AMD Opteron™ 6200 Series processors
AMD Opteron™ 4200 Series processors
Compiler Options Quick Reference Guide

Open64

Latest release: 4.5.1, December 2011
http://developer.amd.com/open64

Architecture	
Generate instructions specific to Interlagos	-march=bdver1
Generate instructions for the local machine	-march=auto
Optimization Levels	
Disable all optimizations	-O0
Local optimizations	-O1
Global optimizations (default)	-O2
Additional aggressive optimizations	-O3
Maximize performance	-Ofast
Additional Optimizations	
Autoparallelization	-apo
Feedback directed optimization	-fb-create -fb-opt
Huge pages	-HP
Interprocedural Analysis and Optimizations	-ipa
Link to ACML 5.0	-L/opt/acml5.0.0/ open64_64_fma4/ lib -lacml
Loop nest optimizations, vectorization, prefetch, fission, fusion	-LNO:fission=n -LNO:fusion=n
Multicore scalability	-mso
OpenMP	-mp
Prefetch (disabled by default)	-LNO:prefetch -LNO:prefetch_ahead
Floating point accuracy	
Floating point accuracy	-fp-accuracy

gcc

Latest release: 4.7, March 2012
http://gcc.gnu.org

Architecture	
Generate instructions specific to Interlagos	-march=bdver1
Optimization Levels	
Disable all optimizations (default)	-O0
Local optimizations	-O1
Global optimizations	-O2
Additional aggressive optimizations	-O3
Maximize performance	-Ofast
Additional Optimizations	
AVX tuning	-mprefer-avx128
Enable unrolling	-funroll-all-loops
Generate prefetch instructions for loops	-fprefetch-loop-arrays --param prefetch-latency=300 (300-700)
Inline string operations	-minline-all-stringops
Link to ACML 5.0	-L/opt/acml5.0.0/ gfortran64_fma4/ lib -lacml
OpenMP	-fopenmp
Profile guided optimization	-fprofile-generate -fprofile-use
Turn off partial redundancy elimination	-fno-tree-pre
Vectorization	-ftree-vectorize
Floating point accuracy	
Enable generation of code that follows IEEE arithmetic	-mieee-fp
Enable faster, less precise math operations	-ffast-math

For more information, visit http://developer.amd.com

Advanced Micro Devices
One AMD Place
P.O. Box 3453
Sunnyvale, CA 94088-3453

AMD Opteron™ 6200 Series processors
AMD Opteron™ 4200 Series processors
Compiler Options Quick Reference Guide

ICC

Latest release: 12.1, January 2012
http://software.intel.com

Architecture	
Generate instructions that run on Interlagos	-msse3 (avoid –mAVX)
Optimization Levels	
Disable all optimizations	-O0
Speed optimization without code growth	-O1
Enable optimization including vectorization	-O2
Aggressive optimization	-O3
Maximize performance	-fast
Additional Optimizations	
Aggressive unrolling	-unroll-aggressive
Disable improved precision floating divides	-no-prec-div
Enable vectorization	-simd
Interprocedural Optimization	-ipo
Link to ACML to 5.0	–L/opt/acml5.0.0/ifort64_fma4/lib –lacml
OpenMP	-openmp
Prefetch optimization	-opt-prefetch
Profile generated optimization	-prof-gen -prof-use
Use optimized header definitions	–use-intel-optimized-headers
Floating point accuracy	
Floating point accuracy	-fp-model

PGI

Latest release: 12.3, March 2012
http://www.pgroup.com

Architecture	
Generate instructions that run on Interlagos	-tp bulldozer
Optimization Levels	
Disable all optimizations	-O0
Local optimization	-O1
Global optimization	-O2
Aggressive global optimization	-O3
Hoist guarded invariant floating point expressions	-O4
Maximize performance	-fast
Additional Optimizations	
Huge pages	-Msmartalloc=huge
Autoparallelize loops	-Mconcur
Enable vectorization	-Mvect
Interprocedural Optimization	-Mipa=fast,inline
Link to ACML to 5.0	-L/opt/acml5.0.0/ pgi64_fma4/lib -lacml
OpenMP	-mp
Prefetch instructions	-Mvect=prefetch
Profile guided optimization	-Mpfi -Mpfo
Unroll loops	-Munroll
Floating point accuracy	
Generate relaxed precision code	-Mfprelaxed
Perform floating point operations in conformance with IEEE standard	-Kieee

For more information, visit http://developer.amd.com

Advanced Micro Devices
One AMD Place
P.O. Box 3453
Sunnyvale, CA 94088-3453