

Computer Systems Performance Analysis and Benchmarking (37-235)

Analytic Modeling Simulation Measurements / Benchmarking

Lecture by:

Prof. Thomas Stricker

Assignments/Projects:

Christian Kurmann

Textbook:

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

Topic of Today:

- **Self Scaling Benchmarks**
SIGMETRICS Paper 1993 by P.Chen
- **Performance Modeling & Prediction**
SUPERCOMPUTING Paper 1998 by M.Taufer

24.1.02 - 1 37-235 Perf.Eval.&Benchmarking © Stricker

What's wrong w. I/O benchmarks

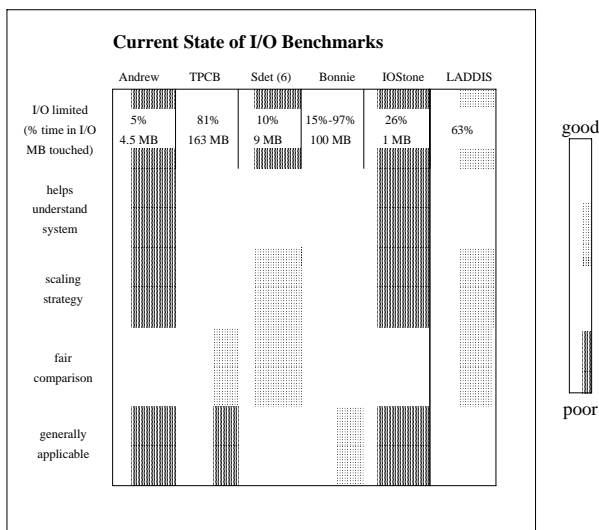
- Become quickly obsolete
i.e. Have no scaling strategy
- Do not stress the I/O system
- Do not help to understand IO
- Comparison for one workload only
i.e. don't allow conclusions about relative performance.
- Narrow application range
i.e. TPC-B is only debit/credit

Look at

- IO Percentage
- Maximal Working Set
- Scaling strategy
- Insight into system
- Suitability for comparison

24.1.02 - 2 37-235 Perf.Eval.&Benchmarking © Stricker

Current State of Art (1993)



24.1.02 - 3 37-235 Perf.Eval.&Benchmarking © Stricker

Workload Model

- uniqueBytes
 - working set
- sizeMean
 - 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
- processNum
 - multithreading, number of processes issuing I/O.


How accurately can we model real workloads?

24.1.02 - 4 37-235 Perf.Eval.&Benchmarking © Stricker

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



Terms

- Single Parameter Graphs
- Focal Point
- Focal Vector

Example

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


Problems:

- 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]

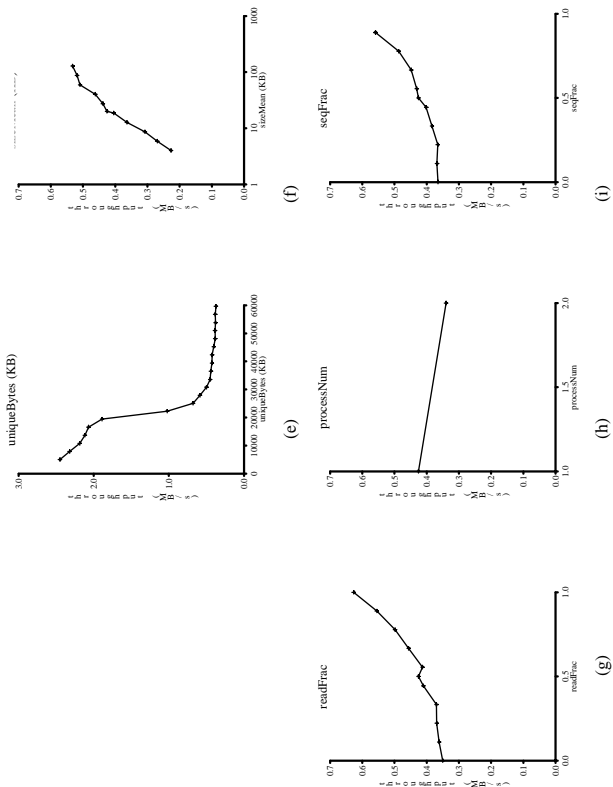
Scaling all Parameters [Sparc1]



Plateau 1 [Sparc1]



Plateau 2 [Sparc1]



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.
- Simple linear interpolation
- Multiplicative linear interpolation for multiple parameters

Clean evaluation of the error in the paper.

Predicting Performance




Accuracy of Prediction [Sparc1]



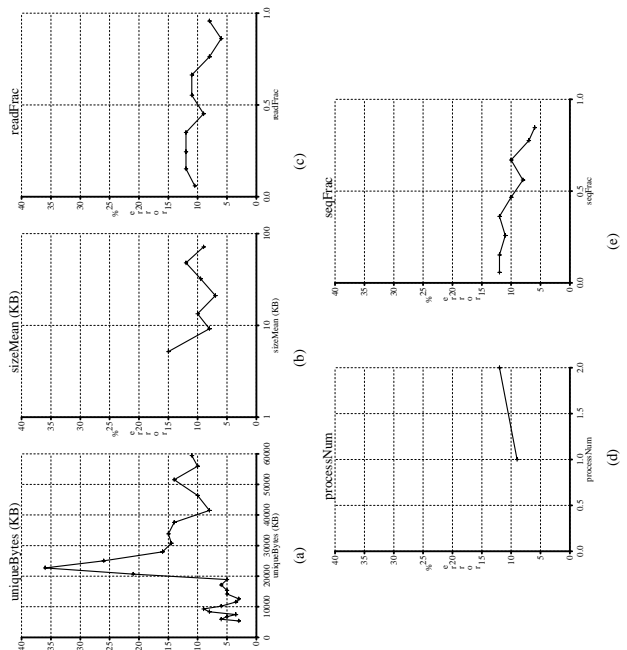
(a)

Error of Prediction [Sparc1]



(b)

Parameters with error




Discussion of Machines

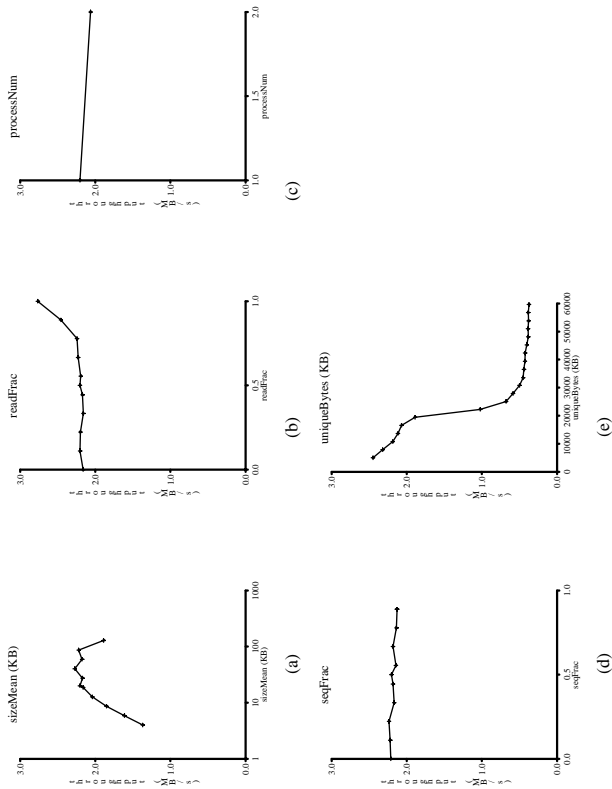
| System Name | SPARCstation 1+ | DECstation 5000/200 | HP 730 |
|---------------------|-----------------|----------------------|--------------|
| Year Released | 1989 | 1990 | 1991 |
| CPU | SPARC 8.3 | MIPS R3000 | PA-RISC 76.8 |
| SPECmarks | CDC Wren IV | 3 disk (Wren) RAID 0 | HP I350SX |
| Disk System | SCSI-I | 100 MB/s | Fast SCSI-II |
| I/O Bus | 28 MB | 32 MB | 32 MB |
| Mem. Bus Peak Speed | SunOS 4.1 | Sprite LFS | HP/UX 8.07 |
| Memory Size | | | |
| Operating System | | | |

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


Storage Hierarchies




Plateau 1 [Sparc1]



Plateau 2 [Sparc1]

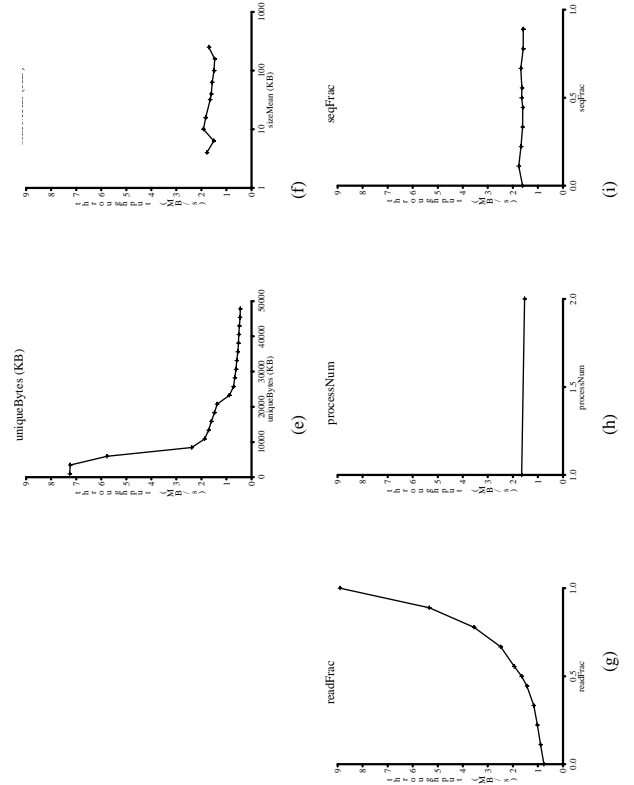


Plateau 1 [DEC 5000]



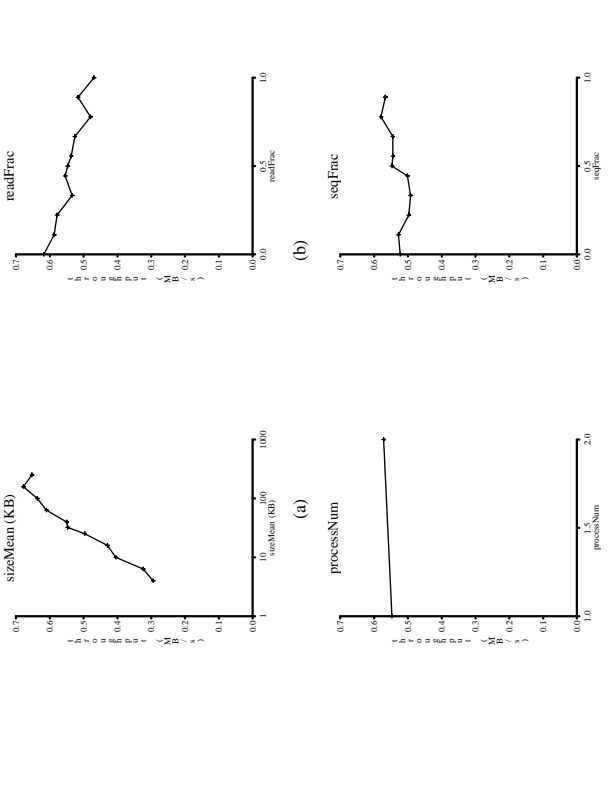
24.1.02 - 21 37-235 Perf.Eval.&Benchmarking © Stricker

Plateau 2 [DEC 5000]



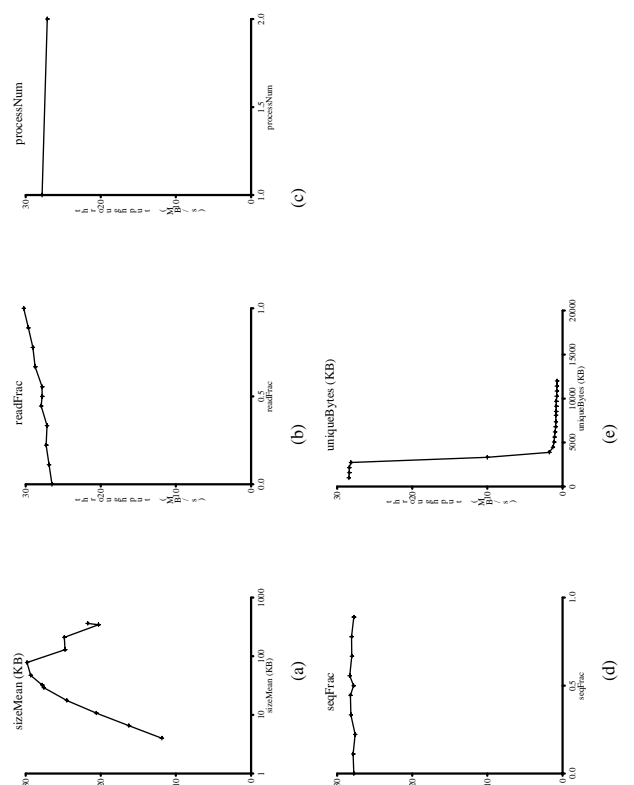
24.1.02 - 22 37-235 Perf.Eval.&Benchmarking © Stricker

Plateau 3 [DEC 5000]



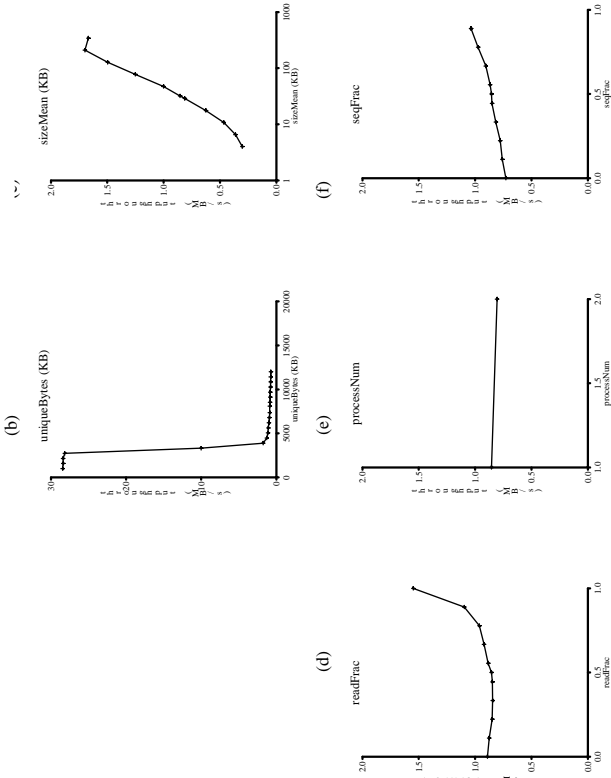
24.1.02 - 23 37-235 Perf.Eval.&Benchmarking © Stricker

Plateau 1 [HP 730]



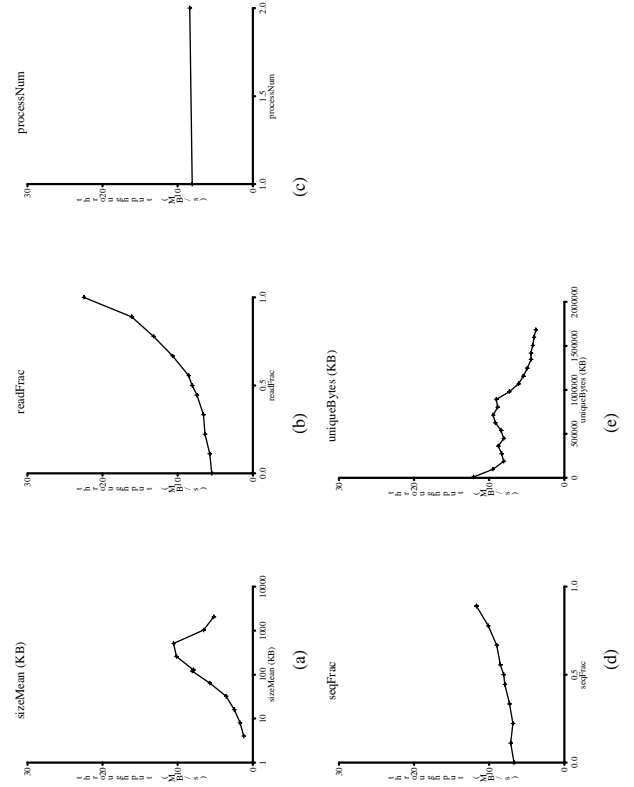
24.1.02 - 24 37-235 Perf.Eval.&Benchmarking © Stricker

Plateau 2 [HP 730]



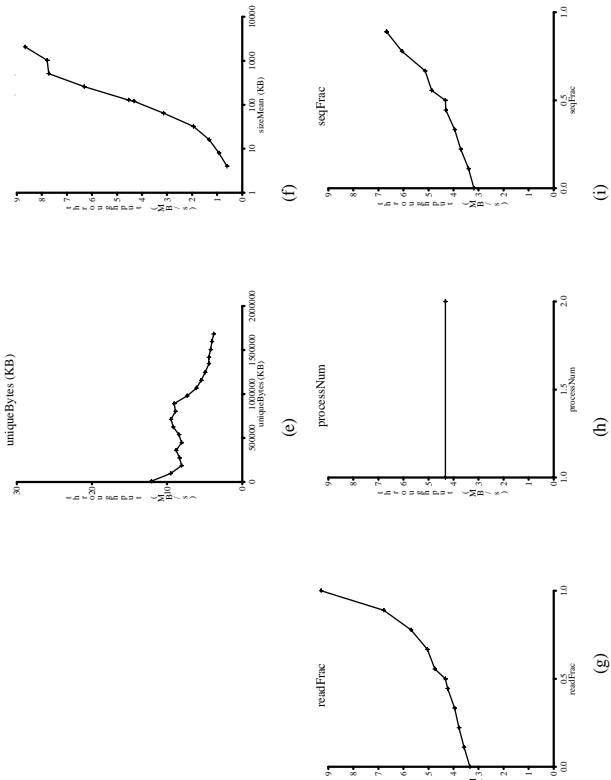
24.1.02 - 25 37-235 Perf.Eval.&Benchmarking © Stricker

Plateau 1 [Convex C240]



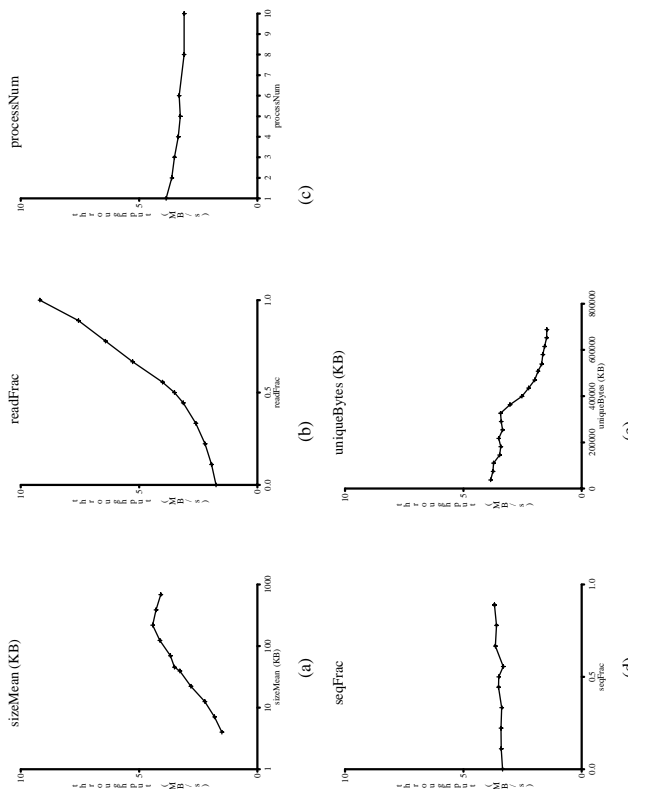
24.1.02 - 26 37-235 Perf.Eval.&Benchmarking © Stricker

Plateau 2 [Convex C240]




24.1.02 - 27 37-235 Perf.Eval.&Benchmarking © Stricker

Plateau 1 [Solbourne]



24.1.02 - 28 37-235 Perf.Eval.&Benchmarking © Stricker

Plateau 2 [Solbourne]



Prediction Accuracy

