

CS184c: Computer Architecture [Parallel and Multithreaded]

Day 4: April 12, 2001
Network Interface



CALTECH cs184c Spring2001 -- DeHon

Parallelism/Concurrency Background?

- ..see where people are coming from

CALTECH cs184c Spring2001 -- DeHon

Today

- Message Handling Requirements
- Active Messages
- Processor/Network Interface Integration

CALTECH cs184c Spring2001 -- DeHon

What does message handler have to do?

- Send
 - allocate buffer to compose outgoing message
 - figure out destination
 - address
 - routing?
 - Format header info
 - compose data for send
- put out on network
 - copy to privileged domain
 - check permissions
 - copy to network device
 - checksums?

CALTECH cs184c Spring2001 -- DeHon

What does message handler have to do?

- Receive
 - queue result
 - copy buffer from queue to privileged memory
 - check message intact (checksum)
 - message arrived right place?
 - Reorder messages?
 - Filter out duplicate messages?
- Figure out which process/task gets message
- check privileges
- allocate space for incoming data
- copy data to buffer in task
- hand off to task
- decode message type
- dispatch on message type
- handle message

CALTECH cs184c Spring2001 -- DeHon

1990 Message Handling

- nCube/2 160 μ s (360ns/byte)
- CM-5 86 μ s (120ns/byte)

CALTECH cs184c Spring2001 -- DeHon

Functions: Destination Addressing/Routing

- Prefigure and put in pointer/object
 - hoist out of inner loop of computation
 - just a lookup
- TLB-like translation table
 - provide hardware support for common case

CALTECH cs184c Spring2001 -- DeHon

Functions: Allocation

- In general,
 - messages may be arbitrarily long
 - many dynamic sized objects...
 - may not be expecting message ?
 - Remote procedure invocation
 - Remote memory request
- Hand off to OS
 - OS user/consumption asynchronous to user process (lifetime unclear)

CALTECH cs184c Spring2001 -- DeHon

Functions: Allocation

- Pre-allocate outside of messaging
 - when expecting a message
 - ? Standard sizes for common cases?
- Avoid copying
 - shared memory
 - issues with synchronization
 - direct from user space

CALTECH cs184c Spring2001 -- DeHon

Functions: Ordering

- Network may reorder messages
 - multiple paths
 - with different lengths, congestion
- Not all tasks require ordering
 - (may be ordered at higher level in computation)
 - dataflow firing example
- What requires ordering?

CALTECH cs184c Spring2001 -- DeHon

Functions: Idempotence

- Failed Acknowledgments
 - may lead to multiple delivery of same message
- Idempotent operations
 - bit-set, bit-unset,
- Non-idempotent operations
 - increment, exchange
- How make idempotent
 - TCP example

CALTECH cs184c Spring2001 -- DeHon

Functions: Protection

- Don't want messages from other processes/entities
 - give away information (get)
 - destroy state (put)
 - perform operation (transfer funds)

CALTECH cs184c Spring2001 -- DeHon

Functions: Protection

- How manage?
 - Treat network as IO
 - OS mediates
 - (can) trust message stamps on network
 - Give network to user
 - messaging hardware tags with process id
 - filter messages on process tags
 - (can) trust message stamps because of hardware
 - Cryptographic packet encoding

CALTECH cs184c Spring2001 -- DeHon

Functions: Checksum

- Message could be corrupted in transit
 - likely with high-bit rate, long interconnect
 - (multiple chips...multiple boxes...)
- Wrong bits
 - in address
 - in message id
 - in data
- Typically solve in hardware

CALTECH cs184c Spring2001 -- DeHon

What does message handler have to do?

- Send
 - allocate buffer to compose outgoing message
 - figure out destination
 - address
 - routing?
 - Format header info
 - compose data for send
 - put out on network
 - copy to privileged domain
 - check permissions
 - copy to network device
 - checksums
- Not all messages require
- Hardware support
- Avoid (don't do it)

CALTECH cs184c Spring2001 -- DeHon

Not all messages require
Hardware support
Avoid (don't do it)

What does message handler have to do?

- Receive
 - queue result
 - copy buffer from queue to privilege memory
 - check message intact (checksum)
 - message arrived right place?
 - Reorder messages?
 - Filter out duplicate messages?
- Figure out which process/task gets message
- check privileges
- allocate space for incoming data
- copy data to buffer in task
- hand off to task
- decode message type
- dispatch on message type
- handle message

CALTECH cs184c Spring2001 -- DeHon

End-to-End

- Variant of the primitives argument
- Applications/tasks have different requirements/needs
- Attempt to provide in the network
 - mismatch
 - unnecessary
- Network should be minimal
 - let application do just what it needs

CALTECH cs184c Spring2001 -- DeHon

Active Messages

- Message contains PC of code to run
 - destination
 - message handler PC
 - data
- Receiver pickups PC and runs
- [similar to J-Machine, conv. CPU]

CALTECH cs184c Spring2001 -- DeHon

Active Message Dogma

- Integrate the data directly into the computation
- Short Runtime
 - get back to next message, allows to run directly
- Non-blocking
- No allocation
- Runs to completion
- ...Make fast case common

CALTECH cs184c Spring2001 -- DeHon

Stopped Here

4/12/01

CALTECH cs184c Spring2001 -- DeHon

User Level NI Access

- Avoids context switch
- Viable if hardware manage process filtering

CALTECH cs184c Spring2001 -- DeHon

Hardware Support I

- Checksums
- Routing
- ID and route mapping
- Process ID stamping/checking
- Low-level formatting

CALTECH cs184c Spring2001 -- DeHon

What does AM handler do?

- Send
 - compose message
 - destination
 - receiving PC
 - data
 - copy/queue to NI
- Receive
 - pickup PC
 - dispatch to PC
 - handler dequeues data into place in computation
 - [maybe more depending on application]
 - idempotence
 - ordering
 - synchronization

CALTECH cs184c Spring2001 -- DeHon

Example: PUT Handler

- Message:
 - remote node id
 - put handler (PC)
 - remote adder
 - data length
 - (flag adder)
 - data
- No allocation
- Idempotent
- Receiver:
 - poll
 - r1 = packet_pres
 - beq r1 0 poll
 - r2=packet(0)
 - branch r2
 - put_handler
 - r3=packet(1)
 - r4=packet(2)
 - r5=packet+r4
 - r6=packet+3
 - mdata
 - *r3=packet(r6)
 - r6++
 - blt r6,r5 mdata
 - consume packet
 - goto poll

CALTECH cs184c Spring2001 -- DeHon

Example: GET Handler

- Message Request
 - remote node
 - get handler
 - local addr
 - data length
 - (flag addr)
 - local node
 - remote addr
- Message Reply can just be a PUT message
 - put into specified local address

CALTECH cs184c Spring2001 -- DeHon

Example: GET Handler

```
get_handler
- out_packet(0)=packet(6)
- out_packet(1)=put_handler
- out_packet(2)=packet(3)
- out_packet(3)=packet(4)
- r6=4
- r7=packet(7)
- r5=packet(4)
- r5=r5+4

mdata
- packet(r6)=*r7
- r6++
- r7++
- blt r6,r5 mdata
- consume packet
- goto poll
```

CALTECH cs184c Spring2001 -- DeHon

Example: DF Inlet synch

- Consider 3 input node (e.g. add3)
 - “inlet handler” for each incoming data
 - set presence bit on arrival
 - compute node when all present

CALTECH cs184c Spring2001 -- DeHon

Example: DF Inlet Synch

- inlet message
 - node
 - inlet_handler
 - frame base
 - data_addr
 - flag_addr
 - data_pos
 - data
- Inlet
 - move data to addr
 - set appropriate flag
 - if all flags set
 - enable DF node computation
- ? Care not enable multiple times?

CALTECH cs184c Spring2001 -- DeHon

Interrupts vs. Polling

- What happens on message reception?
- Interrupts
 - cost context switch
 - interrupt to kernel
 - save state
 - force attention to the network
 - guarantee get messages out of input queue in a timely fashion

CALTECH cs184c Spring2001 -- DeHon

Interrupts vs. Polling

- Polling
 - if getting many messages to same process
 - message handlers short / bounded time
 - may be fine to just poll between handlers
 - requires:
 - user-level/fine-grained scheduling
 - guarantee will get back to
 - avoid context switch cost

CALTECH cs184c Spring2001 -- DeHon

Interrupts vs. Polling

- Can be used together to minimize cost
 - poll network interface during batch handling of messages
 - interrupt to draw attention back to network if messages sit around too long
 - polling works for same process
 - interrupt if different process
 - common case is work on same process for a while

CALTECH cs184c Spring2001 -- DeHon

AM vs. JM

- J-Machine handlers can fault/stall
 - touch futures...
- J-Machine fast context with small state
 - not get to exploit rich context/state
- AM exploits register locality by scheduling together larger block of data
 - processing related handlers together (same context)
 - more next week (look at TAM)

CALTECH cs184c Spring2001 -- DeHon

Active Message Results

- CM5 (user-level messaging)
 - send $1.6\mu\text{s}$ [50 instructions]
 - receive/dispatch $1.7\mu\text{s}$
- nCube/2 (OS must intervene)
 - send $11\mu\text{s}$ [21 instructions]
 - receive $15\mu\text{s}$ [34 instructions]
- Myrinet (GM)
 - $6.5\mu\text{s}$ end-to-end GMs
 - $1\text{-}2\mu\text{s}$ host processing time

CALTECH cs184c Spring2001 -- DeHon

Hardware Support II

- Roll presence tests into dispatch
- compose message data from registers
- common case
 - reply support
 - message types
- Integrate network interface as functional unit

CALTECH cs184c Spring2001 -- DeHon

Presence Dispatch

- Handler PC in common location
- Have hardware supply null handler PC when no messages current
- Poll:
 - read MsgPC into R1
 - branch R1
- Also use to handle cases and priorities
 - by modifying a few bits of dispatch address
 - e.g. queues full/empty

CALTECH cs184c Spring2001 -- DeHon

Compose from Registers

- Put together message in registers
 - reuse data from message to message
 - compute results directly into target
 - user register renaming and scoreboarding to continue immediately while data being queued

CALTECH cs184c Spring2001 -- DeHon

Common Case Msg/Replies

- Instructions to
 - fill in common data on replies
 - node address, handler?
 - Indicate message type
 - not have to copy

CALTECH cs184c Spring2001 -- DeHon

Example: GET handler

- Get_handler
 - R1=i0 // address from message register
 - R2=*R1
 - o2=R2 // value into output data register
 - SEND -reply type=reply_mesg_id
 - NEXT

CALTECH cs184c Spring2001 -- DeHon

AM as primitive Model

- Value of Active Messages
 - articulates a model of what primitive messaging needs to be
 - identify key components
 - then can optimize against
 - how much hardware to support?
 - What should be in hardware/software?
 - What are common cases?
 - Should get special treatment?

CALTECH cs184c Spring2001 -- DeHon

Big Ideas

- Primitives/Mechanisms
 - End-to-end/common case
 - don't burden everything with features needed by only some tasks
- Abstract Model
- Separate essential/useful work from overhead

CALTECH cs184c Spring2001 -- DeHon

[MSB-1] Ideas

- Minimize Overhead
 - moving data around is not value added operation
 - minimize copying
- Overlap Compute and Communication
 - queue
 - don't force send/receive rendezvous
- Get the OS out of the way of **common operations**

CALTECH cs184c Spring2001 -- DeHon