

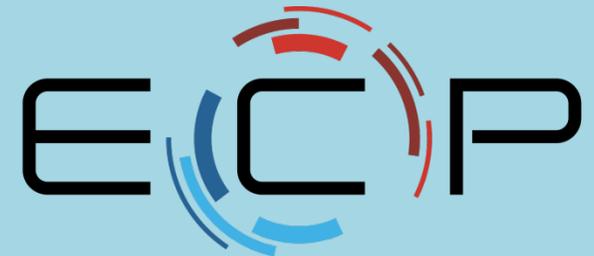
ECP Software: Meeting the Exascale >50X FOM

Carleton DeTar (University of Utah)

USQCD All-Hands Meeting

April 27, 2019

Brookhaven National Laboratory



EXASCALE COMPUTING PROJECT



U.S. DEPARTMENT OF
ENERGY

Office of
Science



Application Figure of Merit (FOM)

- Three fermion actions: DWF, Wilson-Clover, HISQ
 - Two measures each
 - HMC time for one trajectory.
 - Time for analysis.
 - Scale times to whole machine.
 - FOM = geometric mean of times
 - Ratio of FOM to baseline FOM (Titan, Mira) measures improvement
 - Step scaling: increase problem size -> weak scaling factor
- Goal is 50X for Aurora or Frontier

What the FOM leaves out

- Accelerated HMC (Critical slowing down group)
- More efficient treatment of contractions
- Doesn't track improvements in eigensolvers, trace calculations

First Summit result

Item	Component	Time(h)	Machine	Size(nodes)	Fraction	New Time x fraction	Old Time x fraction	ratio
1	DWF gauge generation	1.43	Summit	128	0.02783	0.03979130	0.111166667	2.8
2	DWF analysis	5.26	Mira	32768	0.66667	3.50666667	3.506666667	1
3	HISQ gauge generation	2.32	Summit	96	0.02087	0.04841739	0.395	8.2
4	HISQ analysis	1.11	Summit	108	0.02348	0.02606087	0.4325	16.6
5	Clover gauge generation	0.109	Summit	32	0.00696	0.00075826	0.060821918	80.2
6	Clover analysis	0.0165	Summit	16	0.00348	0.00005739	0.000323973	5.6

- First Summit FOM: 7.5 X improvement over baseline

Short-term strategies for improvement on Summit

- Summit has 6 GPUs per node which doesn't work well when lattice dimensions don't divide by 3. So change/increase lattice size (DWF and Clover-Wilson).
- Fix GPU Direct on Summit
- Further QUDA improvements coming: NVSHMEM
- DWF: Analysis benchmark will have additive-Schwarz preconditioning and deflation.
- HISQ (MILC): More support for multiple rhs and mixed precision
- QMP support for node sub-topology?

Strategies for Aurora and Frontier

- Most likely we will be dealing with accelerators, so need GPU support that is not vendor specific
 - Grid on GPU (Boyle) or a more flexible QUDA (Kate?)
 - OpenMP offloading (Meifeng)
- Communication will very likely be a bottleneck
 - Need communication-avoiding algorithms (Eller, Stathopoulos)
 - Split grid for eigensolvers (Jung et al)
- Critical slowing down
 - For solvers: Multigrid algorithms (Weinberg et al)
 - For HMC: Fourier acceleration? (Christ et al)

Discussions related to FOM this afternoon

- Grid on GPUs (Peter)
- CPS improvements (Chulwoo)
- Chroma improvements (Balint)
- OpenMP (Meifeng)
- QUDA improvements (Kate)
- HISQ Multigrid (Rich)