High Energy Physics Division, ANL

Lattice QCD in extreme environments

D. K. Sinclair (HEP, Argonne)

J. B. Kogut (Physics, Illinois)

D. Toublan (Physics, Illinois)

Lattice QCD

Quantum chromodynamics(QCD) describes Hadrons and their strong interactions. Hadrons consist of quarks held together by gluons.

Lattice QCD is QCD on a 4-dimensional (space-time) lattice. Allows numerical simulation of the functional integrals which define this quantum field theory, and non-perturbative QCD calculations.

Physics — properties of hadrons (masses, etc.), hadronic matrix elements (HEP), hadronic matter at finite temperature and/or densities (RHIC, early universe, neutron stars).

Computational Methods

- Functional integral is mapped to the partition function for a classical system. Molecular-dynamics methods are used to calculate the observables for this classical system.
- The equations of motion for the fields defined on the sites and links of the lattice are integrated numerically using a modified Verlet(leapfrog) method.
- At each update we need to invert the sparse Dirac operator (matrix) operating on the vector of quark fields. We use the conjugate gradient algorithm to solve this system of (typi-

cally $\sim 10^4$—10^7) linear equations.

- On MPP systems we parallelize by dividing the lattice into N pieces, each having the same number of contiguous sites and the same shape. In the MPI implementation 1 piece is given to each of N MPI tasks. This leads to 100% load balancing. The communications are homogeneous send-receive operations between nearest neighbour tasks on the lattice. In addition there are 2 global sums per conjugate gradient iteration.
- Communications can be overlapped with computations when the machine supports this.
- Our task codes are in fortran.

Benchmarks of Production Codes

These are benchmarks for production codes on Seaborg, for lattice sizes used in production running. Note that here we maintain a fixed lattice size while we vary the number of processors used, so that as the number of processors increases, the size of the sublattice on each node decreases.

This contrasts with what has been presented in the past by other groups who have presented benchmarks where they assign a fixed size (and shape) sublattice to each processor, so that the size of the whole lattice increases linearly with the number of processors.

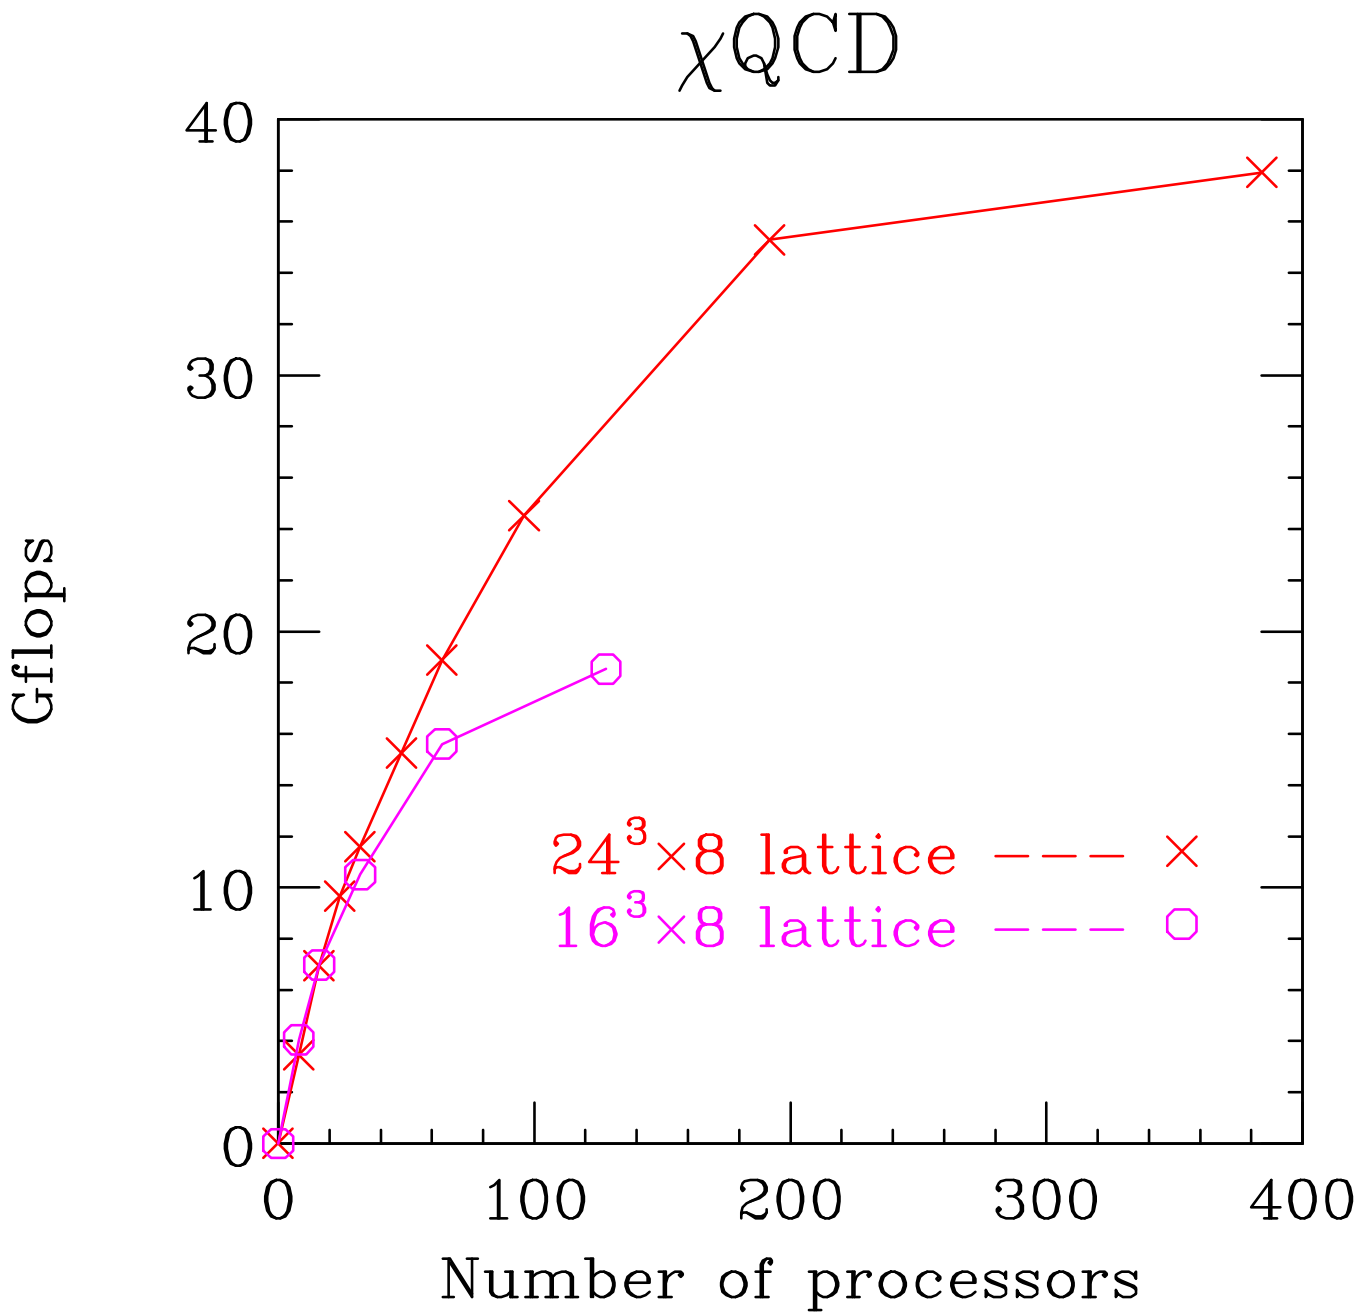

Figure 1: Performance of $N_t = 8$ QCD + 4-fermion code on Seaborg

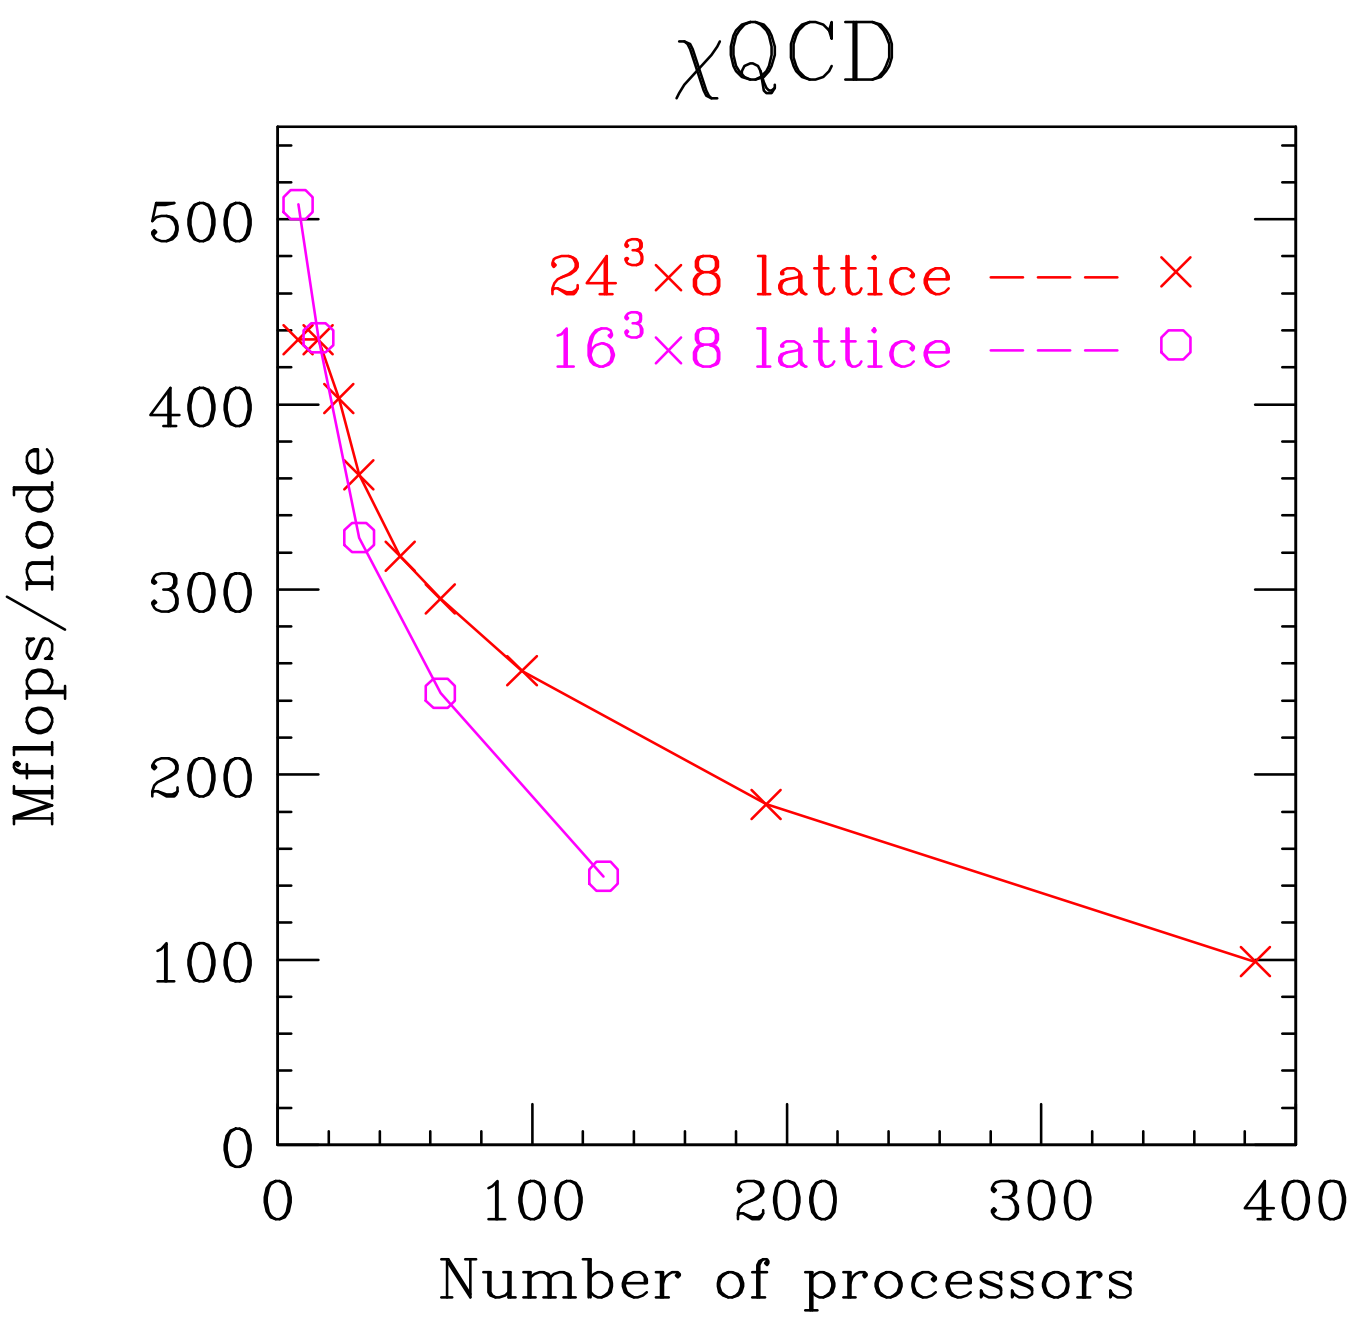

Figure 2: Per processor performance of $N_t = 8$ QCD + 4-fermion code on Seaborg

Lattice QCD with irrelevant chiral 4-fermion interactions at finite temperature

One difficulty with the standard formulations of lattice QCD is that the Dirac operator becomes singular at zero quark mass.

If a is the lattice spacing, the discretization error is $O(a^2)$. We add to the standard action a 4-fermion interaction which preserves its symmetries, contributes $O(a^2)$, and renders the Dirac operator non-singular at zero quark mass.

We are currently using this action to simulate 2-flavour lattice QCD with massless quarks at finite temperatures

to determine the nature of the transition from hadronic matter to a quark gluon plasma, and measure the critical exponents at this transition.

The measurement of hadronic properties (masses, decay rates...) on the lattice with standard formulations is limited by the small masses of the u and d quarks. Simulations are typically performed with unphysically large quark masses and the results extrapolated to the physical quark masses. We plan to apply our methods which are unaffected by the smallness of quark masses to such measurements, using the physical u and d quark masses.

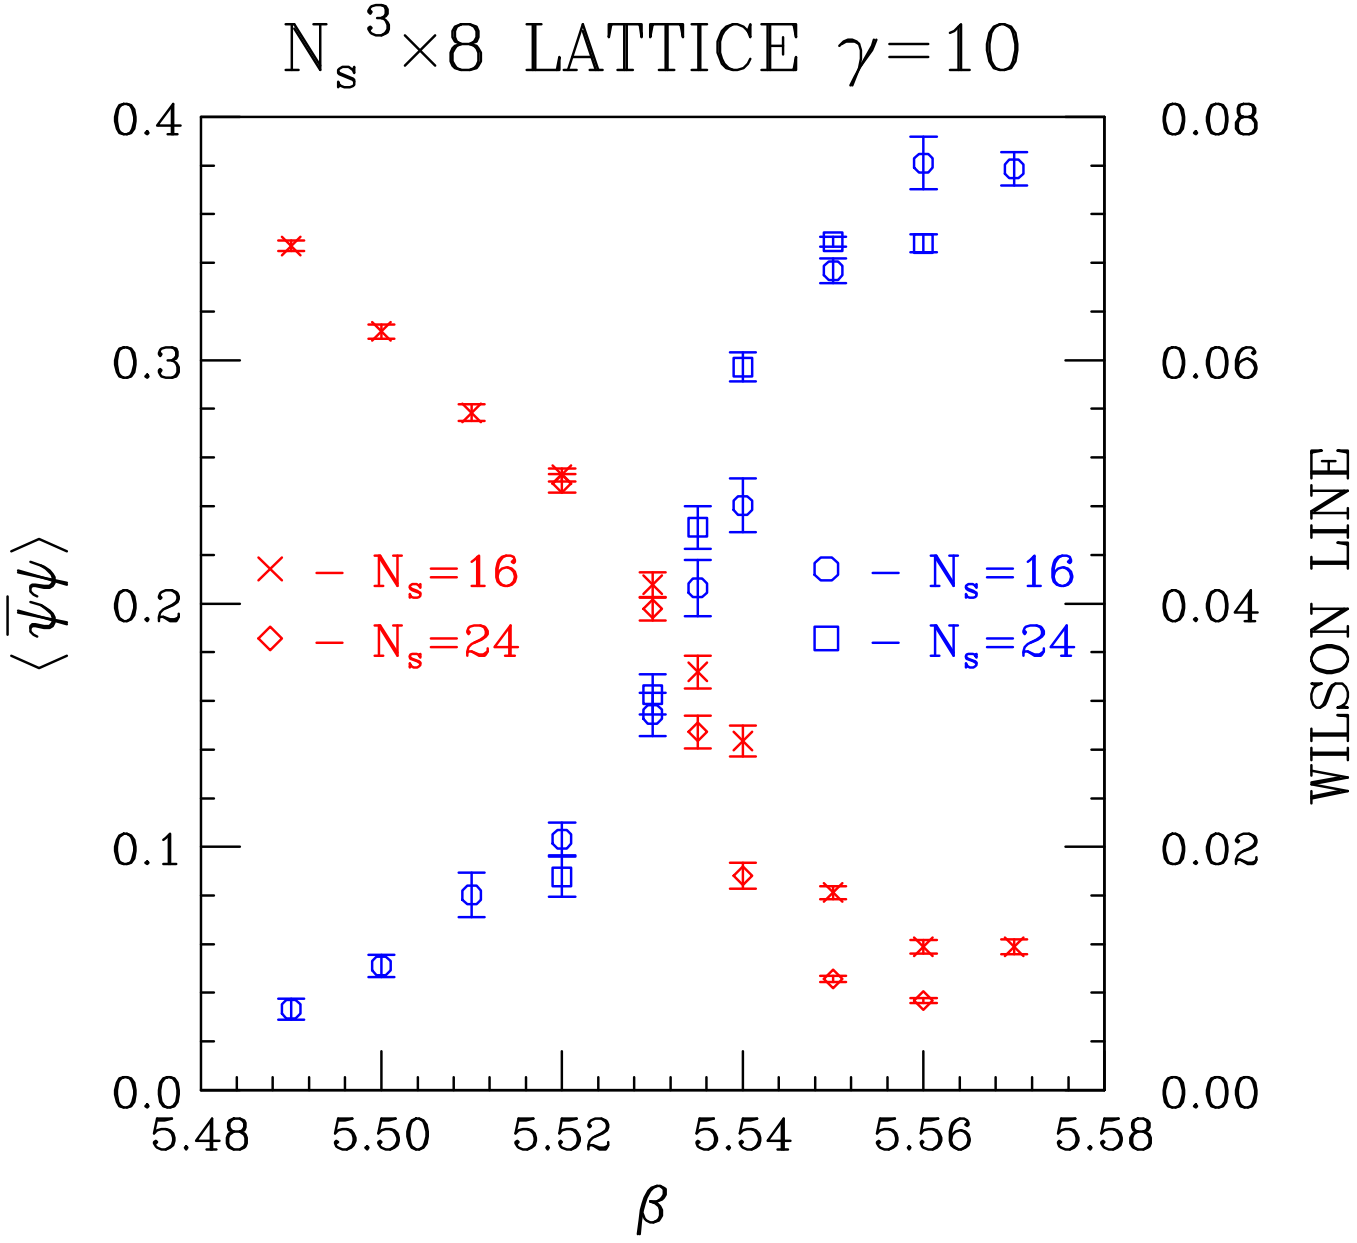

Figure 3: The chiral order parameter $\langle \bar{\psi}\psi \rangle$, and the Wilson Line (Polyakov Loop) as functions of $\beta = 6/g^2$.

Lattice QCD at finite densities

- QCD at finite baryon/quark-number, isospin(I_3) and/or strangeness density (nuclear matter).
- QCD at finite temperature and density.
- QCD at finite baryon-number chemical potential, has a complex fermion determinant. Standard simulation methods based on importance sampling fail.
- We study QCD at finite isospin density, and 2-color QCD at finite quark-number density, which have real positive fermion determinants so stan-

dard methods work.

- Lattice QCD with a chemical potential μ_I for isospin(I_3) has a mean-field second order phase transition at $\mu_I = m_\pi$ to a superfluid state with a charged pion condensate and a massless Goldstone pion.
- We measure the μ_I dependence of the finite temperature transition at small μ_I, where the μ_I and μ_q dependence should be identical. This gives us access to RHIC physics. At large μ_I and high temperature we observe the pion condensate evaporate at a first order transition.

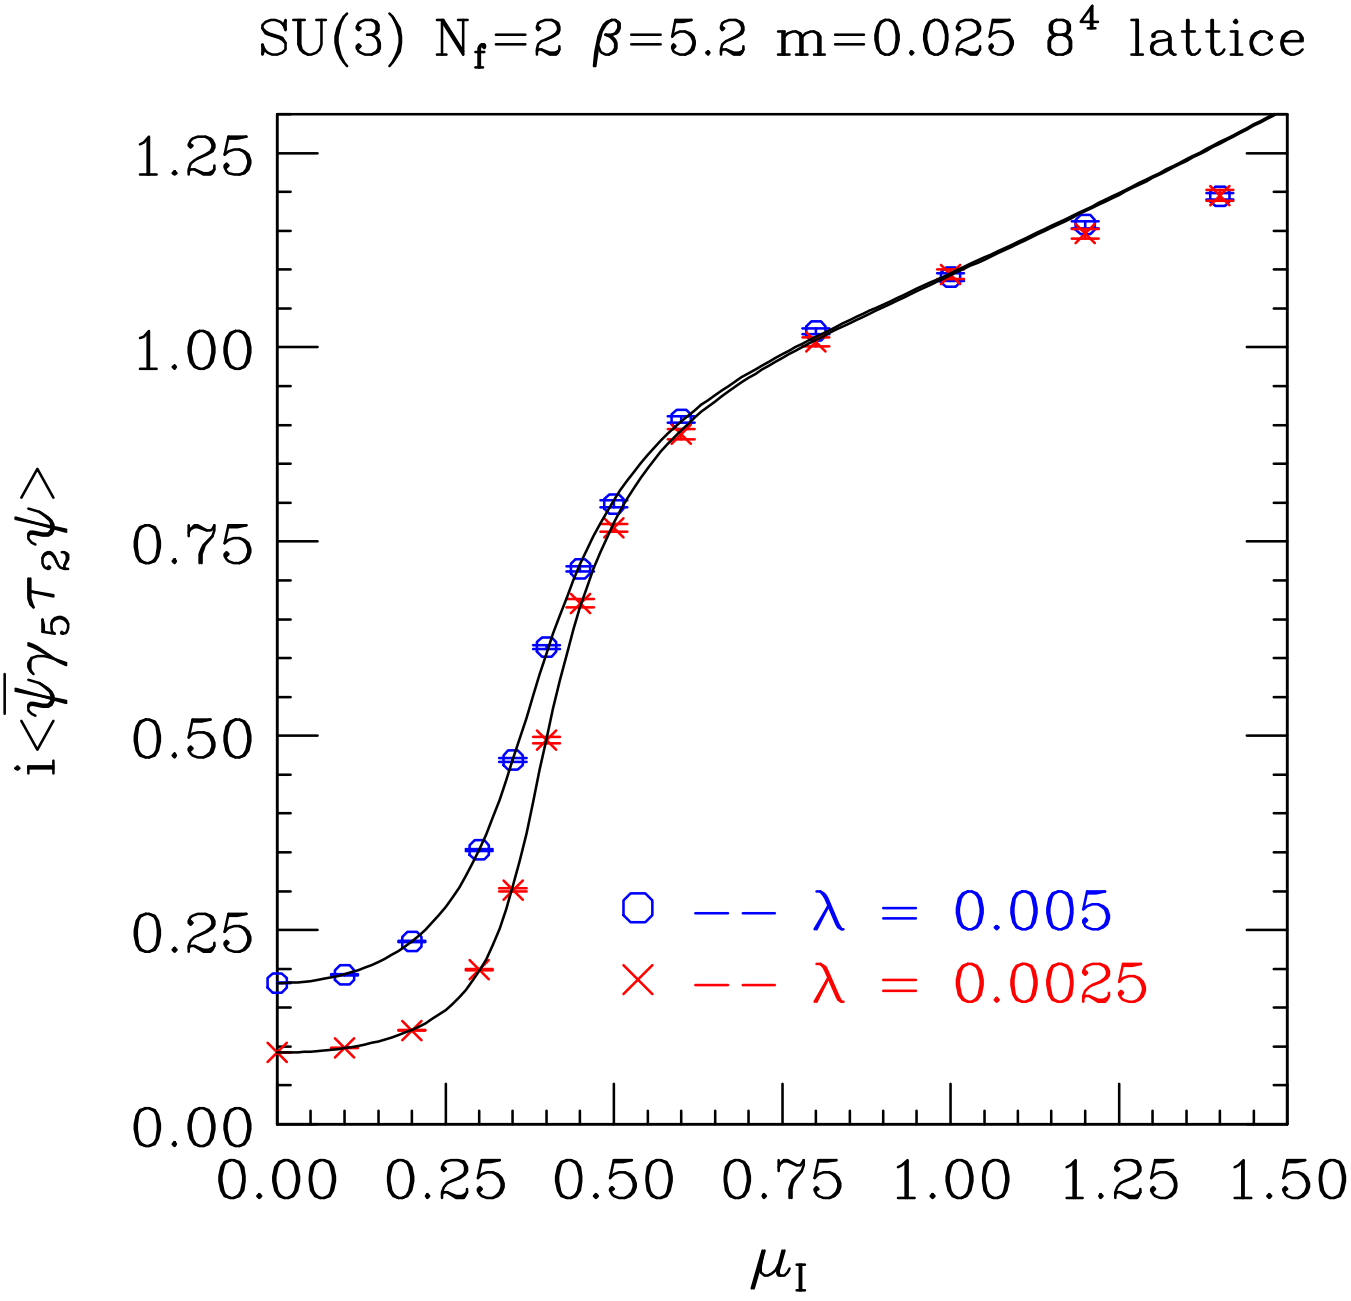

Figure 4: The charged pion condensate as a function of μ_I.

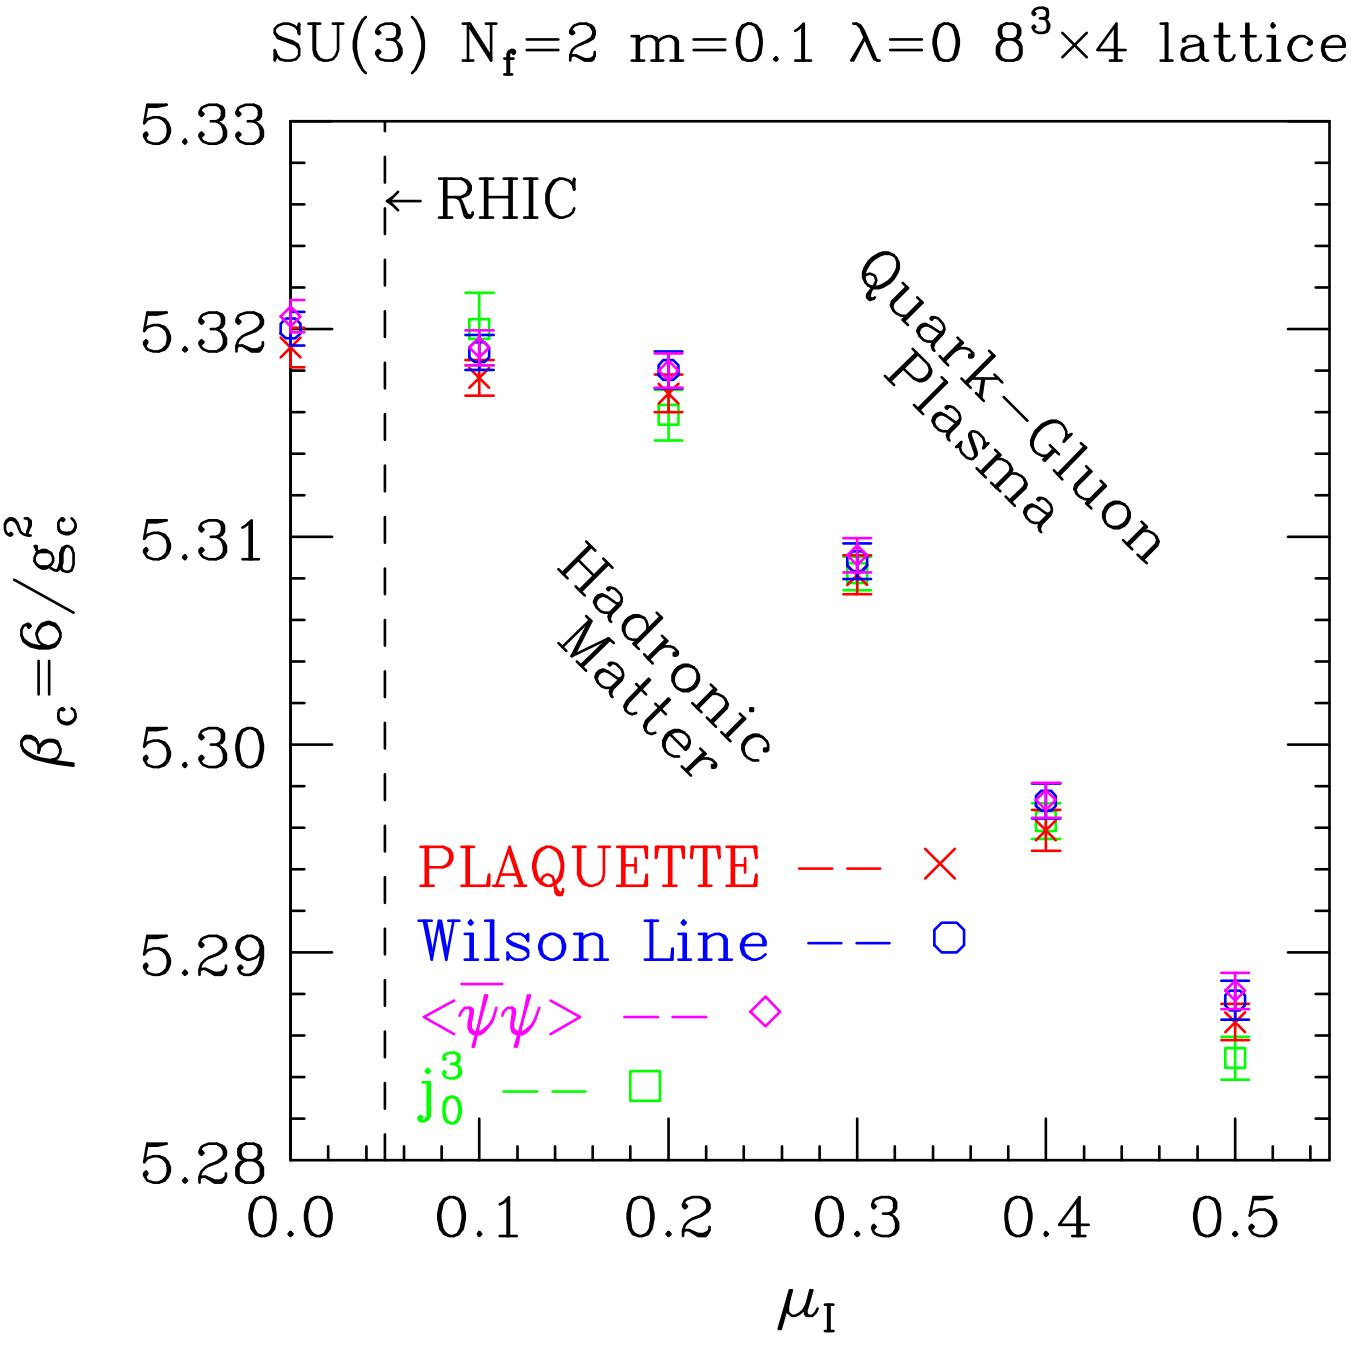

Figure 5: The transition $\beta = 6/g^2$ as a function of μ_I.

Computing Requirements and Observations.

- To make more precise predictions of the properties of hadrons and of hadronic matter in extreme environments, requires larger lattices, smaller lattice spacings, quark masses closer to the physical u and d quark masses and higher statistics.
- It has been estimated that the next stage of Lattice QCD calculations will require computers with sustained performance in the multi-teraflop range. To perform the ultimate Lattice QCD simulations will probably need petaflop computing.

- Seaborg does not readily accommodate fine grain parallelism. This restricts the number of processors which can be used effectively on small lattices. Optimally, if we use the MPI paradigm, we would need a computer which can efficiently handle codes where each task performs only $\sim$ 100–1000 floating point operations between MPI calls.
- Although we have managed to parallelize serial codes to make use of the SMP capabilities of the individual nodes, we have had little success with hybrid execution of our MPI codes.

- Global reduction operations appear to be implemented relatively inefficiently on the IBM SP's.
- One positive aspect of our codes is that they spend most of their time in a single subroutine — the conjugate gradient inversion routine (typically 95–100% in production), which in its simplest serial form is only $\sim$ 100 fortran statements! This simplifies optimization.