Detection of Strong Unstable Predicates in Distributed Programs [1],

Vijay K. Garg
Electrical and Computer Engineering Dept
University of Texas at Austin
Austin, TX 78712-1084
garg@ece.utexas.edu

Brian Waldecker
Austin Systems Center
Schlumberger Well Services
P.O. Box 200015
Austin, TX 78720-0015

Abstract

This paper discusses detection of global predicates in a distributed program. A run of a distributed program results in a set of sequential traces, one for each process. These traces may be combined to form many global sequences consistent with the single run of the program. A strong global predicate is true in a run if it is true for all global sequences consistent with the run. We present algorithms which detect if the given strong global predicate became true in a run of a distributed program.

1 Introduction

Detection of global predicates is a fundamental problem in distributed computing. It arises in the designing, debugging and testing of distributed programs. Global predicates can be classified into two types - stable and unstable. A stable predicate is one which never turns false once it becomes true. An unstable predicate is one without such a property. Its value may alternate between true and false. Detection of stable predicates has been addressed in the literature by means of global snapshots of a distributed computation [ChaLam85, SpeKea86, Bouge87]. Any stable property can be detected by taking global snapshots periodically. This approach does not work for an unstable predicate which may turn true only between two snapshots and not at the time when the snapshot is taken. An entirely different approach is required for such predicates [WalGar91, CooMar91, GarWal92, SchMat92, GarWal94, BabMar93, TomGar93, HPR93].

We have earlier presented an approach to detect a class of unstable predicates called weak predicates [GarWal94]. In this paper, we continue our investigation of detection for a different class of unstable predicates. The reader is referred to [GarWal94] for a discussion of related work and the background. Two types of predicates are discussed in this paper. The first type, called strong linked predicates, refers to a causal sequence of local predicates. The second type, called strong conjunctive predicates, correspond to existence of a global state in which all local predicates are true simultaneously. We introduce the notion of overlapping intervals which is used to detect predicates of this type. Cooper and Marzullo [CooMar91] also describe strong predicate detection (they call such predicates *definitely*). However, they deal with general predicates, i.e., they propose detection of *definitely* : p where p is any predicate defined on a global state. In this paper, we have restricted p to conjunction of local predicates. Detection of general predicates is intractable since it involves a combinatorial

[1]This work was supported in part by the NSF Grant CCR 9110605, the Navy Grant N00039-91-C-0082, a TRW faculty assistantship award, and IBM Agreement 153.

explosion of the state space. For example, the algorithm proposed by Cooper and Marzullo [CooMar91] has complexity $O(k^n)$ where k is the maximum number of events a monitored process has executed and n is the number of processes. The fundamental difference between our algorithm and their algorithm is that their algorithm explicitly checks all possible global states, whereas our algorithm does not.

Spezialetti and Kearns [SpeKea89] also discuss a notion of simultaneity which, however, is different from the one discussed in this paper. They use simultaneity in the sense of a possible consistent global state. Their notion is closer to weak predicates discussed in [GarWal94].

This paper is organized as follows: Section 2 presents our logic for specification of global predicates in a distributed program. It describes the notion of a distributed run, a global clock, a global sequence and the logic. Section 3 discusses detection of linked predicates. Section 4 discusses strong conjunctive predicates. It gives necessary and sufficient conditions for strong conjunctive predicates to hold. It also describes algorithms for detecting strong conjunctive predicates. Section 5 presents techniques to decentralize these algorithms.

2 Our Model

We assume a loosely-coupled message-passing system without any shared memory or a global clock. A distributed program consists of a set of n processes denoted by $\{P_1, P_2, ..., P_n\}$ communicating solely via asynchronous messages. We do not make any assumptions on the ordering or reliability of messages.

A. Run

We will be concerned with a single run r of a distributed program. Each process P_i in that run generates an execution trace $s_{i,0}s_{i,1}\ldots s_{i,l}$, which is a finite sequence of local *states* in the process P_i. A *run* r is a vector of traces with $r[i]$ as the trace of the process P_i.

We define a relation *locally precedes* denoted by $\prec_{im}$ between states in the trace of a single process P_i as follows: $s \prec_{im} t$ if and only if s *immediately* precedes t in the trace $r[i]$. We also say that $s.next = t$ or $t.prev = s$ whenever $s \prec_{im} t$. We use $\prec$ for irreflexive transitive closure and $\preceq$ for for reflexive transitive closure of $\prec_{im}$. States s and t in the traces $r[i]$ and $r[j]$ respectively are defined to be related by $\leadsto$ if and only if a message is sent by P_i resulting in the state s which is received by P_j resulting in the state t. Figure 1 illustrates a run.

We also define a *causally precedes* relation as the transitive closure of union of $\prec_{im}$ and $\leadsto$. That is,
$s \rightarrow t$ iff
1. $(s \prec_{im} t) \vee (s \leadsto t)$, or
2. $\exists u : (s \rightarrow u) \wedge (u \rightarrow t)$

Our $\rightarrow$ is similar to Lamport's happened-before relation [Lamp78] except that *causally precedes* is defined between states rather than events. We say that s and t are concurrent (denoted by $s||t$) if $\neg(s \rightarrow t) \ \wedge \ \neg(t \rightarrow s)$.

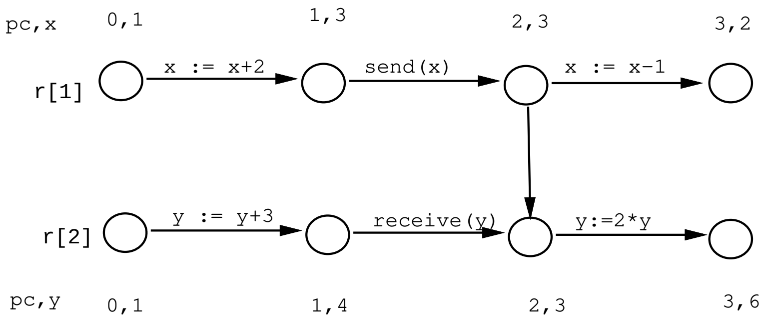

Figure 1: An example of a run

We extend the run r to r' by adding artificial states $\bot_i$ and $\top_i$ at the beginning and the end of each trace $r[i]$ respectively. The event at $\bot_i(\top_i)$ corresponds to the beginning (termination) of the execution of P_i. The addition of these artificial states model the fact that processes begin their execution asynchronously. Thus, in absence of any synchronization (external events) it may be possible that one process may have terminated while the other one has yet to begin the execution. Let S_i be the set of all states in the sequence $r[i]$, and $S_i' = S_i \cup \{\bot_i, \top_i\}$. Our definitions imply that $\forall s \in S_i : \bot_i \rightarrow s \ \wedge \ s \rightarrow \top_i$. We define $S = \bigcup_i S_i$ and $S' = \bigcup_i S_i'$. We also use $s.p$ to denote the process in whose trace s occurs. That is, $s.p = i$ if and only if $s \in S_i'$.

B. Global Clock

A *global clock* C is a map from S' to $\mathcal{N}$ (the set of natural numbers) with the following constraint:

$$\forall s, t \in S : s \prec_{im} t \ \vee \ s \rightsquigarrow t \Rightarrow C(s) < C(t)$$

We use $\mathcal{C}$ to denote the set of all global clocks which satisfy the above constraint. The interpretation of $C(s)$ for any $s \in S$ is that the process $s.p$ enters the state s when the clock value is $C(s)$. Thus, it stays in the state s from time $C(s)$ to $C(s.next) - 1$. This constraint models the sequential nature of execution at each process and the physical requirement that any message transmission requires a non-zero amount of time. From the definition of $\rightarrow$, it is equivalent to

$$\forall s, t \in S : s \rightarrow t \Rightarrow \forall C \in \mathcal{C} : C(s) < C(t) \qquad (\mathbf{CC})$$

The condition (CC) is widely used as the definition of a logical clock since its proposal by Lamport[Lamp78]. It can be shown that the set $\mathcal{C}$ also satisfies the converse of (CC), i.e.,

$$\forall s, t \in S : s \not\rightarrow t \Rightarrow \exists C \in \mathcal{C} : \neg(C(s) < C(t))$$

The reader is referred to [GarWal94b] for the proof. This leads to the following pleasant characterization of $\rightarrow$:

$$\forall s, t \in S : (s \rightarrow t \Leftrightarrow \forall C \in \mathcal{C} : C(s) < C(t))$$

Intuitively, the above formula says that s causally precedes t in a run r if and only if all possible observers of the run agree that s happened before t.

C. Global Sequence

A global state is a vector of local states. This definition of global state is different from that of Chandy and Lamport which includes states of channels. In our model, a channel is just the set of all those messages that have been sent but not received yet. Since this set can be deduced from all the local states, we do not require the state of channels to be explicitly included in the global state. Given a run r, and a global clock C, $seq(r, C)$ defines a sequence of global states called global sequence $g = g_0 g_1 \dots g_m$ for some m where

$$g_k[i] = max\{\{s \in S_i' | C(s) \leq k\} \cup \{\perp_i\}\}, \quad 1 \leq k \leq m, 1 \leq i \leq n$$

Note that $g_k[i]$ is well defined as the argument of max is a non-empty totally ordered (under $\prec$) finite set. It may evaluate to $\perp_i$ which would mean that the process P_i has not begun its execution. Similarly, if it evaluates to $\top_i$, then process P_i has already terminated its execution. The kth prefix of g, i.e. $g_0 g_1 \dots g_{k-1}$ is denoted by g^k.

We define $findex(g, u) = min\{k | g_k[u.p] = u\}$, i.e., the first index in g which has u in its global state. We define the set of global sequences consistent with a run r as $linear(r)$, i.e.

$$g \in linear(r) \Leftrightarrow \exists C \in \mathcal{C} : g = seq(r, C)$$

The following theorem gives an alternative characterization of the set $linear(r)$. Given any $g \in linear(r)$ if the observer restricts his attention to a single process P_i, then he would observe $r[i]$ or a stutter of $r[i]$. A *stutter* of $r[i]$ is a finite sequence where each state in $r[i]$ may be repeated a finite number of times.

Lemma 1 *For any run r, $g \in linear(r)$ if and only if the following constraints hold:*
(S1): $\forall i : g$ *restricted to* $P_i = r[i]$ *(or a stutter of* $r[i]$*)*
(S2): $\forall s, t \in S : s \to t \Rightarrow findex(g, s) < findex(g, t)$.

Proof: ($\Rightarrow$)
Let $G \in linear(r)$. This implies that $\exists C : g = seq(r, C)$. Since $C(s)$ is greater than 0 for all $s \neq \perp_i$, we get that $g_0[i] = \perp_i$. Further, $g_k[i] = s \Rightarrow g_{k+1}[i] \in \{s, s.next\}$ by definition of $seq(r, C)$. Finally, $g_{C[\top_i]}[i] = \top_i$. Thus (S1) holds. To see (S2), let $s \to t$. From (CC), we get that $C(s) < C(t)$. This implies that $findex(g, s) < findex(g, t)$.
($\Leftarrow$)
We define C as follows:

$$C(s) = findex(g, s)$$

C satisfies (CC) due to (S2).

■

In our earlier paper[GarWal94], we have directly defined the notion of global sequences. In this paper, we have chosen the condition (CC) based on global clocks as it is intuitively easier to justify.

From the above two properties of global sequences, we can also deduce **(S3)**:

$$\forall X \subseteq S : (\forall u, v \in X : u||v) \Rightarrow \exists g \in linear(r)\ \exists k \forall u \in X : (g_k[u.p] = u)\ \wedge\ (g_{k-1}[u.p] \neq u)$$

(S3) says that for any set X of concurrent states there exists a global sequence g which goes through a global state g_k such that all local states in X occur in g_k, and none occur in g_{k-1}.

D. Logic for Global Predicates

We now describe our logic for specification of global predicates. There are three syntactic categories in our logic - *bool*, *lin* and *form*. The syntax of our logic is as follows:

form	::=	$\underline{\mathrm{A}}$: lin | $\underline{\mathrm{E}}$: lin
lin	::=	$\diamond$ lin | lin $\hookrightarrow$ lin | lin $\wedge$ lin | $\neg$lin | bool
bool	::=	a predicate over a global state

A *bool* is a boolean expression defined on a single global state of the system. Its value can be determined if the global state is known. For example, if the global state has $(x = 3, y = 6)$, then the *bool* $(x \leq y)$ is true. Here x and y could be part of different processes. A *lin* is a temporal formula defined over a global sequence. A *bool* is true in a global sequence if it is true in the last state of g. $\diamond$ *lin* means that there exists a prefix of the global sequence such that *lin* is true for the prefix. We also use $\Box$ and $\vee$ as duals of $\diamond$ and $\wedge$. We have introduced a binary operator ($\hookrightarrow$) to capture sequencing directly. $p \hookrightarrow q$ means that there exist prefixes g^i and g^j of the global sequence such that p is true of prefix g^i, q is true of prefix g^j, and $i < j$.

A *form* is defined for a run and it is simply a *lin* qualified $\underline{\mathrm{A}}$, and $\underline{\mathrm{E}}$ quantify over the set of global sequences that a run may exhibit, given the traces for each process. $\underline{\mathrm{A}}$:p means that predicate p holds for all global sequences and $\underline{\mathrm{E}}$:p means that predicate p holds for some global sequence. We call formulas starting with $\underline{\mathrm{A}}$: as *strong* formulas and formulas starting with $\underline{\mathrm{E}}$: as *weak* formulas. The intuition behind the term *strong* is that a strong predicate is true no matter how fast or slow the individual processes in the system execute so long as the execution is consistent with the run. That is, it holds for all execution speeds which generate the same trace for an individual process. A *weak* predicate is true if and only if there exists at least one global sequence in which it is true. In other words, the predicate can be made true by choosing appropriate execution speeds of various processors.

Semantics defined in this paper is slightly different from that in [GarWal94]. In [GarWal94] *bool* is defined to be true on a global sequence if it is true in the first global state in the sequence. In this paper, *bool* is required to be true in the last global state. The current version is more useful and easier to understand. Intuitively, the logic in [GarWal94] is based on future while the logic in this paper is based on past. Since the past is known at any point of execution, it is easier to evaluate the formula in the current state.

Following are some examples of the strong formulas detectable by our algorithms.

1. Suppose we have developed an algorithm which works in phases. Assume that the system has three nodes and that there are three phases in the algorithm. Let predicate

$phase_{i,j}$ denote that the process P_i is in phase j. The following formula ensures that the process P_2 is in phase 3, only after all the processes have been through phase 2.

$$(\underline{A} : phase_{1,2} \hookrightarrow phase_{2,3}) \wedge (\underline{A} : phase_{2,2} \hookrightarrow phase_{2,3}) \wedge (\underline{A} : phase_{3,2} \hookrightarrow phase_{2,3})$$

2. Suppose we were testing a commit protocol. Let $Ready_i$ denote the local predicate that the process P_i is ready to commit. Then, the following formula would check that there was a certain point in the execution when all processes were ready to commit.

$$\underline{A} : \Diamond(Ready_1 \wedge Ready_2 \ldots \wedge Ready_n)$$

3. Suppose we wanted to test a distributed minimum spanning tree algorithm. Let K_i represent the local predicate that the process P_i knows its parent. Then, the following formula would indicate that the system has reached a state in which all nodes in the network know their parents.

$$\underline{A} : \Diamond(K_1 \wedge K_2 \wedge \ldots \wedge K_n).$$

3 Linked Predicates

This class of predicates is useful in detecting a sequence of events in a distributed program. We use LP_i to denote a local predicate in some process, and $LP_i(s)$ to denote that the predicate LP_i is true in the state s. We assume that the local predicate LP_i is constructed from only the local variables of that process. This means that the truthness of LP_i can change only through an internal event. In other words, external events cannot make any local predicate change from true to false or vice-versa. Thus, a predicate such as "a message has been sent from P to Q" is not considered a valid LP. Although this appears to be a limitation, the above predicate can be easily modeled in our framework by assuming that an internal event records the send of the message in some boolean variable such as msg_sent. The condition msg_sent is a valid local predicate. The above assumption is equivalent to the following:

(A1) If $(LP(s) \wedge \neg LP(s.next)) \vee (\neg LP(s) \wedge LP(s.next))$ then
1. $t \rightarrow s$ iff $t \rightarrow s.next$, and
2. $s \rightarrow t$ iff $s.next \rightarrow t$ for all t different from s and $s.next$.

(A1) says that if s and $s.next$ differ in their evaluation of LP, then their causal relationships with other states is identical.

We also use the following assumption.

(A2) All LP's evaluate to false in the artificial states $\perp_i$ and $\top_i$ for all i. This assumption is also not a restriction. It just captures the intent of defining $\perp_i$ and $\top_i$ states.

A predicate of the form $\underline{\text{A}}$:$LP_i \hookrightarrow LP_j$ means that for all global sequences, there exists an instance where LP_i is true before LP_j. $\underline{\text{A}}$:$(LP_i \hookrightarrow LP_j) \hookrightarrow LP_k$ means that for all global sequences there exists an instance where LP_i is true before LP_j which is true before LP_k. We treat $\hookrightarrow$ as a left associative operator and leave out the parentheses. We call a formula of the form $\underline{\text{A}}$: $LP_1 \hookrightarrow LP_2 \hookrightarrow \cdots \hookrightarrow LP_m$ a *strong* linked predicate. The following theorem is used in designing the algorithm for the detection of such predicates. Note that one side of the proof ($\Leftarrow$) is obvious. The converse, which is more difficult, has not been addressed in the literature. This is one of the main results of this section.

Theorem 2 *Let LP_1 and LP_2 be local predicates on processes i and j respectively. Then, for any run r, there exist states s_i in $r[i]$ and s_j in $r[j]$, such that $s_i \rightarrow s_j$, $LP_1(s_i)$ and $LP_2(s_j)$ if and only if $\underline{A} : LP_1 \hookrightarrow LP_2$.*

Proof: ($\Rightarrow$) Since $s_i \in P_i$ and $s_j \in P_j$, from (S1) we conclude that any global sequence $g \in linear(r)$ has states g_k and g_l such that $g_k[i] = s_i$ and $g_l[j] = s_j$. From (S2), we know that $findex(g, s_i) < findex(g, s_j)$. Thus, $g \models LP_1 \hookrightarrow LP_2$ is true.

($\Leftarrow$) We show that if such states do not exist, then the formula $\underline{A} : LP_1 \hookrightarrow LP_2$ is false (that is $\neg LHS \Rightarrow \neg RHS$). If LP_1 (or LP_2) is not true for any state in $r[i]$ ([$r[j]$], respectively), then the formula is trivially false. Consider the first state in $r[i]$ in which LP_1 is true. We call this state s_i. Similarly, s_j is defined using the last state in $r[j]$ in which LP_2 is true. The negation of left hand side implies that $s_i \not\rightarrow s_j$. Consider the state $s_j.next$. This state exists by (A2); it may be $\top_j$. Let t be defined as

$$t = min_{s \in S_i'}\{s \| s_j.next\}$$

Note that t could possibly be $\bot_i$. See Fig. 2.

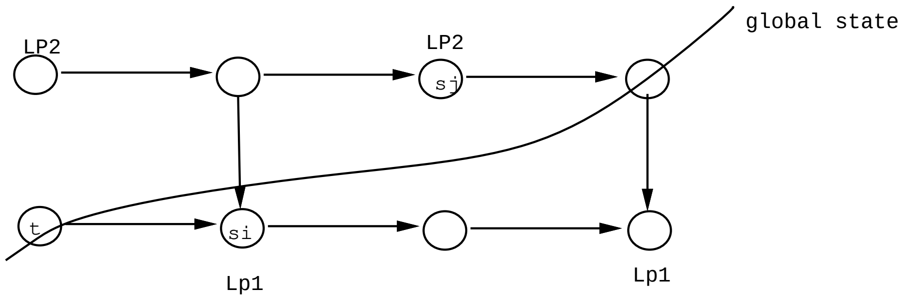

Figure 2: Linked Predicates

We now do case analysis.

Case 1: $s_i \prec t$
This means that s_i is not concurrent with $s_j.next$ by the definition of t. Since $s_i \not\rightarrow s_j$, we get that $s_i \not\rightarrow s_j.next$ from (A1). This implies that $s_j.next \rightarrow s_i$ which in turn implies that $s_j.next \rightarrow t$, a contradiction.

Case 2: $t \preceq s_i$
Since t is concurrent with $s_j.next$, by Lemma (S3), there exists a global sequence g in which t and $s_j.next$ occur in the same global state for the first time. The predicate LP_1 is not true for all preceding global states and LP_2 is false for all following global states. Thus, there are no two global states x, y such that $LP_1(x)$, $LP_2(y)$ and x occurs before y in g.

■

The above result can be generalized to a sequence of more than two local predicates [GarWal94b].

The intuition behind the algorithm to detect the strong linked predicate (in Fig. 3) is as follows. $\underline{\mathrm{A}}$:$LP_1 \hookrightarrow LP_2$ is true only if the state in which LP_1 has occurred happened before ($\rightarrow$) the state in which LP_2 occurs. If both predicates are in the same process, then occurrence of LP_1 would be known when LP_2 occurs. If LP_2 is in the different process then by the definition of $\rightarrow$ we know that there must be a message path to the second process. We use the same message path to inform the second process about occurrence of LP_1.

P_i::
var
detectflag : boolean **always** (true iff $curpred = m + 1$);
pred_list: list of {index:1..m; pred:local predicate }
/* *predicates local to this process*;*/
curpred: integer **initially** 1;
□ **Upon** $(head(pred\_list).index = curpred) \wedge (head(pred\_list).pred = true)$
begin /* *update what predicate is the next one this process is to detect* */
curpred++;
pred_list := tail (pred_list);
end;
□ **Upon rcv** (prog, hiscurpred, ...) from P_s
curpred:=max(curpred, hiscurpred);
□ **To send** /* *we include curpred in message* */
send(prog, curpred, ...) to destin;

Figure 3: Algorithm for strong linked predicates

The implementation of the algorithm is as follows. The variable *pred_list* in each process keeps the list of logical predicates local to that process in the increasing order of indices in which they appear in strong linked predicate. The variable *curpred* keeps the index of the next local predicate in the strong linked predicate which needs to be detected (as currently known by the process). If curpred becomes $m + 1$ in any process, then the strong linked predicate is detected.

We now show the correctness of the above algorithm. Let $link(s, j) = \exists s_1, s_2, ..., s_{j-1} : (s_1 \rightarrow s_2) \wedge (s_2 \rightarrow s_3) \wedge ... \wedge (s_{j-1} \rightarrow s) \wedge LP_1(s_1) \wedge \ldots \wedge LP_{j-1}(s_{j-1})$ for $j > 1$. The predicate $link(s, 1)$ is defined to be true for all s. We also use $s.x$ to refer to the value of the variable x in the state s.

The following lemma describes an assertion on the variable *curpred*.

Lemma 3 *For all local states s:*
$s.curpred = max\{j \mid link(s, j)\}$

Proof: We show that the above assertion is true for the initial state and is maintained by the program. Since *curpred* is initially 1, the assertion is trivially true for the initial state of any process. For the induction case, let $s \prec_{im} t$. We assume that the assertion holds for s, and show it to be true for t. we consider two cases:

Case 1: The event executed at s is not a receive
Since assertion holds for s, $link(s, s.curpred)$ holds. Further, if $LP_{s.curpred}(s)$ is true then $link(t, s.curpred + 1)$ holds as $s \prec_{im} t$. It is also easy to see that $s.curpred + 1$ is the maximum j such that $link(t, j)$ holds. By incrementing $s.curpred$ the assertion is maintained for t.

Case 2: A message is received at s which was sent from the state u
This part of the proof follows from the observation that $link(s, j) \wedge s \rightarrow t$ implies that $link(t, j)$. If the assertion holds for s and u, then it is maintained by taking max of $s.curpred$ and $u.curpred$.

■

Theorem 4 *At the termination of the algorithm, there exists a process for which detectflag = true if and only if the* $\underline{\mathrm{A}}$: $LP_1 \hookrightarrow LP_2 \hookrightarrow \cdots \hookrightarrow LP_m$ *is true.*

Proof: We first show that if the strong linked predicate is true, then it is detected by the algorithm. Let the strong linked predicate be true. This means that at the termination there exists a state s such that $link(s, m + 1)$ is true. From Lemma 3, $s.curpred = m + 1$. It follows that the process which has state s will have its detectflag set.

Conversely, assume that $s.detectflag$ is true, i.e., $s.curpred = m + 1$. Again from Lemma 3, this means that $link(s, m + 1)$ holds. From Theorem 2, this can happen only if $\underline{A} : LP_1 \hookrightarrow LP_2 \ldots \hookrightarrow LP_m$ is true.

■

The above algorithm requires no extra messages but does require each message to contain the value of *curpred*. Hence, each message grows in size by $O(\log m)$ bits where m is the number of local predicates in the linked predicate.

The above algorithm can also be used to detect $\underline{A} : DP_1 \hookrightarrow DP_2 \hookrightarrow \cdots \hookrightarrow DP_m$, where each DP_i is a disjunction of local predicates. The only difference in detection of such a predicate from the strong linked predicate is that an index may occur in *pred_list* of more than one process.

Miller and Choi [Mill88] have also proposed a similar algorithm for linked predicates. In their algorithm, a process p sends out a predicate marker along each channel directed away from p on detecting the local predicate. Thus, the algorithm assumes that underlying communication channels are FIFO. Note that this assumption is also exploited in stopping the program in a consistent state using an algorithm similar to that of Chandy and Lamport [ChaLam85].

4 Strong Conjunctive Predicates

Conjunctive predicates form the most interesting class of predicates in our logic. A strong conjunctive predicate is true if and only if the system will always reach a global state such that all of the given local predicates are true in that state. Formally, a strong conjunctive predicate is of the form: $\underline{A} : \Diamond(LP_1 \wedge \ldots \wedge LP_m)$, where LP_i for $1 \leq i \leq m$ are local predicates. Practically speaking, strong conjunctive predicates are most useful for good or desirable predicates (i.e. predicates which the programmer would like to be true at some point in the program). For example, in the case of a distributed two-phase commit protocol, if the master decides to commit a transaction, then it must be true that the program was in a global state where all the slaves were "ready" to commit. If the program is executed and commits, but a global state where all slave processes are "ready" does not occur, then the program has an error in it.

In this section, we present the conditions that are necessary and sufficient for a strong conjunctive predicate to hold. This is one of the main results of this paper. These conditions use the notion of *intervals*. An interval, I, is defined as a sequence of consecutive states of a trace having a beginning state (designated as $I.lo$) and an ending state (designated as $I.hi$). It is convenient to assume that $I.lo$ and $I.hi$ are distinct such that $I.lo \prec I.hi$. This is not a restriction. To model an interval with a single state it is sufficient to stutter that state once. A set of intervals, $I_{1,} \ldots, I_m$, each belonging to a different process trace is said to *overlap*, represented by, $overlap(I_1, I_2, \ldots I_m)$, if and only if the following holds:

$$\forall i, j : i, j \in \{1 \ldots m\} : I_i.lo \rightarrow I_j.hi$$

Intuitively, the notion of overlapping intervals means that all the interval *lo*'s are ordered before all the interval *hi*'s.

We assume that $m \leq n$ and $LP_1, \ldots, LP_m$ are local predicates in different processes, $P_1, \ldots, P_m$. (because $LP_1 \wedge LP_2$ is just another local predicate if LP_1 and LP_2 belong to the same process). We use $LP(I)$ to denote that the local predicate LP is true for the entire interval I.

The following Lemma shows that existence of overlapping intervals is sufficient to ensure that all global sequences go through a global state in which $(LP_1 \wedge \ldots \wedge LP_m)$ is true.

Lemma 5 $\exists I_1, \ldots, I_m : LP_1(I_1) \wedge \ldots \wedge LP_m(I_m) \wedge overlap(I_1, \ldots, I_m) \Rightarrow \underline{A} : \Diamond(LP_1 \wedge \ldots \wedge LP_m)$

Proof: Using the definition for overlapping intervals we know that:

$$\forall i, j : i, j \in \{1 \ldots m\} : I_i.lo \rightarrow I_j.hi$$

This means that all lo's must appear before all hi's in any global sequence. Therefore, every possible global sequence has a state greater than or equal to all lo's and less than or equal to all hi's. In this state, the boolean expression $LP_1 \wedge \ldots \wedge LP_m$ is true. Hence, the strong conjunctive predicate $\underline{\text{A}}$: $\Diamond(LP_1 \wedge \ldots \wedge LP_m)$ is true.

■

We now show that these conditions are also necessary. Our obligation is to show that if these conditions are violated, then there exists a global sequence in which the strong conjunctive predicate is false. Our proof of the existence of such a global sequence is constructive. The global sequence we construct will have the property that it does not go through any global state in which all LP_i are true. We call such a global sequence *pure*. Formally,

Definition 6 *A global sequence $g = g_1, g_2, ..., g_m$ is* pure *iff* $\forall k : \neg g_k \models LP_1 \wedge LP_2 \wedge \ldots \wedge LP_m$

We will construct a pure global sequence by concatenating together multiple pure global subsequences. Let g be a global sequence of the run from a consistent global state x to a consistent global state y (i.e. x is the first global state in g and y is the last global state in g) and h be a global sequence from the global state y to a global state z. Then, it is easy to see that g concatenated with h is also a global sequence from x to z. In constructing a pure global sequence we use intermediate states which satisfy certain properties.

Let x be any global state. We denote by $first(x)$ the m-tuple of intervals $(I_1(x), I_2(x), \ldots, I_m(x))$ where $I_k(x)$ is the first interval in $r[k]$ which ends after the state $x[k]$ in which LP_k is true. $first(x)$ may not exist if for some process P_k, LP_k never becomes true after $x[k]$. A global state is called *consistent* if $\forall i, j : x[i]||x[j]$. We will use only consistent global states in our description. The intermediate (consistent) global states that we use to construct our pure global sequence satisfy an admissibility property.

Definition 7 *x is an* admissible *intermediate global state if and only if either $first(x)$ does not exist or there exist k, l such that $I_k(x).lo \not\to I_l(x).hi \wedge \neg LP_k(x[k])$.*

In other words, there exists two intervals I_k and I_l such that they do not overlap ($I_k(x).lo \not\to I_l(x).hi$) and LP_k is not true in the state $x[k]$.

Now, we are ready to show that:

Lemma 8

$$\neg\exists I_1, \ldots, I_m : LP_1(I_1) \wedge \ldots \wedge LP_m(I_m) \wedge overlap(I_1, \ldots, I_m) \Rightarrow \neg\underline{A} : \Diamond(LP_1 \wedge \ldots \wedge LP_m)$$

Proof:

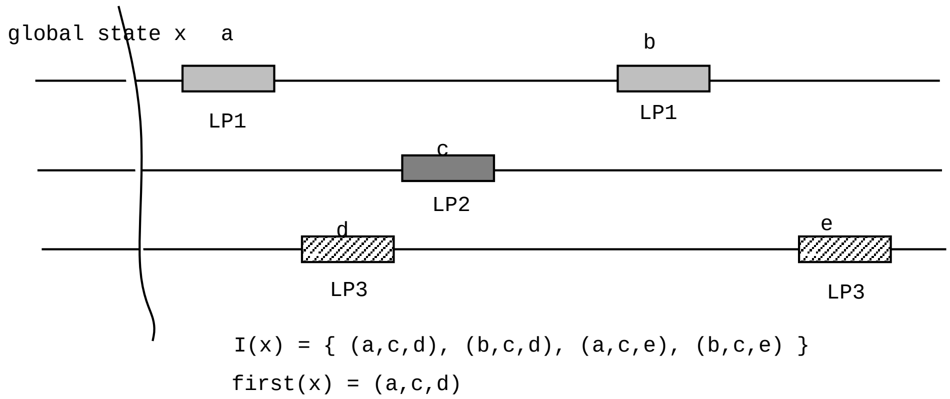

Figure 4: $\mathcal{I}(x)$ and $first(x)$

Let $\mathcal{I}(x)$ for any global state x be the set of all m-tuple of intervals in which LP_i is true for the i^{th} trace (see Fig. 4) after the local state $x[i]$. We show that if none of these m-tuple of intervals satisfy overlapping condition, then there exists a global sequence in which the distributed program is never in any m-tuple in X.

Our aim is to construct a pure global sequence g from $start = (\bot_1, \bot_2, ...\bot_m)$ to $stop = (\top_1, \top_2, ...\top_m)$. Let x be any global state such that we have built a pure global sequence from $start$ to x, and the remaining task is to build a pure global sequence from x to $stop$. Initially, we choose $x = start$. We will show a pure global sequence from x to y such that $|\mathcal{I}(y)| < |\mathcal{I}(x)|$. Thus, by continuing in this manner we will reach a global state z in which $|\mathcal{I}(z)| = 0$. From, that point all global sequences will be pure.

The $start$ state is admissible, because by assumption either $first(start)$ does not exist or there exist k, l such that $I_k(start).lo \not\to I_l(start).hi$. Moreover, LP_k is false in $\perp_k$ by (A2).

Now suppose that we are given an admissible global state x such that $|\mathcal{I}(x)| > 0$. For $1 \leq j \leq m$, let I_j be the first interval in $r[j]$ in which LP_j is true that ends after the state $x[j]$. As x is admissible and $first(x)$ exists ($|\mathcal{I}(x)| > 0$), there exist k, l such that $I_k(x).lo \not\to I_l(x).hi$ and $\neg LP_k(x[k])$. We define s to be the local state $I_l(x).hi.next$. The state s exists because of (A2) (it may be $\top_l$). We construct a global sequence from x to another admissible global state y, where y is defined as the minimum consistent global state such that $y[l] = s$. Such a global state exists because the set of all consistent cuts (ideals) is a lattice [Matt89] and that ideals grow by adding one element at a time [DJR93]. We first show that LP_k is never true between $x[k]$ and $y[k]$. It is sufficient to show that $y[k] \prec I_k(x).lo$. We know that $I_k(x).lo \not\to I_l(x).hi$. Applying (A1) twice, it follows that $I_k(x).lo.prev \not\to I_l(x).hi.next$. Thus, there exists a consistent global state z containing s such that $z[k] \preceq I_k(x).lo.prev$. As y is the minimum consistent global state with $y[l] = s$, we get that $y[k] \preceq I_k(x).lo.prev$. This implies that LP_k is never true between $x[k]$ and $y[k]$. Thus, all global sequences from x to y are pure. See Figure 5.

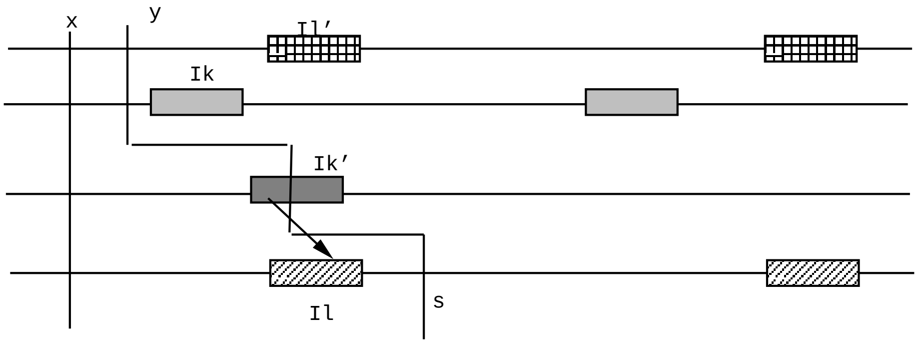

Figure 5: Illustration of the proof of strong conjunctive predicates

We still need to show that y is admissible. If $first(y)$ does not exist, we are done. Otherwise, we know that there exist k', l' such that $I_{k'}(y).lo \not\to I_{l'}(y).hi$ (see Fig. 5). If $\neg LP_{k'}(y[k'])$, then y is an admissible state and we are done. Otherwise, $y[k']$ is inside the interval $I_{k'}(y)$. If $I_{k'}(y).lo.prev \not\to s$, then we can find a consistent global state z such that $z[k'] \preceq I_{k'}(y).lo.prev$ such that $z[l] = s$. Since y is the minimum state with $y[l] = s$ and $I_{k'}(y).lo \preceq y[k']$, z cannot exist. Thus, $I_{k'}(y).lo.prev \to s$. From (A1) this is equivalent to $I_{k'}(y).lo \to I_l(x).hi$.

We now show that y is admissible because $I_l(y).lo \not\to I_{l'}(y).hi \wedge \neg LP_l(y[l])$. The second conjunct is clearly true by the definition of y. We show the first conjunct. From $I_{k'}(y).lo \to I_l(x).hi$, and $I_l(x).hi \prec I_l(y).lo$ it follows that $I_{k'}(y).lo \to I_l(y).lo$. Therefore, $I_l(y).lo \to I_{l'}(y).hi$ is inconsistent with $I_{k'}(y).lo \not\to I_{l'}(y).hi$. Thus, the first conjunct $I_l(y).lo \not\to T_{l'}(y).hi$ also holds.

■

We see from the necessary and sufficient conditions for a strong conjunctive predicate to hold that the intervals delimited by lo's (local predicate transitioning from false to true) and hi's (transitions from true to false) must overlap.

At this point, we discuss the role of (A2). Consider a scenario in which two processes P_1 and P_2 are such that LP_1 and LP_2 are true throughout the execution of P_1 and P_2 respectively. If P_1 and P_2 never communicate with each other, then there does not exist overlapping intervals for LP_1 and LP_2. However, it may seem to the reader that for any global sequence there is a global state in which both LP_1 and LP_2 are true. The global sequence for which there does not exist any global state satisfying the strong conjunctive predicate is obtained by running one process to the completion before the other starts. Clearly, unless execution of both processes are synchronized in some manner, the above sequence is a proper global sequence. By (A2) LP_1 and LP_2 are false at the initial state (before the process has begun execution) and at final state (after the process has finished its execution).

4.1 Algorithms for Detecting A Strong Conjunctive Predicate

We now describe algorithms to check whether intervals in which local predicates hold overlap. These algorithms are executed by two kinds of processes - nonchecker processes and checker processes. They are based on a slight modification of timestamp vectors as proposed by Fidge[Fidge88] and Mattern [Matt89]. Each process detects its local predicate and records the timestamp of the interval associated with the predicate. These intervals are sent to a checker process which uses them to decide if the strong conjunctive predicate became true.

Each non-checker process (Fig.6) keeps its own local *lcmvector* of timestamps. For process P_j, lcmvector[i] ($i \neq j$) is the message id of the last message from P_i (to anybody) which has a causal relationship to P_j. lcmvector[j] for process P_j is the next message id that P_j will use. Each time the local predicate of a process changes from false to true, the current value of lcmvector is remembered as an interval lo. At the next true-to-false transition (denoted by $\downarrow$ in the Fig. 6), the process sends the stored lcmvector (interval lo) and the current lcmvector (interval hi) to the checker process in a debug message. We next observe that a process is not required to send its interval every time the local predicate is detected. The interval need not be sent if there has been no message activity since the last time the interval was sent. This is because the lcmvector can change its value only when a message is sent or received. We now show that it is sufficient to send a lcmvector once after any message is received irrespective of the number of messages sent.

Let predicate $firstlmr(I)$ be true iff the local predicate is true in I for the first time since the last message was received (or the beginning of the trace). We say $scp(I_1, I_2, ..., I_m)$ is true if $I_1, I_2, ... I_m$ are the intervals in different processes making the strong conjunctive predicate true (as in Theorem 5).

Theorem 9 $\exists I_1, ..., I_m : scp(I_1, I_2, ... I_m) \Rightarrow \exists J_1, .., J_m, scp(J_1, J_2, ..., J_m) \wedge \forall k : 1 \leq k \leq m : firstlmr(J_k)$

Proof: By symmetry it is sufficient to prove the existence of J_1 such that $scp(J_1, I_2, ..., I_m) \wedge firstlmr(J_1)$. Let J_1 be the first interval in the trace of P_1 such that $LP(J_1)$ is true. Since $firstlmr(J_1)$ is true, our proof obligation is to show that $scp(J_1, I_2, ..., I_m)$. It is sufficient to show that $overlap(J_1, I_k)$ for $2 \leq k \leq m$. For any I_k, $I_1.lo \rightarrow I_k.hi$ and $J_1.lo \rightarrow I_1.lo$; therefore, $J_1.lo \rightarrow I_k.hi$. Also $I_k.lo \rightarrow I_1.hi$, because $overlap(I_k, I_1)$. Moreover, as there is no message received after $J_1.hi$ and before $I_1.hi$, the last causal message that made $I_k.lo \rightarrow I_1.hi$ true must have arrived before $J_1.hi$. Therefore, it is also true that $I_k.lo \rightarrow J_1.hi$. Hence, we conclude that $overlap(J_1, I_k)$.

∎

Process P_{id}::
var
    lcmvector: **array** [1..n] of (0..MAXMID);
    init $\forall i{:}i \neq$id: lcmvector[i] = 0, lcmvector[id] = 1;
        / last causal msg rcvd from process 1 to n, respec. */*
    Current_Interval: **record** lo, hi : (0..MAXMID);**end**;
    firstflag: boolean **init** true;
    local_pred: Boolean_Expression; */*the local pred. to be tested by this process*/*

□ **For** sending **do**
    send (prog, midgen, lcmvector, . . .);
    lcmvector[id]++ ;
□ **Upon** receive (prog, mid, msg_lcmvector, . . .) **do**
    $\forall i$: lcmvector[i] := max (lcmvector[i], msg_lcmvector[i]);
    firstflag := true;
□ **Upon** (local_pred ↑) ∧ firstflag **do**
    Current_Interval.lo := lcmvector;
□ **Upon** (local_pred ↓)∧ firstflag **do**
    Current_Interval.hi := lcmvector;
    send (dbg, Current_Interval) to CHECKERPROC;
    firstflag := false;

Figure 6: Algorithm for strong conjunctive processes - nonchecker process P_{id}

The dominant space complexity of the above algorithm is due to the array "lcmvector" which is O(n). The main time complexity involves detecting the local predicates which is the same as for a sequential debugger. In the worst case, one debug message is generated for each program message received, so the worst case message complexity is O(m_r) where m_r is the number of program messages received.

We now give the algorithm for the checker process which detects the strong conjunctive predicate using the debug messages sent by other processes. The checker process has a separate queue for each process involved in the strong conjunctive predicate. Incoming debug messages from processes are enqueued in the appropriate queue. We ensure that the checker process gets its message from any process in a FIFO order. The required computation to

```
var
    q_1 ... q_m: queue of record lo, hi: timevector;end;
    changed, newchanged: set of {1,2,...,m}
□ Upon recv(elem) from P_k do
    insert(q_k, elem);
    if (head(q_k) = elem) then begin
        changed := { k };
        while (changed ≠ φ) begin
            newchanged := {};
            for i in changed, and j in [1,2,...,m] do begin
                if head(q_j).lo ≮ head(q_i).hi then
                    newchanged:=newchanged ∪ {i};
                if head(q_i).lo ≮ head(q_j).hi then
                    newchanged:=newchanged ∪ {j};
            end; /* for */
            changed := newchanged;
            for i in changed do deletehead(q_i);
        end;/* while */
        if ∀i : ¬empty(q_i) then found:=true;
    end; /* if */
```

Figure 7: Algorithm for strong conjunctive checker process

check if the lcmvector u is less than the vector v in a different process is

$$(u[u.p] \leq v[u.p])$$

Lemma 10 *Let I and J be intervals in processes P_i and P_j with vector pairs x and y, respectively. Then, overlap(I, J)* **iff** *$(x.lo < y.hi) \wedge (y.lo < x.hi)$.*

Proof: The proof follows from the fact that if s and t are states with time vectors u and v, then $s \rightarrow t$ iff $u < v$. See [Matt89, GarTom93].

■

Thus, the task of the checker process is reduced to checking ordering between lcmvectors to determine if the intervals overlap. Because of the above Lemma, we use terms intervals and vector-pairs interchangeably. The following Lemma shows how the checker process can avoid checking all possible combinations of intervals.

Lemma 11 *Let x and y be two vector pairs at the head of their respective queues. If they do not overlap, then at least one of them can be eliminated from further consideration in checking to see if the strong conjunctive predicate is satisfied.*

Proof: In order for the strong conjunctive predicate to be true, there must exist a set of intervals, one from each queue, such that each overlaps with all the others in the set. Let two intervals x and y be at the head of their queues such that they do not overlap. This means that either $x.lo \not\to y.hi$ or $y.lo \not\to x.hi$. Assume the former without any loss of generality. We show that y can be eliminated in this case. If not, let x' be another interval in the queue of x which overlaps with y. This implies that $x'.lo \to y.hi$. Since $x.lo \to x'.lo$, we conclude that $x.lo \to y.hi$, a contradiction.

■

The checker process receives debug messages containing timestamp pairs from the other processes and executes the algorithm in Fig. 7. Each element of the queue is an interval, and the comparisons are done between $hi's$ and $lo's$ of these intervals. The checker process reduces the number of comparisons by deleting any vector-pair at the head of any queue whose hi lcmvector is not greater than lo lcmvector of vector-pairs of head of all other queues. The checker process has detected the strong conjunctive predicate to be true if it finds a set of intervals at the head of queues such that they are pairwise overlapping.

This algorithm requires at most $O(m^2p)$ comparisons where m is the number of queues each of length at most p.

5 Decentralization of the Algorithm

We now show techniques for decentralizing the above algorithm. If a set of intervals S is such that all pairs of intervals overlap, then the following holds:
$\forall x, y \in S : x.lo < y.hi$ $\qquad\qquad (P1)$
We denote this by predicate $overlap(S)$. Our aim is to show that the above condition can be checked in a decentralized manner. For this, we need the concept of greatest lower bound of a set of intervals. Let $\mathcal{X}$ be set of all intervals, where each interval x is defined as a pair of vectors $x.lo$ and $x.hi$ such that $x.lo \le x.hi$. We now define an order $\sqsubseteq$ between elements in this set as follows:

$$x \sqsubseteq y \equiv (x.lo \ge y.lo) \wedge (x.hi \le y.hi)$$

It can be easily checked that $(\mathcal{X}, \sqsubseteq)$ is a partial order. In this partial order, $x \sqcap y = (max(x.lo, y.lo), min(x.hi, y.hi))$. Then,

$$overlap(x, y) \Rightarrow (x \sqcap y) \in \mathcal{X}$$

Further, if $x_1, x_2, ..., x_m$ are such that $\forall i, j : overlap(x_i, x_j)$, then

$$\sqcap_i x_i \in \mathcal{X}$$

The following theorem shows that the process of finding $overlap(X)$ can be decomposed into smaller sets.

Theorem 12 *Let X, Y, and Z be sets of intervals, such that $X = Y \cup Z$. Then, $overlap(X)$ iff $overlap(Y) \wedge overlap(Z) \wedge overlap(\sqcap\{x|x \in Y\}, \sqcap\{x|x \in Z\})$*

Proof:($\Rightarrow$) $overlap(Y)$ and $overlap(Z)$ are clearly true because $Y, Z \subseteq X$. We need to show that

$$overlap(\sqcap\{x|x \in Y\}, \sqcap\{x|x \in Z\})$$

Let $y^* = \sqcap\{x|x \in Y\}$, and $z^* = \sqcap\{x|x \in Z\}$. Since $overlap(Y)$ and $overlap(Z)$, y^* and z^* belong to $\mathcal{X}$. To prove $overlap(y^*, z^*)$, we need to show that $(y^*.lo < z^*.hi) \wedge (z^*.lo < y^*.hi)$. We show just the first conjunct.

From $overlap(X)$, we get that
$\forall y, z \in X : y.lo < z.hi.$
In particular,
$\forall y \in Y, z \in Z : y.lo < z.hi.$
Then, by definition of y^* and z^*, we conclude that
$y^*.lo < z^*.hi$
($\Leftarrow$) We show that $(P1)$ holds for X, i.e.
$\forall y, z \in X : y.lo < z.hi$
If both y and z belong either to Y and Z, then the above is true from $overlap(Y)$ and $overlap(Z)$. Let us assume without loss of generality that $y \in Y$ and $z \in Z$. We need to show that $y.lo < z.hi$. This is true because $y.lo \leq y^*.lo < z^*.hi \leq z.hi$. The first and the last inequality follow from the definition of y^* and z^*; the middle inequality follows from $overlap(y^*, z^*)$.

$\blacksquare$

Using the above theorem and the notions of a hierarchy, the algorithm for checking the strong conjunctive predicate can be decentralized as follows. We may divide the set of processes into two groups. The group checker process checks for the strong conjunctive predicate within its group. On finding one, it sends the greatest lower bound of all intervals to a higher process in the hierarchy. This process checks the last conjunct of the above theorem. Clearly, the above argument can be generalized to a hierarchy of any depth.

6 Applications

The main application of our results are in debugging and testing of distributed programs. We have incorporated our algorithms in a distributed debugger [Hoagla92]. The online debugger is able to detect global states or sequences of global states in a distributed computation. The architecture of this distributed debugger is shown in Figure 8. With each application process, we attach two processes - a *gdb* process and a *monitor* process. *gdb* is a sequential debugger that we use for detecting local predicates. *monitor* processes are responsible for attaching vector time information with all messages. They also report to the centralized *coordinator* process whenever an interval is detected. Monitor processes also detect strong linked predicates using the algorithm outlined earlier. There is one *coordinator* process in the system. It receives all the information from monitor processes and checks for strong and weak conjunctive predicates. The coordinator also provides a single user-interface to the programmer. Our distributed debugger runs on a cluster of SUN workstations running SUNOS.

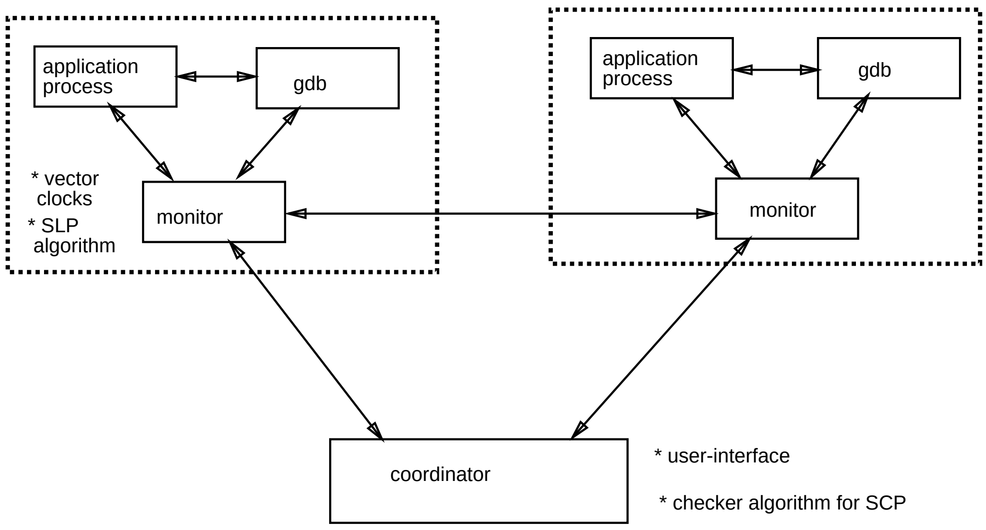

Figure 8: Architecture of Our Distributed Debugger

We have also used our algorithms to implement a trace analyzer for distributed programs [Chin91]. Our analyzer monitors a distributed program and gathers enough information to form a distributed run. The user can then ask whether any global predicate became true.

7 Conclusions

We have discussed detection of global predicates in a distributed program. Earlier algorithms for detection of global predicates proposed by Chandy and Lamport work only for stable predicates. Our algorithms detect even unstable predicates with reasonable time, space and message complexity.

In this paper, we have emphasized conjunctive predicates and not disjunctive predicates. The reason is that disjunctive predicates are quite simple to detect. Disjunctive predicates are of the form $\underline{\mathrm{A}}{:}LP_1 \vee LP_2 \vee \ldots \vee LP_m$, or of the form $\underline{\mathrm{E}}{:}LP_1 \vee LP_2 \vee \ldots \vee LP_m$. It turns out that for the simple case considered here, both expressions are equivalent. To detect a disjunctive predicate $\underline{\mathrm{A}}{:}LP_1 \vee LP_2 \vee \ldots \vee LP_m$, it is sufficient for process P_i to monitor LP_i. If any of the process finds its local predicate true, the disjunctive predicate is true.

We have also not discussed predicates of the form $\underline{A} : \Box bool$. These predicates are duals of $\underline{E} : \Diamond bool$ which have been discussed in [GarWal94].

Algorithms given in this paper detect predicates of the form $\underline{\mathrm{A}}{:}\Diamond bool$, where *bool* is a conjunction of local predicates. It would be of great interest if these algorithms can be generalized to detect predicates when *bool* is any boolean expression of local predicates.

Acknowledgements

We would like to thank Bryan Chin, Mohamed Gouda, Greg Hoagland, Jay Misra, Michel Raynal, and Alex Tomlinson for their comments on earlier versions of this work. We would also like to thank Bryan Chin for implementing offline versions of our algorithms and Greg Hoagland for incorporating our algorithms in a distributed debugger. We also thank anonymous referees for their meticulous review of an earlier version of the paper.

References

[BabMar93] O. Babaoglu and K. Marzullo, "Consistent global states of distributed systems: fundamental concepts and mechanisms", In S. Mullender, editor, *Distributed Systems*, 2nd edition, ACM Press, Frontier Series, 1993.

[Bates88] Peter Bates, "Distributed Debugging Tools for Heterogeneous Distributed Systems", *Proceedings of the 8-th International Conference on Distributed Computing Systems*, 1988, pp. 308 – 315.

[Bouge87] L. Bouge, "Repeated Snapshots in Distributed Systems with synchronous communication and their implementation in CSP", *Theoretical Computer Science*, 49: 145-169, 1987.

[ChaLam85] K.M. Chandy and L. Lamport, "Distributed Snapshots: Determining Global States of Distributed Systems", *ACM TOCS*, February 1985, pp. 63 – 75.

[Chin91] B. Chin, "An Offline Debugger for Distributed Programs", M.S. Thesis, Electrical and Computer Engineering, University of Texas at Austin, December 1991.

[CooMar91] R. Cooper and K. Marzullo, "Consistent Detection of Global Predicates", *Proc. of the ACM Workshop on Parallel and Distributed Debugging*, May 20 – 21, 1991, Santa Cruz, CA, pp. 163 – 173.

[DJR93] C. Diehl, C. Jard, J. Rampon, "Reachability Analysis on Distributed Executions", *Proc. of TAPSOFT' 93: Theory and Practice of Software Development*, April 13 – 17, Orsay, France, 1993, pp. 629 – 643.

[Fidge88] C. Fidge, "Partial Orders for Parallel Debugging", *Proc. of the ACM Workshop on Parallel and Distributed Debugging*, Madison, Wisconsin, May 1988, pp. 130 – 140.

[GarWal92] V.K. Garg and B. Waldecker, "Detection of Unstable Predicate in Distributed Programs," Proc. *12th Conference on the Foundations of Software Technology & Theoretical Computer Science*, Lecture Notes in Computer Science 652, Springer-Verlag, Dec 1992, pp 253-264.

[GarTom93] V.K. Garg, A.I. Tomlinson, "Using Induction to Prove Properties of Distributed Programs", *Proc. Symposium on Parallel and Distributed Processing*, Dallas, Texas, Dec. 1993, pp. 478–485.

[GarWal94] V. K. Garg, B. Waldecker, "Detection of Weak Unstable Predicates in Distributed Programs," *IEEE Transactions on Parallel and Distributed Systems*, Vol. 5, No. 3, March 1994, pp. 299–307.

[GarWal94b] V. K. Garg, B. Waldecker, "Detection of Strong Unstable Predicates in Distributed Programs," Technical Report, Electrical and Computer Engineering Department, University of Texas at Austin, 1994.

[HabWei88] D. Haban and W. Weigel, "Global events and global breakpoints in distributed systems", *Proc. of the 21-st Intl. Conf. on System Sciences*, Vol. 2, Jan 1988, pp 166 – 175. 1990, pp. 134 – 141.

[Hoagla92] G. Hoagland, "A Debugger for Distributed Programs", M.S. Thesis, Electrical and Computer Engineering, University of Texas at Austin, August 1992.

[HPR93] M. Hurfin, N. Plouzeau, and M. Raynal, "Detecting Atomic Sequences of Predicates in Distributed Computations," *Proc. ACM/ONR Workshop on Parallel and Distributed Debugging*, San Diego, California, May 1993, pp. 32–42.

[Lamp78] L. Lamport, "Time, Clocks, and the Ordering of Events in a Distributed System", *Communications of the ACM*, July, 1978, pp. 558 – 565.

[Matt89] F. Mattern, "Virtual time and global states of distributed systems", *Parallel and Distributed Algorithms: Proceedings of the International Workshop on Parallel and Distributed Algorithms*, Elsevier Science Publishers B. V., 1989, pp. 215–226.

[McDow89] C.E. McDowell and D.P. Helmbold, "Debugging Concurrent Programs", *ACM Computing Surveys*, vol. 21, no. 4, Dec. 1989, pp. 593 – 622.

[Mill88] B. P. Miller and J. Choi, "Breakpoints and Halting in Distributed Programs", *Proceedings of the 8-th International Conference on Distributed Computing Systems*, 1988.

[SchMat92] R. Schwartz and F. Mattern, "Detecting Causal Relationships in Distributed Computations: In Search of the Holy Grail", SFB124-15/92, Department of Computer Science, University of Kaiserslautern, Germany, December 1992.

[SpeKea86] M. Spezialetti and P. Kearns, "Efficient Distributed Snapshots", *Proceedings of the 6-th International Conference on Distributed Computing Systems*, 1986, pp. 382 – 388.

[SpeKea89] M. Spezialetti and P. Kearns, "Simultaneous Regions: A Framework for Consistent Monitoring of Distributed Systems," *Proc. of the 9-th Interl Conference on Distributed Computing Systems*, 1988, pp. 61 – 68.

[TomGar93] A.I. Tomlinson and V.K.Garg, "Detecting Relational Global Predicates in Distributed Systems," Proc. *3rd ACM/ONR Workshop on Parallel and Distributed Debugging*, San Diego, California, May 1993, pp. 21 – 31.

[WalGar91] B. Waldecker, and V. K. Garg, "Unstable Predicate Detection in Distributed Programs," *Proc. 2nd ACM/ONR Workshop on Parallel and Distributed Debugging*, Extended Abstract, 1991, Santa Cruz, CA.