SPECTRE, Phi-s and JITs

Bálint Joó[1]
[1]Jefferson Lab, Newport News, VA, USA

NP SciDAC Meeting, April 2014

Who thought up this title?

- SPECTRE? Really? We are not that evil…
- Dr Joó != Dr No
- Although cats are OK I guess…

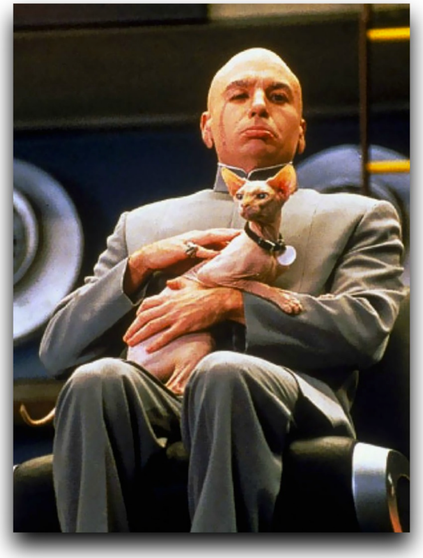


Ladies and Gentlemen… to business…
Miaow!!

Steps of a Custom Calculation

- **Identify Graphs needed**
- **Gauge Generation** (with Chroma)
  - strong scaling, little I/O
- **Compute Propagators** (with Chroma)
  - throughput oriented
  - I/O intensive: solution vectors are large
- **Generate Perambulators** (with Három)
  - I/O intensive: read solution vectors/sources
- **Contract Propagators** (with Redstar)
  - Dense Matrix Multiplication Heavy
  - Scope for sub-expression reuse

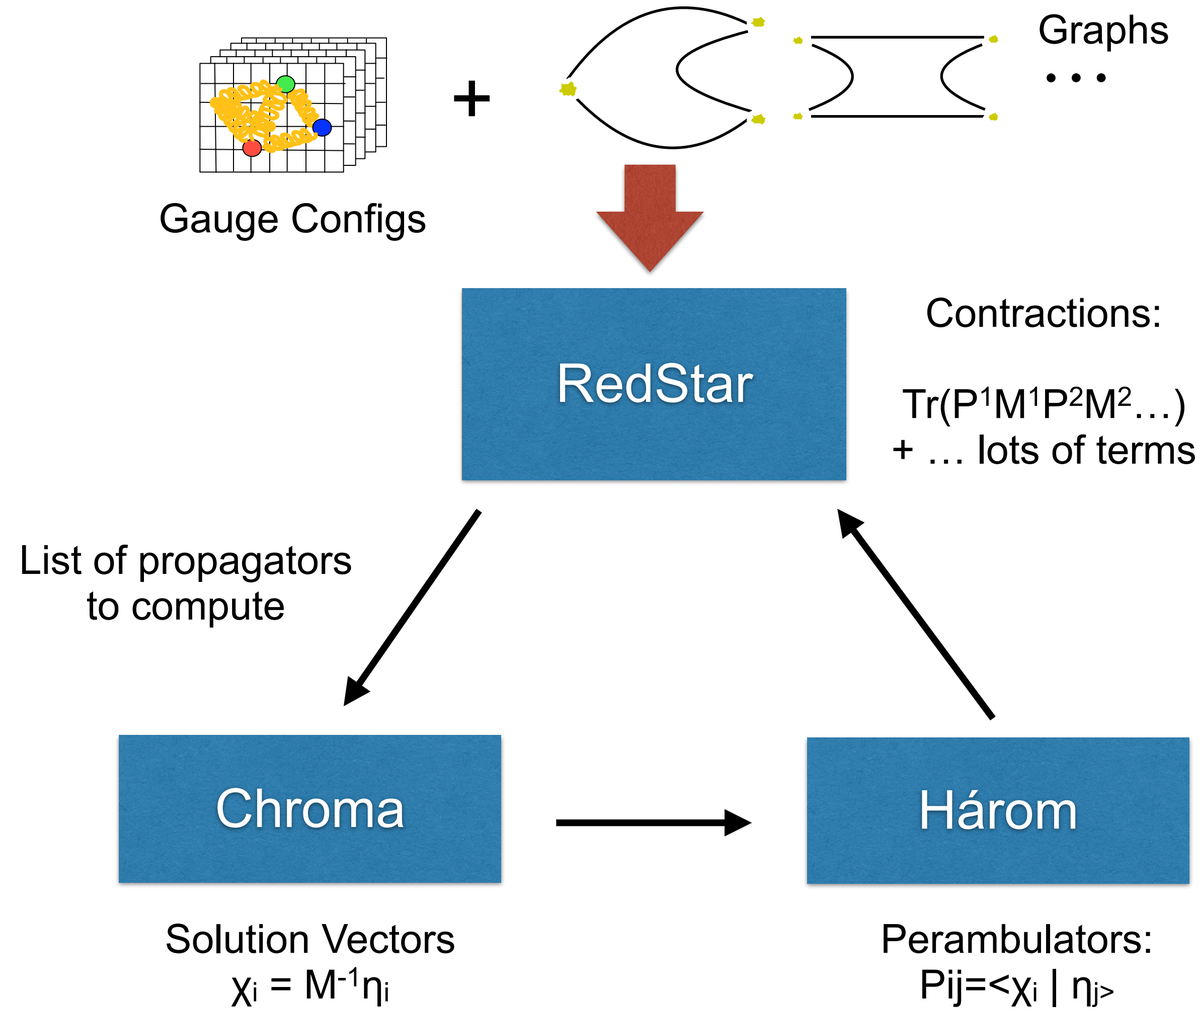

Gauge Generation

- Current Status
  - Anisotropic production done 58M INCITE c/h used
  - Current focus is on isotropic, 64^3x128, 72^3x256 lattices, thermalizing and dialing down light quark mass.
- Code Optimization (detail from Frank, next talk)
  - GPU: Chroma over QDP-JIT + QUDA solvers
    - In production currently.
  - BlueGene/Q: Chroma over QDP-JIT + BAGEL solvers
- Algorithms
  - GPU: RHMC for strange quark is expensive
    - Multi-shift solver scales poorly
  - Light Quarks at Physical Mass
    - Would like to reduce the cost of solves (Multigrid?)

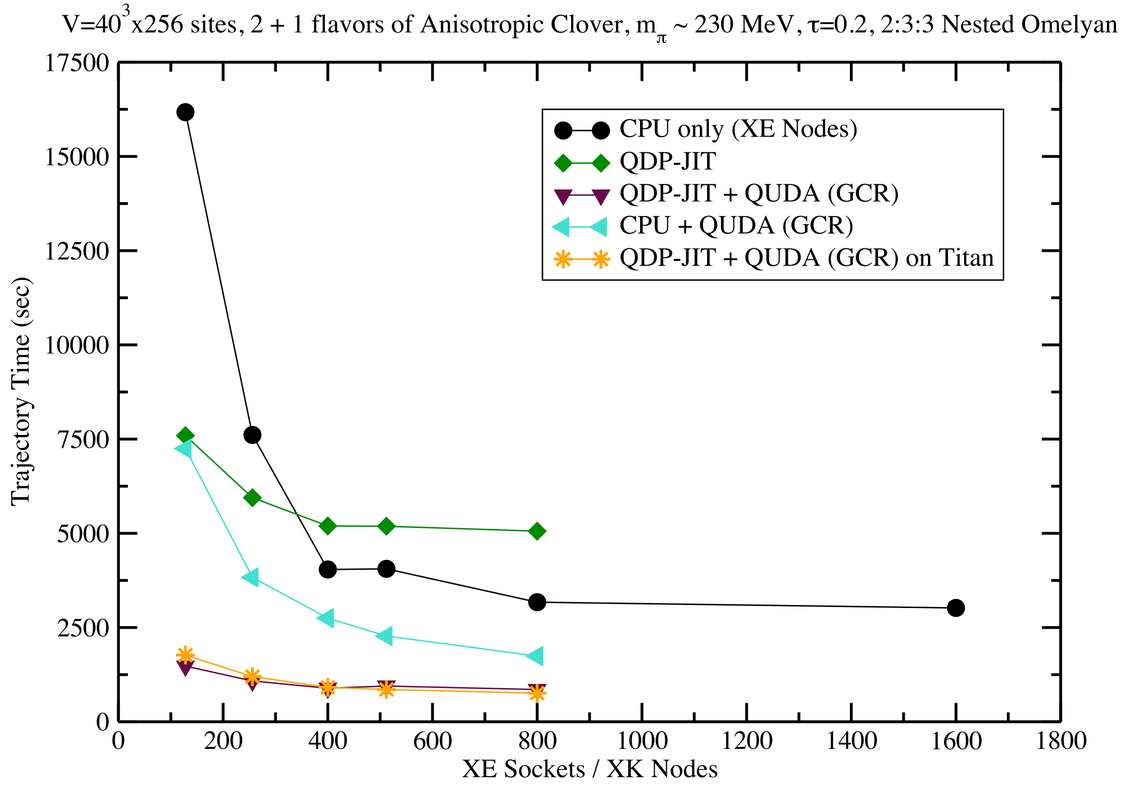

F. Winter, M. A. Clark, R. G. Edwards, B. Joo
to be published at IPDPS'14

Gauge Generation: Time spent in force

- Iso running in infancy, still tuning Monomial structure
- Data here: Aniso Running (400 nodes) m_π~230 MeV
- Terms with Two Flavor Solves
  - 2 Hasenbusch Terms + Cancellation: ~64% of time
- Single Heavy Strange Quark
  - Multi-shift CG, 34% of the time
- Remaining Terms: about 2% of the time.
- Strange quark is expensive
  - and strong scales poorly on GPUs since no DD
  - in iso running DD-Solver ~ 110TF, Multi-shift: ~17TF on 512 XK7 nodes of BlueWaters

40x40x40x256 lattice

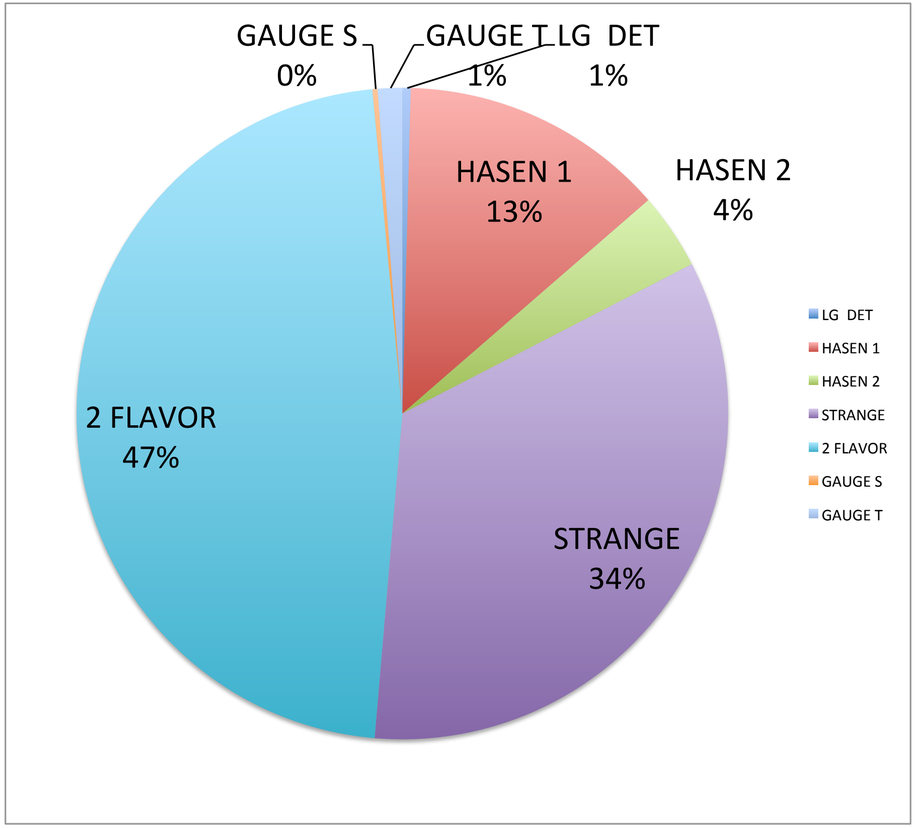

DD vs. non-DD Solvers on Titan

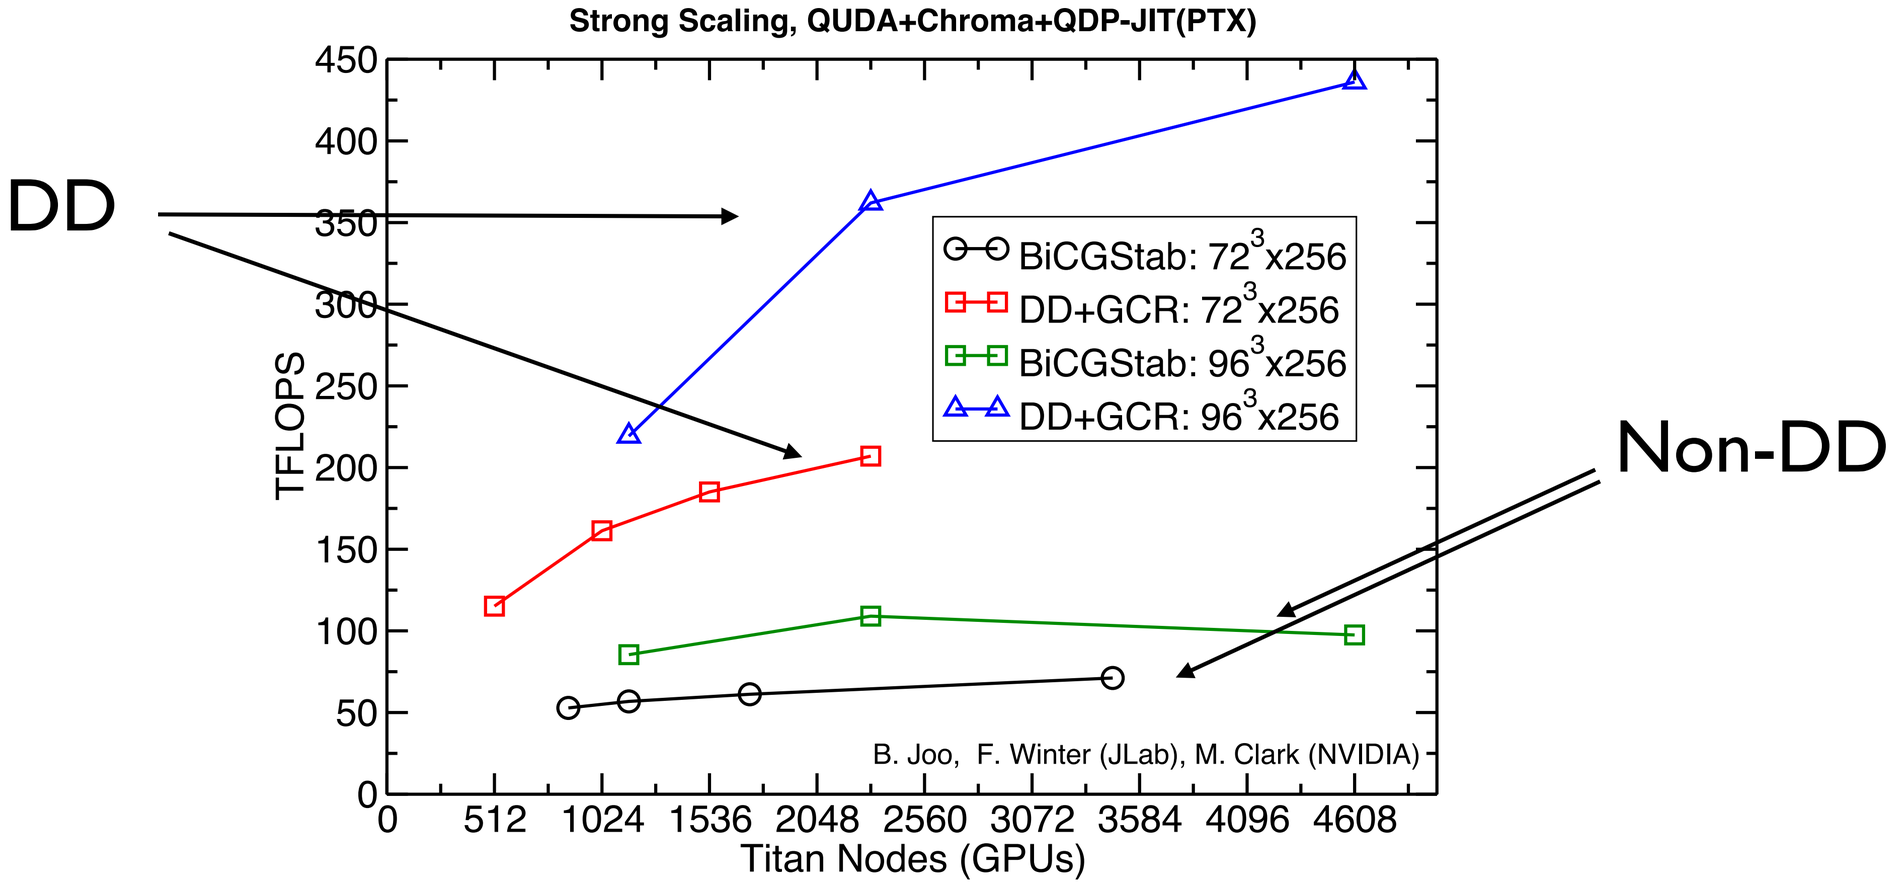

Propagator Computation

- Multi-Grid based solver from QOPQDP (J. Osborn)
- integration by S.D. Cohen with a little help from B. Joo
  - 10-11x speed up on 32^3x256 anisotropic lattices on the CPU at $m_\pi \sim 230$MeV
  - Saul reported 30x speed up at Physical quark mass
    - MILC lattices
  - Currently faster than BiCGStab on GPU
  - Multi-Grid based solvers on GPU/XeonPhi highly desirable

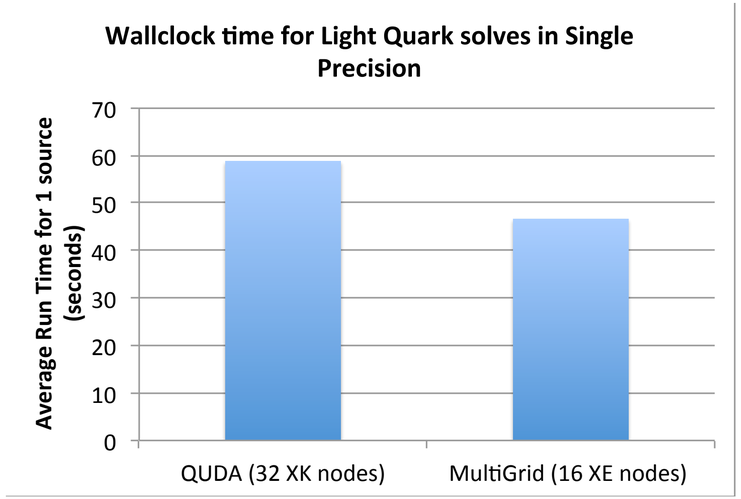

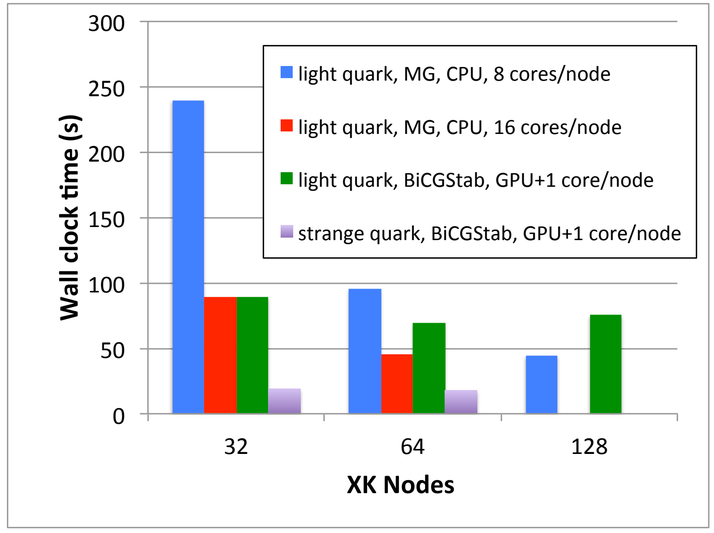

Xeon Phi Developments

- Library now supports
  - single and double precision
  - half precision: up/downcast to/from SP on load/store
  - BiCGStab as well as CG for Clover
  - Simple iterative refinement multi-precision solver
  - 4D communications
  - Xeon Phi & AVX available with AVX2 in the works
- Process to open source
  - Intel is enthusiastic about open sourcing code
  - library, code-generator & CML proxy
  - in process, currently waiting on DOE approvals

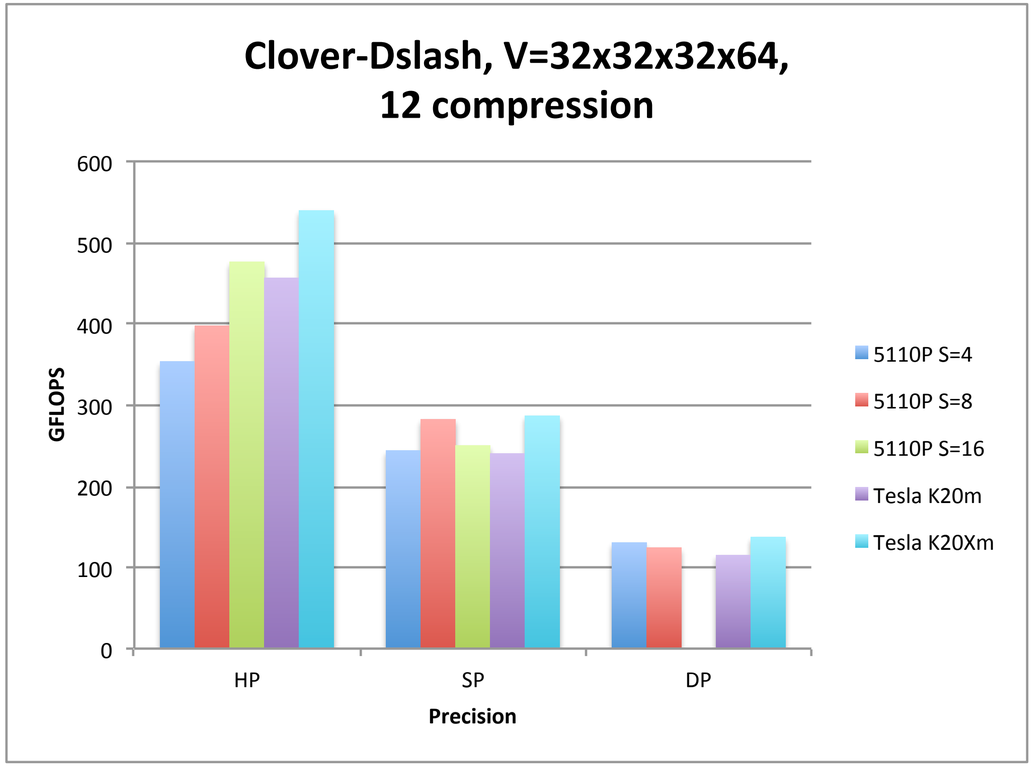

DP=double precision
SP=single precision
HP=half precision

- Xeon Phi: single precision, with 16bit up/down conversion

QCD with Xeon Phi on Stampede

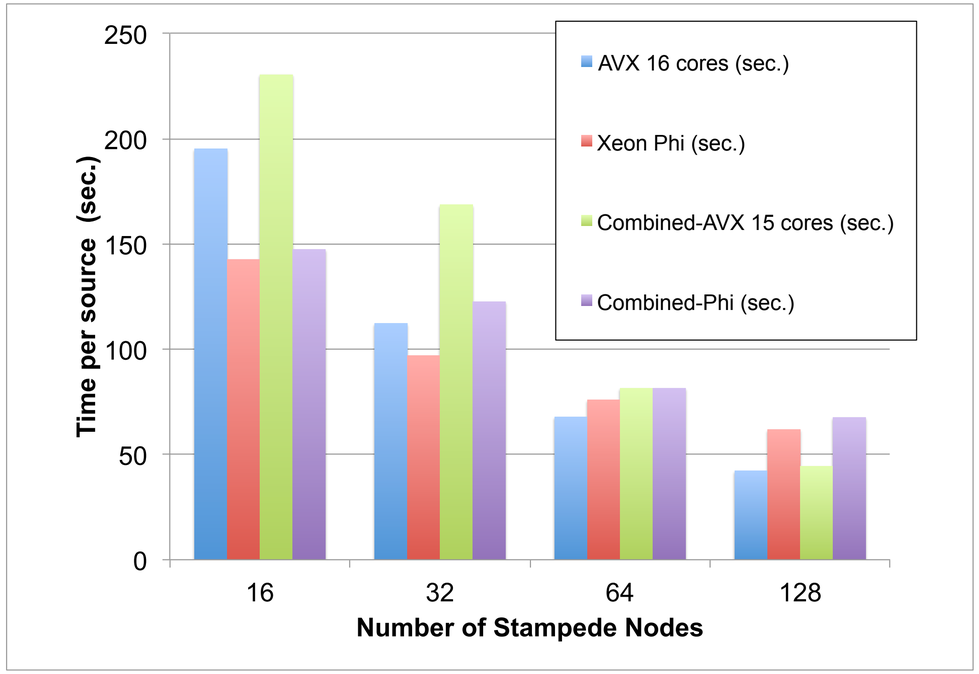

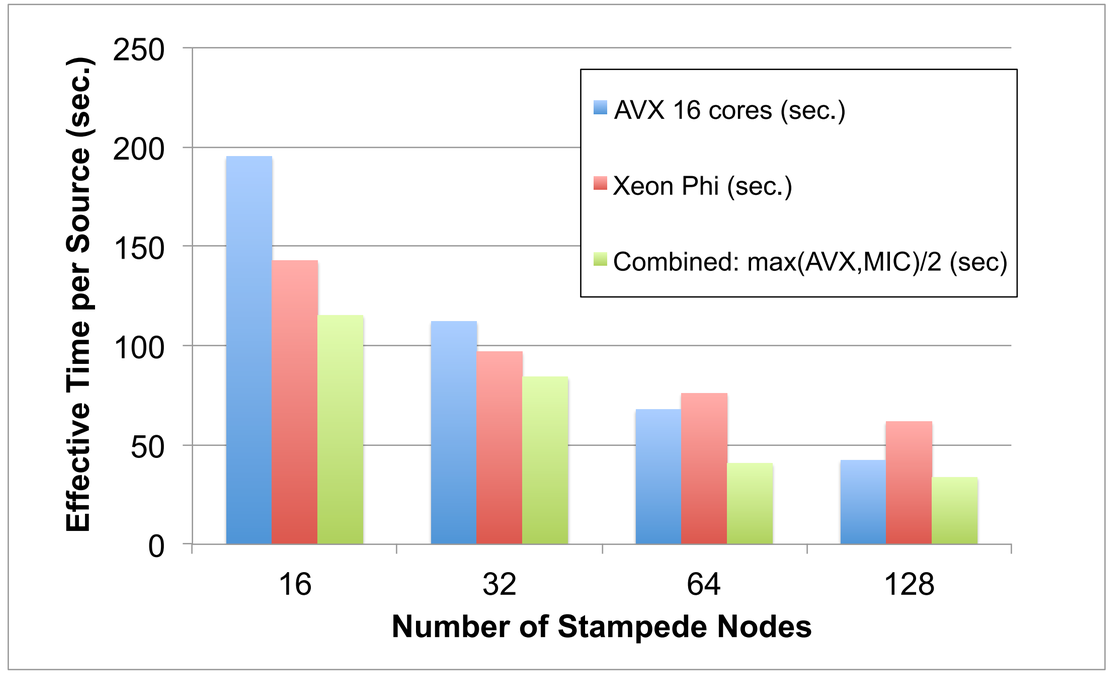

- Performance on Stampede of older SP code, using both Xeon Phi and Xeon
- Using both the Xeons and Xeon Phi results in throughput gain although not full factor of 2

Other Xeon Phi Work

- DD + Flexible GMRES solver developed at University of Regensburg
  - Simon Heybrock & Tilo Wettig
- Collaboration
  - Misha Smelyanskiy, Pradeep Dubey, Intel Parallel Labs Santa Clara
  - Dhiraj Kalamkar, Karthikeyan Vaidyanathan, Bharat Kaul, Intel Parallel Labs, Bangalore
  - B. Joo, Jefferson Lab
- SC'14 paper submission is imminent (submission deadline is Friday)
- Willingness to share the code
  - source code sharing is linked to open sourcing the code generator

Analysis

- Redstar - Chroma - Három combination is in production
  - JLab Clusters, BlueWaters, Stampede
- Contractions challenging
  - all dense matrix multiplies
    - 32^3x256 lattice: 384x384 dense matrices
    - 40^3x256 lattice: 768x768 dense matrices
  - could benefit from acceleration through libraries (CUBLAS, MKL?)
  - Robert is looking at the TCE and considering on how to optimize the contractions
  - algorithmic considerations (stochastic etc?)

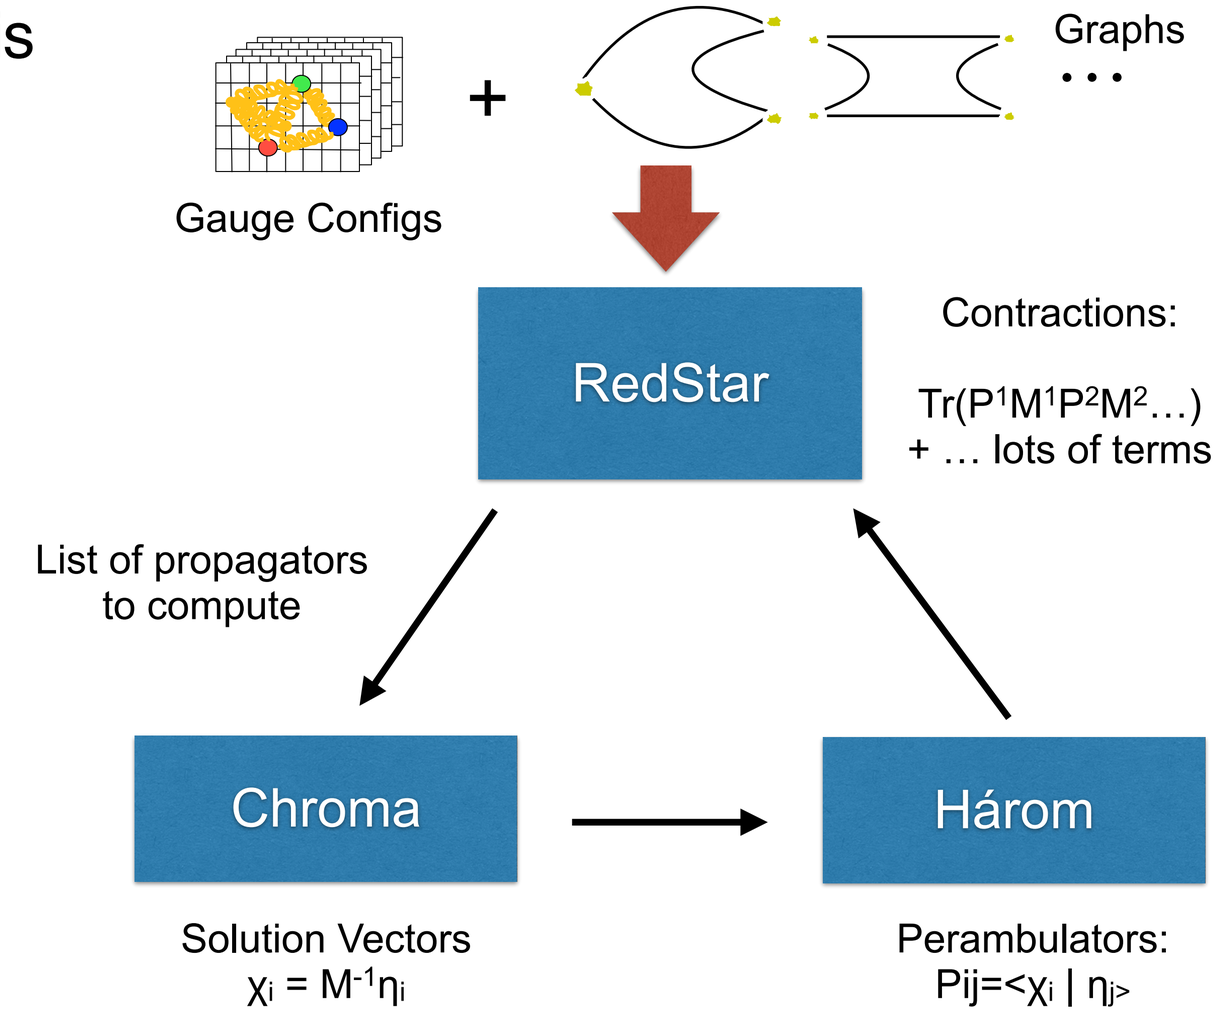

Task Summary

- BlueGene/Q:
  - QDP-JIT + BAGEL Solver integration (Frank)
  - Whole application performance analysis (with SUPER)
- Gauge generation algorithm
  - MG / deflation / condition number awareness in HMC, would like faster strange quarks. They are heavy. The should be cheap.
- Xeon Phi
  - Library cleanup & Open sourcing (Balint)
  - Chroma (re)-integration (Balint)
  - MG (Balint?)
- Propagators
  - MG for GPUs (Mike?)
  - Chroma integration (Balint/Frank with help from Mike?)
- Analysis
  - Acceleration of contractions via libraries (Robert + Jie? + ??)
  - Subexpression elimination and contraction optimization (Robert)
  - Stochastic method exploration, baryons etc (Robert)