Types of Workloads

Adapted by Prof. Gheith Abandah

Overview

- Terminology
- Test Workloads for Computer Systems
  - Addition Instruction
  - Instruction Mixes
  - Kernels
  - Synthetic Programs
  - Application Benchmarks: TPC, SPEC

Terminology

- **Test workload**: Any workload used in performance studies.
  Test workload can be real or synthetic.
- **Real workload**: Observed on a system being used for normal operations.
- **Synthetic workload**:
  - Similar to real workload
  - Can be applied repeatedly in a controlled manner
  - No large real-world data files
  - No sensitive data
  - Easily modified without affecting operation
  - Easily ported to different systems due to its small size
  - May have built-in measurement capabilities.

Test Workloads for Computer Systems

1. Addition Instruction
2. Instruction Mixes
3. Kernels
4. Synthetic Programs
5. Application Benchmarks

Addition Instruction

- Processors were the most expensive and most used components of the system
- Addition was the most frequent instruction

Instruction Mixes

- ❑ Instruction mix = instructions + usage frequency
- ❑ Gibson mix: Developed by Jack C. Gibson in 1959 for IBM 704 systems.

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, etc.	1.6
12.	Instructions Not Using Registers	5.3
13.	Indexing	18.0
	Total	100.0

Instruction Mixes (Cont)

- Disadvantages:
  - Complex classes of instructions not reflected in the mixes.
  - Instruction time varies with:
    - Addressing modes
    - Cache hit rates
    - Pipeline efficiency
    - Interference from other devices during processor-memory access cycles
    - Number of times a conditional branch is taken

Instruction Mixes (Cont)

- Performance Metrics:
  - MIPS = Millions of Instructions Per Second
  - MFLOPS = Millions of Floating Point Operations Per Second

Kernels

- Kernel = nucleus
- Kernel= the most frequent function
- Commonly used kernels: Sieve, Puzzle, Tree Searching, Ackerman's Function, Matrix Inversion, and Sorting.
- Disadvantages: Do not make use of I/O devices

Synthetic Programs

- To measure I/O performance lead analysts to use Exerciser loops
- The first exerciser loop was by Buchholz (1969) who called it a synthetic program.

Synthetic Programs

- Advantage:
  - Quickly developed and given to different vendors.
  - No real data files
  - Easily modified and ported to different systems.
  - Have built-in measurement capabilities
  - Measurement process is automated
  - Repeated easily on successive versions of the operating systems
- Disadvantages:
  - Too small
  - Do not make representative memory or disk references
  - Mechanisms for page faults and disk cache may not be adequately exercised.
  - CPU-I/O overlap may not be representative.
  - Loops may create (synchronizations) better or worse performance.

Application Benchmarks

- For a particular industry: Debit-Credit for Banks
- Some Authors: Benchmark = set of programs taken from real workloads

Example Benchmarks

- TPC
- SPEC
- Whetstone
- LINPACK
- Dhrystone
- Lawrence Livermore Loops

TPC

- Transaction Processing Performance Council
- http://www.tpc.org/
- **TPC-C** simulates a complete computing environment where a population of users executes transactions against a database. TPC-C performance is measured in new-order transactions per minute. The primary metrics are the transaction rate (tpmC), and the associated price per transaction ($/tpmC).

TPC-E

- **TPC-E** uses a database to model a brokerage firm with customers who generate transactions related to trades, account inquiries, and market research. The brokerage firm in turn interacts with financial markets to execute orders on behalf of the customers and updates relevant account information. The TPC-E metric is given in transactions per second (tps).

TPC-H

- **TPC-H** is a decision support benchmark. It consists of a suite of business oriented ad-hoc queries and concurrent data modifications. The performance metric reported by TPC-H is called the TPC-H Composite Query-per-Hour Performance Metric (QphH@Size) for a specific database size. The TPC-H Price/Performance metric is expressed as $/QphH@Size.

TPC-Energy

- **TPC-Energy** contains the rules and methodology for measuring and reporting an energy metric in TPC Benchmarks. This includes the energy consumption of system components associated with typical business information technology environments.

SPEC Benchmark Suite

- Systems Performance Evaluation Cooperative (SPEC): Non-profit corporation formed by leading computer vendors to develop a standardized set of benchmarks.
- www.spec.org
- **SPEC CPU 2017** designed to provide performance measurements that can be used to compare compute-intensive workloads on different computer systems, contains four benchmark suites:
  - **SPECspeed 2017: integer and FP**
  - **SPECrate 2017: integer and FP**

SPEC 2017

Calculating SPECspeed® Metrics	Calculating SPECrate® Metrics
1 copy of each benchmark in a suite is run.	The tester chooses how many concurrent copies to run
The tester may choose how many OpenMP threads to use.	OpenMP is disabled.
For each benchmark, a performance ratio is calculated as: time on a reference machine / time on the SUT	For each benchmark, a performance ratio is calculated as: number of copies * (time on a reference machine / time on the SUT)
Higher scores mean that less time is needed.	Higher scores mean that more work is done per unit of time.

SPEC Examples

Cspeed® Metrics	Calculating SPECrate® Metrics
Example: • The reference machine ran 600.perlbench_s in 1775 seconds. • A particular SUT took about 1/5 the time, scoring about 5. • More precisely: 1775/354.3 = 5.009	Example: • The reference machine ran 1 copy of 500.perlbench_r in 1592 seconds. • A particular SUT ran 8 copies in about 1/3 the time, scoring about 24. • More precisely: 8*(1592/541.5) = 23.5

Summary

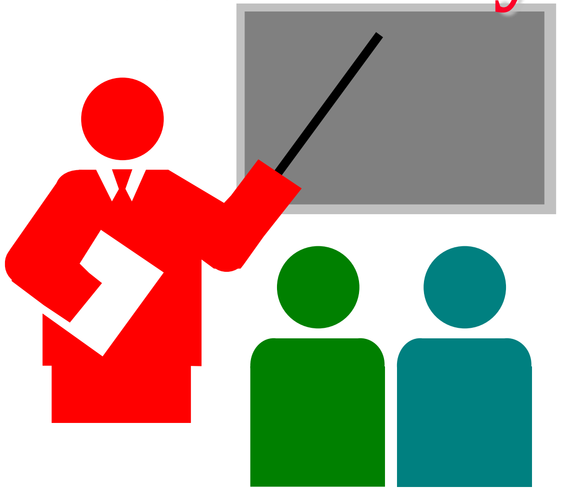

- Synthetic workload are representative, repeatable, and avoid sensitive information
- Add instruction – most frequent instruction initially
- Instruction mixes, Kernels, synthetic programs
- Application benchmarks: Sieve, Ackerman, …
- Benchmark standards: TPC, SPEC