

Cloud Benchmarking:

Fight the black hole

Presented at HPTS September 2013

Anthony F. Voellm
Google Cloud Performance Engineering Manager
voellm@google.com / G+ / @p3rfguy / perfguy.blogspot.com





Abstract:

This talk will cover the history of benchmarking from mainframe to Cloud. The goal is to lay a foundation around where benchmarks started and how they have gotten to where they are. Ideas will be laid out for the future of benchmarking Cloud.

Overview



- Foundation
- History of cloud vs benchmarks
- Off base!
- The black hole

Foundation...

Cloud Computing: Definition



An elastic set of programmable computing resources that can be used on a pay-per-use basis over a network that allows the user to control the computation and focus on the problem being solved and worry less about the physical hardware, software systems, and location.

-Anthony F. Voellm



Watch the talk: http://goo.gl/QAdh5P

Cloud Computing: 2012 Quote from Eric Schmidt





Guillaume Paumier / Wikimedia Commons, CC-BY-3.0.

I don't know if anyone knows what Cloud Computing is... but I do know one thing. It is a marketing term.

Benchmarking: Definition



In computing, a benchmark is the act of running a computer program, a set of programs, or other operations, in order to assess the relative performance of an object, normally by running a number of standard tests and trials against it. The term ...

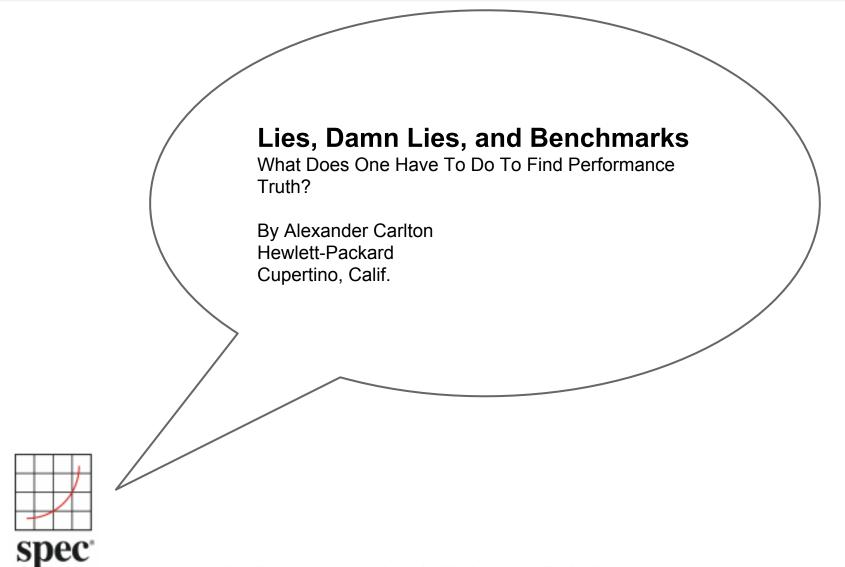
Benchmarks provide a method of comparing the performance of various subsystems across different chip/system architectures.



http://en.wikipedia.org/wiki/Computer_benchmark

Benchmarking: Definition





http://www.spec.org/osg/news/articles/news9412/lies.html

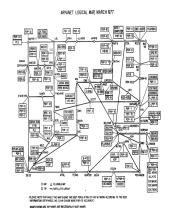
History of benchmarks

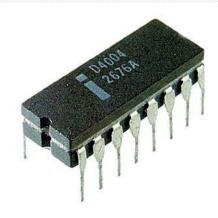
History of Cloud vs Benchmarking





http://en.wikipedia.org/





1961

6000 Computers

John McCarthy -"computation
may someday be organized
as a public utility."

1969

ARPANET 50 kbits/sec

(Mobile phones are 2000x faster in 2013)

1971

Intel markets its first
Microprocessor. 108KHz
6000 transistors

Instruction level Profiling (1957 - 1971)
Gibson Mix, ADP Mix, Process Control Mix

The Gibson Mix (1957 - 1971) Instructions



Percentage of instruction mix

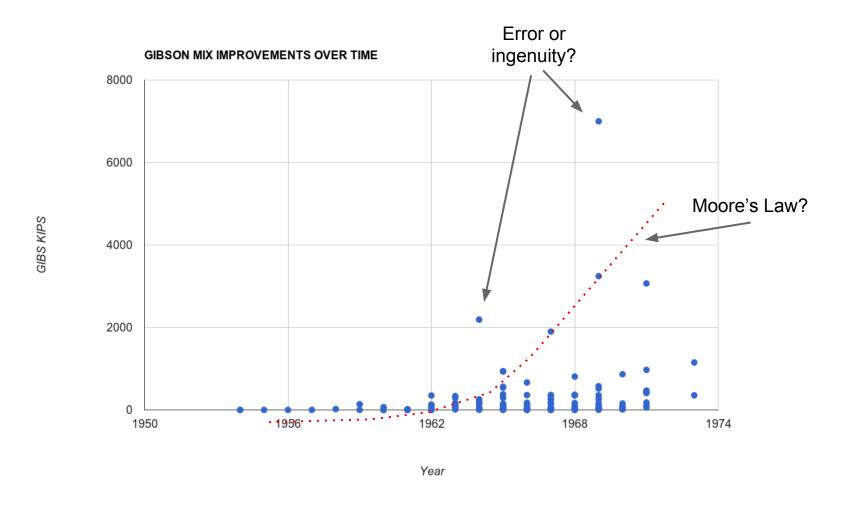
Fixed Point Add/Subtract	0.330
Fixed Point Multiply	0.006
Fixed Point Divide	0.002
Branch	0.065
Compare	0.040
Transfer 8 characters	0.175
Shift	0.046
Logical	0.017
Modification	0.190
Floating Point Add	0.073
Floating Point Multiply	0.040
Floating Point Divide	0.016

Chart From: http://roylongbottom.org.uk/cpumix.htm

The Gibson Mix

Instructions

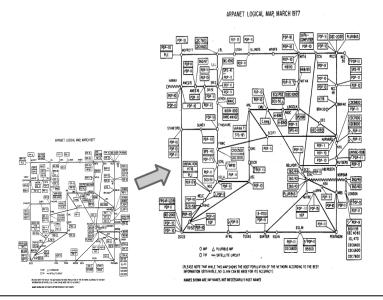




Data From: http://roylongbottom.org.uk/cpumix.htm

History of Cloud vs Benchmarking







http://pocketcalculatorshow.com/magicalgadget/index3.htm



Warrenlead

1970's

Arpanet grows to dozens of machines, 75% of traffic is email, computational power grows from 60k to 330k ops+Loop Unrolling / Compiler Optimization

1980's

BBS Era / Compuserve / GENie / Prodigy Luggables are popular.

1983 - 82720 Graphics Display Controller

1985 - Commodore Amiga

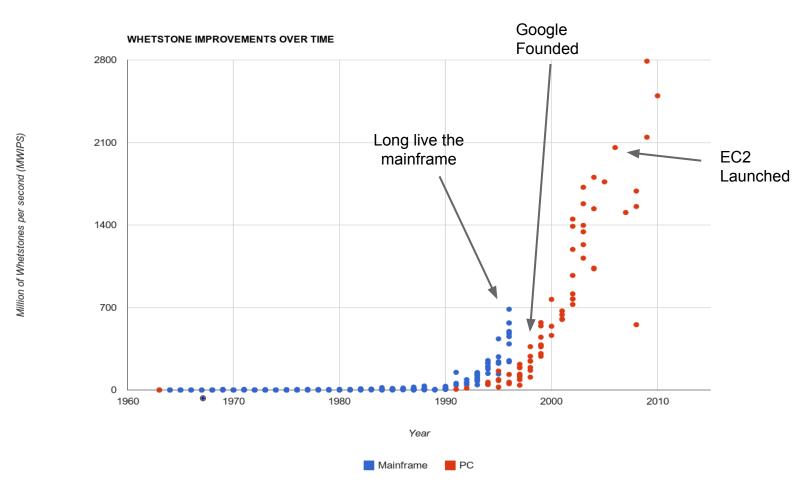
1989 - Parallel Virtual Machine - PVM

- 1972 Enter Whetstone Benchmark use to compare languages First in ALGOL60
- 1972 LINPACK benchmarks first appears
- 1984 Dhrystone is born Integer and Indirection First in Ada
- 1985 Computer Graphics benchmarks are mosty BitBlt, Line, Ellipse, Box, and Text
- 1988 Dhrystone V2.0 Working around compiler optimizations

The Whetstone Benchmark

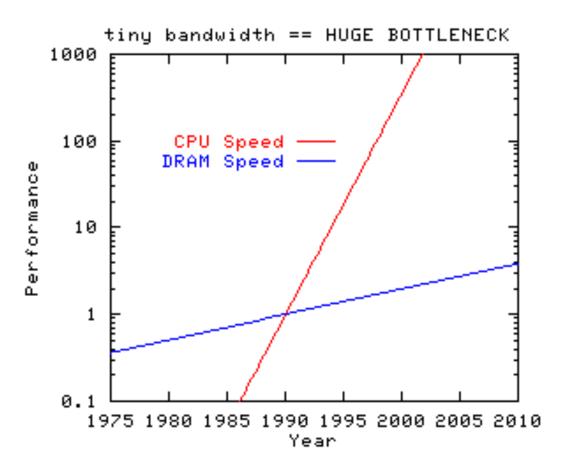
150 Lines, 8 Loops





Data From: http://roylongbottom.org.uk/cpumix.htm





http://www.cs.virginia.edu/stream/ref.html



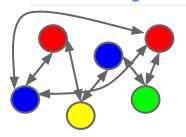
1985

An Empirical Study of Distributed Application Performance Keith A. Lantz, William 1. Nowicki, and Marvin M. Theimer Report No. STAN-CS-86- 1117 Also numbered CX-85-257

A major reason for the rarity of distributed applications, despite the proliferation of networks, is the sensitivity of their performance to various aspects of the network environment. We demonstrate that distributed applications can run faster than local ones, using common hardware ...

History of Cloud vs Benchmarking













1991	1992	1994	1995	1997
MPI created	Mosaic Browser born	Beowulf cluster	Internet has ~6M hosts	Ramnath K. Chellappa defines Cloud Computing

1992 - TPC-C Benchmark

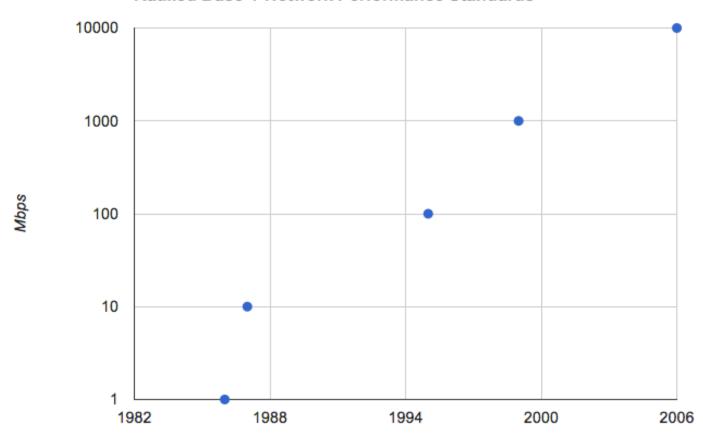
1994 - VideoWin Benchmark measures graphics speeds in terms of millions of pix per second

1995 - Spec95 - Go/Compiler/Lisp/Image/DB/Perl/Mesh/Water/DiffyQ/Gauss/Maxwell

1997 - Specglperf - points, lines, triangles, pixels, etc



Ratified Base-T Network Performance Standards



Data From: http://en.wikipedia.org/wiki/10BASE-T

Latency you can understand

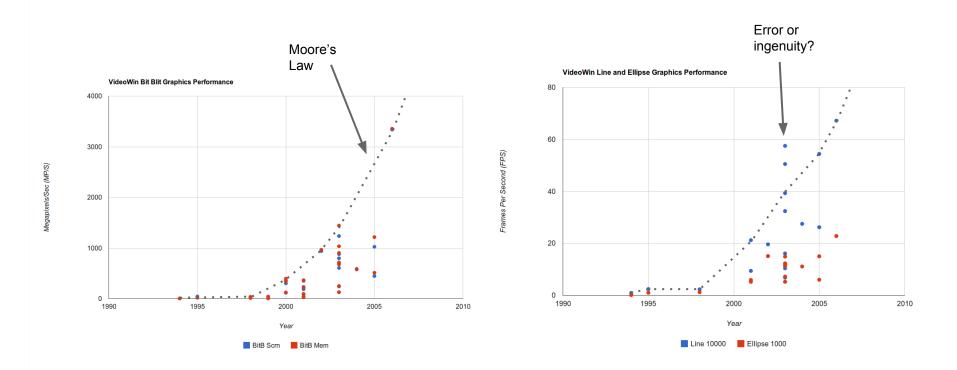




http://www.eecs.berkeley.edu/~rcs/research/interactive_latency.html

VideoWin Benchmark





Data From: http://roylongbottom.org.uk/videowin%20results.htm

History of Benchmarking













2000	2001	2003	2006	2008
The internet has ~72M hosts	461 million people are connected to the internet SaaS is born	Memcache Born	Eric Schmidt uses "cloud computing" to refer to SaaS. (ref) The internet has ~395M hosts Amazon Web Services (AWS) is launched	Google App Engine Beta Chrome in Beta

2000 - [My]SqlBench / DBHammer / SQLIO

2000 - Spec2000 - compression/OODB/3G Graphics/High energy physics/...

2004 - HPL v1.0 - High-Performance Linpack Benchmark for Distributed-Memory Computers

2006 - Spec2006 - perl/xml/physics/weather/speech/ray tracing/...

2006 - 3DMark2006

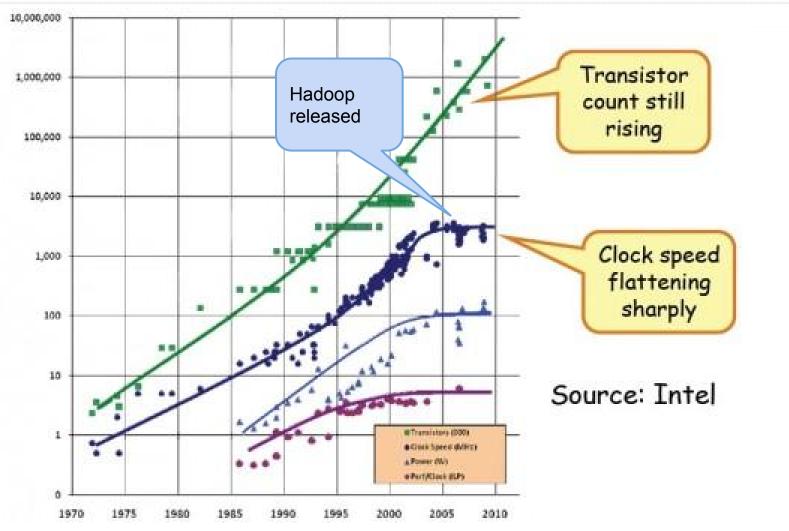
2007 - vConsolidate - The first virtualization benchmark

2007 - Hadoop Sort

2009 - First CoreMark result published

Core Performance





From: http://bertrandmeyer.com/2011/06/20/concurrent-programming-is-easy/. Blue annotations added for this preso.















2010

Google Launches
Google Cloud
Storage (laaS)
Price Per Watt
(PPW)

2011

September Google
App Engine
releases

2012

Google Compute Engine
(GCE) comes out in
preview (laaS)
Google BigQuery
Google Cloud SQL

2013

GCE Load Balancer
Google Cloud Datastore

2010 - SpecViewperf - Model based Graphics Benchmark

2010 - SpecVirt_sc2010 - SPECweb2005, SPECjAppServer2004, and SPECmail2008

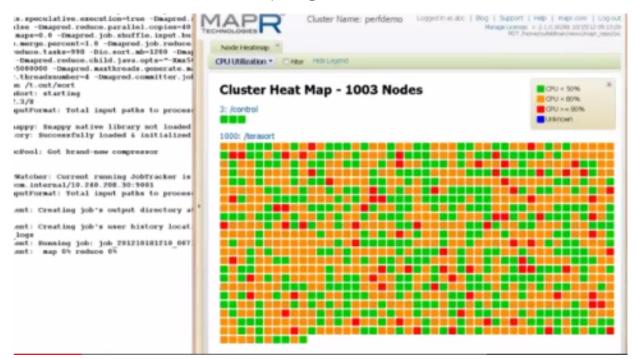
2011 - Cassandra - NoSQL performance

2013 - SpecVirt_sc2013 - SPECweb2005, SPECjAppServer2004, SPECmail2008, SPEC CPU2006



54 Seconds

1,003 servers 4,012 cores 1,003 disks



http://insights.wired.com/video/mapr-google-compute-engine-set-new-world-record-for-hadoop-teraso#axzz2fmZrtfxF

Off base!

600916

Performance Measurement Steps



Step 1: Setup

Step 2: Warm up

Step 3: Pre-execute

Step 4: Execute

Step 5: Post-Execute

Step 6: Cleanup

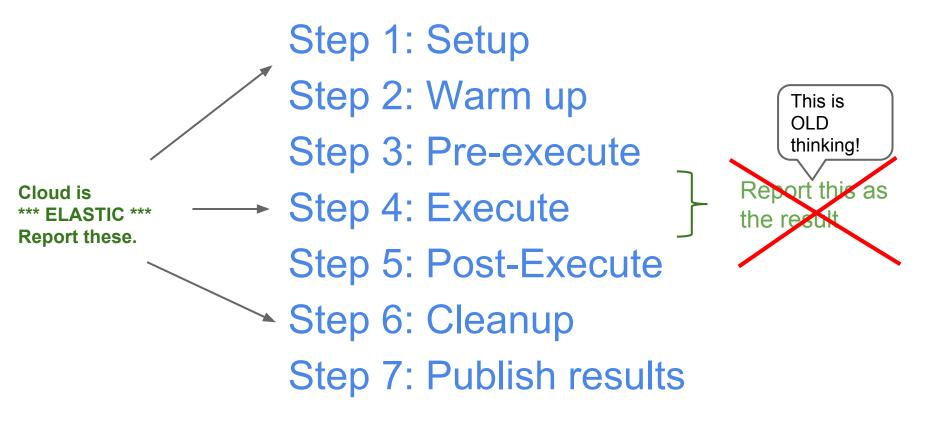
Step 7: Publish results

Report this as the result

*** Thanks to Chris Ashton who came up with this with me years ago.

Performance Measurement Steps





Performance Measurement Key Metrics Google



- Throughput
- Latency
- Overhead (CPI + Power Consumption)
- Cost!

Fight the black hole

Fight the black hole





Playbook



The Basics

- Network
- Storage
- Memory
- CPU
- [Graphics]

Meta Operations

- Start / Stop VM + Cluster
- Snapshot a workload
- Provision a disk
- Query status

New Primitives

- No SQL
- Memcache
- Hadoop [Sort]
- Boundary+Cluster network ops

The Metrics

- Throughput
- Latency
- Overhead (CPI + Power Consumption)
- Cost!

End - Questions?



Name: Anthony F. Voellm (aka Tony)

Contact: voellm@google.com

Blog: http://perfguy.blogspot.com

G+: http://goo.gl/mPXcX

Twitter: @p3rfguy

Appendix

References



- http://www.roylongbottom.org.uk/whetstone.htm
- http://www.roylongbottom.org.uk/whetstone.pdf
- http://en.wikipedia.org/wiki/Graphics processing unit
- http://en.wikipedia.org/wiki/Benchmark (computing)
- ftp://reports.stanford.edu/pub/cstr/reports/cs/tr/86/1117/CS-TR-86-1117.pdf
- http://en.wikipedia.org/wiki/Graphics_processing_unit
- http://en.wikipedia.org/wiki/Beowulf_cluster
- http://www.spec.org
- http://www.eembc.org/
- http://www.netlib.org/benchmark/hpl/
- http://www.opengl.org/resources/benchmarks/
- http://www.intel.com/pressroom/archive/releases/2007/20070417gloc1.htm
- http://www.spec.org/virt_sc2010/
- http://en.wikipedia.org/wiki/Frame_rate
- http://techblog.netflix.com/2011/11/benchmarking-cassandra-scalability-on.html
- http://www.clusters4all.com/tech/scale.html
- http://en.wikipedia.org/wiki/10BASE-T
- http://en.wikipedia.org/wiki/No_sql
- http://en.wikipedia.org/wiki/Memcache
- http://en.wikipedia.org/wiki/100 Gigabit Ethernet
- http://bertrandmeyer.com/2011/06/20/concurrent-programming-is-easy/

Cloud Benchmarking: Fight the black hole



Anthony F. Voellm, is currently leading the Google Cloud Performance Team and has a wide range of experience from kernel and database engines to graphics and automated image and map extraction from satellite images. Anthony is an avid inventor with 7 technology patents issued. In his current role at Google Anthony is focused on delivering Prise Performance to existing products like Google Compute Engine and Google Cloud Storage while also innovating new offerings. Prior to joining Google in 2011, Anthony held multiple roles at Microsoft leading the Microsoft Windows Reliability, Security, and Privacy test team working on Windows7+; Microsoft Hyper-V Performance Team; and SQL Server Performance team. He has also been a developer and tester on the Windows Filesystem, SQL Server Engine, and SGI IRIX networking teams. Anthony has taught performance testing to over 2000 people worldwide and given dozens of informative talks on software fundamentals. He keeps a personal technology blog on software fundamentals at **perfguy.blogspot.com**. In addition to computer interests his passions lie in growing engineers, building things, and doing anything outdoors. Anthony holds a Master of Science from George Washington University, BA in Physics and a BS in Computer Science and Mathematics from the University of Vermont.

