

CPS311 Lecture: Performance

Last revised December 5, 2017

Objectives:

1. To introduce basic concepts regarding performance
2. To introduce use of benchmarks
3. To introduce various sorts of means

Materials:

1. Web access
2. Projectable of comparison of SPEC results for two systems
3. projectable of Null table 11.6

I. Introduction

A. Measuring the performance of a computer system is important for a number of reasons.

ASK for examples

1. Deciding which system to purchase
2. Determining if a specific application is feasible.
3. "Tuning" a system to optimize performance

B. But how do we measure the performance of a computer system?

1. Manufacturers tend to publish numbers that make their products appear impressive

ASK for examples

- a. CPU clock speed
- b. bus speed

c. FLOPS (floating operations per second), MFLOPS, GFLOPS

d. MIPS (millions of instructions per second)

2. However, such numbers are not a good way of comparing performance.

ASK Why?

Requirements of different tasks vary greatly:

a. Compute bound versus memory bound versus IO bound tasks.

Of course, the notion of compute bound versus memory bound versus IO bound depends on relative speed of components. The same task may be IO bound on one system, memory bound on another system, and compute bound on another. Here, the notion of the "VonNeumann bottleneck" (memory) comes in to play. Hierarchical memory attempts to yield a system that is close to as fast as its fastest component, but degree of success depends on proper sizing of component parts and may be application specific (e.g. applications with greater inherent locality do much better than those with little locality).

"A chain is as strong as its weakest link"

b. Tasks requiring extensive FP computation (scientific tasks) versus those requiring little or none.

c. In the case of measures like FLOPS and MIPS, the instruction mix is critical - e.g. on any machine floating multiply is more complex than floating add; on CISCs different instructions can take widely varying numbers of cycles.

d. Other issues such as network communication ...

II. Benchmarks

A better way to measure performance is through use of benchmarks that test a system's performance in a workload that is perceived to be similar to that in which the intended system will be used.

A.ASK if you've ever seen these

1. Trade magazines typically have suites they use for product reviews.

Example: <http://www.macworld.com/article/2082044/how-we-test-speedmark-9-mac-benchmarks.html>

(Note: this article is somewhat dated, but the same basic approach is used for newer products - maybe with a different suite of tests)

SHOW VIDEO

2. There are industry-standard benchmark suites. The most-widely known are the SPEC suites (Standard Performance Evaluation Corporation), which has been developing suites since 1988 (continually being revised, of course!)

Example: <http://www.spec.org/> - then navigate to graphics/workstations, then published results for SPECviewperf@12.1, then results summary

B.Two issues with benchmark suites

1. No system is typically best on every test.

Example: look at example SPEC results - compare Dell 5810 with two different CPU's (Intel Xeon 1650 at 3.50 GHz; Intel Xeon 1680 at 3.20 GHz on various suites). Note that which does best varies from suite to suite.

PROJECT

2. A host of number can get confusing, especially when system price is also considered

3. A manufacturer may actually optimize with a particular benchmark in mind. This has happened - sometimes in very significant ways (e.g. special compiler flags used when compiling a particular benchmark program in order to optimize the code in a benchmark-specific way).

This sort of thing is still allowed with suites such as SPEC - but SPEC publishes two numbers for a given system - with and without optimization flags

C. Some things one may do to address such issues:

1. Calculation of an overall score

(SPEC does this for each of the individual benchmark suites, but does not attempt to produce an overall score considering all suites)

Example: On SPEC site: navigate to results for one suite - note individual tests, weighted average which is then carried forward into overall.

2. Calculation of cost / performance ratio SPEC used to do this, but does not do this anymore.

3. Normalization to a specific system

Example: MacWorld data on page looked at earlier does this for overall summary (Show discussion in 2nd paragraph - all scores normalized to score for 2.4 GHz 2010 MacMini Core Duo processor w/2 GB of RAM as 100)

III. Calculating Means

How one calculates a mean is also an issue.

A. A simple arithmetic average of raw data is usually not very good.

ASK Why?

Suppose we had the following raw data (high numbers better)

	Test 1	Test 2	Test 3	Test 4
System A	1000	5	5	10
System B	980	10	10	20

System A is slightly better than B on Test 1, but B is twice as good as A on all the other tests. What is the simple arithmetic average?

A: $1020/4 = 255$

B: $1020/4 = 255$

Averages suggest both are equal, which is really misleading in this case.

B. Often, a weighted average is better than a simple average

1. Example: SPEC details - note weights

2. However, this has a problem, too.

ASK

Assumes some knowledge of workload mix in order to properly assign weights.

C.A geometric mean may be better.

1. Geometric means are calculated using normalized values
2. The geometric mean of n numbers is $(x_1 * x_2 * \dots * x_n)^{1/n}$ (nth root)
3. Example: using above data, normalizing to A

	Test 1	Test 2	Test 3	Test
System A	1	1	1	1
System B	.98	2	2	2

A: 4th root of 1 = 1

B: 4th root of $.98 * 2 * 2 * 2 = 7.84 = 1.67$

Ratio of geometric means = 1.67:1

4. The ratio is independent of choice of base system - e.g. normalizing to B

	Test 1	Test 2	Test 3	Test
System A	1.02	0.5	0.5	0.5
System B	1	1	1	1

A: 4th root of $1.02 * 0.5 * 0.5 * 0.5 = .1275 = 0.6$

B: 4th root of 1 = 1

Ratio of geometric means = $1:0.6 = 1.67:1$

5. The geometric mean is not useful for predicting performance - i.e. the above data does not mean that typically system B is 1.67 times faster than A!

D. When rates are involved, one must use the harmonic mean

Example:

Suppose we made a 30 mile trip, going 10 miles at 10 mph, 10 at 40 mph, and 10 at 70 mph. What is the average speed?

It is not $(10 + 40 + 70) / 3 = 40$ mph!

1. Clearly, going 10 miles at 10 mph takes 60 minutes.

Going 10 miles at 40 mph takes 15 minutes

Going 10 miles at 70 mph takes about 8.6 min.

So the total trip takes 83.6 minutes, for an average speed of only 21.5 mph.

2. When we have rates, the harmonic mean is calculated as 1 over the average of the reciprocals - e.g. for the example above

$$1 / ((1/10 + 1/40 + 1/70) / 3) = 1 / ((.1 + .025 + .014) / 3) = 1 / .0463 = 21.5$$

E. Summary table on measures of central tendency: project Null table 11.6 page 593 in Null book