

Standard Performance Evaluation Corporation

Kaivalya M. Dixit

***Chief Cook & Bottlewasher, SPEC
Server Performance, IBM***


HPCA - 7

***Performance SPECulations - Benchmarks,
Friend or Foe***

dixit@us.ibm.com

512-838-2081


**Monterrey, Mexico
January 23, 2001**



Standard Performance Evaluation Corporation

Topics

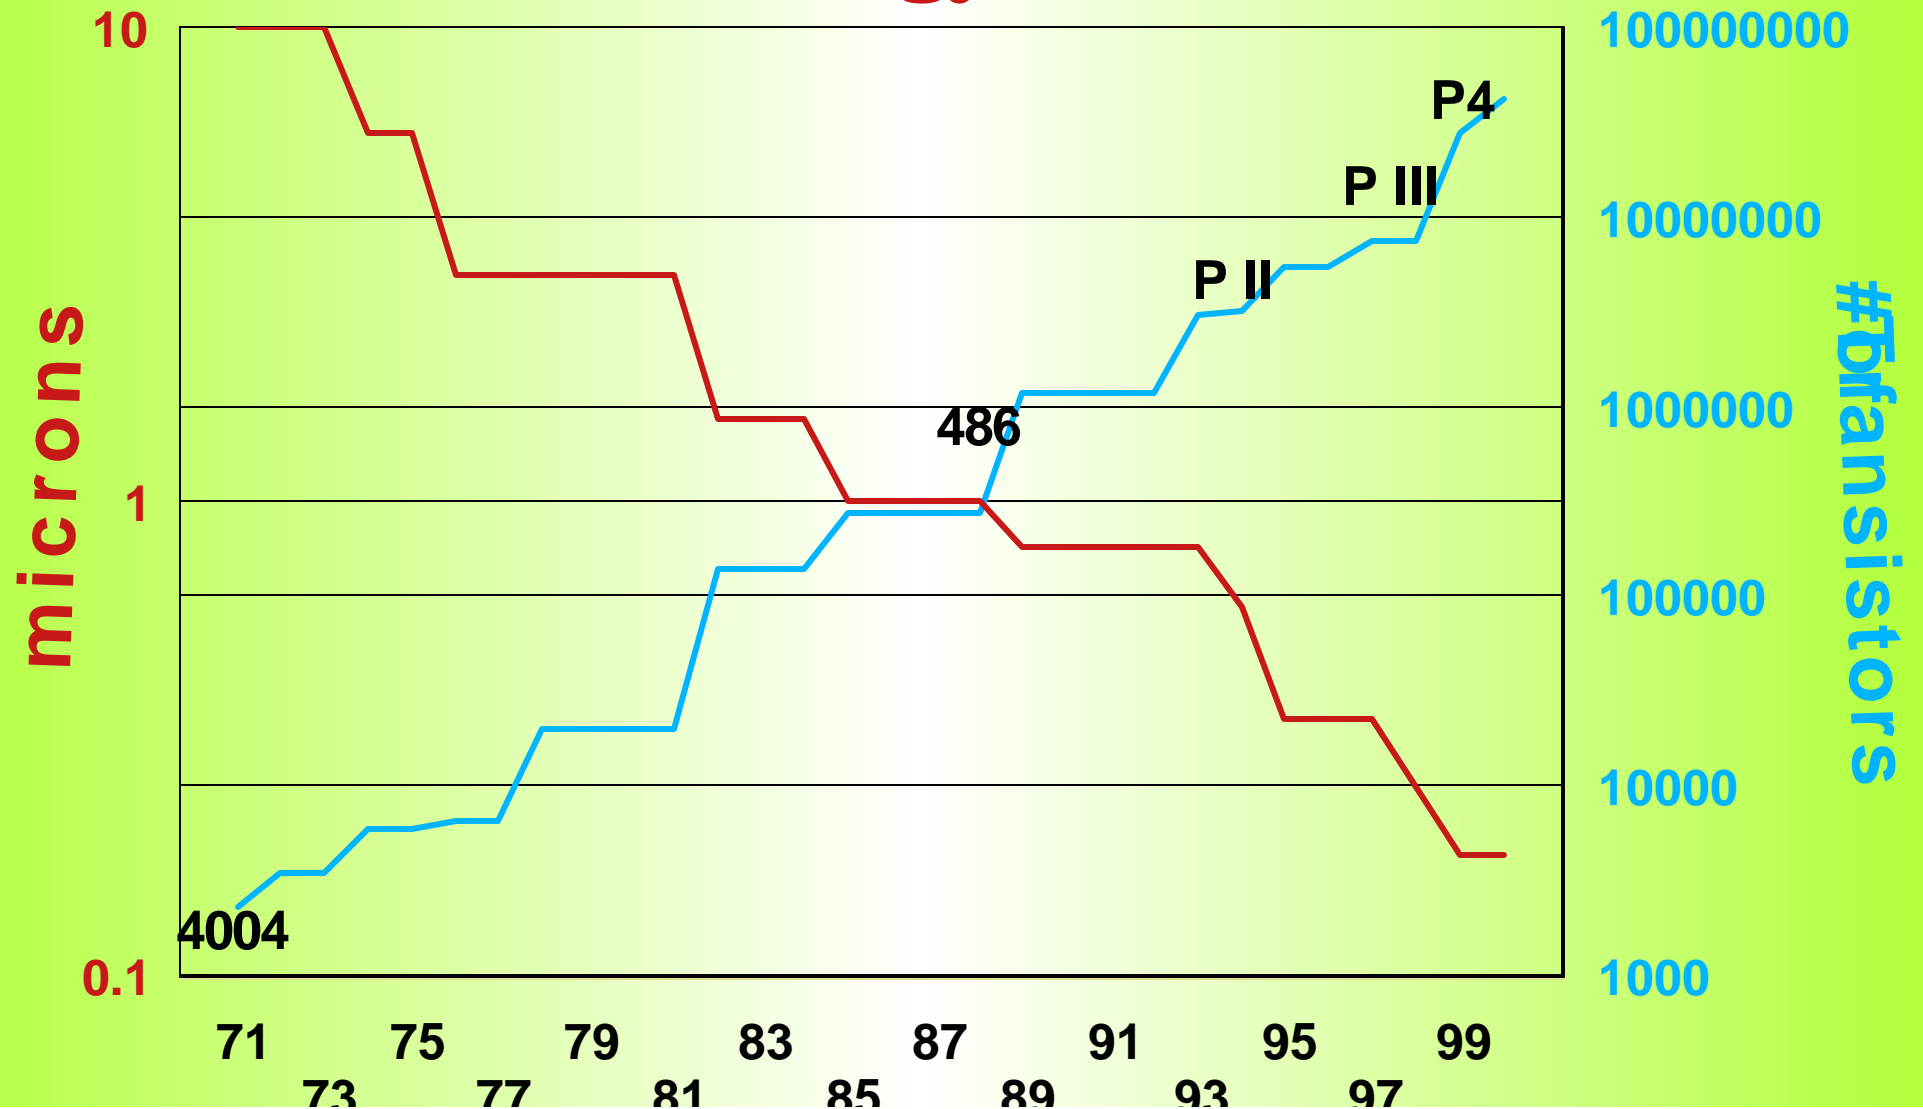
- ▶ Technology Trends
- ▶ Popular Benchmarks -- will they ever die?
- ▶ Applications vs. Benchmarks - one perspective
- ▶ Top 500 Super Computers
- ▶ Performance Problem
- ▶ Benchmark viewed by: Engr, Mktg, Standards, Boss & Users
- ▶ SPEC - infrastructure
- ▶ SPEC benchmarks
- ▶ CPU2K MPG - blind submission
- ▶ CPU2k MPG - one year later
- ▶ Coming attractions
- ▶ Issues
- ▶ Summary




Standard Performance Evaluation Corporation

Transistors (1971-2000)

Technology (1971-2000)







Standard Performance Evaluation Corporation

Frequency (1971-2000)

Technology (1971-2000)






Standard Performance Evaluation Corporation

Other Architectures - Today

Chip	#Transistors (Million)	Frequency (MHz)
Alpha	15.2	833
AMD	37	1200
HP	116	552
IBM	44	600
Sun	29	900
SGI	7.2	400


IBM POWER 4 has announced 170+ million transistors



Standard Performance Evaluation Corporation

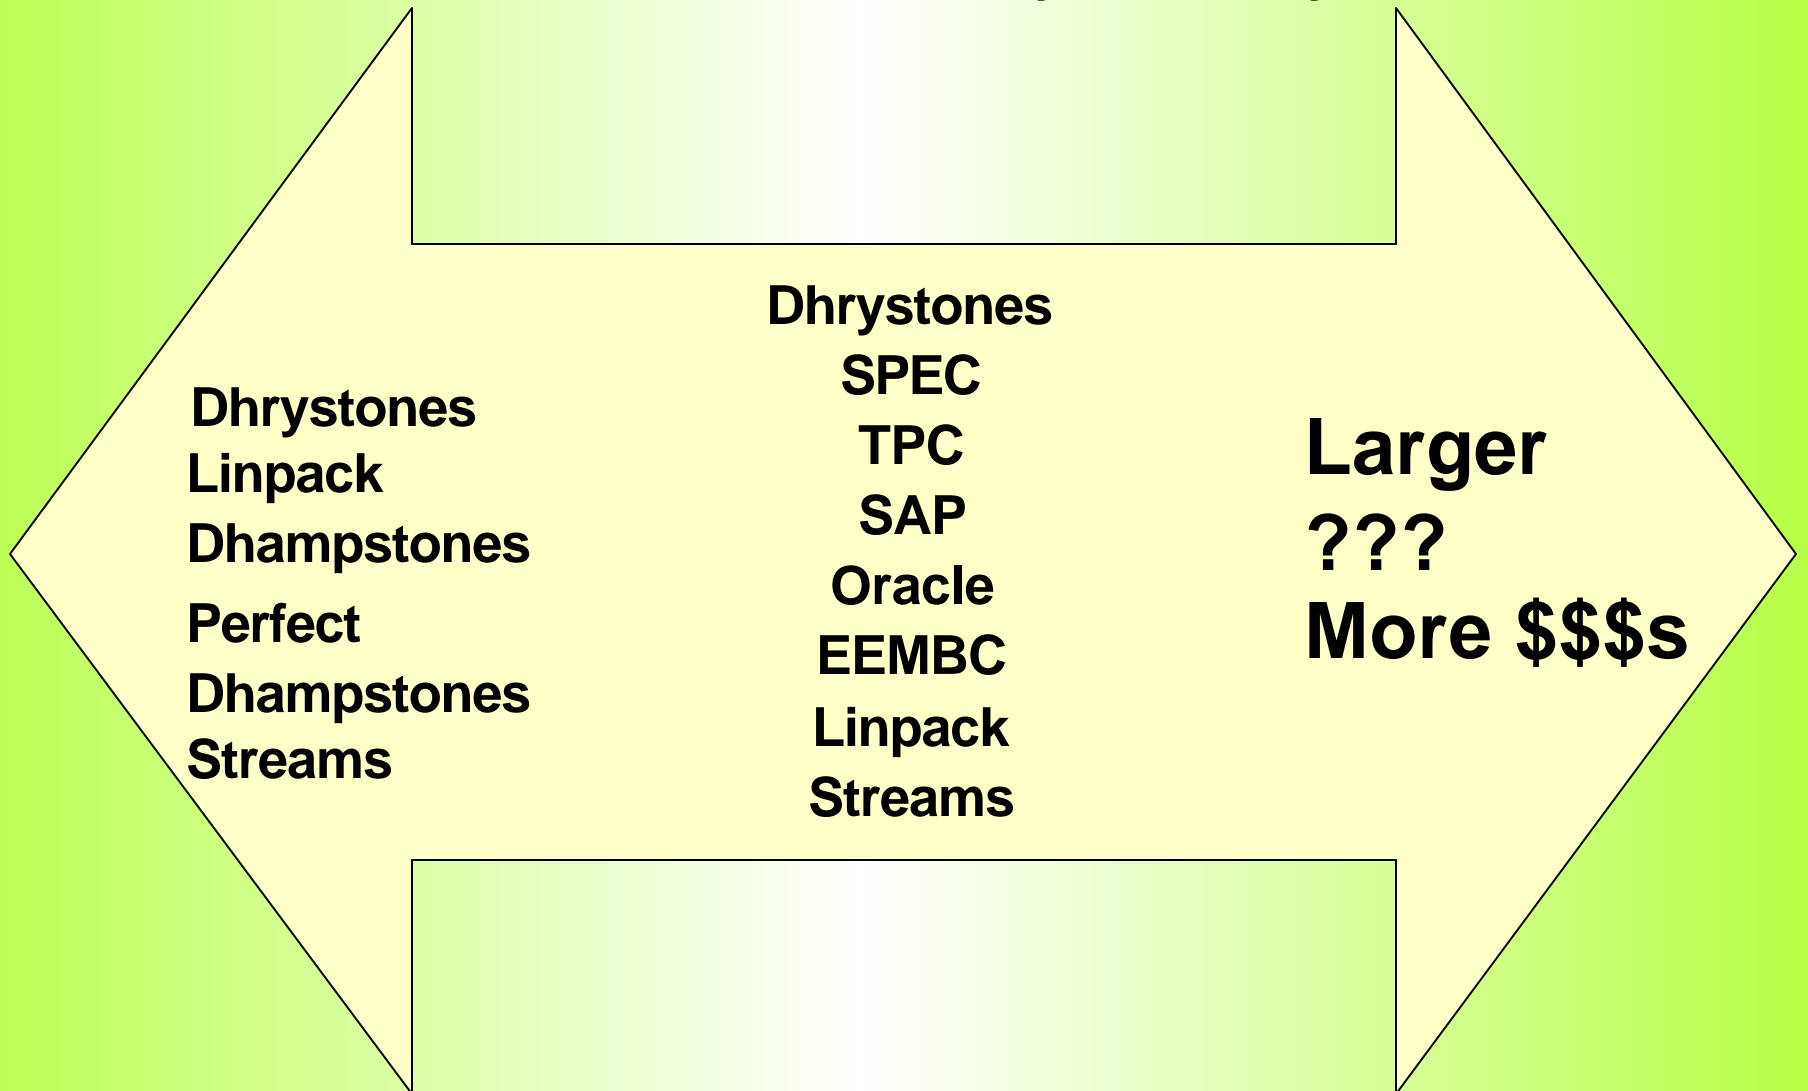
Tomorrow


- ▶ **Copper/silver, Silk, SOI, Electron Beam and other technologies may extend Moore's law for next 10 to 12 years**
- ▶ **Small Transistors (3 atoms)**
- ▶ **High Clockrates (10-20 GHz)**
- ▶ **System on a chip**
- ▶ **Custom processors**
- ▶ **High performance and Low power (oxymoron?)**
- ▶ **GIPs, Gflops and GDB available to everyone**
- ▶ **Challenge Amdahl's Law?**



Standard Performance Evaluation Corporation

Benchmarks - Yesterday, Today & Tomorrow






Standard Performance Evaluation Corporation

Applications vs. Benchmarks


- ▶ Linpack poor predictor ($2 \times$ Linpack = $\sim 10\%$ App. performance)
- ▶ SPECfp95 ($2 \times$ SPECfp95 = $2 \times$ App. performance)
- ▶ BW demand on SPECfp2000 higher than App.
- ▶ Lcache misses of SPECfp2000 an order of magnitude smaller than App.
- ▶ Mispredicted branches similar to App.
- ▶ Ratio of MemOPs to FPOPS 1-4 with mean of 2 for Apps. but 0.75 to 2 with mean of 1 for Benchmark
- ▶ We need to gather and share a lot more data from Apps.
- ▶ Reference:
<http://www.cs.varginia.edu/~mccalpin/wwc-keynote.html>



Standard Performance Evaluation Corporation

Highlights from the Top 100:


- ▶ **ASCI White is new #1 with 4.9 TFlops on Linpack**
- ▶ **DOE ASCI systems hold the first 4 positions**
- ▶ **7 systems have Linpack performance above 1 TFlops**
- ▶ **18 systems have peak performance above 1 TFlops including one commercial system (#15 at Charles Schwab)**
- ▶ **0.89 Tflops is the entry point for the Top 10**
- ▶ **Super Computer Rating either MIPS or MFLOPS -- why?**



Standard Performance Evaluation Corporation

Performance Equation


- ▶ **performance = f (MHz, Path Length, CPI,..)**
- ▶ **Perception: High Performance = High MHz**
- ▶ **CISC, RISC, VLIW, SMT, CMP, ...**
- ▶ **ISA, pipe-lining, super scalar, branch predictors, Reg. renaming, OOO, caches, load/store reordering, prefetch, memory, ...**
- ▶ **Technology, compiler, language,**
- ▶ **SMP, NUMA, MPP, Clusters, ...**
- ▶ **Benchmarks vs. Applications**



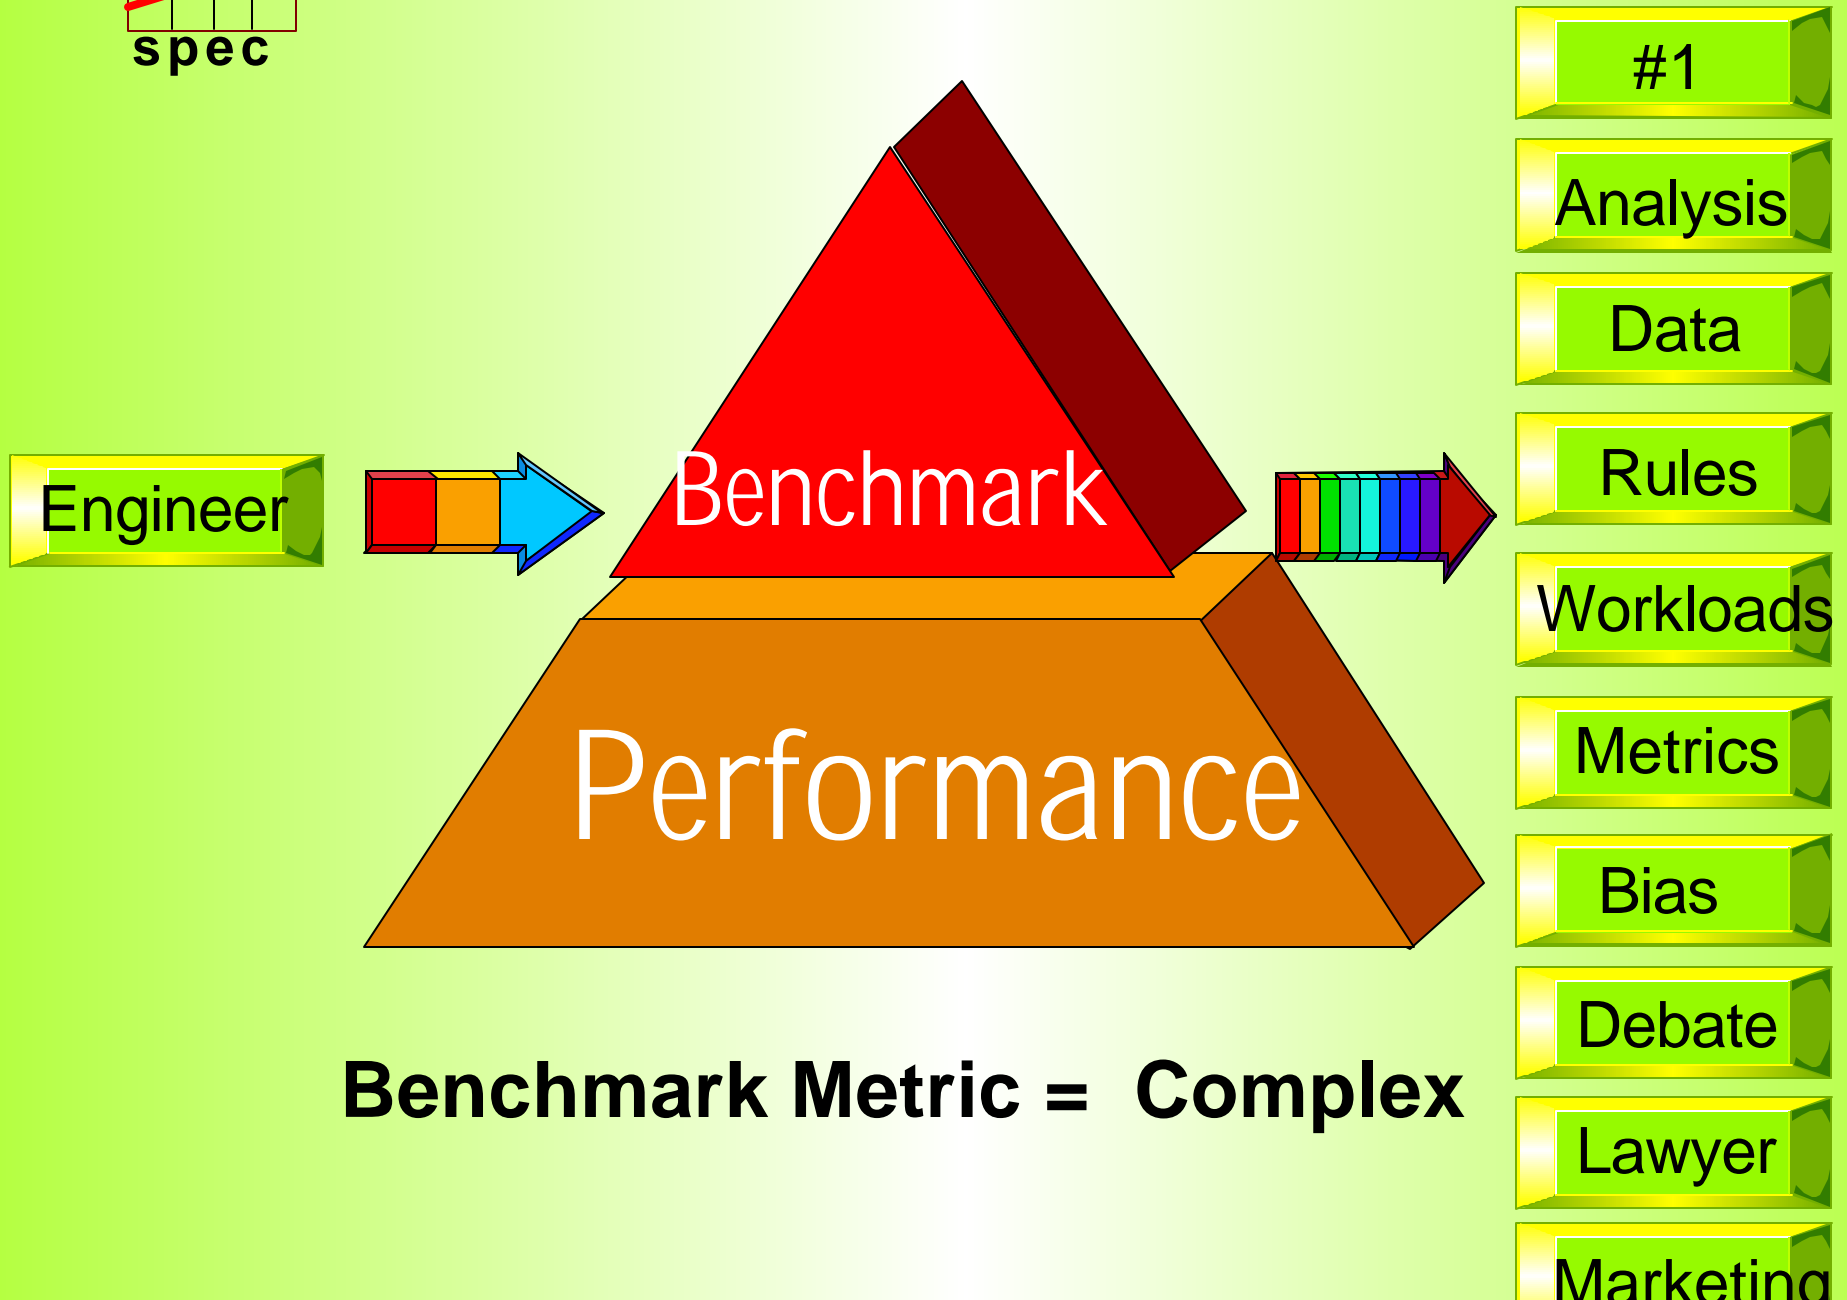
Standard Performance Evaluation Corporation

Benchmark Influence


- ▶ Linpack inspired fma, load / store units, & compiler tricks
- ▶ Dhrystone inspired string instructions
- ▶ Stream influenced CPU to memory designs
- ▶ SPEC inspired preprocessors, improved cache, branch predictors, OOO, ISAs, optimizers for graphics, web, compilers, Libs and FDOs
- ▶ TPC drove improvement of OS, database & I/O
- ▶ Both TPC & SPEC have helped improve H/W and SW scalability
- ▶ All Benchmark Specials are BAD




Standard Performance Evaluation Corporation




Benchmark Metric = Complex



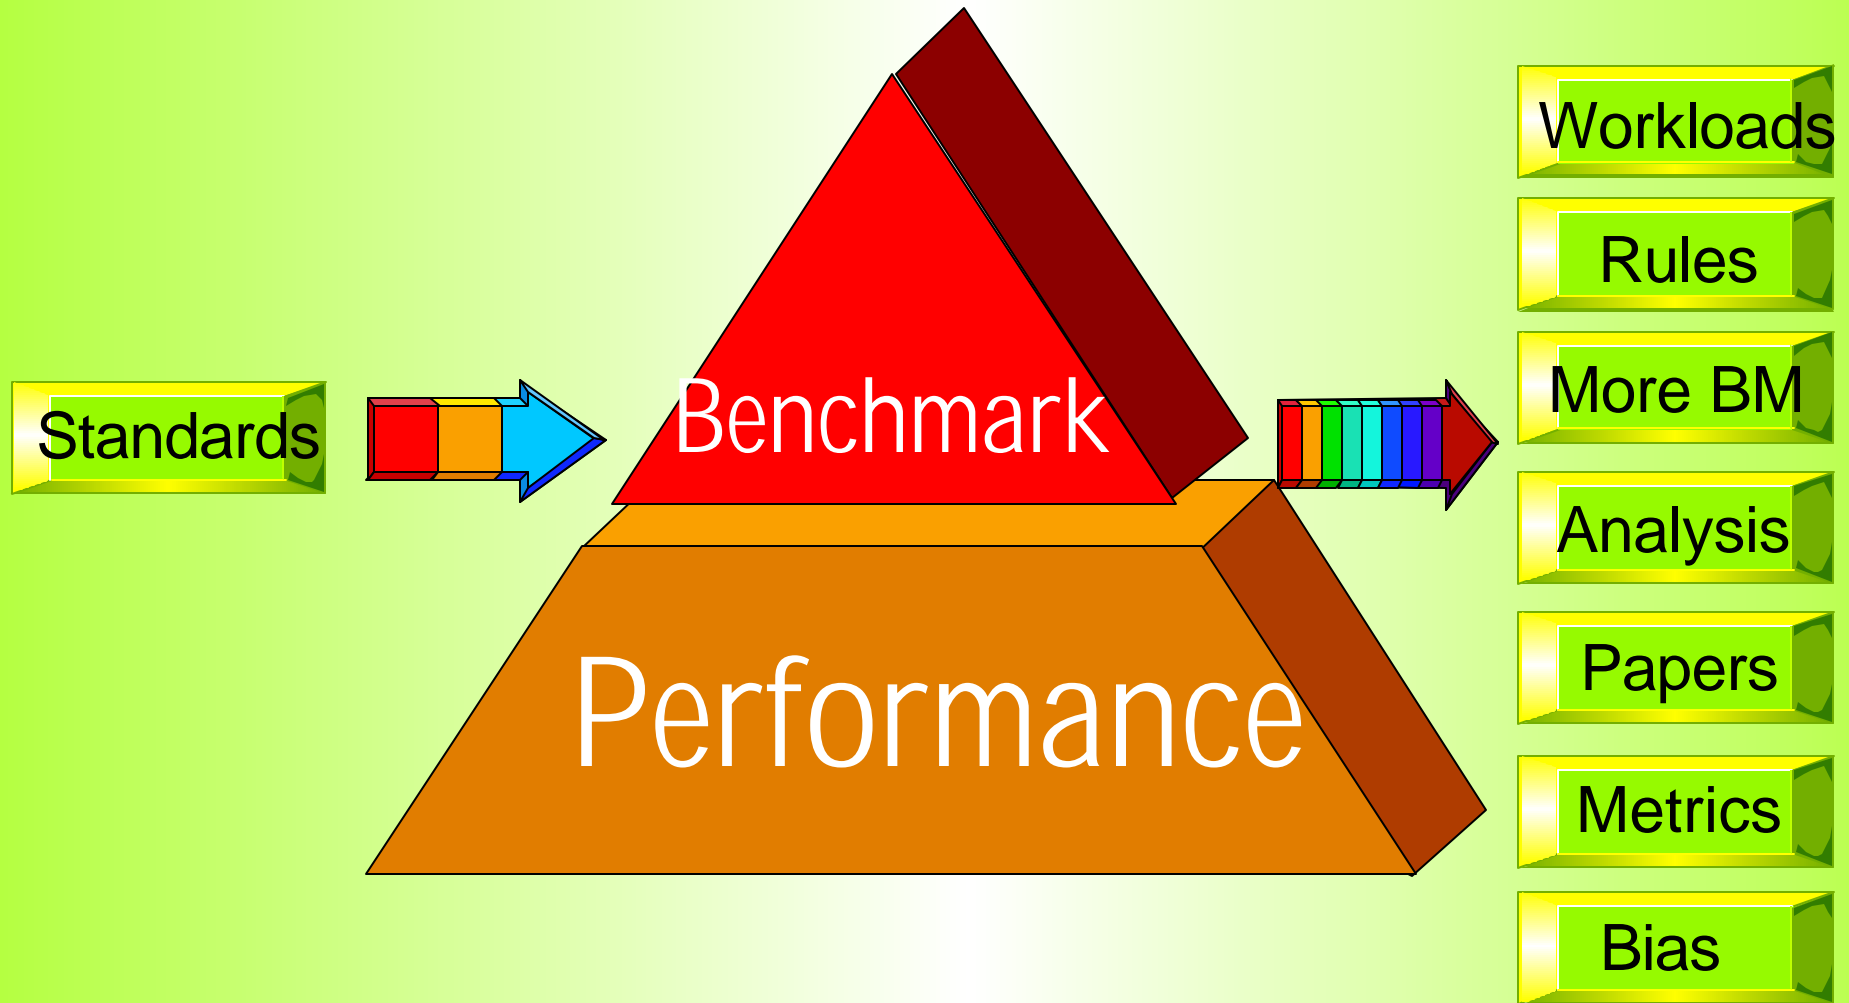
Standard Performance Evaluation Corporation




Benchmark Metric = Single Number




Standard Performance Evaluation Corporation




Benchmark Metric = Full Disclosure




Standard Performance Evaluation Corporation



Benchmark Metric = Don't care




Standard Performance Evaluation Corporation



Benchmark Metric =





Standard Performance Evaluation Corporation

Infrastructure

Board of Directors
Legal, HQ, PR


OSG

GPC

HPG

**Committees & People Who Do Real
Work!**

60+ Members and Associates




Standard Performance Evaluation Corporation

Popular Benchmarks

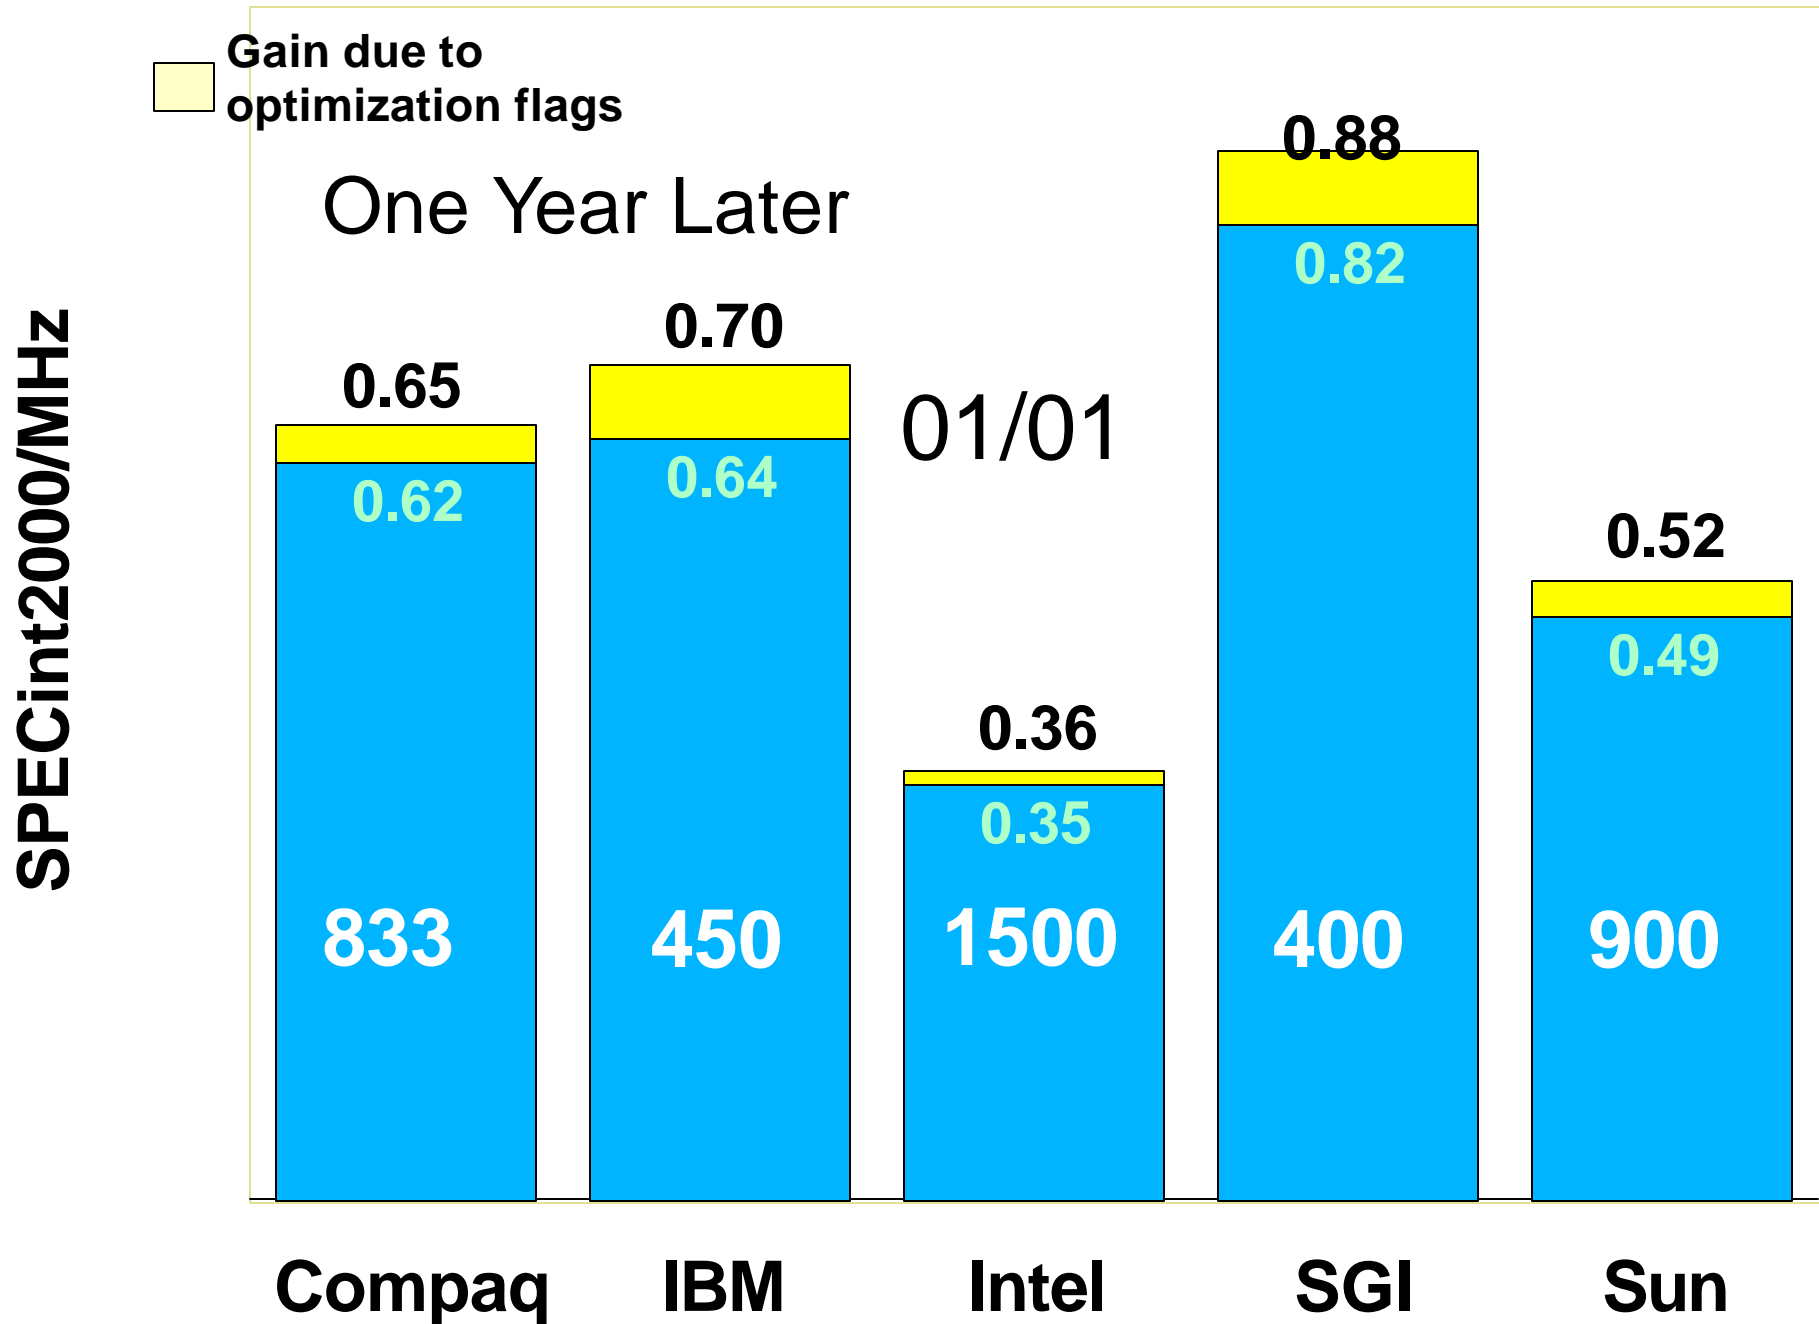
OSG	GPC	HPG
SPECjbb2000	SPECviewperf 6.11	SPEC HPC96 V1.10
SPEC CPU2000	SPECglperf 3.1.2	
SPECweb99	SPECapc for Pro/ENGINEER 2000i	
SPECjvm98	SPECapc for SolidWorks99	
SPECsfs97	SPECapc for Unigraphics V15	

Retired: SPEC89, 92, 95, SPECweb96 ..

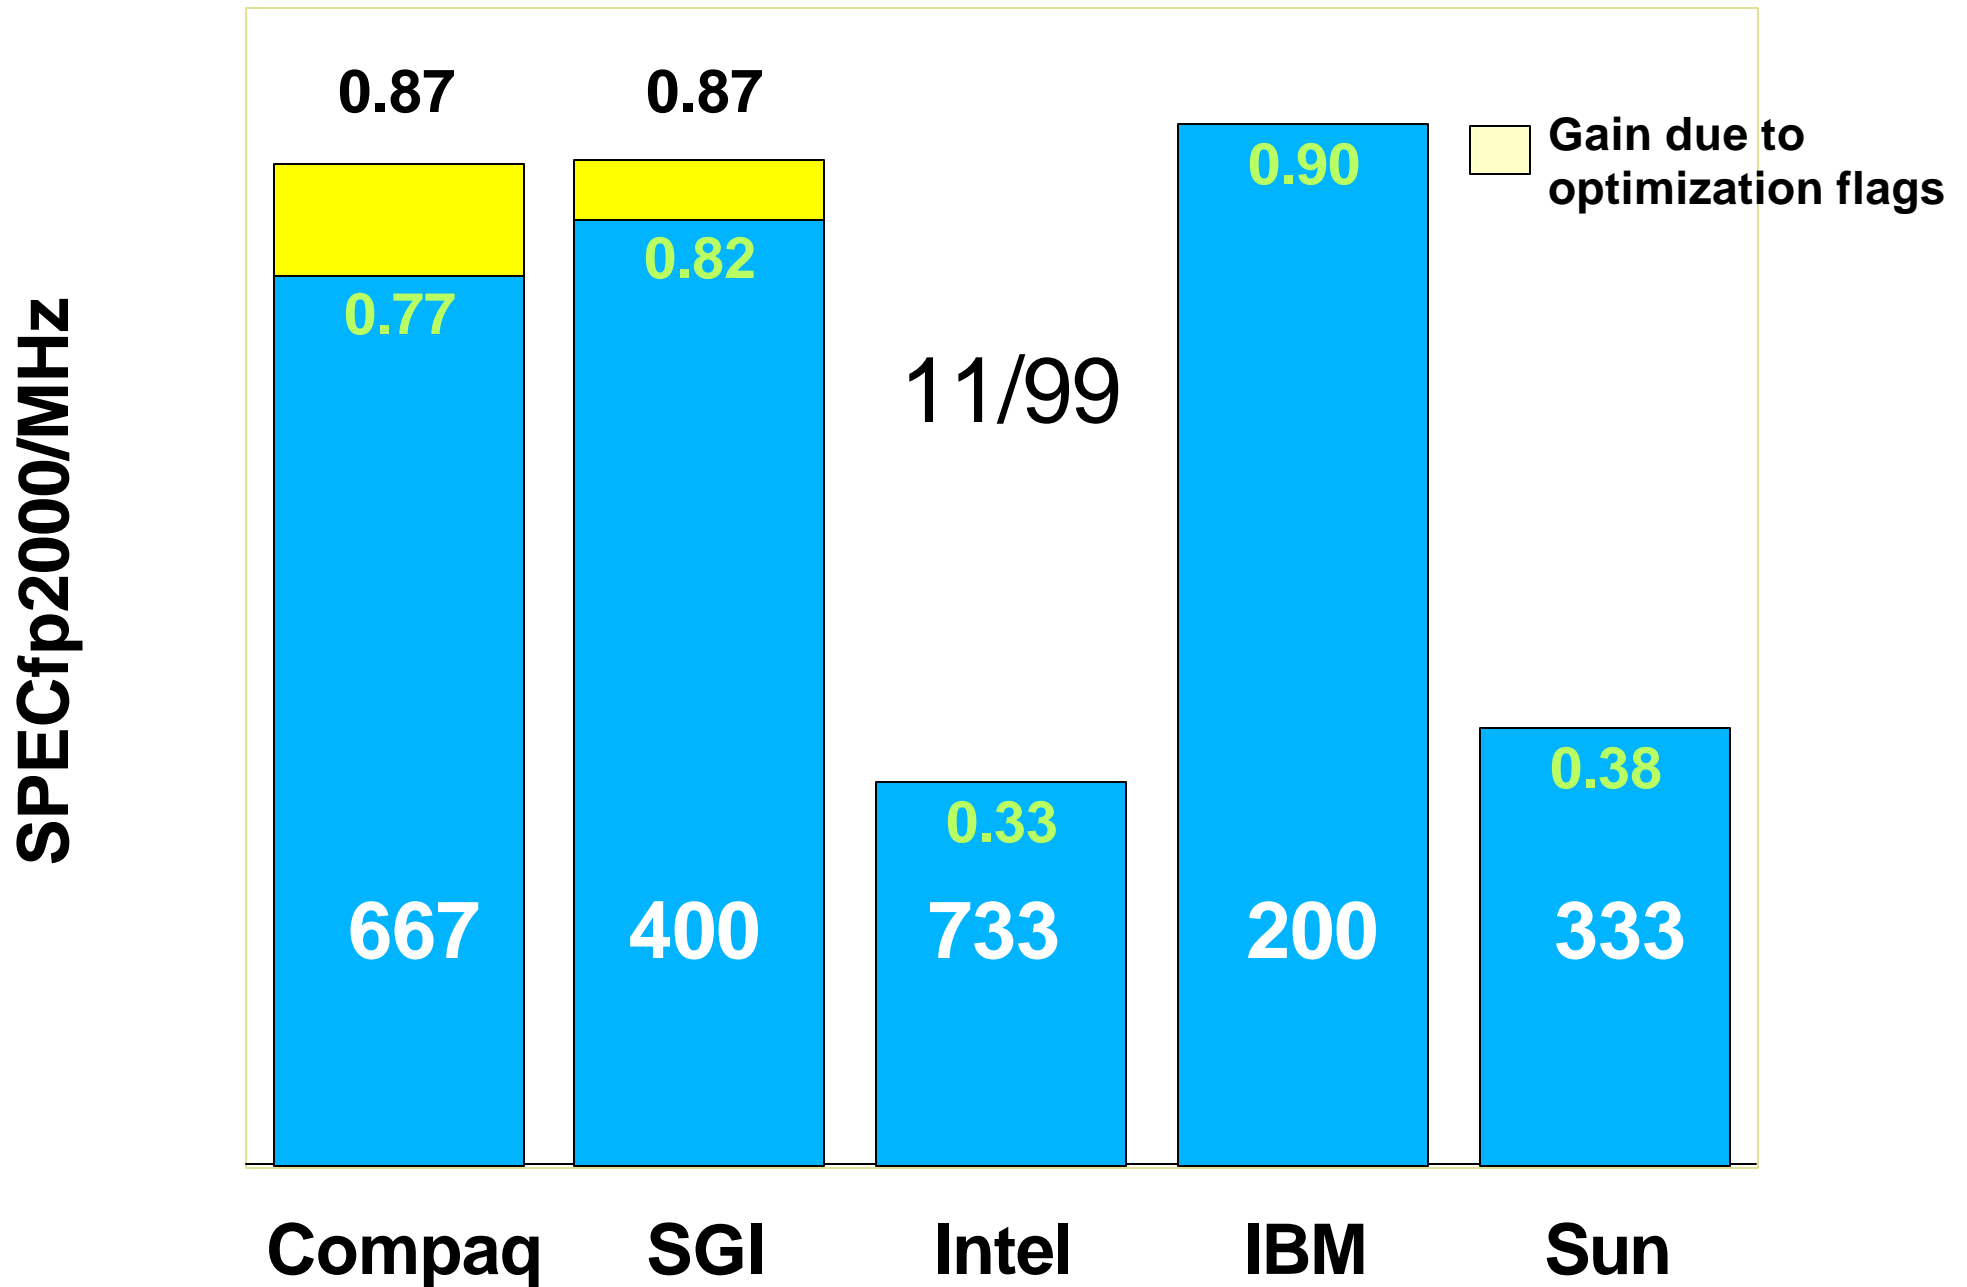
SPECint2000/MHz



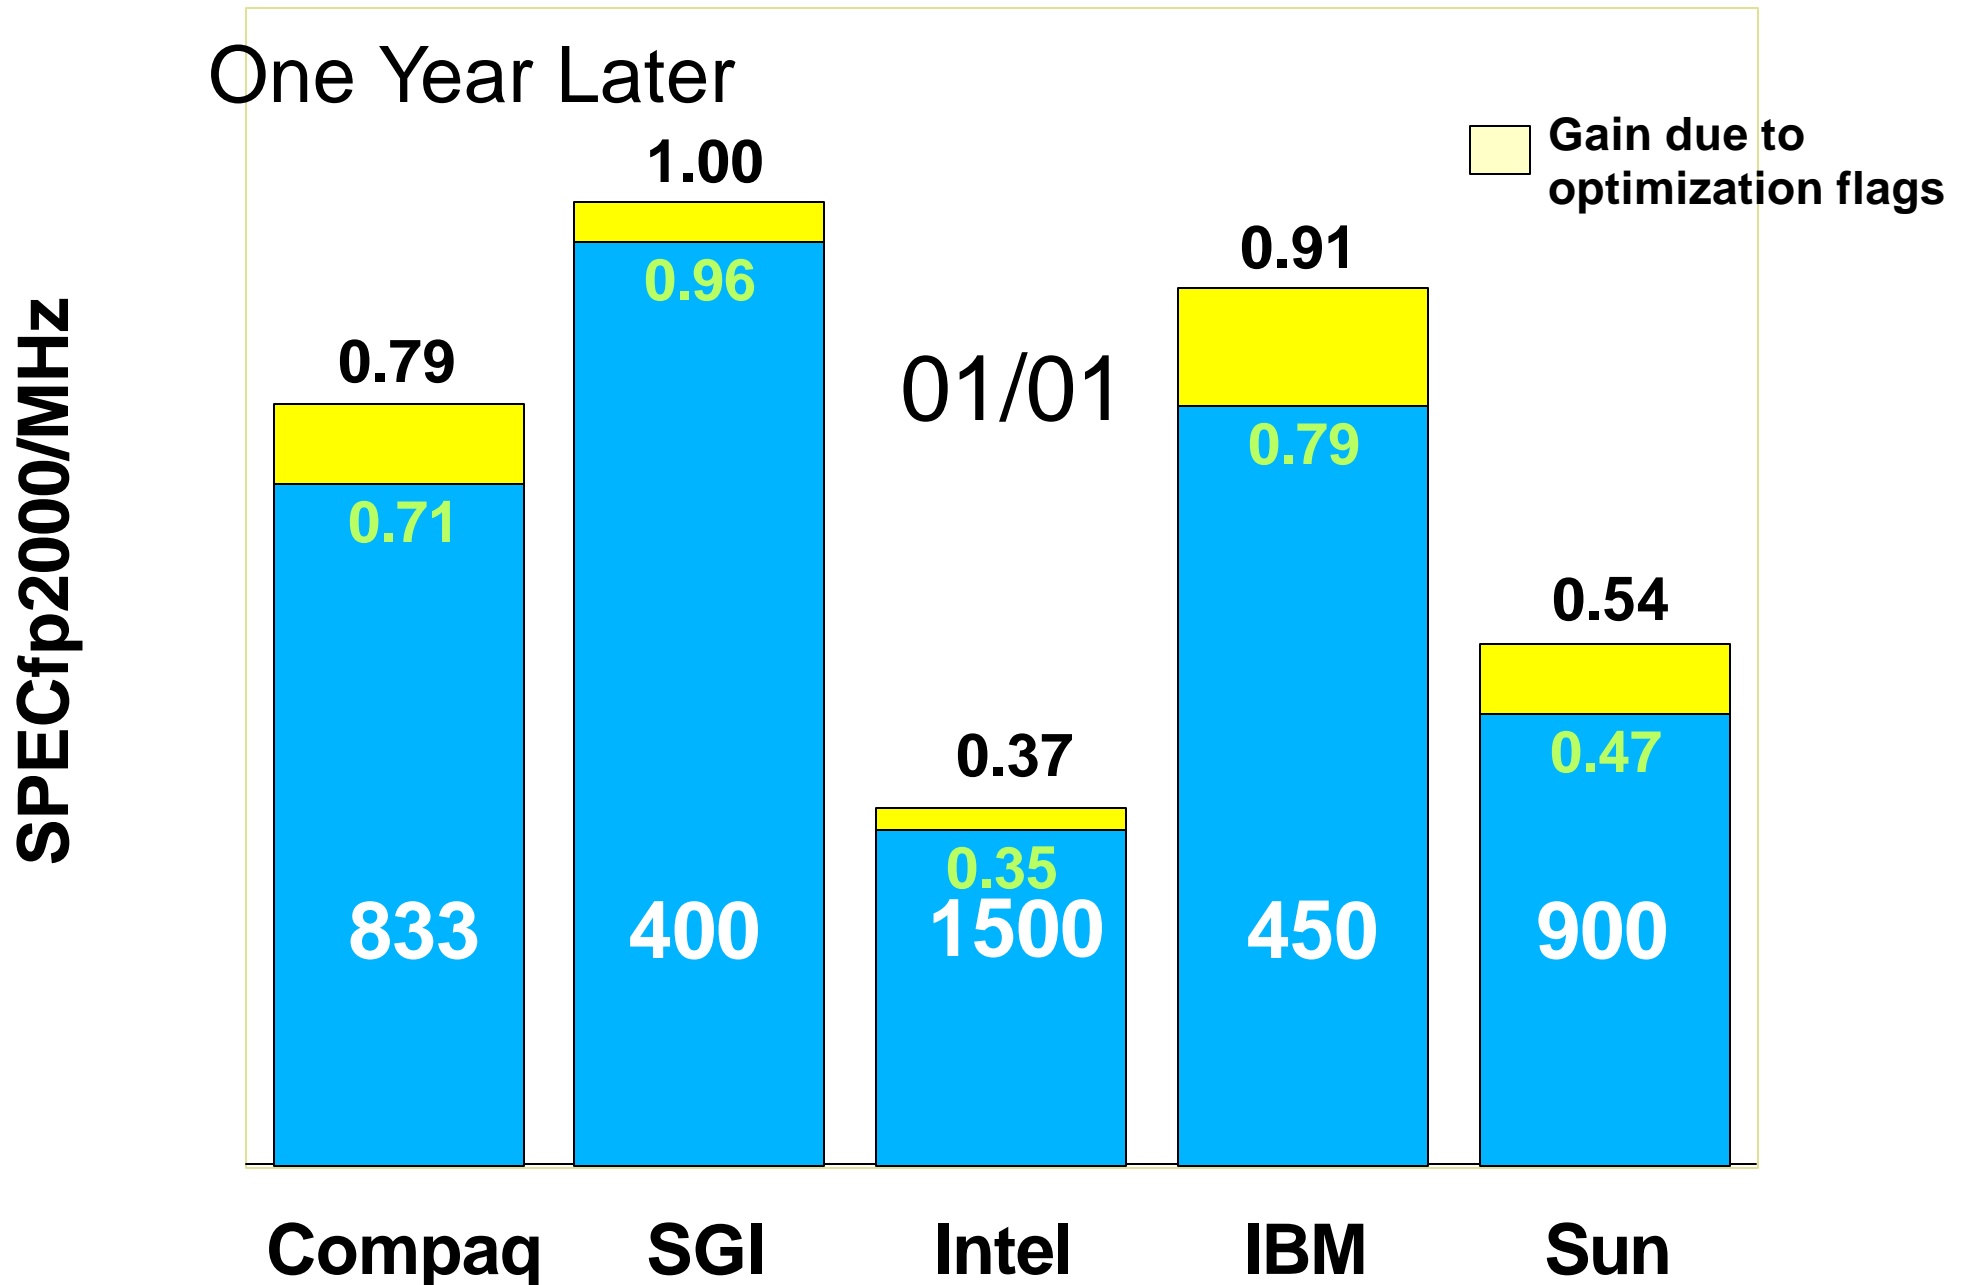
SPECint2000/MHz




SPECfp2000 / MHz



SPECfp2000 / MHz






Standard Performance Evaluation Corporation

Coming Attractions


- ▶ Released 01/19/2001 - SPECmail2000
- ▶ Server side Java - J2EE Benchmark
- ▶ High Performance Computing - SPECComp2000
- ▶ SPEC Annual Meeting (Jan 24, 2001) will be held in Austin, TX
- ▶ Searching for CPU 200X benchmarks
- ▶ For latest information for SPEC and TPC go to <http://www.spec.org> & <http://www.tpc.org>



Standard Performance Evaluation Corporation

Issues

- ▶ **Benchmarking is getting more and more EXPENSIVE & DIFFICULT**
- ▶ **Benchmark development cycles are getting longer - technology racing ahead**
- ▶ **Biggest problem with benchmarks is they represent yesterday's workloads**
- ▶ **Benchmarks help to MOVE bottlenecks not REMOVE them**
- ▶ **Benchmarking takes over benchmarking**
- ▶ **Old benchmarks are good members of NRA (Not Relevant Application) club - vigilance necessary**



Standard Performance Evaluation Corporation

Summary

- ▶ **GIPs, GFLOPs & GDB used for simpler MMI**
- ▶ **Good Benchmarks will improve H/W & S/W**
- ▶ **Biggest problem with benchmarks is that often times they represent yesterday's workloads**
- ▶ **Benchmarks help to MOVE bottlenecks not REMOVE them**
- ▶ **Best way to help is provide what you think is a good benchmark to standards organization**
- ▶ **SPEC has a bounty program and we are working on future benchmarks**
- ▶ **Q & A**