

Scribing - September 16th

News

- AMD Fx - “Fastest” CPU (in terms of frequency)
- 8.722 GHz overclocked
- FX-8370, 4.3GHz turbo, TDP 125W, Piledriver x86 core architecture
- High voltage (2V)
- Integrated liquid cooling
 - Pre-filled, micro-channels, integrated pump

Performance Analysis and Its Impact on Design

- IEEE 1998
- Importance of performance analysis now and throughout history
 - Analysis vs. design, which is superior?
- “Architecture should not write checks that designers cannot cash”
 - Do architects know what’s possible and what constraints are there?

Design Process

- Lead architects consider bounds of potential design
- Performance team creates performance model
- Performance architects create test cases
 - Test model

Terminology

- Path length = instruction count
- Separable components of CPU (Phil emma)
 - $CPI = \text{Infinite-cache-CPI} + FCE$
 - $\text{Infinite-cache-CPI (no stall)} = e_busy + e_idle$
 - FCE (finite cache effect, miss penalty)
 - $\text{Miss rate} = \text{cycles per miss} \times \text{misses per instruction}$
 - Linear relationship between CPI and misses/instruction

Performance Validation

- Generate performance test cases
- Early tests are random
- After failing tests are below threshold, move to focused and handwritten tests
- Latency tests
- Cost estimate for large basic-blocks is based on program dependence graphs
- Best/worst case timings for a basic-blocks can be used as test cases
- Bandwidth tests

Performance Signature Dictionary

- Apart from specifications for cycle count
- Steady state loop performance
- Derive more elaborate performance signature
- Signatures are plots of various quantities that follow characteristic patterns
 - e.g. Periodic pattern of pipeline state

Trace Distortion

- Instruments can cause distortion due to invasiveness (alters behavior)
 - e.g. mtrace (tracing tool)
 - 60 times slower than un-instrumented
 - Collects I/D addresses
 - AIX clock interrupt occurs 100 times per second to wake scheduler (real world time)
 - However, since 60x slower, 60x more interrupts
 - Interrupts have to be slowed by 60x as well

Analysis of Redundancy and Application Balance in the SPEC CPU2006 Benchmark Suite

- Written when 2006 was being created/compiled
- The Laboratory for Computer Architecture at UT Austin advised/analyzed current benchmark choices
- Final paper written on final program selection
- SPEC
 - Open submission, wide-audience for programs (e.g. GCC)
 - Initial tests to eliminate programs
 - Performance data given to UT researchers, who did cluster analysis
 - SPEC made choices based on data

Motivation

- Many benchmarks are similar
- Running more benchmarks that are similar produces no gains
- Good benchmark suite should have good representations of different clusters
 - Reduces experimentation effort
 - Increases coverage

Reduce

- Measure properties (K)
 - microarchitecture independent properties
 - microarchitecture dependent properties
- Display properties (K)
- Workload space consist of clusters of benchmarks

- Choose 1 per cluster
 - Explore space
 - Might be impossible to have uniform distribution throughout the space
- Statistical techniques to combine data and dimensions
 - High dimensionality
 - Principal components (PC)
 - Principal component analysis (PCA)
 - Some variables/components have different variability, reduce to those which give the most variability
 - Capture most of the information
 - Look in the PC space
- Clustering analysis
 - K-means
 - Create K-clusters (or iterate through and find k with best error)
 - Procedure (k=3)
 - Randomly select 3 cluster centers
 - Assign benchmarks to cluster centers
 - Move centers to the center of clusters
 - Repeat until convergence
 - Hierarchical clustering
 - Linkage strategies
 - Farthest distance
 - Center
 - Euclidian distance ($\sqrt{a^2+b^2}$)
 - Manhattan distance (90 degree angles, $a+b$)
 - Procedure
 - Iteratively join clusters (each benchmark is a cluster)
 - Join the two closest by linkage strategy
 - Repeat
 - Dendrogram
 - Indicates how benchmarks are similar
 - Tells what benchmarks to use if only a limited number of benchmarks can be run
- SPEC removed most similar benchmarks

Software

- Statistica
- R
- MATLAB