

# A scalable deadlock detection algorithm for UPC collective operations

Indranil Roy, Glenn R. Luecke, James Coyle and Marina Kraeva

Iowa State University's High Performance Computing Group, Iowa State University,  
Ames, Iowa 50011, USA  
iroy@iastate.edu, grl@iastate.edu, jjc@iastate.edu, kraeva@iastate.edu.

## Abstract

Unified Parallel C (UPC) is a language used to write parallel programs for shared and distributed memory parallel computers. Deadlock detection in UPC programs requires detecting deadlocks that involve either locks, collective operations, or both. In this paper, a distributed deadlock detection algorithm for UPC programs that uses run-time analysis is presented. The algorithm detects deadlocks in collective operations using a distributed technique with  $O(1)$  run-time complexity. The correctness and optimality of the algorithm is proven. For completeness, the algorithm is extended to detect deadlocks involving both locks and collective operations by identifying insolvable dependency chains and cycles in a shared wait-for-graph (WFG). The algorithm is implemented in the run-time error detection tool UPC-CHECK and tested with over 150 functionality test cases. The scalability of this deadlock detection algorithm for UPC collective operations is experimentally verified using up to 8192 threads.

## 1 Introduction

Deadlocks in complex application programs are often difficult to locate and fix. Currently UPC-CHECK [3] and UPC-SPIN [4] are the only tools available for the detection of deadlocks in UPC programs. UPC-SPIN employs a model-checking method which inherently does not scale beyond a few threads. In addition, every time the program is modified, the model has to be updated. In contrast, UPC-CHECK uses the algorithm presented in this paper to automatically detect deadlocks at run-time for programs executing on thousands of threads.

This new algorithm not only detects deadlocks involving UPC collective operations, but also verifies the arguments passed to the collective operation for consistency. The run-time complexity of this algorithm is shown to be  $O(1)$ . The algorithm has been extended to detect deadlocks involving both collective operations and locks. The run-time complexity of the extended algorithm is  $O(T)$ , where  $T$  is the number of threads. Using this deadlock detection algorithm UPC-CHECK detects all deadlock error test cases from the UPC RTED test suite [2].

The rest of this paper is organized as follows. Section 2 provides the background of various existing deadlock detection techniques. In Section 3, a new algorithm to detect potential deadlocks due to incorrect usage of UPC collective operations is presented. The correctness and run-time complexity analysis of the algorithm are also provided. Section 4 describes the extended algorithm to detect deadlocks involving both locks and collective operations. The scalability of this deadlock detection algorithm is experimentally confirmed in Section 5. Finally, Section 6 contains the concluding remarks.

## 2 Background

Out-of-order calls to collective operations on different threads may create a deadlock. Even when the calls to collective operations are in-order, various non-local semantics dictate that consistent arguments need to be used in all participating threads. Non-adherence to these semantics could lead to a deadlock or departure from intended behavior of the program. However, building scalable tools to detect such errors remains a challenge.

Model-checking tools like MPI-SPIN [10] and UPC-SPIN [4] can detect all possible deadlock conditions arising from all combination of parameters in all possible control-flows. However, such tools cannot scale beyond a few threads due to the combinatorial state-space explosion. Tools employing dynamic formal verification methods do not check all the control flows and hence can be used for larger programs. Such tools ISP [13], MODIST [16] and POE [12] generally employ centralized deadlock detection schemes which limit them to verifying executions using a small number of processes. Execution time of such methods is also usually high. DAMPI [15] is a dynamic formal verification tool which overcomes this limitation by using a distributed heuristics-based deadlock detection algorithm.

The most practical method for detecting deadlocks in terms of scalability is run-time analysis. Tools using this kind of analysis detect only those deadlocks that would actually occur during the current execution of a program. Marmot [7] and MPI-CHECK [8] employ synchronized time-out based strategies to detect deadlock conditions. Time-out based strategies may report false-positive error cases and generally cannot pinpoint the exact reason for the error. On the other hand, the run-time analysis tool, Umpire [14] uses a centralized WFG based on the generalized  $AND \oplus OR$  model developed by Hilbrich et al. [5]. However, MPI-CHECK, Marmot and Umpire are all based on the client-server model, which limits their scalability to a few hundred threads. In order to overcome this limitation, MUST [6] utilizes a flexible and efficient communication system to transfer records related to error detection between different processes or threads.

Our algorithm uses a different approach to detect deadlocks involving collective operations. We exploit two properties of operations in UPC which make deadlock detection easier than in MPI. Firstly, communication between two processes is non-blocking and secondly, non-determinism of point-to-point communication operations in terms of any `_source` cannot occur in UPC.

## 3 Detecting deadlocks due to collective errors in collective operations

Terms used throughout the rest of this paper are:

1. *THREADS* is an integer variable that refers to the total number of threads with which the execution of the application was initiated.
2. A *UPC operation* is defined as any UPC statement or function listed in the UPC specification.
3. The *state* of a thread is defined as the name of the UPC operation that the thread has reached. In case the thread is executing an operation which is not a collective or lock-related UPC operation, the *state* is set to `unknown`. If the thread has completed execution, the *state* is set to `end_of_execution`.

4. A *single-valued* argument is an argument of a UPC collective operation which must be passed the same value on every thread.
5. The *signature* of a UPC operation on a thread consists of the name of the UPC operation and the values which are about to be passed to each of the single-valued arguments of the UPC collective operation on that thread.
6. For any thread  $k$ ,  $s_k$  is a shared data structure which stores the state of thread  $k$  in field  $s_k.op$ . In case state is the name of a UPC collective operation,  $s_k$  also stores the single-valued arguments of the operation on that thread.
7. To *compare* the signatures of UPC operations stored in  $s_i$  and  $s_j$  means to check whether all the fields in  $s_i$  and  $s_j$  are identical.
8. If all the fields in  $s_i$  and  $s_j$  are identical, the result of the comparison is a *match*, otherwise there is a *mismatch*.
9.  $C(n, k)$  denotes the  $n^{th}$  collective operation executed by thread  $k$ .

The UPC specification requires that the order of calls to UPC collective operations must be the same for all threads [11]. Additionally, each ‘*single-valued*’ argument of a collective operation must have the same value on all threads. Therefore deadlocks involving only collective UPC operations can be created if:

1. different threads are waiting at different collective operations,
2. values passed to single-valued arguments of collective functions do not match across all threads, and
3. some threads are waiting at a collective operation while at least one thread has finished execution.

An algorithm to check whether any of the above 3 cases is going to occur must compare the collective operation which each thread is going to execute next and its single-valued arguments with those on other threads. Our algorithm achieves this by viewing the threads as if they were arranged in a circular ring. The left and right neighbors of a thread  $i$  are thread  $(i - 1) \% THREADS$  and thread  $(i + 1) \% THREADS$  respectively. Each thread checks whether its right neighbor has reached the same collective operation as itself. Since this checking goes around the whole ring, if all the threads arrive at the same collective operation, then each thread will be verified by its left neighbor and there will be no mismatches of the collective operations. However, if any thread comes to a collective operation which is not the same as that on the other thread, its left neighbor can identify the discrepancy, and issue an error message. The correctness of this approach is proven in Section 3.1.

On reaching a collective UPC operation, a thread  $k$  first records the signature of the collective operation in  $s_k$ . Thread  $k$  sets  $s_k.op$  to **unknown** after exiting from a operation. Let  $a$  and  $b$  be the variables that store signatures of collective operations. The assign ( $\leftarrow$ ) and the compare ( $\neq$ ) operations for the signatures of collective operation stored in  $a$  and  $b$  are defined as follows:

1.  $b \leftarrow a$  means
  - (a) assign value of variable  $a.op$  to variable  $b.op$ , and
  - (b) if  $a.op \neq end\_of\_execution$ , copy values of single-valued arguments recorded in  $a$  to  $b$

2.  $b \neq a$  is true if

- (a)  $b.op \neq a.op$ , or
- (b) if  $a.op \neq end\_of\_execution$ , any of the single-valued arguments recorded in  $a$  is not identical to the corresponding argument recorded in  $b$ .

Let thread  $j$  be the right neighbor of thread  $i$ . During execution, thread  $i$  or thread  $j$  could reach their respective  $n^{th}$  collective operation first. If thread  $i$  reaches the operation first, then it cannot compare  $C(n, i)$  recorded in  $s_i$  with  $C(n, j)$ , since  $s_j$  does not contain the signature of the  $n^{th}$  collective operation encountered on thread  $j$ , i.e.  $C(n, j)$ . The comparison can be delayed until thread  $j$  reaches its  $n^{th}$  collective operation. In order to implement this, another shared variable  $ds_k$  is used on each thread  $k$  to store the desired signature. For faster access, both shared variables  $s_k$  and  $ds_k$  have *affinity*<sup>1</sup> to thread  $k$ . If thread  $i$  finds that thread  $j$  has not reached a collective operation ( $s_j.op$  is **unknown**), then it assigns  $s_i$  to  $ds_j$ . When thread  $j$  reaches a collective operation it first records the signature in  $s_j$  and then compares it with  $ds_j$ . If they do not match, then thread  $j$  issues an error message, otherwise it sets  $ds_j.op$  to **unknown** and continues.

If thread  $i$  reaches the collective operation after thread  $j$  ( $s_j.op$  is assigned a name of a collective UPC operation), then thread  $i$  compares  $s_j$  with  $s_i$ . If they match, then there is no error, so execution continues.

The UPC specification does not require collective operations to be synchronizing. This could result in one or more state variables on a thread being reassigned with the signature of the next collective operation that it encounters before the necessary checking is completed. To ensure that the signature of the  $n^{th}$  collective operation encountered on thread  $i$  i.e.  $C(n, i)$  is compared with the signature of the  $n^{th}$  collective operation encountered on thread  $j$ , i.e.  $C(n, j)$ , the algorithm must ensure that:

1. If thread  $i$  reaches the  $n^{th}$  collective operation before thread  $j$  and assigns  $ds_j$  the signature of  $C(n, i)$ , it does not reassign  $ds_j$  before thread  $j$  has compared  $ds_j$  with  $s_j$ , and
2. If thread  $j$  reaches the  $n^{th}$  collective operation before thread  $i$  and assigns  $s_j$  the signature of  $C(n, j)$ , it does not reassign  $s_j$  before either thread  $i$  has a chance to compare it with  $s_i$  or thread  $j$  has a chance to compare it with  $ds_j$ .

In order to achieve the behavior described above, two shared variables  $r_sj$  and  $r_dsj$  are used for every thread  $j$ . Variable  $r_sj$  is used to prevent thread  $j$  from reassigning  $s_j$  before the necessary comparisons described above are completed. Similarly, variable  $r_dsj$  is used to prevent thread  $i$  from reassigning  $ds_j$  before the necessary comparisons are completed. Both  $r_sj$  and  $r_dsj$  have affinity to thread  $j$ .

For thread  $j$ , shared data structures  $s_j$  and  $ds_j$  are accessed by thread  $i$  and thread  $j$ . To avoid race conditions, accesses to  $s_j$  and  $ds_j$  are guarded using lock  $L[j]$ .

Our deadlock algorithm is implemented via the following three functions:

- *check\_entry()* function which is called before each UPC operation to check whether executing the operation would cause a deadlock,
- *record\_exit()* function which is called after each UPC operation to record that the operation is complete and record any additional information if required, and

---

<sup>1</sup>In UPC, shared variables that are stored in the physical memory of a thread are said to have *affinity* to that thread.

- *check\_final()* function which is called before every **return** statement in the **main()** function and every **exit()** function to check for possible deadlock conditions due to the termination of this thread.

---

**Algorithm A1**


---

```

1 On thread  $i$ :
   // Initialization
2  $s_i.op \leftarrow ds_i.op \leftarrow unknown, r_{s_i} \leftarrow 1, r_{ds_j} \leftarrow 1$ 
   // Function definition of check_entry(f_sig):
3 if  $THREADS = 1$  then
4   Exit check. ;
5 end
6 Acquire  $L[i]$  ;
7  $s_i \leftarrow f\_sig$  ;
8  $r_{s_i} \leftarrow 0$  ;
9 if  $ds_i.op \neq unknown$  then
10  if  $ds_i \not\cong s_i$  then
11    Print error and call global exit function. ;
12  end
13   $r_{s_i} \leftarrow 1$  ;
14   $r_{ds_i} \leftarrow 1$  ;
15   $ds_i.op \leftarrow unknown$  ;
16 end
17 Release  $L[i]$  ;
18 Wait until  $r_{ds_j} = 1$  ;
19 Acquire  $L[j]$  ;
20 if  $s_j.op = unknown$  then
21   $ds_j \leftarrow s_j$  ;
22   $r_{ds_j} \leftarrow 0$  ;
23 else
24  if  $s_j \not\cong s_i$  then
25    Print error and call global exit function ;
26  end
27   $r_{s_j} \leftarrow 1$  ;
28 end
29 Release  $L[j]$ 
   // Function definition of check_exit():
30 Wait until  $r_{s_i} = 1$  ;
31 Acquire  $L[i]$  ;
32  $s_i.op \leftarrow unknown$  ;
33 Release  $L[i]$ 
   // Function definition of check_final():
34 Acquire  $L[i]$  ;
35 if  $ds_i.op \neq unknown$  then
36  Print error and call global exit function. ;
37 end
38  $s_i.op \leftarrow end\_of\_execution$  ;
39 Release  $L[i]$  ;

```

---

The pseudo-code of the distributed algorithm<sup>2</sup> on each thread  $i$  to check deadlocks caused by incorrect or missing calls to collective operations<sup>3</sup> is presented. Function *check\_entry()* receives as argument the signature of the collective operation that the thread has reached, namely  $f\_sig$ .

### 3.1 Proof of Correctness

Using the same relation between thread  $i$  and thread  $j$ , i.e. thread  $i$  is the left neighbor of thread  $j$ , the proof of correctness is structured as follows. Firstly, it is proved that the algorithm is free of deadlocks and livelocks. Then Lemma 3.1 is used to prove that the left neighbor of any thread  $j$  does not reassign  $ds_j$  before thread  $j$  can compare  $s_j$  with  $ds_j$ . Lemma 3.2 proves that the right neighbor of any thread  $i$ , does not reassign  $s_j$  before thread  $i$  can compare  $s_i$  with  $s_j$ . Using Lemma 3.1 and Lemma 3.2 it is proven that for any two neighboring threads  $i$  and  $j$ , signature of  $C(n, j)$  is compared to the signature of  $C(n, i)$ . Finally, using Lemma 3.3 the correctness of the algorithm is proven by showing that : 1) no error message is issued if all the threads have reached the same collective operation with the same signature and 2) an error message is issued if at least one thread has reached a collective operation with a signature

<sup>2</sup>As presented, the algorithm forces synchronization even for non-synchronizing UPC collective operations. However, if forced synchronization is a concern, this can be handled with a queue of states. This will not change the O(1) behavior.

<sup>3</sup>UPC-CHECK treats non-synchronizing collective operations as synchronizing operations because the UPC 1.2 specification says that "Some implementations may include unspecified synchronization between threads within collective operations" (footnote; page 9).

different from the signature of the collective operation on any other thread. Case 1 is proved by Theorem 3.4 and Case 2 is proved by Theorem 3.5.

There is no hold-and-wait condition in algorithm A1, hence there cannot be any deadlocks in the algorithm. To show that the algorithm is livelock-free, we show that any given thread must eventually exit the waits on line 18 and 30. For any thread  $i$  reaching its  $n^{\text{th}}$  collective operation  $C(n, i)$ , thread  $i$  can wait at line 18 if thread  $i$  itself had set  $r\_ds_j$  to 0 on line 22 on reaching  $C(n-1, i)$ . This is possible only if thread  $i$  found that  $s_j.op = \text{unknown}$  on line 20, i.e. thread  $j$  is not executing an UPC collective operation. Eventually thread  $j$  either reaches the end of execution or a UPC collective operation. In the former case, a deadlock condition is detected, an error message is issued and the application exits. In the second case, thread  $j$  finds conditional statement on line 9 to be true and sets  $r\_ds_j$  to 1 on line 14. Since only thread  $i$  can set  $r\_ds_j$  to 0 again, thread  $i$  would definitely exit the wait on line 18. Similarly, for thread  $j$  to be waiting at line 30 after executing  $C(n, j)$ , it must not have set  $r\_s_j$  to 1 at line 13. This means that  $ds_j.op$  must be equal to  $\text{unknown}$  at line 9, implying that thread  $i$  has still not executed line 21 and hence line 20 (by temporal ordering) due to the atomic nature of operations accorded by  $L[j]$ . When thread  $i$  finally acquires  $L[j]$ , the conditional statement on line 20 must evaluate to false. If thread  $i$  has reached a collective operation with a signature different from that of  $C(n, j)$ , a deadlock error message is issued, otherwise  $r\_s_j$  is set to 1. Since only thread  $j$  can set  $r\_s_j$  to 0 again, it must exit the waiting at line 30.

**Lemma 3.1.** *After thread  $i$  assigns the signature of  $C(n, i)$  to  $ds_j$ , then thread  $i$  does not reassign  $ds_j$  before thread  $j$  compares  $s_j$  with  $ds_j$ .*

*Proof.* This situation arises only if thread  $i$  has reached a collective operation first. After thread  $i$  sets  $ds_j$  to  $s_i$  (which is already set to  $C(n, i)$ ) at line 21, it sets  $r\_ds_j$  to 0 at line 22. Thread  $i$  cannot reassign  $ds_j$  until  $r\_ds_j$  is set to 1. Only thread  $j$  can set  $r\_ds_j$  to 1 at line 14 after comparing  $s_j$  with  $ds_j$ .  $\square$

**Lemma 3.2.** *After thread  $j$  assigns the signature of  $C(n, j)$  to  $s_j$ , then thread  $j$  does not reassign  $s_j$  before it is compared with  $s_i$ .*

*Proof.* After thread  $j$  assigns the signature of  $C(n, j)$  to  $s_j$  at line 7, it sets  $r\_s_j$  to 0. Thread  $j$  cannot modify  $s_j$  until  $r\_s_j$  is set to 1. If thread  $i$  has already reached the collective operation, then thread  $j$  sets  $r\_s_j$  to 1 at line 13 only after comparing  $s_j$  with  $ds_j$  at line 10. However, thread  $i$  must have copied the value of  $s_i$  to  $ds_j$  at line 21. Alternatively, thread  $j$  might have reached the collective operation first. In this case, thread  $i$  sets  $r\_s_j$  to 1 at line 27 after comparing  $s_i$  to  $s_j$  at line 24.  $\square$

**Lemma 3.3.** *For any neighboring threads  $i$  and  $j$ , the signature of  $C(n, i)$  is always compared with the signature of  $C(n, j)$ .*

*Proof.* This is proved using induction on the number of the collective operations encountered on threads  $i$  and  $j$ .

*Basis.* Consider the case where  $n$  equals 1, i.e. the first collective operation encountered on thread  $i$  and thread  $j$ . The signature of  $C(1, i)$  is compared with the signature of  $C(1, j)$ . If thread  $i$  reaches collective operation  $C(1, i)$  first, then it assigns  $ds_j$  the signature of  $C(1, i)$ . Using Lemma 3.1, thread  $i$  cannot reassign  $ds_j$  until  $ds_j$  is compared with  $s_j$  by thread  $j$  on reaching its first collective operation,  $C(1, j)$ . Alternatively, if thread  $j$  reaches its collective operation first, then Lemma 3.2 states that after thread  $j$  assigns the signature of  $C(1, j)$  to  $s_j$ , thread  $j$  cannot reassign  $s_j$  before it is compared with  $s_i$ . The comparison between  $s_j$  and  $s_i$  is

done by thread  $i$  after it reaches its first collective operation and has assigned  $s_i$  the signature of  $C(1, i)$ .

*Inductive step.* If the signature of  $C(n, i)$  is compared with the signature of  $C(n, j)$ , then it can be proven that the signature of  $C(n+1, i)$  is compared with the signature of  $C(n+1, j)$ . If thread  $i$  reaches its next collective operation  $C(n+1, i)$  first, then it assigns  $ds_j$  the signature of  $C(n+1, i)$ . Using Lemma 3.1, thread  $i$  cannot reassign  $ds_j$  until  $ds_j$  is compared with  $s_j$  by thread  $j$  on reaching its next collective operation, i.e.  $C(n+1, j)$ . Alternatively, if thread  $j$  reaches its next collective operation first, then Lemma 3.2 states that after thread  $j$  assigns  $C(n+1, j)$  to  $s_j$ , thread  $j$  cannot reassign  $s_j$  before it is compared with  $s_i$ . The comparison of  $s_j$  with  $s_i$  is done by thread  $i$  after it reaches its next collective operation and has assigned  $s_i$  the signature of  $C(n+1, i)$ .  $\square$

Using Lemma 3.3, it is proven that for any neighboring thread pair  $i$  and  $j$ , the signature of  $n^{\text{th}}$  collective operation of thread  $i$  is compared with the signature of  $n^{\text{th}}$  collective operation of thread  $j$ . As  $j$  varies from 0 to  $THREADS - 1$ , it can be said that when the  $n^{\text{th}}$  collective operation is encountered on any thread, it is checked against the  $n^{\text{th}}$  encountered collective operation on every other thread before proceeding. Thus in the following proofs, we need to only concentrate on a single (potentially different) collective operation on each thread. In the following proofs, let the signature of the collective operation encountered on a thread  $k$  be denoted by  $S_k$ . If a state or desired state  $a_i.op$  is **unknown**, then it is denoted as  $a = U$  for succinctness. Then in algorithm A1, after assigning the signature of the encountered collective operation, i.e. line  $s_i \leftarrow f\_sig$ , notice that for thread  $i$ :

$s_i$  must be  $S_i$ ,  
 $ds_i$  must be either  $U$  or  $S_{i-1}$ ,  
 $s_j$  must be either  $U$  or  $S_j$ , and  
 $ds_j$  must be  $U$ .

**Theorem 3.4.** *If all the threads arrive at the same collective operation, and the collective operation has the same signature on all threads, then Algorithm A1 will not issue an error message.*

*Proof.* If  $THREADS$  is 1, no error message is issued, so we need to consider only cases of execution when  $THREADS > 1$ . If all threads arrive at the same collective operation with the same signature, then during the checks after  $s_i \leftarrow f\_sig$ , is the same for all  $i$ . Let  $S$  denote this common signature. We will prove this theorem by contradiction. An error message is printed only if:

1.  $ds_i \neq U$  and  $ds_i \neq s_i \Rightarrow ds_i = S$  and  $ds_i \neq S \Rightarrow S \neq S$  (contradiction) or
2.  $s_j \neq U$  and  $s_j \neq s_i \Rightarrow s_j = S$  and  $s_j \neq S \Rightarrow S \neq S$  (contradiction)

So Theorem 3.4 is proved.  $\square$

**Theorem 3.5.** *If any thread has reached a collective operation with a signature different from the signature of the collective operation on any other thread, then a deadlock error message is issued.*

*Proof.* There can be a mismatch in the collective operation or its signature only if there is more than one thread.

Since the signatures of the collective operations reached on every thread are not identical, there must be some thread  $i$  for which  $S_i \not\cong S_j$ . For these threads  $i$  and  $j$ , the following procedures are made to be atomic and mutually exclusive through use of lock  $L[j]$ :

- Action 1: Thread  $i$  checks  $s_j$ . If  $s_j = U$ , then thread  $i$  executes  $ds_j \leftarrow s_i$ , else, computes  $s_j \not\cong s_i$  and issues an error message if true.
- Action 2: Thread  $j$  assigns the signature of the collective operation it has reached to  $s_j$ . Thread  $j$  checks  $ds_j$ . If  $ds_j \neq U$ , the thread  $j$  computes  $ds_j \not\cong s_j$  and issues message if true.

There are only two possible cases of execution: either action 1 is followed by action 2 or vice versa.

In the first case, in action 1, thread  $i$  finds  $s_j = U$  is true, executes  $ds_j \leftarrow S_i$  and continues. Then in action 2, thread  $j$  executes  $s_j \leftarrow S_j$ , finds that  $ds_j \neq U$  and hence computes  $ds_j \not\cong s_j$ . Now, since  $ds_j = S_i$  and  $s_j = S_j$  and  $S_i \neq S_j$  (by assumption) implies that  $ds_j \not\cong s_j$  is true. Therefore thread  $j$  issues an error message.

In the second case, in action 2, thread  $j$  assigns  $s_j \leftarrow S_j$ , finds  $ds_j = U$  and continues. Before thread  $i$  initiates action 1 by acquiring  $L[j]$ , it must have executed  $s_i \leftarrow S_i$ . If  $ds_i \neq U$  and  $ds_i \not\cong s_i$ , then an error message is issued by thread  $i$ , otherwise it initiates action 1. Thread  $i$  finds  $s_j \neq U$  and computes  $s_j \not\cong s_i$ . Now, since  $s_i = S_i$  and  $s_j = S_j$  and  $S_i \neq S_j$  (by assumption) implies that  $s_j \not\cong s_i$  is true. Therefore thread  $i$  issues an error message.

Since the above two cases are exhaustive, an error is always issued if  $S_i \neq S_j$  and hence Theorem 3.5 is proved. □

**Theorem 3.6.** *The complexity of the Algorithm A1 is  $O(1)$ .*

*Proof.* There are two parts to this proof.

1. The execution-time overhead for any thread  $i$  is  $O(1)$ . Any thread  $i$  computes a fixed number of instructions before entering and after exiting a collective operation. It waits for at most two locks  $L[i]$  and  $L[j]$  each of which can have a dependency chain containing only one thread, namely thread  $i - 1$  and thread  $j$  respectively. Thread  $i$  synchronizes with only two threads, i.e. its left neighbor thread  $i - 1$  and right neighbor thread  $j$ . There is no access to variables or locks from any other thread. Therefore the execution time complexity of the algorithm in terms of the number of threads is  $O(1)$ .
2. The memory overhead of any thread  $i$  is independent of the number of threads and is constant. □

### 3.2 Detecting deadlock errors involving `upc_notify` and `upc_wait` operations

The compound statement  $\{upc\_notify; upc\_wait\}$  forms a split barrier in UPC. The UPC specification requires that firstly, there should be a strictly alternating sequence of `upc_notify` and `upc_wait` calls, starting with a `upc_notify` call and ending with a `upc_wait` call. Secondly, there can be no collective operation between a `upc_notify` and its corresponding `upc_wait` call. These conditions are checked using a private binary flag on each thread which is set when a `upc_notify` statement is encountered and reset when a `upc_wait` statement is encountered. This binary flag is initially reset. If any collective operation other than `upc_wait` is encountered when the flag is set, then there must be an error. Similarly, if a `upc_wait` statement is encountered when the flag is reset, then there must be an error. Finally, if the execution ends, while the flag is set,

then there must be an error. These checks are performed along with the above algorithm and do not require any communication between threads. Also modifying and checking private flags is an operation with complexity of  $O(1)$ .

If all the threads issue the `upc_notify` statement, then the next UPC collective operation issued on all the threads must be a `upc_wait` statement. Therefore algorithm *A1* working in unison with the above check needs to only verify the correct ordering of `upc_notify` across all threads. The correct ordering of the `upc_wait` statements across all threads is automatically guaranteed with the above mentioned checks. This is reflected in Algorithm *A2*.

## 4 Detecting deadlocks created by hold-and-wait dependency chains for acquiring locks

In UPC, acquiring a lock with a call to the `upc_lock()` function is a blocking operation. In UPC program, deadlocks involving locks occur when there exists one of the following conditions:

1. a cycle of hold-and-wait dependencies with at least two threads, or
2. a chain of hold-and-wait dependencies ending in a lock held by a thread which has completed execution, or
3. a chain of hold-and-wait dependencies ending in a lock held by a thread which is blocked at a synchronizing collective UPC operation.

Our algorithm uses a simple edge-chasing method to detect deadlocks involving locks in UPC programs. Before a thread  $u$  tries to acquire a lock, it checks if the lock is free or not. If it is free, the thread continues execution. Otherwise, if the lock is held by thread  $v$ , thread  $u$  checks  $s_v.op$  to check if thread  $v$ :

1. is not executing a collective UPC operation or `upc_lock` operation ( $s_v.op$  is *unknown*),  
or
2. is waiting to acquire a lock, or
3. has completed execution, or
4. is waiting at a synchronizing collective UPC operation.

If thread  $v$  is waiting to acquire a lock, then thread  $u$  continues to check the *state* of the next thread in the chain of dependencies. If thread  $u$  finally reaches thread  $m$  which is not executing a collective UPC operation or `upc_lock` operation, then no deadlock is detected. If thread  $u$  finds itself along the chain dependencies, then it reports a deadlock condition. Similarly, if thread  $u$  finds thread  $w$  which has completed execution at the end of the chain of dependencies, then it issues an error message.

When the chain of dependencies ends with a thread waiting at a collective synchronizing operation, the deadlock detection algorithm needs to identify whether the thread will finish executing the collective operation or not. Figure 1 illustrates these two cases. Thread  $u$  is trying to acquire a lock in a chain of dependencies ending with thread  $w$ . When thread  $u$  checks the  $s_w.op$  of thread  $w$ , thread  $w$  may (a) not have returned from the  $n^{th}$  synchronizing collective operation  $C_s(n, w)$ , (b) have returned from the  $n^{th}$  synchronizing collective operation but has not updated the  $s_w.op$  in the `check_exit()` function, (c) have completed executing `check_entry()` function for the next synchronizing collective operation  $C_s(n + 1, w)$ , or (d)

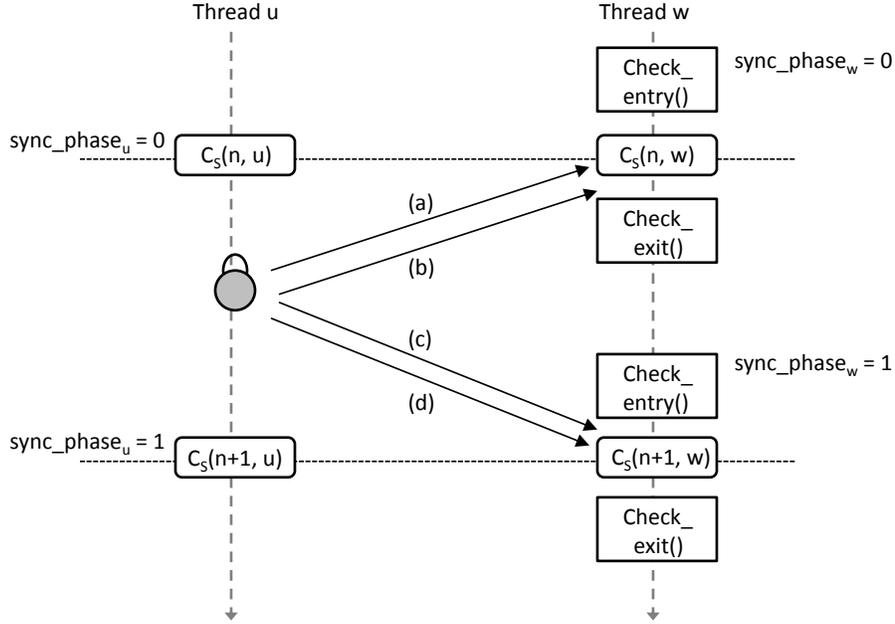


Figure 1: Possible scenarios when detecting deadlocks involving chain of hold-and wait dependencies. Scenario (a) or (b) is not a deadlock condition, while scenario (c) or (d) is.

waiting at the  $(n + 1)^{th}$  synchronizing collective operation  $C_s(n + 1, w)$ . The  $n^{th}$  synchronizing collective operation encountered on thread  $w$  must be a valid synchronization operation that all threads must have called (otherwise the  $check\_entry()$  function would have issued an error message). Therefore scenarios (a) and (b) are not deadlock conditions, while (c) and (d) are. To identify and differentiate between these scenarios, a binary shared variable  $sync\_phase_k$  is introduced for each thread  $k$ . Initially  $sync\_phase_k$  is set to 0 for all threads. At the beginning of each  $check\_entry()$  function on thread  $k$ , the value  $sync\_phase_k$  is toggled. Thread  $u$  can now identify the scenarios by just comparing  $sync\_phase_u$  and  $sync\_phase_w$ . If they match (are *in-phase*), then it is either scenario (a) or (b) and hence no deadlock error message is issued. If they do not match (are *out-of-phase*), then it is either scenario (c) or (d) and hence a deadlock error message is issued.

#### 4.1 The complete deadlock detection algorithm

The complete algorithm to detect deadlocks created by errors in collective operations and hold-and-wait dependency chains for acquiring locks is presented below. The  $check\_entry()$  and  $check\_exit()$  functions receive two arguments: 1) the signature of the UPC operation that the thread has reached, namely  $f\_sig$  and 2) the pointer  $L\_ptr$ .  $L\_ptr$  points to the lock which the thread is trying to acquire or release if the thread has reached a  $upc\_lock$ ,  $upc\_lock\_attempt$  or  $upc\_unlock$  statement.

Checking for dependency chains and cycles adds only a constant amount of time overhead for each thread in the chain or cycle. This means that the overhead is  $O(T)$  where  $T$  is the number of threads in the dependency chain.

## Algorithm A2

---

```

1 On thread  $i$ :
   // Initialization
2 Create empty list of acquired and requested locks ;
3  $s_i.op \leftarrow ds_i.op \leftarrow unknown$ ;  $r_{-s_i} \leftarrow 1$ ;  $r_{-ds_j} \leftarrow 1$ ;
   ( $sync\_phase_i \leftarrow 0$ )
   // Function definition of  $check\_entry(f.sig, L\_ptr)$ :
4 Acquire  $L[i]$  ;
5  $s_i \leftarrow f.sig$  ;
6 Release  $L[i]$  ;
7 if  $f.sig.op = at\_upc\_wait\_statement$  then
8   Exit check ;
9 end
10 if  $f.sig.op = at\_upc\_lock\_operation$  then
11   Acquire  $c.L$  ;
12   Check status of  $L\_ptr$  ;
13   if  $L\_ptr$  is held by this thread or is part of a cycle
      or chain of dependencies then
14     Print suitable error and call global exit ;
15   else
16     Update list of requested locks ;
17     Release  $c.L$  ;
18     Exit check ;
19   end
20 end
21 if  $f.sig.op = at\_upc\_unlock\_operation$  then
22   if  $L\_ptr$  is not held by this thread then
23     Print suitable error and call global exit ;
24   else
25     Update list of acquired locks ;
26     Exit check
27   end
28 end
   // Thread has reached a collective operation
29 if  $THREADS = 1$  then
30   Exit check ;
31 end
32 Acquire  $c.L$  ;
33 if this thread holds locks which are in the list of
      requested locks then
34   Print suitable error and call global exit ;
35 end
36 Release  $c.L$  ;
37 Acquire  $L[i]$  ;
38  $r_{-s_i} \leftarrow 0$  ;
39 if this is a synchronizing collective operation then
40    $sync\_phase_i \leftarrow (sync\_phase_i + 1)\%2$  ;
41 end
42 if  $ds_i.op \neq unknown$  then
43   if  $ds_i \neq s_i$  then
44     Print error and call global exit function ;
45   end
46    $r_{-s_i} \leftarrow 1$  ;
47    $r_{-ds_i} \leftarrow 1$  ;
48    $ds_i.op \leftarrow unknown$  ;
49 end
50 Wait until  $r_{-ds_j} = 1$  ;
51 Acquire lock  $L[j]$  ;
52 if  $s_j.op = unknown$  then
53    $ds_j \leftarrow s_i$  ;
54    $r_{-ds_j} \leftarrow 0$  ;
55 else
56   if  $s_j \neq s_i$  then
57     Print error and call global exit function ;
58   end
59    $r_{-s_j} \leftarrow 1$  ;
60 end
61 Release lock  $L[j]$ 
   // Function definition of  $check\_exit(f.sig, L\_ptr)$ :
62 Wait until  $r_{-s_i} = 1$  ;
63 Acquire  $L[i]$  ;
64  $s_i \leftarrow unknown$  ;
65 Release  $L[i]$  ;
66 if  $f.sig.op = at\_upc\_lock\_operation$  then
67   Acquire  $c.L$  ;
68   Remove  $L\_ptr$  from the list of requested locks ;
69   Add  $L\_ptr$  to the list of acquired locks ;
70   Release  $c.L$  ;
71 end
72 if  $f.sig.op = at\_upc\_lock\_attempt\_operation$  then
73   if  $L\_ptr$  was achieved then
74     Acquire  $c.L$  ;
75     Remove  $L\_ptr$  from the list of requested locks
76     ;
77     Add  $L\_ptr$  to the list of acquired locks ;
78     Release  $c.L$  ;
79   end
80 Continue execution ;
   // Function definition of  $check\_final()$ :
81 Acquire  $L[i]$  ;
82  $s_i \leftarrow end\_of\_execution$  ;
83 if  $ds_i.op \neq unknown$  then
84   Print error and call global exit function ;
85 end
86 Release  $L[i]$  ;
87 Acquire  $c.L$  ;
88 if this thread holds locks which are in the list of
      requested locks then
89   Print suitable error and call global exit ;
90 end
91 if this thread is still holding locks then
92   Print suitable warning ;
93 end
94 Release  $c.L$  ;

```

---

## 5 Experimental verification of scalability

This deadlock detection algorithm has been implemented in the UPC-CHECK tool [3]. UPC-CHECK was used to experimentally verify the scalability of this algorithm on a Cray XE6 machine running the CLE 4.1 operating system. Each node has two 16-core Interlagos processors. Since we are interested in the verification of scalability, the authors measured the overhead of our deadlock detection method for 2, 4, 8, 16, 32, 64, 128, 256, 512, 1024, 2048, 4096 and 8192 threads. The verification of scalability was carried out by first measuring the overhead incurred when calling a UPC collective operation and then measuring the overhead when running the CG and IS UPC NAS Parallel Benchmarks (NPB) [1]. The Cray C 8.0.4 compiler was used with the `-hupc` option. To pin processes and memory the `aprun` command was used with the following options:

```
-ss -cc cpu.
```

The authors first measured the overhead of checking for deadlocks involving the `upc_all_broadcast` operation with a message consisting of one 4 byte integer. Since deadlock checking is independent of the message size, the small message size was used so that the checking overhead could be easily measured. To measure the time accurately, 10,000 calls to `upc_all_broadcast` were timed and an average reported.

Overhead times ranged from 76 to 123 microseconds for multiple nodes, i.e. 64, 128, 256, 512, 1024, 2048, 4096 and 8192 threads. When replacing `upc_all_broadcast` with `upc_all_gather_all`, overhead times ranged from 73 to 119 microseconds. In both cases, a slight increase is observed as we increase the number of threads. The authors attribute this to the fact that, in general, not all pairs of UPC threads can be mapped to physical processors for which the communication between UPC threads  $i$  and  $(i + 1) \% THREADS$

is the same for all  $i$ . The maximal communication time for optimally placed UPC threads still grows slowly as the total number of UPC threads grows. The deviation from constant time in the above experiment is only a factor of 1.5 for 128 times as many UPC threads.

```
time (t1);
for (i = 0; i < 10000; i++)
{
    upc_all_broadcast;
}
time {t2};
bcast_time = (t2 - t1)/10000;
```

Number of threads	Class B			Class C		
	Without checks	With checks	Overhead	Without checks	With checks	Overhead
2	77.2	77.6	0.4	211.2	211.8	0.6
4	41.4	41.7	0.3	112.7	112.8	0.1
8	28.1	28.7	0.6	73.9	74.2	0.3
16	15.3	16.0	0.6	39.4	40.0	0.6
32	8.6	9.5	0.9	21.1	22.1	0.9
64	5.5	6.6	1.1	13.1	14.0	1.0
128	3.3	4.7	1.3	8.3	9.7	1.4
256	NA	NA	NA	5.6	7.2	1.6

Table 1: Time in seconds of the UPC NPB-CG benchmark with and without deadlock checking

Number of threads	Class B			Class C		
	Without checks	With checks	Overhead	Without checks	With checks	Overhead
2	4.56	4.59	0.03	20.00	20.11	0.11
4	2.18	2.18	0.00	9.50	9.52	0.01
8	1.34	1.34	0.00	5.28	5.28	0.00
16	0.79	0.79	0.00	3.46	3.46	0.00
32	0.42	0.43	0.01	1.89	1.89	0.00
64	0.29	0.30	0.01	1.30	1.31	0.01
128	0.21	0.22	0.01	0.82	0.82	0.00
256	0.26	0.27	0.01	0.57	0.57	0.00

Table 2: Time in seconds of the UPC NPB-IS benchmark with and without deadlock checking

UPC-CHECK was tested for correctness using 150 tests from the UPC RTED test suite [2]. Each test contains a single deadlock. For all the tests, UPC-CHECK detects the error, prevents the deadlock from happening and exits after reporting the error correctly [3]. Since these tests are very small, the observed overhead was so small that we could not measure them accurately.

Timing results for the UPC NPB CG and IS benchmarks are presented in Tables 1 and 2 using 2, 4, 8, 16, 32, 64, 128, and 256 threads. Timings using more than 256 threads could not be obtained since these benchmarks are written in a way that prevents them from being run with more than 256 threads. These results also demonstrate the scalability of the deadlock detection algorithm presented in this paper. Timing data for the class B CG benchmark using 256 threads could not be obtained since the problem size is too small to be run with 256 threads.

## 6 Conclusion

In this paper, a new distributed and scalable deadlock detection algorithm for UPC collective operations is presented. The algorithm has been proven to be correct and to have a run-time complexity of  $O(1)$ . This algorithm has been extended to detect deadlocks involving locks with a run-time complexity of  $O(T)$ ,  $T$  is the number of threads involved in the deadlock. The algorithm has been implemented in the run-time error detection tool UPC-CHECK and tested with over 150 functionality test cases. The scalability of this deadlock detection algorithm has been experimentally verified using up to 8192 threads.

In UPC-CHECK, the algorithm is implemented through automatic instrumentation of the application via a source-to-source translator created using the ROSE toolkit [9]. Alternatively, such error detection capability may be added during the precompilation step of a UPC compiler. This capability could be enabled using a compiler option and may be used during the entire debugging process as the observed memory and execution time overhead even for a large number of threads is quite low.

## Acknowledgment

This work was supported by the United States Department of Defense & used resources of the Extreme Scale Systems Center at Oak Ridge National Laboratory.

## References

- [1] UPC NAS Parallel Benchmarks.
- [2] James Coyle, James Hoekstra, Marina Kraeva, Glenn R. Luecke, Elizabeth Kleiman, Varun Srinivas, Alok Tripathi, Olga Weiss, Andre Wehe, Ying Xu, and Melissa Yahya. UPC run-time error detection test suite. 2008.
- [3] James Coyle, Indranil Roy, Marina Kraeva, and Glenn Luecke. UPC-CHECK: a scalable tool for detecting run-time errors in Unified Parallel C. *Computer Science - Research and Development*, pages 1–7. 10.1007/s00450-012-0214-4.
- [4] Ali Ebnenasir. UPC-SPIN: A Framework for the Model Checking of UPC Programs. In *Proceedings of Fifth Conference on Partitioned Global Address Space Programming Models*, PGAS '11, 2011.
- [5] Tobias Hilbrich, Bronis R. de Supinski, Martin Schulz, and Matthias S. Müller. A graph based approach for MPI deadlock detection. In *Proceedings of the 23rd international conference on Supercomputing*, ICS '09, pages 296–305, New York, NY, USA, 2009. ACM.
- [6] Tobias Hilbrich, Martin Schulz, Bronis R. Supinski, and Matthias S. Müller. MUST: A scalable approach to runtime error detection in mpi programs. In Matthias S. Müller, Michael M. Resch, Alexander Schulz, and Wolfgang E. Nagel, editors, *Tools for High Performance Computing 2009*, pages 53–66. Springer Berlin Heidelberg, 2010. 10.1007/978-3-642-11261-4\_5.
- [7] Bettina Krammer, Matthias Müller, and Michael Resch. MPI application development using the analysis tool MARMOT. In Marian Bubak, Geert van Albada, Peter Sloot, and Jack Dongarra, editors, *Computational Science - ICCS 2004*, volume 3038 of *Lecture Notes in Computer Science*, pages 464–471. Springer Berlin / Heidelberg, 2004. 10.1007/978-3-540-24688-6\_61.
- [8] Glenn R. Luecke, Yan Zou, James Coyle, Jim Hoekstra, and Marina Kraeva. Deadlock detection in MPI programs. *Concurrency and Computation: Practice and Experience*, 14(11):911–932, 2002.
- [9] Daniel J. Quinlan and et al. ROSE compiler project.
- [10] Stephen Siegel. Verifying parallel programs with MPI-Spin. In Franck Cappello, Thomas Herault, and Jack Dongarra, editors, *Recent Advances in Parallel Virtual Machine and Message Passing Interface*, volume 4757 of *Lecture Notes in Computer Science*, pages 13–14. Springer Berlin / Heidelberg, 2007. 10.1007/978-3-540-75416-9\_8.
- [11] The UPC Consortium. UPC Language Specifications (v1.2). 2005.
- [12] Sarvani Vakkalanka, Ganesh Gopalakrishnan, and Robert M. Kirby. Dynamic verification of MPI programs with reductions in presence of split operations and relaxed orderings. In *Proceedings of the 20th international conference on Computer Aided Verification*, CAV '08, pages 66–79, Berlin, Heidelberg, 2008. Springer-Verlag.
- [13] Sarvani S. Vakkalanka, Subodh Sharma, Ganesh Gopalakrishnan, and Robert M. Kirby. ISP: a tool for model checking mpi programs. In *Proceedings of the 13th ACM SIGPLAN Symposium on Principles and practice of parallel programming*, PPOPP '08, pages 285–286, New York, NY, USA, 2008. ACM.
- [14] Jeffrey S. Vetter and Bronis R. de Supinski. Dynamic software testing of MPI applications with Umpire. In *Proceedings of the 2000 ACM/IEEE conference on Supercomputing (CDROM)*, Supercomputing '00, Washington, DC, USA, 2000. IEEE Computer Society.
- [15] A. Vo, S. Ananthakrishnan, G. Gopalakrishnan, B.R. de Supinski, M. Schulz, and G. Bronevetsky. A scalable and distributed dynamic formal verifier for MPI programs. In *High Performance Computing, Networking, Storage and Analysis (SC), 2010 International Conference for*, pages 1–10, nov. 2010.
- [16] Junfeng Yang, Tisheng Chen, Ming Wu, Zhilei Xu, Xuezheng Liu, Haoxiang Lin, Mao Yang, Fan Long, Lintao Zhang, and Lidong Zhou. MODIST: transparent model checking of unmodified distributed systems. In *Proceedings of the 6th USENIX symposium on Networked systems design and implementation*, NSDI'09, pages 213–228, Berkeley, CA, USA, 2009. USENIX Association.