Fig 1

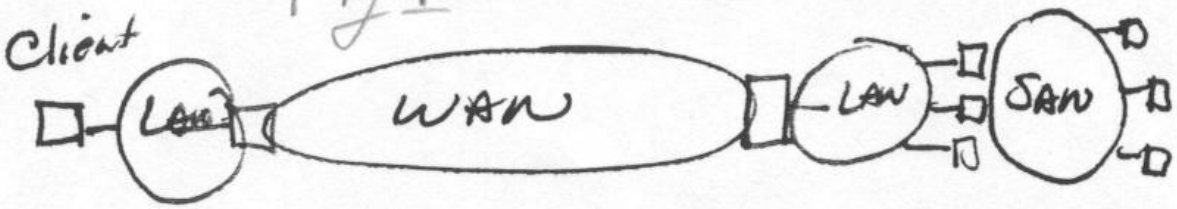

What's a network system?

Fig 2

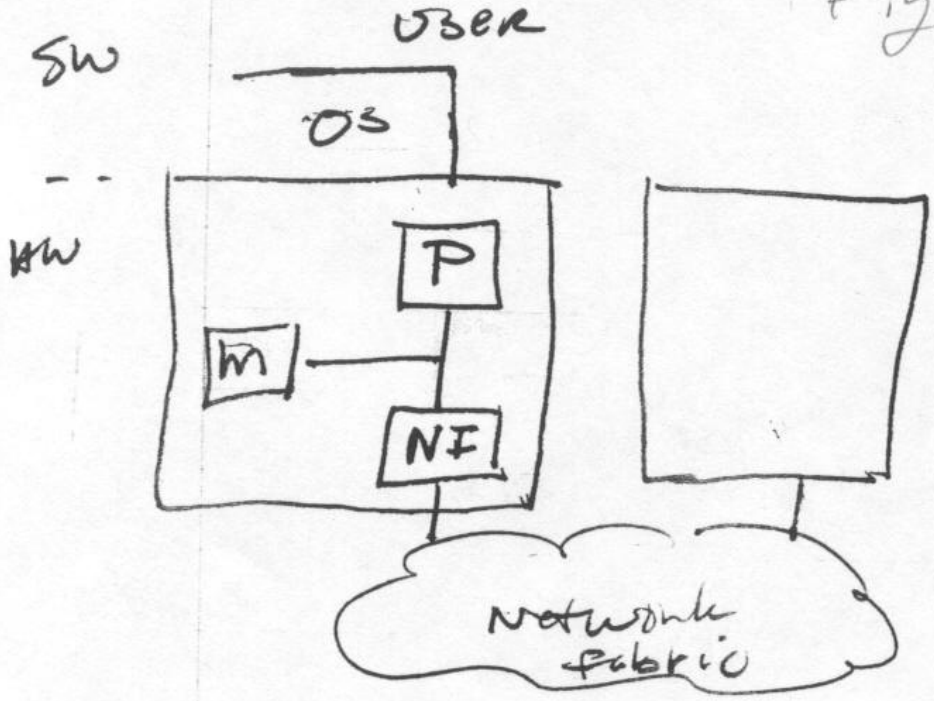

Fig 3

Typical LAN NIC

Issues:

- juggling events
- accessing data & control info
- copy (blind data receive)

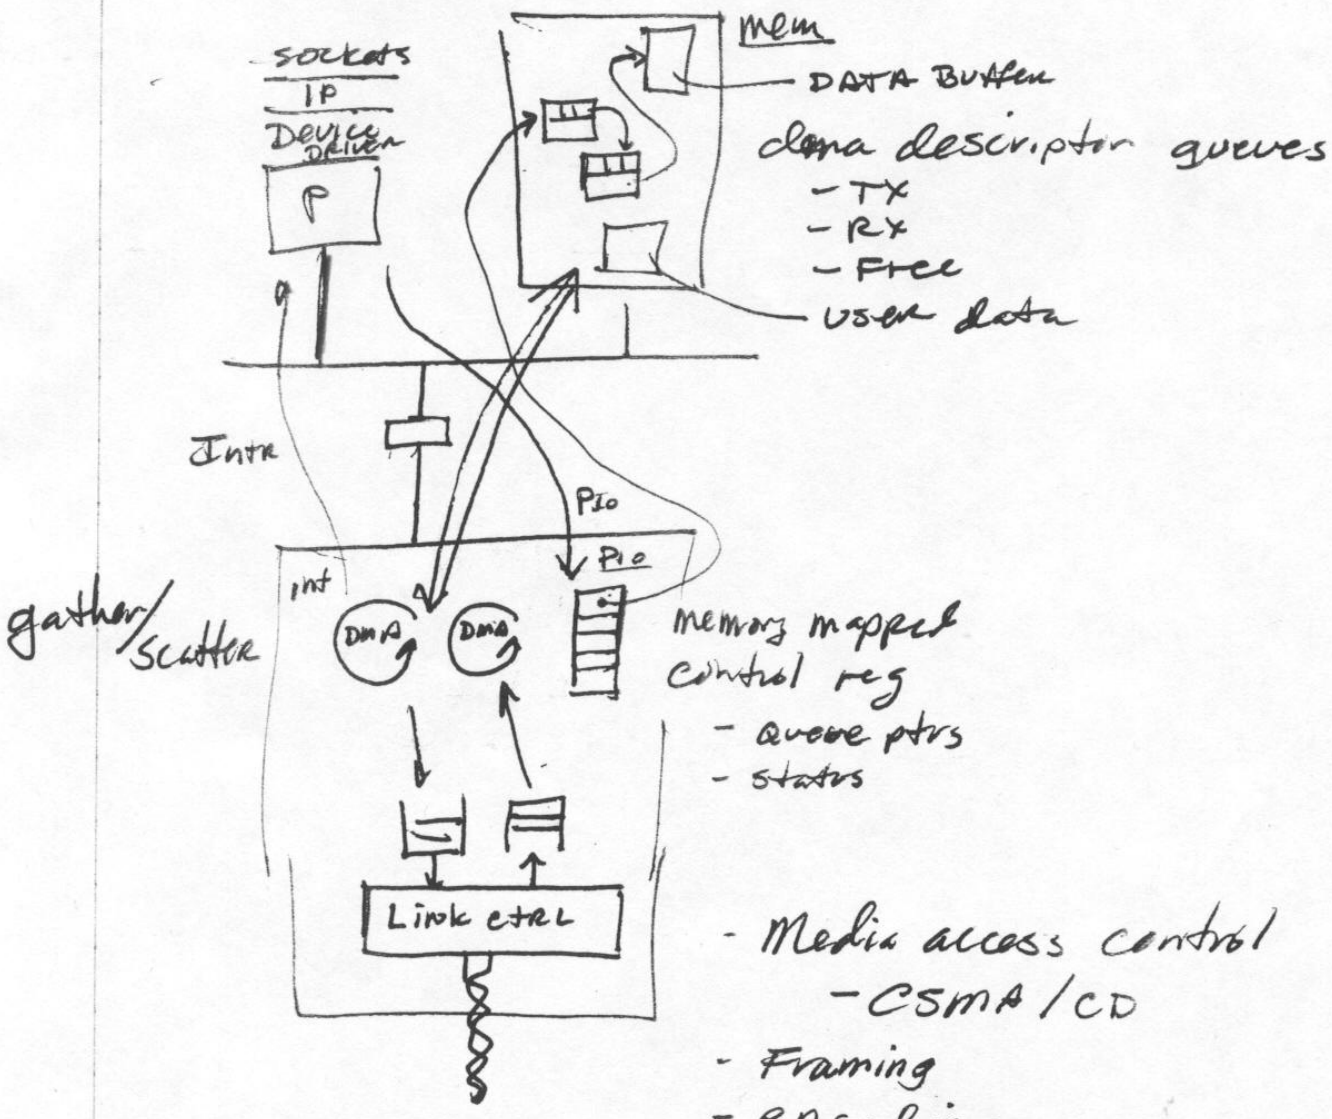

PARALLEL Machine Networks

Fig 4

MSG DRIVEN PROCESSORS

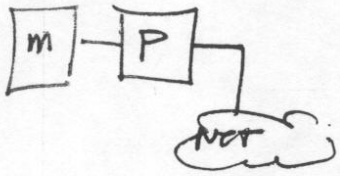

- DATAFLOW
- J-machine
- iWARP
- X-T

Shared Physical Addr

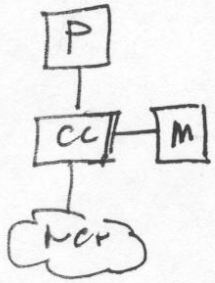

Local physical address space

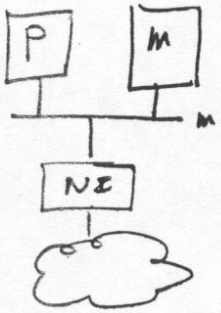

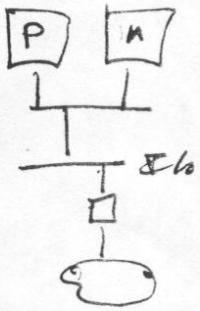

Simple DMA

Fig 5

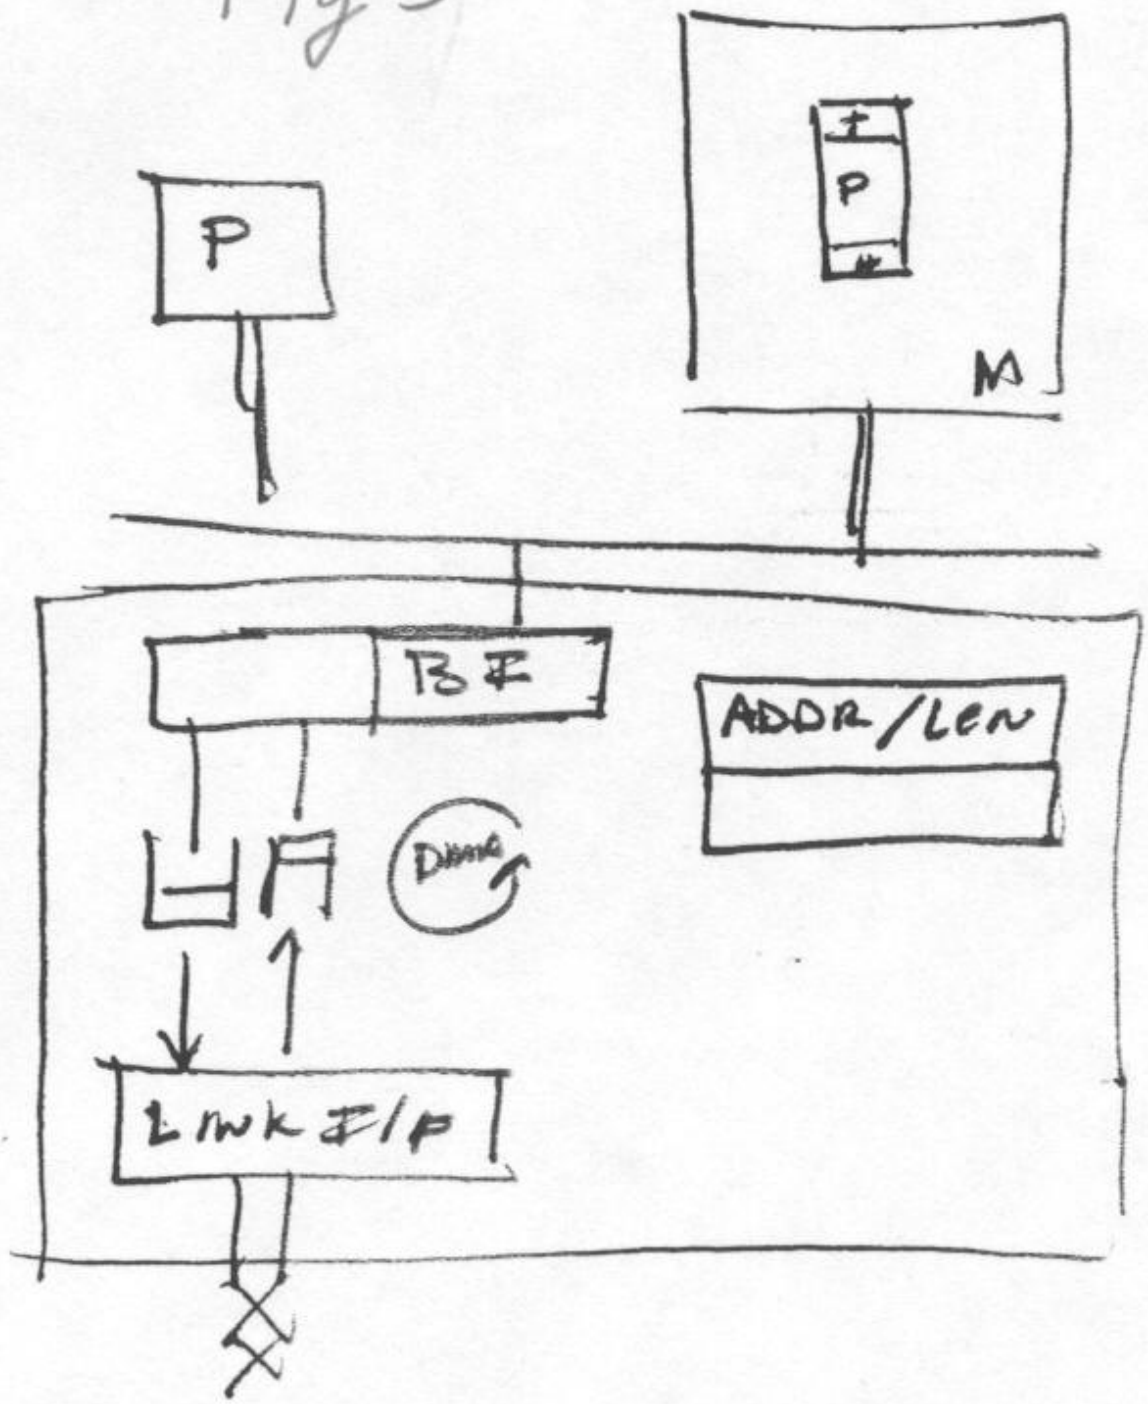

Fig 6

SIMPLE PIO

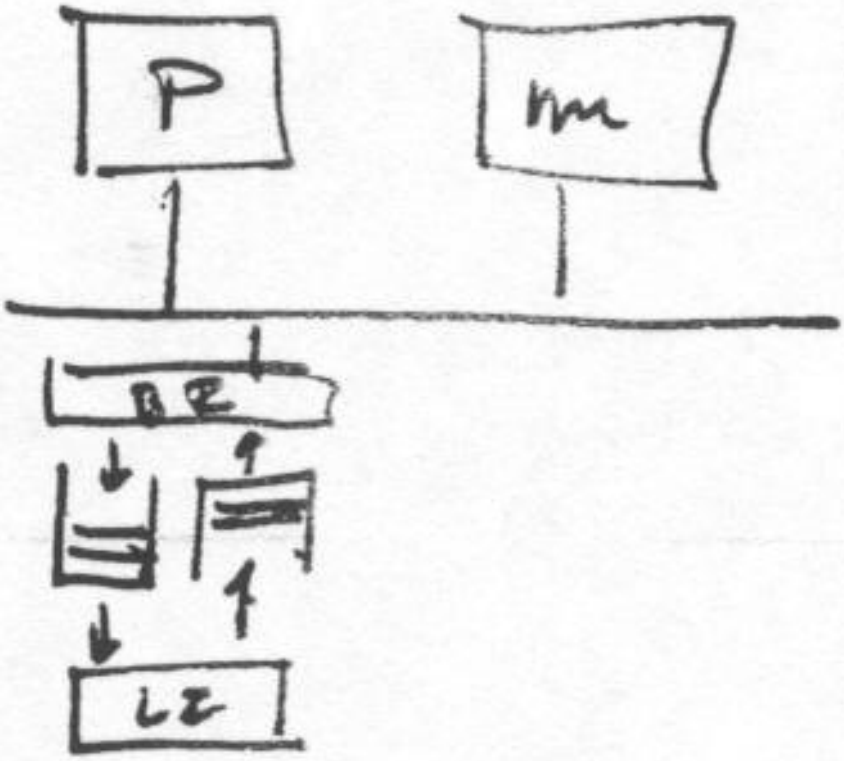

Fig 7

Physical Parallelism

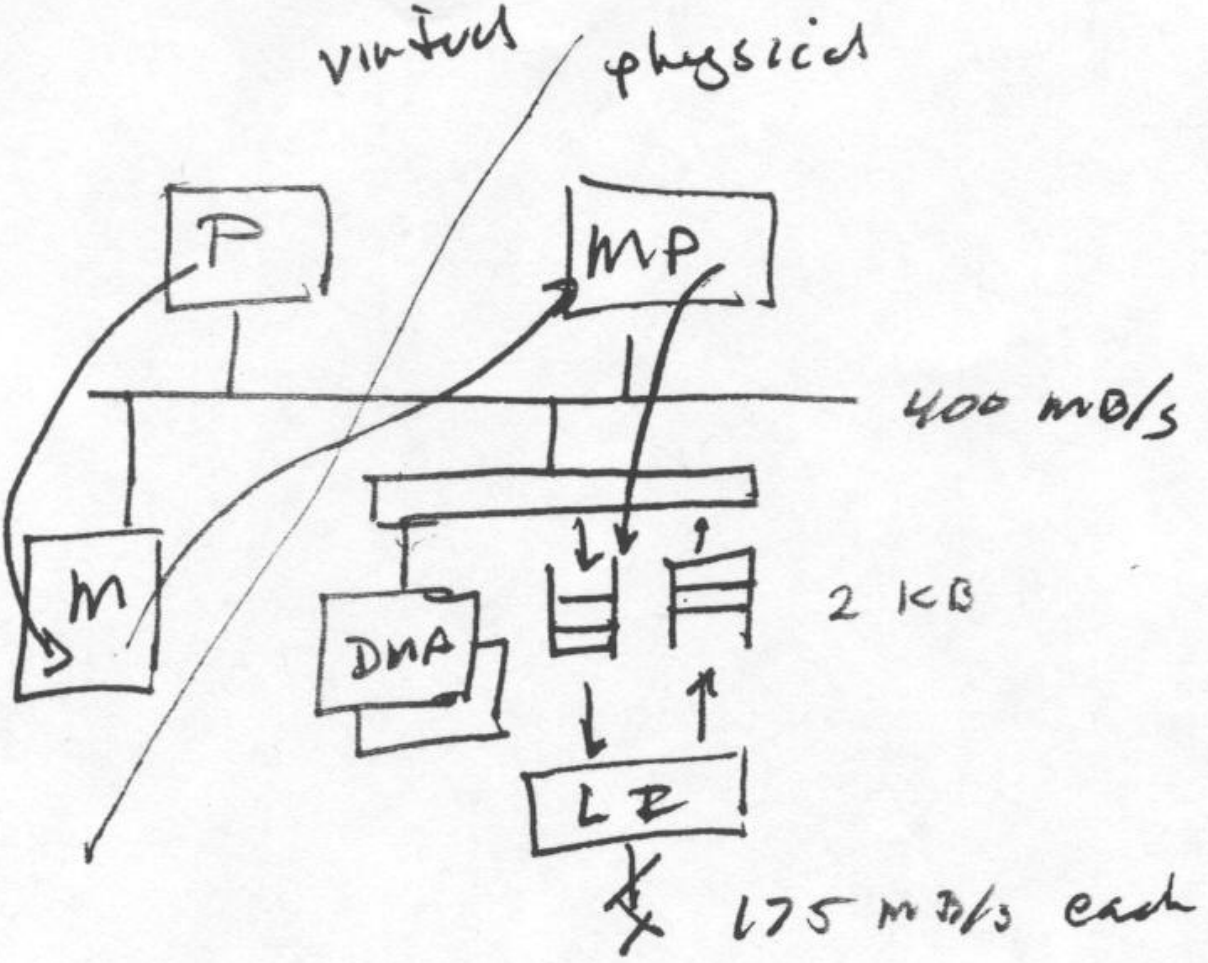

Fig 8

Virtual (Multithreaded) Parallelism

Ex: Meiko

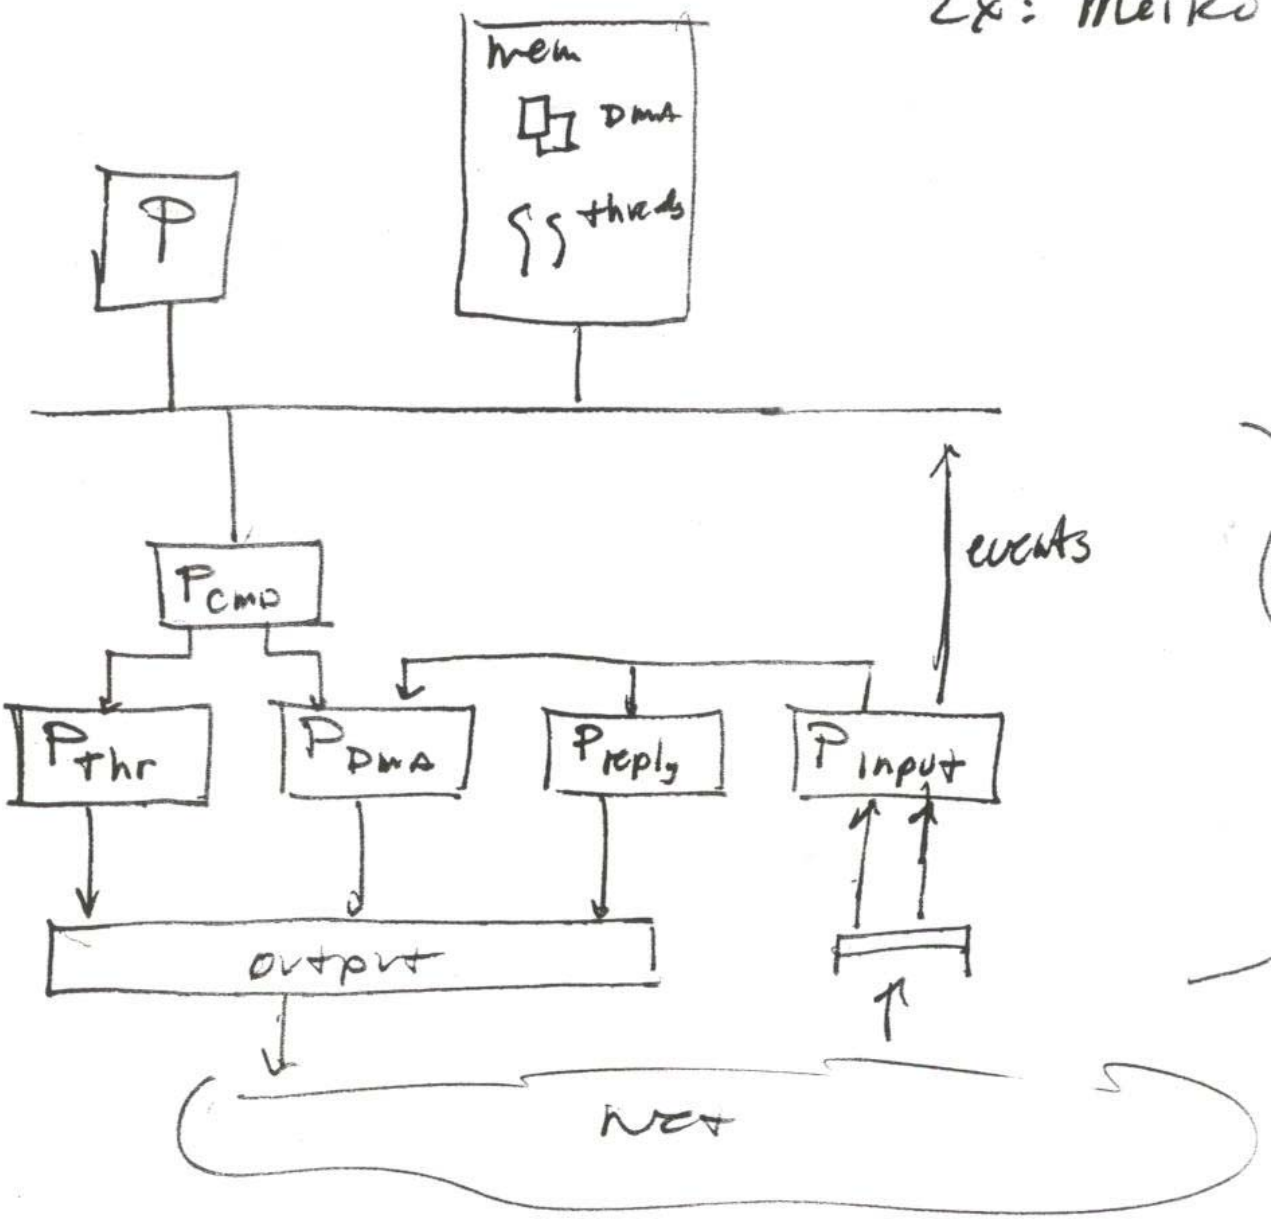

Fig 9

Stanford Flash Magic

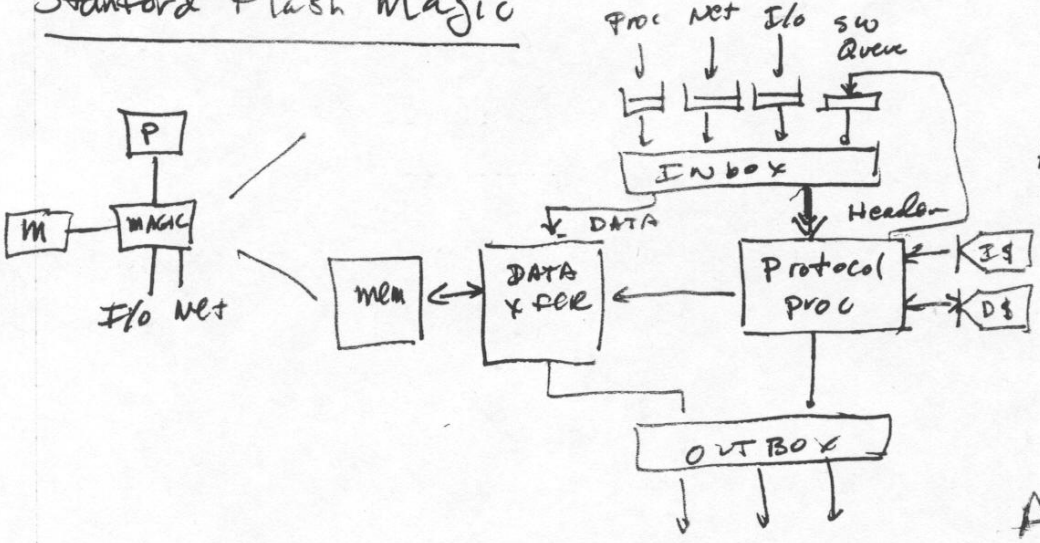

Classic View

Fig 10

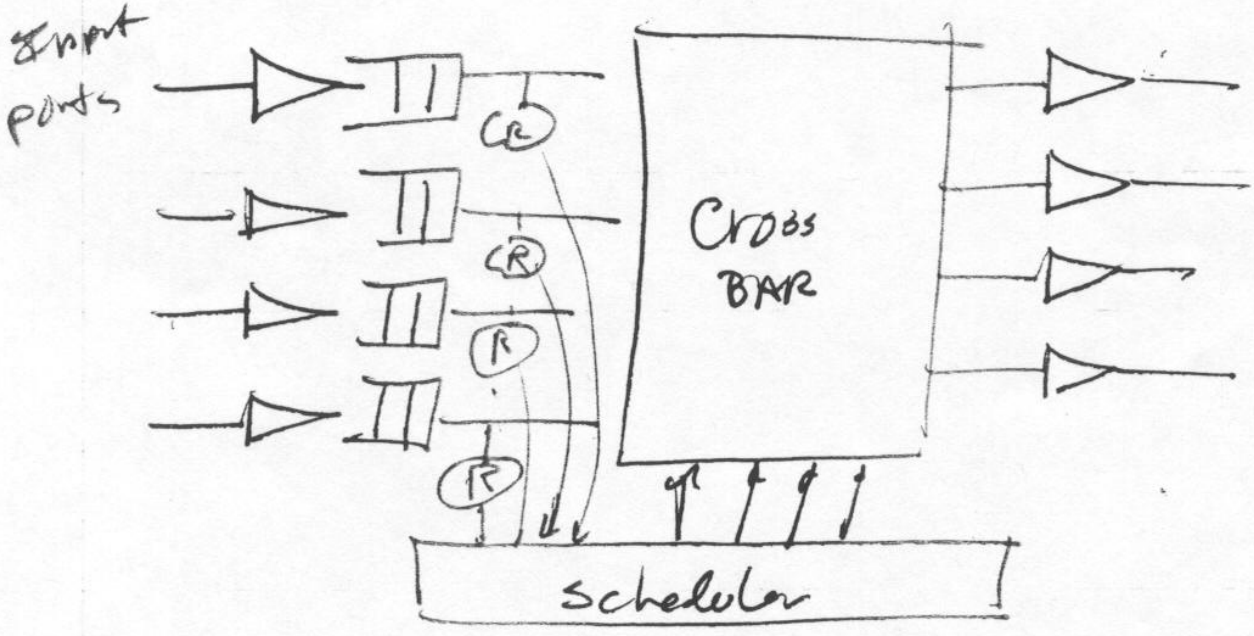

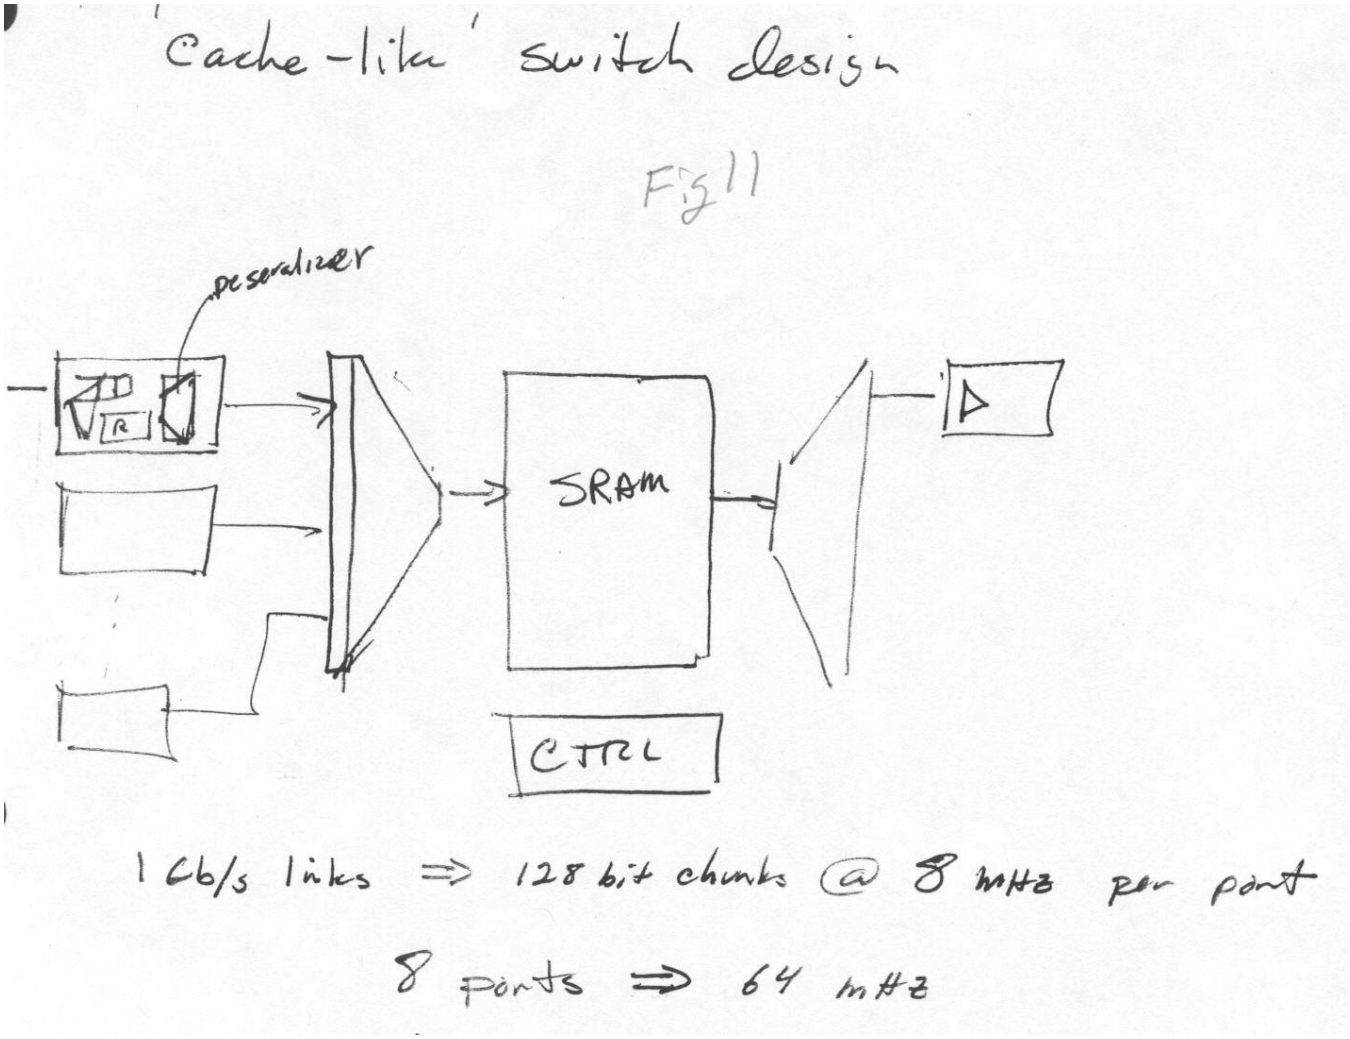

Scans of original lecture notes (Culler)

Network Interface Architecture CS252 2/21

Motivation

- Networks becoming the commonest I/O device
- Blow the machine apart
- Networked appliances

- Internet architectures are all about networks

Fig 1

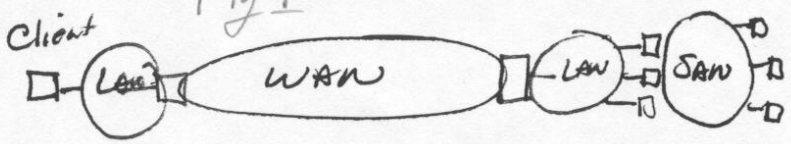

- Parallel Machine Networks ⇒ LAN/WAN Breakthrough
  - design approach
  - "killer switch"

- Different Kind of Architecture Challenge
  - juggling events & data movement

Today:

LANs vs PARALLEL MACHINES vs WIRELESS

What's a network system?

Fig 2

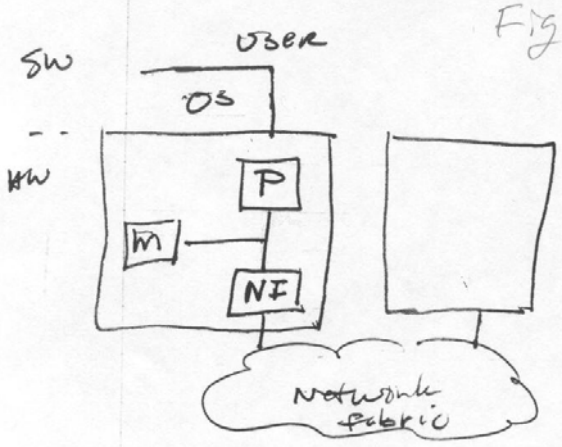

Issues

- How is communication integrated into node arch?
  - processing & storage
- What is the communication abstraction?
  - O/S calls, library, PIO, Load/store
- What is the NI architecture?
- How does data get from place to place?

Parallel Machines SANS LANS WANS

Performance framework

$$\text{Send overhead} + \text{Time of flight} + \frac{\text{Size}}{\text{BW}} + \text{Rcv overhead}$$

($\frac{\text{Size}}{\text{BW}}$: Transfer time)

If BW = Bw · $\frac{\text{Xfer Time}}{\text{Total Time}}$

Classical LAN

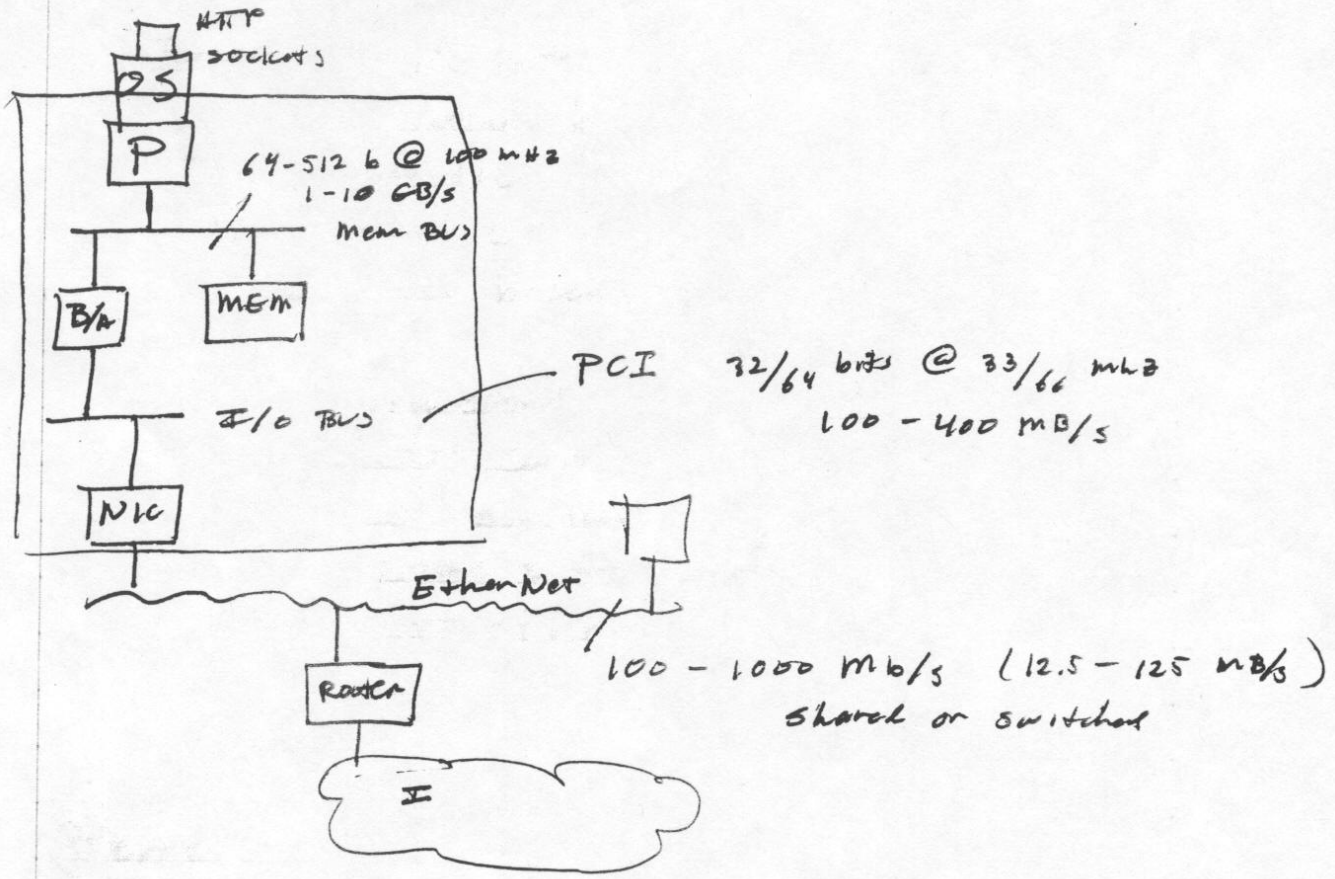

Key Abstractions

- digital symbols : bits
  - channel coding (manchester)
- Framing - information is serialized

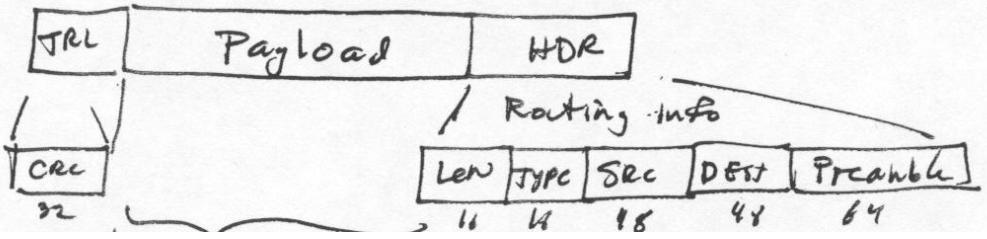

Kinds of routing

Determing the route
- destination based
- arithmetic
- virtual circuit
- source based

Mode of transport
- circuit switched
- packet switched
  - store & forward
  - cut-through

Performance

FB3

Typical LAN NIC

Issues:

- juggling events
- accessing data & control info
- copy (blind data receive)

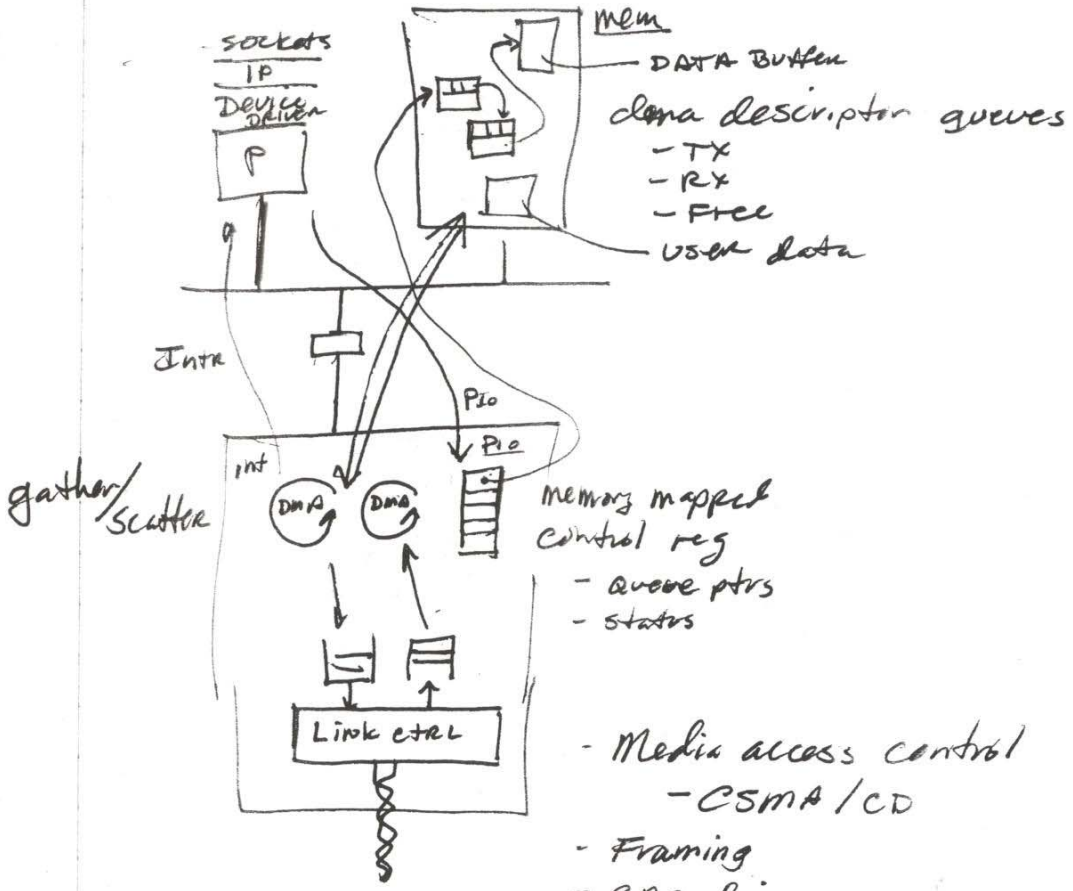

- Media access control
  - CSMA/CD
- Framing
- encoding

} completely standardized

Typical costs

Total time
Overhead
Transfer time

} ⇒ Reduce Overhead
- zero copy
User level Interface
- protection

Parallel Machine Networks

Fig 4

MSG DRIVEN PROCESSORS

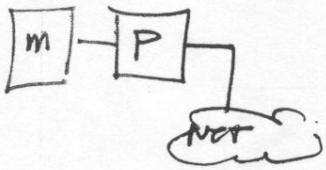

- DATAFLOW
- J-machine
- iWARP
- X-T

Shared Physical Addr

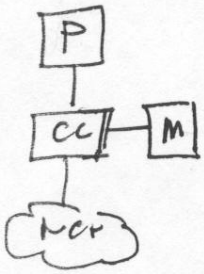

Local physical address space

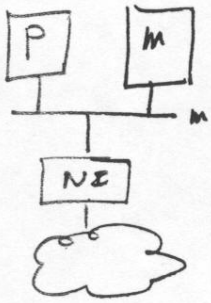

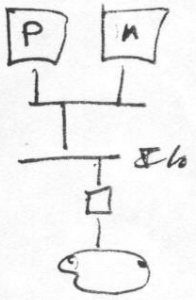

Very low latency network
High BW
Regular topology
Highly reliable

} ⇒ minimize overhead

SIMPLE DMA

Fig 5

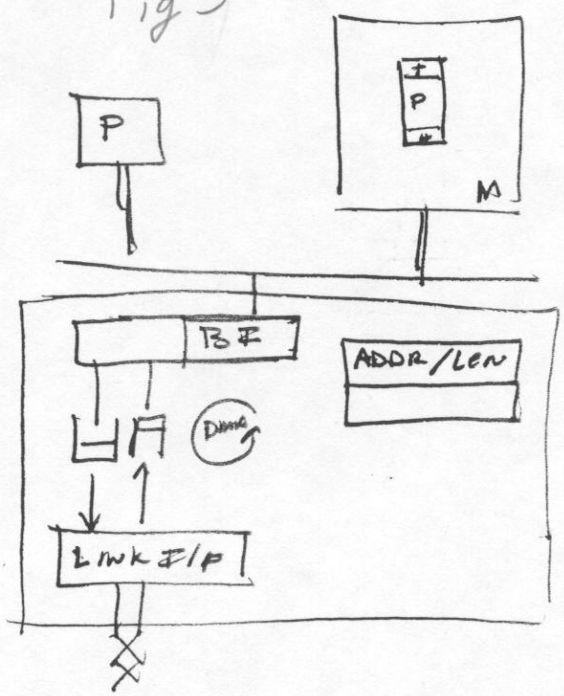

Ex: nCube, ipsc/2

- Simple Framing
- ~~Buffer~~
- few stores to start xfer
- Copy or swap receive data.
- virtual ↔ physical

Fig 6

SIMPLE PIO

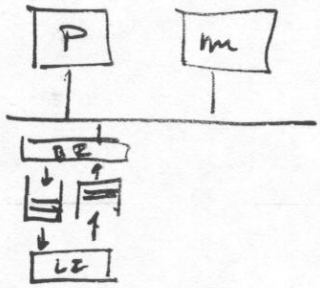

Ex: CM-5

- No copy
- Fast small msgs
- limited BW
- Proc. intensive
- OS manages NI
  - process mapping
  - context switch
- Host code dispatches on many conditions

Fig 7

Physical Parallelism

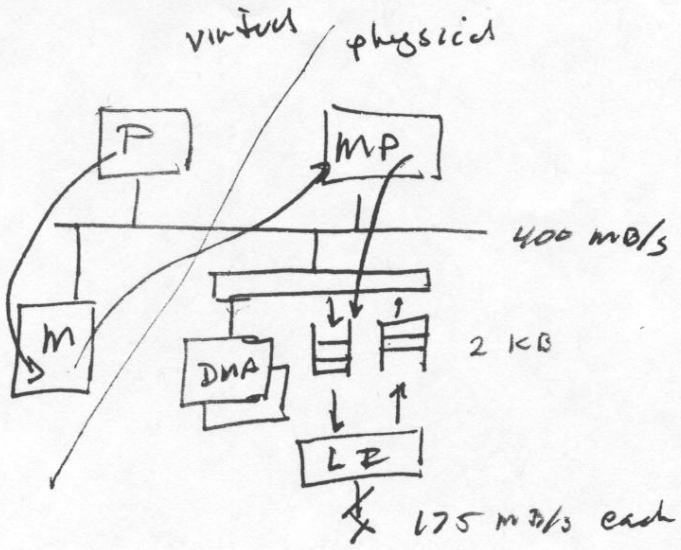

Ex: Intel Paragon

- Concurrency intensive
  - Proc events
  - NI events
  - DMA events
- BW intensive
  2-3 x bus crossing
- Synch overhead

Fig 8

Virtual (Multithreaded) Parallelism

Ex: Meiko CS-2

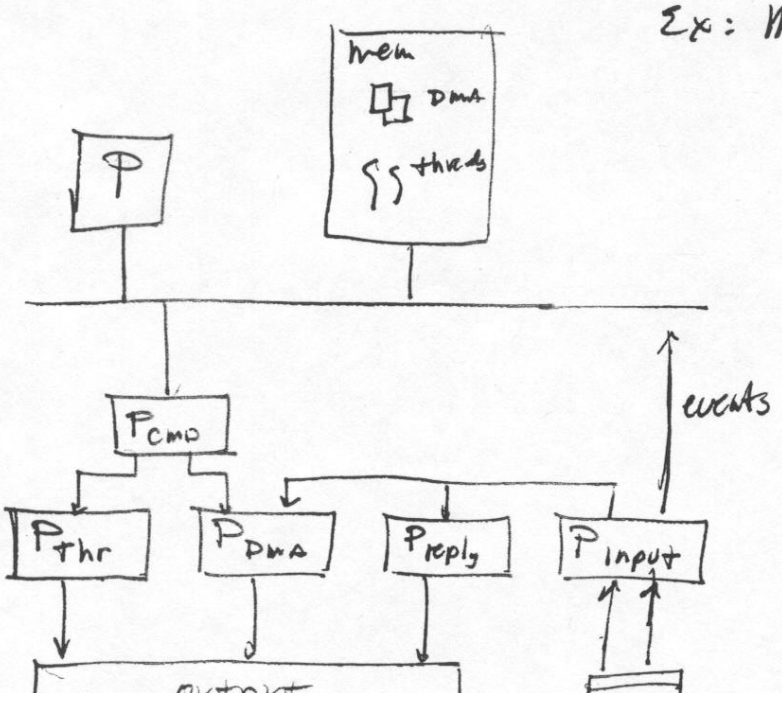

Implemented on single μ-coded engine

- zero data transfer cost

Network Transaction

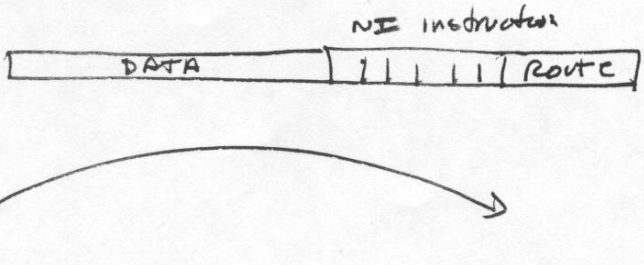

DMA TX - builds "store data" cmd
Thr - straight-line code
- abort on error (inc page fault)
Remote Rd - DMA reply

Flexible protocols

Fig 9

Stanford Flash Magic

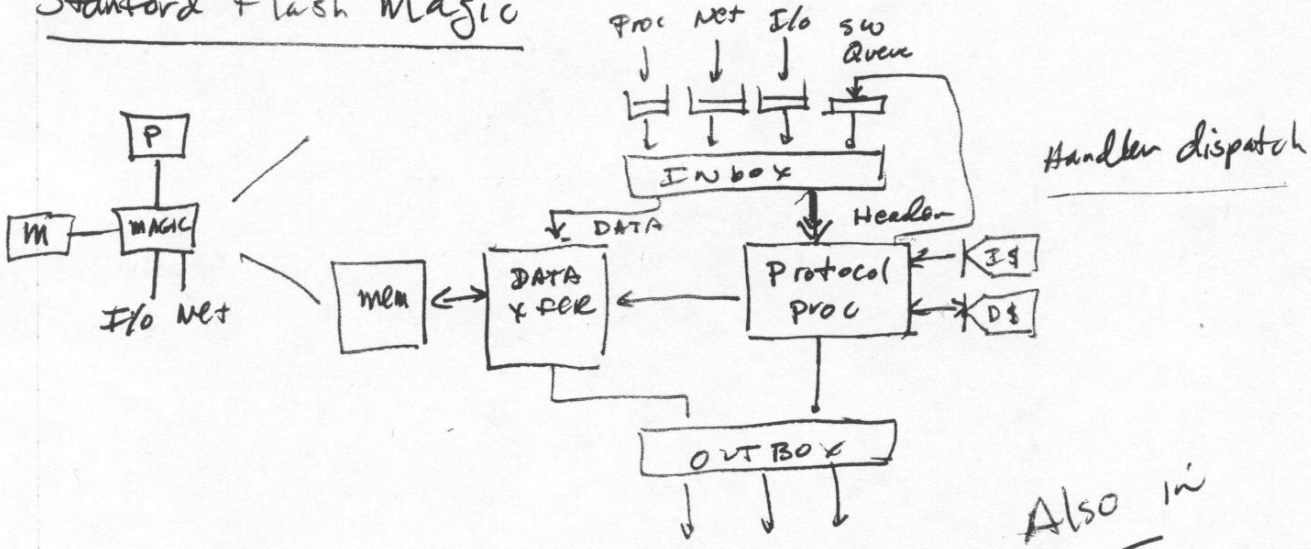

Routing / Switching

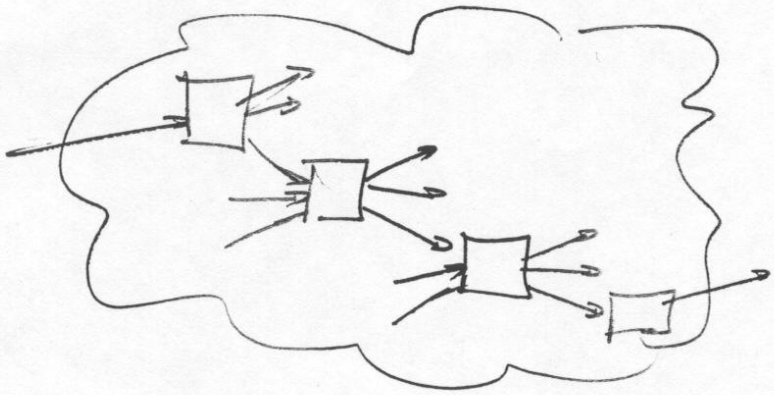

Classic View

Fig 10

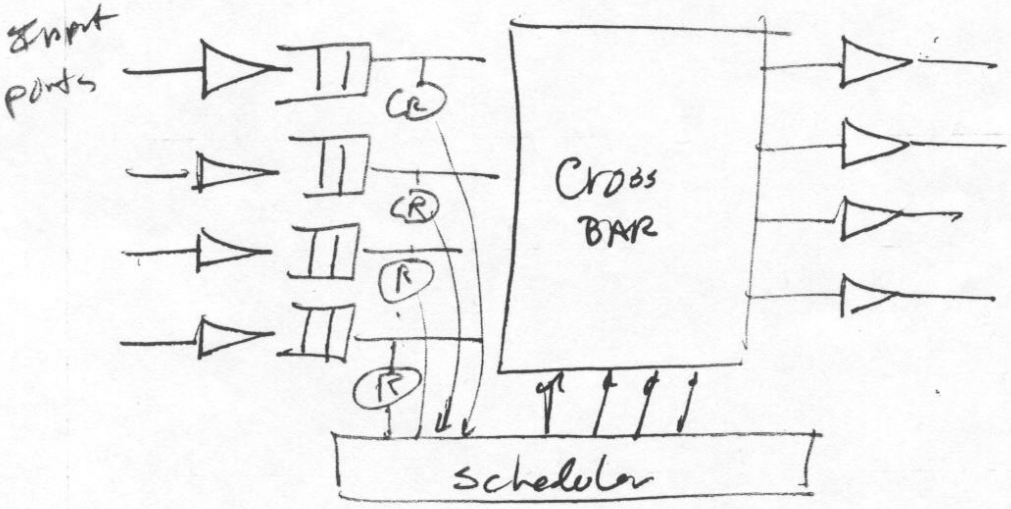

- Recognize Packet
- Compute output port
  - Arithmetic in mesh, butterfly, etc
  - select in source based
  - table look up in cloud or virtual circuit

'Cache-like' switch design

Fig 11

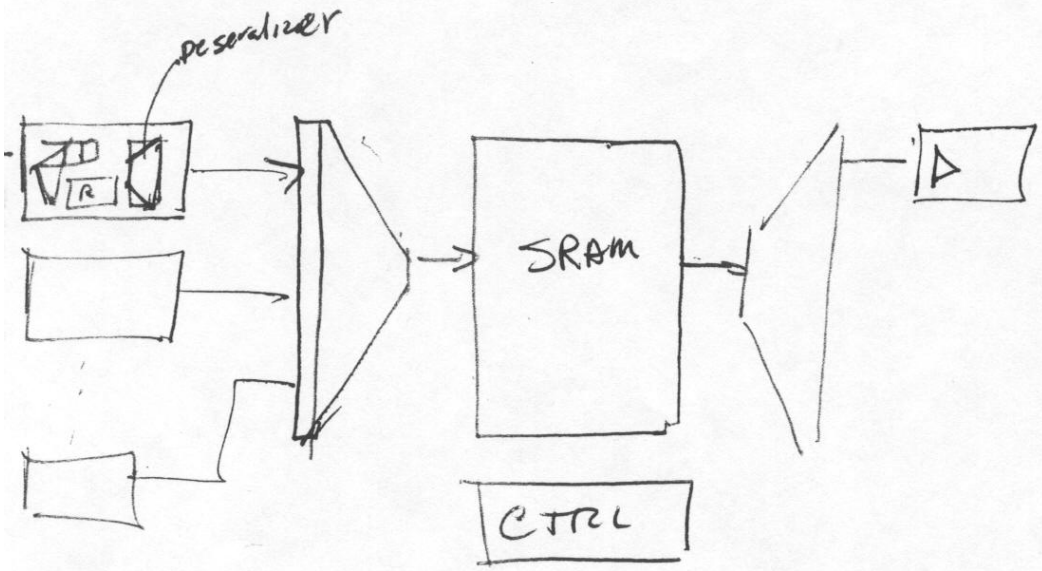

1 Gb/s links ⇒ 128 bit chunks @ 8 MHz per port

8 ports ⇒ 64 MHz

⋮

⇒ make routing & scheduling decision on many bits

- Ethernet header
- IP header
- HTTP request

What about wireless

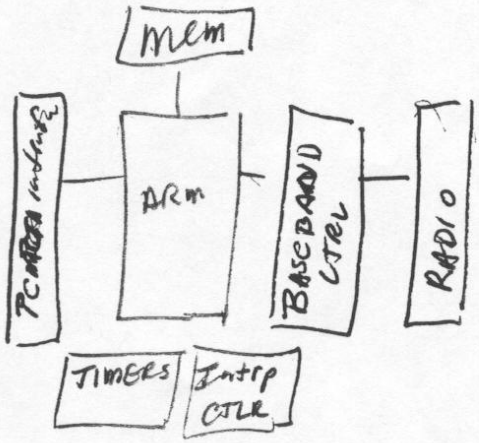

Present fairly conventional dma dq interface to host

Federation of devices to support low level operations

Programmable controller to implement
- MAC

sometimes a dedicated baseband controller
- Hopping, sequencing

Discussion

- Physical Parallelism vs Virtual Parallelism for concurrency intensive architecture
- HW support for protocol processing in deeply embedded networks
- Perspective on Jason's lecture