

Scalable Algorithms for Global Snapshots in Distributed Systems

Rahul Garg,
IBM India Research Lab,
New Delhi, India
grahul@in.ibm.com

Vijay K. Garg*
ECE Department
The University of Texas at Austin
Austin, TX 78712-1084, USA
garg@ece.utexas.edu

Yogish Sabharwal,
IBM India Research Lab,
New Delhi, India
ysabharwal@in.ibm.com

Abstract

Existing algorithms for global snapshots in distributed systems are not scalable when the underlying topology is complete. In a network with N processors, these algorithms require $O(N)$ space and $O(N)$ messages per processor. As a result, these algorithms are not efficient in large systems when the logical topology of the communication layer such as MPI is complete. In this paper, we propose three algorithms for global snapshot: a grid-based, a tree-based and a centralized algorithm. The grid-based algorithm uses $O(N)$ space but only $O(\sqrt{N})$ messages per processor. The tree-based algorithm requires only $O(1)$ space and $O(\log N \log w)$ messages per processor where w is the average number of messages in transit per processor. The centralized algorithm requires only $O(1)$ space and $O(\log w)$ messages per processor. We also show a matching lower bound for this problem. Our algorithms have applications in checkpointing, detecting stable predicates and implementing synchronizers. We have implemented our algorithms on top of the MPI library on the BlueGene/L supercomputer. Our experiments confirm that the proposed algorithms significantly reduce the message and space complexity of a global snapshot.

Keywords: Checkpointing, Global Snapshot Algorithms, Fault-tolerance, Stable Predicates, BlueGene/L

Corresponding Author: Vijay K. Garg, garg@ece.utexas.edu

*This work was performed when the author was visiting IBM India Research Lab

1 Introduction

Computing the global snapshot of a system is a fundamental problem in distributed computing. It has applications in fault-tolerance of long-running programs by providing an intermediate *checkpoint* of the system. In case of a failure, the system can restart from the checkpoint instead of the beginning of the program.

Global snapshots are also useful in monitoring stable properties of the system. A property is stable if, once it becomes true, it stays true. Some examples of stable properties are termination, deadlock, loss-of-a-token etc. By repeatedly computing the global snapshot and evaluating the property on the computed snapshot, one can detect any stable property.

The definition and the first algorithm to compute a consistent global snapshot for a system with FIFO channels was given by Chandy and Lamport in [CL85]. For a colorful description of Chandy and Lamport's algorithm see [Dij85]. This elegant algorithm has the property that it does not freeze the underlying computation during global snapshot computation. Thus, the underlying application is not stopped from sending or receiving any messages when the snapshot algorithm is in progress. We will restrict ourselves to algorithms with this property.

Spezialetti and Kearns have given efficient algorithms to disseminate a global snapshot to processes initiating the snapshot computation [SK86]. Bouge [Bou87] has given an efficient algorithm for repeated computation of snapshots for synchronous computations. In the absence of the FIFO assumption, as shown by Taylor [Tay89], any algorithm for a snapshot is either inhibitory (that is, it may delay actions of the underlying application) or requires piggybacking of control information on basic messages. Lai and Yang [LY87] and Mattern [Mat93] have given snapshot algorithms that require only the piggybacking of control information.

A different protocol has been implemented more recently by Schulz, Bronevetsky, Fernandes, Marques, Pingali and Stodghill in [SBF⁺04]. We will refer to this algorithm as SBF+ algorithm. The reader is referred to a [KRS95] for a survey of global snapshot algorithms.

All the existing global snapshot algorithms for non-FIFO channels [Mat93, SBF⁺04], require at least one message and one integer to be stored for every channel in the system. In a system with N processors and a completely connected topology, this translates to $O(N)$ messages and $O(N)$ space per processor. In massively parallel computers, this overhead can be quite significant. For example, the IBM Blue Gene/L computer has $64K$ processors. Furthermore, at the application level a process at any processor may send message to any other processor. Thus, the topology at the application level is that of a completely connected graph. Using existing algorithms on such a system would result in $64K$ messages per processor per snapshot.

In this paper, we propose three algorithms with different characteristics that make them much more scalable than existing algorithms. We call these algorithms grid-based, tree-based and centralized algorithms.

The grid-based algorithm assumes a logical grid-like structure of the system. It requires $O(\sqrt{N})$ messages per processor with each message of size $O(\sqrt{N})$ integers. For Blue Gene/L, this algorithm would use 256 messages per node instead of 64K messages. Although each message of our algorithm is bigger than existing algorithms, the total overhead of communication and computation is significantly reduced by using fewer messages. The space complexity of the grid-based algorithm is similar to existing algorithms. It uses $O(N)$ space at every processor.

Our tree-based algorithm reduces the space complexity of the global snapshot problem. It uses a constant number of variables at each processor and $O(\log N \log w)$ messages per processor where w is the average number of in-transit messages when the snapshot is taken. A crucial difference between the earlier algorithms and our tree-based algorithm is that earlier algorithms relied on informing each process the number of in-transit messages. The tree-based algorithm avoids this step to reduce the space complexity at each process. The tree-based algorithm uses $O(N \log N \log w)$ messages in all.

The total message complexity of the system can be further reduced to $O(N \log w)$ messages by using a centralized algorithm. The disadvantage of the centralized algorithm is that it may require a single node to process as many as $O(N \log w)$ messages.

We also show a matching lower bound on the message complexity of any global snapshot algorithm that is based on detecting when all in-transit messages have been received. We show that detecting whether W messages have been received in a system with N processes require at least W control messages when W is at most N and $\Omega(N \log W/N)$ when W is greater than N .

The characteristics of the algorithms proposed in this paper are summarized in Figure 1. The algorithm CLM refers to Chandy and Lamport's algorithm with Mattern's modification for non-FIFO channels. It requires $O(N^2)$ messages in all when the underlying topology is completely connected. The marker message contains a single integer denoting the number of white messages sent on that channel. We will assume that a single integer is sufficient to count the total number of messages in the system. Thus, message size denotes the total number of integers sent in the message. The space requirement of this algorithm is $O(N)$ as a process has to store the number of white messages it has sent to any other process in the system. The grid-based algorithm reduces the number of messages from $O(N^2)$ to $O(N^{3/2})$ but each message is of size $O(\sqrt{N})$. The tree-based algorithm reduces the space complexity to $O(1)$. Its message complexity depends on the average number of in-transit messages. The centralized algorithm reduces the message complexity; however, it requires the coordinator node to process $O(N \log w)$ messages.

This paper is organized as follows. In Section 2, we briefly describe the problem and the existing

Algorithm	Message Complexity	Message Size	Space
CLM	$O(N^2)$	$O(1)$	$O(N)$
Grid-based	$O(N^{3/2})$	$O(\sqrt{N})$	$O(N)$
Tree-based	$O(N \log N \log w)$	$O(1)$	$O(1)$
Centralized	$O(N \log w)$	$O(1)$	$O(1)$

Notation: N : Number of processes, w : Average Number of in-transit messages/process

Figure 1: Summary of Global Snapshot Algorithms

work. Section 3, 4 and 5 discuss the grid-based algorithm, the tree-based algorithm and the centralized algorithm respectively. In Section 6, we describe the lower bound for this problem. Section 7 describes our implementation of the global snapshot algorithm on the BlueGene/L computer and performance analysis of the algorithms proposed in the paper. Section 7 describes other applications of techniques described in the paper.

2 Model and Background

We model a distributed system as an asynchronous message-passing system without any shared memory or a global clock. A *distributed program* consists of a set of N processes denoted by $\{P_1, P_2, \dots, P_N\}$ and a set of unidirectional channels. A channel connects two processes. Thus the topology of a distributed system can be viewed as a directed graph in which vertices represent the processes and the edges represent the channels. A channel is assumed to have infinite buffer and to be error-free. We do not make any assumptions on the ordering of messages. Any message sent on the channel may experience arbitrary but finite delay. The state of the channel at any point is defined to be the sequence of messages sent along that channel but not received.

A *computation* is defined as a tuple (E, \rightarrow) where E is the set of all events that are generated during the execution, and \rightarrow is the happened-before relation [Lam78] on E . We define a *consistent cut*, or a global snapshot, as any subset $F \subseteq E$ such that

$$f \in F \wedge e \rightarrow f \Rightarrow e \in F.$$

Chandy and Lamport's algorithm, which works only for FIFO channels, uses the concept of *color* of messages and processes. We associate with each process a variable called *color* that is either white or red. Intuitively, the computed global snapshot corresponds to the state of the system just before the processes turn red. All processes are initially white. After recording the local state, a process turns red. Thus the state of a local process is simply the state just before it turned red. One or more processes initiate the

snapshot algorithm by recording their local state and turning red. Once a process turns red, it is required to send a special message called *marker* along all its outgoing channels before it sends out any message. A process is required to turn red on receiving a marker if it has not already done so. Since channels are FIFO, the above mentioned rule guarantees that no white process ever receives a message sent by a red process. This property ensures that the global snapshot recorded is consistent.

It is easy to define the state of the channels for the snapshot. These are the messages sent by a white process received by a red process. These messages cross the global snapshot in the forward direction and form the state of the channel in the global snapshot because they are in transit when the snapshot is taken. To record the state of the channel from P_i to P_j , the process P_j starts recording all messages it receives from P_i after turning red. Since P_i sends a marker to P_j on turning red, the arrival of the marker at P_j from P_i indicates that there will not be any further white messages from P_i sent to P_j . It can, therefore, stop recording messages once it has received the marker. When a process has received a marker on all its incoming channels, it has finished recording all incoming channels and its component of the snapshot algorithm is finished.

The entire algorithm takes exactly one message (marker) along each link. Although this algorithm is efficient for sparse networks it is not scalable for large dense networks. The message complexity of this algorithm is $O(N^2)$.

We now briefly describe an extension to this algorithm due to Mattern[Mat93] that removes the assumption that channels are FIFO. When channels do not satisfy FIFO property, the marker message cannot be used to distinguish between white and red messages. Therefore, Mattern's algorithm includes the color in all the outgoing messages for any process besides sending the marker. Further, even after P_i gets a red message from P_j or the marker, it cannot be sure that it will not receive a white message on that channel. A white message may arrive later than a red message due to the overtaking of messages. To solve this problem Mattern's algorithm includes in the marker the total number of white messages sent by that process along that channel. The receiver keeps track of the total number of white messages received and knows that all white messages have been received when this count equals the count included in the marker. The message complexity of this algorithm is also $O(N^2)$. The SBF+ algorithm [SBF⁺04] also uses a message count per channel that indicates the number of white messages sent on the channel.

Before we discuss our algorithms it is important to understand the source of quadratic message complexity in existing algorithms. In Chandy and Lamport's algorithm (and that of Mattern's), a marker message is sent along every channel. These marker messages serve three purposes. We discuss how we achieve these three functions without sending an explicit message for the channel.

First, a marker message indicates to a process that one or more other processes have initiated the global

snapshot algorithm. Thus, marker messages serve to disseminate the knowledge that a global snapshot is required. This task corresponds to broadcast of some information in a network. We would like the entire network to know that the global snapshot algorithm has been initiated. Assuming that there is a pre-defined spanning tree in the network, it is a simple matter to broadcast any information in $O(N)$ messages. Any white node that wants to initiate the snapshot algorithm, sends “initiate” messages to all its neighbors that are part of the spanning tree. On receiving such a message a white process turns red and sends “initiate” message to all its neighbors in the tree except for the node that sent initiate message to it. Observe that any information about initiation of a global checkpoint algorithm is sent over only along the links that are part of the spanning tree. Since there are $N - 1$ links in the tree, there are at most $2(N - 1)$ messages of type (one message in each direction). In all three algorithms proposed in the paper, we will assume such an initial step that turns all processes red.

Second, in a FIFO channel marker messages serve to distinguish white messages from the red messages. Any message received before the marker is a white message and any message received after the marker is a red message. This ability to distinguish between messages is useful in recording in-transit messages. This technique is applicable only for FIFO channels and for Non-FIFO channels one has to piggyback a bit with each message indicating whether it is a white or a red message. This technique is employed in Mattern’s as well as SBF+ algorithm. We also use the same technique in all our algorithms.

Finally, we come to the third function of the marker. A marker also serves as the indicator of the end of white messages from a given process. When process P receives a marker on the channel from process Q , it knows that all the messages sent from Q to P before Q took its local checkpoint have been received. Since this property is true only for FIFO channels, Mattern and SBF+ algorithm achieve this functionality by explicitly sending the total number of white messages sent. By counting the number of white messages received a process can determine if it has received all the white messages sent on that channel. When a process has determined that this has become true, it *closes* that channel. When all the incoming channels are closed, the process has finished its component of the global snapshot algorithm and can report to the initiator. Note that explicitly sending along each channel the number of white messages sent contributes not only to the overhead in terms of messages but also in terms of state maintained at each process. Every process is forced to maintain $O(N)$ integers recording the number of white messages sent along each channel. It can be observed that even when process P_i has not communicated with P_j , it still has to send out a message to P_j with 0 as the message count. All our algorithms either avoid or optimize on this step of a global snapshot algorithm.

3 Grid-based Algorithm

The grid-based algorithm reduces the message complexity of sending the white message count by making use of the following two observations. First, a process does not need a separate count of the white messages sent to it by a specific process. All it needs is the total number of white messages that have been sent to it. Whereas Mattern’s algorithm and SBF+ algorithm communicate to every process the number of white messages sent, our algorithm communicates the total number of white messages sent to a process. The second observation is that a process can use one message to send out information about the number of messages it has sent for multiple processes. We use this observation to reduce the number of messages at the expense of increasing the size of messages.

Our algorithm uses the notion of color of processes similar to Chandy and Lamport’s algorithm. A process may be white, red, or black. It is white if it has not recorded its local state. It is red if it has recorded its local state but not the state of incoming channels, and black if it has recorded its local state and all the transit messages. Initially all processes are white.

There are three main components in our algorithm. The first component assumes that there is a pre-defined spanning tree in the system. It uses the spanning tree to broadcast the fact that one or more processes want to take the global snapshot. It ensures that all processes in the system turn red with $O(N)$ messages by using the spanning tree. Furthermore, by including the color of the sending process in all messages, the algorithm also guarantee that no white process ever acts on a red message. On receiving a red message a white process immediately turns red. This component is fairly standard, but for completeness sake it is shown in Figure 2.

```
initiate() // enabled only if (color=white)
  take local checkpoint;
  color = red;
  send "init" to all processes connected by tree edges;

On receiving "init" message or a red message on edge e
  if (color = white)
    take local checkpoint;
    color = red;
    send "init" to all tree edges except e
```

Figure 2: Common Initiation Component

The second component is required for processes to determine the total number of white messages it is supposed to receive. This component assumes a 2D grid structure to compute the total number of white

messages sent to a specific process. We will use $P_{i,j}$ to denote the process at the co-ordinate (i, j) of the grid. For simplicity, we will assume that the grid is a perfect square. The algorithm can be applied to grid with any number of rows and columns. Each process maintains a matrix $whiteSent[i, j]$ which records the number of white messages sent to process $P_{i,j}$. Further, any process $P_{r,c}$ will also compute $rowCount[j]$ which is the total number of white messages sent by processes in row r to the process $P_{c,j}$. Note that the vector $rowCount$ has size \sqrt{N} and is maintained by each process. Further, all diagonal processes in the grid $P_{c,c}$ maintain an additional vector $totalCount$ such that $totalCount[j]$ equals the total number of white messages sent by all processes to the process $P_{c,j}$.

This component is shown in Figure 3. In step 1, each process $P_{r,c}$ sends i^{th} row of $whiteSent$ to $P_{r,i}$. Note that this message has $O(\sqrt{N})$ integers and $O(\sqrt{N})$ such messages are sent by each process. In step 2, $P_{r,c}$ receives c^{th} row from all processes in row r . It maintains a cumulative sum of all the vectors received. When it has received values from all the processes in its row, $rowCount[j]$ is equal to the number of white messages sent to $P_{c,j}$ by processes in row r . It then sends $rowCount[.]$ to $P_{c,c}$. In step 3, a process participates only if it is a diagonal processor in the grid. Nodes that correspond to $P_{i,i}$ are responsible for computing the $totalCount$ for all processes in row i . Once $P_{i,i}$ has calculated $totalCount$ for each process in its row, it informs that process with a message of size $O(1)$.

The message complexity of this component is $O(\sqrt{N})$ messages sent/received per node, each of size $O(\sqrt{N})$. This component results in every process knowing $whiteCount$, the total number of white messages that have been sent to the process.

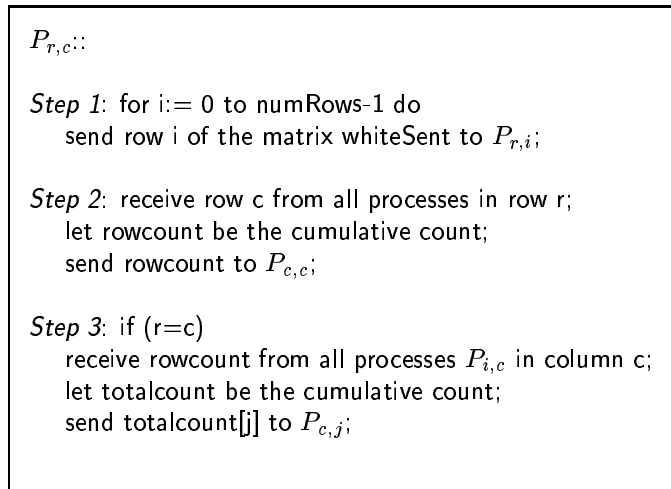


Figure 3: The Grid-based Algorithm to Compute whiteCount

The third component is responsible for turning all processes black and detecting when that has happened. A process keeps track of all the white messages it has received in the counter $whiteReceived$. When this

count equals $whiteCount$, the process knows that it has recorded all possible in-transit messages and turns black. When the entire tree has turned black, the snapshot algorithm is complete. Detecting that the entire tree is black can be performed with $O(N)$ messages by using standard convergecast on the spanning tree[Gar04].

Based on the preceding discussion, we have the following Theorem.

Theorem 1 *There exists a global snapshot algorithm that requires $O(\sqrt{N})$ messages per node where each message contains $O(\sqrt{N})$ integers.*

Remark: In the above discussion, we assumed that the grid was $\sqrt{N} \times \sqrt{N}$. The algorithm can be generalized to any grid of size $numR \times numC$. When the number of rows is less than the number of columns, the algorithm does not need any change. When the number of rows is greater than the number of columns, then a process $P_{r,c}$ sends $whiteSent[i, \cdot]$ to $P_{r, i \bmod numcols}$ instead of sending it to $P_{r,i}$. The algorithm uses $O(numR)$ messages each of size $O(numC)$. Note that when the grid is N by 1, the algorithm reduces to a completely distributed algorithm with $O(N^2)$ messages of size $O(1)$. The other extreme case is when the grid is 1 by N . This case reduces to the centralized algorithm with $O(N)$ messages each of size $O(N)$.

4 Tree-based Algorithm

The grid-based algorithm requires every process to maintain the number of white messages sent to any process. This information requires $O(N)$ integers to be maintained at every process. This overhead may be prohibitive for high performance computing applications. We now show an algorithm that reduces this overhead to $O(1)$ integers.

In the tree-based algorithm each process does not maintain individual counts for the number of white messages sent. Instead, it maintains the *deficit* which denotes the total number of white messages sent minus the total number of white messages received. Note that the deficit for a process may be negative. Clearly, the total number of in-transit messages that correspond to the global checkpoint equals the sum of all deficits when the processes turn from white to red. Our algorithm is based on computing this number and detecting when all these messages have been delivered.

The tree-based algorithm consists of three components. The first component for initiation is same as the grid-based algorithm. The second component corresponds to computation of the total deficit in the network. This can be accomplished easily using a convergecast in $O(N)$ messages. Let W be the total deficit computed using second component. The third and final component corresponds to detecting when

all in-transit messages have been received. We call this problem the *distributed message counting problem*. The *distributed message counting problem* can be defined as follows. There are W in-transit messages that are destined for N processes. The problem is to detect when all W messages have been received. We are interested in algorithms that are efficient for even large values of W .

We abstract the total deficit as tokens that are distributed in a network. Each token represents a pending message for which the destination is unknown. Whenever a message is received a token is consumed. The goal is to detect when the total number of tokens reduces to zero. A simple algorithm in which a coordinator maintains all the tokens can solve this problem in $O(W)$ messages. Whenever a process receives a white message it simply informs the coordinator who consumes a token.

We now show an algorithm that uses $O(N \log N \log(W/N))$ messages. When W is much larger than N , this algorithm will outperform the coordinator based algorithm. The algorithm works based on rounds. There are at most $\lceil \log W/N \rceil$ rounds. Initially, we divide the total number of tokens W equally among all processes. The maximum number of tokens any process has in the first round is $w = \lceil \log W/N \rceil$. Let w_k be the maximum number of tokens any process has at round k . Our algorithm ensures that $w_{k+1} \leq w_k/2$. Thus the maximum number of tokens owned by a process goes down by a factor of two in every round.

We now describe the algorithm at round k . We use three colors — green, yellow and orange — to label all the processes. A process is orange if it has no tokens and it has received one or more in-transit messages. These in-transit messages have not consumed tokens and such processes would initiate messages to search for tokens. A process that does not have this problem is either yellow or green. A process p_i with r_i tokens is considered green if it has strictly greater than $w_k/2$ tokens and yellow otherwise. Intuitively, green processes are rich, yellow are debt-free, and orange are in debt (and therefore poor).

We organize the processes in a perfect binary tree. The degree and the height of the tree is used only in analyzing the message complexity. The algorithm is correct for any spanning tree.

Our algorithm ensures the following properties: (I1) A yellow process cannot have a green child. (I2) The root is green, and (I3) Any orange node eventually becomes yellow.

It is sufficient to describe the protocol when an in-transit message arrives. This event is detected as arrival of a white message at a red process. For every white message, a token must be consumed. No action is required at the destination to maintain the invariants if consuming a token does not change the color of the process. After all, the invariants are specified only using the color of the process. We now consider two possible transitions that can happen when the color gets changed. The pseudo-code for the algorithm is shown in Figure 4.

First, consider the case when the color changes from green to yellow. This transition can violate invariant (I1) if any of its children is green. To maintain the invariant, the process sends a “swap” message with its

<p>On turning from green to yellow if any child green send (“swap”, tokens) to that child; else if root node reset for the next round</p> <p>On turning from yellow to orange send (“split”, tokens) to the nearest green ancestor;</p>
--

Figure 4: The Tree-based Algorithm for Distributed Message Counting

tokens to its left child (if any). If the left child is green, it replies with an “accept” message and performs the swap operation. Otherwise, it replies with a “reject” message. When a process receives a “reject” message, it tries the swap operation with the other child. If none of the children is green then the process knows that it does not violate (I1). If this node is root, then it also knows that the entire tree does not have any green node. In this case, it initiates a global “reset” operation that takes the algorithm to the next round.

Note that when a child performs a swap operation, it turns yellow and may violate (I1). So after the swap operation, the child, in turn, may initiate its own swap message. The total number of “swap” “accept” and “reject” messages is $O(\log N)$ because they traverse one path down the tree.

Now consider the case when a node changes from yellow to orange. The process sends a “split” message to its parent. Our algorithm will allow only a green node to split. If the parent is green, it splits its tokens with the requesting process. Otherwise, this message is forwarded to its parent. The message is guaranteed to find a green node due to (I2). When a green node splits its token, it is guaranteed to become yellow. This can now fire up the rule for turning from green to yellow.

The “reset” operation is performed as follows. The root requests all nodes to send their tokens to the root. Once the root has received messages from all processes, it calculates the total number of tokens (in-transit messages) and recalculates w_k for the next round.

By using the above algorithm, and repeatedly halving w_k , we are guaranteed to reach the case when w_k is 1. At that point, we continue with one more round with a process to be green if it has the token, yellow if it has no token and orange if it has no token but a pending in-transit message. When the root detects that the the entire tree has turned yellow, it can signal the end of the global snapshot algorithm.

We now show the following claim.

Theorem 2 *The algorithm uses $O(N \log N \log(W/N))$ messages.*

Proof: Initially no node has more than W/N tokens. Let w_k be the maximum number of tokens any process has at round k . We show that $w_{k+1} \leq w_k/2$. The new round begins only when the root turns yellow and

both its children are yellow. This implies that no node in the tree has more than $w_k/2$ tokens. Thus, the root node can start the next round with $w_{k+1} = w_k/2$. This argument shows that there are at most $\log(W/N)$ rounds.

In each round, there can be two types of transitions — from green to yellow and from yellow to orange. The yellow to orange transition results in splitting operations. There can be at most N splits because each split turns a green node into yellow. Thus, there can be at most N splits. Each split takes at most $\log N$ messages before it succeeds in finding a green node. Every split is followed by a transition of a node from green to yellow. This transition also takes at most $\log N$ messages.

The number of transitions from green to yellow is also bounded by N . Thus, each round requires $O(N \log N)$ messages.

■

5 Centralized Algorithm

The centralized algorithm is similar to the tree-based algorithm described in the previous section. The difference lies in the algorithm for solving the distributed message counting problem. Thus, the centralized algorithm also consists of three components. The first component for initiation is the same as the grid-based and tree-based algorithms. The second component for computation of the total deficit in the network is same as the tree-based algorithm. The third component corresponds to the *distributed message counting problem* that detects when all the in-transit messages have been received.

We now give an algorithm that uses $O(N \log(W/N))$ messages. In the next section we will show that any algorithm that solves the *distributed message counting problem* will use at least $\Omega(N \log(W/N))$ messages, implying that we cannot hope to do better in terms of message complexity.

The centralized algorithm for solving the *distributed message counting problem* is similar to the one used in the tree-based algorithm. The difference lies in the logical organization of processes. Unlike the tree-based algorithm in which the processes were logically organized as a perfect binary tree, in the centralized algorithm, the green processes are organized as a list with a fixed tail. The advantage of organizing the processes in such a manner (as we will show) is that we can avoid the factor of $\log n$ message exchanges that occur in the tree-based algorithm when the colour of a process changes, and in fact we can complete the processing required on change of the colour of a process with only a constant number of message exchanges. The disadvantage, however, is that all the processes must interact with the process representing the tail of the list, leading to centralization of message exchanges.

```

If not tail process
  On turning from green to yellow
    send ("swap", tokens) to tail;
  On turning from yellow to orange
    send ("split", tokens) to tail;
  On receiving "swap" request
    retain  $w_k/2$  tokens,
    send accept message with remaining tokens to requestor;
  On receiving "split" request
    send accept message with half the tokens to requestor;

else if tail process
  On turning from green to yellow
    If  $head \neq tail$ 
      send ("swap", tokens) to  $head$ ;
       $head = head - 1$ ;
    else
      reset for the next round;
  On receiving "swap" or "split" request
    If  $head \neq tail$ 
      forward request to  $head$ ;
       $head = head - 1$ ;
    else
      reset for the next round;

```

Figure 5: The Centralized Algorithm for Distributed Message Counting

The centralized algorithm also works in rounds, as does the tree-based algorithm. Again, there are at most $\lceil \log W/N \rceil$ rounds. Initially, we divide the total number of tokens W equally among all processes. The maximum number of tokens any process has in the first round is $w = \lceil \log W/N \rceil$. Let w_k be the maximum number of tokens any process has at round k . The centralized algorithm also ensures that $w_{k+1} \leq w_k/2$. Thus the maximum number of tokens owned by a process goes down by a factor of two in every round.

We now describe the algorithm at round k . We again use three colours — green, yellow and orange — to label all the processes. these colours have the same significance as in the tree-based algorithm. A process is orange if it has no tokens and it has received one or more in-transit messages. A process is considered green if it has strictly greater than $w_k/2$ tokens and yellow otherwise.

The processes are organized as follows. All the green nodes are organized as a list. There is a fixed process that represents the tail of the list. This node always remains green. Let the processes have ranks $0, 1, \dots, N-1$. Let the tail be the process with rank 0. The tail maintains the list using only one variable, *head*. The variable *head* stores the first green node in the list. When *head* equals h , then logically, the list of green nodes is $\{h, h-1, \dots, 0\}$. The algorithm ensures that whenever there is need to remove a green node from the list, only *head* is removed from the list. The list can then be quickly updated simply by decrementing *head*. At the start of every round, *head* is initialized to $N-1$.

We now describe the protocol when an in-transit (white) message arrives. For every white message, a token must be consumed. No action is required if this does not change the colour of the process. We now consider two possible transitions that can happen when the colour gets changed. the pseudo-code for the algorithm is shown in figure 5.

First, consider the case when the color changes from green to yellow. Since nodes are only removed from the head of the green list, this process borrows all except $w_k/2$ tokens from *head*, so that it remains green and *head* turns yellow instead. In order to accomplish this, it sends a “swap” request to the tail. The tail forwards this request to the *head* and removes *head* from the green list, since the *head* is guaranteed to become yellow. The *head*, on receiving the message sends all except $w_k/2$ tokens to the requesting process. If the requestor is the same as *head*, then the tail sends back a reject to the *head*, and removes it from the green list. The *head* on receiving the reject, turns yellow. Note that if the *head* is same as tail, that is there is no other green node, then the tail knows that all the processes have turned yellow. In this case, it initiates a global “reset” operation that takes the algorithm to the next round.

Now consider the case when a node changes from yellow to orange. Since nodes are only removed from the head of the green list, this process borrows half the tokens from *head*. In order to accomplish this, the process sends a “split” message to the tail. The tail forwards this request to the *head* and removes *head* from the green list, since the *head* is guaranteed to become yellow. The *head*, on receiving the message splits its

tokens with the requesting process. In case there is no other green node other than the tail, the tail initiates a global “reset” operation that takes the algorithm to the next round.

The “reset” operation is performed as in the case of the tree-based algorithm. The root requests all nodes to send their tokens to the root. Once the root has received messages from all processes, it calculates the total number of tokens (in-transit messages) and recalculates w_k for the next round.

We now show the following claim.

Theorem 3 *The algorithm uses $O(N \log(W/N))$ messages.*

Proof: Initially no node has more than W/N tokens. Let w_k be the maximum number of tokens any process has at round k . We show that $w_{k+1} \leq w_k/2$. The new round begins only when all the nodes turn yellow. This implies that no node in the tree has more than $w_k/2$ tokens. Thus, the root node can start the next round with $w_{k+1} = w_k/2$. This argument shows that there are at most $\log(W/N)$ rounds.

In each round, there can be two types of transitions — from green to yellow and from yellow to orange. The green to yellow transition results in swapping operations and the yellow to orange transition results in splitting operations. There can be at most N splits/swaps because each split/swap turns a green node into yellow. Each split/swap takes $O(1)$ messages in finding a green node. Thus, each round requires $O(N)$ messages.

■

6 Lower Bound

In this section we show that any algorithm must exchange at least $O(N \log W/N)$ point-to-point control messages to detect termination if $W > N$ and W control messages otherwise.

Lemma 1 *Let W be the number of messages remaining to be delivered. Then there exists a processor p_W that generates a control message on receiving at most $\lceil W/N \rceil$ messages. Therefore any algorithm for termination detection can be forced to exchange a control message on delivery of at most $\lceil W/N \rceil$ messages (by sending these messages to p_W).*

Proof: Suppose there is no such processor. Then the adversary could send $W/N + 1$ messages to $W \bmod N$ processors and W/N messages to the remaining $N - (W \bmod N)$ processors. The algorithm would not exchange any control message and therefore would not be able to detect termination. The adversary can therefore force a control message to be generated by sending $\lceil W/N \rceil$ messages to p_W .

■

Let W be the remaining messages to be delivered. Now, consider the algorithm of the adversary described in figure 6. We define one *round* to be one iteration of the loop of the adversary.

```

 $W :=$  Sum of deficits at all processes;
while (  $W > 0$  ) do
    deliver  $\lceil W/N \rceil$  messages to processor  $p_W$  (as determined in lemma 1) ;
     $W := W - \lceil W/N \rceil$  ;

```

Figure 6: Algorithm of the adversary

Lemma 2 *Any termination detection algorithm must exchange at least r control messages where r is the number of rounds executed by the adversary.*

Proof: Using lemma 1, one control message is forced to be generated by the termination detection algorithm on execution of one round of the adversary’s algorithm. The statement follows. ■

Theorem 4 *The number of control messages exchanged by any algorithm to detect termination is*

$$\begin{aligned} \Omega(N \log(W/N)) & \text{ if } W \geq N \text{ and} \\ \Omega(W) & \text{ if } W < N \end{aligned}$$

Proof: We will show that the adversary executes W rounds when $W < N$ and $\Omega(N \log(W/N))$ rounds when $W \geq N$.

Consider the case when $W < N$. Only one message is delivered in every round until W becomes 0. Therefore the adversary can execute W rounds.

Now consider the case when $W \geq N$. Let $k = \lfloor \log W/N \rfloor$, i.e., $2^k N \leq W < 2^{k+1} N$. We claim that the adversary executes at least $(k + 1)N/4$ rounds. The proof is by induction.

- **Base Case ($k = 0$):** $N \leq W < 2N$

Two messages are delivered in every round until W becomes $< N$. Thereafter one message is delivered in every round until W becomes 0. Therefore at least $N > N/4$ rounds are executed by the adversary.

- **Induction Case :** $2^k N \leq W < 2^{k+1} N$

The number of messages delivered in each round is $\lceil W/n \rceil$ until $W < 2^k N$. Thereafter, the number of messages delivered in each round decreases. Therefore the number of messages delivered in any round

is no more than 2^{k+1} . Consider the first $N/4$ rounds. The number of messages delivered in these $N/4$ rounds is no more than $2^{k+1}N/4$. Therefore the number of messages remaining to be delivered after $N/4$ rounds is at least $W - 2^{k+1}N/4 \geq 2^{k-1}N$. If the number of messages remaining to be delivered after these $N/4$ rounds is $\geq 2^k N$, then the adversary keeps executing rounds and delivering messages until $2^{k-1}N \leq W < 2^k N$. Now, using the induction hypothesis, the adversary executes at least $kN/4 + N/4 = (k+1)N/4$ rounds before termination. Substituting the value of k , the adversary executes at least $N(1 + \log W/N)/4$ rounds.

The theorem follows from the discussions above and lemma 2. ■

Note that the distributed message counting problem is a special case of the termination detection problem where one tries to detect when all processes are passive and all channels are empty [DS80]. If one considers only those messages that are sent before checkpoint and defines a process to be passive after checkpoint, our problem is reduced to that of termination detection. It is well known that termination detection may require at least as many messages as the application messages in the worst case[CM85]. The main difference between the general termination detection problem and distributed message counting problem is that once a process turns passive it can never turn active (even when it receives a message). This difference explains the reduced message complexity.

7 Experimental Results

We implemented the three checkpointing algorithms (grid, tree and centralized) for BlueGene/L using C. The BlueGene/L is a massively parallel supercomputer that scales upto 65,536 dual-processor nodes. Each of the two cores is an embedded PPC440 core, designed to reach a nominal clock frequency of 700 MHz (1.4 giga-operations per second). The BlueGene/L uses five interconnect networks for I/O, debug, and various types of interprocessor communication. The most significant of these interconnection networks is the three-dimensional torus that has the highest aggregate bandwidth and handles the bulk of all communication. The signaling rate for the nearest neighbor links is 1.4 Gbps in each direction. Each node supports six independent, bidirectional nearest neighbor links, with an aggregate bandwidth of 2.1 GBps. The hardware latency to transit a node is approximately 100 ns. The torus network uses both dynamic (adaptive) and deterministic routing with virtual buffering and cut-through capability. Adaptive routing allows packets to follow any minimal path to the final destination, allowing packets to dynamically “choose” less congested routes.

7.1 Setup

Each processor has a unique logical rank. These ranks vary from 0 to $N - 1$.

For the grid based algorithm, the grid is constructed by logically organizing the processors into m rows and n columns, where $n = m$ if the number of nodes is a perfect square and $n = 2m$ otherwise. For instance the 32 nodes system is organized as a 4×8 grid whereas the 64 nodes system is organized as an 8×8 grid. The processors are logically organized in the grid by row major ordering of their ranks, i.e., a processor with rank i is in row $\lfloor i/n \rfloor$ and column $i \bmod n$ in the logical grid.

The convergecast tree is logically organized over the nodes as follows. The root has rank 0, its left and right children are nodes with ranks 1 and 2 respectively, and so on. Thus, the nodes are organized in a binary tree filled in order of ranks from top to bottom and left to right. Thus for a node with rank i , its left child has rank $2 \cdot (i + 1) - 1$, its right child has rank $2 \cdot (i + 1)$ and its parent has rank $\lfloor (i + 1)/2 \rfloor - 1$ (provided these are valid ranks). For the tree based algorithm, the convergecast tree is used as the checkpointing algorithm tree as well.

The algorithm for the application code that is used to analyze the performance of these algorithms is described in Figure 7.

```
ApplicationAlgorithm( W, M )
for  $i = 1$  to  $W$  /*  $W$  Loop */
    send data message to random destination ;

int  $fin = N - 1$  ; // number of finished messages remaining to be received
int  $i = 0$  ; // number of messages sent to other processes

while ( (  $fin > 0$  ) OR (  $i < M$  ) )
    if (  $i < M$  )
        send data message to random destination ;
         $i++$  ;
        if (  $i == M$  )
            send finish message to all other processors ;

    if (  $fin > 0$  )
        rcv message ;
        if ( received finish message )
             $fin--$  ;
endWhile
```

Figure 7: Application algorithm for analyzing performance of the checkpointing algorithms

The algorithm takes as parameters W and M . All processors first perform W sends without performing any receives. This is followed by a loop in which on every iteration, one send and one receive is performed. This is done till M sends have been performed. As the messages are sent to random destinations, the number

of messages to be received is not known. Therefore, a finish logic is included in which every processor on finishing M sends, sends a finish message to every other destination. A processor can therefore proceed to cleanup when it has received $N - 1$ finish messages, where N is the number of processors. After finishing all sends and receives, a processor waits for the checkpointing algorithm to complete (if it has not already completed).

The send and receive calls are blocking calls that in turn use the MPI library for inter-process communication. They first receive all checkpointing messages that are available. These messages are passed to the checkpointing algorithm that in turn processes them (generating and sending checkpointing messages to other processors, if necessary). After the checkpointing messages have been processed, the send call performs the corresponding MPI send routine (*MPI_Send*) and the receive call performs the corresponding non-blocking MPI receive routine (*MPI_Irecv*). The reason that the receive call is implemented using a non-blocking MPI call is to allow processing of checkpointing messages as and when they become available, even when no application messages are available.

Every checkpointing message has a fixed 32 byte header that contains some checkpointing algorithm specific information such as *message type*, *source*, *destination*, *round number* (for tree and centralized algorithms), etc. Therefore the minimum message size for a checkpointing message is 32 bytes. Over and above this fixed header, a checkpointing message may contain additional information, such as *tokens*, *counts of white messages sent*, etc.

We ran this application using all the three algorithm on BlueGene/L partitions of 32, 64, 128, 256 and 512 nodes. We analyzed the algorithms for the fixed values of 40000 for W , and 50000 for M .

The checkpoint initiation time was uniformly picked by every processor from a range T_1 to T_2 . In one set of experiments we set T_1 to 950 milliseconds and T_2 to 1050 milliseconds. The mean checkpoint initiation time was 1 second. In another set of experiments we set T_1 to 285 milliseconds and T_2 to 315 milliseconds. The mean checkpoint initiation time was 300 milliseconds.

We also performed another optimization on the tree and centralized algorithms. We modified these algorithms, so that on the initiation of a round, i.e., when initial deficit for a round is being computed, we receive all the white messages that are pending that have not been absorbed by the algorithm. For example, suppose that a node has 4 pending messages which it has received but which have not been absorbed by the algorithm because it is orange and is awaiting a response for its split request. Suppose that a reset round is initiated at this time. Then, the node absorbs these white messages by incrementing its received count by 4 and computes the deficit with this new value of received count. We collected data on 512 nodes using this optimization to analyze the improvements if any.

We recorded the following information for all the runs:

- *Maximum Processor Latency:* The processor latency is the time elapsed (in microseconds) on a processor from when checkpointing is initiated on that processor to when the algorithm completes checkpointing, i.e., it determines that all the white messages have been received. The maximum processor latency is the processor latency maximized over all the processors.
- *Total Latency:* The total latency is the time elapsed (in microseconds) from when the earliest processor determines that checkpointing has been initiated to when the last processor determines that checkpointing has been completed.
- *Message Size:* the minimum, the maximum and the average message size of a checkpointing message generated during the run.
- *Messages sent/received:* The minimum, the maximum and the average number of checkpointing messages sent/received by any processor.
- *Initial Deficit:* The initial deficit computed for the tree and centralized algorithms.
- *Number of Rounds:* The number of rounds before the tree/centralized algorithms detects completion.

7.2 Results

7.2.1 Latency

The total latencies (in microseconds) observed for the three checkpointing algorithms when the application was executed with $W = 40000$, $M = 50000$ and mean checkpoint timeout of 1 second, are shown in table 1. The maximum processor latency was observed to be only negligibly smaller than the total latency.

N	Initial Deficit	Grid			Tree	Centralized		
		Latency	Max Msgs	Time/Msg	Latency	Latency	Max Msgs	Time/Msg
32	2880992	104	23	4.52	6192	5053	1548	3.26
64	5764032	157	31	5.06	11232	9108	3109	2.93
128	11536256	190	48	3.96	19191	15765	5738	2.75
256	23105280	293	64	4.58	36921	34641	12290	2.82
512	46341632	572	96	5.96	50578	69574	24506	2.84

Table 1: Total latencies (microseconds) for the checkpointing algorithms on different number of nodes

With a mean checkpoint timeout of 1 second, we observed that all the white messages were sent before the checkpointing algorithm started. This is reflected by the initial deficit count for the tree and centralized algorithms. The initial deficit count is exactly $N \cdot (W + M + N - 1)$, accounting for the data and finish messages

generated from all the nodes. Therefore the latency observed reflects the time taken by the checkpointing algorithm from initiation to completion.

The maximum time during active processing is spent in MPI calls. For the grid algorithm, the latencies are very small compared to the tree and centralized algorithms. This is because the number of messages generated is small upto 512 nodes and most of these are processed (sent and received) in parallel on different nodes. There is little to compare between the latencies of the tree and centralized algorithms.

Since the maximum time in processing is spent in MPI calls, we calculated the maximum messages processed (sent + received) by any node for different values of N . For the centralized algorithm, the tail node is almost always busy sending/receiving messages to/from other nodes. Therefore, we expected the latency to be proportional to the number of messages being processed by the tail node. For this, we calculated the average time for processing a message by dividing the latency by the sum of the messages sent and received (assuming send and receive MPI calls take roughly the same time). Our hypothesis was confirmed by the consistency observed for time per message. For the grid algorithm, we expected similar results since the diagonal elements of the grid are almost always busy processing messages. However, the number of messages sent and received are too small to account for noise (initiation over convergecast tree, computations, etc.) to make any meaningful inferences. For the tree based algorithm, the message exchange is more dynamic in nature and there is no single node that is almost always busy making it difficult to analyze the latency numbers.

7.2.2 Message Size

The minimum, maximum and average message sizes observed for the three checkpointing algorithms when the application was executed with $W = 40000$, $M = 50000$ and mean checkpoint timeout of 1 second, are shown in table 2.

N	Grid				Tree			Centralized		
	Layout	Min	Max	Avg	Min	Max	Avg	Min	Max	Avg
32	4 × 8	32	64	53.70	32	36	35.97	32	36	35.95
64	8 × 8	32	64	58.08	32	36	35.98	32	36	35.95
128	8 × 16	32	96	82.70	32	36	35.98	32	36	35.95
256	16 × 16	32	96	89.07	32	36	35.98	32	36	35.95
512	16 × 32	32	160	144.97	32	36	35.98	32	36	35.95

Table 2: Message Sizes (in bytes) for the checkpointing algorithms on different number of nodes

As the checkpoint initiation message contains no additional information other than the message type, which is in the fixed 32 byte header, we observe a minimum message size of 32 bytes for all the algorithms.

For the grid based algorithms, the messages exchanged in steps 1 and 2 contain one integer (4 bytes) for

every column in the grid layout. Therefore the maximum message size is $32 + 4 \cdot (\text{number of columns})$. The average message size also increases with the increase in the number of columns. Most messages are of type 1 or 2. Type 3 messages only carry a single integer.

For the tree and centralized algorithms, message sizes are not dependent on the number of nodes in the system. Some messages carry an additional information element (such as requestor or tokens) that take up an additional 4 bytes. Therefore, the maximum message size observed is 36 bytes. The average message size for these algorithms is also close to the maximum message size as most messages carry an extra information element.

7.2.3 Message Counts

The minimum, maximum and average send message counts observed for the three checkpointing algorithms when the application was executed with $W = 40000$, $M = 50000$ and mean checkpoint timeout of 1 second, are shown in table 3. Similarly, the receive counts are shown in table 4.

N	Grid				Tree			Centralized		
	Layout	Min	Max	Avg	Min	Max	Avg	Min	Max	Avg
32	4 × 8	4	12	5.88	34	485	149.41	29	576	80.16
64	8 × 8	8	16	9.80	37	813	193.25	28	991	78.03
128	8 × 16	8	24	10.09	38	1884	252.68	28	1992	80.72
256	16 × 16	16	32	17.94	35	3603	269.91	27	3894	81.91
512	16 × 32	16	48	18.13	31	3434	190.36	23	6557	69.66

Table 3: Sent Message Counts for the checkpointing algorithms on different number of nodes

N	Grid				Tree			Centralized		
	Layout	Min	Max	Avg	Min	Max	Avg	Min	Max	Avg
32	4 × 8	1	12	5.88	48	429	149.41	44	592	80.16
64	8 × 8	8	15	9.80	51	799	193.25	43	1012	78.03
128	8 × 16	1	24	10.09	53	2120	252.68	42	2100	80.72
256	16 × 16	16	32	17.94	49	4053	269.91	42	4085	81.91
512	16 × 32	1	48	18.13	43	3314	190.36	35	7193	69.66

Table 4: Receive Message Counts for the checkpointing algorithms on different number of nodes

The minimum and maximum messages sent and received for the grid based algorithm is deterministic and follows from the communication pattern of the algorithm.

Even though the number of messages received and sent for the grid algorithm is small compared to the tree and centralized algorithms, this number increases at a much faster rate with an increase in the number of nodes. Comparing the message counts for the tree and centralized algorithm, observe that the average number of messages sent/received for the centralized algorithm is less than that of the tree based algorithm, however, the maximum message counts are much higher. These observations are in agreement with the

theoretical analysis; even though the message count complexity of the centralized algorithm is lower, it suffers from centralization in communication at the tail node.

7.2.4 Optimized Implementation

The initial deficit counts and number of rounds for the tree and centralized algorithms (both optimized and unoptimized) for mean checkpoint timeout values of 300 ms and 1 second, with $W = 40000$ and $M = 50000$ on 512 nodes are shown in table 5.

Algorithm	Mean Checkpoint Timeout=300 ms		Mean Checkpoint Timeout=1 s	
	Initial Deficit	No. of Rounds	Initial Deficit	No. of Rounds
Tree	30318186	13	46341632	16
Tree (Opt)	20469063	5	0	1
Centralized	30203345	13	46341632	9
Centralized (Opt)	20467177	7	0	1

Table 5: Initial Deficit Counts and number of rounds taken for unoptimized and optimized versions of tree and centralized algorithms

When the mean checkpoint timeout is 1 s, checkpointing is initiated after all the data messages have been sent and received. Therefore, in the optimized versions of the algorithms, the initial deficits are straightaway computed as 0 and the algorithms terminate in the first round.

When the mean checkpoint timeout is 300 ms, checkpointing is initiated while the data messages are being exchanged. The initial deficit counts are lower as messages that have already been received are not included in the deficit. It can be observed that the optimized algorithms terminated in lesser number of rounds. Actually, the optimized algorithms are guaranteed to terminate in the first round reset that occurs after all the processes have received all white messages destined for them.

8 Other Applications

One of the key ingredients in the tree-based and the centralized algorithm is an efficient solution of a problem that we call *distributed message counting* problem. The solution of this problem also has applications in implementation of synchronizers [Awe85]. A synchronizer is a layer of software that allows simulation of a synchronous network on asynchronous networks. The mechanism gives each process a logical abstraction of a pulse. A process can start the next pulse if all messages that have been sent to it in the last pulse have been received.

The β synchronizer [Awe85] uses a spanning tree to determine when a process is enabled to start the next pulse. Define a process to be *safe* if all messages that it has sent has been received. It is clear that if

all processes are safe then the process can start the next pulse. To detect that a process has become safe, β synchronizer uses acknowledgements. This results in overhead of a control message per application message. An alternative implementation of a synchronizer is as follows. In each pulse, the root computes the total deficit using broadcast and convergecast on the tree. It then uses distributed message counting to detect when all messages in the current pulse have been received. Finally, a new pulse is generated using broadcast.

Assuming a perfect binary tree, the β synchronizer uses $O(W + N)$ messages. Our algorithm uses $O(N \log N \log(W/N))$ messages which is significantly smaller than $O(W + N)$ when $W \gg N$.

References

- [Awe85] B. Awerbuch. Complexity of network synchronization. *Journal of the ACM*, 32(4):804–823, October 1985.
- [Bou87] L. Bouge. Repeated snapshots in distributed systems with synchronous communication and their implementation in CSP. *Theoretical Computer Science*, 49:145–169, 1987.
- [CL85] K. M. Chandy and L. Lamport. Distributed snapshots: Determining global states of distributed systems. *ACM Transactions on Computer Systems*, 3(1):63–75, February 1985.
- [CM85] K. M. Chandy and J. Misra. How processes learn. In Ray Strong, editor, *Proceedings of the 4th Annual ACM Symposium on Principles of Distributed Computing*, pages 204–214, Minaki, ON, Canada, August 1985. ACM Press.
- [Dij85] E. W. Dijkstra. The distributed snapshot of K.M. Chandy and L. Lamport. In M. Broy, editor, *Control Flow and Data Flow: Concepts of Distributed Programming*, volume F14. NATO ASI Series, Springer-Verlag, New York, NY, 1985.
- [DS80] E. W. Dijkstra and C. S Scholten. Termination detection for diffusing computations. *Information Processing Letters*, 11(4):1–4, August 1980.
- [Gar04] V. K. Garg. *Concurrent and Distributed Computing in Java*. Wiley & Sons, 2004.
- [KRS95] A D Kshemkalyani, M Raynal, and M Singhal. An introduction to snapshot algorithms in distributed computing. *Distributed Systems Engineering*, 2(4):224–233, December 1995.
- [Lam78] L. Lamport. Time, clocks, and the ordering of events in a distributed system. *Communications of the ACM*, 21(7):558–565, July 1978.

- [LY87] T. H. Lai and T. H. Yang. On distributed snapshots. *Information Processing Letters*, pages 153–158, May 1987.
- [Mat93] F. Mattern. Efficient algorithms for distributed snapshots and global virtual time approximation. *Journal of Parallel and Distributed Computing*, pages 423–434, August 1993.
- [SBF⁺04] Martin Schulz, Greg Bronevetsky, Rohit Fernandes, Daniel Marques, Keshav Pingali, and Paul Stodghill. Implementation and evaluation of a scalable application-level checkpoint-recovery scheme for MPI programs. In *SC'2004 Conference CD*, Pittsburgh, PA, November 2004. IEEE/ACM SIGARCH.
- [SK86] M. Spezialetti and P. Kearns. Efficient distributed snapshots. In *Proc. of the 6th International Conference on Distributed Computing Systems*, pages 382–388, 1986.
- [Tay89] K. Taylor. The role of inhibition in asynchronous consistent-cut protocols. In *Workshop on Distributed Algorithms*, pages 280–291. Springer-Verlag, LNCS 392, 1989.