To appear in 2nd European Workshop on Planning, Dec 9-11th, 1993

Exploiting Locality in Temporal Reasoning

Shieu-Hong Lin   Thomas Dean[1]
Department of Computer Science
Brown University, Providence, RI 02912

Abstract. Temporal reasoning with uncertainty about the ordering of events is important in a wide variety of applications. Previous research shows that the associated decision problems are hard even for very restricted cases. In this paper, we investigate using the temporal locality of events and the spatial locality of state spaces to speed inference. We propose a new perspective for representing cause-and-effect relationships based on finite state automata. This perspective exposes the sources of complexity in temporal reasoning more clearly, and facilitates transforming rule-based temporal reasoning problems into graph reachability problems. We present both positive and negative results that provide insight into the complexity of temporal reasoning.

1. Introduction

Given a set of events, a set of ordering constraints on the events, and a description of the cause-and-effect relationships involving the associated event types, it is often useful to determine whether a given set of propositions can all be true after some admissible event sequence. This is the essence of the temporal reasoning problem considered in this paper. In the corresponding planning problem, the planner is free to select an arbitrary set of event instances from a fixed set of event types with no restrictions on the ordering of event instances. In special cases, temporal reasoning is harder than the corresponding planning problem [10]. This happens because in these special cases the type, number, and order of events can be controlled by the planner. In most cases, if the events are not under the planner's control, however, the problems are computationally equivalent. In

[1]This work was supported in part by a National Science Foundation Presidential Young Investigator Award IRI-8957601, by the Advanced Research Projects Agency of the Department of Defense monitored by the Air Force under Contract No. F30602-91-C-0041, and by the National Science foundation in conjunction with the Advanced Research Projects Agency of the Department of Defense under Contract No. IRI-8905436. The first author can be contacted at shl@cs.brown.edu or (401) 863-7668.

particular, temporal reasoning with partially ordered event instances and cause-and-effect relationships described by STRIPS-like propositional operator schemas is hard [4].

In AI, the cause-and-effect relationships governing the consequences of actions are typically represented in terms of a set of rules. Each rule is associated with a set of antecedent conditions (or *preconditions*) and a set of consequent conditions (or *postconditions*). Previous complexity results on temporal reasoning and planning primarily address problems characterized in terms of rule-based cause-and-effect descriptions [1] [2] [3] [4] [6]. In this paper, we explore temporal reasoning from a different perspective; the cause-and-effect relationships involving events are viewed as state-transition functions of finite automata. The conditions in rules define a state space where each condition corresponds to a binary state variable. Events are considered as input symbols. Rules determine how events trigger state transitions in the state space. From this perspective, the task in temporal reasoning is to determine whether a goal state can be reached, given a set of possible event sequences. The possible event sequences are determined by the ordering constraints on the events. Our results show that the ordering constraints on events, the size of the state space, and the structure of the state transitions all contribute to the complexity of temporal reasoning. Temporal reasoning turns out to be hard even if events are totally unordered and the associated state space is polynomial in the size of the event set. This result motivates our effort to exploit the structure inherent in event orderings and state spaces.

The use of localized reasoning to exploit locality in planning has been advocated by Lansky [8], using an event-based approach. In Lansky's work, locality is introduced by the *group* constraints and the *element* constraints [8]. A group constraint bounds the scope of causal effects and simultaneity of a subset of events. An element constraint requires a subset of events to be totally ordered. In this paper, we study temporal locality introduced by three types of ordering constraints. Event subsets are hierarchically encapsulated by the *Encapsulate* constraints and the *Preempt* constraints. The *Interleave* constraints allow the encapsulated event subsets at the same hierarchical level to be partially ordered as a constant number of chains. We present an abstraction technique to exploit temporal locality in a hierarchical way, using a state-based approach. This technique is different from previous work on abstraction [7] in that we transform encapsulated event subsets into abstract events, instead of transforming individual operators into abstract operators. We explore spatial locality where each encapsulated event subset has local conditions that only appear in the causal rules associated with the events in the event subset. We show that spatial locality can reduce the sizes of the state spaces involved in this hierarchical problem solving. More details regarding the temporal reasoning algorithm, the characterization of spatial locality, and the complexity analysis are presented in [9].

In Section 2, we formally define the temporal reasoning problem and contrast the two different perspectives. In Section 3, we demonstrate the complexity trade-offs of temporal reasoning involving the ordering constraints on events, the size of the state space, and the structure of the state transitions. In Section 4, we explore the opportunities for using temporal locality and spatial locality to solve the temporal reasoning problem. Our techniques can hierarchically reduce temporal reasoning to graph reachability. Our results do not depend on severe syntactic restrictions on the causal rules.

2. Representations and Perspectives

2.1. Temporal Projection and Temporal Reachability

The *temporal projection problem* was studied by Dean and Boddy [4]. A problem instance is defined by (1) a set of event types, $\mathcal{T} = \{type_1, type_2, \ldots, type_n\}$; (2) a set of conditions, $\mathcal{P} = \{p_1, p_2, \ldots, p_m\}$; (3) a set of causal rules, $\mathcal{R} = \{r_1, r_2, \ldots, r_o\}$, of the form $\langle t, \varphi, \alpha, \delta \rangle$ composed of: a triggering event type, $t \in \mathcal{T}$; a set of antecedent conditions, $\varphi \subseteq \mathcal{P}$; a set of added conditions, $\alpha \subseteq \mathcal{P}$; a set of deleted conditions, $\delta \subseteq \mathcal{P}$; (4) a set of actual events, $\mathcal{E}$, where $\forall e \in \mathcal{E}, type(e) \in \mathcal{T}$; (5) a set of initial conditions, $\mathcal{I} \subseteq \mathcal{P}$; and (6) a partial order $\prec$ on $\mathcal{E}$.

The task in temporal projection is to determine whether a specified condition $p \in \mathcal{P}$ is true immediately following a specified event $e \in \mathcal{E}$ in some total order consistent with the partial order $\prec$. When an instance of a given event type occurs, the status, true or false, of a condition immediately following an instance is determined as follows. A condition is true immediately following an instance if and only if either (1) the condition appears in the added conditions of a rule (associated with the event type) all of whose antecedent conditions are true immediately prior to the instance, *or* (2) the condition is true immediately prior to the instance *and* it is not the case that the condition appears in the deleted conditions of a rule (associated with the event type) all of whose antecedent conditions are true immediately prior to the instance.

Events in $\mathcal{E}$ occur as a sequence. A partial order is a set of ordering constraints on the events. For example, given $e_i \prec e_j$, e_i must appear before e_j in any possible event sequence. Ordering constraints $\prec$ on $\mathcal{E}$ determine a set of possible event sequences. In this paper, we exploit the temporal locality introduced by three types of ordering constraints: *Preempt*, *Encapsulate* and *Interleave*. We define the *Preempt* constraints and the *Encapsulate* constraints here, and we shall investigate the *Interleave* constraints in Section 4.

An event e is *preemptible* if there exists a $Preempt(X)$ constraint such that $e \in X$, $X \subseteq \mathcal{E}$, and there are no $Encapsulate(X)$ constraints such that $e \in Y$, $Y \subset X$, $Y \neq X$; otherwise, e is *non-preemptible.* In other words, each event e in X is allowed to be preempted by a constraint $Preempt(X)$ without actually occurring, unless e is encapsulated by another constraint $Encapsulate(Y)$ where $e \in Y, Y \subset X, Y \neq X$. We assume that given an arbitrary constraint $Preempt(X)$ or $Encapsulate(X)$ and an arbitrary constraint $Preempt(Y)$ or $Encapsulate(Y)$, either X and Y are disjoint or one is a proper subset of another.
For a constraint $Preempt(X)$ or a constraint $Encapsulate(X)$ where $X \subseteq \mathcal{E}$:
The events not in X can only occur either before or after the events in X. All the non-preemptible events in X must occur, while the preemptible events in X may or may not occur.
We define the following terms regarding event sequences and ordering constraints.

- An event sequence $\mathbf{q}$ (of length h) over a set of events $\mathcal{E}$ is a sequence $\langle e_1, e_2, \ldots, e_h \rangle$ where $e_i \in \mathcal{E}$, $1 \leq i \leq h \leq |\mathcal{E}|$, and $e_i \neq e_j$ if $i \neq j$. An event sequence $\mathbf{q}' = \langle e_1, e_2, \ldots, e_{h'} \rangle$, $h' \leq h$, is a prefix of $\mathbf{q} = \langle e_1, e_2, \ldots, e_{h'}, \ldots, e_h \rangle$.
- Given a set of ordering constraints $\mathcal{O}$ on $\mathcal{E}$, $\mathbf{Q}^{\mathcal{E}}_{\mathcal{O}}$ is composed of the event sequences over $\mathcal{E}$ (1) which are consistent with the ordering constraints in $\mathcal{O}$ and (2) in which all the non-preemptible events occur while the preemptible events may or may not occur.
  Note that in case no *Preempt* constraints are present in $\mathcal{O}$, $\mathbf{Q}^{\mathcal{E}}_{\mathcal{O}}$ is composed of the event sequences of length $|\mathcal{E}|$ over $\mathcal{E}$ which are consistent with the constraints in $\mathcal{O}$.
- $\overline{\mathbf{Q}}^{\mathcal{E}}_{\mathcal{O}} = \{\mathbf{q}' \mid \exists \mathbf{q} \in \mathbf{Q}^{\mathcal{E}}_{\mathcal{O}} \text{ where } \mathbf{q}' \text{ is a prefix of } \mathbf{q}\}$. $\mathbf{Q}^{\mathcal{E}}_{\mathcal{O}} \subset \overline{\mathbf{Q}}^{\mathcal{E}}_{\mathcal{O}}$.

In this paper, we focus on the following *temporal reachability problem.* The problem instance is defined by $\langle \mathcal{T}, \mathcal{P}, \mathcal{R}, \mathcal{I}, \mathcal{G}, \mathcal{E}, \mathcal{O} \rangle$ where (1) $\mathcal{T}$ is a set of event types, (2) $\mathcal{P}$ is a set of conditions, (3) $\mathcal{R}$ is a set of causal rules, (4) $\mathcal{I}$ is a set of initial conditions where $\mathcal{I} \subseteq \mathcal{P}$, (5) $\mathcal{G}$ is a set of goal conditions where $\mathcal{G} \subseteq \mathcal{P}$, (6) $\mathcal{E}$ is a set of actual events, and (7) $\mathcal{O}$ is a set of ordering constraints on $\mathcal{E}$. The set of ordering constraints $\mathcal{O}$ on $\mathcal{E}$ implicitly defines a set of event sequences $\mathbf{Q}^{\mathcal{E}}_{\mathcal{O}}$. *Given the set of initial conditions $\mathcal{P}$, the task in temporal reachability is to determine the existence of an event sequence $\mathbf{q} \in \overline{\mathbf{Q}}^{\mathcal{E}}_{\mathcal{O}}$ immediately following which the conditions in $\mathcal{G}$ are all true.* The following lemma shows that the temporal reachability problem is closely related to the temporal projection problem.

Lemma 1 *Temporal reachability regarding a partial order on events can be reduced to temporal projection in polynomial time.*

Proof. We first transform an instance of temporal reachability in the following way. A new condition p_0 and an event instance e_0 of a new event type t_0 are added. Event type t_0 is associated with a unique causal rule $r_0 = \langle t_0, \mathcal{G}, \{p_0\}, \mathcal{G}\rangle$. Event e_0 is allowed to occur after any event in $\mathcal{E}$. Temporal projection is then applied to determine whether p_0 can be true immediately following event e_0. □

By Lemma 1, the NP-Completeness results for the temporal reachability problem in section 3 (regarding partial orders on events) can also be applied to the temporal projection problem.

2.2. A Perspective based on Finite Automata

An alternative perspective on cause-and-effect relationships involves representing them in terms of state transitions in a finite state space. *The causal rules in $\mathcal{R}$ and the event types in $\mathcal{T}$ can determine a state transition function* $\Delta : 2^{\mathcal{P}} \times \mathcal{T} \rightarrow 2^{\mathcal{P}}$ in the following way. The set of conditions $\mathcal{P} = \{p_1, p_2, \ldots, p_m\}$ determines a state space $S = 2^{\mathcal{P}}$. A state is represented as a vector $u = (x_1, x_2, \ldots, x_m)$ in $\{0, 1\}^m$. If condition p_i is true at state u, $x_i = 1$; otherwise $x_i = 0$. We say that a rule $r = \langle t, \varphi, \alpha, \delta\rangle$ is *effective* at state u if all antecedent conditions in φ are true at state u. Event instances of type t trigger state transitions at a state u in $2^{\mathcal{P}}$ in the following ways: transitioning from state u to any state v if there exists an effective rule r $= \langle t, \varphi, \alpha, \delta\rangle$ at state u where (1) the conditions in α are true at state v, (2) the conditions in δ are false at state v, and (3) the conditions in $\mathcal{P} - (\alpha \bigcup \delta)$ are true at v if and only if they are true at u; transitioning from state u to state u itself if there is no effective rule at state u associated with the event type t.

Given an instance $\langle \mathcal{T}, \mathcal{P}, \mathcal{R}, \mathcal{I}, \mathcal{G}, \mathcal{E}, \mathcal{O}\rangle$, we can construct a deterministic finite automaton $M = (2^{\mathcal{P}}, \mathcal{T}, \Delta, s_0, F)$ where (1) $2^{\mathcal{P}}$ is the state space, (2) the set of event types $\mathcal{T}$ is the alphabet of M, (3) $\Delta : 2^{\mathcal{P}} \times \mathcal{T} \rightarrow 2^{\mathcal{P}}$ is the state transition function determined by the rules in $\mathcal{R}$ and the event types in $\mathcal{T}$, (4) s_0 is an initial state in $2^{\mathcal{P}}$ where the conditions in $\mathcal{I}$ are true and the conditions in $\mathcal{P} - \mathcal{I}$ are false, and (5) F denotes the set of states in $2^{\mathcal{P}}$ where the conditions in $\mathcal{G}$ are all true. An event e (of type t_e) in $\mathcal{E}$ is considered as a symbol t_e from the alphabet $\mathcal{T}$. In this way, the event sequences in $\mathbf{Q}^{\mathcal{E}}_{\mathcal{O}}$ correspond to *input strings* to the finite automaton M. *Based on this finite automata representation, the task in temporal reachability can be restated as determining whether there exists an event sequence* $\mathbf{q}$ *in* $\overline{\mathbf{Q}}^{\mathcal{E}}_{\mathcal{O}}$ *accepted by* M.

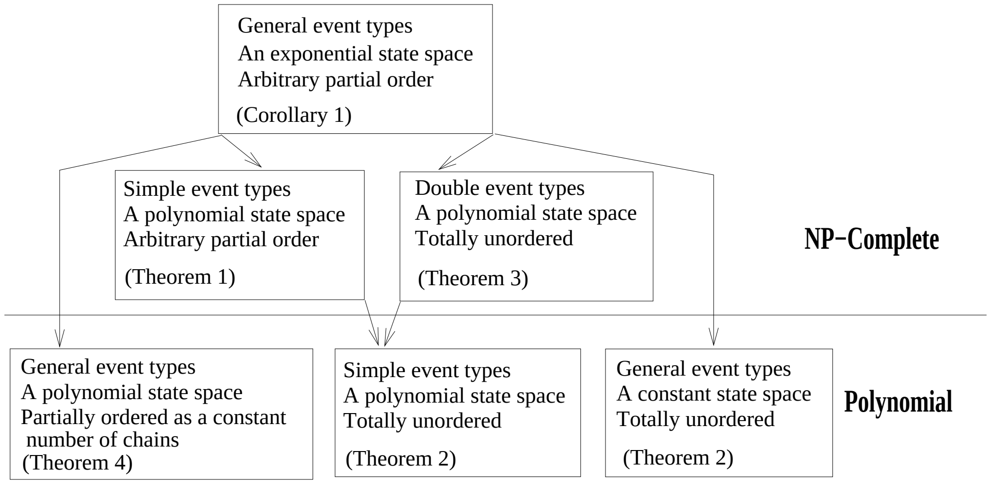

Figure 1: Complexity trade-offs for temporal reachability

3. Complexity Trade-offs

Previous research on the complexity of temporal projection has primarily focused on the structure of causal rules [4] [10]. Causal rules determine how events trigger state transitions. In this paper, we show that the ordering constraints on events and the size of the involved state space also play important roles. Figure 1 depicts the complexity trade-offs for temporal reachability involving these three factors. Ordering constraints on events are classified as one of (1) arbitrary partial order, (2) totally unordered (no ordering constraints), or (3) partially ordered as a constant number of chains. We say that events are partially ordered as a constant number of chains if (1) events are partitioned into a constant number of disjoint event subsets, and (2) in each event subset, events are totally ordered (forming a chain). The size of the state space $2^{\mathcal{P}}$ is specified as one of constant, polynomial, or exponential, in the size of the event set $\mathcal{E}$. (Namely $|\mathcal{P}| = O(1)$, $|\mathcal{P}| = O(\log|\mathcal{E}|)$, $|\mathcal{P}| = O(|\mathcal{E}|)$ respectively.) Event types are categorized as *simple*, *double* or *general* according to how they trigger state transitions. A simple event type maps exactly one state to some other state and all others to themselves, which means that exactly one rule is associated with the event type and the rule is effective at exactly one state. A double event type maps no more than two states to states other than themselves, which means that no more than two rules are associated with the event type and each rule is effective exactly at a (distinctive) state. A general event type allows arbitrary deterministic state transitions, which means that we may have an arbitrary number of effective rules for each state. Note that a simple event type is a special case of the double event types, and that a double event type is a special case of the general event types. For convenience, we shall refer to the events of simple (double) event types as "simple events" ("double events"), and the term "events" is equivalent to "events of general event types".

The temporal reachability problem is in NP. Since state transitions are deterministic, we can determine the state that is reached immediately after a given event sequence applied at an initial state. We can then verify whether the set of goal conditions are all true at the state in polynomial time. In the following, we provide proofs and algorithms to justify the claims in Figure 1. These proofs and algorithms are all based on the state-transition perspective regarding the cause-and-effect relationships. First, we define the following terms. An event sequence $\mathbf{q}$ is *coherent* from state s to state t regarding an instance of the temporal reachability problem if (1) starting from state s each event maps a state u to a state u' ($u \neq u'$), and (2) state t is reached after the last event, and (3) the event ordering in $\mathbf{q}$ is consistent with the ordering constraints on the events. The *trajectory* of a coherent event sequence $\mathbf{q}$ is the sequence of states visited when the events trigger state transitions.

Theorem 1 *Temporal reachability regarding double events, totally unordered, in a state space of polynomial size is NP-Complete.*

Proof. Given a directed graph with a set of forbidden pairs of arcs, the forbidden pairs (of arcs) problem [5] determines whether there is a path between two nodes without using both arcs in any forbidden pair of arcs. It is NP-Complete even if all the forbidden pairs are disjoint. An instance of the disjoint forbidden pairs of arcs problem can be reduced to an instance of the temporal reachability problem where there is only a final state. Given an instance of disjoint forbidden pairs, the non-isolated nodes in the directed graph are considered as states. Each directed edge (u, v) in the graph corresponds to state transition from u to v triggered by some event. We map each forbidden pair of arcs to the state transitions involving a distinct double event, and each of the remaining arcs to the state transition involving a distinct simple event. Note that the number of non-isolated nodes (states) is polynomial in the size of the number of edges (events). The initial node and the goal node are mapped to the initial state s_0 and the final state f respectively. If we can find a coherent event sequence from the initial state s_0 to the final state f, then the answer to the forbidden pairs problem is yes. Since an event corresponds to a forbidden pair of edges, it can only map one of its two possible state transitions. The trajectory of the coherent event sequence then corresponds to a feasible directed path from s_0 to f. Similarly, a directed path from s_0 to f without using both arcs in any forbidden pair corresponds to a coherent event sequence from s_0 to f $\square$

Theorem 2 *Temporal reachability regarding simple events, totally unordered, in a polynomial-size state space, or regarding general events, totally unordered, in a constant-size state space are both solvable in polynomial time.*

Proof. We derive a directed graph in the following way by (1) considering states as nodes, (2) considering a state transition (u, v) as a directed edge for every (u, v) appearing in the state transition function determined by the given event type. In the case of simple events, a simple event then corresponds to a directed edge in the graph. Temporal reachability regarding simple events in a polynomial-size state space can be directly reduced to graph reachability, since a coherent event sequence from the initial state s_0 to a final state f corresponds to a directed path from s_0 to f and vice versa.

In the case of general events, there is only a constant number of simple directed paths from the initial state s_0 to a final state f, given the constant-size state space. For each directed path, we determine whether there is an event sequence which can trigger state transitions along such a trajectory path. This problem is then reduced to the bipartite matching problem regarding the set of directed edges in a given simple directed path and the given set of events. Since the events are totally unordered, an event can match a directed edge (u, v) iff the event can trigger a state transition from state u to state v. The answer is yes if we can match distinct events to the directed edges in the simple directed path. □

Nebel and Bäckström [10] prove the NP-Completeness of a special case of the temporal projection problem by a reduction from the forbidden pairs of arcs problem on directed acyclic graphs [5]. Their proof is adapted to prove the following theorem.

Theorem 3 *Temporal reachability regarding simple events with an arbitrary partial order on events, in a polynomial-size state space is NP-Complete.*

Proof. The forbidden pairs of arcs problem on directed acyclic graphs is NP-Complete [5]. An instance of the forbidden pairs of arcs problem on a directed acyclic graph can be transformed to an instance of the temporal reachability problem where there is only a final state. The non-isolated nodes in the given directed acyclic graph are considered as states. Each arc is mapped to a unique simple event. Note that the number of non-isolated nodes (states) is polynomial in the size of the number of edges (events). For each forbidden pair of arcs (e_1, e_2) assign precedence $e_1 \prec e_2$ if there is a directed path with arc e_2 preceding arc e_1. Otherwise, assign precedence $e_2 \prec e_1$ if there is a directed path with arc e_1 preceding arc e_2. Since the graph is acyclic, the precedence relation among events is a partial order. The trajectory of a coherent event sequence from the initial state s_0 to the final state f never contains both arcs of a forbidden pair. Such a trajectory corresponds to a feasible directed path from s_0 to f. Similarly, edges in a feasible directed path from s_0 to f correspond to a coherent event sequence from s_0 to f. □

Corollary 1 *Temporal reachability regarding general events with an arbitrary par-*

tial order on events, in an exponential-size state space is NP-Complete.

Proof. Note that simple events and double events are the special cases of general events, and we can add in redundant conditions to augment a polynomial-size state space into an exponential one. Therefore the reductions in Theorem 1 and Theorem 3 still hold in this case, and this implies the NP-Completeness of the temporal reachability problem in its general form. □

The temporal reachability problem regarding events partially ordered as chains can be solved using the following procedure. Temporal reachability is reduced to graph reachability in this case. This procedure is also used in Section 4 for hierarchically solving the temporal reachability problem.

Procedure MultiChain($\mathcal{E}, \mathcal{O}, I$)

input: (1) $\mathcal{E}$, a set of events; (2) $\mathcal{O}$, a set of ordering constraints on $\mathcal{E}$ where events are partially ordered as chains; (3) I, a set of initial states.

output: (1) the set of states reached after the event sequences in $\overline{\mathbf{Q}}^{\mathcal{E}}_{\mathcal{O}}$; (2) the set of states reached after the event sequences in $\mathbf{Q}^{\mathcal{E}}_{\mathcal{O}}$; (3) the sets of states that we may reach immediately before a specific event e occurs, for all e in $\mathcal{E}$.

- Suppose events are partially ordered as c chains with l_i events in the ith chain. Construct the following directed acyclic graph $G = (V, E)$.
  - Each state s in S is expanded into $\prod_{1 \le i \le c}(1 + l_i)$ nodes of the form $(s, x_1, x_2, \ldots, x_c)$ where $0 \le x_i \le l_i$. These nodes comprise the vertex set V of G.
    Reaching a node $(s, x_1, x_2, \ldots, x_c)$ has the semantics that the state s may be reached immediately after the first x_i events in the ith chain have occurred for every i.
  - Construct a directed edge from $(s, x_1, x_2, \ldots, x_k - 1, \ldots, x_c)$ to $(t, x_1, x_2, \ldots, x_k, \ldots, x_c)$ if the x_kth event in the kth chain can trigger a state transition from s to t. These edges comprise the edge set V of G.
- We apply the standard graph reachability algorithms to derive the sets of states we want.
  - State t is reachable from state s ($s \in I$) after an event sequence in $\overline{\mathbf{Q}}^{\mathcal{E}}_{\mathcal{O}}$ if and only if there exists $(t, x_1, x_2, \ldots, x_c)$ reachable from $(s, 0, 0, \ldots, 0)$ in G.
  - State t is reachable from state s ($s \in I$) after an event sequence in $\mathbf{Q}^{\mathcal{E}}_{\mathcal{O}}$ if and only if $(t, l_1, l_2, \ldots, l_c)$ is reachable from $(s, 0, 0, \ldots, 0)$ in G.
  - Suppose e is the rth event in the kth chain. Immediately before event e occurs, we can reach state t from state s ($s \in I$) if and only if there exists $(t, x_1, x_2, \ldots, x_k = r - 1, \ldots, x_c)$ reachable from $(s, 0, 0, \ldots, 0)$ in G.

Theorem 4 *Temporal reachability regarding events partially ordered as a constant number of chains in a polynomial-size state space can be solved in polynomial time.*

Proof(Sketch). We apply **MultiChain**($\mathcal{E}, \mathcal{O}, \{s_0\}$) to determine the set of states reachable after the event sequences in $\overline{\mathbf{Q}}^{\mathcal{E}}_{\mathcal{O}}$. The number of nodes in the graph constructed by procedure **MultiChain** is equal to $|S| \prod_{1 \le i \le c}(1 + l_i)$ where $|S|$ is the size of the state space and l_i is the number of events in the ith chain. Since

$l_i = O(|\mathcal{E}|)$ and c is a constant, the graph G is polynomial in the size of $\mathcal{E}$. Graph reachability and thus temporal reachability is solvable in polynomial time. □

4. Exploiting Locality

4.1. Temporal Locality in Event Ordering

As shown in Section 3, temporal reachability regarding general events is hard if the events are totally unordered and the state space is polynomial in the size of the event set $\mathcal{E}$. When totally unordered, events can occur in arbitrary order. However, there is temporal locality in event ordering if the occurrences of events or subsets of events closely relate to one another. For example, a subset of events may always occur as an atomic unit, such that all the other events either occur before or after the subset of events. Similarly, several atomic units may always occur as an atomic unit. In this way, temporal locality can be introduced at many levels. Suppose we have a set of events $\mathcal{E}$. The *Encapsulate* and *Preempt* types of ordering constraints defined in Section 2 can introduce temporal locality in an event subset $X = \{e_1, e_2, \ldots, e_k\}$ in $\mathcal{E}$. We say that X is a *local event subset* in $\mathcal{E}$ if X is constrained by the constraints $Encapsulate(X)$ or $Preempt(X)$. The event set $\mathcal{E}$ and each event in $\mathcal{E}$ are local event subsets in $\mathcal{E}$, since they are implicitly constrained by the encapsulation constraints. For any two local event subsets X and Y in $\mathcal{E}$, we require either that X and Y are disjoint or that one is a proper subset of the other. Suppose X is a subset of Y. We say that X is *maximal* in Y if X is not a proper local event subset of any other local event subset in Y. In this way, the local event subsets in $\mathcal{E}$ form a hierarchy.

Given a local event subset Y, the following *Interleave* type of ordering constraint introduces a partial order on the maximal local event subsets in Y. These event subsets are partially ordered as chains.
$Interleave(\langle X_1^1, X_1^2, \ldots\rangle, \langle X_2^1, X_2^2, \ldots\rangle, \ldots, \langle X_c^1, X_c^2, \ldots\rangle)$: X_i^j's are the maximal local event subsets in Y. $\forall i$, $X_i^j \prec X_i^k$ if $j < k$. $\langle X_i^1, X_i^2, \ldots\rangle$ forms a chain where the events in X_i^j must occur before the events in X_i^k if $j < k$.

Figure 2 shows the following ordering constraints on ten events where $X = \{e_2, e_3, e_4, e_5\}$ and $Y = \{e_6, e_7, e_8, e_9\}$: $Preempt(X)$, $Interleave(\langle\{e_2\}, \{e_3\}\rangle, \langle\{e_4\}, \{e_5\}\rangle)$, $Encapsulate(Y)$, $Interleave(\langle\{e_6\}, \{e_7\}\rangle, \langle\{e_8\}, \{e_9\}\rangle)$, $Encapsulate(\{e_1\} \bigcup X \bigcup Y \bigcup \{e_{10}\})$, $Interleave(\langle\{e_1\}, X\rangle, \langle Y, \{e_{10}\}\rangle)$. $X, Y, \{e_1\}, \{e_{10}\}$ are the maximal local event subsets in the event set composed of these ten events. The events in X always occur together, but they can be preempted. In X, e_3 never occurs before e_2 and

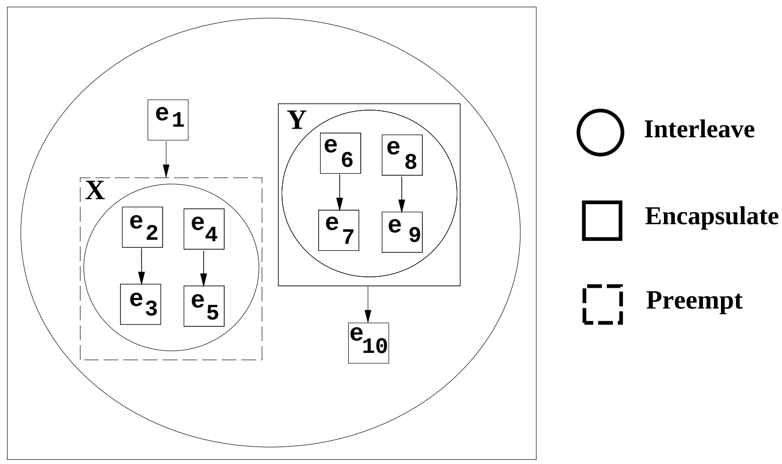

Figure 2: Temporal locality in local event subsets

e_5 never occurs before e_4. The events in Y always occur together. In Y, e_6 always occurs before e_7. e_8 always occurs before e_9. e_1 always occurs before the events in X. e_{10} always occurs after the events in Y.

4.2. Hierarchical Problem Solving

Given a local event subset X, a set of ordering constraints $\mathcal{O}$ on X, and a subset of conditions $\mathcal{V}$ ($\mathcal{V} \subseteq \mathcal{P}$), X can be abstracted as an *abstract event* $e_X^{\mathcal{V}}$. Suppose that we can reach a state v in $2^{\mathcal{V}}$ from a state u in $2^{\mathcal{V}}$ after an event sequence $\mathbf{q} \in \mathbf{Q}_{\mathcal{O}}^X$. Then the abstract event $e_X^{\mathcal{V}}$ may trigger a state transition from u to v. This describes a *state-transition relation* since we may reach another state v' from state u after another event sequence $\mathbf{q}' \in \mathbf{Q}_{\mathcal{O}}^X$. In other words, $e_X^{\mathcal{V}}$ summarizes the possible state transitions in $2^{\mathcal{V}}$ caused by the events in X as a whole. In this subsection, we assume $\mathcal{V} = \mathcal{P}$ and abbreviate $e_X^{\mathcal{V}}$ as e_X.

When the ordering of events is governed by the *Encapsulate*, *Preempt*, and *Interleave* ordering constraints, the temporal reachability problem can be solved hierarchically. In hierarchical problem solving, local event subsets correspond to abstract events. For example, temporal reachability regarding events in Figure 2 can be solved in two levels. At the top level, we determine the possibility of reaching any goal state, given (abstract) events e_1, e_{10}, e_X, e_Y. If the answer is negative, it shows that the goal states can neither be reached before nor after the events in X (Y) occur. In this case, we first determine the possible states that we may reach immediately before the events in X (Y) occur. Then at the second level, we determine the possibility of reaching any goal state when the events in X (Y) occur, given the possible states that we may reach immediately before events in X (Y).

In the following, the procedure **MultiChain** in Section 3 is used as a subroutine in hierarchical problem solving. The following procedure recursively derives the abstract events for the local event subsets in a set of events.

Procedure Event-Abstraction$(\mathcal{E}, \mathcal{O})$

input: (1) $\mathcal{E}$, a set of events; (2) $\mathcal{O}$, a set of ordering constraints on $\mathcal{E}$.

output: the abstract events for the local event subsets in $\mathcal{E}$.

- For each maximal local event subset X in $\mathcal{E}$, we call **Event-Abstraction**$(X, \mathcal{O}_X)$ to derive the abstract events for the local event subsets in X where $\mathcal{O}_X$ is composed of the ordering constraints on X, $\mathcal{O}_X \subset \mathcal{O}$.
- Let $\mathcal{A}$ be composed of the abstract events for the maximal local event subsets in $\mathcal{E}$. Let $\prec$ be the partial order on $\mathcal{A}$ introduced by the *Interleave* ordering constraint on the maximal local event subsets in $\mathcal{E}$. We derive the abstract event $e_{\mathcal{E}}$ for $\mathcal{E}$ as follows.
  - For each state s, the abstract event $e_{\mathcal{E}}$ can trigger a state transition from state s to (1) any state reachable after the event sequences in $\mathbf{Q}_{\prec}^{\mathcal{A}}$ if $\mathcal{E}$ is constrained by an *Encapsulate* ordering constraint, (2) any state reachable after the event sequences in $\overline{\mathbf{Q}}_{\prec}^{\mathcal{A}}$ if $\mathcal{E}$ is constrained by a *Preempt* ordering constraint. We call **MultiChain**$(\mathcal{A}, \prec, \{s\})$ to determine this state-transition relation.

The following procedure recursively decomposes the temporal reachability problem and then solves it level by level.

Procedure Problem-Decomposition$(\mathcal{E}, \mathcal{O}, I, G)$

input: (1) $\mathcal{E}$, a set of events; (2) $\mathcal{O}$, a set of ordering constraints on $\mathcal{E}$;
(3) I, a set of initial states; (4) G, a set of goal states.

output: yes if there exist s in I, t in G, $\mathbf{q}$ in $\overline{\mathbf{Q}}_{\mathcal{O}}^{\mathcal{E}}$ such that t can be reached from s after the event sequence $\mathbf{q}$; no, otherwise.

- We call **Event-Abstraction**$(\mathcal{E}, \mathcal{O})$ to derive the abstract events for all the local event subsets in $\mathcal{E}$.
- Let $\mathcal{A}$ be composed of the abstract events for the maximal local event subsets in $\mathcal{E}$. Let $\prec$ be the partial order on $\mathcal{A}$ introduced by the *Interleave* ordering constraint on the maximal local event subsets in $\mathcal{E}$. At the top level, we call **MultiChain**$(\mathcal{A}, \prec, I)$ to determine the set of states reachable after the event sequences in $\overline{\mathbf{Q}}_{\prec}^{\mathcal{A}}$.
- If we can reach states in G at this level, the answer to the problem is yes. Otherwise, we call **MultiChain**$(\mathcal{A}, \prec, I)$ to determine I_X, the set of states that we may reach immediately before e_X occurs ($\forall e_X \in \mathcal{A}$). I_X corresponds to the states that we may reach immediately before any event in X occurs.
- At the next level, we call **Problem-Decomposition**$(X, \mathcal{O}_X, I_X, G)$ for each maximal local event subset X in $\mathcal{E}$ where $\mathcal{O}_X$ is composed of the ordering constraints on X, $\mathcal{O}_X \subset \mathcal{O}$.

4.3. State Space Reduction in Temporal Reachability

In subsection 4.2, the procedure **Problem-Decomposition** deals with the global state space $2^{\mathcal{P}}$. However, spatial locality in local event subsets may enable us to

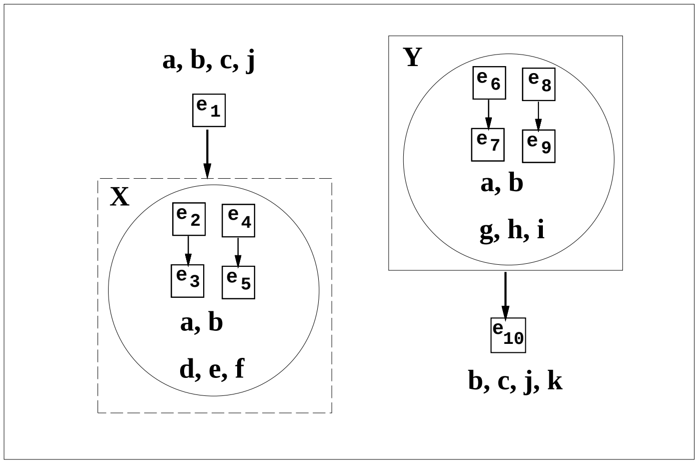

Figure 3: Spatial locality in local event subsets

deal with local state spaces in the form of $2^{\mathcal{V}}$ where $\mathcal{V}$ is a proper subset of $\mathcal{P}$, rather than the global state space $2^{\mathcal{P}}$. We say that event e involves condition p if p appears in a causal rule associated with the event type of e. In Figure 3, eleven conditions are involved. Event e_1 involves the conditions a, b, c, j. The events in $X = \{e_2, e_3, e_4, e_5\}$ involve the conditions a, b, d, e, f. The events in $Y = \{e_6, e_7, e_8, e_9\}$ involve the conditions a, b, g, h, i. Event e_{10} involves the conditions b, c, j, k. We want to determine whether we can reach a state where the subset of conditions $G = \{a, b, c\}$ are all true. Let $\mathcal{U} = \{a, b, d, e, f\}$, $\mathcal{V} = \{a, b, g, h, i\}$, $\mathcal{W} = \{a, b, c, j, k\}$, $\mathcal{Z} = \{a, b\}$, and $\mathcal{P} = \{a, b, c, d, e, f, g, h, i, j, k\}$. It turns out that we only need to deal with local state spaces $2^{\mathcal{U}}$, $2^{\mathcal{V}}$, $2^{\mathcal{W}}$ rather than the global state space $2^{\mathcal{P}}$.

At the second level, a local state space $2^{\mathcal{U}}$ ($2^{\mathcal{V}}$) where $\mathcal{V} = \{a, b, d, e, f\}$ ($\mathcal{U} = \{a, b, g, h, i\}$) is considered for deriving the abstract event $e_X^{\mathcal{Z}}$ ($e_Y^{\mathcal{Z}}$) where $\mathcal{Z} = \{a, b\}$, since only conditions in $\mathcal{U}$ ($\mathcal{V}$) are involved in X (Y). At the top level, we deal with the events $e_1, e_{10}, e_X^{\mathcal{Z}}, e_Y^{\mathcal{Z}}$ in the local state space $2^{\mathcal{W}}$. This works, because (1) the conditions d, e, f (g, h, i) are only locally involved in X (Y), and (2) the effects of conditions d, e, f (g, h, i) only propagate to e_1, e_{10} through the possible status changes of the conditions $\mathcal{Z} = \{a, b\}$ caused by the events in X and Y.

At the top level, if we can reach a state in $2^{\mathcal{W}}$ where the conditions a, b, c are all true, the answer to the problem is yes. Otherwise, we must examine the possibility that (1) c is made true at the top level by $e_1, e_{10}, e_X^{\mathcal{Z}}, e_Y^{\mathcal{Z}}$, and (2) a, b are made true at the second level by the events in the event subsets X or Y. This is because the status of condition c cannot be changed by the events in X (Y). We determine the possible statuses of the conditions in $\mathcal{U} = \{a, b, d, e, f\}$ ($\mathcal{V} = \{a, b, g, h, i\}$) immediately before the events in X (Y) occur in the following way. First, we derive the set of states in $2^{\mathcal{W}}$ possibly reached at the top level immediately before

$e_X^{\mathcal{Z}}$ ($e_Y^{\mathcal{Z}}$). Second, we find out the possible statuses of the conditions in $\mathcal{Z} = \{a, b\}$ at these states, given that the condition c must be true at these states too. The statuses of the conditions d, e, f (g, h, i) are the same as their initial statuses, since they are only locally involved in the event subset X (Y). At the second level, we determine whether a, b, c may all be true when the events in X (Y) occur, given the possible initial statuses of the conditions in $\mathcal{U} = \{a, b, d, e, f\}$ ($\mathcal{V} = \{a, b, g, h, i\}$).

5. Conclusion

In this paper, we describe an alternative perspective on temporal reasoning, representing the cause-and-effect relationships as finite state automata. We show that the ordering constraints on events, the size of the state space, and the state-transition rules all contribute to the problem complexity. Our results indicate that we have very limited ability to deal with general events if there is no temporal or spatial locality at all. We characterize the temporal locality induced by interleaving, encapsulating, and preempting subsets of events, and provide an algorithm that takes advantage of the locality in events in performing temporal reasoning.

References

[1] C. Bäckström and I. Klein. Parallel non-binary planning in polynomial time. In *Proceedings IJCAI 12*, pages 268–273. IJCAII, 1991.

[2] Tom Bylander. Complexity results for planning. In *Proceedings IJCAI 12*, pages 274–279. IJCAII, 1991.

[3] David Chapman. Planning for conjunctive goals. *Artificial Intelligence*, 32:333–377, 1987.

[4] Thomas Dean and Mark Boddy. Reasoning about partially ordered events. *Artificial Intelligence*, 36(3):375–399, 1988.

[5] Michael R. Garey and David S. Johnson. *Computers and Intractibility: A Guide to the Theory of NP-Completeness.* W. H. Freeman and Company, New York, 1979.

[6] Naresh Gupta and Dana S. Nau. Complexity results for blocks-world planning. In *Proceedings AAAI-91*, pages 629–633. AAAI, 1991.

[7] Craig A. Knoblock. Learning abstraction hierarchies in problem solving. In *Proceedings AAAI-90*, pages 923–928. AAAI, 1990.

[8] Amy L. Lansky. Localized event-based reasoning for multiagent domains. *Computational Intelligence*, 4(4), 1988.

[9] Shieu-Hong Lin and Thomas Dean. Rational reconstruction of localized reasoning. In preparation.

[10] Bernhard Nebel and Christer Bäckström. On the computational complexity of temporal projection, planning, and plan validation. *Artificial Intelligence*, 1993. To appear.