

Types of Workloads

Raj Jain

Washington University in Saint Louis

Saint Louis, MO 63130

Jain@cse.wustl.edu

These slides are available on-line at:

<http://www.cse.wustl.edu/~jain/cse567-11/>



- ❑ Terminology
- ❑ Test Workloads for Computer Systems
 - Addition Instruction
 - Instruction Mixes
 - Kernels
 - Synthetic Programs
 - Application Benchmarks: Sieve, Ackermann's Function, Debit-Credit, SPEC

Part II: Measurement Techniques and Tools

Measurements are not to provide numbers but insight
- Ingrid Bucher

1. What are the different types of workloads?
2. Which workloads are commonly used by other analysts?
3. How are the appropriate workload types selected?
4. How is the measured workload data summarized?
5. How is the system performance monitored?
6. How can the desired workload be placed on the system in a controlled manner?
7. How are the results of the evaluation presented?

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: Usage frequency
3. Kernels: Key functions, e.g., sorting. Matches order
4. Synthetic Programs: Bigger programs
5. Application Benchmarks: Matches I/O

SPEC Benchmark Suite

- ❑ Systems Performance Evaluation Cooperative (SPEC): Non-profit corporation formed by leading computer vendors to develop a standardized set of benchmarks.
- ❑ Release 1.0 consists of the following 10 benchmarks: GCC, Espresso, Spice 2g6, Deduc, LI, Eqntott, Matrix300, Fpppp, Tomcatv
- ❑ Primarily stress the CPU, Floating Point Unit (FPU), and to some extent the memory subsystem \Rightarrow To compare CPU speeds.
- ❑ Benchmarks to compare I/O and other subsystems may be included in future releases.


SPEC (Cont)

- ❑ The elapsed time to run two copies of a benchmark on each of the N processors of a system (a total of $2N$ copies) is measured and compared with the time to run two copies of the benchmark on a reference system (which is VAX-11/780 for Release 1.0).
- ❑ For each benchmark, the ratio of the time on the system under test and the reference system is reported as **SPECthruput** using a notation of #CPU@Ratio. For example, a system with three CPUs taking $1/15$ times as long as the the reference system on GCC benchmark has a SPECthruput of 3@15.
- ❑ Measure of the per processor throughput relative to the reference system

SPEC (Cont)

- ❑ The aggregate throughput for all processors of a multiprocessor system can be obtained by multiplying the ratio by the number of processors. For example, the aggregate throughput for the above system is 45.
- ❑ The geometric mean of the SPECthroughputs for the 10 benchmarks is used to indicate the overall performance for the suite and is called **SPECmark**.

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: Debit-Credit, SPEC

Exercise 4.1

Select an area of computer systems (for example, processor design, networks, operating systems, or databases), review articles on performance evaluation in that area and make a list of benchmarks used in those articles.

Exercise 4.2

Implement the Sieve workload in a language of your choice, run it on systems available to you, and report the results.

Homework 4

- Make a list of latest workloads from www.spec.org