

Constructive Computer Architecture

Cache Coherence

Arvind

Computer Science & Artificial Intelligence Lab.
Massachusetts Institute of Technology


Plan

- ◆ The invalidation protocol
- ◆ Non-blocking L1

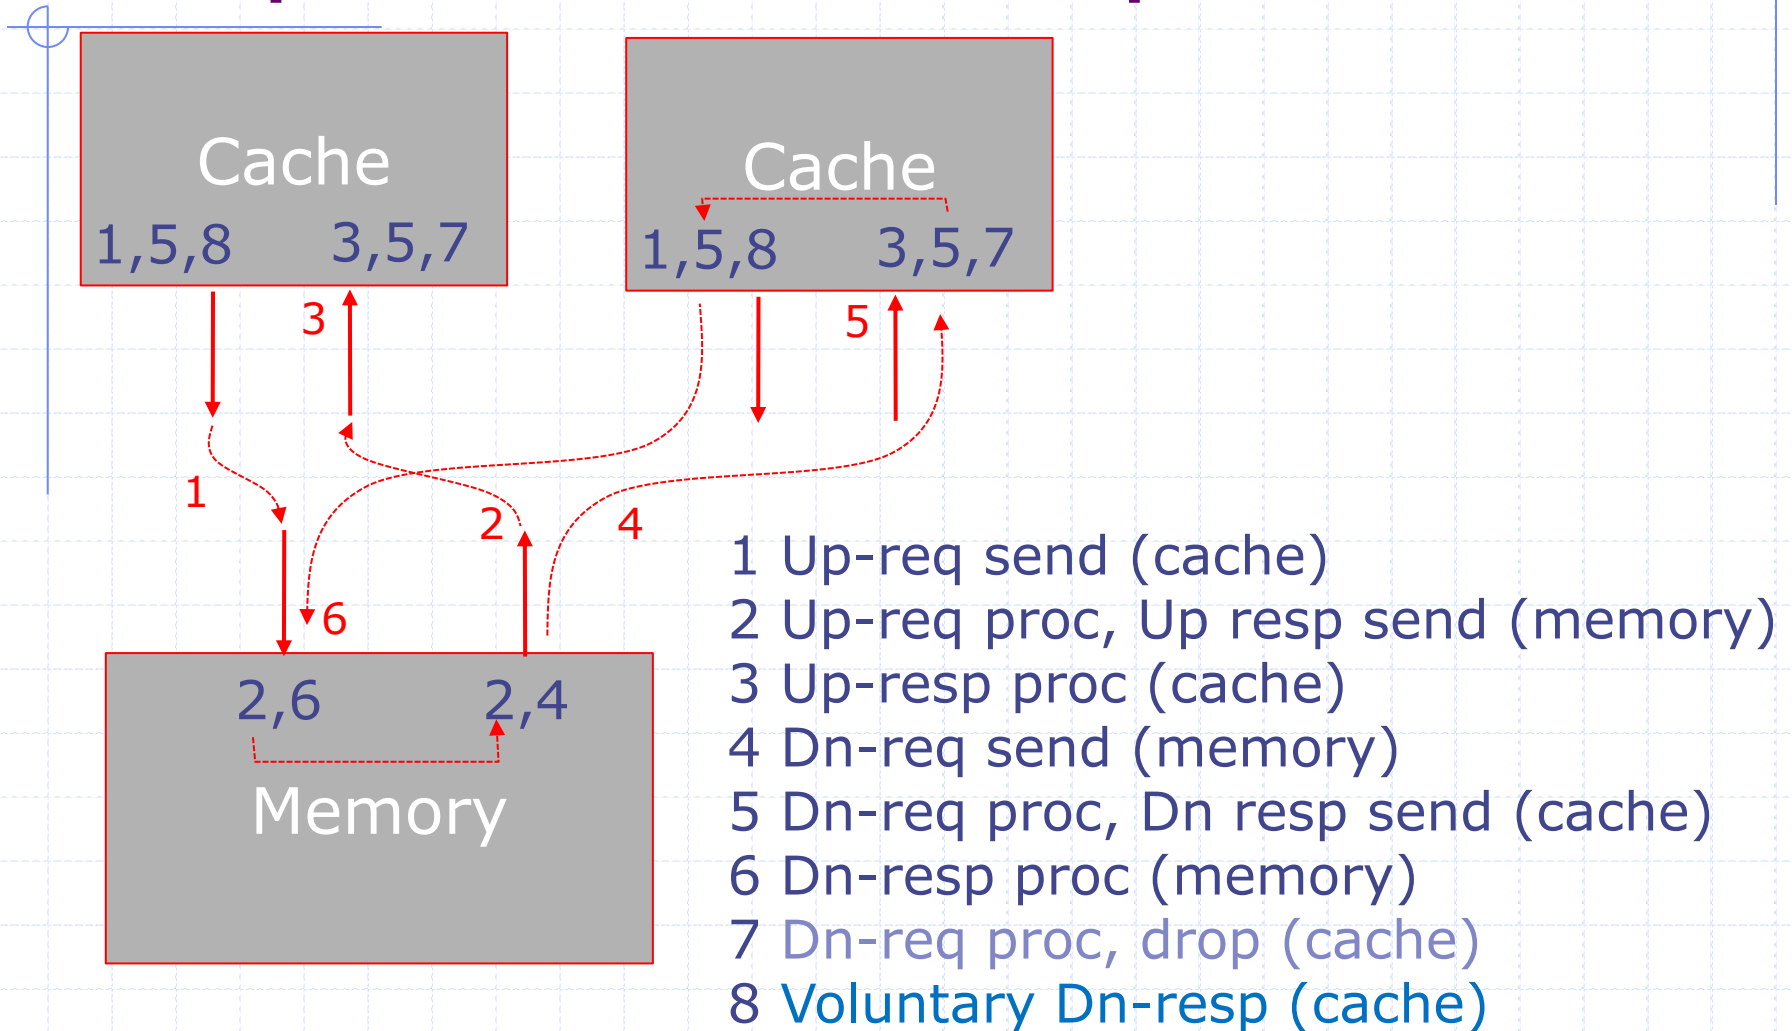
Processor Hit Rules

◆ Load-hit rule
p2m.msg=(Load a) &
c.tag[cs(a)]=tag(a) &
c.state[cs(a)]>I
→ p2m.deq;
m2p.enq(c.data[cs(a)]);

◆ Store-hit rule
p2m.msg=(Store a v) &
c.tag[cs(a)]=tag(a) &
c.state[cs(a)]=M
→ p2m.deq;
m2p.enq(Ack);
c.data[cs(a)]:=v;



Processing misses: Requests and Responses



Invariants for a CC-protocol design

- ◆ Directory state is always a conservative estimate of a child's state
 - E.g., if directory thinks that a child cache is in S state then the cache has to be in either I or S state
- ◆ For every request there is a corresponding response, though sometimes it is generated even before the request is processed
- ◆ Communication system has to ensure that
 - responses cannot be blocked by requests
 - a request cannot overtake a response for the same address
- ◆ At every merger point for requests, we will assume fair arbitration to avoid starvation

Child Requests

1. Child to Parent: Upgrade-to-y Request

$c.tag[cs(a)] \neq tag(a) \ \& \ c.state[cs(a)] = I \ \&$

$c.waitp[cs(a)] = No$

$\rightarrow c.tag[cs(a)] := tag(a);$

$c.waitp[cs(a)] := Yes \ y;$

$c2m.enq(\langle Req, c \rightarrow m, a, y, - \rangle);$

A request is never sent unless the cache has a slot and the slot contains the tag

$cs(a)$ is the cache slot address

$c.tag[cs(a)] = tag(a) \ \& \ (c.state[cs(a)] < y) \ \&$

$c.waitp[cs(a)] = No$

$\rightarrow c.waitp[cs(a)] := Yes \ y;$

$c2m.enq(\langle Req, c \rightarrow m, a, y, - \rangle);$

These rules are mutually exclusive and can be combined. This rule would normally be triggered by a cache miss.

Parent Responds

2. Parent to Child: Upgrade-to-y response

```
( $\forall j$ , m.waitc[j][a]=No) & c2m.msg=<Req,c $\rightarrow$ m,a,y,-> &  
( $\forall i \neq c$ , IsCompatible(m.child[i][a],y))  
 $\rightarrow$  m2c.enq(<Resp, m $\rightarrow$ c, a, y,  
            (if (m.child[c][a]=I) then m.data[a] else -)>);  
m.child[c][a]:=y; c2m.deq;
```

Child receives Response

3. Child receiving upgrade-to-y response

```
m2c.msg = <Resp, m → c, a, y, data>
```

```
→ m2c.deq;
```

```
if(c.state[cs(a)] = I) c.data[cs(a)] := data;
```

```
c.state[cs(a)] := y;
```

```
c.waitp[cs(a)] := No;
```

```
// the child must be waiting for state y
```


Parent Requests

4. Parent to Child: Downgrade-to-y Request

```
c2m.msg = <Req, c → m, a, y, - > &  
(m.child[i][a] > y) & (m.waitc[i][a] = No)  
→ m.waitc[i][a] := Yes y;  
m2c.enq(<Req, m → c, a, y, - >);
```

Child Responds

5. Child to Parent: Downgrade-to-y response

(m2c.msg=<Req,m→c,a,y,->) &

c.state[cs(a)]>y &

c.tag[cs(a)]=tag(a)

→ c2m.enq(<Resp, c->m, a, y,

(if (c.state[cs(a)]=M) then c.data[a] else -)>);

c.state[cs(a)]:=y; m2c.deq

Parent receives Response

6. Parent receiving downgrade-to-y response

```
c2m.msg = <Resp, c → m, a, y, data>
```

```
→ c2m.deq;
```

```
if(m.child[c][a] = M) m.data[a] := data;
```

```
m.child[c][a] := y;
```

```
if(m.waitc[c][a] = (Yes x) & x ≥ y)
```

```
    m.waitc[c][a] := No;
```

Child receives served Request

7. Child receiving downgrade-to-y request
(m2c.msg=<Req, m→c, a, y, - >) &
((c.tag[cs(a)]=tag(a) & c.state[cs(a)]≤y)
|| c.tag[cs(a)]!=tag(a))
→ m2c.deq;

Child Voluntarily downgrades

8. Child to Parent: Downgrade-to-y response (vol)
(c.waitp[cs(a)]=No) & (c.state[cs(a)]>y)
→ c2m.enq(<Resp, c->m, a, y,
 (if (c.state[cs(a)]=M) then c.data[a] else -)>);
 c.state[cs(a)]:=y;

Rules 1 to 8 are complete - cover all possibilities
and cannot deadlock or violate cache invariants

Non-blocking Cache

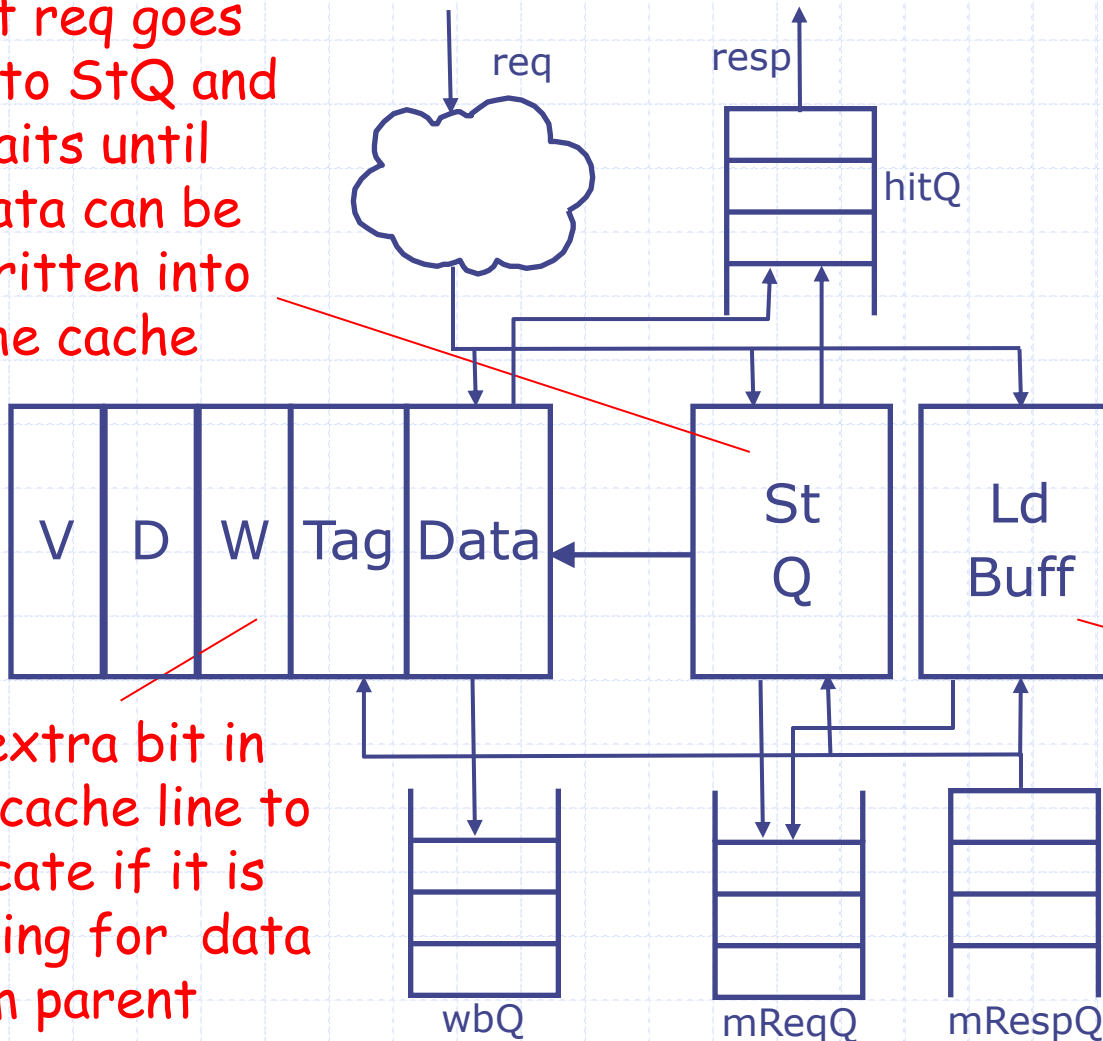
single processor

St req goes into StQ and waits until data can be written into the cache

An extra bit in the cache line to indicate if it is waiting for data from parent

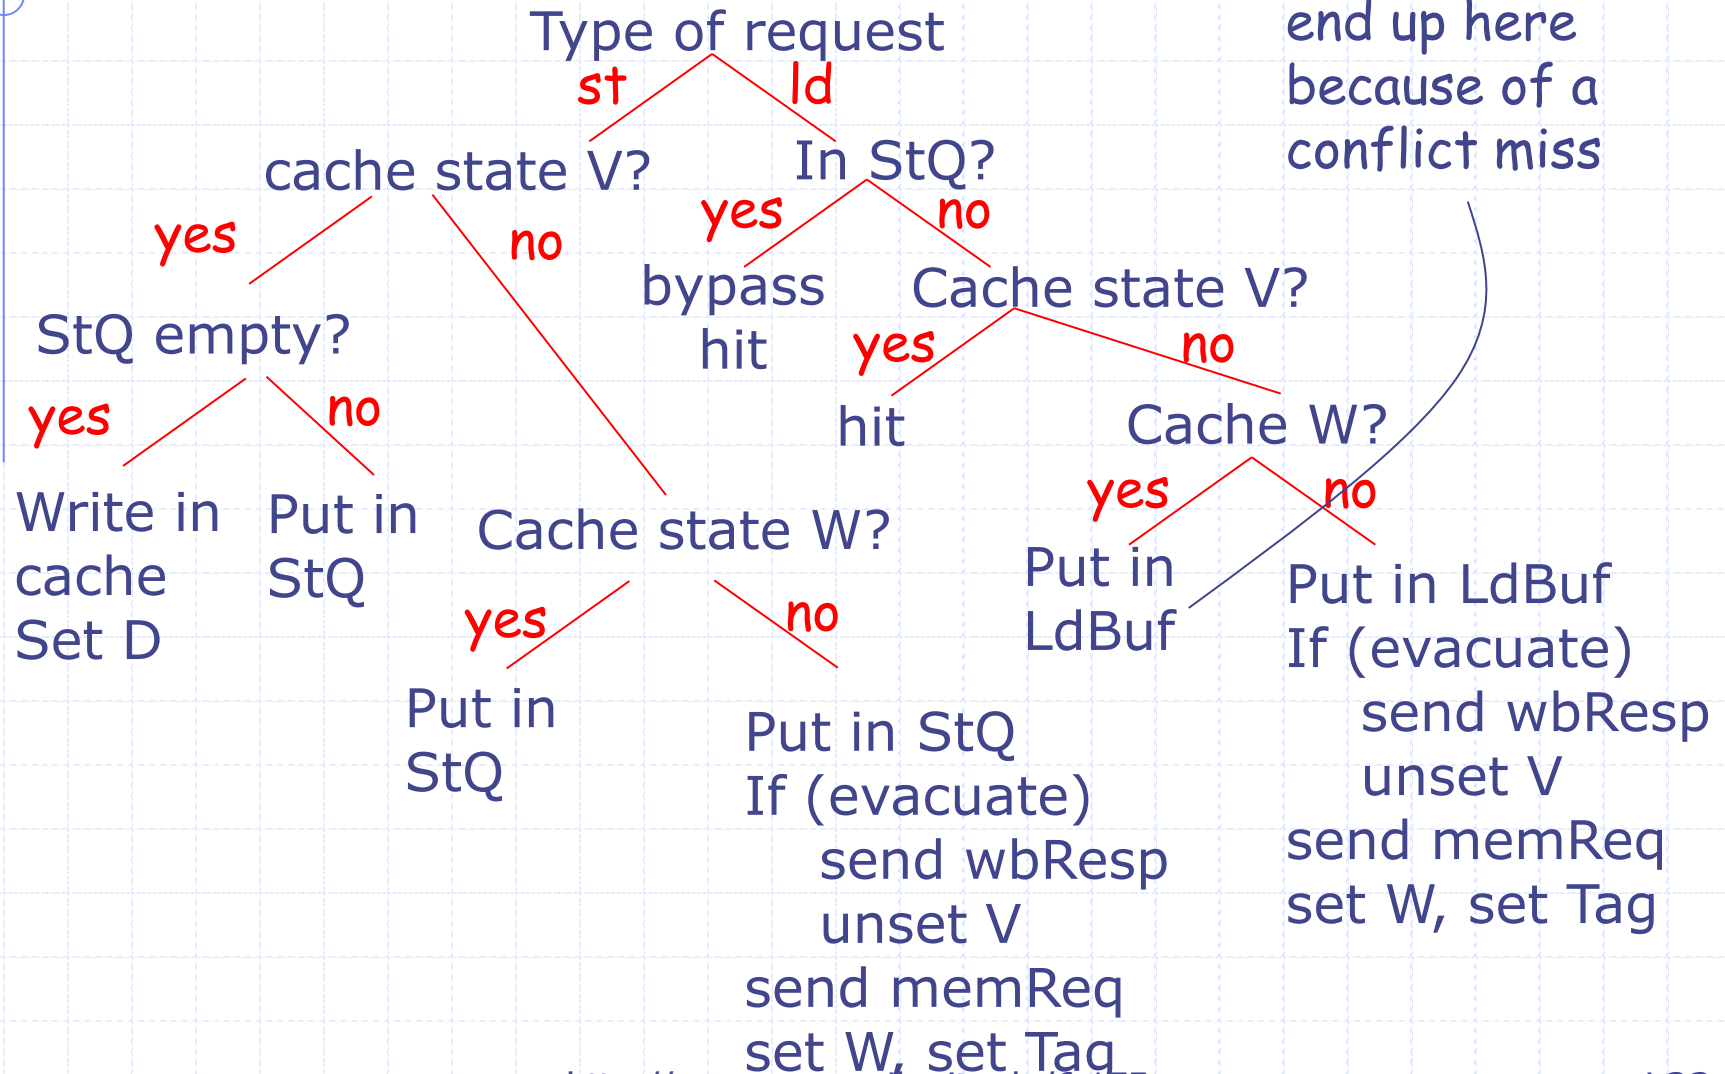
Behavior to be described by 2 concurrent FSMs to process input requests and memory responses, respectively

load reqs waiting for data



Incoming req

single processor



a request may end up here because of a conflict miss

Mem Resp (for line cl)

single processor

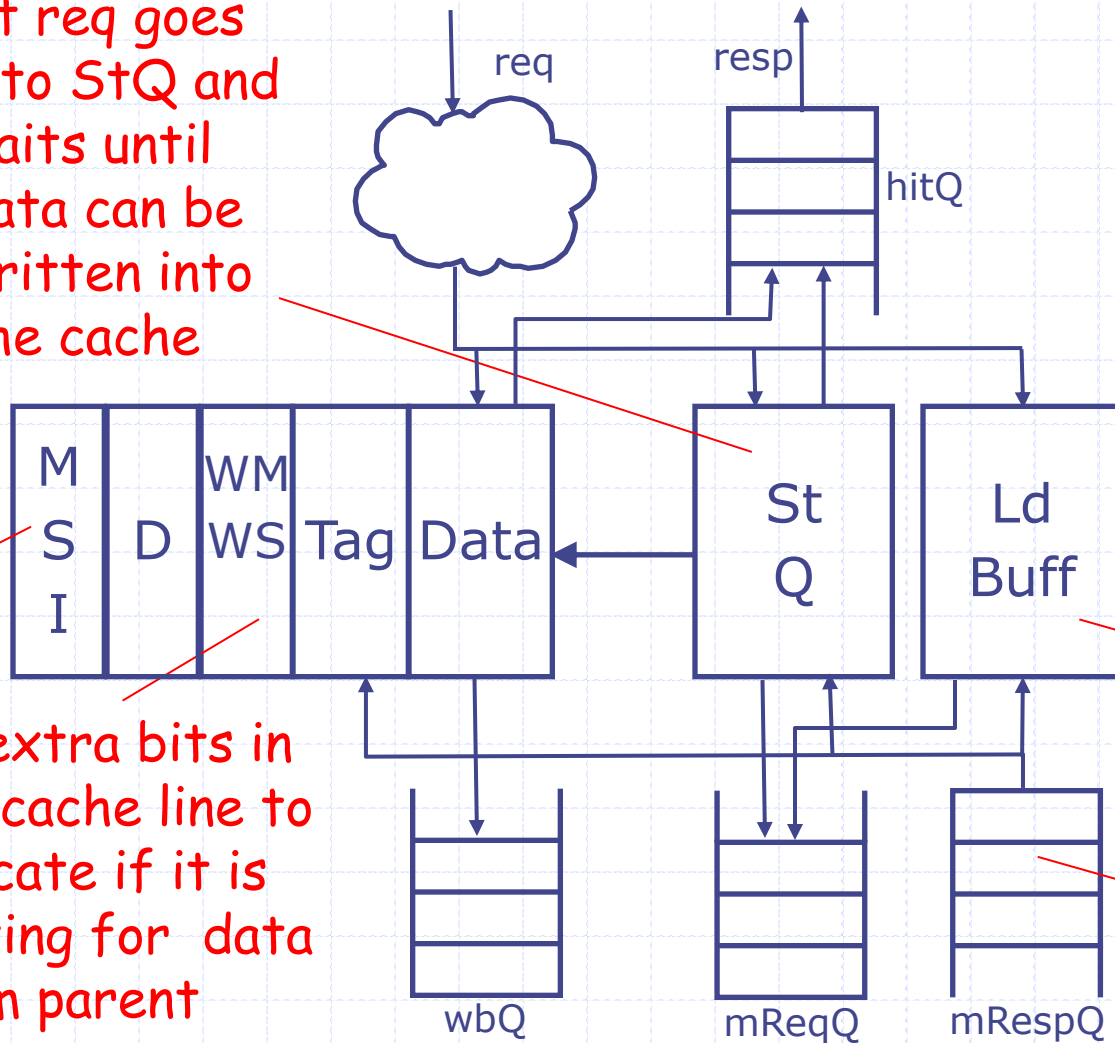
1. Update cache line (set V, unset D, and unset W)
2. Process all matching ldbuf entries and send responses
3. L: If cachestate(oldest StQ entry address) = V
then
 update the cache word with StQ entry; set D
 remove the oldest entry;
 Loop back to L
else if there is a ldbuf entry for cl // process conflict misses
 then if(evacuate) wbResp; unset V
 memReq for the address in ldbuf;
 set W, set Tag
else if cachestate(oldest StQ entry address) = !W
 then if(evacuate) wbResp; unset V
 memReq for this store entry;
 set W, set Tag

Non-blocking Cache

multi-processor

St req goes into StQ and waits until data can be written into the cache

Behavior to be described by 2 concurrent FSMs to process input requests and memory responses, respectively



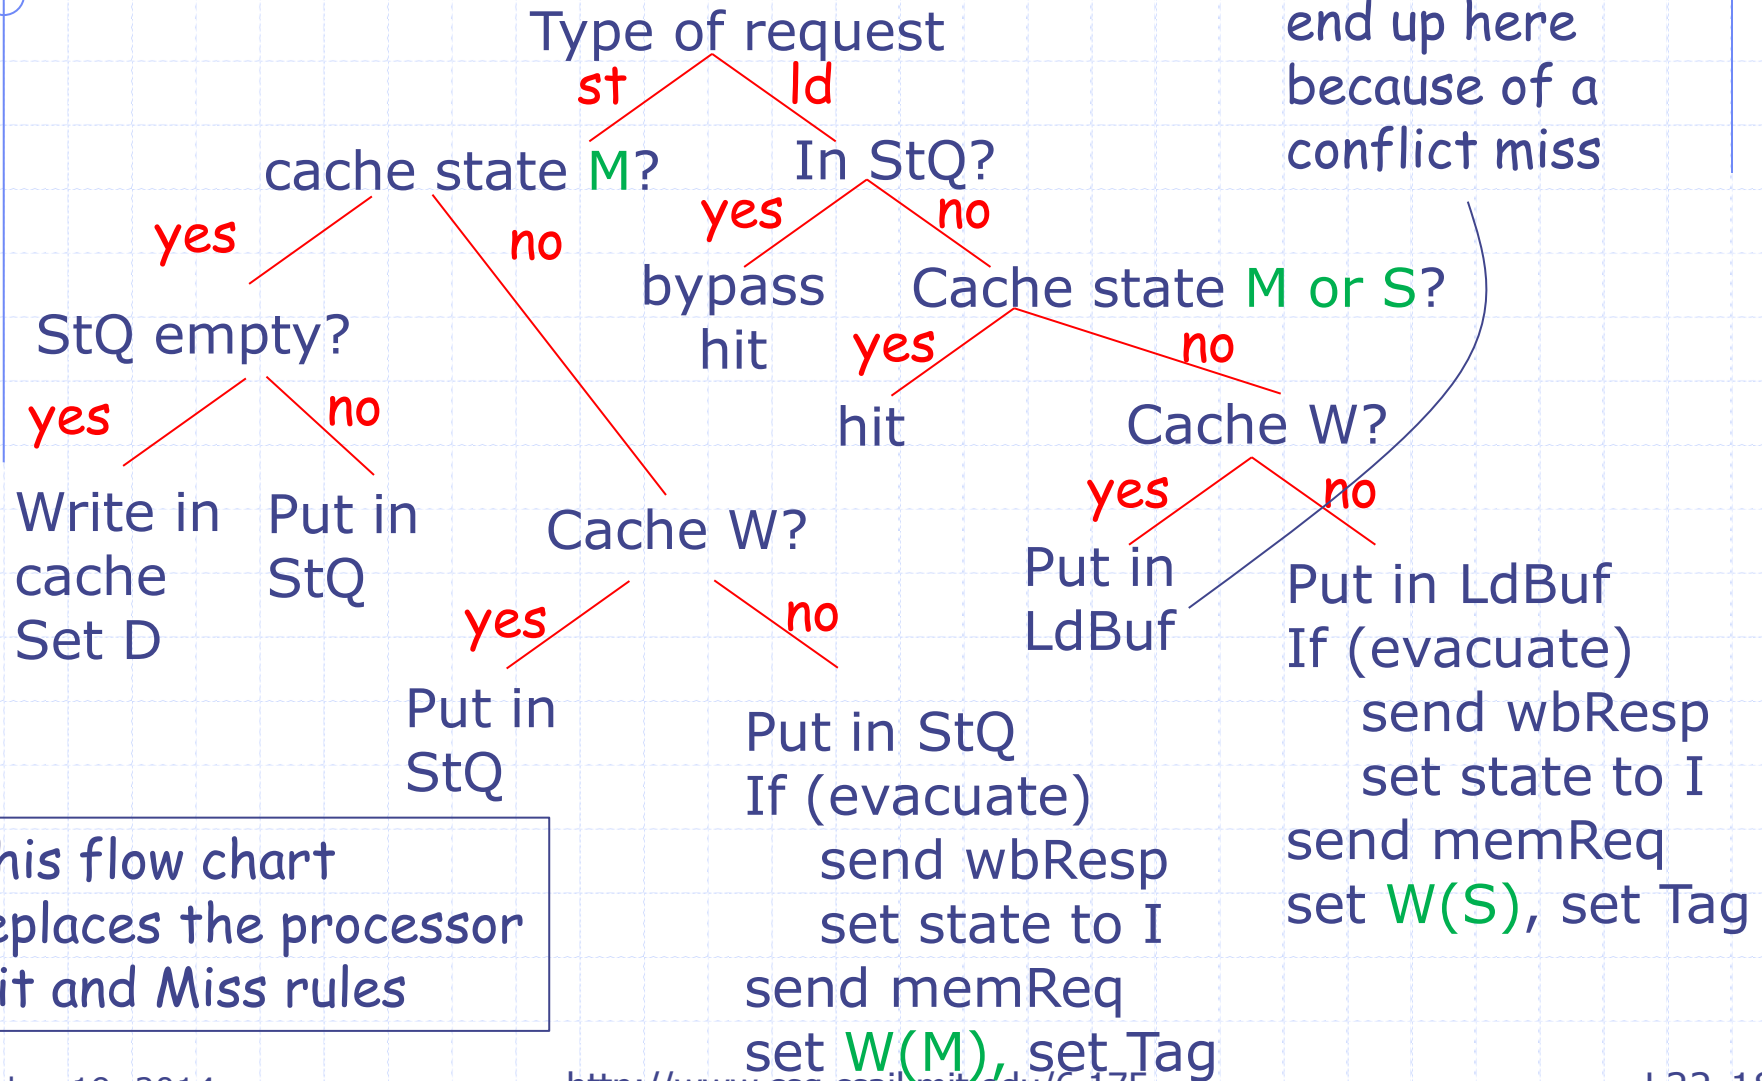
load reqs waiting for data

Includes invalidation messages

An extra bits in the cache line to indicate if it is waiting for data from parent

Incoming req

multi-processor



This flow chart replaces the processor Hit and Miss rules

Mem Resp (for line cl)

multi-processor

1. Update cache line (set state to M or S based on W, unset D, unset W)
2. Process all matching ldbuf entries and send responses
3. L: If cachestate(oldest StQ entry address) = M
then
 update the cache word with StQ entry; set D
 remove the oldest entry;
 Loop back to L
else if there is a ldbuf entry for cl
 then if(evacuate) wbResp; set state to I
 memReq for the address in ldbuf;
 set W(S), set Tag
else if cachestate(oldest StQ entry address) = !W
 then if(evacuate) wbResp; set state to I
 memReq for this store entry;
 set W(M), set Tag

This flow
chart
replaces
rule 3
(for L1)

*next - Network and buffer issues to
avoid deadlocks*