Laboratory for Computer Architecture
The University of Texas at Austin
and IBM

Performance Characterization of SPEC CPU Benchmarks on Intel's Core Microarchitecture Based Processor

Sarah Bird, Aashish Phansalkar, Lizy K. John,
Alex Mericas, Rajeev Indukuru

Outline

- Motivation
- Objectives
- Methodology
- System Design and Details
- Performance Characterization Results of SPEC CPU Benchmarks
- Fusion Description and Results
- Conclusion
- Questions

Motivation

- Study the design of the Core Microarchitecture and it's new features to learn how they work
- Study the behavior of the SPEC CPU benchmark suites on the Core Microarchitecture
- Study the effect of the new features on the behavior of the SPEC CPU benchmarks

Objectives

- To analyze the behavior of the SPEC CPU2006 suite in comparison to the behavior of the SPEC CPU2000 suite on a Core Microarchitecture processor.

- To determine if fusion (macro and micro-op) contributed noticeably to the improved performance of the Core Microarchitecture processor as compared to its predecessors.

Methodology

- Run SPEC CPU2006 and CPU2000 benchmark suites on a Core Microarchitecture based processor*
- Use performance counters to collect information about the behavior of the benchmarks*
- Use data provided by performance counters to compare CPU2006 and CPU2000
- Use runtimes from the SPEC website for Core predecessors to determine the performance improvement for each benchmark on the Core Microarchitecture
- Compare the amount of fusion (macro and micro-op) measured by the performance counters to the calculated to performance improvement

*Steps performed by IBM

System

- **Woodcrest System**
  - Tyan S5380 Motherboard
  - 2 Xeon 5160 CPU's running at 3.0Ghz
  - 4x1GB memory DIMMS at 667Mhz

- **Benchmark Compilers**
  - Intel C Compiler for 32-bit applications, Version 9.1
  - Intel Fortran Compiler for 32-bit applications, Version 9.1

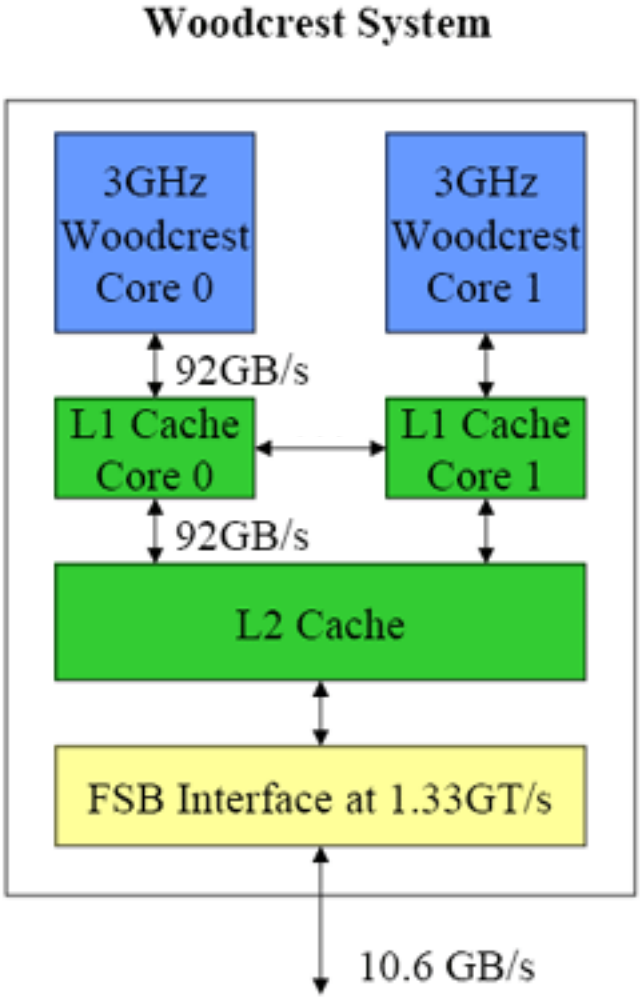

*Image taken from Real World Technologies "Intel's Next Generation Microarchitecture Unveiled"

System Details

- **L1 Cache**
  - 2 32KB caches
  - 8 way associativity
- **L2 Cache**
  - 1 unified 4MB cache
  - 16 way associativity
- **Macro-Fusion**
  - Fuses 2 x86 instructions
  - Compare and jump instructions
- **Micro-op Fusion**
  - Fuses 2 micro-ops
  - Store address and data micro-ops

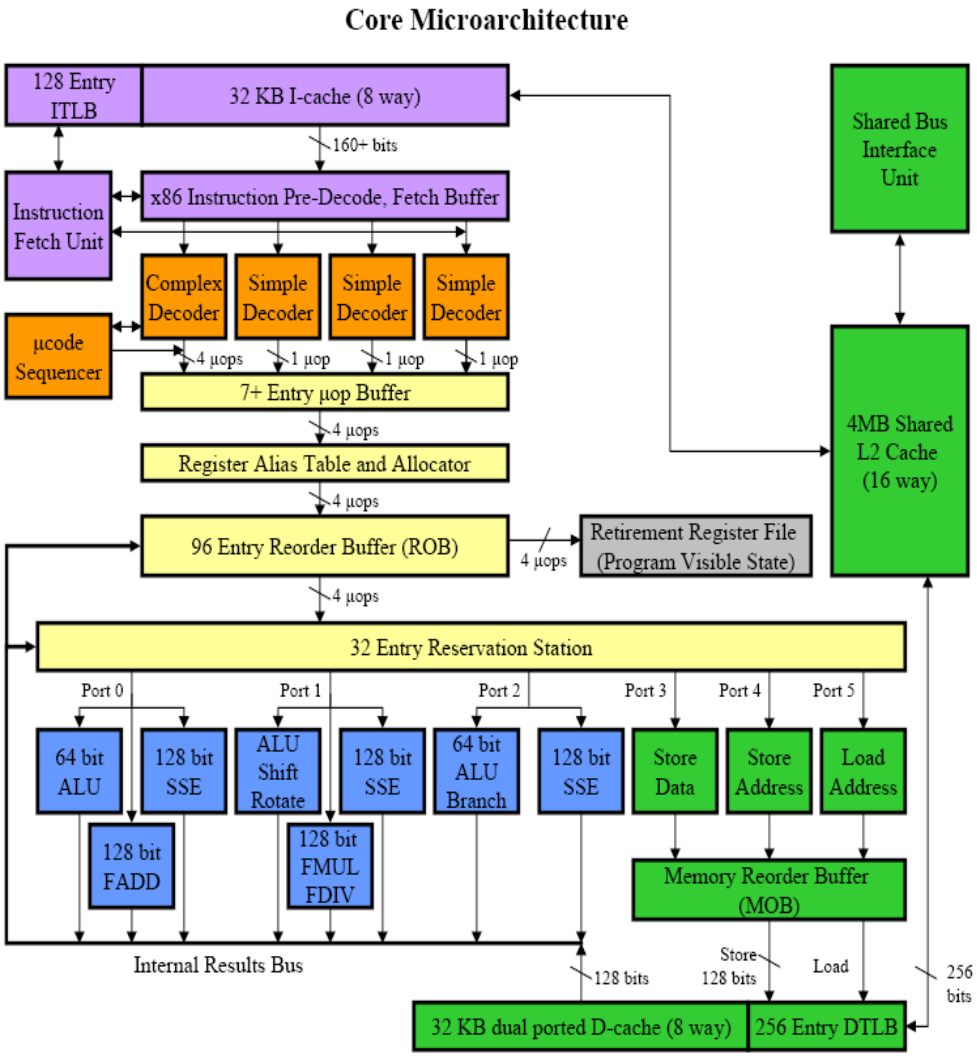

*Image taken from Real World Technologies "Intel's Next Generation Microarchitecture Unveiled"

Performance Characterization Results

Instruction Mix for CPU2006

SPEC CPU2006 Integer

Benchmark	% Branches	% Loads	% Stores
400.perlbench	23.3%	23.9%	11.5%
401.bzip2	15.3%	26.4%	8.9%
403.gcc	21.9%	25.6%	13.1%
429.mcf	19.2%	30.6%	8.6%
445.gobmk	20.7%	27.9%	14.2%
456.hmmer	8.4%	40.8%	16.2%
458.sjeng	21.4%	21.1%	8.0%
462.libquantum	27.3%	14.4%	5.0%
464.h264ref	7.5%	35.0%	12.1%
471.omnetpp	20.7%	34.2%	17.7%
473.astar	17.1%	26.9%	4.6%
483.xalancbmk	25.7%	32.1%	9.0%

SPEC CPU2006 Floating Point

Benchmark	% Branches	% Loads	% Stores
410.bwaves	0.7%	46.5%	8.5%
416.games	7.9%	34.6%	9.2%
433.milc	1.5%	37.3%	10.7%
434.zeusmp	4.0%	28.7%	8.1%
435.gromacs	3.4%	29.4%	14.5%
436.cactusADM	0.2%	46.5%	13.2%
437.leslie3d	3.2%	45.4%	10.6%
444.namd	4.9%	23.3%	6.0%
447.dealII	17.2%	34.6%	7.3%
450.soplex	16.4%	38.9%	7.5%
453.povray	14.3%	30.0%	8.8%
454.calculix	4.6%	31.9%	3.1%
459.GemsFDTD	1.5%	45.1%	10.0%
465.tonto	5.9%	34.8%	10.8%
470.ibm	0.9%	26.3%	8.5%
481.wrf	5.7%	30.7%	7.5%
482.sphinx3	10.2%	30.4%	3.0%

L1 data cache misses per 1000 Instructions

SPEC CPU2006

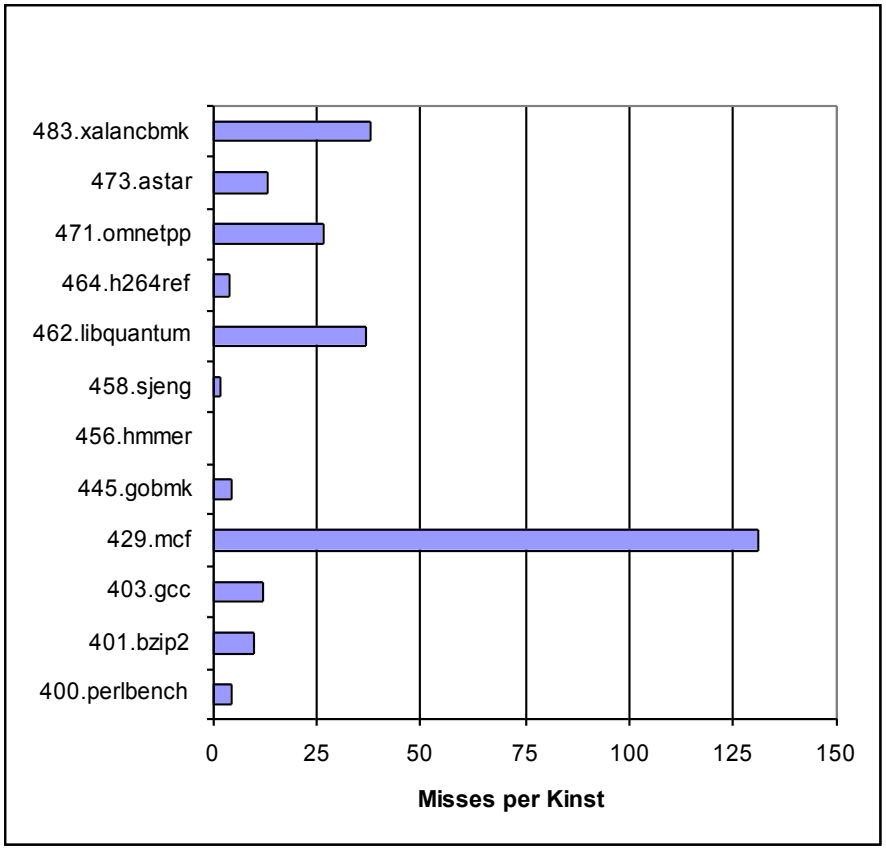

SPEC CPU2000

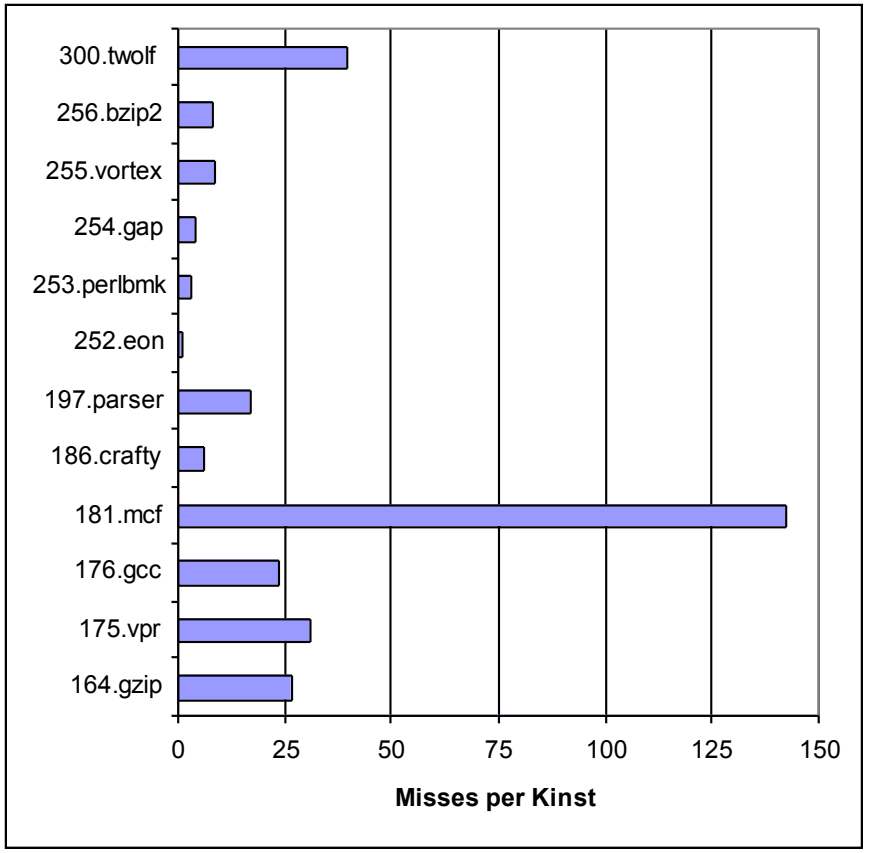

L2 cache misses per 1000 Instructions

SPEC CPU2006

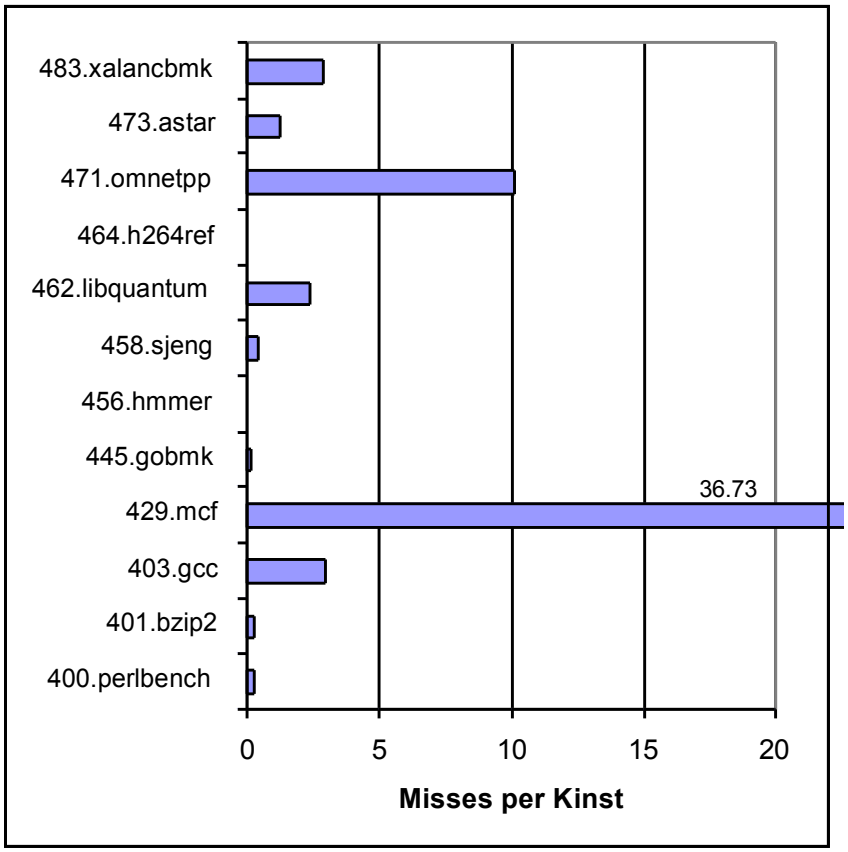

SPEC CPU2000

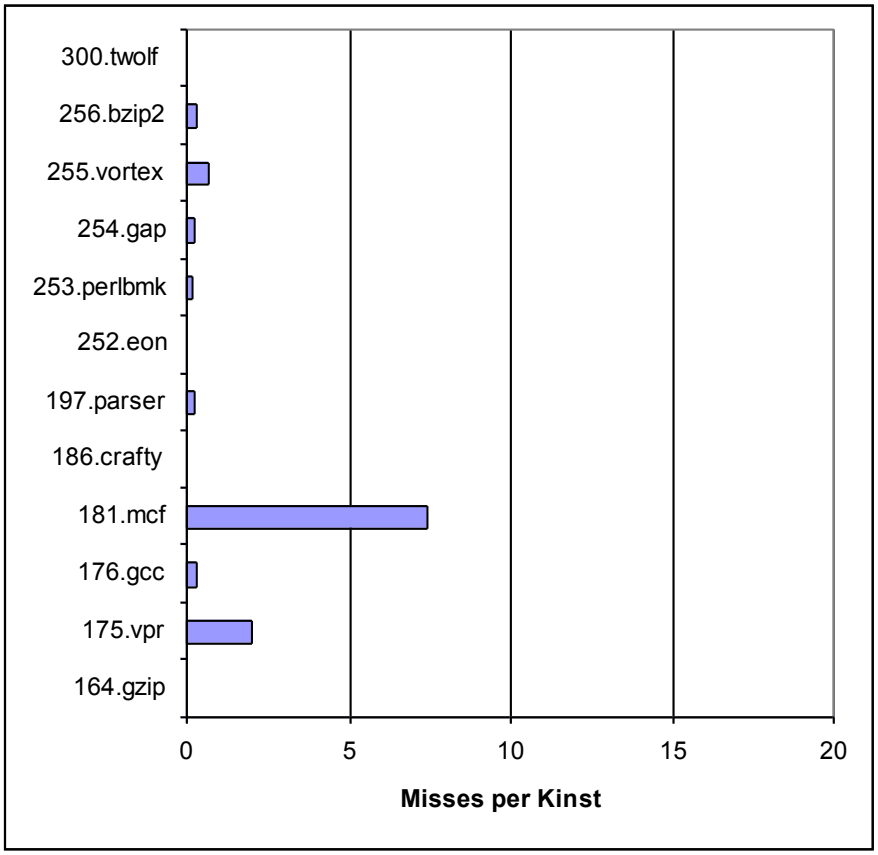

Branch mispredictions per 1000 Instructions

SPEC CPU2006

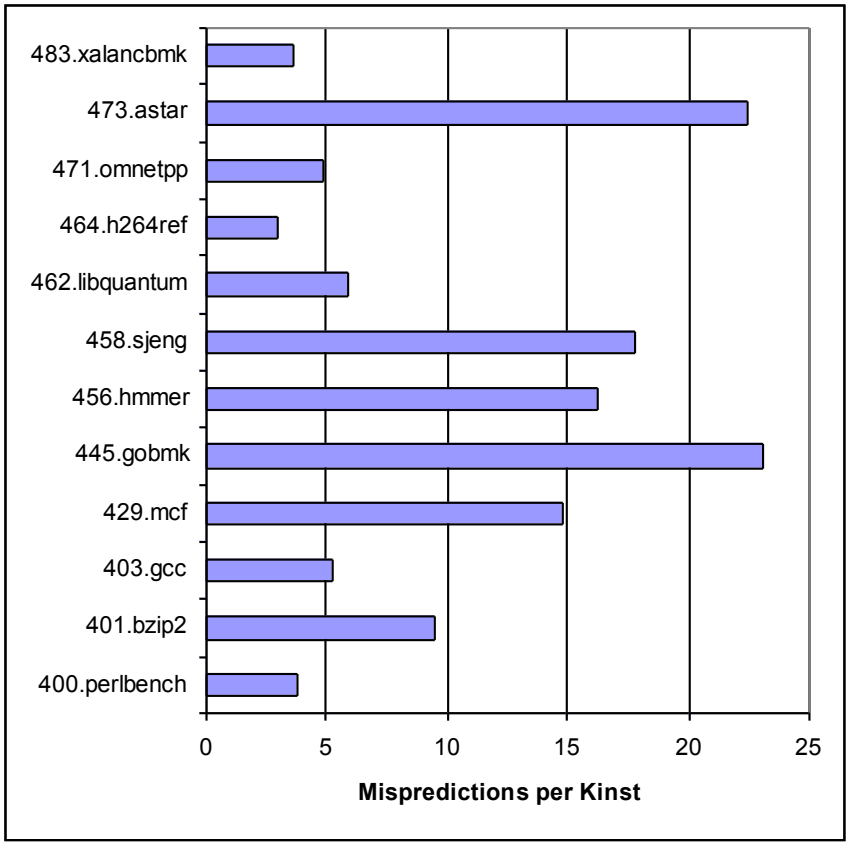

SPEC CPU2000

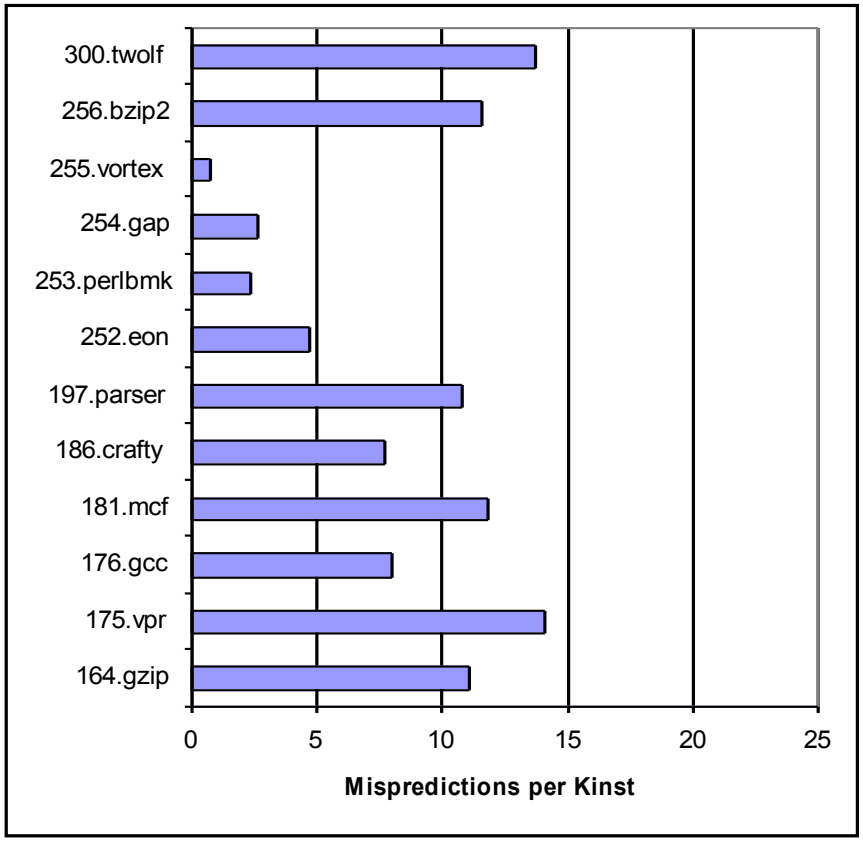

Performance Characteristics Correlation with CPI

Characteristics	Correlation Coefficient
Branch mispresdictions per KI and CPI	0.150
L1-D cache misses per KI and CPI	0.918
L2 misses per KI and CPI	0.964

- L2 misses have the strongest correlation with CPI
- L1-D caches misses also show significant correlation with CPI
- Branch mispredictions do not appear to directly impact CPI

Fusion

Macro-Fusion

- New feature for Core
- Decreases number of micro-ops
- One macro-fusion per cycle
- Fused in pre-decode phase
- Fuses branch and compare instructions

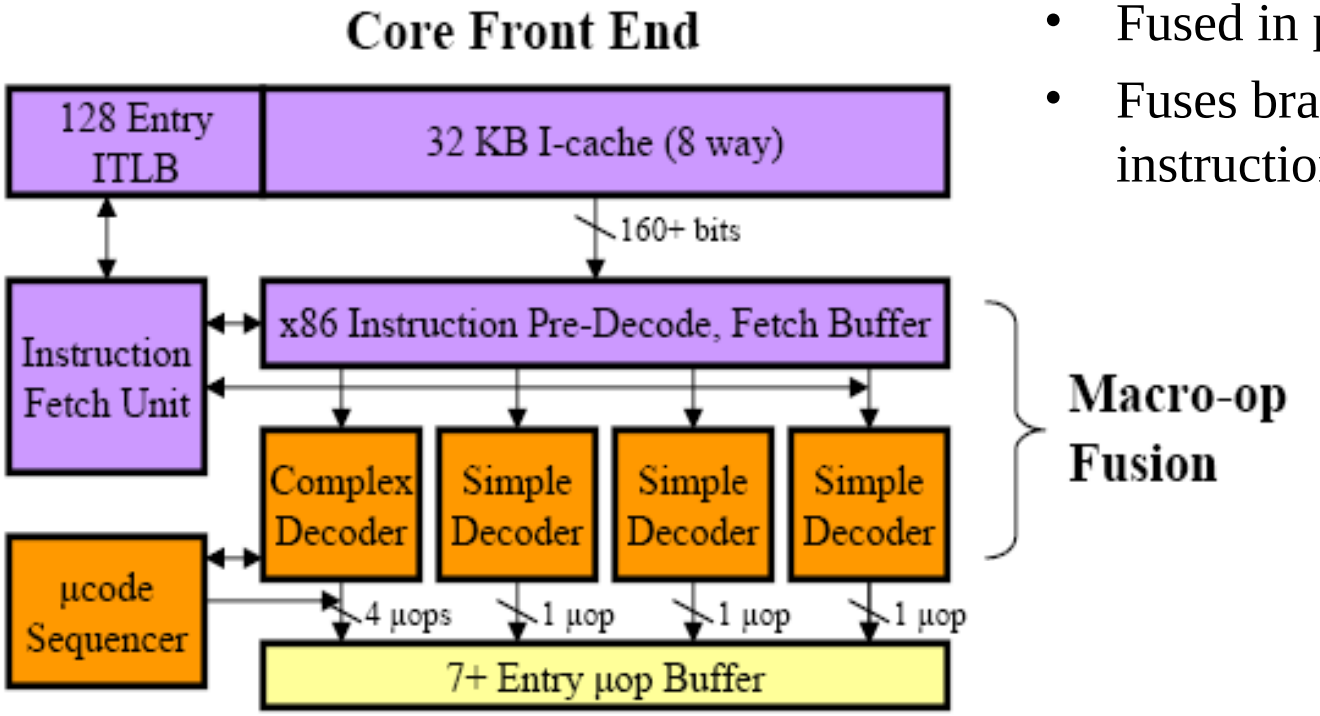

*Image taken from Real World Technologies "Intel's Next Generation Microarchitecture Unveiled"

Micro-op Fusion

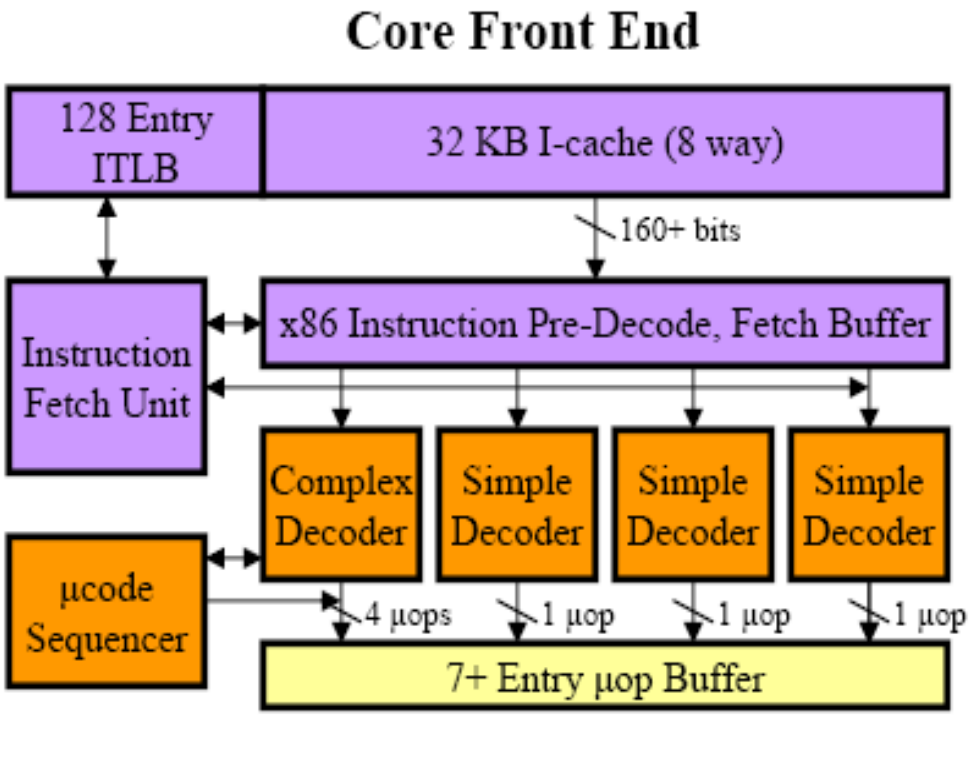

- Enhanced version of Pentium M feature
- Occurs in the decode phase
- Fused pair is issued/executed separately, but tracked by the reorder buffer as one micro-op
- Typically fuses store address micro-op and a data micro-op
- Increases space in reorder buffer

*Image taken from Real World Technologies "Intel's Next Generation Microarchitecture Unveiled"

Performance Improvement Calculation Details

$$\text{Increase in Performance} = \frac{\text{(Predecessor Runtime) *(Predecessor Frequency) - (Core Cycles)}}{\text{(Predecessor Runtime) *(Predecessor Frequency)}}$$

Core a[illegible]essor C[illegible]rison Chart

Architecture	Core	Yonah	NetBurst
Processor	Xeon 5160	Core Duo T2500	Pentium Extreme Edition 965
Macro-Fusion	Yes	No	No
Micro-op Fusion	Yes	Yes	No
Data Source	Performance Counter Data	SPEC Website Reported Results	SPEC Website Reported Results

Percentage of fused operations for SPEC CPU2006

SPEC CPU2006 Integer

Benchmark	Macro	Micro	Total
400.perlbench	12.18%	19.68%	31.86%
401.bzip2	11.84%	18.95%	30.79%
403.gcc	16.18%	18.39%	34.57%
429.mcf	13.93%	21.40%	35.33%
445.gobmk	11.33%	20.19%	31.52%
456.hmmer	0.13%	23.30%	23.43%
458.sjeng	14.33%	18.32%	32.65%
462.libquantum	1.59%	13.92%	15.15%
464.h264ref	1.51%	23.45%	24.96%
471.omnetpp	8.19%	23.87%	32.06%
473.astar	12.86%	14.40%	27.06%
483.xalancbmk	15.98%	21.12%	37.10%

SPEC CPU2006 Floating Point

Benchmark	Macro	Micro	Total
416.games	2.18%	23.58%	25.76%
433.milc	0.35%	17.82%	18.17%
434.zeusmp	0.09%	16.32%	16.41%
435.gromacs	0.71%	15.58%	16.29%
436.cactusADM	0.00%	24.14%	24.14%
437.leslie3d	0.70%	24.12%	24.82%
444.namd	0.36%	10.02%	10.38%
447.dealII	8.03%	22.98%	31.01%
450.soplex	5.10%	13.59%	18.69%
453.povray	4.40%	21.88%	26.28%
454.calculix	0.44%	19.67%	20.11%
459.GemsFDTD	0.38%	18.66%	19.04%
465.tonto	1.69%	24.35%	26.04%
470.ibm	0.22%	19.56%	19.78%

Percentage of increase in performance and percentage fused macro-ops for Woodcrest over Yonah

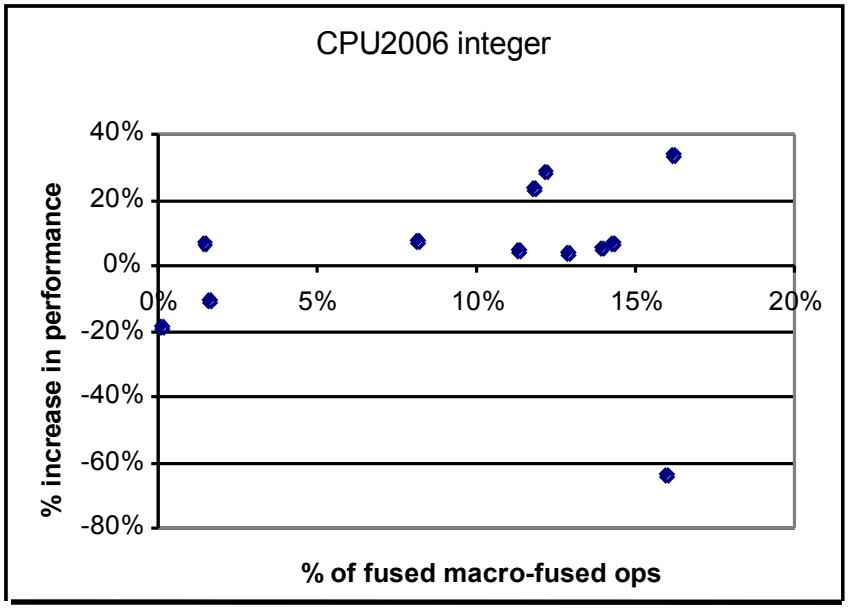

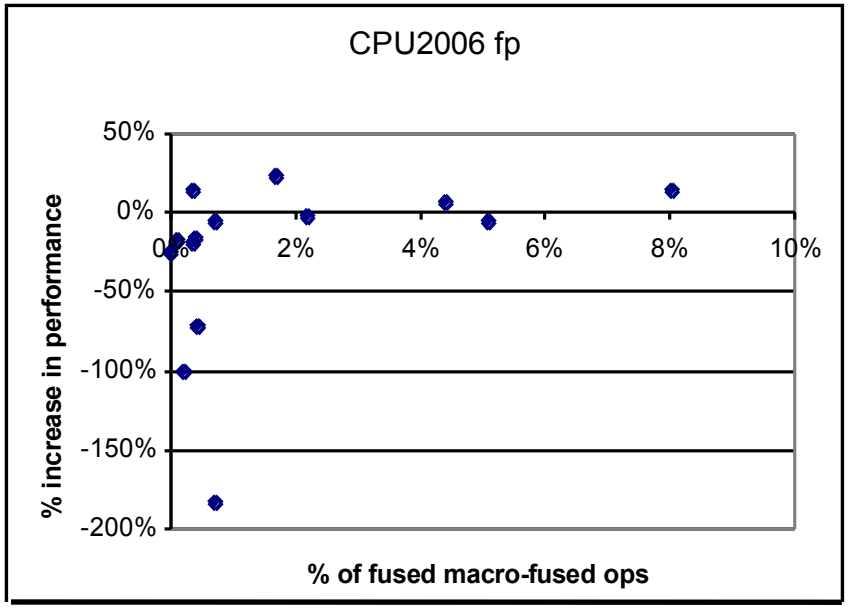

Percentage of increase in performance and percentage fused macro-ops for Woodcrest over Netburst

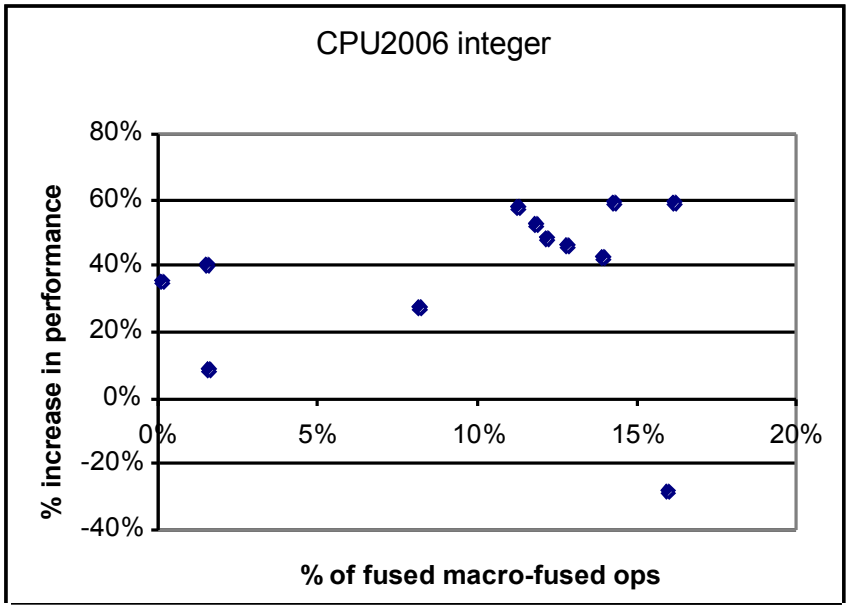

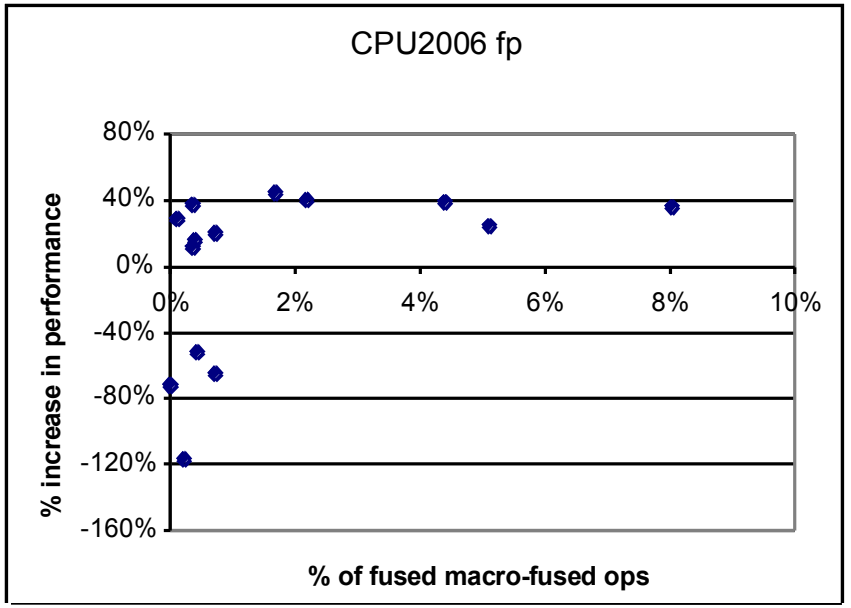

Percentage increase in performance compared with measured micro-op fusion for Netburst to Core

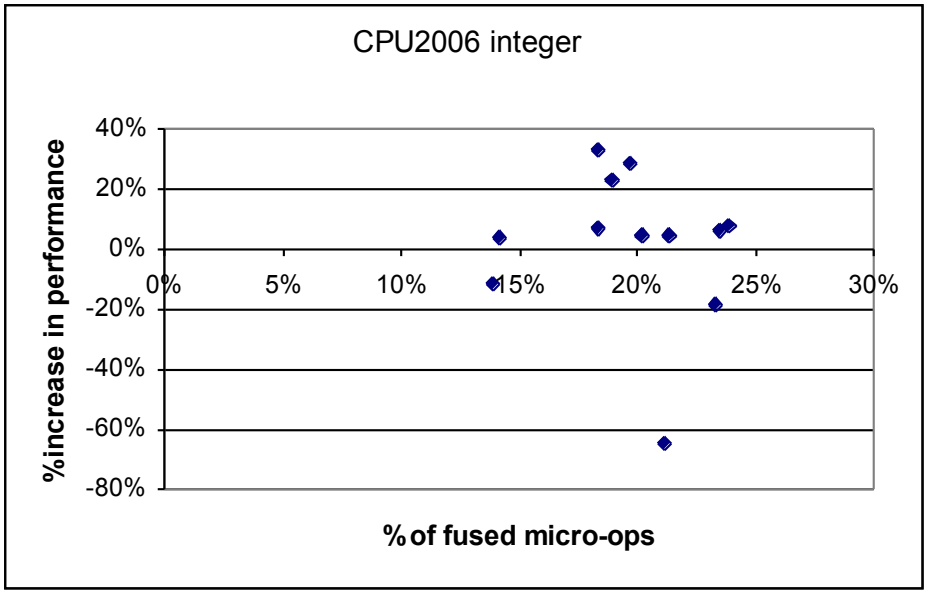

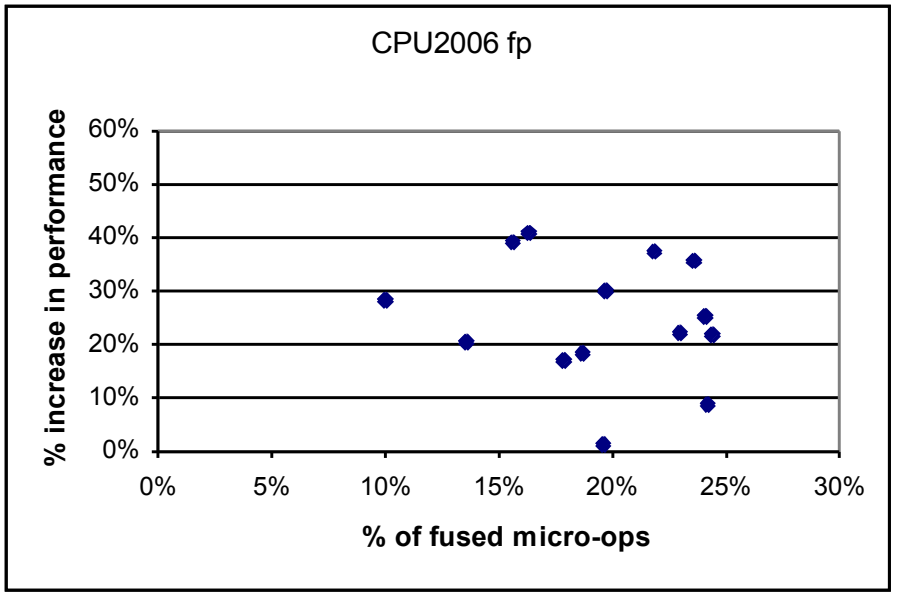

Conclusion

- SPEC CPU2006 stresses L2 Cache more
- The measured amount of Macro-fusion in floating point benchmarks is very low and does not correlate well with increase in performance
- The measured amount of Macro-fusion in integer benchmarks shows significant correlation with the increase in performance
- Micro-op fusion does not correlate well with increase in performance

Questions?

For more information please see
the Laboratory for Computer Architecture website
http://lca.ece.utexas.edu