Distributed Shared Memory - Read Replication Model

Pravin Paratey (99005009) **Santosh Bag (99005035)**

Read replication is a multiple reader, single writer protocol that minimizes latency by allowing more than one processor to have read-only copies of memory blocks.

Status of Assignment

Very Basic structure implemented. Our program forks to create n virtual nodes. Rest will be implemented before the final submission date.

Implementation Details

1. We have assumed that all addresses are statically distributed and that each node knows where a particular memory block resides.
2. We are using *fork()* to create multiple virtual nodes.
3. Read is being implemented as follows: If a read request comes for a block which is already shared then then requesting node is added to the share set and the block is sent back. Otherwise the block is exclusively owned by the requesting node and it is migrated to the requesting node.
4. Write will be implemented as follows: Consider the following scenario. node1 wants to write to a block held by node2. node2 has given a copy of the block to node3.
   (a) node1 will check if the block comes under its address. It finds that node2 is the one with the memory block.
   (b) So node1 will send a write request to node2 for the memory block.
   (c) On receiving the request, node2 will check whether it has lent read-only copies to other block. Yes it has given a read-only copy to node3.

(d) node2 will then send a invalidate block message to node3.

(e) node2 will reply to node1's request and hand over the block for write access

(f) node1 now has read/write access to the block.
