An Efficient Algorithm for Detecting Conjunctions of General Global Predicate

V. K. Garg, J. Roger Mitchel

TR TR-PDS-1996-005 July 1996

Parallel & Distributed Systems group
Department of Electrical & Computer Engineering
University of Texas at Austin
Austin, Texas 78712

An Efficient Algorithm for Detecting Conjunctions of General Global Predicates

Vijay K. Garg* J. Roger Mitchell†

Parallel and Distributed Systems Laboratory
email: pdslab@ece.utexas.edu
Electrical and Computer Engineering Department
The University of Texas at Austin,
Austin, TX 78712
http://maple.ece.utexas.edu

July 15, 1996

*supported in part by the NSF Grants ECS-9414780, CCR-9520540, a TRW faculty assistantship award, a General Motors Fellowship, and an IBM grant

†supported in part by a Virginia & Ernest Cockrell fellowship

Abstract

The ability to detect global predicates is important for distributed fault monitoring systems as well as distributed debugging. In this paper, we present an efficient algorithm to detect if the conjunction of two nonlocal predicates is possibly true in a distributed computation. For such global predicates, our algorithm is significantly more efficient than the previous algorithms by Cooper and Marzullo, and by Stoller and Schneider.

1 Introduction

The detection of global conditions is a fundamental problem in an asynchronous distributed system. In such a system a process cannot know the state of other processes at any given time. This creates difficulty in detecting conditions, or predicates, spread across the system. Deadlock and termination are two such global predicates. Detection of these predicates is useful for distributed debugging, monitoring distributed systems for faults, and other areas of distributed computing.

There are two interpretations of truthness of a global predicate in a distributed computation. A global predicate denoted by Φ is *possibly* true [CM91], (or true in the weak sense [GW92]) if there exists a consistent global state in the computation in which the global predicate is true. A global predicate is *definitely* true [CM91], (or true in the strong sense [GW92]) if all sequential observations of the computation go through a consistent global state in which the global predicate is true. In this paper, we will restrict ourselves to the first interpretation.

Cooper and Marzullo [CM91] discuss methods of detecting any global predicate. However their method involves searching the entire lattice of possible global states, a search which is $O(m^n)$, where m is the number of states in a process, and n is the number of processes. By restricting the class of predicates detected, algorithms have been proposed which are of polynomial time. For example, the class of conjunctive predicates can be detected in $O(n^2m)$ time using the algorithms proposed by Garg and Waldecker [GW92].

Stoller and Schneider [SS95] combine the approach of Garg and Waldecker with that of Cooper and Marzullo. Their algorithm, like the algorithm of Cooper and Marzullo, can detect any general global predicate Φ; and like the algorithm of Garg and Waldecker, can exploit the structure of the predicate Φ to reduce the computational complexity of detection. They assume that the global predicate is a conjunction of predicates of the form

$$\Phi(x_1, ..x_j), \tag{1}$$

where $\Phi()$ is a predicate with variables, x_i, from different processes. An example of a predicate in this form is

$$(x_1 = x_2) \wedge (x_3 > x_4) \wedge (x_5 < 10), \tag{2}$$

where $x_i's$ are variables on different processes. When the global predicate has only one conjunct, their approach reduces to that of Cooper and Marzullo; and when each of the conjuncts is local to a single process, their approach reduces to that of Garg and Waldecker. They introduce the notion of a *fixed set* of a global predicate. A fixed set of a global predicate is a set of variables such that by fixing these variables in the predicate, the predicate reduces to a conjunction of local predicates. The sets $\{x_1, x_3\} and \{x_2, x_3\}$ are some examples of fixed sets for the global predicate in equation 2. The algorithm proposed by Stoller and Schneider is exponential in the size of the fixed set.

In this paper we present a significantly faster algorithm for a restricted class of predicates. Specifically, Stoller and Schneider's algorithm requires $O(m^{j+k-1})$ time for a predicate of the form

$$\Phi_1(x_1, ..., x_j) \wedge \Phi_2(y_1, ..., y_k),$$

where m is the number of states at one process. We present an algorithm which requires $O(m^l)$ time, where $l = max(j, k)$. We use many novel ideas in the algorithm and therefore make the following contributions:

- We identify a property of predicates called *monotonicity* which is crucial for efficient evaluation of global predicates in real-life applications. This property allows detection algorithms to consider at most as many states as the number of messages sent and received by the process as opposed to the total number of states from a process.
- We introduce the notion of *successor* vector clocks. These clocks in conjunction with *predecessor* vector clocks [Mat89, Fid89] can be used to identify states concurrent with the given state on any process.
- We introduce a novel technique which transforms the problem of predicate detection into a computational geometric problem. This transformation may have applications in other off-line analysis problems of distributed computations. For predicate detection, the corresponding computational geometric problem reduces to checking whether any of the given points lies inside any of the given boxes.
- We give an efficient algorithm to the computational geometric problem. Our algorithm for this problem may also be of independent interest.

This paper is organized as follows. Section 2 presents our model of distributed computation and the notation that is used in the paper. Section 3 describes our algorithm and its correctness. Section 4 presents some extensions and optimizations of the basic algorithm.

2 Model

We assume a loosely-coupled message-passing system without 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. In this paper, we will be concerned with a single run r of a distributed program. Each process P_i in that run generates a single execution trace $r[i]$ which is a finite sequence of *states*. The state of a process is defined by the value of all its variables including its program counter. The state transition occurs in any process due to an internal action, the send of a message, or the receive of a message.

We define the usual *causally precedes* relation ($\rightarrow$) on states as follows. A state $s \in r[i]$ causally precedes ($\rightarrow$) a state $t \in r[j]$ if and only if one of the following conditions holds.

1. $i = j$ and s occurs before t in $r[i]$. We say that $s \prec t$ (s locally precedes t) when this is true.
2. The action following s is the send of a message and the action preceding t is the receive of that message.

3. There is a state u such that $s \rightarrow u$ and $u \rightarrow t$.

We use $s \| t$ to denote that states s and t are concurrent. That is, $s \| t \equiv s \not\rightarrow t \wedge t \not\rightarrow s$. We use g and h to represent global states. A global state, g is a set of states $\cup_i g[i]$ for which $g[i] \in r[i]$ and $\forall i, j : i \neq j : g[i] \| g[j]$. The global state g is also called a consistent cut. We also use the notion of sub-cuts - a sub-cut is a cut across a subset of processes in the system. We use $g(\Phi)$ to represent a consistent sub-cut across a minimal set of processes in which the predicate Φ is true. This notation uniquely specifies the local states used to satisfy Φ. The notation $g(\Phi_1) \| g(\Phi_2)$ means that every pair of states in the two sub-cuts satisfying Φ_1 and Φ_2 are concurrent; we say, then, that $g(\Phi_1)$ and $g(\Phi_2)$ are concurrent. A set of cuts is written as G.

2.1 State Interval

A state interval is a sequence of states between two *external* events where an external event is the sending or receiving of a message, the beginning of the process and the termination of the process. Formally, the k^{th} interval in P_i (denoted by (i, k)) is the subsequence of $r[i]$ between the $(k-1)^{th}$ and k^{th} external events. For a given interval (i, k), if k is out of range then (i, k) refers to $\perp$ which represents a sentinel value (or a "null" interval). The notion of intervals is useful because the relation of two states belonging to the same interval is a congruence with respect to $\rightarrow$. That is, the relation of two states being in the same interval is an equivalence relation and for any two states s, s' in the same interval and any state u from a different interval: $(s \rightarrow u \Leftrightarrow s' \rightarrow u)$ and $(u \rightarrow s \Leftrightarrow u \rightarrow s')$. We exploit this congruence in our algorithms by assigning a single timestamp to all states belonging to the same interval. An example of a distributed computation and its state intervals is shown in Figure 1.

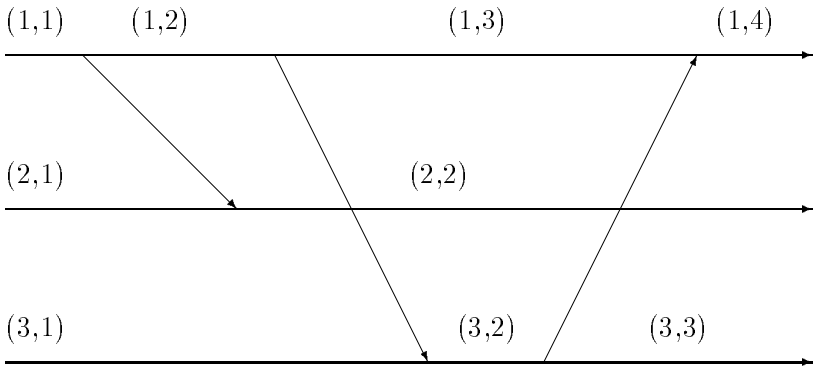

Figure 1: State intervals

2.2 Predecessor and Successor functions

The predecessor and successor functions are defined as follows for any state u and $1 \leq i \leq n$:

$$pred.u.i = \max\{v \in r[i] | v \rightarrow u\}$$

$$succ.u.i = \min\{v \in r[i] | u \rightarrow v\}$$

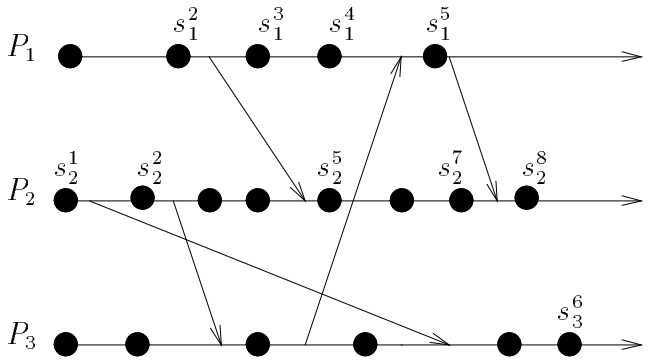

Figure 2: Demonstration of the *pred* function: $pred.s_1^5.2 = s_2^2$, $pred.s_3^6.2 = s_2^2$, $pred.s_1^4.2 = \perp$.

Consider a state $s_i \in r[i]$. The predecessor of s_i in $r[j]$, denoted $pred.s_i.j$, is the latest state in $r[j]$ which causally precedes s_i. Thus, if $(pred.u.i = v)$ then v is the maximum element in $r[i]$ which causally precedes u. If no element in $r[i]$ causally precedes u, then $pred.u.i = \perp$. Notice that the predecessor or s_i^n on $r[i]$ is just the previous state s_i^{n-1}. Figure 2 shows an example of the predecessor function.

Due to the congruence for states belonging to the same interval, the *pred* and *succ* functions and the $||$ relation are well defined on intervals.

The following lemma establishes some useful properties of *pred* and *succ*.

Lemma 1 *Let u and v be states in $r[i]$ and $r[j]$ respectively.*

1. $(p,i) = pred.(q,j).p \quad \Leftrightarrow \quad (\forall k : k > i : (p,k) \not\rightarrow (q,j)) \wedge (p,i) \rightarrow (q,j)$
2. $(p,i) = succ.(q,j).p \quad \Leftrightarrow \quad (\forall k : k < i : (q,j) \not\rightarrow (p,k)) \wedge (q,j) \rightarrow (p,i)$
3. $v \preceq pred.u.j \Leftrightarrow succ.v.i \preceq u.$
4. $succ.(pred.u.j).i \preceq u \preceq pred.(succ.u.j).i$

Proof:

1. Follows from the definition of *pred*.
2. Follows from the definition of *succ*.
3. First assume that $v \preceq pred.u.j$. This implies $v \rightarrow u$ which in turn implies that $succ.v.i \preceq u$. The converse is similar.
4. First substitute $pred.u.j$ for v in part (1). We get that

$$pred.u.j \preceq pred.u.j \Leftrightarrow succ.(pred.u.j).i \preceq u$$

which is equivalent to $succ.(pred.u.j).i \preceq u$. The proof for the second inequality is similar.

□

In addition to *pred* and *succ*, we also use *prev* and *next*. For any state u in $r[i]$, $prev.u$ returns the state preceding u in $r[i]$ if there exists one, else it returns $\perp$. The definition of $next.u$ is similar.

3 The algorithm

For simplicity, we first give an algorithm for predicates of the form

$$\Phi_1(x_1, x_2) \wedge \Phi_2(y_1, y_2).$$

This algorithm is of complexity $O(m^2)$. This is faster than the previously known algorithm by Stoller and Schneider [SS95] which requires $O(m^3)$ time. Later we briefly explain how to extend the algorithm to the more general case.

There are three steps in the algorithm - computing the subcuts that satisfy Φ_1 and Φ_2, transformation of the problem to a computational geometric problem, and finally, solution of the transformed problem. We discuss each of these steps next.

3.1 Solving sub-problems

In the first step, we compute the set of consistent sub-cuts in process P_1 and P_2 that satisfy $\Phi_1(x_1, x_2)$. For example, if $\Phi_1(x_1, x_2) = (x_1 < x_2)$, then we find pairs of states $s \in P_1$ and $t \in P_2$ such that s and t are concurrent and the value of x_1 in s is less than or equal to x_2 in t. If there are at most m distinct states in each of P_1 and P_2, and evaluation of the predicate Φ_1 requires $O(1)$ time, then this step can be done in $O(m^2)$ time. This step can be viewed as an application of Cooper and Marzullo's algorithm for detection of Φ_1.

It is important to observe that in most applications, m as defined earlier is too large for this step to be practical. However, we now show that for most predicates Φ_1, m can be taken as the total number of state intervals. The total number of state intervals is much smaller than the total number of states.

The property of any global predicate Φ that allows this reduction is considered next.

Definition 2 *Assume that x_1 takes its value from a set totally ordered with respect to a relation $<$. We say that Φ is monotone with respect to x_1 if it satisfies the following equation:*

$$\forall a, x_2 : \Phi(a, x_2) \Rightarrow (\forall b : b < a : \Phi(b, x_2))$$

or,

$$\forall a, x_2 : \Phi(a, x_2) \Rightarrow (\forall b : a < b : \Phi(b, x_2)).$$

Informally, a predicate is monotone with respect to a variable x_1 if replacing the variable by a larger value (or by a smaller value) while keeping all other variables the same does not violate the truthness of the predicate. For example, consider the predicate $\Phi = x_1 < x_2$ where x_1 and x_2 are integers. Then, Φ is monotone with respect to x_1 because if $\Phi(x_1, x_2)$ holds for a certain value of x_1, then it would continue to do so for any smaller value of x_1. Φ is also monotone with respect to x_2 because if $\Phi(x_1, x_2)$ holds for a certain value of x_2, then it would continue to do so for any

larger value of x_2. As another example, consider the conjunctive predicate $x_1 \wedge x_2 \ldots \wedge x_n$ where each x_i is a boolean variable. By viewing the boolean domain as a totally ordered set with *false* to be defined as smaller than *true*, it can be easily seen that the conjunctive predicate is also monotone with respect to all variables. An example of a predicate that is not monotone with respect to x_1 or x_2 is $(x_1 = x_2)$.

Monotonicity of a predicate allows us to restrict our attention to state intervals rather than states. For example, for the predicate $(x_1 < x_2)$, it is sufficient to keep the state with the smallest value of x_1 for each state interval. Similarly, it is sufficient to keep the state with the largest value of x_2 for each state interval in P_2. This reduction also exploits the fact that if two local states say s and s' on the same process are separated only by internal events, then they are indistinguishable to other processes so far as consistency is concerned. That is, if t is a local state on some other process, then $s||t$ if and only if $s'||t$. Note that these observations have also been exploited implicitly in [GW94, TG93].

We denote by $G(\Phi_1)$ and $G(\Phi_2)$ the set of all consistent sub-cuts which satisfy Φ_1 and Φ_2. In the worst case, $G(\Phi_1)$ and $G(\Phi_2)$ are of size $O(m^2)$. However, as mentioned above, for monotonic predicates the value of m is at most the number of messages sent or received by any process. From now on we will use state and state interval interchangeably with the understanding that for monotonic predicates, at most one state will be used from each state interval.

3.2 Transformation of the Predicate Detection Problem

We now describe the transformation of the predicate detection problem. Each element in $G(\Phi_1)$ can be viewed as a point in the 2-dimensional space. For example, if the third interval on process P_1 (denoted by $(1,3)$) and the fifth interval on P_2 (denoted by $(2,5)$) together satisfy Φ_1, then we view this consistent cut as a point $(3,5)$ in the plane generated by states in processes P_1 and P_2. For our example of the predicate $(x_1 < x_2)$ the point $(3,5)$ signifies that the state interval $(1,3)$ is concurrent to the state interval $(2,5)$ and the minimum value of x_1 in the interval $(1,3)$ is less than the maximum value of x_2 in the interval $(2,5)$.

Next, we construct a rectangle in the above plane for each sub-cut $(u,v) \in G(\Phi_2)$. Intuitively, the rectangle represents the set of all sub-cuts in P_1 and P_2 which are consistent with the sub-cut (u,v). This is done as follows. Let

$$s_l = next(max(pred.u.1, pred.v.1)).$$

Further, let

$$s_h = prev(min(succ.u.1, succ.v.1)).$$

The following lemma shows the purpose of states s_l and s_h.

Lemma 3 *Any state s such that $s_l \preceq s \preceq s_h$ is concurrent with u and v. Conversely, any state s in P_1 that is concurrent with both u and v satisfies $s_l \preceq s \preceq s_h$.*

Proof: From the definition of s_l and Lemma 1, it follows that $pred.u.1 \prec s_l$. Since $s_l \preceq s$ we get that $pred.u.1 \prec s$. From Lemma 1, it follows that $s \not\rightarrow u$. Similarly, $s \preceq s_h$ and $s_h \prec succ.u.1$, implies that $u \not\rightarrow s$. Therefore, $s||u$. By repeating the argument for v, we get $s||v$.

Conversely, let s be concurrent with u and v. Since $s \not\rightarrow u$ and $s \not\rightarrow v$, it follows that $max(pred.u.1, pred.v.1) \prec s$. Thus, $s_l \preceq s$. The argument for the upper bound is similar. □

Similarly, we find

$$t_l = next(max(pred.u.2, pred.v.2))$$

and

$$t_h = prev(min(succ.u.2, succ.v.2)).$$

The rectangle is then given by the following two endpoints - (s_l, t_l) representing the lower left point of the rectangle and (s_h, t_h) representing the upper right point of the rectangle. Let $R(\Phi_2)$ be the set of all rectangles constructed from $G(\Phi_2)$.

The following result shows that the problem of checking whether any sub-cut in $G(\Phi_1)$ is concurrent with any sub-cut in $G(\Phi_2)$ is equivalent to checking whether any point corresponding to a sub-cut in $G(\Phi_1)$ lies in any of the rectangles in $R(\Phi_2)$.

Theorem 4 *For any global predicate* $\Phi_1(x_1, x_2) \wedge \Phi_2(y_1, y_2)$,
$\exists (s,t) \in G(\Phi_1), (u,v) \in G(\Phi_2) : (s,t) \| (u,v)$ *iff* $\exists (s,t) \in G(\Phi_1)$ *which lies in one of the rectangles in* $R(\Phi_2)$.

Proof: Using Lemma 3. □

We now turn our attention to the complexity of this transformation. Since there are $O(m^2)$ sub-cuts in $G(\Phi_2)$, the operation of defining these rectangle needs to be done in $O(1)$ time per rectangle. We achieve this bound by use of vector clocks for predecessors as well as successors. The traditional vector clock implements the predecessor function. We now show a construction for successor vector clocks.

We describe an off-line version of a successor vector clock denoted by sv. Let the interval of the final state on process P_i be (i, f_i). Then, the final state of any process P_i is timestamped with a vector sv defined as $sv[i] = f_i$ and $sv[j] = \infty$ for all $j \neq i$. The rules to construct the sv vector of other state intervals are as follows. The intervals are timestamped with sv in a backward manner. Assume that the state s on P_i immediately precedes t which has already been timestamped. If the transition from s to t was a receive, then the following action is taken.

$$s.sv[j] = \begin{cases} t.sv[j] & \text{if } j \neq i \\ t.sv[j] - 1 & \text{otherwise} \end{cases}$$

If the transition was a send of a message which was received at state u, then the following action is taken

$$s.sv[j] = \begin{cases} min(t.sv[j], u.sv[j]) & \text{if } j \neq i \\ t.sv[j] - 1 & \text{otherwise} \end{cases}$$

The above rules are the simple dual of predecessor vector clocks. The following states that the vector sv indeed implements the *successor* function. In the following lemma, we identify any state s on P_i with $s.sv[i]$, the component i of its successor vector clock.

Lemma 5 *Let* s *be any state in* P_i. *Then, for any* $j \neq i$,

$$s.sv[j] = succ.s.j$$

Proof: Follows easily from the rules of updating sv and induction. □

Given predecessor and successor clocks, it is now easy to construct a rectangle for each sub-cut in $G(\Phi_2)$. Observe that end points of rectangles are constructed using functions $pred$, max, $next$, $succ$, min and $prev$ each of which can be computed in constant time.

3.3 Solution of the transformed problem

From Theorem 4, checking whether a sub-cut in $G(\Phi_1)$ is consistent with a sub-cut in $G(\Phi_2)$ is equivalent to checking whether there exists a rectangle in $R(\Phi_2)$ which includes a point corresponding to a sub-cut in $G(\Phi_1)$. Since there are $O(m^2)$ points and $O(m^2)$ rectangles, a simple brute force algorithm of checking each point for each rectangle will lead to inefficient $O(m^4)$ algorithm. We now describe a $O(m^2)$ algorithm for this problem. The algorithm shown in Figure 3 detects whether any of up to m^2 points on an $m \times m$ plane lie within any of up to m^2 rectangles. For the case when the number of points and the number of rectangles is much less than m^2, a more efficient algorithm can be used. This will be given in a following section.

```
boolean is_pt_in_rectangle() {
    istrue[m, m]: 0..1; /* istrue[s, t] = 1 iff (s, t) ∈ G(Φ1) */
    R: list of record /* rectangles specified by two corners. */
        {int s_l, t_l, s_h, t_h};

    int A[m + 1, m + 1]; /* number of points that lie between (0,0) and (s,t) */

    for i = 0 to m do
        A[0,i]=0; A[i,0]=0;

    /* Array assignment: coding A with the points */
    for i = 1 to m do
        for j = 1 to m do
            A[i, j] = A[i − 1, j] + A[i, j − 1] - A[i − 1, j − 1] + istrue[i, j];

    /* Rectangle/point overlap check */
    for (s_l, t_l, s_h, t_h) ∈ R do begin
        /* exclusion inclusion principle */
        if A[s_h, t_h] − A[s_l − 1, t_h] − A[s_h, t_l − 1] + A[s_l − 1, t_l − 1] > 0 then
            return true;
    return false;
}
```

Figure 3: Detecting if any point in an $m \times m$ plane is contained in any rectangle.

The algorithm *is_pt_in_rectangle* operates as follows. The arrays *istrue* and R store the coordinate values of the points and rectangles created in steps 2 and 3. The values stored by R specify two opposite corners as (s_l, t_l) and (s_h, t_h). These arrays are assigned before calling the algorithm, an operation of complexity $O(m^2)$. The algorithm uses another array A of size $m \times m$. The array $A[s, t]$ is used to store the total number of points that are in the rectangle defined by $(0, 0)$ and (s, t). The first column and row in A is first initialized to zeroes. The "array assignment" loop then assigns the correct value to other entries of A. The "Rectangle/point overlap check" checks for overlaps as follows. For all rectangles, the values in A at the four corners of the rectangle are used to determine

if a point is contained within. Using set-theoretic inclusion exclusion principle it follows that the number of points in the rectangle (s_l, t_l, s_h, t_h) is given by

$$A[s_h, t_h] - A[s_l - 1, t_h] - A[s_h, t_l - 1] + A[s_l - 1, t_l - 1]$$

The algorithm can check m^2 points against m^2 rectangles with time complexity $O(m^2 + m^2)$ or $O(m^2)$. The algorithm uses $O(m^2)$ space; however, we later show how the space overhead can also be reduced to $O(m)$. The correctness of the *is_pt_in_rectangle* algorithm is proven next.

First we show that $A[s,t]$ equals the number of points which are contained in the rectangle defined by $(0,0)$ and (s,t).

Lemma 6 *After the array assignment of A,*
$A[s,t] = |B(s,t)|$ *where* $B(s,t) = \{(i,j) | istrue(i,j) \wedge 0 \leq i \leq s \wedge 0 \leq j \leq t\}$

Proof: We use induction on (s,t). If either of them is 0, then by initial assignment $A[s,t]$ is 0. Since $B(s,t)$ is empty for this case, the assertion holds. Now consider the case when both s and t are strictly greater than 0. The program calculates $A[s,t]$ using:

$$A[s,t] = A[s-1,t] + A[s,t-1] - A[s-1,t-1] + istrue[s,t]$$

But using induction hypothesis, and the fact that the points in $B(s-1,t-1)$ are common in $B(s-1,t)$ and $B(s,t-1)$ we get that $A[s,t]$ equals cardinality of $B(s,t)$. □

Theorem 7 *The function* is_pt_in_rectangle *returns true iff there exists a point and a rectangle such that the point is contained in the rectangle.*

Proof: Using Lemma 6, $A[s,t]$ contains the number of points in the rectangle $(0,0,s,t)$. It follows that $A[s_h, t_h] - A[s_l - 1, t_h] - A[s_h, t_l - 1] + A[s_l - 1, t_l - 1]$ is equal to the number of points in the rectangle (s_l, t_l, s_h, t_h). □

Finally, we consider the complexity of this algorithm. It is easy to verify that each step in *is_pt_in_rectangle* takes at most $O(m^2)$ time.

4 Discussion: Optimization and Extensions

4.1 Reducing the Space Complexity

One drawback of the above algorithm is that it requires $O(m^2)$ space. This may be unacceptable when m is large. We now show that the space requirements of the above algorithm can be reduced to $O(m)$. First, we assume that the set of points are provided as a list of coordinates rather than a $m \times m$ boolean array. This is similar to the way the rectangles are provided in the algorithm of Fig. 3. Next, we observe that the algorithm uses four values from the array A for each rectangle (s_l, t_l, s_h, t_h):

$$A[s_h, t_h], A[s_l - 1, t_h], A[s_h, t_l - 1], A[s_l - 1, t_l - 1].$$

We will store these values with each rectangle. The only problem is how to determine these values. We cannot use the entire array A since that itself is $O(m^2)$. The next important observation is that

calculation of the i^{th} row of A does not need any entry from any row lower than $i-1$. A similar observation applies to columns. As a result, keeping the previous row of A is sufficient to calculate the next row of A. With above observations, it is easy to develop an $O(m)$ space algorithm. We leave the details to the reader.

4.2 Finding k points in l rectangles, when $\max\{k,l\} \ll m$

Since there are m state intervals in a process, in the worst case there can be $O(m^2)$ consistent subcuts for processes P_1 and P_2. This implies that if almost all subcuts satisfied Φ_1, then the size of $G(\Phi_1)$ is also $O(m^2)$. Similar reasoning can be applied for the number of rectangles in $R(\Phi_2)$. However, in practice, it can be expected that the number of points in $G(\Phi_1)$ as well as the number of rectangles in $R(\Phi_2)$ is much smaller. In this case, we show that the above algorithm can be modified to reduce its computational time. We first encode each rectangle using its bottom left corner and the upper right corner. With this convention, we have $k+2l$ points in all - k for the points themselves and $2l$ for the l rectangles. These $k+2l$ points can be sorted for each of the two dimensions, an operation which is not more than $O(m \log m)$ where $m = \max\{k,l\}$. The coordinates of each point are now the rank of each dimension. A $(k+2l) \times (k+2l)$ space is created in which to place all points. The *is_pt_in_rectangle* algorithm is now used to detect any of the k points within any of the l rectangles. This approach requires $O((k+2l)^2)$ operations.

As an example, suppose that there are two points (3,15) and (6,12) and a rectangle specified by opposite corners (5,7) and (10,21). The coordinates can be ordered for the two dimensions as follows: 3,5,6,10 and 7,12,15,21. Mapping these sequences to the sequence 1,2,3,4 results in a mapping of the points to (1,3) and (3,2) and the rectangle corners to (2,1) and (4,4). Running the algorithm on this reduced space will give the same result as the original: a point is contained in the rectangle.

4.3 The algorithm for $j+k$ variables

We next demonstrate how to detect a predicate of the form

$$\Phi_1(x_1, ..., x_j) \wedge \Phi_2(y_1, ..., y_k)$$

where $x_1, ..., x_j$ are on $P_1, ..., P_j$ and $y_1, ..., y_k$ are on $P_{j+1}, ..., P_{j+k}$. The algorithm works as follows:

1. All $g(\Phi_1)$ and $g(\Phi_2)$ are found, and placed in the sets $G(\Phi_1)$ and $G(\Phi_2)$. These operations are $O(m^j)$ and $O(m^k)$ respectively.

2. The sub-cuts $g(\Phi_1) \in G(\Phi_1)$ found in step 1 are now used to create points for a j-dimensional space. Each sub-cut specifies one point. As for the case of 2 variables, we also determine a j-dimensional box for each of the sub-cut in $G(\Phi_2)$. By using predecessor and successor clocks, the box corresponding to a sub-cut in $G(\Phi_2)$ can be determined in $O(jk)$ operations. Since there are $O(m^k)$ sub-cuts, finding all the boxes requires $O(jkm^k)$ operations.

3. Following the procedure of the *is_pt_in_rectangle* algorithm, all points and boxes found in steps 2 and 3 are checked for overlap. Since the space is $O(m^j)$, this check can be done with $O(m^j)$ operations. If any point lies in a box, then the predicate is true.

Therefore, for the general case, the complexity is $O(m^j + m^k + jkm^k + m^j) = O(m^j + jkm^k)$.

5 Conclusions

We have presented an algorithm for detecting the conjunction of two nonlocal predicates of the form $\Phi_1(x_1, ..., x_j) \wedge \Phi_2(y_1, ..., y_k)$ which has time complexity of $O(m^l)$ where $l = max\{j, k\}$. Our algorithm is significantly faster than any previously known algorithm. Specifically, Stoller and Schneider's algorithm requires $O(m^{j+k-1})$ operations while Cooper and Marzullo's algorithm requires $O(m^{j+k})$ operations for this predicate.

References

[CM91] R. Cooper and K. Marzullo. Consistent detection of global predicates. In *Proc. of the Workshop on Parallel and Distributed Debugging*, pages 163–173, Santa Cruz, CA, May 1991. ACM/ONR.

[Fid89] C. J. Fidge. Partial orders for parallel debugging. *Proceedings of the ACM SIGPLAN/SIGOPS Workshop on Parallel and Distributed Debugging, published in ACM SIGPLAN Notices*, 24(1):183–194, January 1989.

[GW92] V. K. Garg and B. Waldecker. Detection of unstable predicates in distributed programs. In *Proc. of 12th Conference on the Foundations of Software Technology & Theoretical Computer Science*, pages 253–264. Springer Verlag, December 1992. Lecture Notes in Computer Science 652.

[GW94] V. K. Garg and B. Waldecker. Detection of weak unstable predicates in distributed programs. *IEEE Transactions on Parallel and Distributed Systems*, 5(3):299–307, March 1994.

[Mat89] F. Mattern. Virtual time and global states of distributed systems. In *Parallel and Distributed Algorithms: Proc. of the International Workshop on Parallel and Distributed Algorithms*, pages 215–226. Elsevier Science Publishers B.V. (North-Holland), 1989.

[SS95] S. D. Stoller and F. B. Schneider. Faster possibility detection by combining two approaches. In *Proc. of the Workshop on Distributed Algorithms*, France, September 1995.

[TG93] A. I. Tomlinson and V. K. Garg. Detecting relational global predicates in distributed systems. In *Proc. of the Workshop on Parallel and Distributed Debugging*, pages 21–31, San Diego, CA, May 1993. ACM/ONR.