Handout 16

Computer Benchmarking

Benchmark

- The act of running a computer program, a set of programs, or other operations, in order to assess the relative performance.
- Benchmarks provide a method of comparing the performance of various subsystems across different chip/system architectures.

Why Benchmarks?

- It has became more difficult to compare the performance of various computer systems simply by looking at their specifications.
- A slower processor, with regard to clock frequency, can perform as well as a processor operating at a higher frequency.
- For example, RISC vs CISC.

Megahertz Myth

- The error of using clock rate to compare the performance of different microprocessors.
- While clock rates are a valid way of comparing the performance of different speeds of the same model and type of processor.
  - pipelines and instruction sets can greatly affect the performance when considering different processors.
- One processor may take one clock cycle to add two numbers and another clock cycle to multiply by a third number, whereas another processor may do the same calculation in one clock cycle.

Benchmark Types

- Benchmarks are designed to mimic a particular type of workload on a component or system.
- **Synthetic benchmarks** do this by specially created programs that impose the workload on the component.
- **Application benchmarks** run real-world programs on the system.

Standard Performance Evaluation Corporation (SPEC)

- A non-profit organization that aims to produce "fair, impartial and meaningful benchmarks for computers.
- Results are informally referred to as "SPECmarks", results or even just "SPEC".
- http://www.spec.org/
- Pricing: http://www.spec.org/order.html

Other Benchmarks

- CPU benchmark database
- DEISA – Distributed European Infcrastructure for Supercomputer Applications.
- Futuremark
- Embedded Microprocessor Benchmark Consortium

FLOPS

- Acronym meaning **FL**oating point **O**perations **P**er **S**econd.

Computer Performance	
Name	flops
megaflop	10^6
gigaflop	10^9
teraflop	10^{12}
petaflop	10^{15}
exaflop	10^{18}
zettaflop	10^{21}
yottaflop	10^{24}

Top 500

- The TOP500 table shows the 500 most powerful commercially available computer systems known.
- Statistics of interest:
  - R_{max} - Maximal LINPACK performance achieved
  - R_{peak} - Theoretical peak performance
  - values are in TFlops.

Linpack Benchmark

- LINPACK is a collection of Fortran subroutines that analyze and solve linear equations and linear least-squares problems.
- It was intended for use on supercomputers in the 1970s and early 1980s.
- LINPACK makes use of the BLAS (Basic Linear Algebra Subprograms) libraries.
- LAPACK can be seen as the successor to the original LINPACK.

Top 500

- In the TOP500 List table, the computers are ordered first by their Rmax value.
- http://www.top500.org

Types of Synthetic Benchmarks

- **Whetstone benchmark**
- Primarily measures the floating-point arithmetic performance
  - Originally measured computing power in units of kilo-Whetstone Instructions Per Second (kWIPS).
  - Later changed to Millions of Whetstone Instructions Per Second (MWIPS).

Types of Synthetic Benchmarks

- **Dhrystone benchmark**
  - Characterized numerous programs in terms of various common constructs – procedure calls, pointer indirections, assignments, etc.
  - From this the Dhrystone benchmark was written to correspond to a representative mix.
  - Output from the benchmark is the number of Dhrystones per second.

Types of Synthetic Benchmarks

- **Livermore loops** (also known as the Livermore Fortran kernels or LFK)
  - Is a benchmark for parallel computers. It was run on computers at Lawrence Livermore National Laboratory.
  - It consists of 24 Do loops, some of which can be vectorized, and some of which cannot.

Roy Longbottom's PC Benchmark Collection

- http://freespace.virgin.net/roy.longbottom
- http://www.roylongbottom.org.uk/
- Worked for CCTA
  - Central Computer and Telecommunications Agency
  - UK government agency providing computer and telecoms support to Government departments
  - Official design authority for Whetstone.

BenchNT

DHRY1NNT.EXE	Dhrystone benchmark 1.1	non-optimised
DHRY1ONT.EXE	Dhrystone benchmark 1.1	optimised
DHRY2NNT.EXE	Dhrystone benchmark 2.1	non-optimised
DHRY2ONT.EXE	Dhrystone benchmark 2.1	optimised
LINPCNNT.EXE	Linpack "C" double precision benchmark	non-optimised
LINPCONT.EXE	Linpack "C" double precision benchmark	optimised
LIVECNNT.EXE	Livermore Kernels "C" double precision	non-optimised
LIVECONT.EXE	Livermore Kernels "C" double precision	optimised
WHETCNNT.EXE	Whetstone "C" benchmark single precision	non-optimised
WHETCONT.EXE	Whetstone "C" benchmark single precision	optimised
WHETDCNT.EXE	Whetstone "C" benchmark double precision	optimised