

Distributed Snapshots

Justin Pearson

February 9, 2010

1 Introduction

The thing to remember is that the algorithm is channel based. The idea is to record the state of all the channels in the system as well as the states of the nodes.

2 The Algorithm

function MARKER SENDING RULE FOR PROCESS i

Process i records its state

For each outgoing channel C on which a marker has not been sent, i sends a marker along C before i sends further messages along C .

end function

function MARKER RECEIVING RULE FOR PROCESS j

On receiving a marker along Channel C

if j has not recorded its state **then**

Record the state of C as the empty list and Follow the 'Marker Sending Rule'

else

Record the state of C as the set of messages received along C after j 's state was recorded and before j received the marker along C .


end if

end function


- The algorithm starts by a single node decided to record the state and executing the Marker sending Rule. marker.
- The algorithm terminates when each node has received a marker along all its incoming channel.

3 Example

Consider the system with 3 processes and 4 channels. The aim of the algorithm is to record the state of the channels.




- Suppose that process A executes the marker sending rule, its outgoing edges is c_1 and c_3 , so A sends a marker to B via c_1 and C via c_3



- When B receives the marker it executes the Marker Receiving rule. It has never received a marker along channel c_1 so it begins recording the state of c_1 as the empty list and executes the marker sending rule to its outgoing channel c_2 to A
- When A has received a marker along c_2 it can stop recording its state. Note that it would of recorded all the messages that where sent along c_2 after it had sent M along c_1 .
- When C receives the marker it executes the Marker Receiving rule. It has never received a marker along channel c_3 so it begins recording the state of c_3 as the empty list and executes the marker sending rule to its

outgoing channel c_4 to B .



- Note that C has received a marker along all its input channels. So it can stop recording its state.
- When B receives a marker along c_4 from C it can stop recording its state.