SciDAC Software Infrastructure for Lattice Gauge Theory

Richard C. Brower

All Hands Meeting BNL, March 22-23 , 2007

SciDAC-2 kickoff workshop Oct27-28, 2006 http://super.bu.edu/~brower/workshop

Progress report: Sept 15, 2006 to Feb 1, 2007 http://super.bu.edu/~brower/scc.html

Code distribution see http://www.usqcd.org/software.html

QUIZZ

THIS IS THE 50th ANIVERSITY OF WHAT?

FORTRAN IS 50 YEARS OLD!

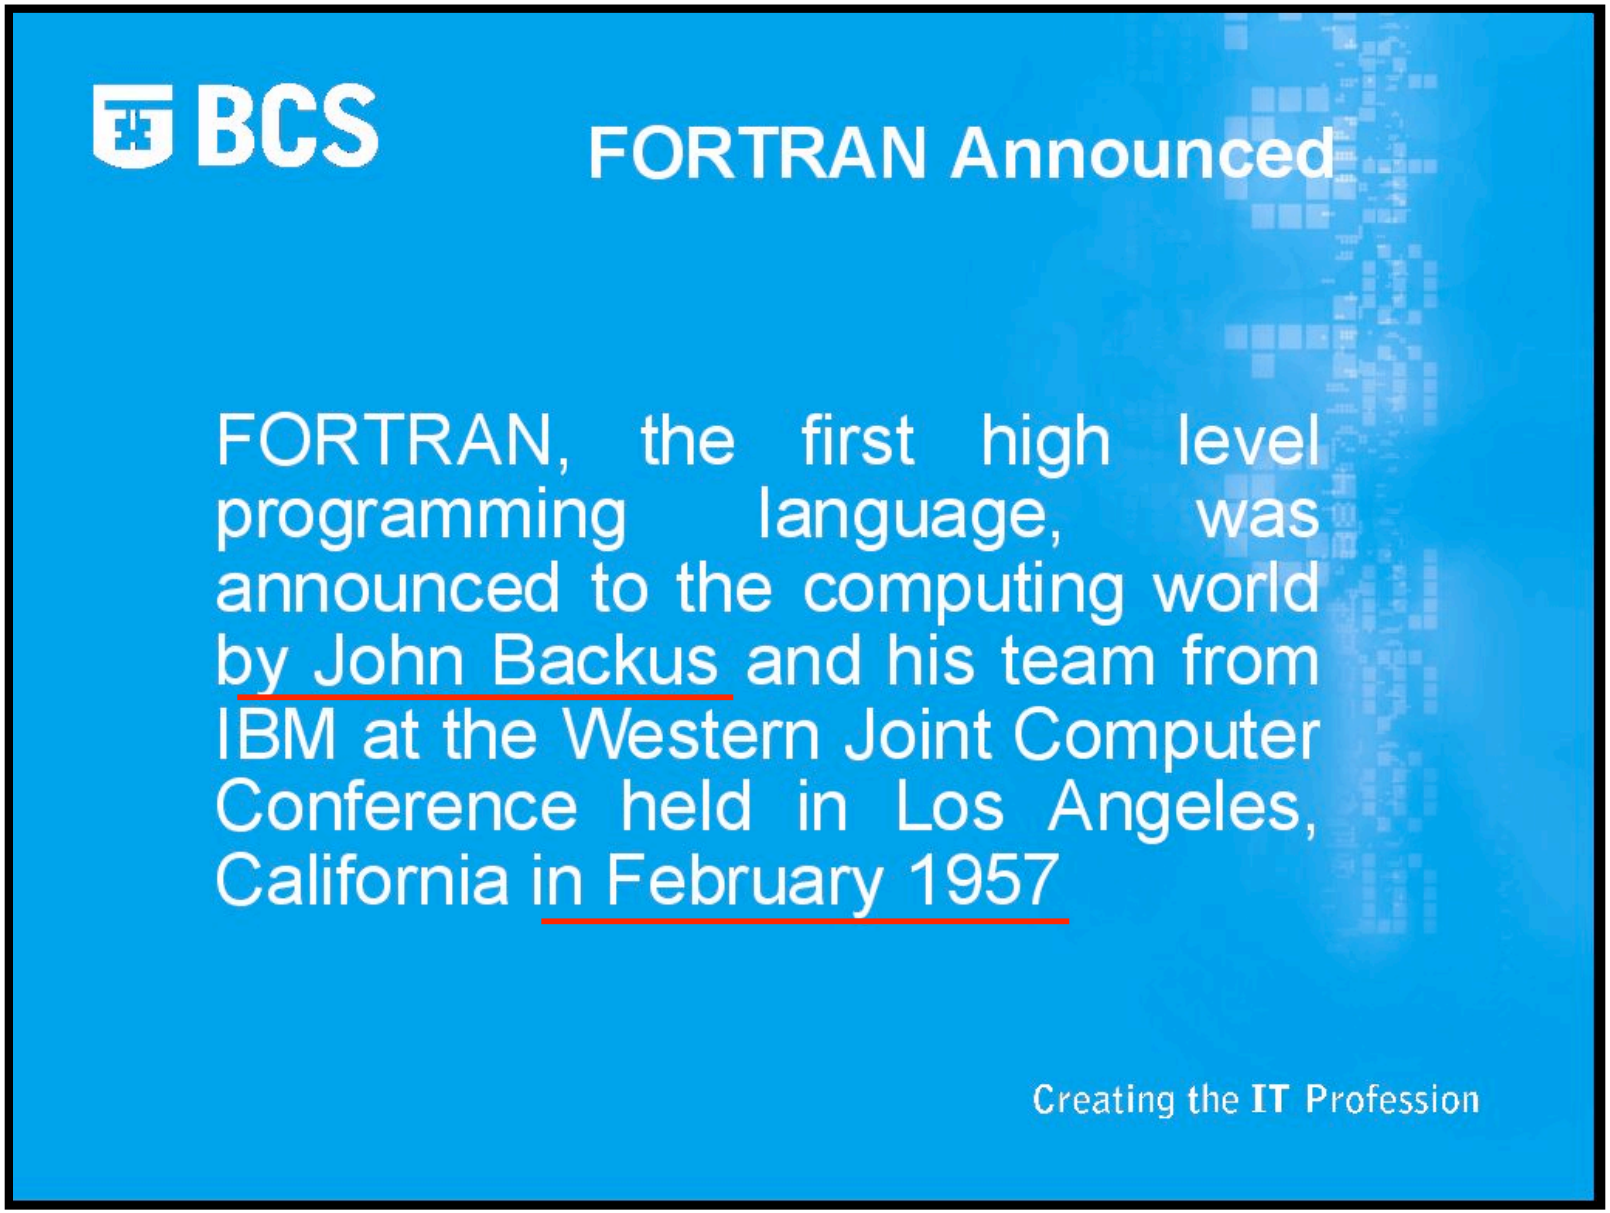

BCS

John Backus' team in the late 1950s

BCS

An IBM 704 mainframe

BCS

FORTRAN Conceived

In late 1953, John Backus sent a brief letter to his boss at IBM, asking that he be allowed to search for a "better way" of programming computers, with a project timescale of six months. He got the nod and began the research project that would eventually produce FORTRAN.

As John Backus says in the film, “project completion was always six months away”!

BCS

A FORTRAN anecdote

Frank Engel of Westinghouse, Pittsburg was concerned about the efficiency of the tape operations with the first FORTRAN compiler. He asked IBM if he could have a copy of the source code. They replied "IBM does not supply source code." So Frank worked his way through an octal dump of the compiler and optimised the tape operations. The improvement so impressed IBM that they asked for a copy of the code, to which Frank replied "Westinghouse does not supply source code."

Major Participants in SciDAC Project

Arizona	Doug Toussaint	MIT	Andrew Pochinsky
	Dru Renner		Joy Khoriaty
BU	Rich Brower *	North Carolina	Rob Fowler
	James Osborn		Ying Zhang *
	Mike Clark	JLab	Chip Watson *
BNL	Chulwoo Jung		Robert Edwards *
	Enno Schloz		Jie Chen
	Efstratios Efstathiadis		Balint Joo
Columbia	Bob Mawhinney *	IIT	Xien-He Sun
DePaul	Massimo DiPierro	Indiana	Steve Gottlieb
FNAL	Don Holmgren *		Subhasish Basak
	Jim Simone	Utah	Carleton DeTar *
	Eric Neilsen		Ludmila Levkova
	Amitoj Singh	Vanderbilt	Ted Bapty

* Software Committee: Participants funded in part by SciDAC grant

Institutions	*Oversight*
• BNL/Columbia	Mawhinney/ Chulwoo Jung
• JLab	Edwards/Watson
• FNAL/ITT/Vanderbuilt	Holmgren/Simone
• BU/MIT	Brower/Pochinsky
• DePaul/NorthCarolina	DiPierro/Zhang
• Arizona/Indiana/Utah	DeTar/Gottlieb/Toussaint

SciDAC-1 QCD API

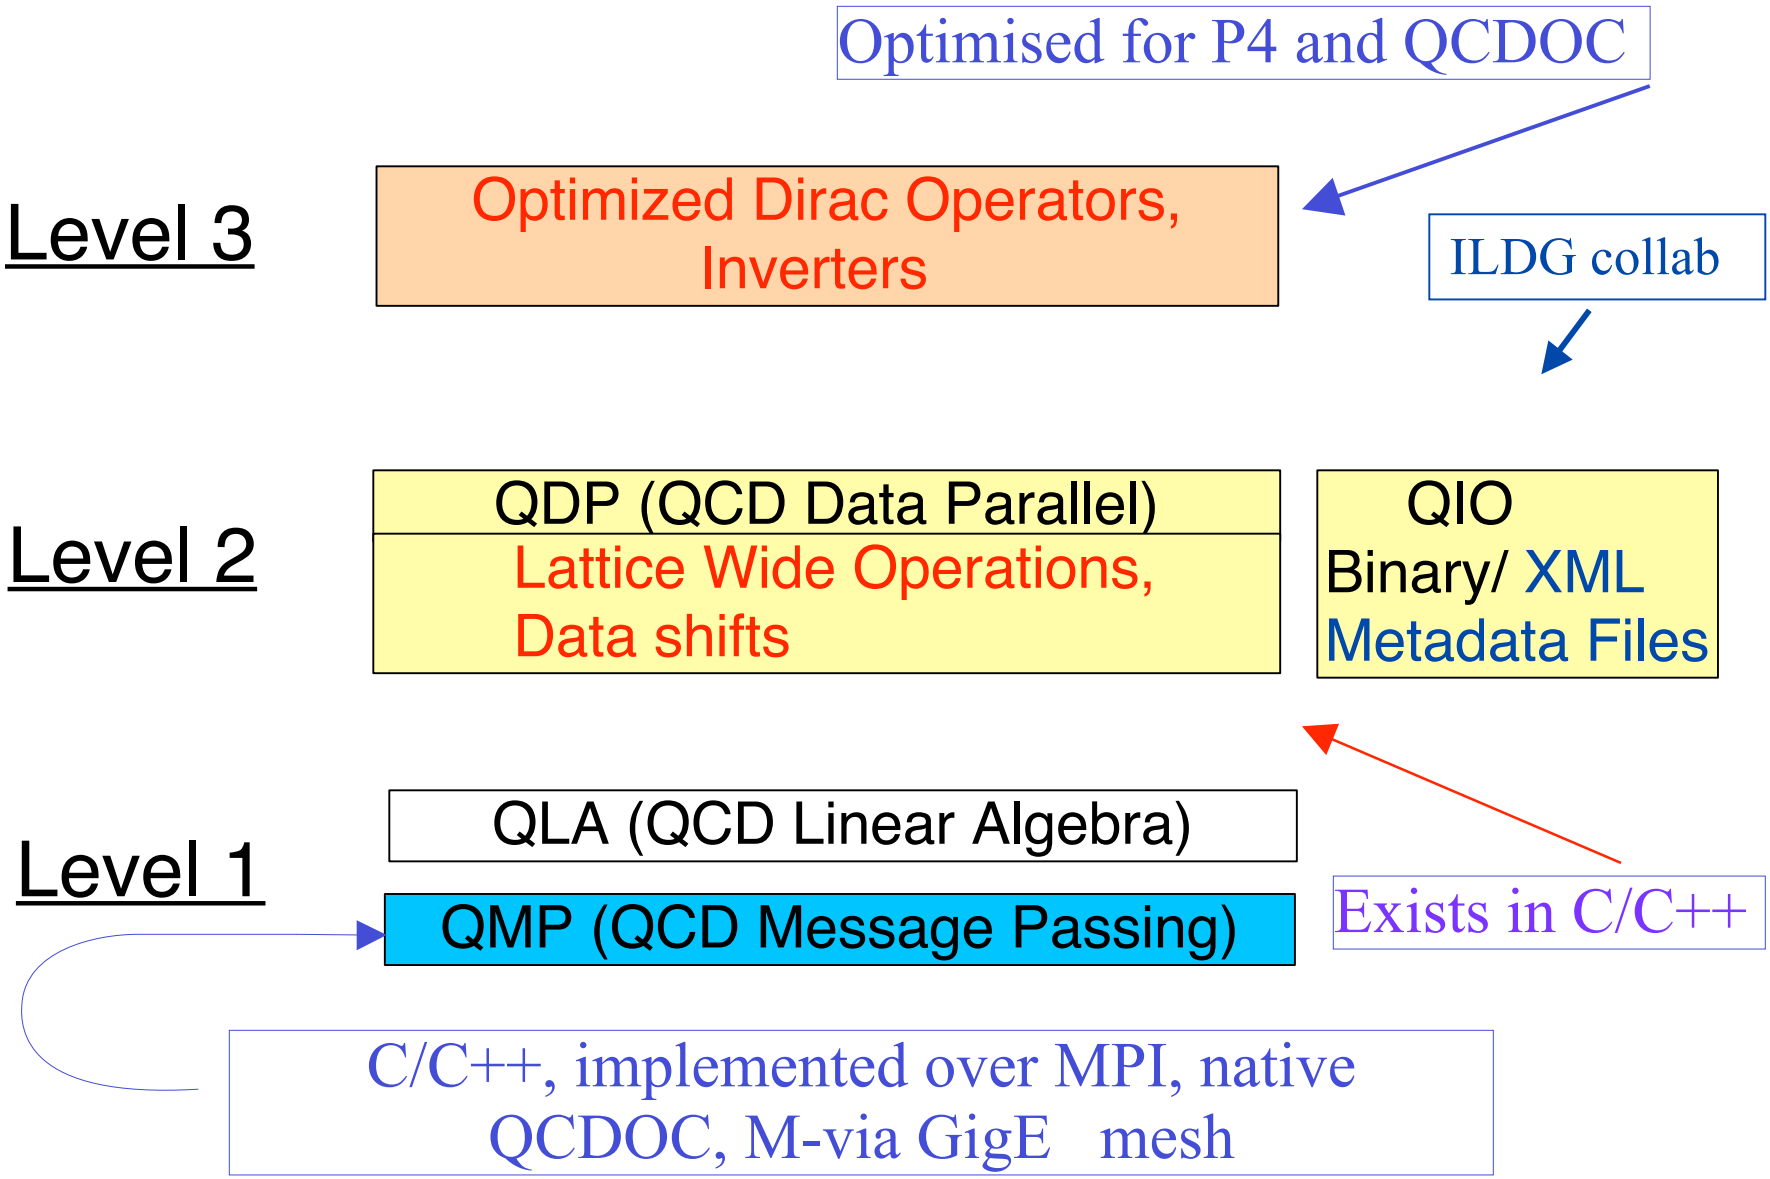

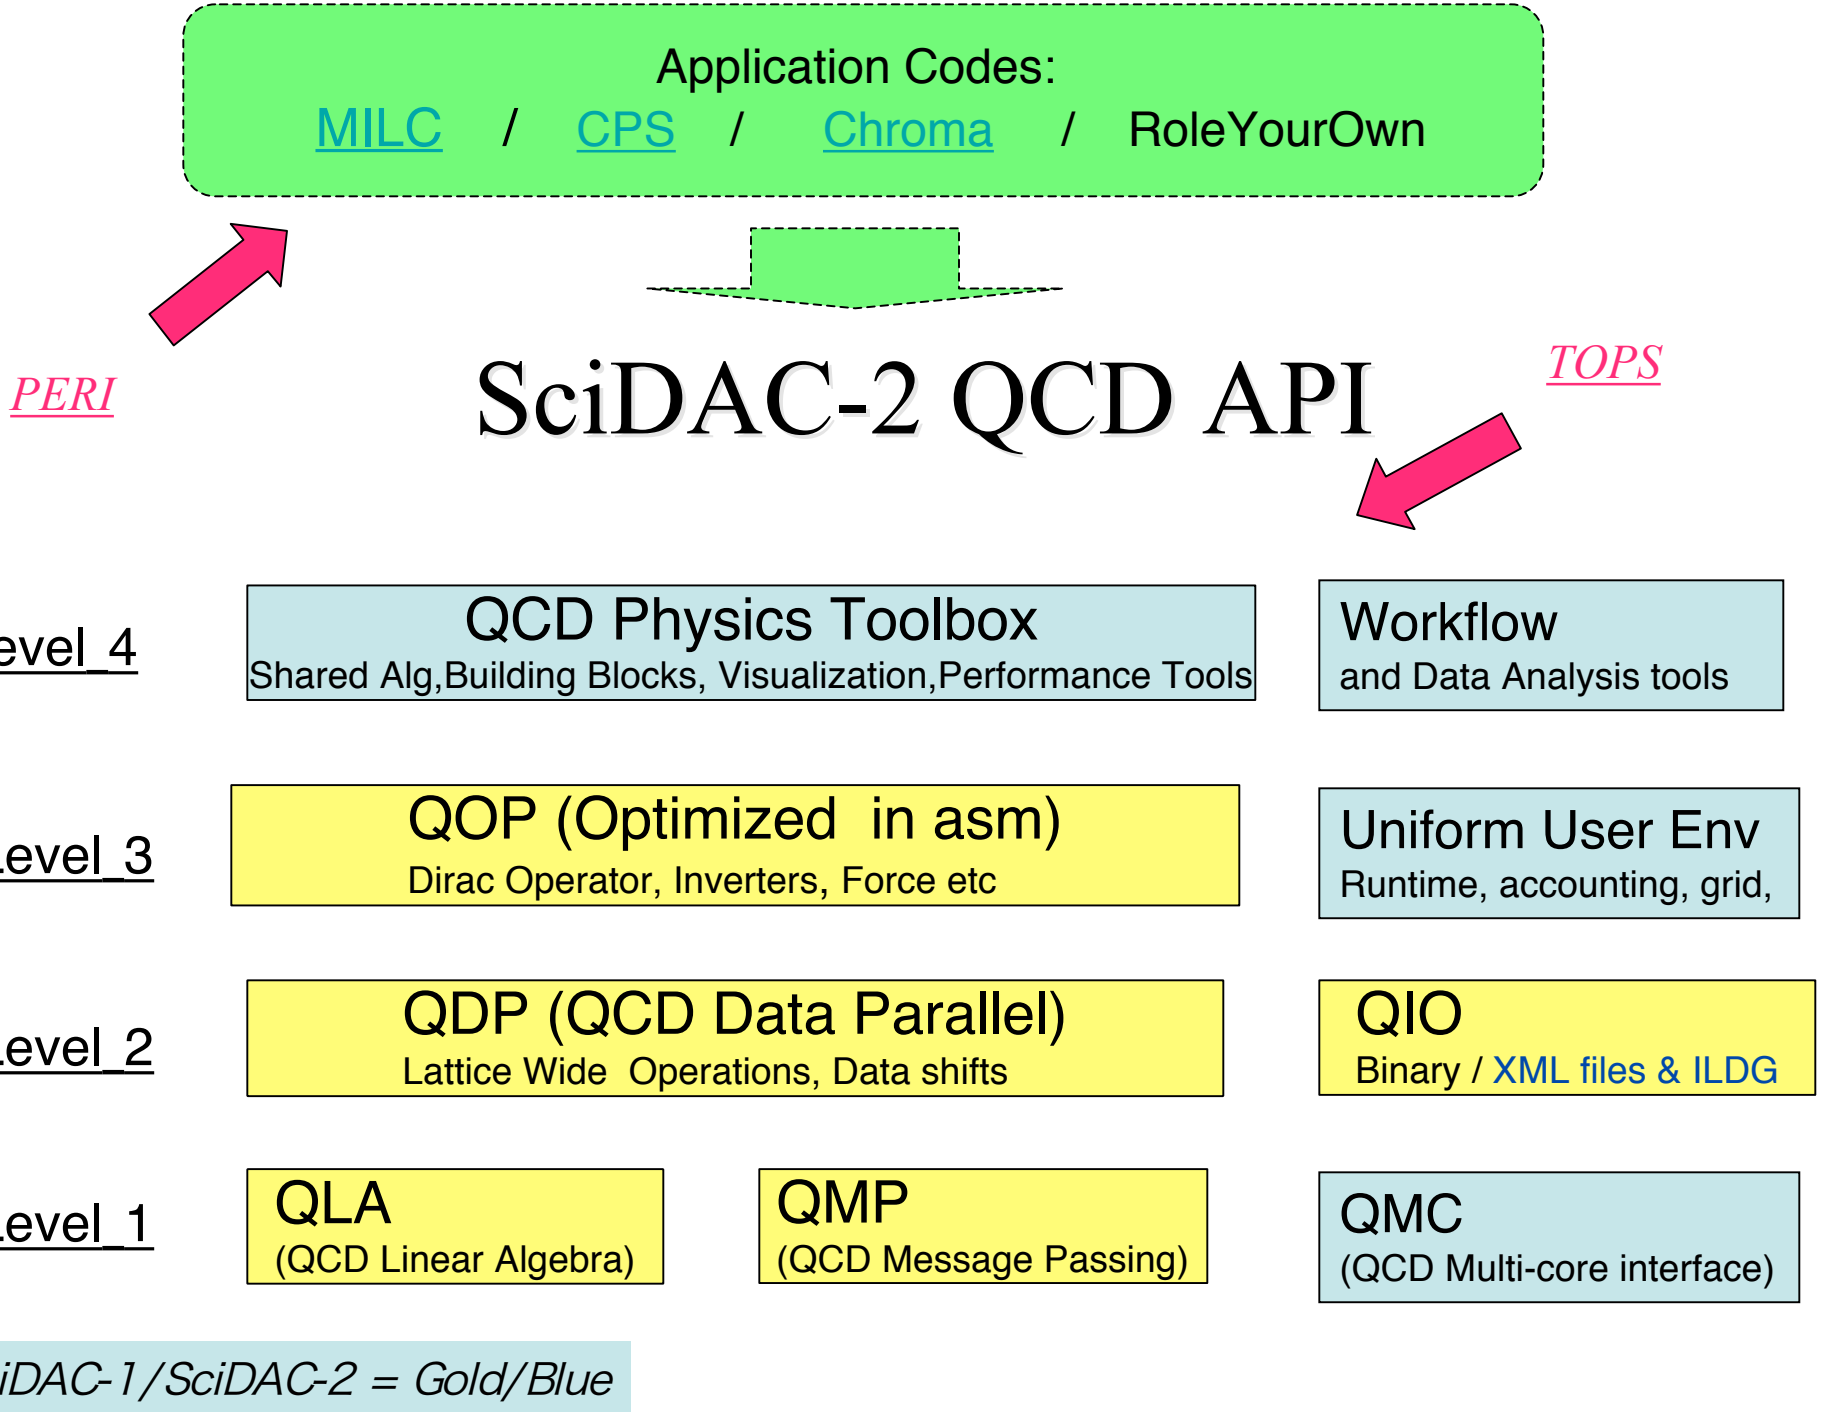
Application Codes:
MILC / CPS / Chroma / RoleYourOwn
PERI
SciDAC-2 QCD API
TOPS
Level_4
QCD Physics Toolbox
Shared Alg,Building Blocks, Visualization,Performance Tools
Workflow
and Data Analysis tools
Level_3
QOP (Optimized in asm)
Dirac Operator, Inverters, Force etc
Uniform User Env
Runtime, accounting, grid,
Level_2
QDP (QCD Data Parallel)
Lattice Wide Operations, Data shifts
QIO
Binary / XML files & ILDG
Level_1
QLA
(QCD Linear Algebra)
QMP
(QCD Message Passing)
QMC
(QCD Multi-core interface)
SciDAC-1/SciDAC-2 = Gold/Blue

Some current activities & Priorities

- Fuller use of API in application code.

  Round table: Software vs software

- Porting API to new Machines

  BG/L & BG/P: QMP and QLA using XLC & Perl script

  Cray XT3 & XT4: Opteron, 32 bit SSE, etc.

- Common Runtime Env. "Practical Meta-facility"

  File transfer, Batch scripts, Compile targets

- Workflow and Data Analysis

  Automate campaign to combine lattices, propagators
  to extract physical parameters. (FNAL Jim Simone & ITT)

- Tool Box (shared algorithms / building blocks)

  RHMC, eigenvector solvers, etc

  Visualization and Performance Analysis

- Exploitation of Multi-core

  Plans for a QMC API (JLab Jie Chen/ Edwards)

Status of QMP on BG/L

- based on QMP/MPI code base
- added --with-qmp-comms-type=BGL option
- native BG/L point-to-point (send/receive)
- uses MPI for everything else (collectives)
- requires barriers (MPI_Barrier) around some collectives (broadcast, binary_reduction)
- mostly done -- still needs cleanup & testing & (more)optimization

James Osborn

Performance of QMP on BG/L (contiguous quad-aligned buffers)

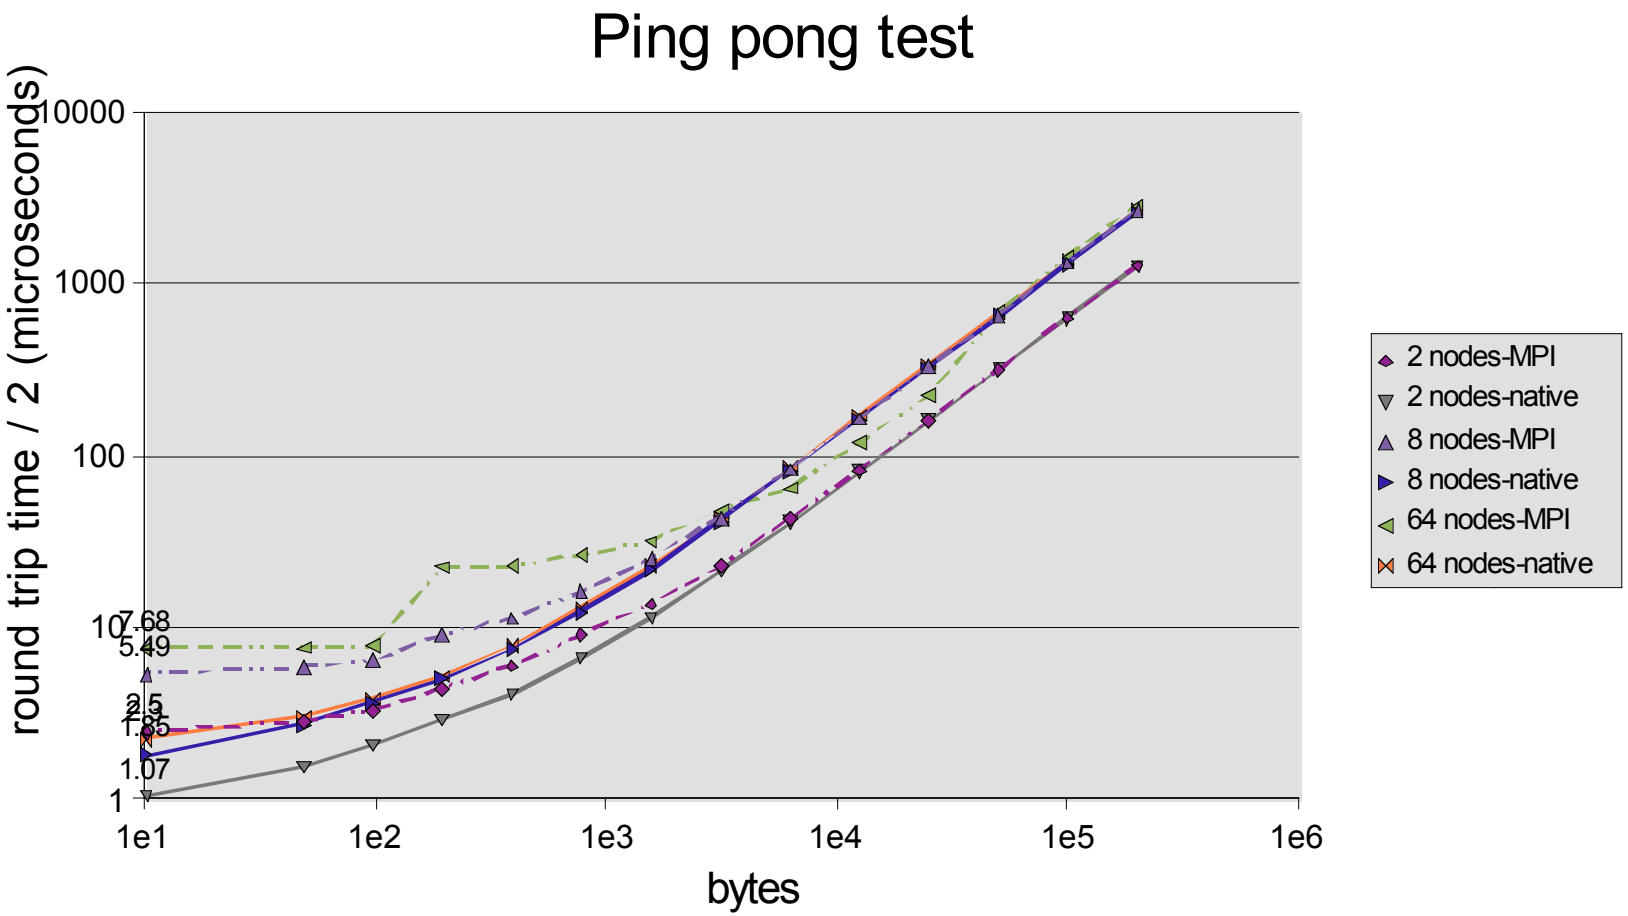

Status of QLA on BG/L

- previous version had a single 440 asm routine
- now has a 440d asm version of same routine
- development version now uses XLC v8 and C99 complex types (along with necessary alignment and disjoint hints) to make use of 440d
- has passed full testsuite running on BG/L
- BAGEL routines may still be useful

James Osborn, Joy Khoriaty & Andrew Pochinsky

Performance of QLA on BG/L (QOPQDP – asqtad inverter)

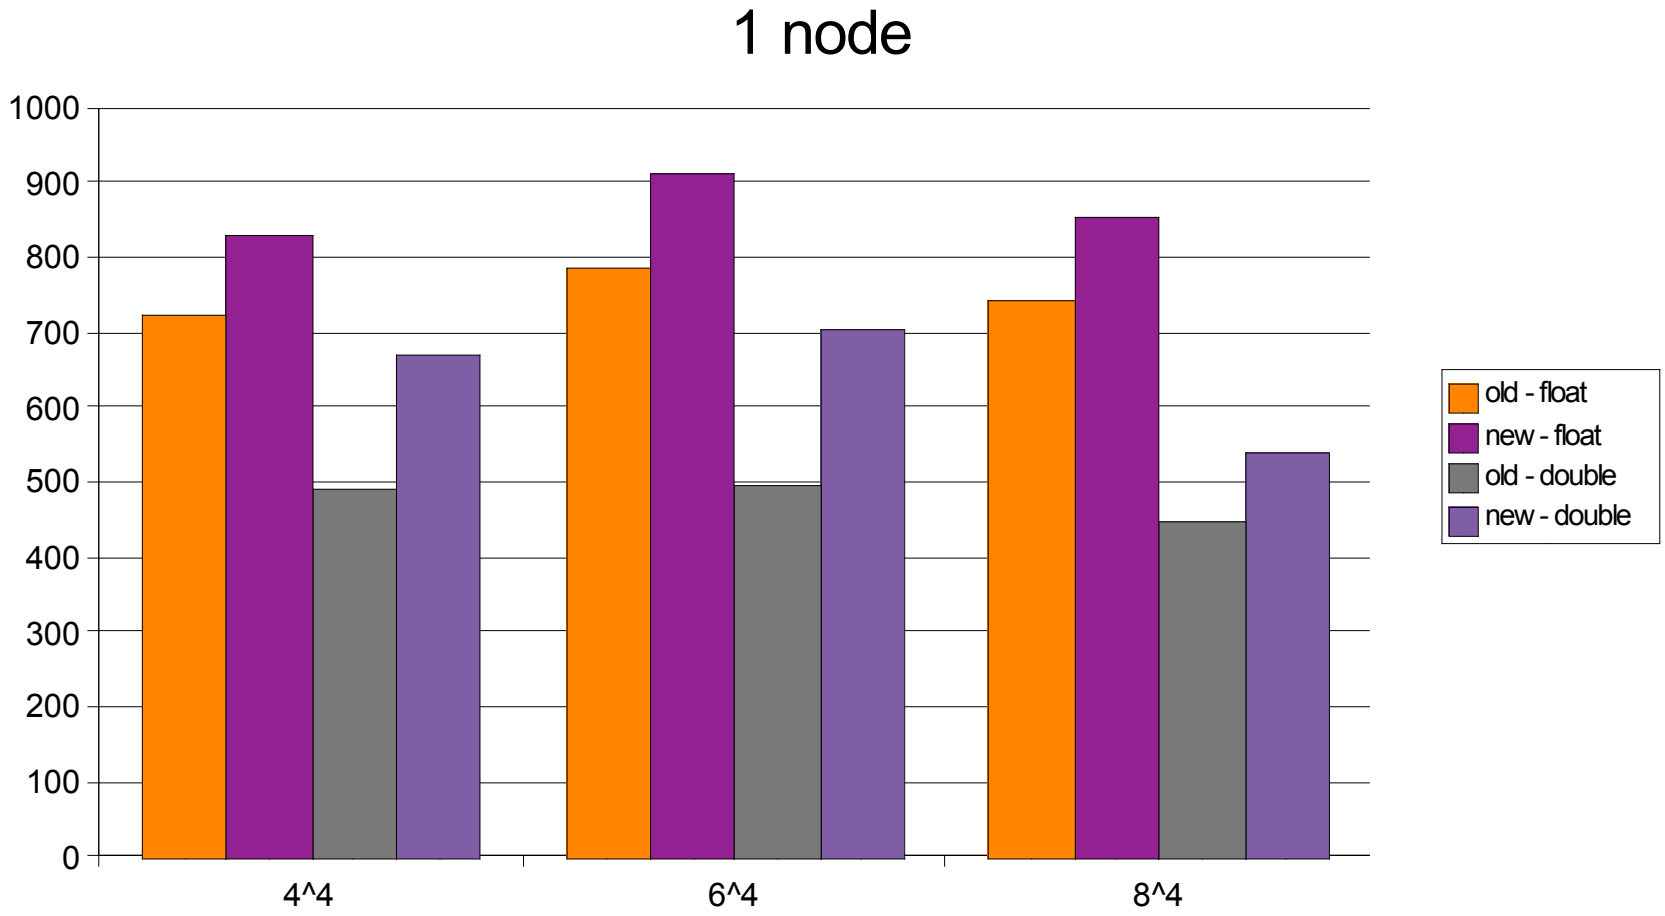

Performance of QLA on BG/L (QOPQDP – Wilson inverter)

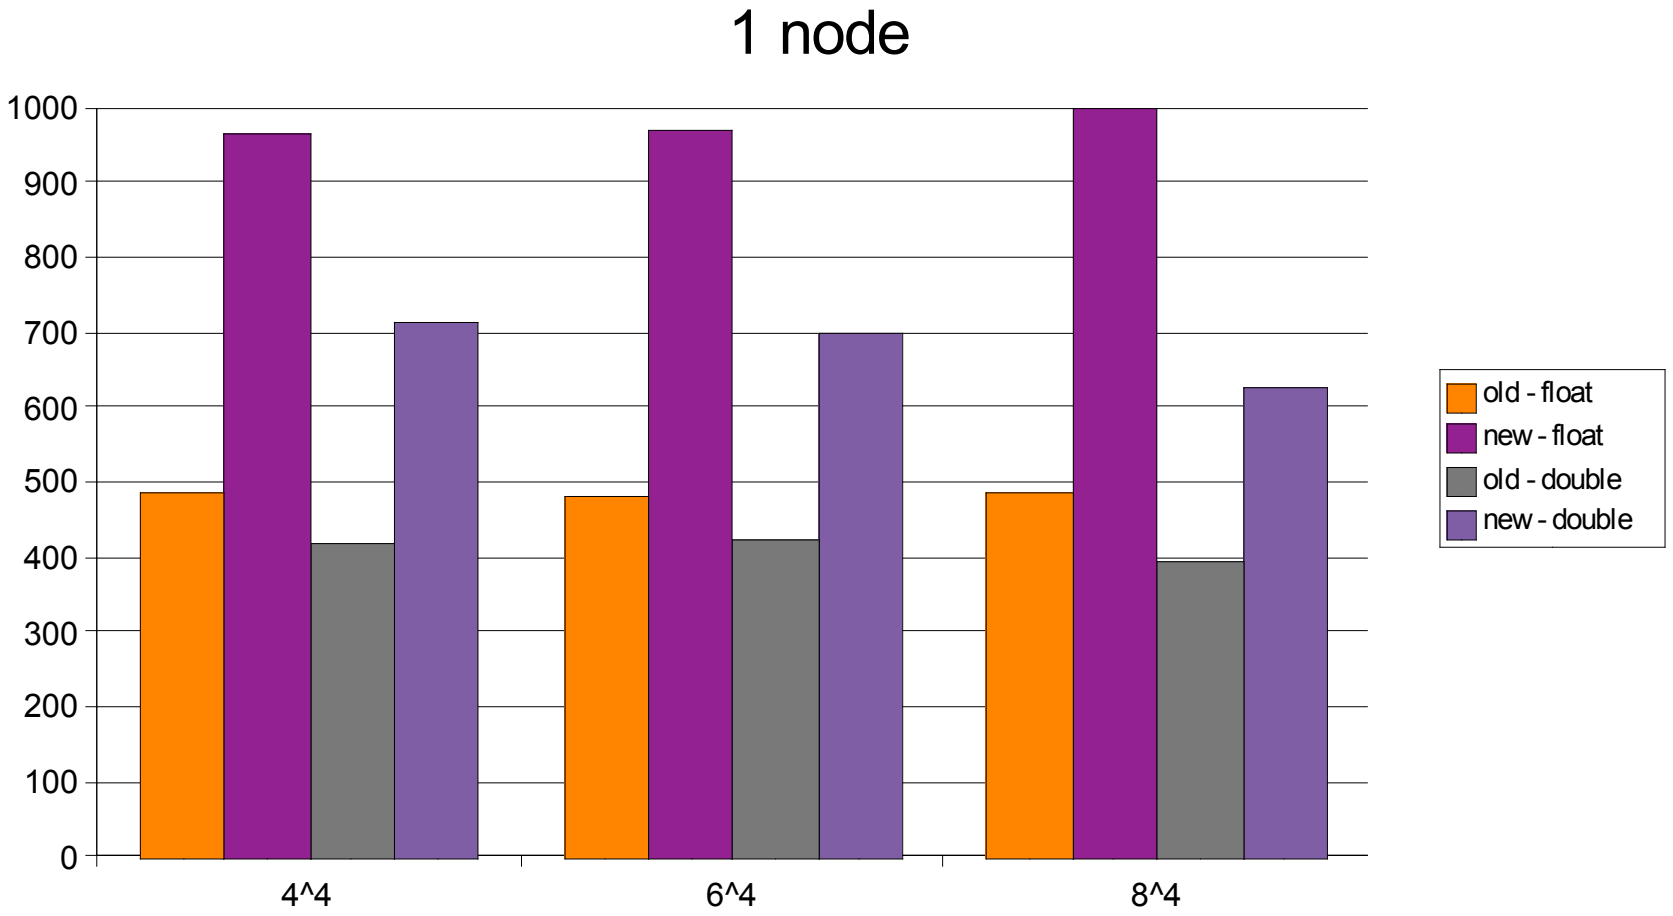

Performance of QMP+QLA on BG/L (QOPQDP – asqtad inverter)

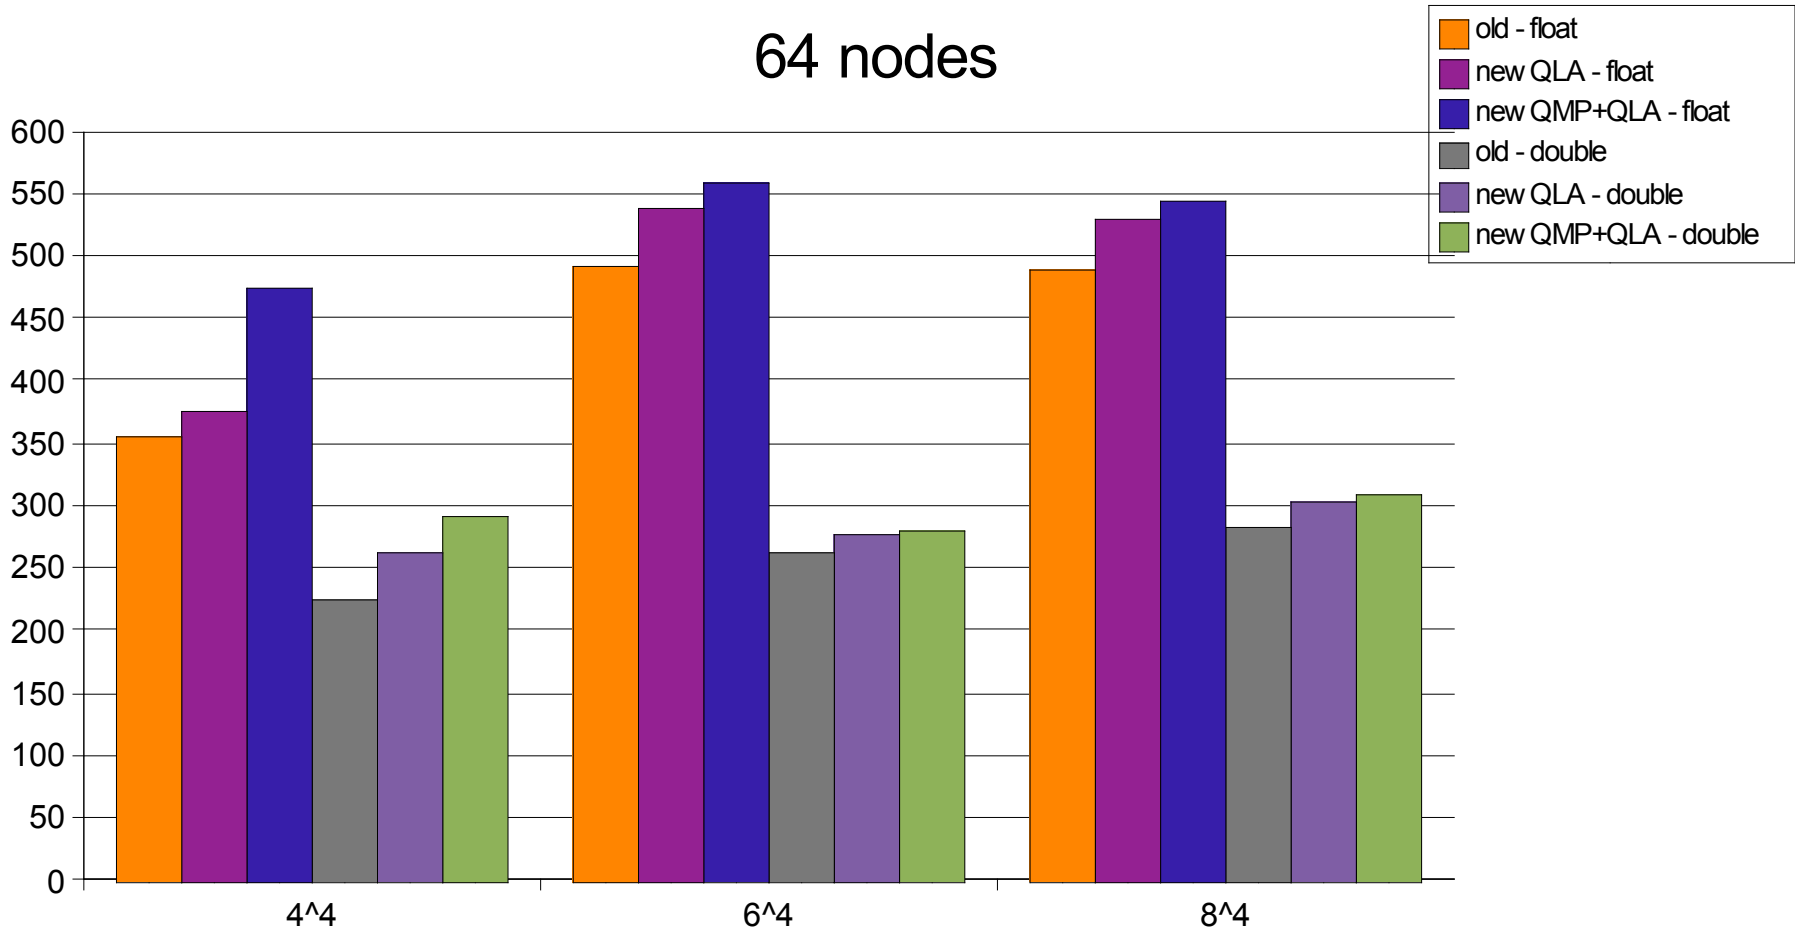

Performance of QMP+QLA on BG/L (QOPQDP – Wilson inverter)

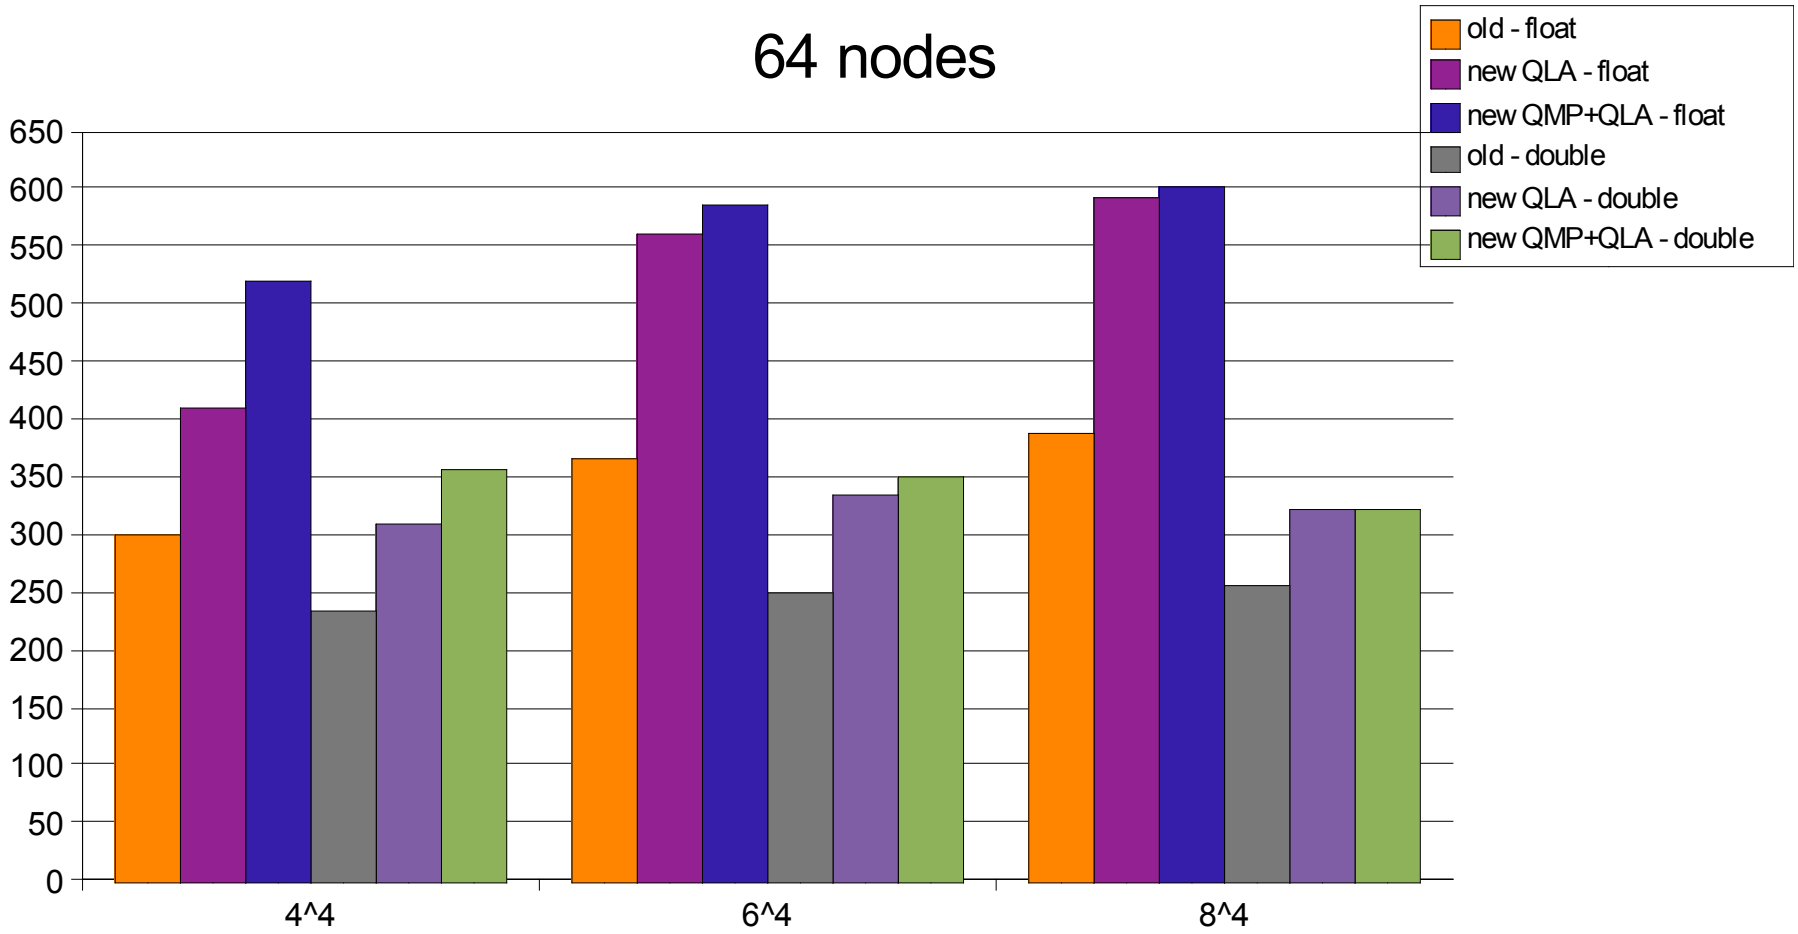

Software Committee

- Rich Brower (chair) brower@bu.edu
- Carleton DeTar detar@physics.utah.edu
- Robert Edwards edwards@jlab.org
- Don Holmgren djholm@fnal.gov
- Bob Mawhinney rdm@phys.colmubia.edu
- Chip Watson watson@jlab.org
- Ying Zhang zhang@cs.uiuc.edu

QLA on Opterons (kaon)