

BG/Q Performance Project: porting status of

- Performance Tools
- Debuggers
- Programming Models
- Libraries

ALCF Performance Engineering

BG/Q Performance Tools

- Early efforts were initiated to bring widely used performance tools to the BG/Q
- A variety of tools providers are currently working with IBM and Argonne to port and test tools on the Q
- **BG/Q provides a hardware & software environment that supports many standard performance tools:**
 - Software:
 - Environment similar to 64 bit PowerPC Linux
 - provides standard GNU binutils
 - New performance counter API bgpm
 - Performance Counter Hardware:
 - BG/Q provides 424 64-bit counters in a central node counting unit
 - Counter for all cores, prefetchers, L2 cache, memory, network, message unit, PCIe, DevBus, and CNK events
 - Provides support for hardware threads and counters can be controlled at the core level
 - Countable events include: instruction counts, flop counts, cache events, and many more
- **Argonne working with IBM on providing POMP features for tools developers. Also collaborating with LLNL on OpenMP performance interface standard in OpenMP ARB**



BG/Q Tools status on VEAS

Tool Name	Source	Provides	VEAS status
gprof	GNU/IBM	Timing (sample)	Version available on Veas
TAU	Univ. Oregon	Timing (inst), MPI	Version available on Veas
HPCToolkit	Rice University	Timing (sample), HPC (sample)	Version available on Veas
IBM HPCT	IBM	MPI, HPC	IBM commercial product. Early version available on Veas.
mpiP	LLNL	MPI	Version ported and available on Veas.
PAPI	UTK	HPC API	PAPI 4.4.2 supports BG/Q. Available on Veas
Darshan	ANL	IO	Port done to Veas. Not yet turned on for testing.
Open Speedshop	Krell Institute	Timing (sample), HCP, MPI, IO	Developers on Veas.
MPE/Jumpshot	Argonne	MPI	Developers on Veas.
Scalasca	Juelich	Timing (inst), MPI	Version available on Veas
DynInst	UMD/Wisc./IBM	Binary rewriter	Paradyn porting to BG/Q. Part of Mira SOW contract.
ValGrind	ValGrind/IBM	Memory & Thread Error Check	Funded by LLNL for BG/Q port to start this summer. Collaboration with LLNL.



BG/Q Parallel Debuggers Status

- **IBM CDTI (Code Development and Tools Interface)**
 - Collaboration of IBM/LLNL/ANL resulted in update v1.7 (August 2011)
 - Refined interface for multiple tool support, breakpoint handling, stepping, and signal handling
- **Rogue Wave TotalView**
 - Ported to BG/Q (Q32 at IBM) with basic functionality in August 2011
 - Pre-release testing by LLNL December 2011
 - Status
 - Tested working: basic ops (step, breakpoint, stack), QPX instructions, fast conditional breakpoints, job control for C/C++/Fortran with MPI/OMP/threads.
 - Still testing: **scalability**, fast conditional watchpoints, debugging in TM/SE
 - Working on research license on VEAS. Will present at second ESP workshop.
- **Allinea DDT**
 - Preparation via ANL scalability research contract on BG/P to address I/O node bottlenecks
 - Multiplexed debug daemons – complete and tested (Nov 2011)
 - Multiplexed gdbserver processes – complete and tested for single threading (Dec 2011)
 - BG/P Beta release (March 2012)
 - Status
 - Exploring performance of multiplexed gdbserver with multiple threads/process.
 - BG/Q port commenced on VEAS. Will present at second ESP workshop in April.



Debuggers status on VEAS

Debugger	Vendor	Provides	VEAS status
Totalview	RogueWave	Parallel debugger with GUI	Initial port to BG/Q available. Scalability enhancements in progress. Argonne negotiating research license for VEAS . Plan to present and work with users at April ESP workshop.
DDT	Allinea	Parallel debugger with GUI	Port to BG/Q in progress. Scalability work completed on BG/P will enable debugging the larger number of ranks/threads per I/O node of BG/Q. Research license currently in place.
gdb	IBM	Non-GUI debugger. Each gdb client can connect to a single MPI rank. At most 4 ranks may be debugged using separate clients.	Currently available. Porting system software to support interactive jobs on VEAS in progress.



Programming Models status on VEAS

Programming Model	Source	Provides	VEAS status
GA Toolkit	PNNL	One sided communication and high level library for array computations	Functional port using MPI-2 RMA. Optimized PAMI implementation is in progress (IBM+ANL+PNNL).
Berkeley UPC	Berkeley	PGAS variant of C	Uses GASnet. Port available with GCC and XLC.
GASnet	Berkeley	Active message communication	Port using PAMI available. Generic MPI version also available.
Charm++	UIUC	Dynamic load balancing and task virtualization	Functional port using PAMI. Generic MPI port is also available.
CoArray Fortran	Rice University	PGAS variant of Fortran	Uses GASnet. Debugging IBM compiler issues with translated Fortran source



BG/Q Libraries

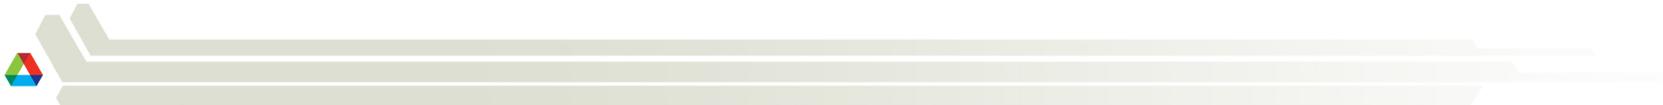
- ESSL available through IBM
- PETSc is being optimized as part of BG/Q Tools ESP project
- Porting and tuning 3rd party libraries (FFTW, BLAS, LAPACK, ScaLAPACK, ParMETIS, ...) using compiler optimizations
- Reporting and fixing library bugs found on BG/Q (e.g., ScaLAPACK 2.0.2)
- Collecting actual library usage data; libraries are stamped with a detectable string id
- Collaborating with Robert van de Geijn's group on rewriting Goto-BLAS so that it can be easily ported and tuned to new architectures like BG/Q (BLIS)
- Exploring an optimized FFT library with Spiral Gen



ALCF Supported BG/Q Libraries on VEAS

- Located in `/soft/libraries/alcf`
- Maintained in-house, frequently updated
- GNU and XL built versions of each library

Library Name	Source	Provides	VEAS status
ESSL	IBM	Dense Linear Algebra & FFT Kernels	5.1.1-0 beta version
BLAS	NETLIB (UTK) & ESSL	Dense Linear Algebra Kernels	Veas version is based on ESSL GEMM. All level-3 routines are GEMM-based.
BLIS	Univ. Texas & ANL	Framework for a successor to GotoBLAS	Under design; plans for hand-tuned BG/Q version
CBLAS	UTK	C wrappers to BLAS	On Veas
LAPACK	UTK	Dense Linear Algebra Solver	3.4.1
ScaLAPACK	UTK	Parallel Dense Linear Algebra Solver	2.0.2
ARPACK & PARPACK	Rice Univ.	Eigenvalues & Eigenvectors	2.1 (last version released in 2001)
FFTW	MIT	Fast Fourier Transform	2.1.5, 3.3.1. No hand-tuning for BG/Q yet.
METIS	UMN	Graph Partitioning	5.0.2
ParMETIS	UMN	Graph Partitioning	4.0.2



ALCF Unsupported BG/Q Libraries on VEAS

- Located in `/soft/libraries/unsupported`
- Provided as a convenience
- Not actively maintained by ALCF (for now)

Library Name	Source	Provides	VEAS status
Boost	OSS Community	C++ data structures & algorithms	1.49.0
HDF5	NCSA	Hierarchical data structures	1.8.8
NETCDF	UCAR	Portable data structures	4.1.3
P3DFFT	SDSC	Parallel FFT	2.4 (patched by ANL)
Tcl	OSS Community	Tool Command Language	8.4.14
zlib	Mark Adler	Compression	1.2.6

