CORAL Procurement Benchmarks

CORAL Vendor Meeting

May 31, 2013

Argonne NATIONAL LABORATORY

Lawrence Livermore National Laboratory

OAK RIDGE National Laboratory

U.S. DEPARTMENT OF ENERGY

LLNL-PRE-637694
This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC

CORAL

Contents

- Benchmark Categories
- Marquee and Elective Benchmarks
- Benchmark Procedures
- Code Modifications

The performance references for the CORAL procurement are the DOE production system Sequoia/Mira and Titan

- Target 4-8X performance improvement for full system science runs
- Target 6-12X performance improvement for large ensemble/throughput simulations
- CORAL systems will require disruptive application changes
- CORAL seeks to minimize the changes that are not part of standard programming models
- CORAL benchmarks have been selected to represent a broad range of applications

Argonne NATIONAL LABORATORY Lawrence Livermore National Laboratory OAK RIDGE National Laboratory

CORAL Benchmark Categories range from complex applications to single node tests

- Scalable Science Benchmarks
  - Expected to run at full scale of the CORAL systems
  - Target 4-8X performance improvement with combination of strong and weak scaling
- Throughput Benchmarks
  - Represent large ensemble runs
  - Target 6-12X performance improvement with combination of strong and weak scaling
- Skeleton Benchmarks
  - Used to investigate various platform characteristics including network performance, threading overheads, I/O patterns, memory access patterns, system software requirements, and new programming models
- Micro Benchmarks
  - Small code fragments that represent expensive compute portions of some of the scalable science and throughput applications
  - Useful for testing programming methods and performance at the node level
  - Ideal for early evaluations and explorations on hardware emulators and simulators

CORAL Benchmarks code and procedures are available at https://asc.llnl.gov/CORAL/benchmarks

CORAL Marquee and Elective Benchmarks

	Scalable Science	Throughput	Skeleton	Micro-kernels
Marquee TR-1	LSMS Qbox NEKbone HACC	CAM-SE UMT2013 AMG2013 MCB	CLOMP IOR CORAL MPI Memory CORAL loops	
TR-2		QMCPACK NAMD LULESH SNAP miniFE	Pynamic HACC I/O FTQ XSBench	
TR-3			miniMADNESS	NEKbonemk HACCmk UMTmk AMGmk MILCmk GFMCmk

Benchmarks used in overall performance metric
Optional benchmarks used in overall performance metrics

CORAL Marquee Benchmark Coverage Matrix

	Scalable Science				Throughput				Skeleton				
	LSMS	Qbox	Nekbone	HACC	CAM-SE	UMT 2013	AMG 2013	MCB	CLOMP	IOR	MPI bm	STREAMS /STRIDE	Loops suite
High FP intensity	X	X	X	X		X							X
SIMD / vectorization			X	X									X
High Integer / branching				X				X					
High memory bandwidth		X				X						X	
Regular memory accesses			X	X	X							X	X
Irregular memory accesses				X			X	X				X	X
Large memory footprint						X							
Non-local P2P comm	X		X	X	X						X		
Small messages (interconnect latency)			X	X	X			X			X		
Large messages (interconnect BW)					X	X					X		
MPI Collectives		X	X								X		
Bi-Section bw		X									X		
OpenMP threading				X			X		X			X	X
System software / compilers									X		X		X
I/O										X			

CORAL system performance targets will be projected for both scalable science and throughput

Marquee Scalable Science Benchmarks
(each run/projected at full machine scale)

Qbox
LSMS
NEKbone
HACC

Marquee Throughput Benchmarks
(Allow 6-9 copies of each benchmark to fill machine)

UMT2013	AMG2013	MCB	CAM-SE
UMT2013	AMG2013	MCB	CAM-SE
UMT2013	AMG2013	MCB	CAM-SE
UMT2013	AMG2013	MCB	CAM-SE
UMT2013	AMG2013	MCB	CAM-SE
UMT2013	AMG2013	MCB	CAM-SE

$$S_S = \frac{\sum_{i=1}^{N_S}(FOM_i * W_i)}{N_S}$$

$$S_T = \frac{\sum_{i=1}^{N_{TP}}(FOM_i * W_i)}{N_{TP}}$$

CORAL will provide

- Benchmark weights determined by normalizing performance of each benchmark to 1.0
  - Throughput apps: Based on job using 1/24 of Sequoia or Titan
  - Science apps: Based on full-system Sequoia or Titan run
- All Marquee benchmarks are of equal importance
- FOM's and weights subject to change between now and final RFP release

Offeror asked to:

- Estimate Speed Up = w_{App} x FOM_{App} for each Marquee Science and Throughput benchmark
- Provide results for Marquee Skeleton benchmarks
  - CLOMP, IOR, MPI, memory, Loop suite

CORAL Micro Benchmarks are available to aid vendors in performance projections

Micro Benchmarks	TR-x	Aspect of Platform Exercised	LOC	Parallelism	Language			Application Domains						
				OpenMP/p threads	F	C	C++	FFT	Dense LA	Sparse LA	Particles	Monte Carlo	Struct. Grid	Unstruct. Grid
NEKbonemk	3	SIMD vectorization	2000		X				X				X	
HACCmk	3	Compute intensity	250	X			X							
UMTmk	3													
AMGmk	3													
MILCmk	3	Compute intensity, Memory performance	5,000	X		X			X	X			X	
GFMCmk	3	Random memory access	150	X	X					X				

Allowed Code Modification to CORAL Benchmarks

- Benchmarks may be modified as necessary to get them to compile and run
- A full set of benchmark runs must be reported with this “as is” source code
- Offeror may also report optimized results
  - Most valuable optimization are from compiler flags and hints
  - Next in value is pragma-style guidance
  - Changes to MPI and OpenMP implementations are allowed
  - Wholesale algorithmic changes that are strongly architecture specific have less value
  - All benchmark code modification will be documented and provided to CORAL
- CORAL and Offeror will continue to improve the efficiency and scalability of all benchmarks between award of the contracts and delivery of the systems
  - Emphasis on higher level optimizations as well as compiler optimization technology improvements while maintaining readable and maintainable code

CORAL Benchmark Website

- CORAL Benchmark website contains additional information not found in the technical requirements document
  - Benchmark descriptions
  - Directions for compiling
  - Procedures for running CORAL benchmarks
  - Input parameters
  - Benchmark weights and figures of merit (FOM)
  - Benchmark scaling data on CORAL reference systems
  - CORAL Benchmark Spreadsheet for reporting results
- Any updates or answers to questions will be posted on the website

https://asc.llnl.gov/CORAL/benchmarks

The CORAL Benchmark Suite

				Parallelism		Language				Application Domains						
CORAL Benchmarks	**TR-x**	Reference Platform	LOC	MPI	OpenMP /pthreads	F	Py	C	C++	FFT	Dense LA	Sparse LA	Part-icles	Monte Carlo	Struct. Grid	Unstruct. Grid
Scalable Science Benchmarks																
LSMS	1	Titan/GPU	200,000	X	X	X			X		X			X		
QBOX	1	BG/Q	47,000	X	X				X	X			X			
NEKbone	1	BG/Q	48,000	X		X		X			X				X	
HACC	1	BG/Q	34,803	X	X				X	X			X			
Throughput Benchmarks																
CAM-SE	1	Titan/CPU only	150,000	X	X	X		X		X	X	X			X	
UMT2013	1	BG/Q	51,000	X	X	X	X	X	X							X
AMG2013	1	BG/Q	75,000	X	X			X								
MCB	1	BG/Q	13,000	X	X			X						X		
QMCPACK	2	Titan GPU	200,000	X	X			X	X	X	X		X	X	X	
NAMD	2	BG/Q	180,190	X	X				X	X			X			
LULESH	2	BG/Q	5,000	X	X			X								X
SNAP	2	BG/Q	3,000	X	X	X									X	
miniFE	2	BG/Q	50,000	X	X				X			X				X
Skeleton Benchmarks																
CLOMP	1				X			X								
IOR	1															
MPI Benchmark suite	1			X				X								
Memory subsystem (STREAMS or STRIDE)	1				X			X								
Loop suite	1		5,000						X							
Pynamic	2			X			X		X							
HACC IO	2															
FTQ	2				X			X								
XSBench (mini OpenMC)	2		1,000		X			X						X		
MiniMADNESS	3		10,000	X	X				X		X					X

Questions?