SPEC* CPU2006 on Itanium®2 Systems

Gerolf Hoflehner

Intel® Itanium Compiler Lab

Agenda

- SPEC CPU2006 Overview
- Counter Data
- Performance
  - IPF relevant features
  - Comparison with GCC
- Conclusions

CPU2006: Background

- Released on August 24th 2006
- Replaces CPU2000
  - CPU2000 retired since February 24th
- Real-world applications
- Increased focus on C++
- 64bit/Mixed/32 bit results
- No PGO in Base

CINT2006 Characteristics

#Benchmarks	12(12)
Language	3 C++ (1)
Memory Footprint	5X
Lines of Code	2.3X
Hot Functions (90%)	19 (18)
Compile Time	1.6X
Run Time	9X
Reporting	64bit/Mixed/32bit

No PGO in Base

CINT2006 Benchmarks

Benchmark	Language	LOC	#HOT Routines	Description
400.perlbench	C	155432	33	Perl v5.8.7
401.bzip2	C	8293	8	Bzip v1.0.3
403.gcc	C	518781	61	GCC v3.2
429.mcf	C	2685	2	Combinatorial Optimizations
445.gobmk	C	197215	39	Go Game
456.hmmer	C	35992	1	Sensitive Database Searching
458.sjeng	C	13487	14	Chess Game
462.libquantum	C	4805	3	Quantum Computer Simulation running Shor's alg
464.h264ref	C	51578	14	Video compression Standard
471.omnetpp	C++	48159	34	Discrete Event Simulation of an Ethernet Network
473.astar	C++	5842	3	2D path finding library used in game AI
483.xalancbmk	C/C++	326504	21	XML Converter (Xalan-C++)

CFP2006 Characteristics

#Benchmarks	17(14)
Language	4 C++ (0)
Memory Footprint	5X
Lines of Code	8.2X
Hot Functions (90%)	11(4)
Compile Time	7.8X
Run Time	11.6X
Reporting	64bit/Mixed/32bit

No PGO in Base

CFP2006 Benchmarks

Benchmark	Language	LOC	#HOT Routines	Description
410.bwaves	Fortran	918	2	Computes 3D transonic transient laminar viscous flow
416.gamess	Fortran	932818	10	General Atomic and Molecular Electronic Structure System
433.milc	C	15042	5	Particle Physics - Gauge Fields
434.zeusmp	Fortran	37326	8	Astrophysics - Simulation
435.gromacs	C	87736	8	Newton Equation Solver - many particle motion
436.cactusADM	Fortran/C	104047	1	Einstein Equation Solver
437.leslie3d	Fortran	3807	6	Large Eddy Simulations
444.namd	C++	5315	9	Large Biomolecular Systems Simulation
447.dealII	C++	199654	17	Adaptive Finite Elements Solver
450.soplex	C++	41417	14	Simplex Algorithm
453.povray	C++	157825	18	Ray Tracing
454.calculix	Fortran,C	49927	18	3D Finite Element Solver
459.GemsFDTD	Fortran90	11580	4	3D Maxwell Equations Solver
465.tonto	Fortran90	143152	31	Quantum Chemistry Package
470.lbm	C	1176	1	3D incompressible fluid simulations
481.wrf 2	Fortran90,C	217896	33	Weahter Research and Forecasting Modeling System
482.sphinx3	C	207732	5	Speech recognition system

CPU2006 Dual-Core Itanium® 2 Published Results

- No. 1 on 64bit CINT2006
- No. 2 on CINT2006
- No. 1 on CFP2006
- No. 1 scalability
  - No. 1 for 32, 64, 128, 256 cores rate results (INT, FP)

Agenda

- SPEC CPU2006 Overview
- Counter Data
- Performance
  - IPF relevant features
  - Comparison with GCC
- Conclusions

Data Collection Setup

- **Platform:** 400MHz bus, Dual Core Itanium® 2 1.6GHz, 12MB L3/Core, 4-Sockets, Dual Core Enabled, 4GB RAM, RedHat* EL4 U3
- **Compilers:**
  - ICC 9.1 Fortran/C++, Base and O2
- **Tools:** Caliper*
- **CPU2006 Kit1.0**

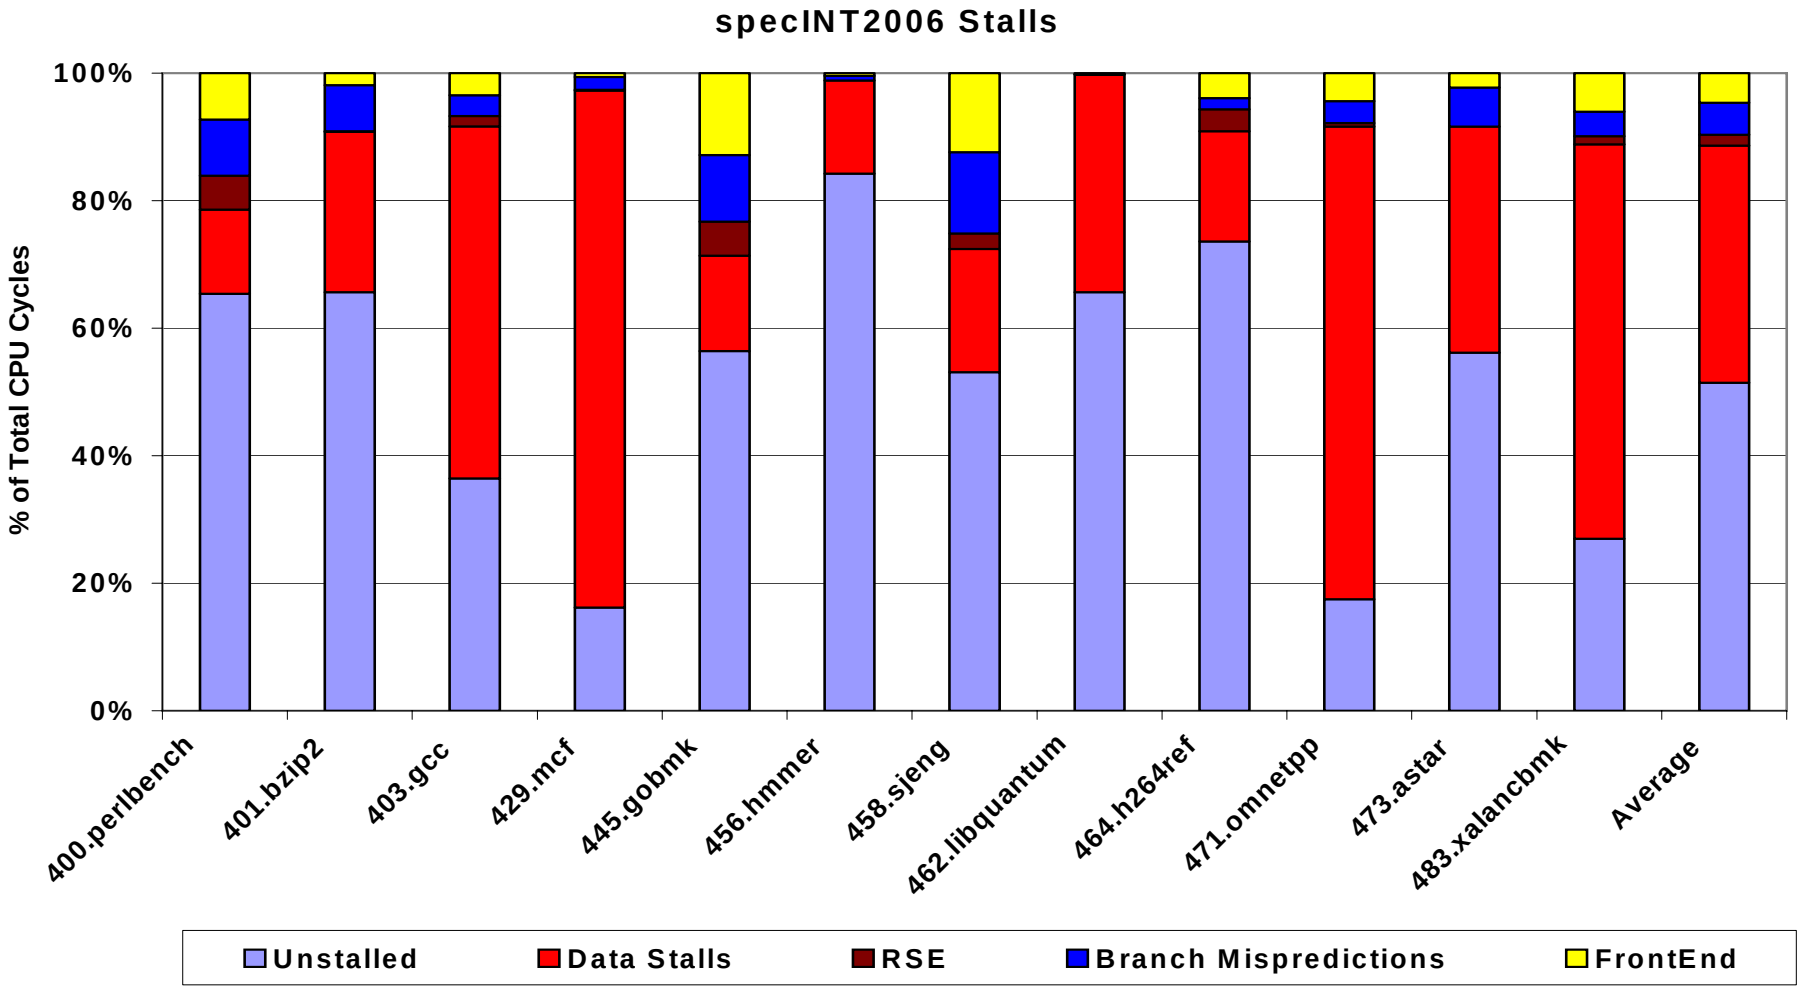

- 8/12 benchmarks have >50% unstalled execution time
- Data stalls are the performance barrier
- 471.omnetpp, 473.astar and 483.xalancbmk (80%) benefit from better memory management libraries
- 471.omnetpp did not move for more than 1 year

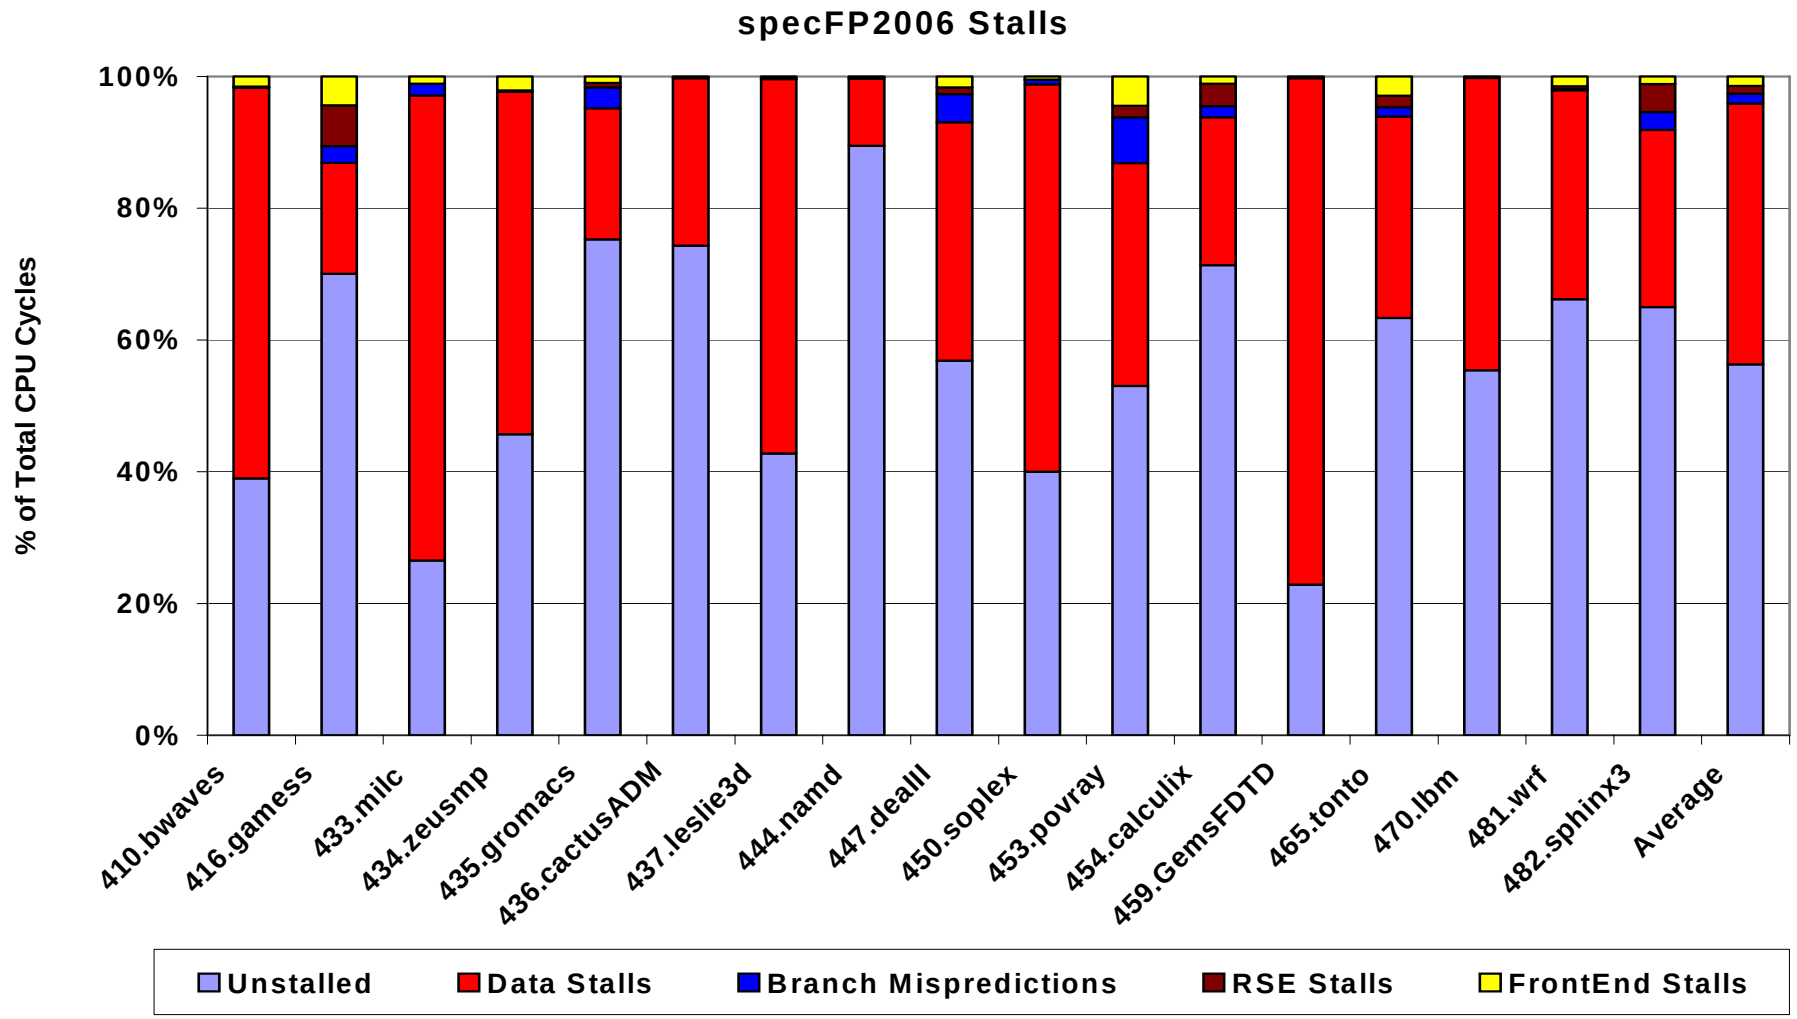

- 10/17 >50% unstalled execution time.
- Data stalls are the performance barrier
- No impact from memory management libraries

Agenda

- SPEC CPU2006 Overview
- Counter Data
- Performance
  - IPF relevant features
  - Comparison with GCC
- Conclusions

CINT2006: Optimization Gains

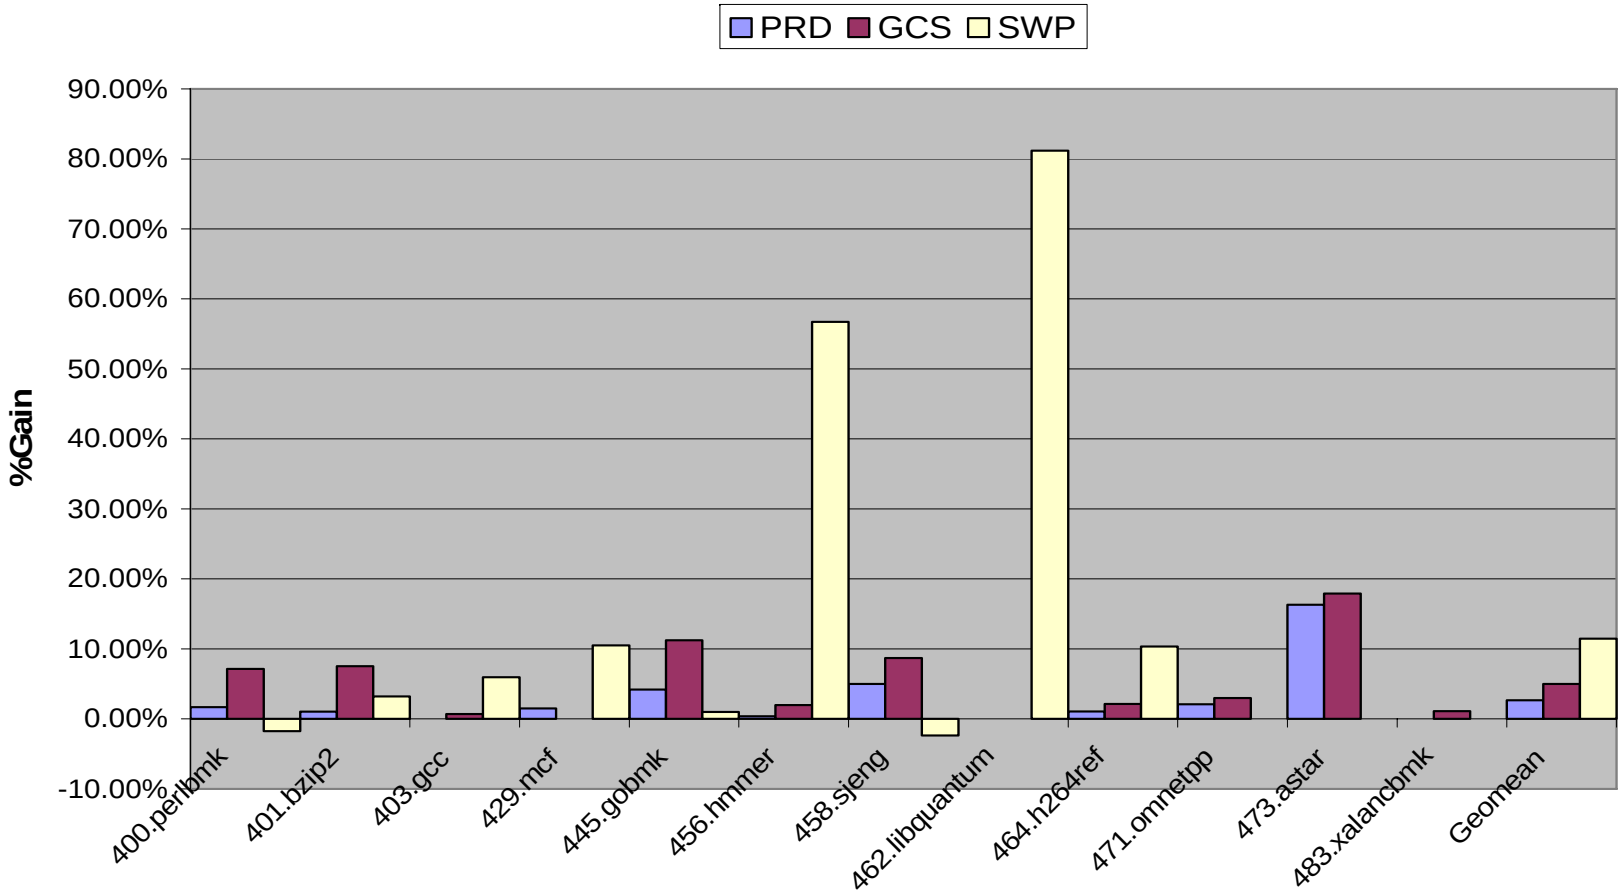

-10% gains from modulo scheduling (unlike CPU2000)

-Gains from global code scheduling and predication similar to CPU2000

-Predication gain is for non-pipelined regions only

CFP2006: Optimization Gains

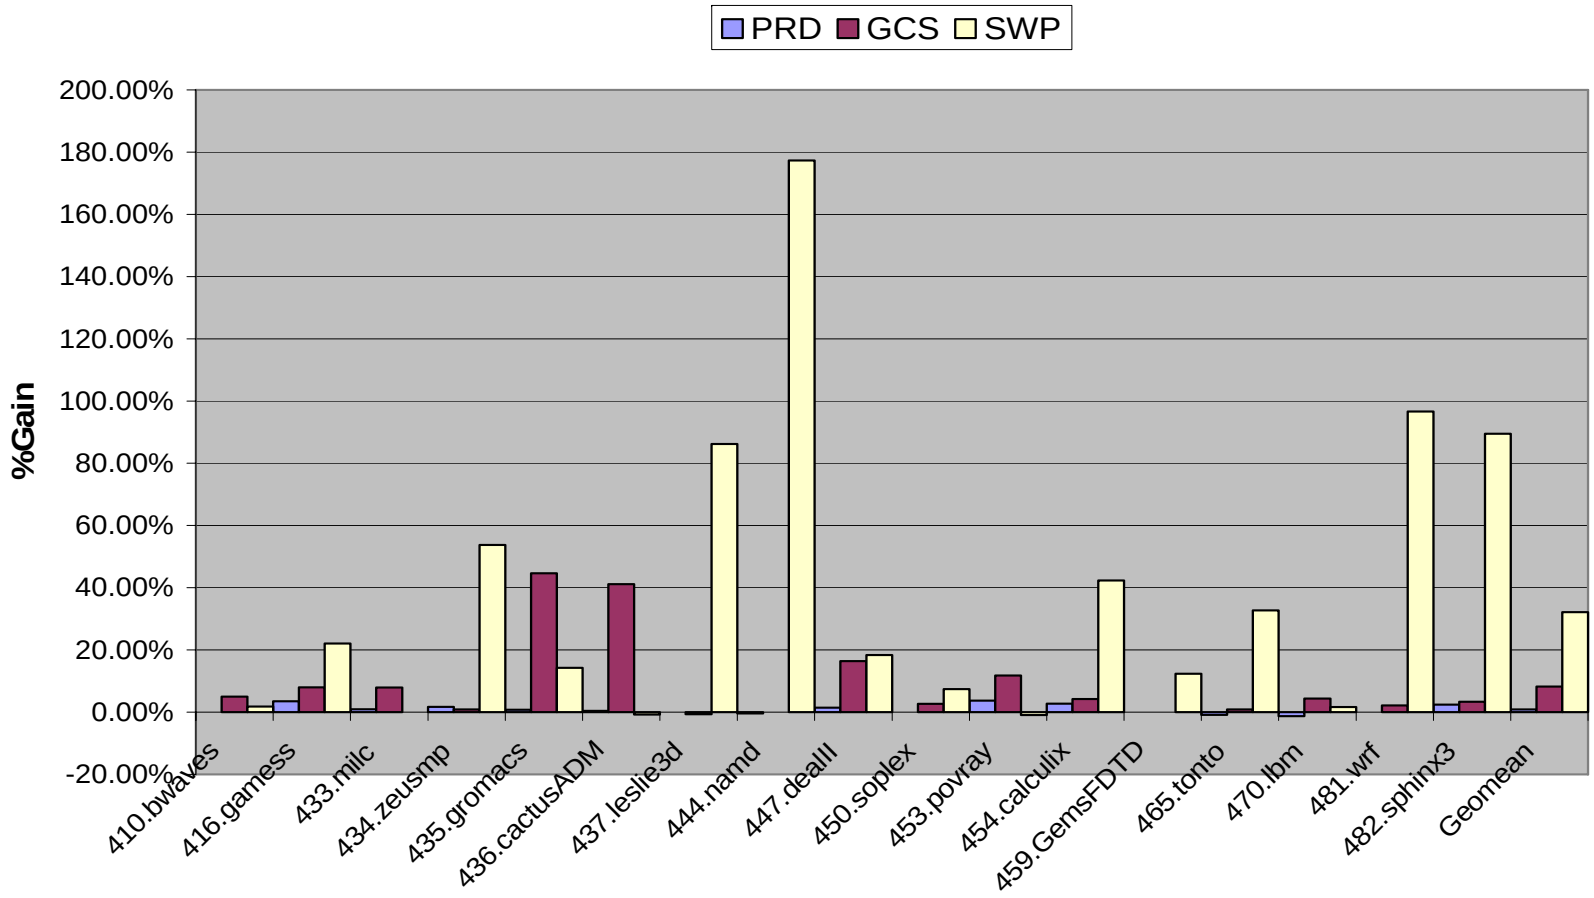

-Significant gains from modulo scheduling

-Global code scheduling is more important for 6/17 benchmarks

-Predication gain is for non-pipelined regions only

Agenda

- SPEC CPU2006 Overview
- Counter Data
- Performance
  - IPF relevant features
  - Comparison with GCC
- Conclusions

Compiler options

- ICC10.0 Beta Base:
  - fast –IPF_fp_relaxed –ansi-alias
- GCC 4.2 Base:
  - OPTIMIZE=-O3 -funroll-all-loops -ffast-math –static

CINT2006 Base: ICC10.0 Beta vs GCC4.2

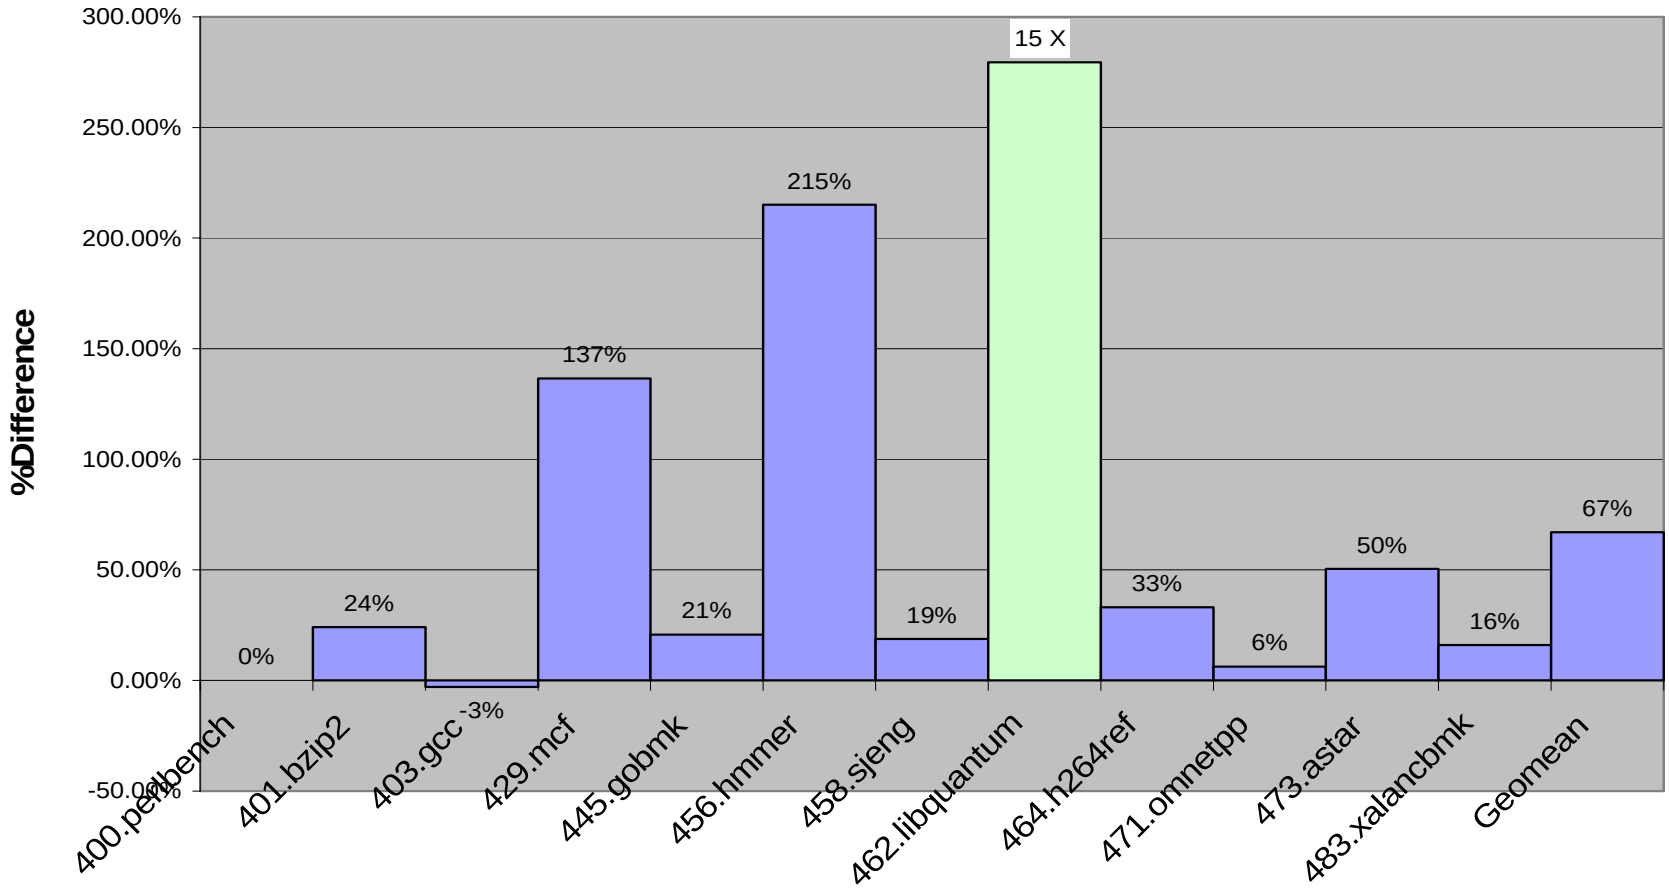

-Base Delta: 429.mcf (>2X), 456.hmmer(>3X), 462.libquantum(>15X)

-Exceptions: 403.gcc and 471.omnetpp

** Disclaimer: Numbers based on internal measurements

CINT2006 O2: ICC10.0 Beta vs GCC4.2

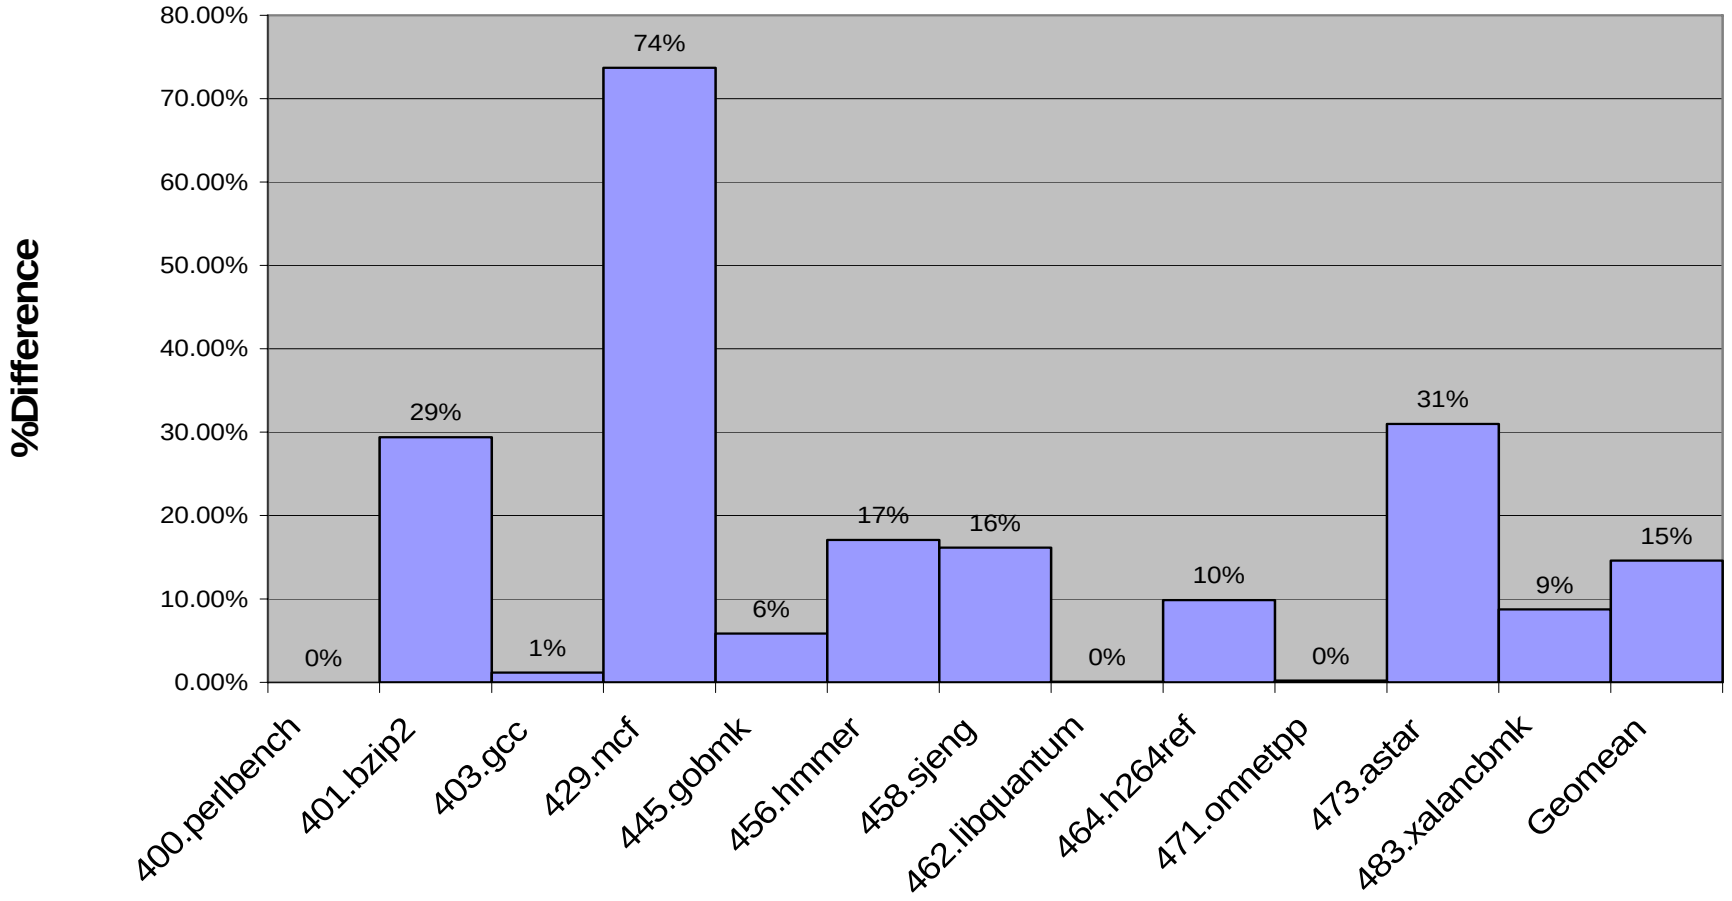

- Biggest difference for429.mcf (>70%). Most benchmarks >10% faster, except 403.gcc, 462.libquantum and 471.omnetpp.

** Disclaimer: Numbers based on internal measurements

CFP2006 Base: 10.0Beta vs GCC4.2

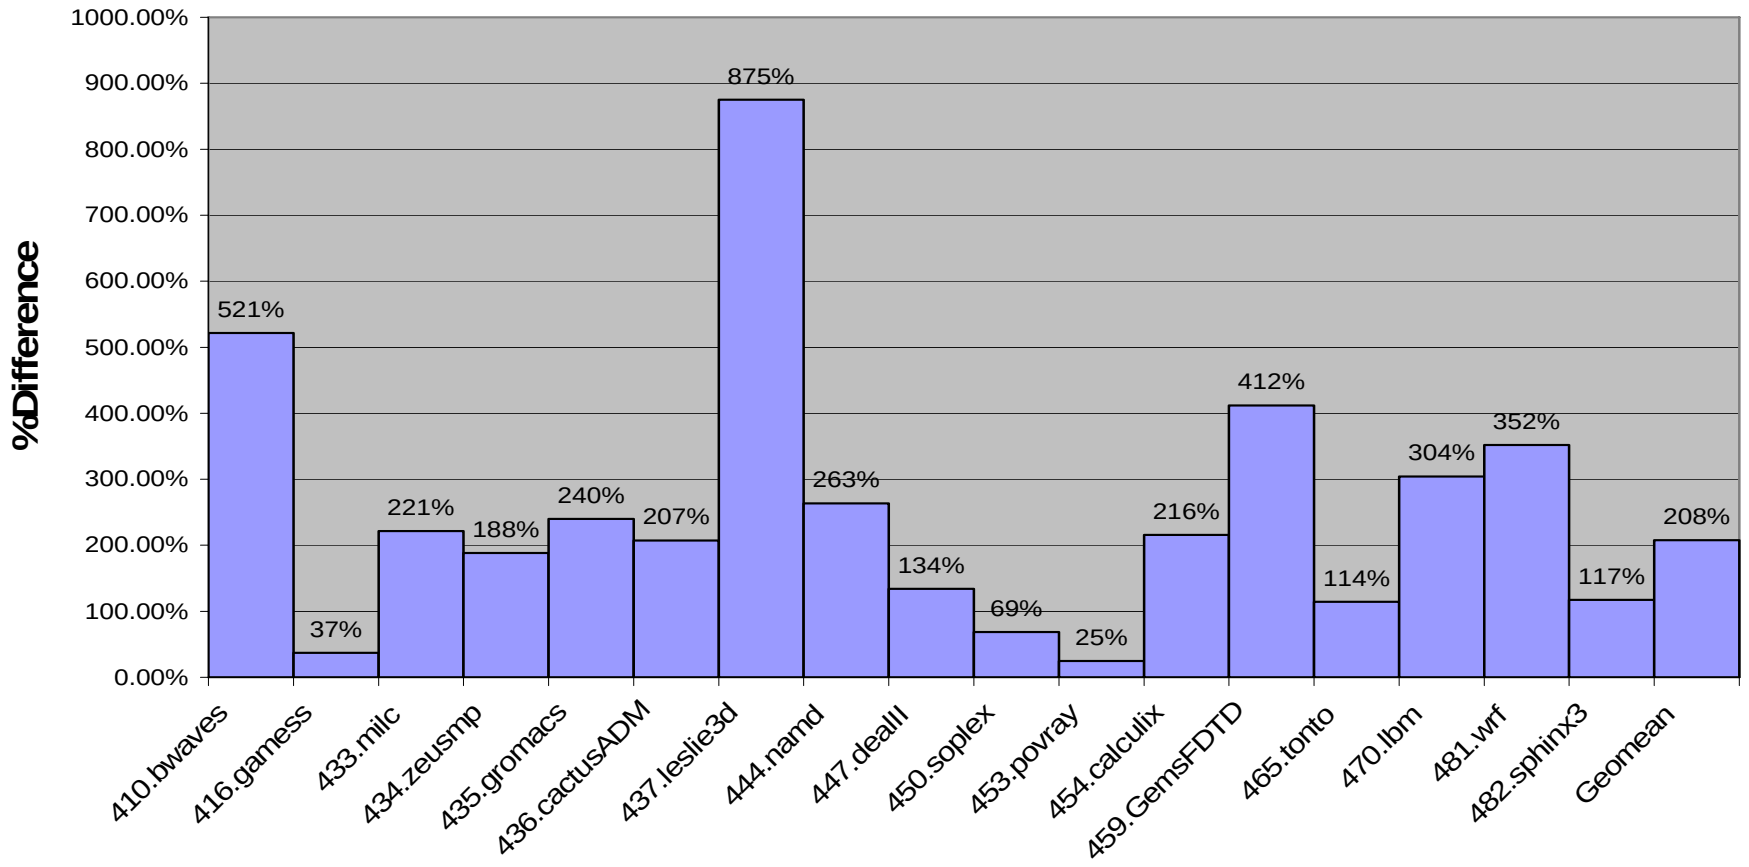

- FP is a bigger challenge for GCC
- Set of optimizations have to work together
- Flush to zero mode could improve GCC by > 20% on average

** Disclaimer: Numbers based on internal measurements

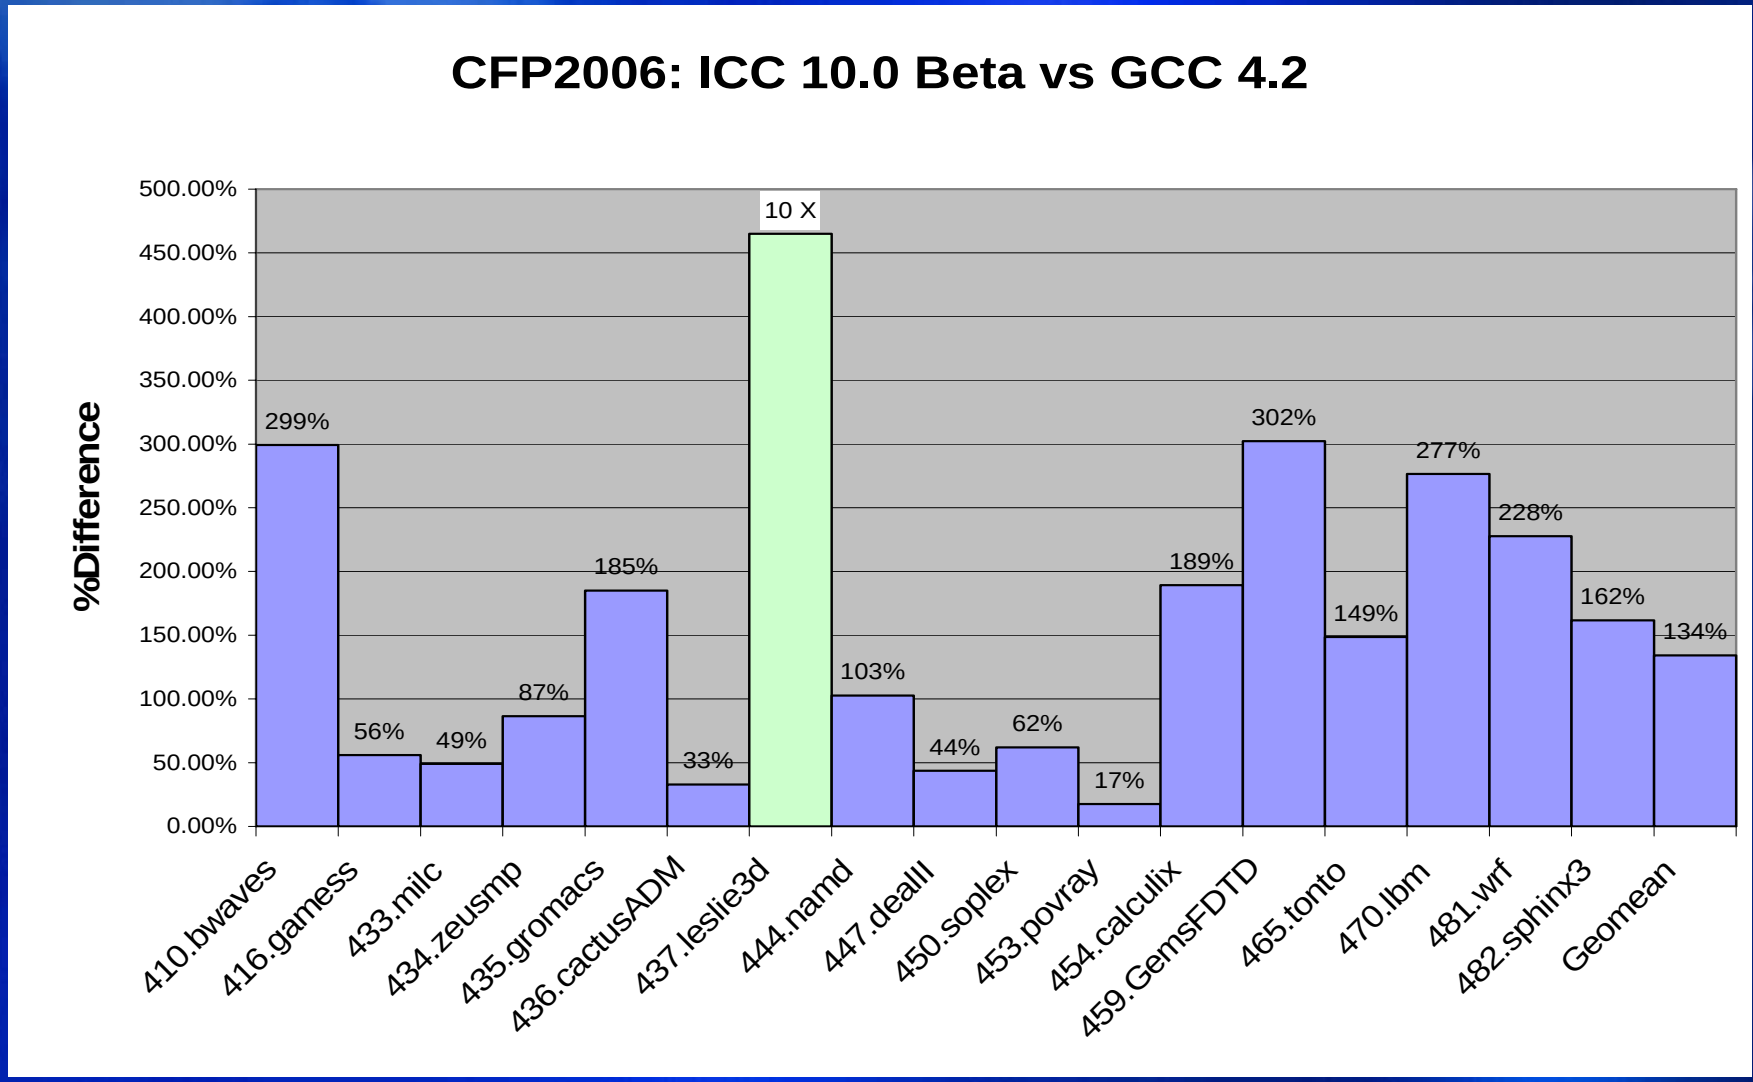

- 437.Leslie: 10X faster with ICC

** Disclaimer: Numbers based on internal measurements

Suggestions for GCC

- Memory Disambiguation and Points-To Analysis
- Data Prefetching
- Software Pipelining
- Global Code Scheduling

Watch Compile Time

Conclusions

- CPU2006 is challenging for compiler technology
- Compile Time can be a limiter for performance work
- Focus on data stalls
- Dual-Core Itanium® 2 gives leadership results on both, INT and FP
- GCC has room to improve

Acknowledgements

- Intel® Itanium® Compiler Team
- Special thanks to Darshan Desai, Howard Chen, Alexander Isaev, Alexander Ostanevich, Arun Kejariwal (Univ of Irvine, CA)

THANK YOU!