

Shared Virtual Memory

Context

- Parallel architectures & programming models
 - Bus-based shared memory multiprocessors
 - h/w support for coherent shared memory
 - can run both shared memory & message passing
 - scalable to 10's of nodes
 - Distributed memory machines/clusters of workstation
 - provides message passing interface
 - scalable up to 1000s of nodes
 - cheap! economies of scale, commodity shelf h/w

Distributed Shared Memory

- Radical idea: let us not have the hardware dictate what programming model we can use
- Provide a shared address space abstraction even on clusters
- Is this a good idea? What are the upsides/downsides of this approach?

DSM Issues

- What is the granularity of sharing?
- What is the additional overhead?
- What is the consistency model?

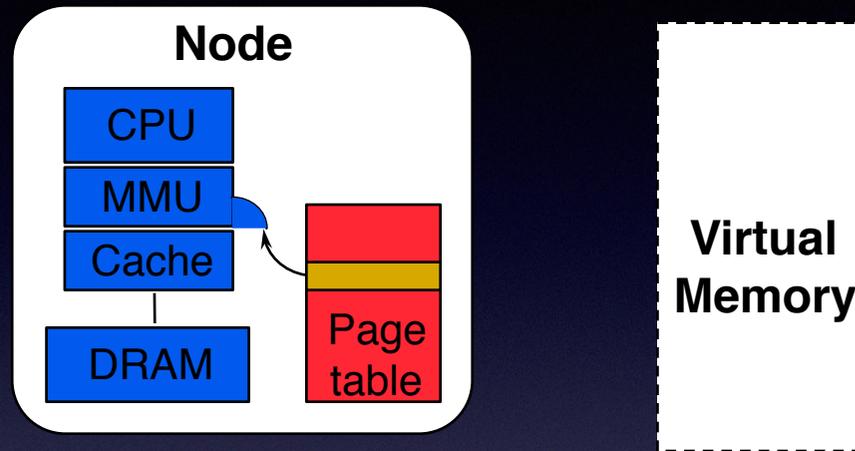
How do we provide this abstraction?

- Operating system support
 - page level granularity, hardware protection
 - e.g., Ivy, Treadmarks, Munin
- Compiler support (Shasta)
 - minimize overhead through compiler analysis
 - object granularity as opposed to byte granularity
 - notions of immutable data, sharing patterns
- Limited hardware support (Wisconsin Wind Tunnel, DEC memory channel)

IVY Shared Virtual Memory

- Seminal system that sparked the entire field of DSM (distributed shared memory)
- Motivations:
 - sharing things on a network
 - “embassy” system to support a network file system between two different OSes
 - parallel *scheme* run time system on a cluster
- Focus: parallel computing and not distributed computing
 - less emphasis on request-reply, fault-tolerance, security

Traditional Virtual Memory



- Page Table entry:

Virt. page #	physical page #	valid
--------------	-----------------	-------

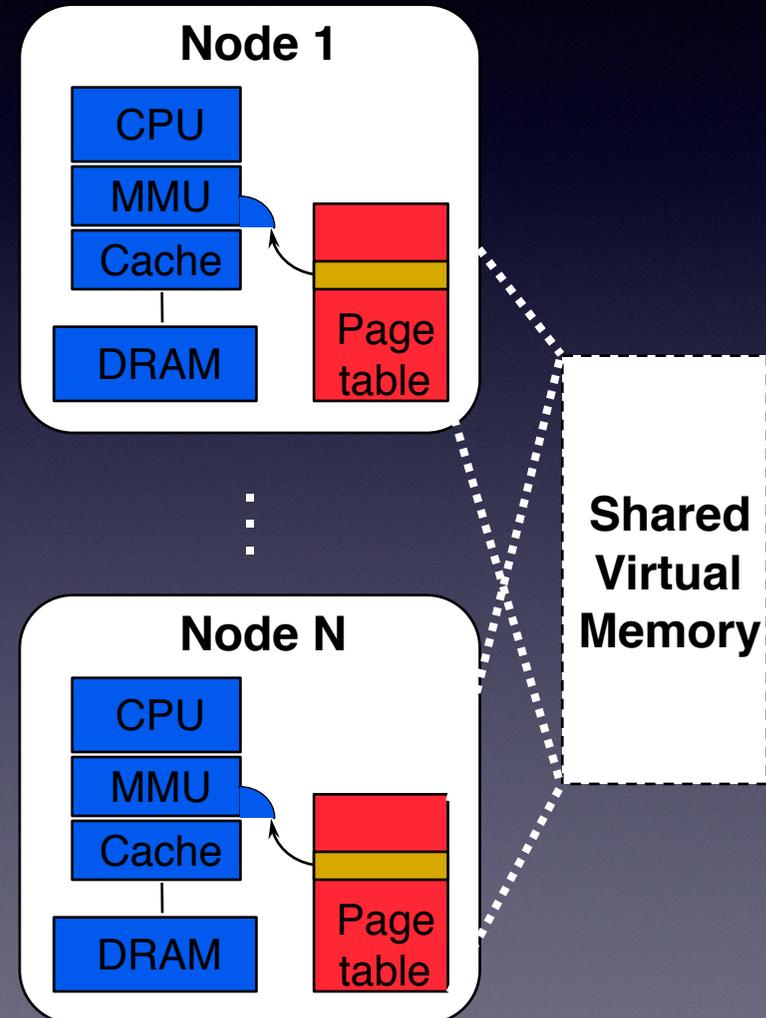
- If “valid”, translation exists
- If “not valid”, traps into the kernel, gets the page, re-executes trapped instruction
- Check is made for every access; TLB serves as a cache for the page table entries

Shared Virtual Memory

- Pool of “shared pages”: if not local, page is not mapped
- Page table entry access bits

Virt. page #	physical page #	valid	access
--------------	-----------------	-------	--------

- H/w detects read access to invalid page
 - read faults
- H/w detects writes to mapped memory with no write access
 - write faults
- OS maintains consistency at VM page level
 - copying data
 - setting access bits



Issues

- Programming model (as in coherence, consistency, etc.)
- Correctness of implementation
- Performance related issues

Programming Model

- Contract between programmer and h/w
- Ordering guarantees can vary from weak to strong
 - No guarantees
 - Coherence (i.e., ensures consistent ordering of operations on a per-variable basis)
 - Consistency model (e.g., sequential consistency that takes into account program order and relates operations across variables)
 - Linearizability (strong guarantee)

Sequential Consistency

- “The result of any execution is the same as if the operations of all the processes were executed in some sequential order and the operations of each individual process appear in this sequence in the order specified by its program” (Lamport, 1979)

$p_1: W(x)a$

$p_2: W(x)b$

$p_3: R(x)b^1 \quad R(x)a^2$

$p_4: R(x)b^1 \quad R(x)a^2$

Is this data store sequentially consistent?

Sequential Consistency

- “The result of any execution is the same as if the operations of all the processes were executed in some sequential order and the operations of each individual process appear in this sequence in the order specified by its program” (Lamport, 1979)

$p_1: W(x)a$

$p_2: W(x)b$

$p_3: R(x)b^1 \quad R(x)a^2$

$p_4: R(x)a^1 \quad R(x)b^2$

Is this data store sequentially consistent?

Sources of “reorderings”

- Multiple caches & cache update policy
- Write buffer and reordering of writes
- Compiler rewrites

- Doing something smart to hide latency ==> weaker consistency guarantees

Linearizability

“The result of any execution is the same as if the operations of all the processes were executed in some sequential order and the operations of each individual process appear in this sequence in the order specified by its program.

In addition, if $ts_{OP1}(x) < ts_{OP2}(y)$, then operation $OP1(x)$ should precede $OP2(y)$ in this sequence (Herlihy & Wing, 1991)

Weakening Sequential Consistency: Causal Consistency

- Writes that are potentially causally related must be seen by all processes in the same order. Concurrent writes may be seen in a different order on different machines. (Hutto and Ahamad, 1990)

$p_1: W(x)a \longrightarrow W(x)c$

$p_2: R(x)a \longrightarrow W(x)b$

$p_3: R(x)a \qquad R(x)c \qquad R(x)b$

$p_4: R(x)a \qquad R(x)b \qquad R(x)c$

Is this data store sequentially consistent?

Causally consistent?

More Weakening: FIFO Consistency

- “Writes done by a single process are seen by all other processes in the order in which they were issued, but writes from different processes may be seen in a different order by different processes” (PRAM consistency, Lipton and Sandberg 1988)

$p_1 : W(x)a$

$p_2 : R(x)a \rightarrow W(x)b \rightarrow W(x)c$

$p_3 : R(x)b \quad R(x)a \quad R(x)c$

$p_4 : R(x)a \quad R(x)b \quad R(x)c$

Is this data store causally consistent?

Is this data store FIFO consistent?

- What do you make out of these consistency models?

Ivy DSM

- Goal: provide sequentially consistent shared memory
- Distributed protocol for coherence
- Baseline Implementation:
 - centralized manager
 - manager maintains the “owner” and the set of readers (“copyset”)

Read Faults

- Handler on client:
 - asks manager
 - manager forwards request to owner
 - owner sends the page
 - requester sends an ACK to manager

Pseudocode

Read Fault Handler:

```
Lock (Ptable [p] .lock) ;  
ask manager for p ;  
receive p ;  
send confirmation to manager ;  
Ptable [p] .access = read ;  
Unlock (Ptable [p] .lock) ;
```

Manager:

```
Lock (Info [p] .lock) ;  
Info [p] .copyset =  
    Info [p] .copyset U {reqNode} ;  
ask Info [p] .owner to send p ;  
receive confirmation from reqNode ;  
Unlock (Info [p] .lock) ;
```

Read Server:

```
Lock (Ptable [p] .lock) ;  
Ptable [p] .access = read ;  
send copy of p ;  
Unlock (Ptable [p] .lock) ;
```

Write Faults

- Handling includes invalidations:
 - make request to manager
 - copies are invalidated
 - manager forwards request to owner
 - owner relinquishes page to requester
 - requester sends an ACK to the owner

Write Pseudocode

Write Fault Handler:

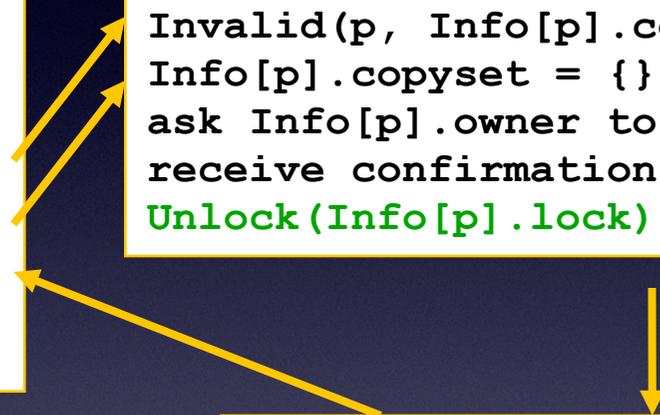
```
Lock(Ptable[p].lock);  
ask manager for p;  
receive p;  
send confirmation to manager;  
Ptable[p].access = write;  
Unlock(Ptable[p].lock);
```

Manager:

```
Lock(Info[p].lock);  
Invalid(p, Info[p].copyset);  
Info[p].copyset = {};  
ask Info[p].owner to send p;  
receive confirmation from reqNode;  
Unlock(Info[p].lock);
```

Write Server:

```
Lock(Ptable[p].lock);  
Ptable[p].access = nil;  
send copy of p;  
Unlock(Ptable[p].lock);
```



Improved Manager

- Owner serves as the manager for each page

Read Fault Handler:

```
Lock (Ptable [p] .lock) ;  
ask manager for p ;  
receive p ;  
Ptable [p] .access = read ;  
Unlock (Ptable [p] .lock) ;
```

Read Server:

```
Lock (Ptable [p] .lock) ;  
If I am owner {  
    Ptable [p] .access = read ;  
    Ptable [p] .copyset =  
        Ptable [p] .copyset U {reqNode} ;  
    send copy of p ;  
} else {  
    forward request to probable owner ;  
}  
Unlock (Ptable [p] .lock) ;
```

Performance Questions

- In what situations will IVY perform well?
- How can we improve IVY's performance?