Active Disks For Large-Scale Data Mining and Multimedia

Erik Riedel
Parallel Data Laboratory
Carnegie Mellon University
www.pdl.cs.cmu.edu/Active

NSIC/NASD Workshop
June 8th, 1998

Outline

Opportunity

Active Disks

Applications

Performance Model

Speedups in Prototype

Opportunity

Large database systems - lots of disks, lots of power

System	Processing (MHz)		Data Rate (MB/s)	
	CPU	Disks	I/O Bus	Disks
Compaq TPC-C	4 x 200=**800**	*113* x 75=**8,475**	**133**	**1,130**
Microsoft Terraserver	4 x 400=**1,600**	*320* x 75=**24,000**	**532**	**3,200**
Digital 500 TPC-C	1 x 500=**500**	*61* x 75=**4,575**	**266**	**610**
Digital 4100 TPC-D	4 x 466=**1,864**	*82* x 75=**6,150**	**532**	**820**

- **assume disk offers equivalent of 75 host MHz**
- **assume disk sustained data rate of 10 MB/s**

Lots more cycles and MB/s in disks than in host

Advantage - Active Disks

Basic advantages of an Active Disks system

- **parallel processing - lots of disks**
- **bandwidth reduction - filtering operations common**
- **scheduling - little bit of computation can go a long way**

Appropriate applications

- **execution time dominated by data-intensive core**
- **allows parallel implementation of core**
- **small memory footprint**
- **small number of cycles per byte of data processed**

Example Application

Data mining - association rules [Agrawal95]

- **frequent sets summary counts**
- **count of *1-itemsets* and *2-itemsets***
- **milk & bread => cheese**
- **diapers & beer**

Prototype

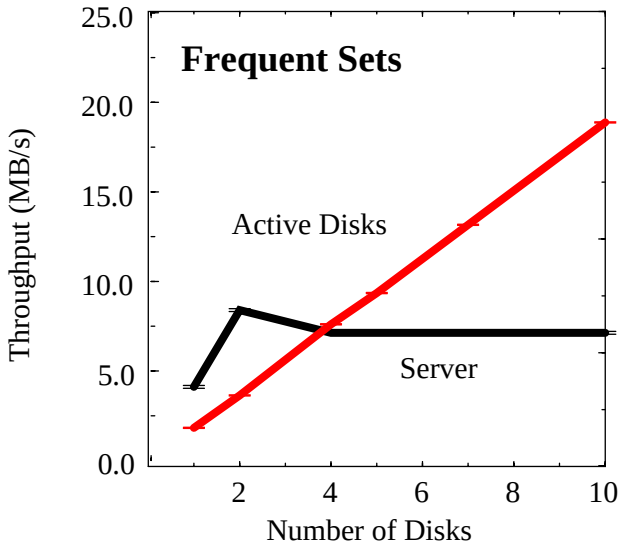

Scaling Up

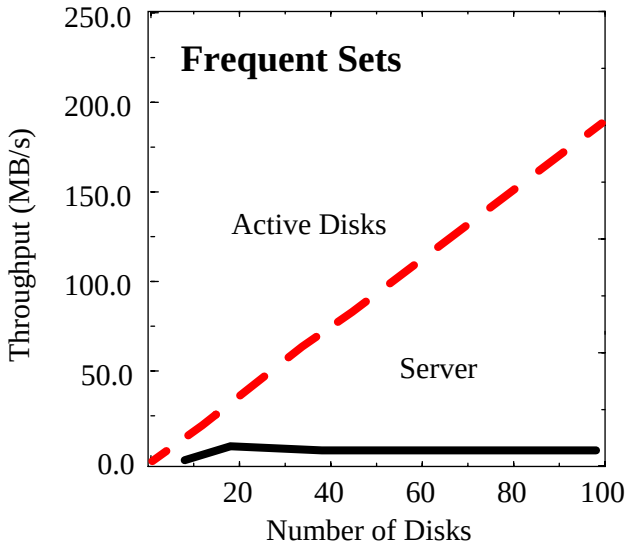

Performance Model

Scalable throughput

- **speedup** = (#disks)/(host-cpu-speed/disk-cpu-speed)
- (host-cpu/disk-cpu-speed) ~ 5 (two processor generations)
- **selectivity** = #bytes-input / #bytes-output

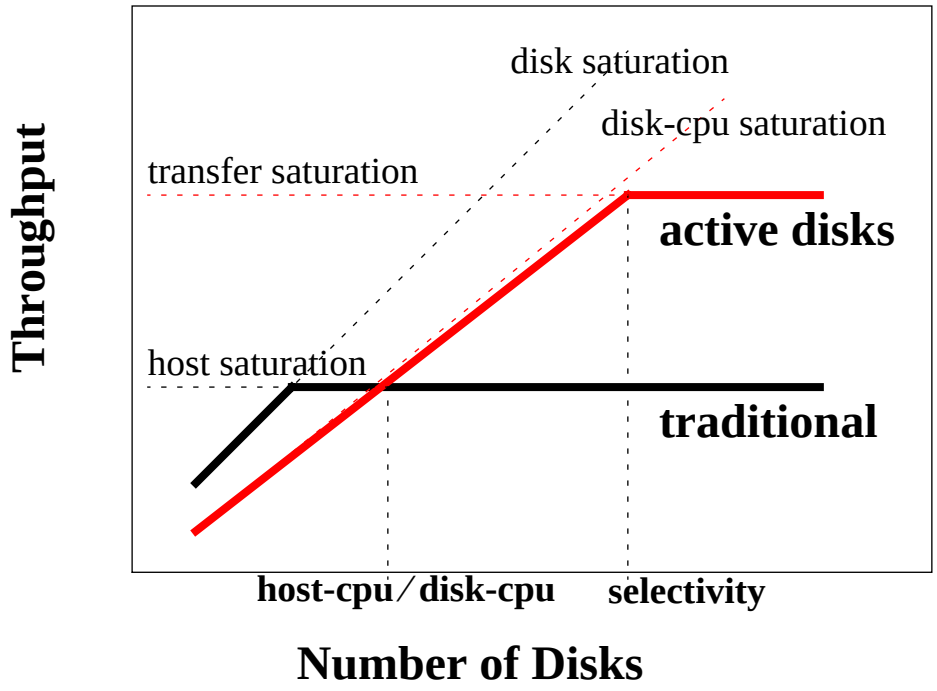

Additional Applications

Database - select

- extract records that match a particular predicate

Database - nearest neighbor search

- *k* records closest to input record
- with large number of attributes, reduces to scan

Multimedia - edge detection [Smith95]

- detect edges in an image

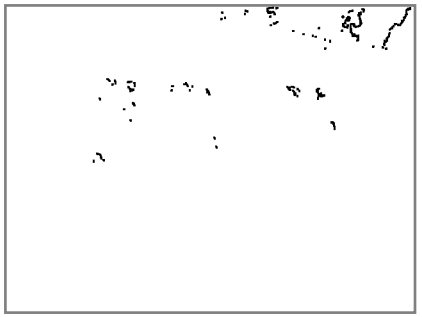

Multimedia - image registration [Welling97]

- find rotation and translation from reference image

Traditional Server

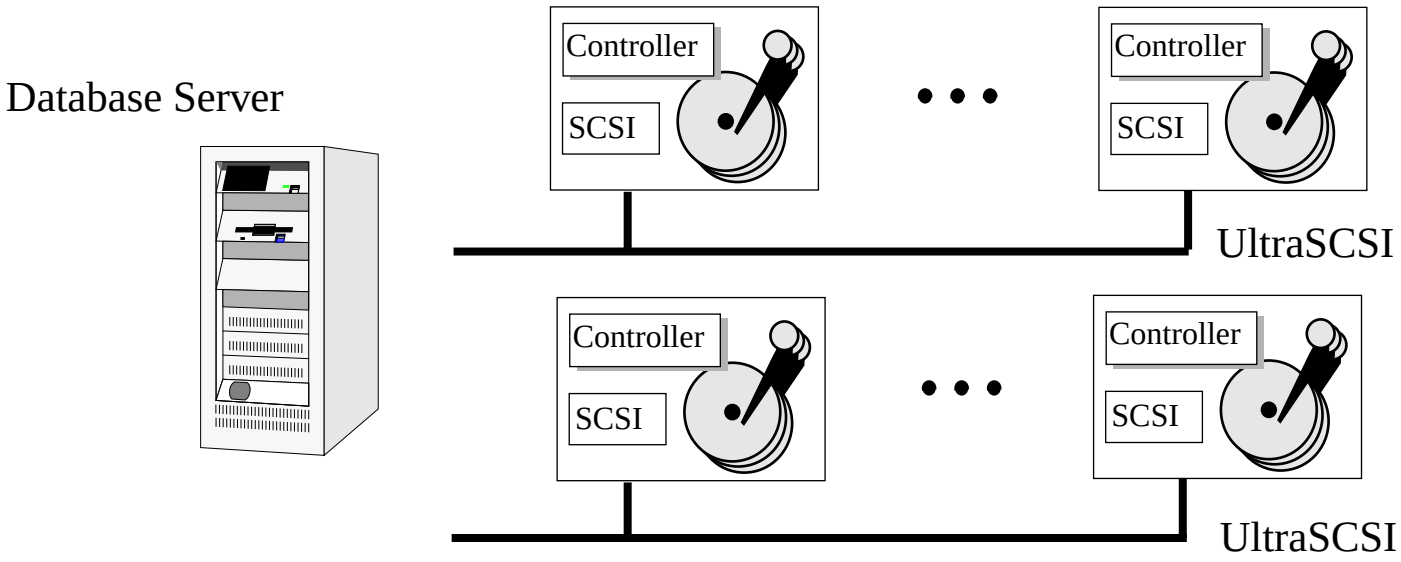

Digital AlphaServer 500/500

- 500 MHz, 256 MB memory
- disks - Seagate Cheetah
- 4.5 GB, 10,000 RPM, 11.2 MB/s

Server with Active Disks

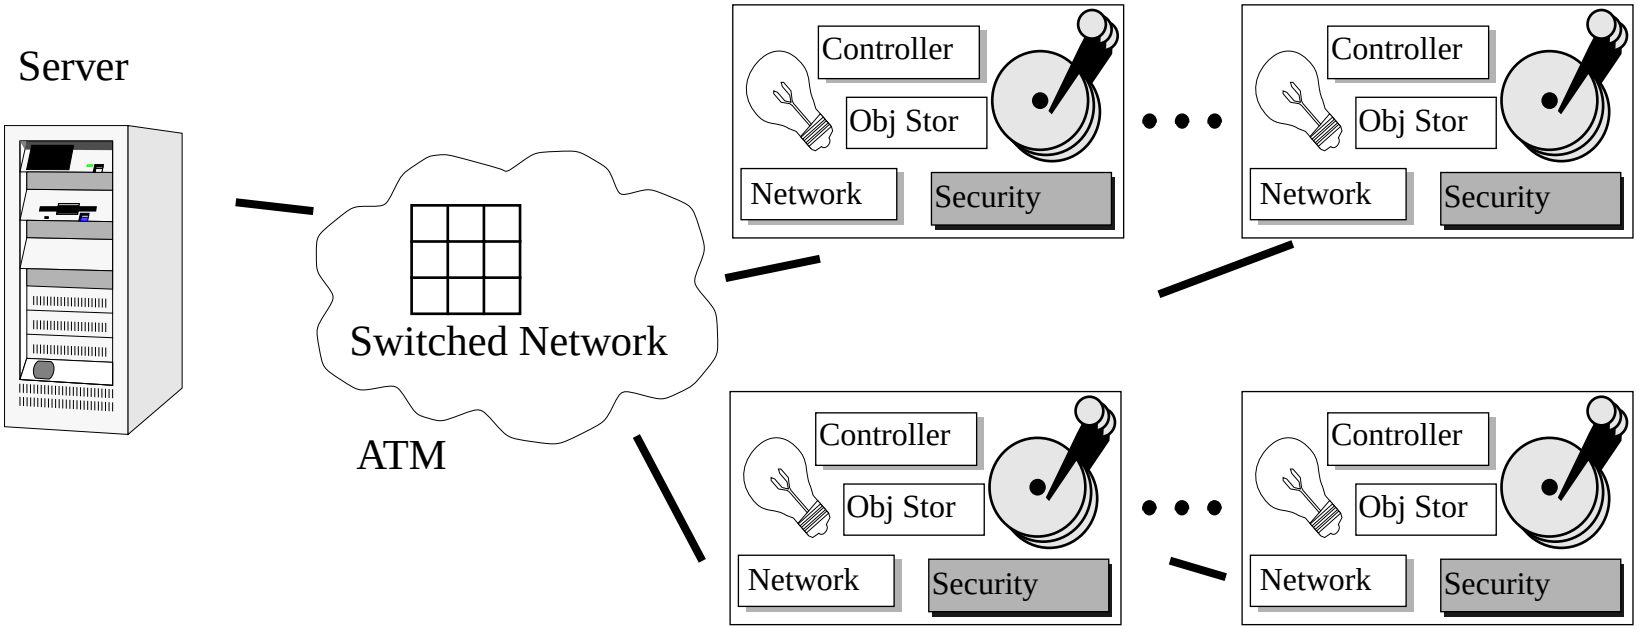

Prototype Active Disks

- Digital AXP 3000/400 workstation
- 133 MHz, software NASD prototype
- Seagate Medallist disks

Performance with Active Disks

application	input	computation (inst/byte)	throughput (MB/s)	memory (KB)	selectivity (factor)	bandwidth (KB/s)
Select	m=1%	7	28.6	-	100	300
Search	k=10	7	28.6	72	80,500	0.1
Frequent Sets	s=0.25%	16	12.5	620	15,000	1
Edge Detection	t=75	303	0.67	1776	110	2
Image Registration	-	4740	0.04	672	150	2

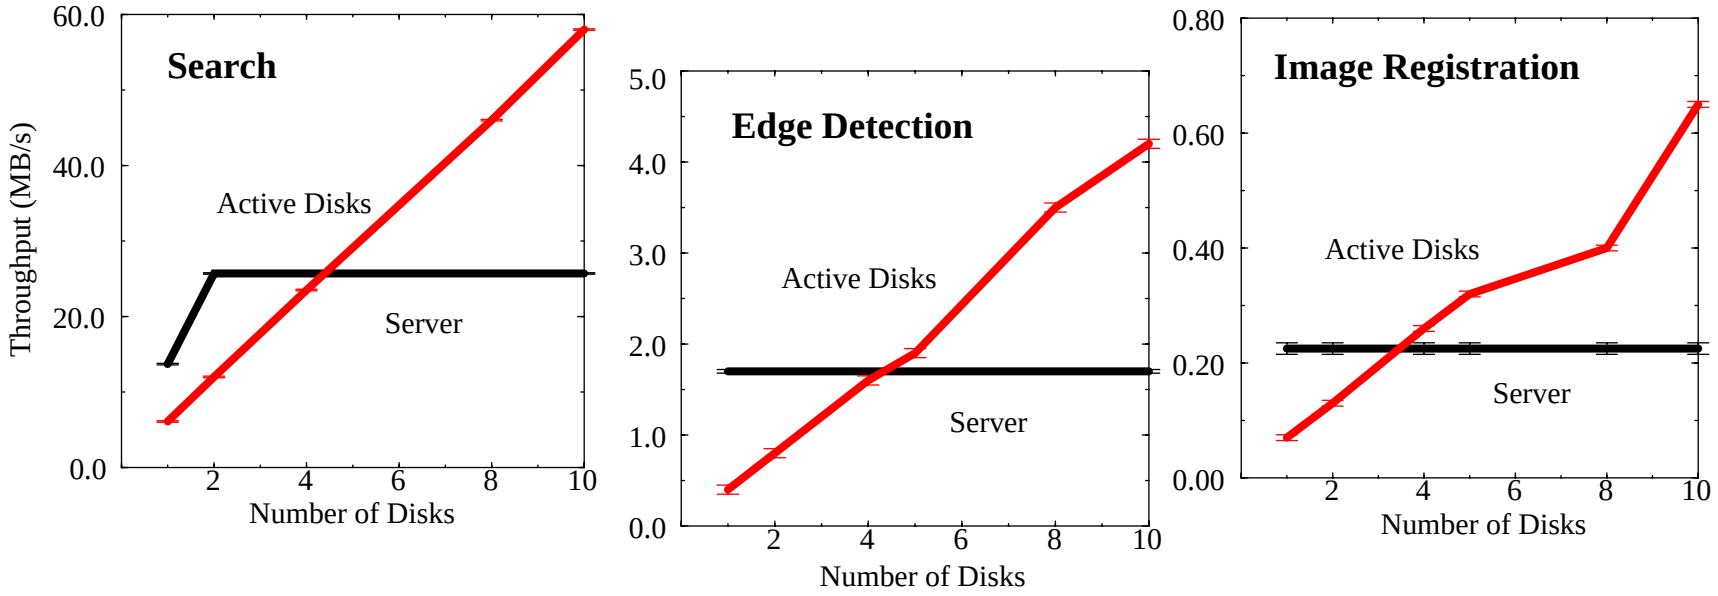

Why Isn't This Parallel Programming?

It is

- parallel cores
- distributed computation
- serial portion needs to be small

Disks are different

- must protect the data
- must continue to serve demand requests
- memory/CPU ratios driven by cost, reliability, volume
- come in boxes of ten

Opportunistically use this power

- e.g. data mining possible on an OLTP system

Summary

Technology trends provide the opportunity

- "excess" cycles
- large systems => lots of disks => lots of power

Dramatic benefits possible

- data mining and multimedia
- parallelism, selectivity

Scales down as well as up

- about 4 disks match a single host processor
  (two VLSI generations)
- factors of 2-3 speedup with "PC" server and 10 disks