Computer Systems Performance Analysis and Benchmarking (37-235)

Analytic Modeling

Simulation

Measurements / Benchmarking

Lecture/Assignments/Projects:
Dr. Christian Kurmann

Textbook:
Raj Jain, "The Art of Computer Systems Performance Analysis", 1991 Wiley & Sons, New York

Topic of Today:

- **Self Scaling Bechmarks**
  SIGMETRICS Paper 1993
  by Peter M. Chen

What's wrong w. I/O benchmarks

- Become quickly obsolete
  i.e. have no scaling strategy
- Do not stress the I/O system
- Do not help to understand I/O perf.
- Comparison for one workload only
  i.e. don't allow conclusions about relative performance.
- Narrow application range
  i.e. TPC-B is only debit/credit

Proposal

- *Self-scaling* benchmark that dynamically adjusts workload to system being measured.
- *Predicted* performance estimates the performance of unmeasured workloads based on a small set of measured WLs.
- Works now and in the future.

Ideal I/O Benchmark:

- should be I/O limited (I/O percentage)
- helps to understand the system performance
- scales over a wide range of current and future machines
- allows fair comparisons across machines
- is relevant for a wide range of applications
- is tightly specified.

Current State of the Art (1993)

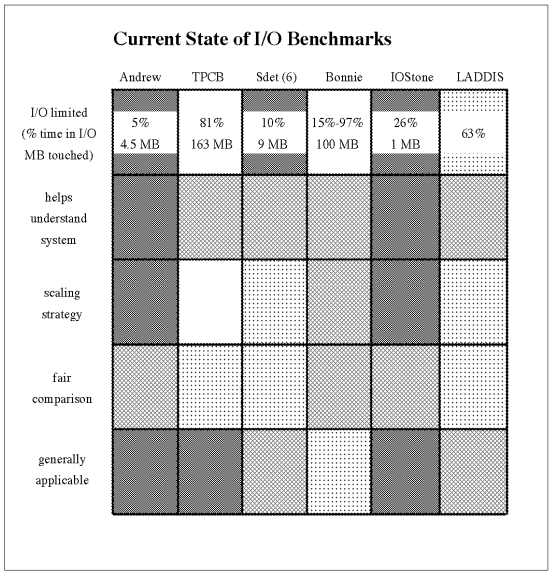

Workload Model of Self-Scaling Benchmarks

- uniqueBytes
  - total working set size
- sizeMean
  - average size of I/O request normal distribution with C.O.V of 1
- readFrac
  - percentage of reads vs. writes
- seqFrac
  - probability that next request follows the prior request (sequential access)
- processNum
  - multithreading, number of processes issuing I/O (concurrency)

How accurately can we model real workloads?

Representativeness of workload

Application	Throughput	Read Response Time	Write Response Time	Average Response Time
Sort	.20 MB/s	19.7 ms	1.6 ms	11.7 ms
Workload Model	.20 MB/s	20.0 ms	1.9 ms	11.0 ms
TPC-B	.13 MB/s	25.6 ms	1.3 ms	14.0 ms
Workload Model	.13 MB/s	22.1 ms	1.6 ms	12.3 ms

The Parameter Space

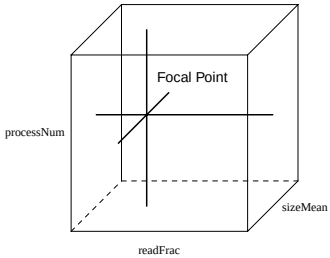

Most benchmarks report perf. for a single WL. Better are multiple WLs in a graph.

Self-scaling benchmarks produce a set of graphs for each parameter.

Terms

- Single Parameter Graphs
- Focal Point
- Focal Vector

Example

- uniqueBytes= 21 MB
- sizeMean= 10 KB
- readFrac= 0
- processNum= 1
- seqFrac= 0

Full Fractional Design

- 6^5 = 8000 points @ 10min = 2 Months

Reduced Design

- 30 points @ 10 min = 5 hours

Basic Version of Benchmark

- Take focal point somewhere at 75% of maximal performance of each factor
- Iterative approach to find focal vector.

Problems:

- Finding focal point iteratively is slow
- Focal points at outer limits
- Focal points set arbitrarily
- Focal point in an unstable area
- Not aware of plateaus

Improved Version of Benchmark

- Aware of plateaus

  multiple focal points for uniqueBytes
- Focal points midway e.g. 0.5 for [0..1]

Scaling all Parameters [Sparc1]

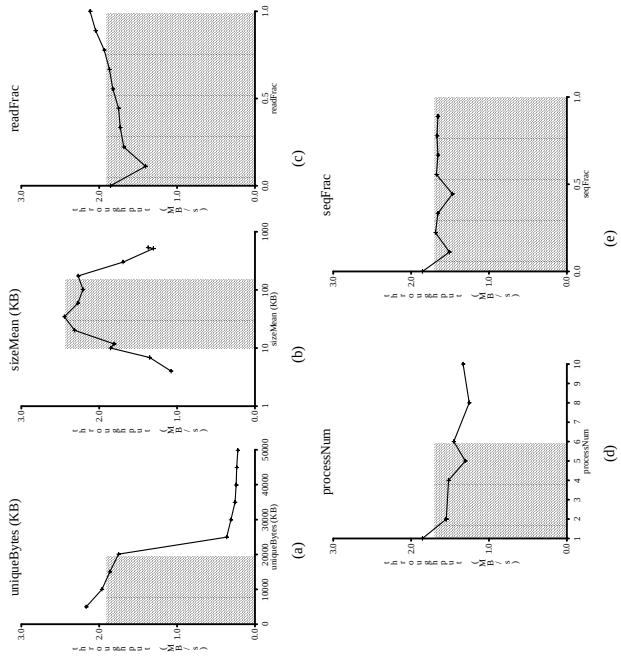

Plateau 1 [Sparc1]

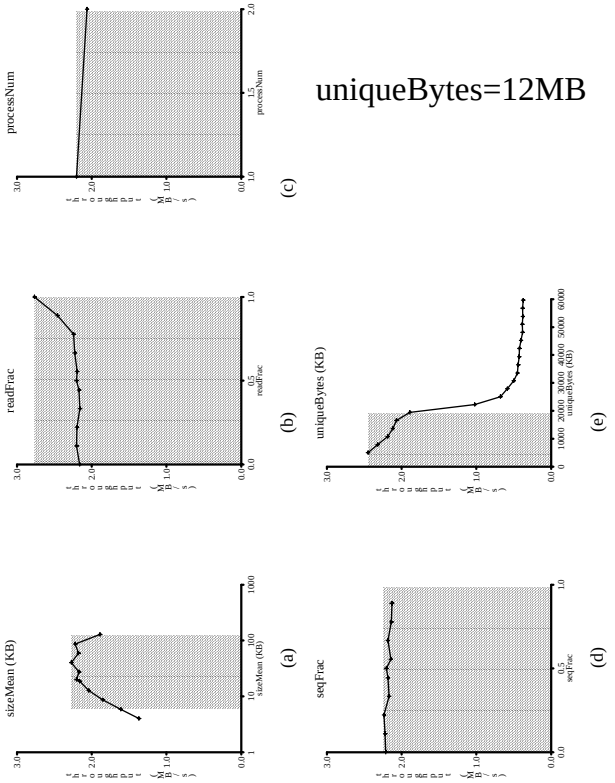

uniqueBytes=12MB

Plateau 2 [Sparc1]

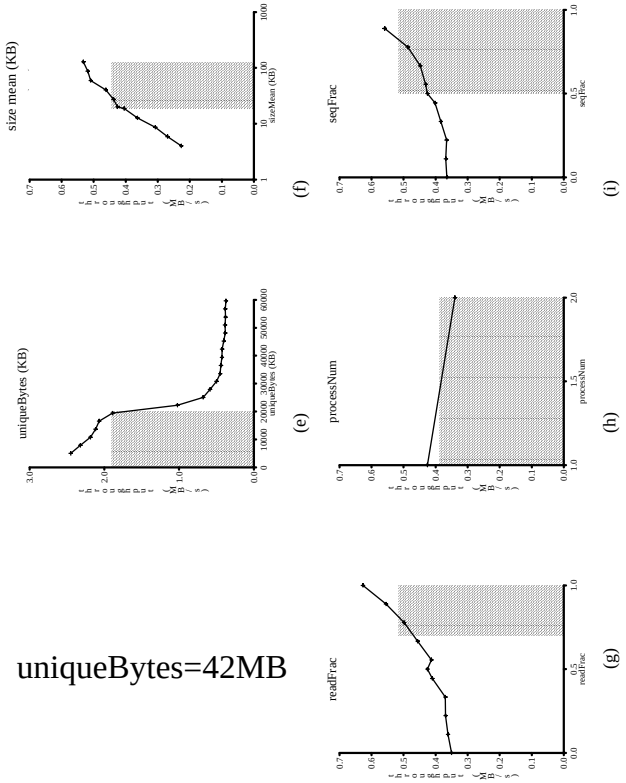

uniqueBytes=42MB

How to allow for comparisons?

Simple prediction component for any workload assumed

Big assumption

- Shape of single parameter curve does not change with focal point of other parameters.
- Simple linear interpolation
- Multiplicative linear interpolation for multiple parameters

Clean evaluation of the error in the paper.

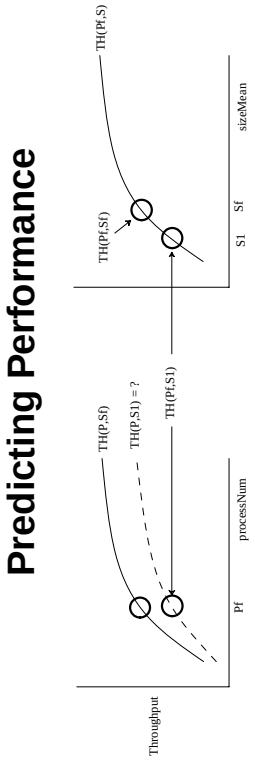

Accuracy of Prediction [Sparc1]

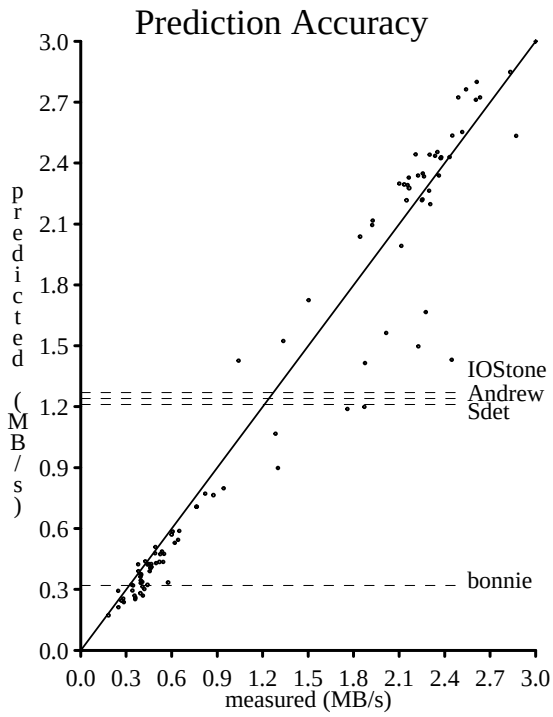

(a)

Error of Prediction [Sparc1]

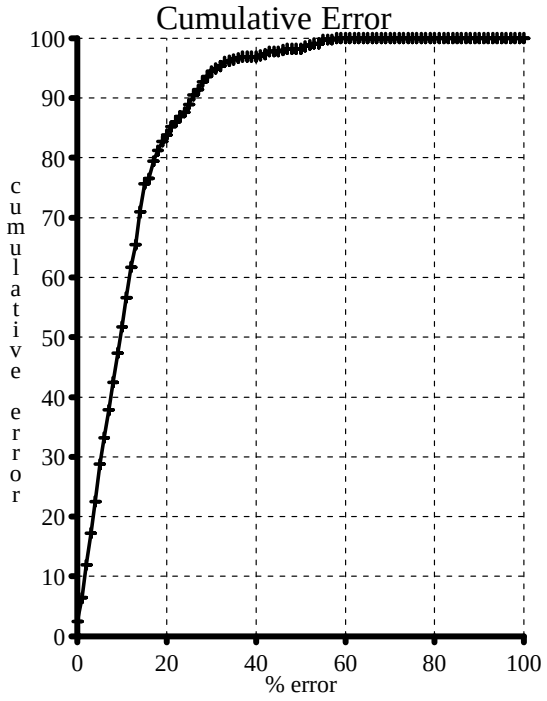

(b)

Parameters with error

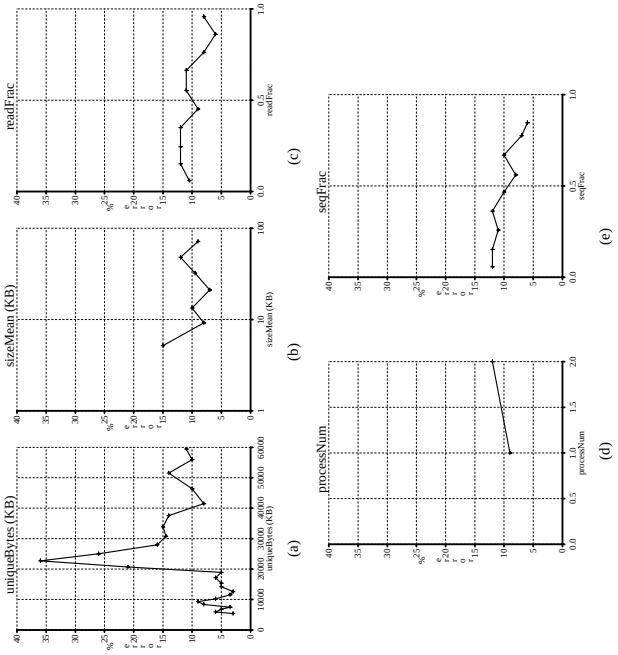

Discussion of Machines

System Name	SPARCstation 1+	DECstation 5000/200	HP 730
Year Released	1989	1990	1991
CPU	SPARC	MIPS R3000	PA-RISC
SPECmarks	8.3	19.9	76.8
Disk System	CDC Wren IV	3 disk (Wren) RAID 0	HP 1350SX
I/O Bus	SCSI-I	SCSI-I	Fast SCSI-II
Mem. Bus Peak Speed	80 MB/s	100 MB/s	264 MB/s
Memory Size	28 MB	32 MB	32 MB
Operating System	SunOS 4.1	Sprite LFS	HP/UX 8.07

System Name	Convex C240	Solbourne 5E/905
Year Released	1990 (???)	???
CPU	C2 (4 processors)	SPARC (??? processors)
Speed	220 MIPS	???
Disk System	4 DKD-502 RAID 5	4
I/O Bus	IPI-2	IPI
Memory Bus Peak Speed	200 MB/s	128 MB/s ??
Memory Size	1024 MB	384 MB
Operating System	ConvexOS 10.1 (BSD derived)	SunOS 4.1A.2 (revised)

Storage Hierarchies

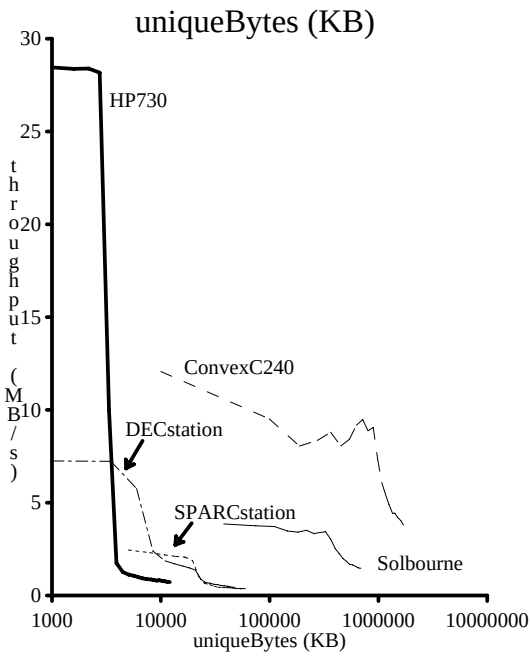

Plateau 1 [Sparc1]

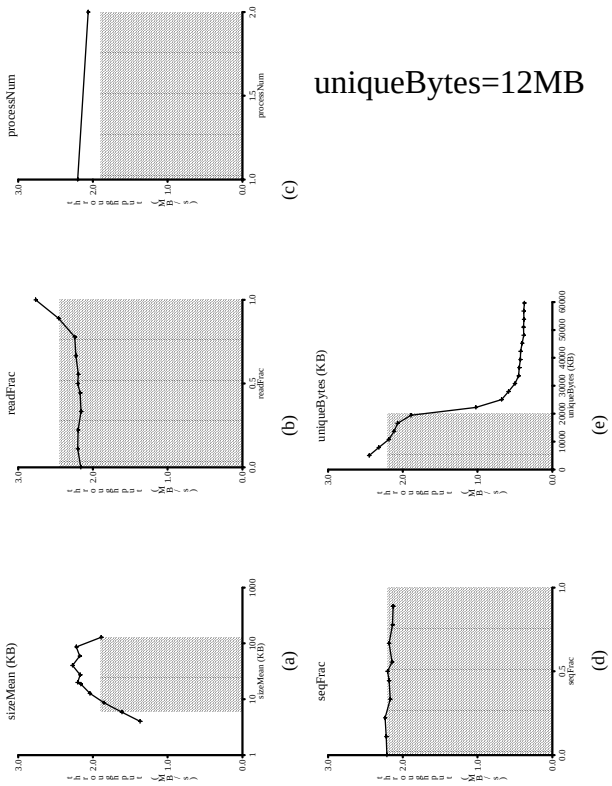

Plateau 2 [Sparc1]

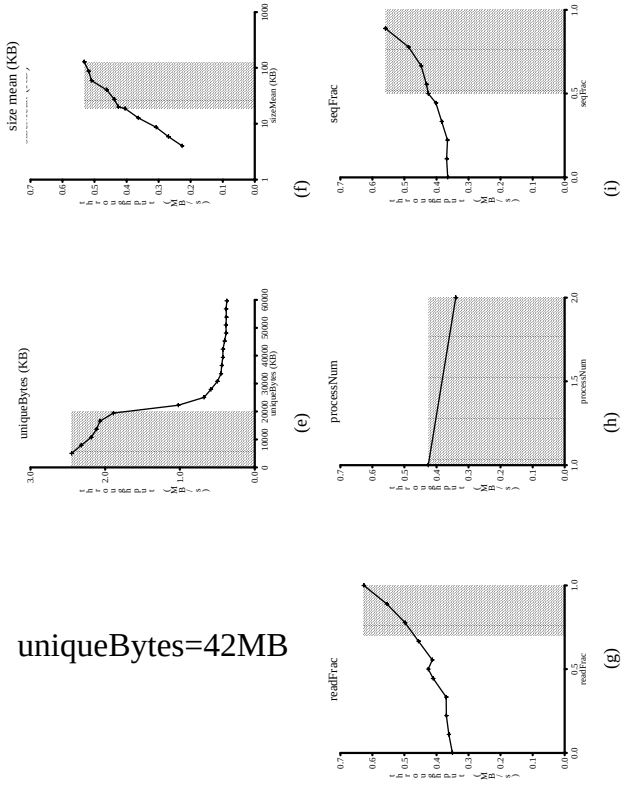

Plateau 1 [DEC 5000]

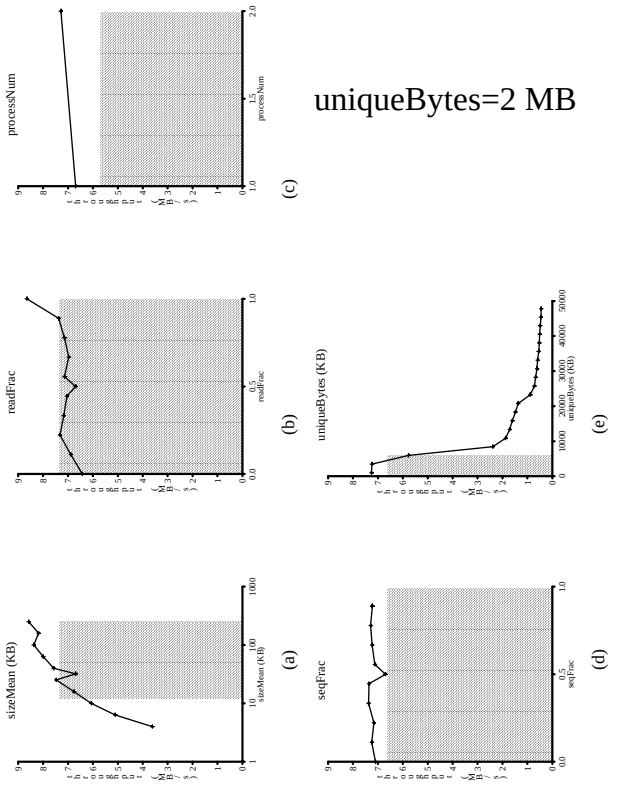

Plateau 2 [DEC 5000]

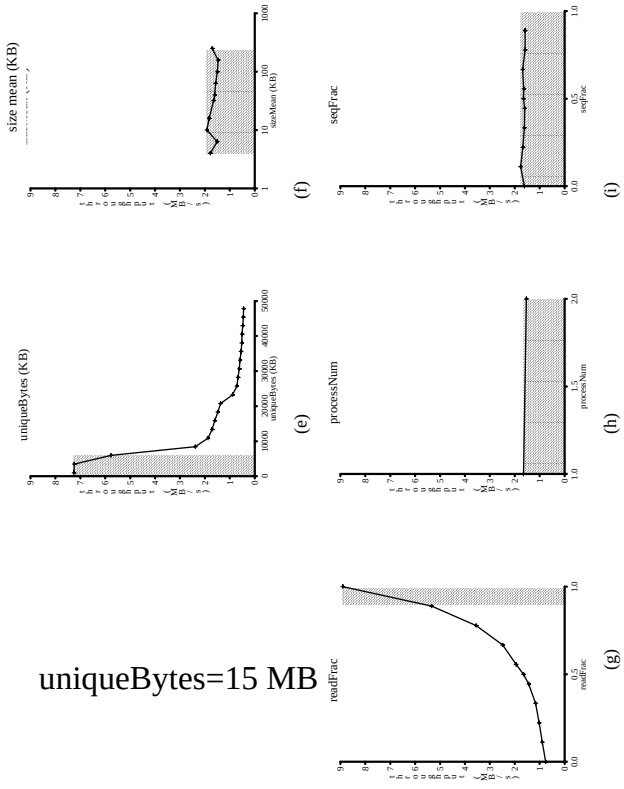

Plateau 3 [DEC 5000]

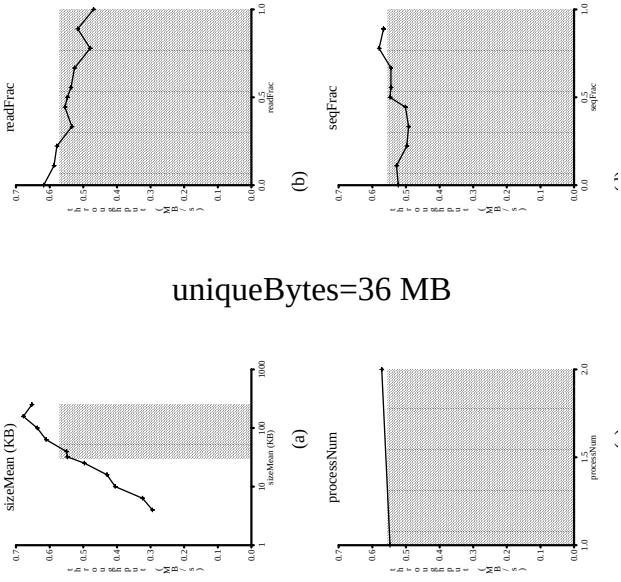

Plateau 1 [HP 730]

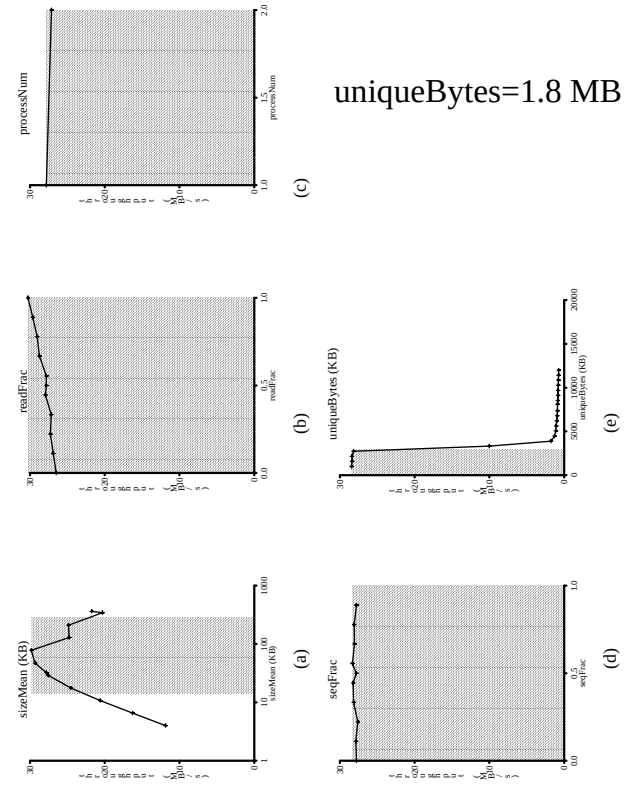

Plateau 2 [HP 730]

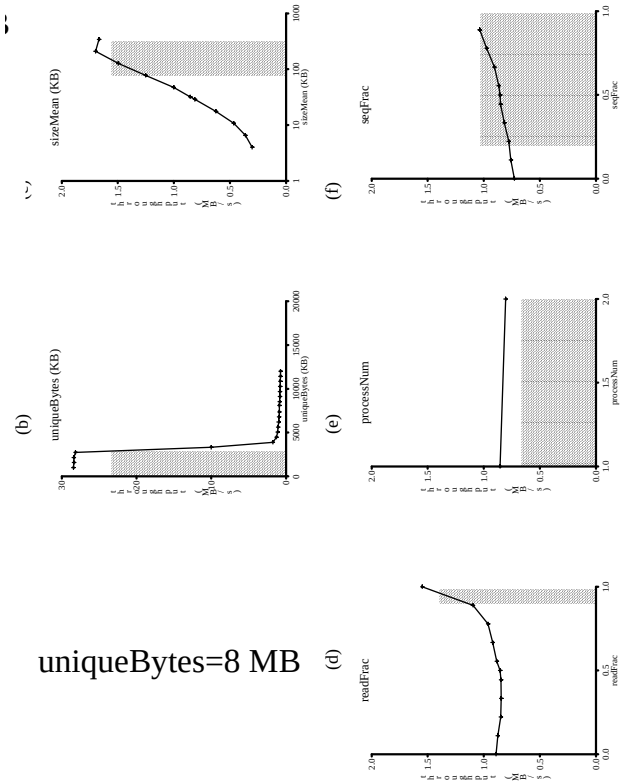

Plateau 1 [Convex C240]

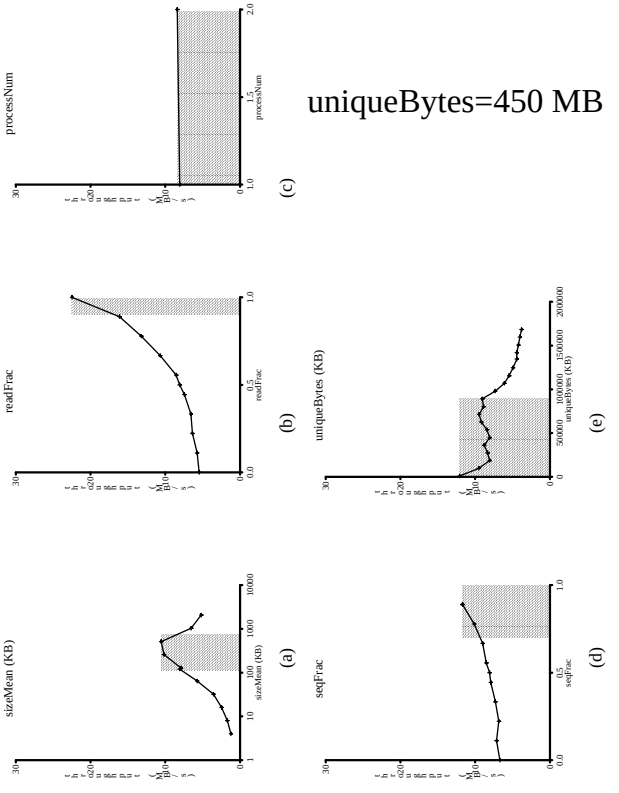

Plateau 2 [Convex C240]

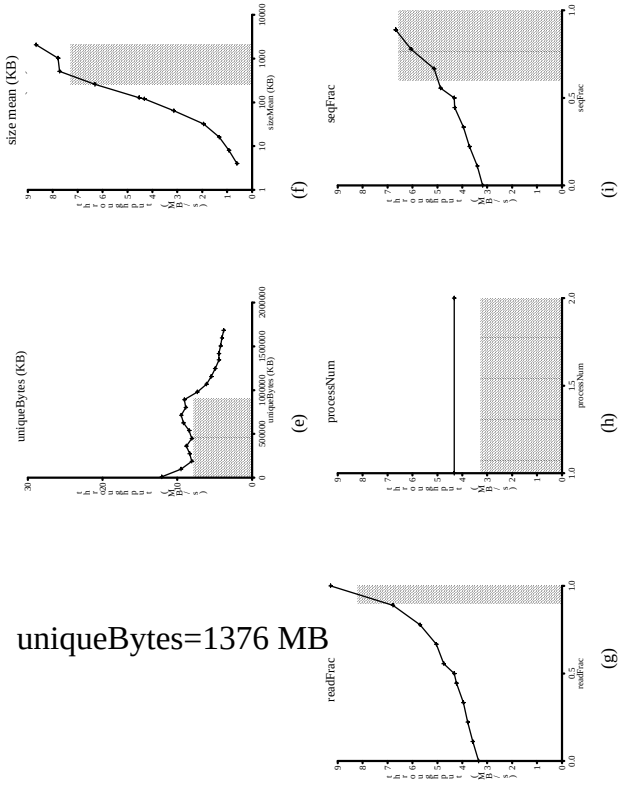

Plateau 1 [Solbourne]

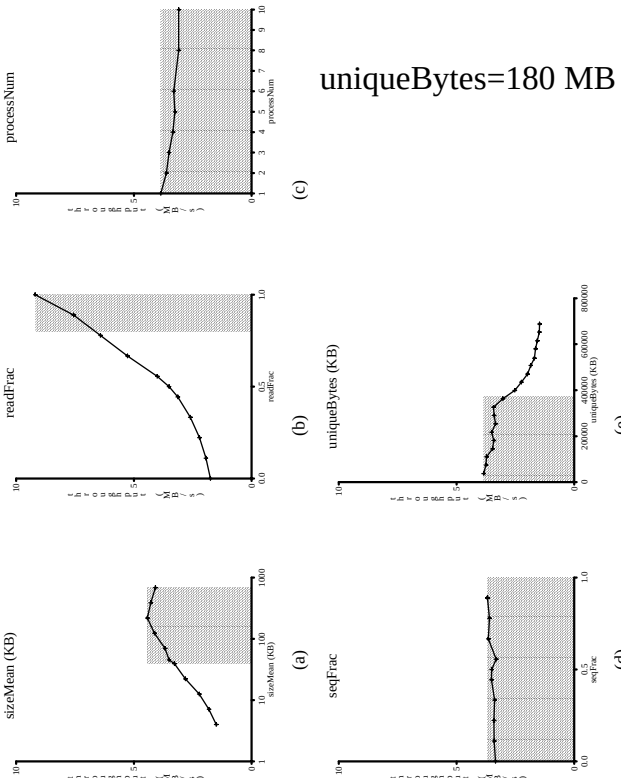

Plateau 2 [Solbourne]

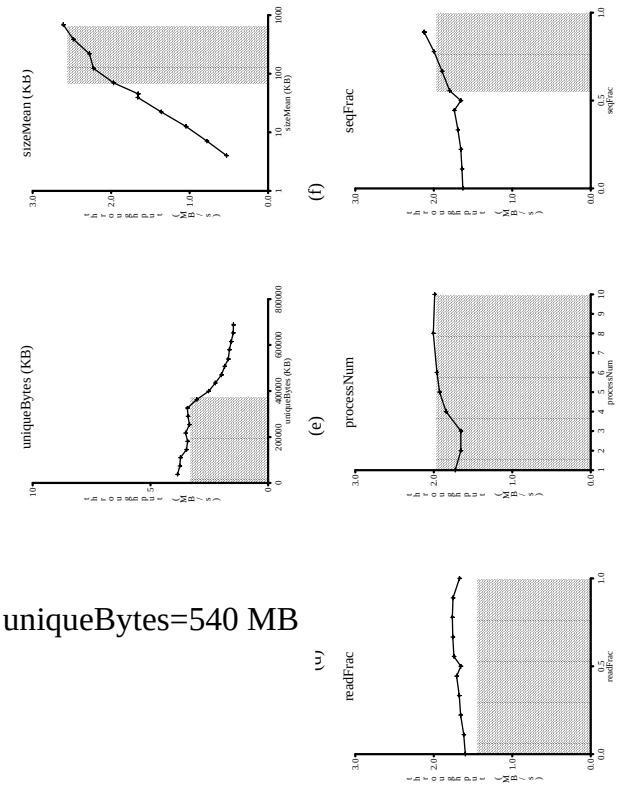

uniqueBytes=540 MB

Prediction Accuracy

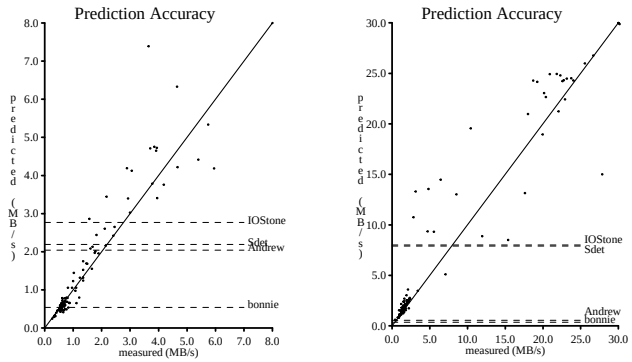

(a) DECstation 5000/200 (b) HP 730

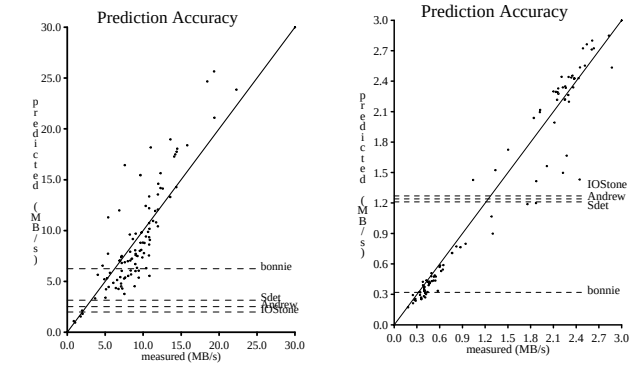

(c) Convex C240 Sparc 1