

# Distributed Shared Memory Service

Shankar

June 1, 2014

# Overview: Distributed shared memory – 1

- Memory shared by multiple threads on different computers
  - implemented via msg-passing
  - may also use hardware support // eg, in caches
- Provides read and write operations at the minimum
  - $val \leftarrow \text{read}(addr)$
  - $\text{void} \leftarrow \text{write}(addr, val)$
- Basic approach to implementation
  - divide memory into pages
  - “move” (a copy of) page to the user currently accessing it
- Fundamental issue:
  - when does a user's write become visible to another user
  - traditional memory: write becomes visible immediately
  - distributed memory: trade-offs

- Traditional “**sequential**” memory
  - atomic reads and writes
  - read returns last write's value
  - **poor** performance on distributed platform
- Distributed memory
  - allow reads to return earlier values
  - improved performance, but weaker correctness
  - **sequential-consistency**: equivalent to sequential wrt correctness
  - **release** consistency: weaker than sequential
  - ...
- Define sequentially-consistent memory service
  - **DsmSeqConService ( ADDR, PNO, PVAL )**

## ■ Parameters

- ADDR: addresses
- PNO: page numbers
- PVAL: possible values of a page

## ■ Input functions at addr $j$

- $j.read(pn)$  and  $j.write(pn, pv)$

## ■ Main

- $Seq\ \alpha \leftarrow []$  // seq of “write-call” and “read-ret” entries
  - $[WCALL, j, pn, pv]$  // write-call
  - $[RRET, j, pn, pv]$  // read-ret
- return  $\{v_j \leftarrow sid()\}$  // access system at  $j$

- Below  $h$  is a seq of WCALL and RRET entries
- Helper fn `lastWrite(h, pn)` // last val written to  $pn$  in  $h$ 

```

    xh ← [ [WCALL, ·, pn, ·] in h ]
    return xh.last.pv

```
- Helper fn `sequential(h)` // true iff every read returns last write

```

    return forall(i in h.keys, hi = [RRET, ·, pn, pv]:
        hi = lastWrite( h0..i-1, pn))

```
- Helper fn `seqConsistent(h)`

```

    return forsome( Seq xh:
        sequential(xh) and
        forall(j: [[·, j, ·, ·] in xh] = [[·, j, ·, ·] in h]) )

```

- `vj.read(pn)`
  - `ic { no ongoing vj.read(.) or vj.write(.) }`
  - output `pv`
    - `oc { seqconsistent( $\alpha \circ [[\text{RRET}, j, pn, pv]]$ )}`
    - `$\alpha$ .append([RRET, j, pn, pv])`
    - return `pv`
  
- `vj.write(pn, pv)`
  - `ic { no ongoing vj.read(.) or vj.write(.) }`
    - `$\alpha$ .append([RRET, j, pn, pv])`
  - `oc { true }`
    - return

- atomicity assumption: input parts and output parts
- progress assumption:
  - ongoing  $j.read(.)$  *leads-to* no ongoing  $j.read(.)$
  - ongoing  $j.write(.)$  *leads-to* no ongoing  $j.write(.)$

# Coherent shared memory

- Remote sequential memory // accessed via fifo channel
- Equivalent to seq-consistent memory with additional constraint
  - for every read or write return  $r$ :  
permutation must not move any part after  $r$  to before  $r$
- Need hardware support to implement sequentially-consistent but not coherent