} e'_{r11}$. If this edge were excluded, no event in E'_r would have both incoming and outgoing edges. Therefore, the cycle X must include this edge. From the definition of $\leadsto$, there are two possible cases for the next edge in the cycle X:

Case 1: $e'_{r11} \leadsto e_{i(2p-1)1}$ is the next edge in cycle X, for some i and p:

In this case, the literal $l_{ip} = u_r$ and, therefore, is *true* under t. Therefore, by the definition of $\leadsto$ and $\leadsto^{c1}$, there are no outgoing edges from $E_{i(2p-1)}$ in $\leadsto \cup \leadsto^{c1}$. Therefore, $CS_{i(2p-1)}$ is maximal in $\langle \mathcal{C}_i, \mapsto_i^{c1} \rangle$.
Since the cycle X contains at least one maximal critical section in $\langle \mathcal{C}_i, \mapsto_i^{c1} \rangle$, consider the critical section CS_{ik} such that: $p_i(k) = max\{p_i(j) : CS_{ij}$ is maximal in $\langle \mathcal{C}_i, \mapsto_i^{c1} \rangle \wedge E_{ij} \cap X \neq \emptyset\}$ (i.e. CS_{ik} is the last maximal critical section in the topological sort of $\langle \mathcal{C}_i, \mapsto_i^{c1} \rangle$ that is intersected by cycle X). Since CS_{ik} is maximal in $\langle \mathcal{C}_i, \mapsto_i^{c1} \rangle$, E_{ik} can have no outgoing edges in $\leadsto \cup \leadsto^{c1}$. Therefore the outgoing edge from E_{ik} in the cycle must be: $e_{ik2} \leadsto^{c2} e_{iq1}$ such that $p(q) = p(k)+1$. Since, CS_{iq} is also maximal in $\langle \mathcal{C}_i, \mapsto_i^{c1} \rangle$ (by TOP2), this contradicts the choice of $p_i(k)$.

Case 2: $e'_{r11} \leadsto e_{i(2p)1}$ is the next edge in cycle X, for some i and p:
In this case, the literal $l_{ip} = \overline{u_r}$ and, therefore, is *false* under t. Therefore, using $CS_{i(2p)}$ as the maximal critical section in X we arrive at a contradiction in a similar manner to Case 1.

End of Proof (Claim 1)

Therefore, $\langle E, \rightarrow^c \rangle$ is a computation. It is easy to verify that for any k such that $1 \leq k \leq (x + w)$, all of the k-critical sections can be arranged in a sequence

$CS_1, \cdots, CS_m$ such that $\forall i : CS_i.next.first \rightarrow^c CS_{i+1}.first$ (where m is 6 or 2 depending on whether $1 \leq k \leq x$ or $x < k \leq (x+w)$). Therefore, we can prove that $\langle E, \rightarrow^c \rangle$ is a controlling computation along similar lines to the proof of Theorem 5.

Part 2: Let $\langle E, \rightarrow^c \rangle$ be a controlling computation of ϕ_{ind} in $\langle E, \rightarrow \rangle$. Before constructing a truth assignment on U, we prove:

Claim 2: For any two distinct k-critical sections CS and CS', either

(1) $CS.next.first \rightarrow^c CS'.first$, or

(2) $CS'.next.first \rightarrow^c CS.first$,

but not both.

Proof: Clearly both (1) and (2) cannot be true since otherwise there would be a cycle in $\rightarrow^c$. Suppose both (1) and (2) are false. Let C be the consistent cut that is the least upper bound (lub) of $lcc(CS.first)$ and $lcc(CS'.first)$ (using Lemma 4). Since (1) is false, $CS.next.first \notin lcc(CS'.first).past$ and since $CS.first \prec CS.next.first$, we have $CS.next.first \notin lcc(CS.first).past$ (using Lemma 5). Therefore, $CS.next.first \notin C.past$ (by the definition of lub). Together with $CS.first \preceq C[CS.proc]$ (by the definition of lub), we have: $C[CS.proc] \in CS$. In a similar way, we can show that $C[CS'.proc] \in CS'$. Therefore, $k\_critical(C[CS.proc]) \wedge k\_critical(C[CS'.proc])$, and so: $\neg\phi_{ind}(C)$. This contradicts the fact that $\langle E, \rightarrow^c \rangle$ is a controlling computation. *End of Proof (Claim 2)*

We construct a truth assignment t on U such that for a variable u_r:

- $t(u_r) = true$, if $CS'_{r2}.next.first \rightarrow^c CS'_{r1}.first$, and
- $t(u_r) = false$, if $CS'_{r1}.next.first \rightarrow^c CS'_{r2}.first$

By Claim 2, this is a valid truth assignment.

Further, let x be a boolean function on literals defined as follows:

- $x(l_{ip}) = true$, if $CS_{i(2p)}.next.first \rightarrow^c CS_{i(2p-1)}.first$, and
- $x(l_{ip}) = false$, if $CS_{i(2p-1)}.next.first \rightarrow^c CS_{i(2p)}.first$

Again, by Claim 2, this is a valid definition.

Claim 3: $(\ x(l_{ip}) = true \ ) \ \equiv \ (\ l_{ip}$ is *true* under t)

Proof:

Case 1: $x(l_{ip}) = true$**:**

Therefore, substituting in the definition of x, we get:

$e_{i(2p)2} \rightarrow^c e_{i(2p-1)1}$ –[A]

Case a: $l_{ip} = u_r$ for some r:

Assume $t(u_r) = \mathit{false}$, then (substituting):

$e'_{r12} \rightarrow^c e'_{r21}$ –[B]

From the definition of $\rightarrow$, we have:

$e'_{r21} \rightarrow e_{i(2p)1}$ and $e_{i(2p-1)1} \rightarrow e'_{r12}$ –[C]

From [A], [B], and [C], we have a cycle contradicting the fact that $\langle E, \rightarrow^c\rangle$ is a computation. Therefore, $t(u_r) = \mathit{true}$, and so, l_{ip} is true under t.

Case b: $l_{ip} = \overline{u_r}$ for some r:

Assume $t(u_r) = \mathit{true}$, then (substituting):

$e'_{r22} \rightarrow^c e'_{r11}$ –[D]

From the definition of $\rightarrow$, we have:

$e'_{r11} \rightarrow e_{i(2p)1}$ and $e_{i(2p-1)1} \rightarrow e'_{r22}$ –[E]

From [A], [D], and [E], we have a cycle contradicting the fact that $\langle E, \rightarrow^c\rangle$ is a computation. Therefore, $t(u_r) = \mathit{false}$, and so, l_{ip} is true under t.

Case 2: $x(l_{ip}) = \mathit{false}$:

Along similar lines to Case 1, we can show that l_{ip} is false under t.

End of Proof (Claim 3)

Now, we show that the truth assignment t satisfies the set of clauses C. Consider any clause c_i. If all of $x(l_{i1})$, $x(l_{i2})$, and $x(l_{i3})$ are *false*, then by the definition of x, $\prec$ and $\rightarrow$, we have a cycle in $\rightarrow^c$: $e_{i12}, e_{i21}, e_{i22}, e_{i31}, e_{i32}, e_{i41}, e_{i42}, e_{i51}, e_{i52}, e_{i61}, e_{i62}$. Therefore, at least one literal, say $x(l_{ip})$ is *true*. It follows from *Claim 3* that l_{ip} is *true* under t. Therefore t satisfies c_i. □

Lemma 18 *The algorithm in Figure 6 is well-specified.*

Proof: To prove that the algorithm is well-specified, the only non-trivial check is for the term $curr.next.first$ at line L37. To show that $curr.next \neq null$ at this point, we prove that, after the termination of loop L13-L19, for each strongly connected component (scc) s in $chain$: $(\exists c \in s : c.next = null) \Rightarrow (s = chain.head)$. Let s be an scc in $chain$ and let $c \in s$ be a critical section such that $c.next = null$. Consider the point when s is about to be added to $chain$ in line L27. Since $s \in valid\_scc$, we have $\forall s' \in scc\_set : (s' \neq s) \Rightarrow (s' \not\hookrightarrow s)$ However, since $c.next = null$, we have $\forall c' \in \mathcal{C} : c' \mapsto c$. Therefore, $\forall s' \in scc\_set : s' \hookrightarrow s$. Therefore, s is the only element of $valid\_scc$ at this point. It follows that after line L29, $CS = \emptyset$ and the loop terminates. Therefore, after the termination of loop L19-32, $s = chain.head$.

We next show that the algorithm terminates. The loop L19-L32 terminates since in each iteration either one I_i advances over a critical section (line L29) or two distinct critical sections merge into one (line L32). The loop L34-L38 terminates since *chain* decreases by one critical section in each iteration. $\Box$

Lemma 20 *If the algorithm in Figure 6 terminates normally, then* $\langle E, \rightarrow^c \rangle$ *is a controlling computation of* ϕ_{gen_mutex} *in* $\langle E, \rightarrow \rangle$*, where* $\rightarrow^c$ *is the transitive closure of the union of* $\rightarrow$ *with the set of edges in* $\mathcal{O}$*.*

Proof: At any point in the algorithm, let $\mathcal{C}$ denote the set of regular (not general) critical sections in the input computation. Note that $\mathcal{C}$ is a constant throughout the algorithm execution. Let $\mathcal{S}$ be the set of strongly connected components (scc) in the graph $\langle \mathcal{C}, \mapsto \rangle$. We use the relations $\mapsto$ and $\hookrightarrow$ for both regular and general intervals and scc's since the usage is clear from context. Let $\mathcal{C}_\mathcal{G}$ denote the set of general critical sections in $\mathcal{I} = \bigcup_i \mathcal{I}_i$ at any stage of the algorithm (note that $\mathcal{C}_\mathcal{G}$ is a variable since it changes at the *merge* step during the algorithm).

Let *scc* be a function that maps each critical section in $\mathcal{C}$ to the strongly connected component in $\mathcal{S}$ that contains it. First, we prove that the following are invariants.

INV1: $\forall g \in \mathcal{C}_\mathcal{G} : \ \forall x, y \in g.set : \ scc(x) = scc(y)$

INV2: g is reachable from g' in the graph $\langle \mathcal{C}_\mathcal{G}, \mapsto \rangle \ \Rightarrow$

$\forall x \in g.set, x' \in g'.set :$ x is reachable from x' in the graph $\langle \mathcal{C}, \mapsto \rangle$

It is easy to prove that INV2 is a consequence of INV1 (by induction on the length of the reachability path). Therefore, we have only to prove that INV1 is an invariant. We do so by structural induction.

Initially, each general critical section in $\mathcal{C}_\mathcal{G}$ contains only one critical section and so INV1 holds. The only step at which $\mathcal{C}_\mathcal{G}$ changes is the *merge* at line L32. Consider the execution point just before line L32. Since $c \in mergeable$ we have:

$c.prev.prev \in candidate$, and –[A]

let $c' \in candidate$ such that $c \mapsto c'$. –[B]

Since *candidate* is an scc of the subgraph $\langle CS, \mapsto \rangle$ of $\langle \mathcal{C}_\mathcal{G}, \mapsto \rangle$, c' and $c.prev.prev$ are in the same scc, say s, of $\langle \mathcal{C}_\mathcal{G}, \mapsto \rangle$. Clearly, $c.prev.prev \mapsto c$. Together with [B], this implies that c is also in s. Therefore:

c and $c.prev.prev$ are reachable from each other in graph $\langle \mathcal{C}_\mathcal{G}, \mapsto \rangle$. –[C]

Since $candidate \subseteq CS$, using [A], we have $c.prev.prev = CS_{c.proc}$. Let x and y be any two critical sections in $g.set$, where g is the general critical section formed by merging $c.prev.prev$ with c. By the inductive hypothesis, the case when both x

and y belong to c or to $c.prev.prev$ is trivial. Therefore, without loss of generality, assume that $x \in c$ and $y \in c.prev.prev$. By the inductive hypothesis, and since INV2 follows from INV1, we use INV2 and [C] to give: x and y are reachable from each other in graph $\langle \mathcal{C}, \mapsto \rangle$. Therefore, $scc(x) = scc(y)$.

Now, consider the point of execution just before line L27 when *candidate* is added to *chain*. Let *chain.set* be the set of general critical sections in the scc's of the sequence *chain*. It is easy to show that the following is an invariant:

INV3: $\forall c \in \mathcal{C}_{\mathcal{G}} : \; c \prec CS_{c.proc} \Rightarrow c \in chain.set$

Suppose there is a general critical section c in $(\mathcal{C}_{\mathcal{G}} - (candidate \cup chain.set))$ such that for some $c' \in candidate$, $c \mapsto c'$. Consider two cases:

Case 1: $CS_{c.proc} \in candidate$**:** Since by definition, $c \notin candidate$, we have $c \neq CS_{c.proc}$. Using INV3, we have $CS_{c.proc} \prec c$ and, since c is a critical section $CS_{c.proc}.next.next \preceq c$. Together with $c \mapsto c'$, this gives us $CS_{c.proc}.next.next \mapsto c'$. This contradicts the fact that $mergeable = \emptyset$ just before line L27 (by the check at L25).

Case 2: $CS_{c.proc} \notin candidate$**:** Using INV3, we have $CS_{c.proc} \preceq c$. Since $c \mapsto c'$, we have $CS_{c.proc} \mapsto c'$. Let $s \in scc\_set$ such that $CS_{c.proc} \in s$. Therefore $s \hookrightarrow candidate$ and $s \neq candidate$. This contradicts the fact that *candidate* was selected from *valid_scc* (lines L21, L22).

Therefore, $\forall c \in (\mathcal{C}_{\mathcal{G}} - (candidate \cup chain.set)) : \; \forall c' \in candidate : \; c \not\mapsto c'$ –[D]
This proves that *candidate* is an scc in $(\mathcal{C}_{\mathcal{G}} - chain.set)$ and, furthermore, that it is a minimal scc in $(\mathcal{C}_{\mathcal{G}} - chain.set)$.

Let $reg(candidate)$ be the set of regular critical sections in *candidate* (i.e. $\{x \in \mathcal{C} \mid \exists g \in candidate : x \in g.set\}$). Similarly, let $reg(chain.set)$ be the set of regular critical sections in *chain.set*. By the definition of *ord* on $\mathcal{C}_{\mathcal{G}}$, we have: $\forall c, c' \in \mathcal{C}_{\mathcal{G}} : \; (\exists x \in c, x' \in c' : \; x \mapsto x') \Rightarrow (c \mapsto c')$. Therefore, by [D], we have:
$\forall x \in (\mathcal{C} - (reg(candidate) \cup reg(chain.set))) : \; \forall x' \in reg(candidate) : \; x \not\mapsto x'$ –[E]
By INV2, we know that all critical sections in $reg(candidate)$ are reachable from one another. Therefore, using [E], $reg(candidate)$ is an scc in $\langle \mathcal{C} - reg(chain.set), \mapsto \rangle$. Further, using [E] again, $reg(candidate)$ is also a minimal scc in $\langle \mathcal{C} - reg(chain.set), \mapsto \rangle$.

Since the minimal scc is chosen at each step, *chain* represents a topological sort of the graph $\langle \mathcal{S}, \hookrightarrow \rangle$. Therefore, from the construction of $\mathcal{O}$ from *chain* in lines L33-L38 of the algorithm, the $\rightarrow^c$ relation is the same as that defined in the proof of Theorem 8 (note that the optimization of connecting general critical sections by the $\rightsquigarrow^c$ edges instead of regular critical sections, as in the proof of Theorem 8, does not affect the equality of $\rightarrow^c$ defined by each). Therefore, in a very similar way to the proof of Theorem 8, we can prove that $\langle E, \rightarrow^c \rangle$ is a controlling computation of ϕ_{gen_mutex} in $\langle E, \rightarrow \rangle$. $\Box$