

Computer Systems Performance Analysis and Benchmarking (37-235)

Analytic Modelling 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:

- **Workloads**
 - **Popular Benchmarks**
 - **Workload Selection**

Workloads

- test workload
 - workload used in a performance evaluation.
- real workload
 - normal operations
 - often not suitable for test
 - not reproduceable
- synthetic workload
 - characteristic similar to real
 - controlled
 - repeatedly applied

Types of workloads

- Single Instruction (e.g. an add)
- Instruction mixes
- Kernels
- Synthetic programs
- Application benchmarks

Addition Instruction

very frequent instruction

good pick

early days of computing

Instruction mix

Gibson Mix

TABLE 4.1 Gibson Instruction Mix

1. Load and Store	31.2
2. Fixed-Point Add and Subtract	6.1
3. Compares	3.8
4. Branches	16.6
5. Floating Add and Subtract	6.9
6. Floating Multiply	3.8
7. Floating Divide	1.5
8. Fixed-Point Multiply	0.6
9. Fixed-Point Divide	0.2
10. Shifting	4.4
11. Logical, And, Or	1.6
12. Instructions not using registers	5.3
13. Indexing	18.0
	100.0

IBM 704 and IBM 650 (dinosaurs!)

- Average instruction time
- Clocks per instruction
- Directly MIPS/MFLOPS

Problem: Pipelining, caching address translation.

Kernels

Specific function like:

- Scalar Product of two Vectors (BLAS1)
- Matrix Vector Multiply (BLAS2)
- Sieve, Puzzle
- Tree Search
- Ackermann
- Matrix Inversion
- Sorting (NAS sort)

Synthetic Programs

see example in the book...

- Used in a microbenchmarks
- Exercises one or two parameter spaces

Application Benchmarks

- Banking
- Airline reservation
- large scientific codes
 - includes everything
 - hardware
 - input/output
 - networks
 - operating system
 - databases
- e.g. debit credit benchmark

Popular Benchmarks

Sieve of Erathostenes:

- Brain dead way to find primes
- Depends on:
 - Memory Systems Performance
 - Implementation Array/Bitset
 - Working Set in Memory

Ackermann Function:

- Silly recursive Funktion
- Product of Infinite Log Shaving by Theoreticians. Notion of $\log^* \sim \text{const.}$ but not quite constant.
- Inverse of $x=\log^*(y)$
 - $A(1, j) = 2^j$
 - $A(i, 1) = A(i-1, 2)$
 - $A(i, j) = A(i-1, A(i, j-1))$

see also: Cormen, Leiserson, Rivest: Algorithms

Benchmarks

Dhrystone

- Reinhold Weicker at Siemens [1984]
- lots of procedure calls
- systems programming workload
- in C, Ada and Pascal
- Integer only: DIPS

Whetstone

- British central computer agency [1975]
- 11 modules designed after 949 Algol programs.
- Fixed and floating point arithmetic WHIPS
- compiler sensitive

both will fit in caches...

Benchmarks

LINPACK

- Jack Dongarra (1983)
- Subroutines to solve dense linear systems.
- Most time consumed in BLAS (Basic Linear Algebra Subroutines)
- MFLOPs

Popular for:

- 100x100 system of equations
- 1000x1000 system of equations

Represent:

- finite element analysis
- calls for high computation speed
- graphics processing ???

Benchmarks

Debit Credit

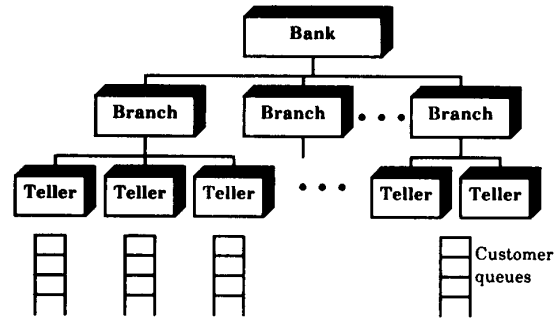


FIGURE 4.4 Banking environment.

Begin-Transaction	
Read message from the terminal	(100 bytes)
Rewrite account	(100 bytes, random)
Write history	(50 bytes, sequential)
Rewrite teller	(100 bytes, random)
Rewrite branch	(100 bytes, random)
Write message to the terminal	(200 bytes)
Commit-Transaction	

FIGURE 4.5 Debit-credit transaction pseudo-code.

TPC Benchmark A, August 1988

Benchmarks

SPEC Suite 1992, 1995, 1998...

Programs submitted by scientists and engineers.

- GCC
- Espresso Electronic Design Automation
- Spice 2g6 Electronic Design Automation
- Doduc Monte Carlo Simulation Nuke
- NASA7 Matrix OPs submitted by NASA
- LI Do 9 queens on 8x8 board in LISP
- Eqntott Boolean Eqn to Truth Table
- Matrix300 LINPACK 300x300 doubles
- Fpppp Quantum Chemistry integral on double floats FORTRAN
- Tomcatv Vectorized Mesh Generator

Benchmarks

- CPU and FPU
- Two copies run on each processor
- Compared to Reference System a VAX 11/780 Release 1.0 ~ 1 MIPS
- Ratio of each benchmarks
- Spec Througput considering the number of processors #cpu@ratio.
- Geometric Mean leads to SPECmark.

New trend:

- Industry standard benchmarks

Workload Selection (More an Art than a Science)

Very important part of every performance evaluation project.

- System under Test
e.g. CPU
- Component under Study
e.g. ALU

System determines selection of workload not component.

Banking

- transaction per second
- not instructions per second


CPU

- vice versa


View of a System

- 1. ALU
- 2. CPU
- 3. OS
- 4. Application

Hierarchical



Example Network



Choice of Workload depends on layer at which networks are compared.

Example Backup System

- Backup System
- Tape Data System
- Tape Drives
- Read/Write Subsystem
- Read/Write Heads

A backup system

1. *Backup System:*
 - (a) Services: Backup files, backup changed files, restore files, list backed-up files.
 - (b) Factors: File system size, batch or background process, incremental or full backups.
 - (c) Metrics: Backup time, restore time.
 - (d) Workload: A computer system with files to be backed up. Vary frequency of backups.
2. *Tape Data System:*
 - (a) Services: Read/write to the tape, read tape label, autoloader tapes.
 - (b) Factors: Type of tape drive.
 - (c) Metrics: Speed, reliability, time between failures.
 - (d) Workload: A synthetic program generating representative tape I/O requests.
3. *Tape Drives:*
 - (a) Services: Read record, write record, rewind, find record, move to end of tape, move to beginning of tape.
 - (b) Factors: Cartridge or reel tapes, drive size.
 - (c) Metrics: Time for each type of service, for example, time to read record and to write record, speed (requests per unit time), noise, power dissipation.
 - (d) Workload: A synthetic program exercising generating various types of requests in a representative manner.
4. *Read/Write Subsystem:*
 - (a) Services: Read data, write data (as digital signals).
 - (b) Factors: Data-encoding technique, implementation technology (CMOS, TTL, and so forth).
 - (c) Metrics: Coding density, I/O bandwidth (bits per second).
 - (d) Workload: Read/write data streams with varying patterns of bits.
5. *Read/Write Heads:*
 - (a) Services: Read signal, write signal (electrical signals).
 - (b) Factors: Composition, interhead spacing, gap sizing, number of heads in parallel.
 - (c) Metrics: Magnetic field strength, hysteresis.
 - (d) Workload: Read/write currents of various amplitudes, tapes moving at various speeds. □

Level of Detail

- Most frequent request
- Frequency of request types
- Trace of requests (with timestamps)
- Average resource demand
- Distribution of resource demands
- Concept of a trace

Workload descriptions

- executable
- non-executable

Representativeness

- Arrival Rate
- Resource Demands
- Resource Usage Profile

Timeliness

- demands change
- workloads have to change

Example: SPEC 89,92,95,98...

Problem: caches get bigger and bigger

Considerations

- Loading Level (best, typical, worst)
- External Components
- Repeatability