SciDAC Software Infrastructure for Lattice Gauge Theory

Richard C. Brower

Sept 23, ILFT04 Shuzenji, Japan

(under construction) http://usqcd.org

(see lqcd.org for general gateway)

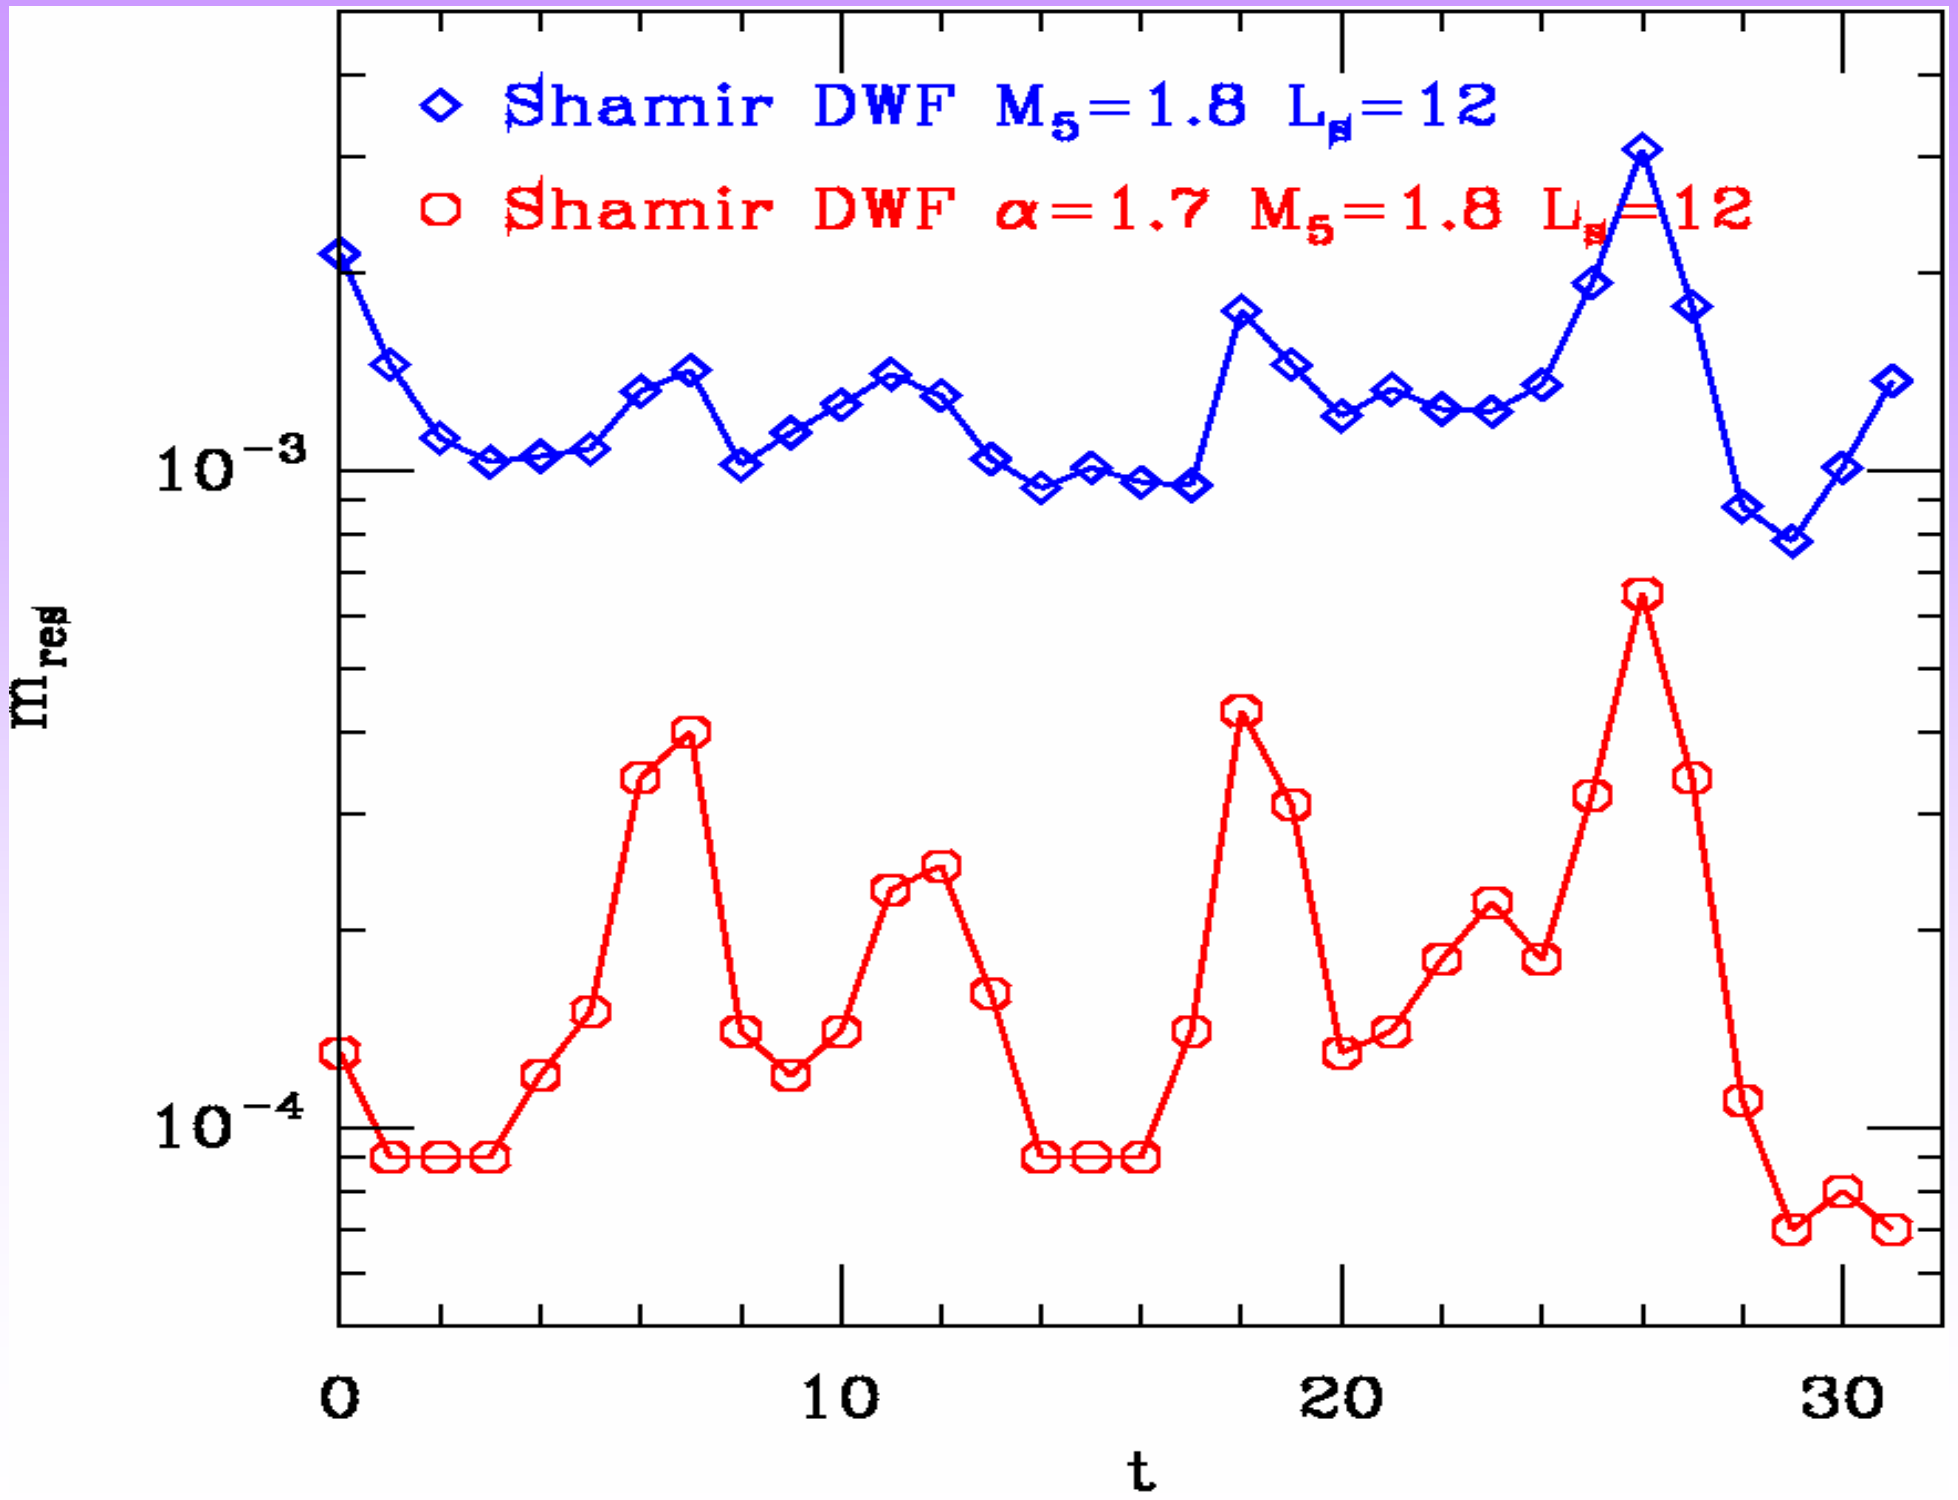
Shamir DWF $M_5=1.8$ $L_s=12$
Shamir DWF $\alpha=1.7$ $M_5=1.8$ $L_s=12$
m_{res}
10^{-3}
10^{-4}
0
10
20
30
t

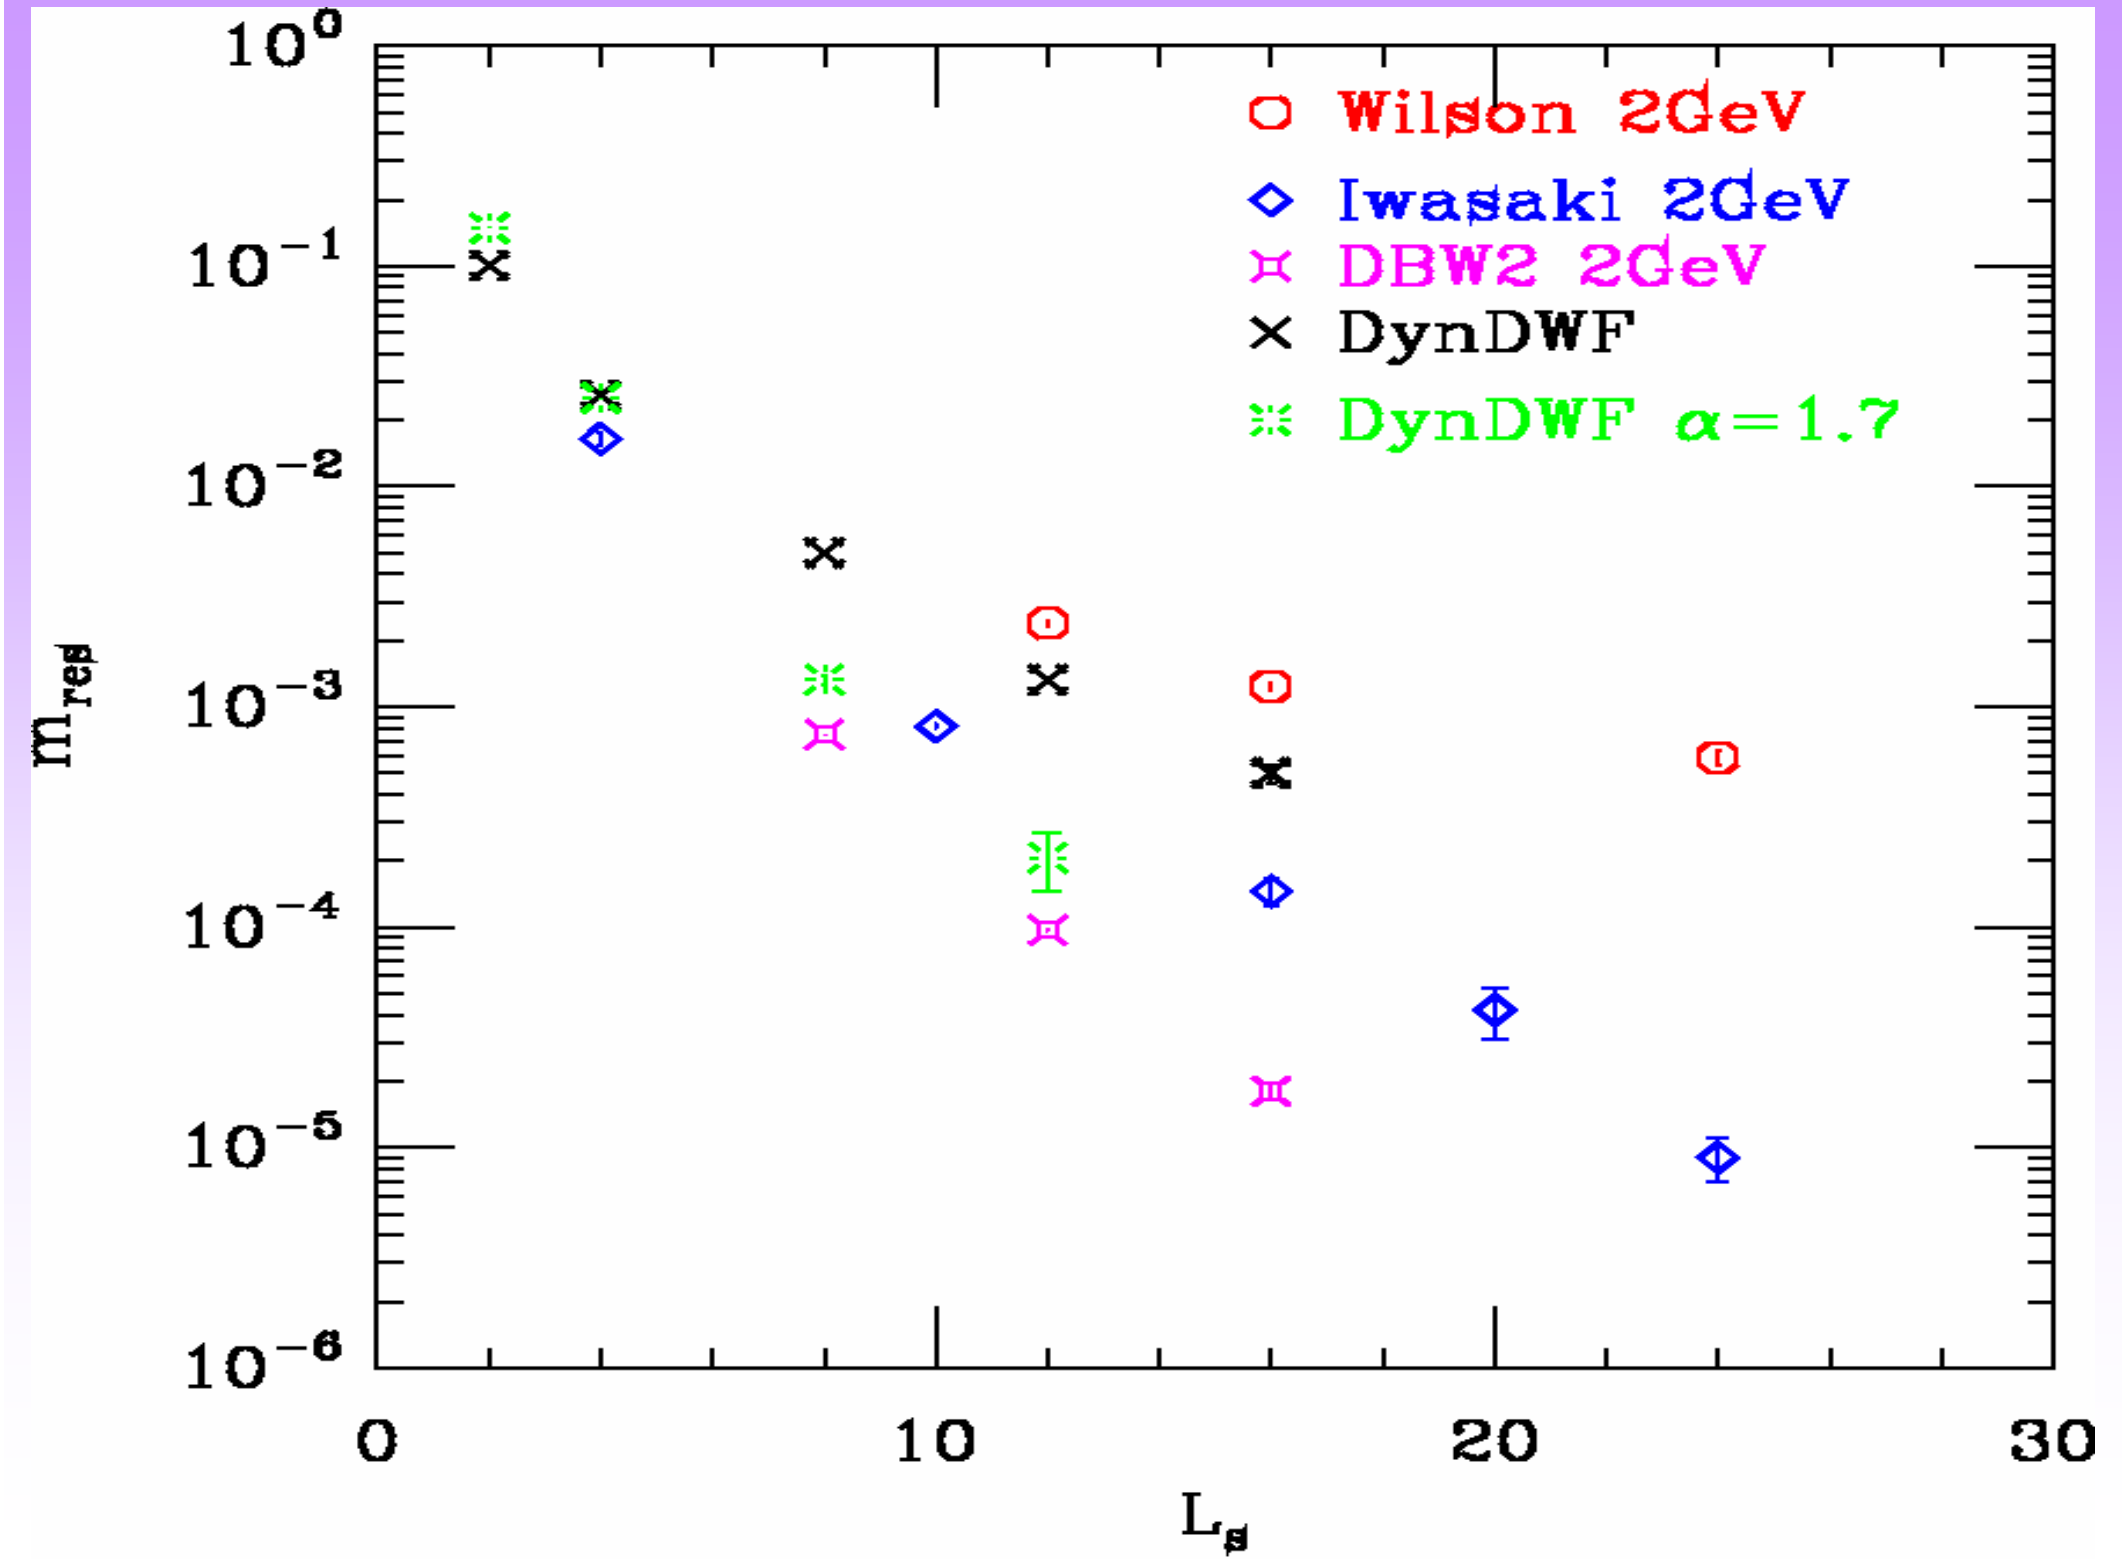
Wilson 2GeV
Iwasaki 2GeV
DBW2 2GeV
DynDWF
DynDWF α=1.7
m_{res}
L_s
10^0
10^{-1}
10^{-2}
10^{-3}
10^{-4}
10^{-5}
10^{-6}
0
10
20
30

Participants in Software Project

Arizona	Doug Toussaint		FNAL	Amitoj Sing	
Arizona	Dru Renner (Eric Gregory)		Illinois	Ying Zhang (Celso Mendes)	*
BU	Rich Brower	*	Illinois	Daniel Reed	
BU	James Osborn (Hartmut Neff)		JLab	Robert Edwards	*
BNL	Chulwoo Jung		JLab	Chip Watson	*
BNL	Chris Miller		JLab	Jie Chen	
BNL	Kostantin Petrov		JLab	Walt Akers	
Columbia	Bob Mawhinney	*	MIT	Andrew Pochinsky	
FNAL	Don Holmgren	*	Utah	Carleton DeTar	*
FNAL	Jim Simone		Utah	Francesca Maresca	
FNAL	Eric Nielsen				

(UK Peter Boyle & Balint Joo)

* Software Coordinating Committee

SciDAC Software Structure

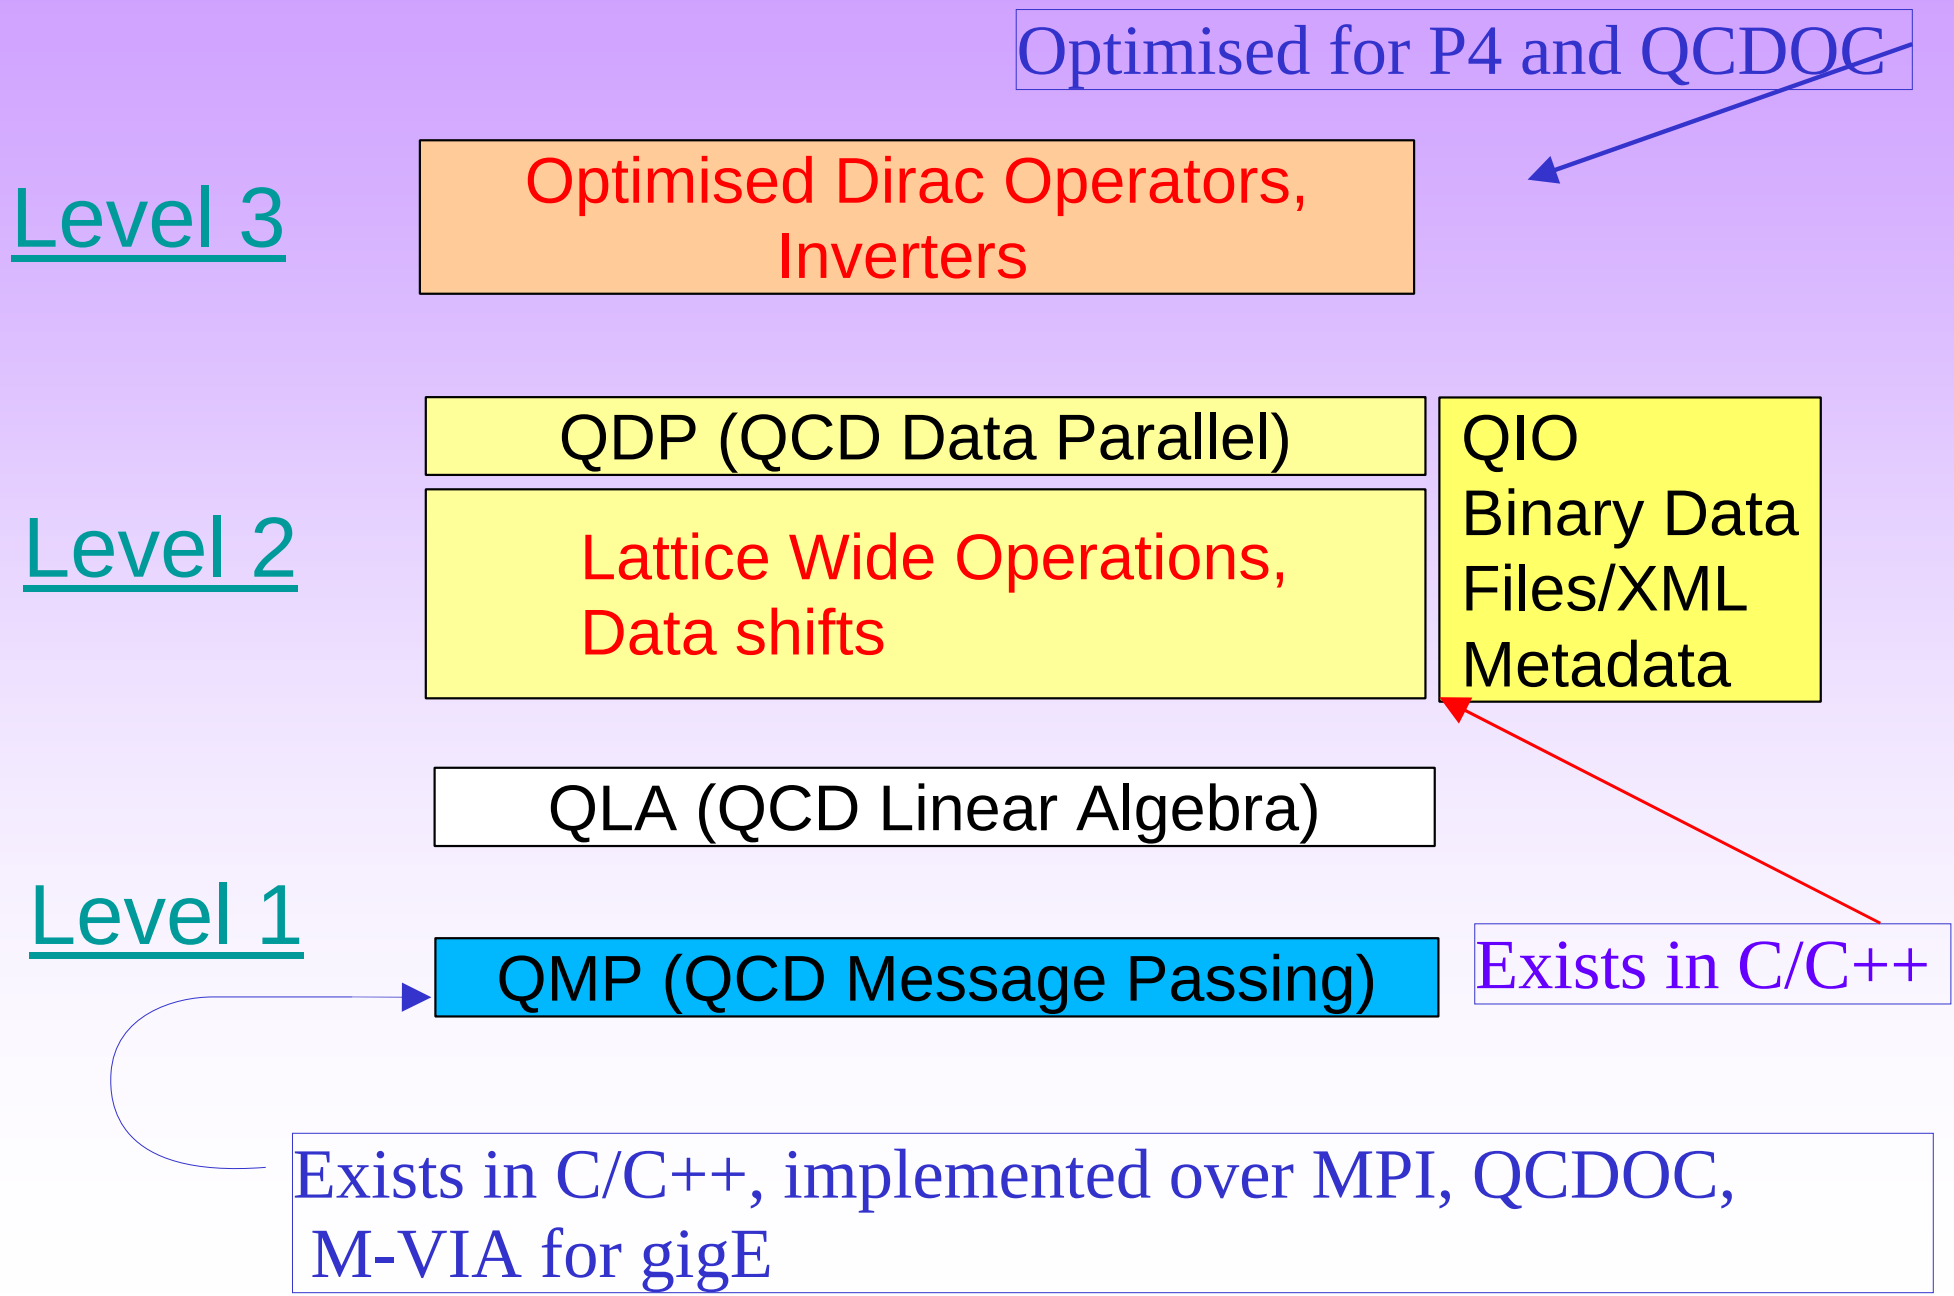

Performance of QCDOC

	Sites/node	Optimized			MILC	MILC/QDP
Precision		Double			Single	
Application		D	D	CG	CG	
Machine size		Sim	128	128	128	16
Asqtad Dirac/CG	2^4		22%	19%		2%
	4^4	43%	42%	40%	9.5%	15%
	6^4		37%	36%	13.6%	19%
	8^4		29%	28%	12.3%	
Machine size		128			128	
Asqtad Force	2^4	19%				
	4^4	35%			10.0%	
	6^4	29%			10.8%	
	8^4	20%			8.2%	
Asqtad fatlink/ Gauge Force*	2^4	23%				
	4^4	57%			7.2%	
	6^4	34%			8.1%	
	8^4	27%			8.5%	

(*):multi-dimensional SU(3) parallel transport

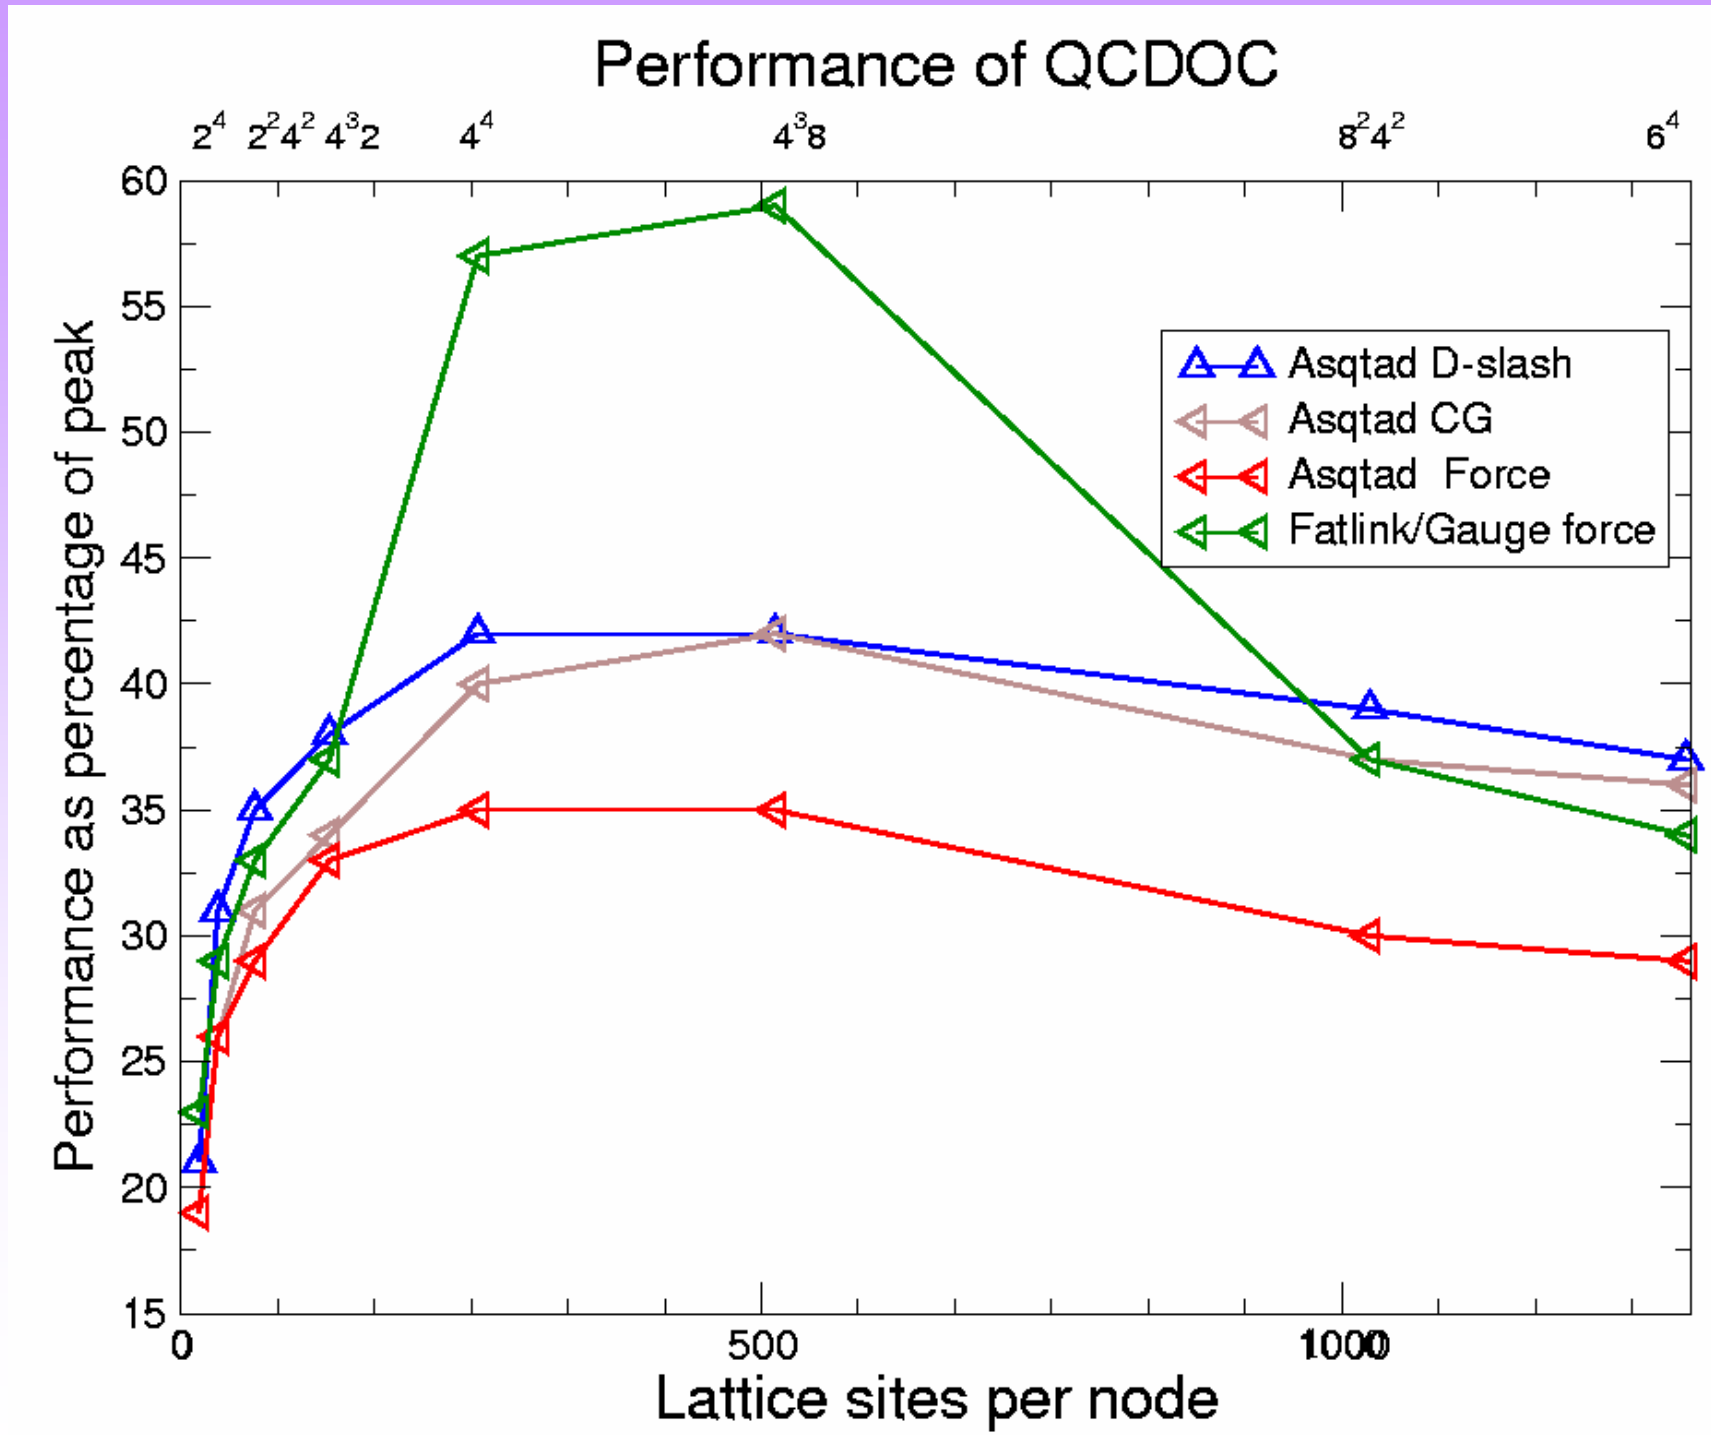
Performance of QCDOC
2^4 $2^2 4^2$ $4^3 2$ 4^4 $4^3 8$ $8^2 4^2$ 6^4
Performance as percentage of peak
60
55
50
45
40
35
30
25
20
15
0
500
1000
Lattice sites per node
Asqtad D-slash
Asqtad CG
Asqtad Force
Fatlink/Gauge force

Level 3 QCDOC- April 29, 2004

(Performance is % of Peak: Run at 420 MHz)

	Sites/node	Optimized Code			MILC Code			MILC Code			MILC/QDP
Precision		Double			Single			Double			Single
Application		D	D	CG	CG			CG			CG
Machine size		Sim	128	128	128			16			16
Memory used					DDR	CG temps EDRAM	CG temps+lat EDRAM[1]	DDR	CG temps EDRAM[1]	CG temps+lat EDRAM[1]	
Wilson[4]	2^4	47%	44%	32%							
	4^4		43%	38%							
	6^4		44%	39%							
	8^4		29%								
ASQTAD[2]	2^4		22%	19%							2%
	4^4	43%	42%	40%	7.6%	8.5%	9.5%	5.9%	6.5%	7.5%	15%
	6^4		35%	33%	9.9%	11.5%	13.6%	7.6%	8.6%		19%
	8^4		29%	28%	10.5%	12.3%		8.1%			
Clover[4]	2^4			31%							
	4^4			47%							
DWF[4]	$2^4 \times 4$			32%							
	$4^4 \times 4$			42%							
Application		ASQTAD force			ASQTAD force[3]			ASQTAD force[3]			
Machine size		Sim	128	128	128			16			
Memory used					DDR	temps EDRAM	temps+lat EDRAM[1]	DDR	temps EDRAM[1]	temps+lat EDRAM[1]	
ASQTAD[2]	4^4				7.8%		10.0%	5.5%	5.5%	6.1%	
	6^4				8.2%		10.8%	6.1%	6.1%		
	8^4				8.2%			6.0%			

- Level 3 SciDAC DWF inverter (A. Pochinsky) – minimizes mem bus usage

256-node 3d gigE2.66Ghz P4, 533Mhz mem bus -> 384 Gflops

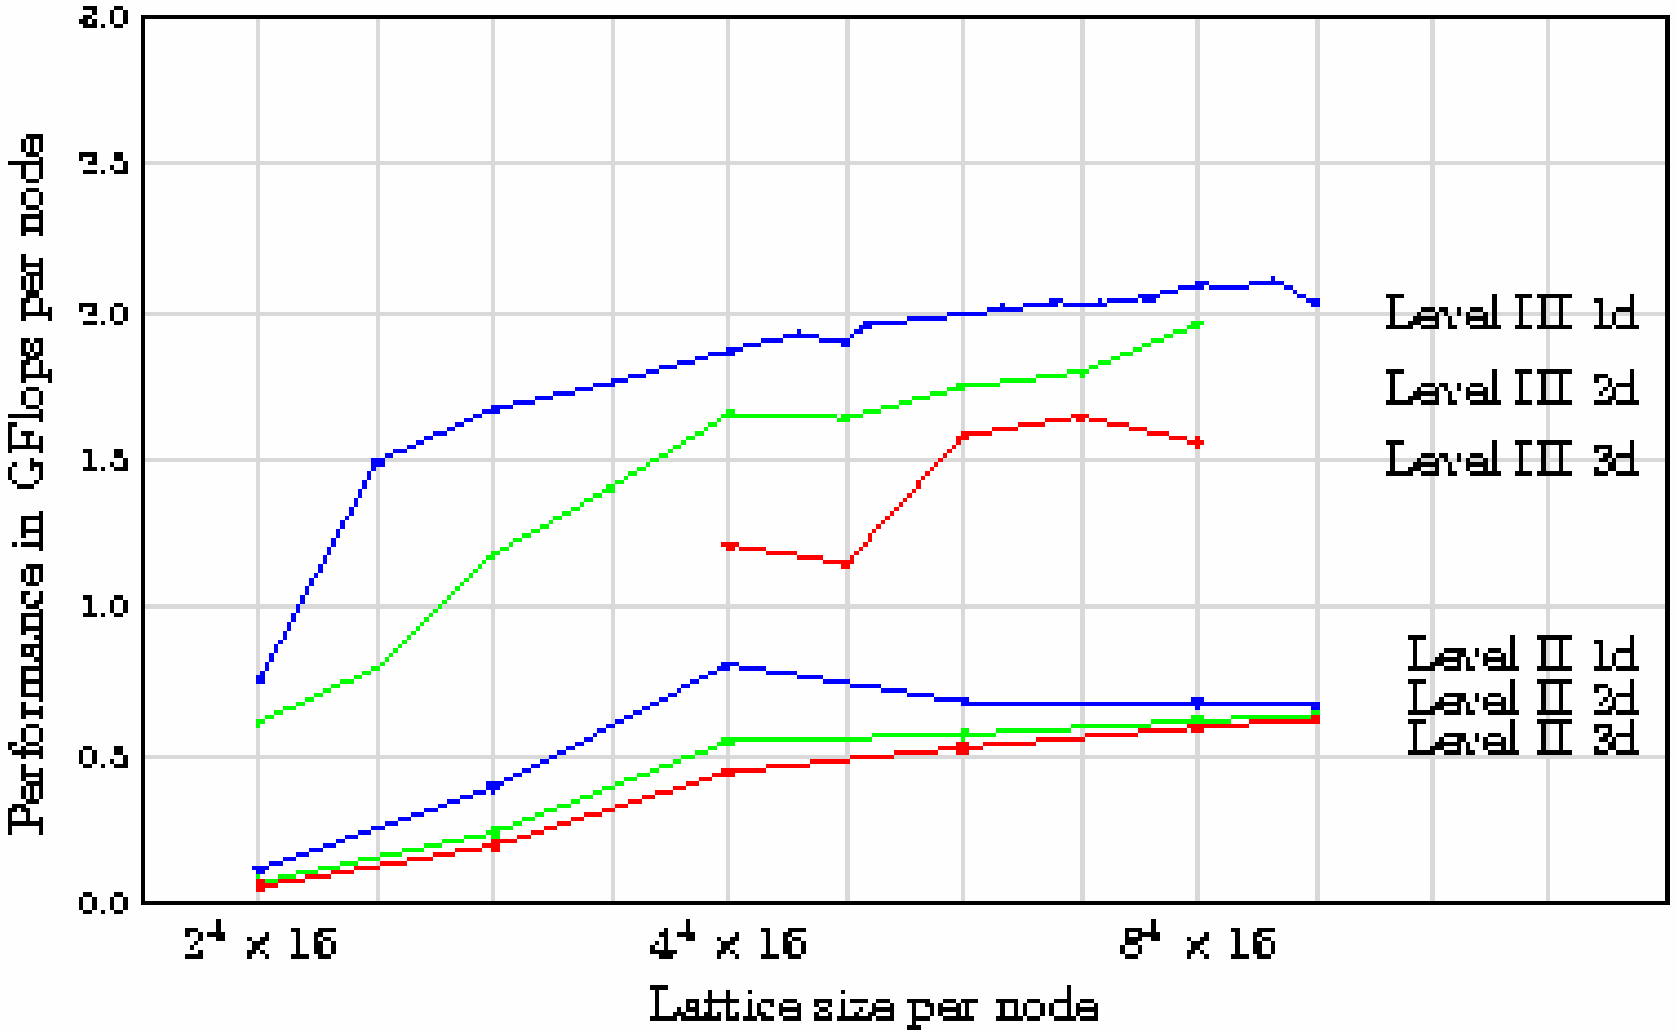

Some Current Tasks

- Level 3 Inverters: Asqtad odd # sites/axis/processor, single precession.
  - (Chulwoo Jung, Adrew Pochinsky, Peter Bolye)
  - Interface to MILC/QDP (Chroma)
- Level 1 QMP: Version 2 Full Compliance/Testing
  - (Chuwloo Jung/Jie Chen/Celso Mendes/Ying Zhang/Chip Watson)
- Level 2 QIO: Implementation of QIO/Host architecture Interface
  - (Carleton Detar/Peter Boyle/Kostya Petrov/Jie Chen)
- Level 2 QDP: opt for MILC QDP++: opt for Chroma CPS
  - (James Osborn/Robert Edwards/Peter Boyle-BAGEL/Balint Joo)
- Runtime Env: Uniform Queuing/Scripting/File Management
  - (Holmgren/FNAL, Petrov/BNL, Watson/Jlab)
- ILDG : Metadata (Edwards/Simone) & Middleware (Watson)

Performance on Non-Scidac Machines ?

Non SciDAC Target Systems for Ken Roche	
System	System Site
IBM SP Power 3	ORNL
IBM SP Power 3	NERSC
IBM p690 (Federated)	ORNL
SGI Altix	ORNL
Bluegene/L	IBM, Yorktown
Cray X1	ORNL
HP Cluster (Itanium2,Longs Peak nodes)	PNL
Cray Red Storm	SNL