Efficient Evaluation of Causality Relations between Nonatomic Events

Ajay D. Kshemkalyani
ECECS Department, P. O. Box 210030, Univ. of Cincinnati, Cincinnati, OH 45221

High-level actions in distributed computations are modeled by nonatomic events which are collections of events at a finer level of granularity of observing the computation. Each such nonatomic event is a partially ordered set of "atomic" events at the finer level of granularity, ordered by the causality relation $\prec$ between those events. The study of causality between nonatomic events is important in analyzing distributed computations [6, 7, 8].

A suite of causality relations between nonatomic events (also termed *intervals*) was proposed in [1, 2] to extend the two relations, "precedes" and "can affect", proposed by Lamport [7]. The suite of relations was used to derive all the possible orthogonal relations that can hold between a pair of linear nonatomic events [1, 2]. (In a suite of orthogonal relations, one and only one relation can exist between a given pair of events, and no relation can be expressed as the disjunction of others.) It was shown that using the dense (non-dense) model of time, there are 29 (40) possible orthogonal relations between any pair of linear intervals [1, 2]. When the relations of [2] are applied to a pair of poset intervals, the hierarchy they form is not complete.

In [1, 3, 4], we extended the results of [2] and formulated an "exhaustive" set of causality relations between nonatomic poset events using first-order predicate logic and only the $\prec$ relation between atomic events. These relations were formulated as follows. The beginning and end subsets of each poset event were identified and played the role of proxies for the poset, similar to the common usage of the endpoints of a linear interval to denote the interval. The proxies could be the antichains of the least and greatest members of the poset, or simply the sets of the earliest and latest events at each process/node. Causality relations were specified between the proxies of X and Y, denoted $\hat{X}$ and $\hat{Y}$, respectively. To specify relations between posets X and Y, there is a choice of two proxies of X and two proxies of Y – hence, four combinations. For each combination, there are 8 ways in which members of $\hat{X}$ and members of $\hat{Y}$ are related by $\prec$; each of $\hat{X}$ and $\hat{Y}$ can be quantified by either the existential or universal quantifier, and the order of quantifications of $\hat{X}$ and $\hat{Y}$ can be swapped. This yields 32 relations between X and Y, of which 24 are distinct. These 24 relations can be ordered by the relation "is a subset of" to form a highly symmetric lattice which can be expressed as a product of two lattices. An axiom system to reason with these relations was given in [1, 5]. This axiom system can be used to derive all the orthogonal relations between the poset events.

The above relations provide applications a fine level of discrimination in specifying causality/synchronization conditions, and for subsequent temporal and spatial reasoning. Complex conditions can be expressed as predicates over these relations. Given a set of nonatomic events of interest to the application, we need to determine efficiently whether a specific relation (or each of the relations) holds between each pair of nonatomic events in the presented set. In the general case, the use of proxies reduced the complexity of evaluating causality between posets X and Y from $|X| \times |Y|$ to $|N_X| \times |N_Y|$, where N_A is defined to be the number of system nodes at which the atomic events in nonatomic event A occur. We derive linear-time evaluation conditions for the proposed relations using properties of partial orders – most relations can be evaluated in $min(|N_X|, |N_Y|)$ integer comparisons, a few in $|N_X|$ or $|N_Y|$ integer comparisons.

The linear-time evaluation conditions are derived as follows. We define various execution prefixes (a.k.a. cuts) associated with a nonatomic poset event, and their timestamps. These cuts capture useful causality information about the past and future of the event in a compact form. We then define a relation $\ll$ between cuts associated with the past of Y and the future of X; intuitively, $C(X) \ll C'(Y)$ says that the projection of cut $C(X)$ at each node is a proper subset of the projection of cut $C'(Y)$ at that node. We show that this relation can be evaluated in $min(|N_X|, |N_Y|)$ integer comparisons. Lastly, we show that each of the fine-grained causality relations between X and Y holds iff the above relation holds between appropriately identified cuts associated with the future of $\hat{X}$ and the past of $\hat{Y}$. We note that by using only cuts associated with the past of $\hat{X}$ and the past of $\hat{Y}$ to express the fine-grained causality relations, the overhead of evaluating the relations is often higher.

The full version of this paper is included in [1].

References

[1] A. Kshemkalyani, Temporal interactions of intervals in distributed systems, *Tech. Rep. 29.1933, IBM*, Sept. 1994.

[2] A. Kshemkalyani, Temporal interactions of intervals in distributed systems, *Journal of Computer and System Sciences*, 52(2), 287-298, April 1996 (a subset of [1]).

[3] A. Kshemkalyani, Relative timing constraints between complex events, *8th IASTED Conf. on Parallel and Distributed Computing and Systems*, 324-326, Oct. 1996.

[4] A. Kshemkalyani, Synchronization for distributed real-time applications, *5th Workshop on Parallel and Distributed Real-time Systems*, IEEE CS Press, 81-90, April 1997.

[5] A. Kshemkalyani, Causality between nonatomic poset events in distributed computations, *5th IEEE Workshop on Future Trends in Distributed Computing Systems*, 276-282, 1997.

[6] A. Kshemkalyani, A framework for viewing atomic events in distributed computations, *Theoretical Computer Science*, 196(1-2), 45-70, April 1998.

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

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

PODC 98 Puerto Vallarta Mexico