Chasing the Unicorn: RHIC and the QGP

Robert D. Pisarski

Nuclear Theory & High Energy Theory Nuclear Theory Groups,
Brookhaven National Lab., Upton, NY, 11973, U.S.A.;
Niels Bohr Institute, Blegdamsvej 17,
2100 Copenhagen, Denmark;
Frankfurt Institute for Advanced Study,
J. W. Goethe Univ., Robert Meyer Str. 10,
D-60054 Frankfurt, Germany

At nonzero temperature, it is expected that QCD undergoes a phase transition to a deconfined, chirally symmetric phase, the Quark-Gluon Plasma (QGP). I review what we expect theoretically about this possible transition, and what we have learned from heavy ion experiments at RHIC. I argue that while there are unambiguous signals for qualitatively new behavior at RHIC, versus experiments at lower energies, that in detail, no simple theoretical model can explain all salient features of the data.

The phase transitions of a nonabelian gauge theory are of intrinsic interest in their own right. It is reasonable to expect that there will be qualitatively new behavior, not seen in the phase transitions of spin systems, nor in the plasma physics of abelian gauge theories.

In this talk, I summarize what we expect about the phase transitions of the theory of strong interactions, Quantum ChromoDynamics (QCD), at nonzero temperature. I then summarize what we have learned so far from experiments, concentrating especially on the results obtained from the Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory (BNL) since its commission in 2000. We now have an absolute wealth of data from RHIC; the outpouring of data has been summarized in a series of "white papers", from each of the four major experiments [1]. A version of this talk is available on line [2]; in this Proceeding I try to emphasize a qualitative discussion of the physics, and leave the number of plots as limited as possible, stressing the most important, qualitative, aspects of the physics.

My perspective in the present talk will be that of an interested but skeptical observer. At the outset, I should confess that my perspective is that of a field theorist, who does not do detailed fits to the data. Thus much of my criticism can simply be dismissed as the ravings of someone who talks, but who doesn't do. There is some validity to this comment.

However, I will try to stress that when considered in total, that there are many things going on in the RHIC data which is much more interesting than in any single model. The difficulty is that while there are many models which explain particular features of the data, that there is no single model which can encompass *all* interesting features of the data. By this, I do not mean simply that calculations are off by a small amount. One of the most impressive features of the RHIC data is that the experiments have obtained a large quantity of *precise* data. The quanities which they measure are, with rare exception which will be noted, measured to within a few percent, and agree between the different experiments. The problem is that not that

we do not have data at to what is going on; the problem is that we do not know what to make of everything, all together.

I do think that it is clear that something, qualitatively new, has happened by RHIC energies, as compared to lower energies, such as at the SPS. In part, this is because phenomenon such as jets only open up at RHIC energies, and are not present at the SPS. This strongly suggests that a transition to a new kind of matter has occured by RHIC energies. Note that I do *not* say that there is a transition to a new "state" of matter, in the thermodynamic sense. What I mean is that the matter behaves differently from that at lower energies. Whether it is, in fact, hadronic matter in thermal equilibrium, above the transition to deconfinement, is still to be established unequivocally. In particular, it is well possible that the transition has occured at lower energies, such as at the SPS (or perhaps even lower), but that it was too difficult — at present! — to disentangle there, at least with the previous experimental probes.

If nothing else, I wish to emphasize the following. Before the RHIC experiments, those in other fields might well have assumed that while one could do heavy ion experiments at RHIC energies, that the systems would simply rip through each other, interacting in some manner, but with no especially notable traces of the interaction.

This is *NOT* true. At least at RHIC energies, the interactions between the nuclei are extremely strong; while the nuclei rip through each other, they leave strong traces of their interaction behind. Further, these traces are strongly suggestive that a thermal system has, in fact, been formed. If it is not a Quark-Gluon Plasma, it is behaving very much like one. What is extraordinary is that the QGP appears to be a different beast from what we had expected.

This brings me to the central image of my talk. In Western literature, the Unicorn is a familiar image of a fantastic and mythical beast. The Unicorn was first used by David Scott, Reinhardt Stock, and Miklos Gyulassy [3] as a metaphor for the QGP. Following medieval tradition, where people widely borrowed from one another, I also use this metaphor. (Recog-

nizing, of course, that skeptics will view any evidence for the QGP as similar to that for sightings of a Unicorn.) Thus my conclusion is that while a Unicorn *has* been found, that it isn't the Unicorn we expected, but is even more subtle and marvelous than we ever had reason to expect. In the end, it is an absolute triumph of Experimental Physics to have discovered whatever beast they seem to have discovered.

I. THEORETICAL EXPECTATIONS

A. Deconfinement

There is a rigorous order parameter for deconfinement, first discussed by G. 't Hooft, A. Polyakov, and L. Susskind. It is well known that in an SU(3) gauge theory, that most observed states are either mesons, formed from a quark and an anti-quark, or baryons, formed from three quarks. Both states are also composed of some interterminate number of gluons, which are, in the end, responsible for confining the quarks and anti-quarks into just these states.

This allows us to describe confinement as a type of three state model. Suppose that we have a clock, where the hands can only point at three places, at 12:00, 4:00, and 8:00. Each time we add a quark, it rotates the hands of the clock forward by 120 degrees; adding a anti-quark rotates the hands in the opposite direction, by minus 120 degrees. We then consider which states are invariant under these transformations. Obviously, they are just mesons and baryons. Mesons correspond to rotation by + 120 - 120 = 0 degrees, and so are invariant. Similarly, baryons correspond to a rotation by 3 x 120 = 360 degrees.

This system of rotations forms an abelian group, which is that of Z(3). Z(3) is a cyclic group, because a rotation by three quarks is equivalent to the identity. Confinement is then the statement that at zero temperature, this Z(3) symmetry is an *un*broken symmetry of the gluonic vacuum.

Such a Z(3) symmetry is almost the simplest possibility one can have for a global symmetry group. For example, if the gauge interactions involved two colors, instead of three, the corresponding symmetry group for confinement would be Z(2), as in the Ising model. In condensed matter, the Z(3) model is known as the three states Potts model.

There is a rigorous order parameter for this Z(3) symmetry. Consider the propagation of an infinitely massive test quark. If one puts the quark down at some point, it will just propagate forward in time, without moving in space. Even so, it can still exchange color, although not momentum (since it is infinitely heavy) with the vacuum. To measure this, we introduce the "Wilson line",

$$\mathbf{L}(\vec{x}) = \mathcal{P} \exp\left(ig \int_0^{1/T} A_0(\tau, \vec{x}) d\tau \right) . \tag{1}$$

I have immediately gone to compute the properties in thermal equilibrium at a temperature T, using the imaginary time formalism. (This is a fancy way of saying that to compute properties in thermal equilibrium, one weights configurations by $\exp(-\mathcal{H}/T)$, where $\mathcal{H}$ is the Hamiltonian. This looks as if there is propagation in a "time" τ which runs from 0 to $1/T$. Normally, propagation is by a factor $\exp(-i\mathcal{H}t)$, where t is the time. Because there is no factor of i, the time τ is imaginary.) Otherwise, the Wilson line is just like a Aharonov-Bohn phase factor: the color rotates with the gauge field, interacting like a path ordered exponential. Thus g is the gauge coupling constant, A_0 is the vector potential for the gauge field in the time direction, and $\mathcal{P}$ denotes path ordering, required because the gauge field is a nonabelian matrix, and one has to specify how to order matrices at different times.

The Wilson line is a color matrix, and so is not gauge invariant. A gauge invariant quantity is formed by taking the color trace of this matrix, which yields the Polyakov loop,

$$\ell(\vec{x}) = \frac{1}{3} \operatorname{tr} \mathbf{L}(\vec{x}) . \tag{2}$$

It is easy to guess how the Polyakov loop behaves at high temperature: by asymptotic freedom, any gauge theory becomes ideal in the limit as $T \to \infty$. That is, the coupling constant vanishes logarithmically, $g^2(T) \sim 1/\log(T)$. Then we can neglect g in the exponential, and the matrix is equal to the unit matrix. Thus at high temperature, the expectation value of the Polyakov loop should approach unity, $\langle \ell \rangle \to 1$.

How the Polyakov loop behaves at low temperatures is less obvious. The insight of 't Hooft was to show that under the Z(3) transformations discussed above, the Polyakov loop transforms by an overall phase factor, which is the same at every point in space,

$$\ell(\vec{x}) \to \exp^{2\pi i j/3} \ell(\vec{x}) \quad , \quad j = 0, 1, 2. \tag{3}$$

The mathematical realization of the transformations of a clock, described above, is multiplication by one of the third roots of unity.

Thus we were somewhat sloppy at very high temperatures: there the Polyakov loop can not only approach one, but it can also approach $\exp(2\pi i/3)$, or $\exp(-2\pi i/3)$. All of these states are equally good vacua at infinitely high temperature. In other words, there are three degenerate, and equivalent, vacua at high temperature in an SU(3) gauge theory. In a spin system, this corresponds to the spontaneous breaking of a global symmetry in the vacuum.

In contrast, at zero temperature I argued that confinement corresponds to an unbroken phase of Z(3) spins. The best way to envision this is to use the spin analogy. When spins "break" a symmetry, they do so because they tend to line up, as in a ferromagnet. When the symmetry is restored, it is because the spins tend to fluctuate. Thus while the spins might tend to be parallel to each other over short distances, due to thermal fluctuations, as one goes to longer and longer distances, this tendency to short range order is washed out. Pictorially, one views the disordered phase as a superposition of many little domains, where the spins are ordered in each domain, but all jumbled up together, so that in infinite volume, the expectation value of the spin vanishes.

Confinement is just like this: the theory is composed of many little domains. In each domain, the Wilson line is near

one, times one of the Z(3) phase factors possible. However, domains have a definite size, so that by going to larger and larger distances, one samples more and more domains. This means that at zero temperature — or in general, at temperatures below that for deconfinement, the expectation value of the Polyakov loop vanishes,

$$\langle \ell \rangle = 0 \quad , \quad T < T_d \ . \tag{4}$$

Here T_d is the temperature for deconfinement. Similarly, the expectation value of the Polyakov loop is nonzero above T_d, approaching one in magnitude as $T \to \infty$.

What is peculiar about the "spins" in a gauge theory is that they order *opposite* to those of ordinary spins: they tend to line up at high temperauture, and to wash out at low temperature.

B.   Chiral Symmetry

The above was a bit of a trick: I assumed that each time a quark was added, that one rotates a phase by one of the third roots of unity. This is only possible when one considers quarks as an external probe in the pure glue theory. In QCD, with dynamical quarks, there is no such distinction possible: quarks and anti-quarks can pop out of the vacuum at any time. Thus at best, the global Z(3) symmetry can only be an *approximate* symmetry in QCD. Even so, I will show evidence shortly as to why this is a reasonable approximation.

There is another approximate symmetry which is well known in QCD: this is the chiral symmetry which is responsible for why pions and kaons are the lightest hadrons in QCD. I will not go into the details of chiral symmetry here, since there are probably more familiar to the reader.   For three flavors of quarks, as are required to describe ordinary and strange mesons and baryons, we have a global symmetry of $SU_L(3) \times SU_R(3)$. (There is also a global symmetry of axial $U(1)$; this is a much more subtle symmetry, which is broken quantum mechanically by the axial anomaly, which makes the η' much heavior than the π, K's, and η.) The subscripts L and R refer to left and right handed quarks, and are special to a symmetry of massless fermions.

However, as a global symmetry, the chiral symmetry behaves like a typical spin system. A condensate, $\langle \overline{q}q \rangle \neq 0$, is responsible for the spontaneous breaking of the chiral symmetry in the hadronic vacuum. As hadrons are heated, fluctuations tend to disorder the chiral symmetry, so that it is restored above some temperature T_{ch}.

C.   Transitions of QCD

What does one expect in general about the deconfining and chiral phase transitions? First consider the case where each symmetry is exact.

For deconfinement, one unavoidably expects a first order transition. This is because as a system with three states, that one can form a cubic invariant. The transition is then like that between a liquid and a gas: there the analogous order

parameter is the density, and terms involving density cubed can always arise. By standard mean field analysis, this implies that the transition is of first order, unless the cubic invariant where to vanish (which one doesn't expect, at least on grounds of symmetry).

For chiral symmetry, for three massless flavors one again has a cubic invariant, and so a first order transition. Things change as the number of flavors goes down: for two massless flavors, the transition can be of second order. This is sensitive to the axial $U(1)$ symmetry, however; if that is restored by the time that the chiral transition occurs (or nearly so), then the transition can become of first order.

Analytic arguments are not sufficient to decide the question. To answer that, we *must* appeal to numerical simulations on the Lattice. The idea of the Lattice is simple: one discretizes the theory, with some lattice spacing $= a$. By asymptotic freedom, we can then tune $a \to 0$, and be assured that however we discretize the theory, that the continuum limit, $a = 0$, will be the same. Of course, we do not have infinite computing power at our disposal, and so Lattice gauge theorists spend much time minimizing the effects of discretization at nonzero lattice spacing.

The effects of discretization are extremely serious for the chiral symmetry of massless quarks. As shown by K. Wilson and F. Wegner, it actually it "easy" to discretize a gauge theory, and still maintain a *local* gauge symmetry at each site of the lattice. What is difficult is implementing the *global* chiral symmetry for massless quarks! This suggests that in some fundamental way, we really do not understand chiral symmetry, but for the present purposes, this is simply a difficult technical problem.

On the other hand, this should not obscure one of the great triumphs of Lattice gauge theory: that for the pure gauge theory, one *can* obtain results which can be reliably extrapolated to the continuum limit, with errors of perhaps a few percent. In this way, the Lattice provides an estimate for the deconfining transition temperature of $T_d \approx 300$ MeV, with errors of about 10%. In fact, the errors in this number arise not from the uncertainty as to where the transition occurs, but to the computation of the string tension, which provides the overall mass scale required to change lattice units into MeV.

For the theory with dynamical quarks, *all* evidence must be treated as preliminary. I quote the evidence as of 2004, but stress its tentative nature. This is not merely a matter of ungrounded skepticism: the order of the transition for two flavors is especially sensitive to what happens to the axial $U(1)$ symmetry. This, in turn, is very sensitive to treating the chiral symmetry properly. It is certainly conceivable that as smaller lattice spacings are probed, with better chiral symmetry, that the transition becomes more first order.

With these caveats, the evidence is that there is no phase transition in QCD. For some reason, which we do not understand at a fundmental level, the deconfining and chiral transitions are equal, $T_d \approx T_{ch}$ to the accuracy of all simulations. While there is no true phase transition, the pressure increases *extremely* rapidly within a narrow region in temperature. This allows one to speak of an approximate transition temperature, even if there is no true transition in the thermodynamic sense.
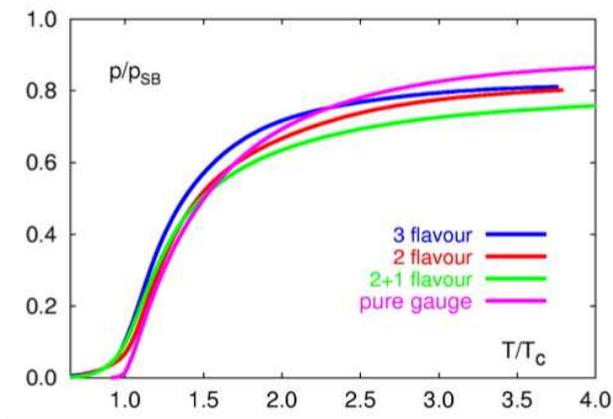
FIG. 1: Pressure divided by ideal pressure

The results are shown in Fig. (1). As first done by the Bielefeld group, it is most interesting to plot the pressure divided by the ideal pressure, versus the temperature, to the "transition" temperature.

One sees that there appears to be a nearly universal curve, with perhaps errors of 20% or less. This has been termed "flavor independence" by the Bielefeld group. This plot obscures the fact that as one goes from the pure glue theory to three massless flavors, that the ideal gas terms increases by about a factor of 3. Similarly, the transition temperature decreases as the number of quark flavors increases (again, for reasons we do not understand), going to about $T_d \approx T_{ch} \approx 175$ MeV for QCD. The errors here cannot be estimated, since they are dominated by the systematics.

Recently, it has also been possible to extract the Polyakov loop from Lattice measurements [4, 5]. The usual quantity measured on the Lattice is a bare loop, which vanishes in the continuum limit. To extract a renormalized loop, which is nonzero in the continuum limit, it is necessary to account for a "mass" renormalization of the loop [4]. Doing so, one obtains Fig. (2)

Remember that in a theory with dynamical quarks, that the Z(3) symmetry is only approximate. Thus it could well be that once quarks are added, that the loop is always large and nonzero for all temperatures. From the results of Petreczky and Petrov [5], this is not what the Lattice finds: instead, the loop with quarks is near that without. Indeed, it strongly suggests that the similarity of (rescaled) pressures, as observed in Fig. (1), is due to the similarity of the values of the Polyakov loop, in Fig. (2). This leads to what is known as the Polyakov loop model [6], where the pressure is assumed to be dominated by the condensate for the loop.

Another surprise is the following. In a perturbative regime, the loop is near one. In the confined phase, it is zero (or small, if there are quarks). The lattice sees a confined phase, below T_d, and a perturbative phase: but only above temperatures of about $\approx 3T_d$! From T_d to $\approx 3T_d$, the value of the loop is nonzero — so that only is manifestly in a deconfined phase of the theory — and yet it is far from unity. This intermediate region can thus be termed a *Non*-perturbative QGP.

This agrees with resummations of perturbation theory, which work down from infinite temperatures by including the effects of the Debye mass [7]. These consistently fail at temperatures on the order of $3T_c$. In particular, they cannot describe the sharp decrease of the pressure.

There is not reason why this should be so. It is perfectly possible that there is a *strong* first order transition from a confined, to a deconfined, phase. As a strongly first order transition, there would need be no close relation between the phases; the free energies only have to match. Instead, the transition appears to be weakly first order, for reasons which are not, at present, will understood.

I also note that this region of the theory, between T_c and something like $3T_c$, has also been described as "sQGP", from a strongly coupled Quark-Gluon Plasma. The difficulty with this is that Lattice measurements of the quark anti-quark potential show that while the coupling gets larger as the temperature decreases — exactly what one expects from the converse to asymptotic freedom — that the increase is relatively modest, by about a factor of two. As I discuss shortly, the experimental data from RHIC, in particular elliptical flow, suggests that the system interacts *very* strongly, close to the limit of ideal hydrodynamics.

While the experimental data on elliptical flow is most suggestive, the Lattice data appears unavoidable. Very large increases of the QCD coupling simply do not appear to happen near T_c. On the other hand, the Lattice data does indicate that while the increase of the coupling constant is relatively modest, that nevertheless, one *does* enter into what one can define as a Non-perturbative QGP.

The author does recognize that NpQGP is not as catchy an acronym as sQGP, and on this ground alone, is destined to fall into the wastebin of history.

II.   EXPERIMENTAL DATA FROM RHIC

In this section I review the experimental data from RHIC. As I stated in the introduction, this will be an *extremely* biased review. Nevertheless, I will try to provide some sort of overview of what one can, and cannot, believe. (Or more precisely, what I do, or do not, believe.)

A.   Multiplicity, average momentum

At high energies, the natural variable is s, the total energy squared in the center of mass system. To compare proton-proton (pp) collisions at a given value of $\sqrt{s}$, to the collisions of two nuclei, each with atomic number A (AA), one uses the energy per nucleon, or $\sqrt{s}/A$. (The experimentalists denote this quantity as $\sqrt{s_{NN}}$, which I find fussy.) As a collider, RHIC at BNL covers energies of $\sqrt{s}/A \approx 20 \rightarrow 200$ GeV. The SPS at CERN is a fixed target machine, and so probes lower energies, $\sqrt{s}/A \approx 5 \rightarrow 17$ GeV. At RHIC, the four major experiments are BRAHMS, PHENIX, PHOBOX, and STAR.

Experimentally, one can distinguish between central collisions, where the nuclei completely overlap, from peripheral
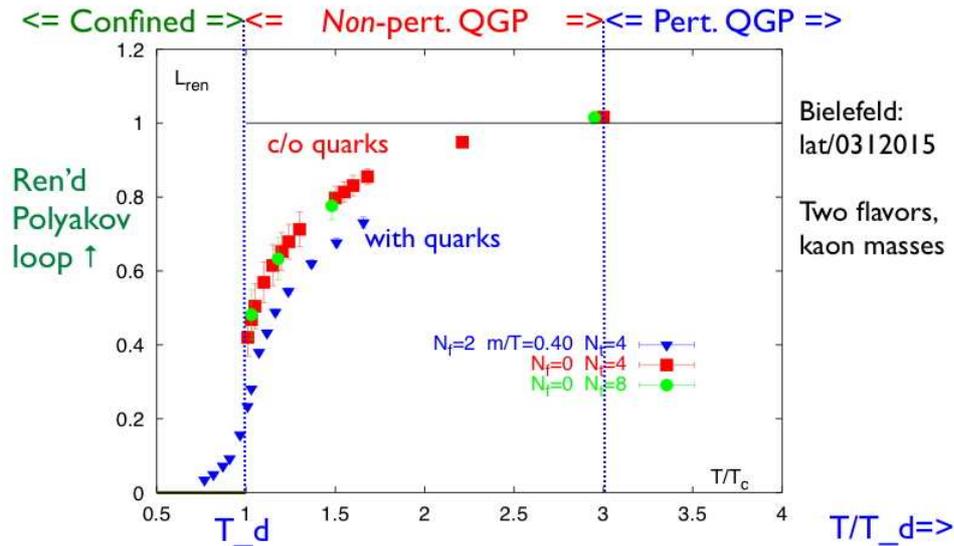
FIG. 2: Renormalized Polyakov Loop versus Temperature

collisions, where the nuclei overlap only partially. While theoretically central collisions are the cleanest situation, data on peripheral collisions are automatically collected. As we shall see, they also provide some crucial and unexpected insights.

Before RHIC, the usual picture of what would happen was based upon an analysis by Bjorken. Extrapolating from *pp* collisions, one expects a "central plateau": a region in which the particle multiplicity is constant versus rapidity. Rapidity is a type of longitudinal momentum for ultrarelativistic particles: in a collider, rapidity of $y = 0$ corresponds to sitting at 90 degrees, perpendicular to the beam. Nonzero rapidities correspond to moving along the beam direction. The advantage of rapidity is that a Lorentz boost simply adds to rapidity. Consequently, if one is in a regime invariant under boosts along the beam axis, then one expects a plateau, in particle multiplicity, along rapidity.

In *pp* collisions, a central plateau is not present at SPS energies, but does appear at RHIC energies. For central collisions of gold nuclei at RHIC, the central plateau is *much* narrower than in *pp* collisions.

(I refer to the central collisions of gold nuclei as *AA* , since these are the largest nuclei which can be collided at RHIC. "*A*" is the atomic number, $A \approx 200$ for gold nuclei.) At energies of 200 GeV/A, the central plateau is very narrow: asking that the particle multiplicity per unit rapidity, dN/dy, of identified particles is flat, the data indicates a plateau for $\Delta y_{CP} \approx \pm 1.0$. In constrast, the total range in rapidity is $\Delta y_{total} \approx \pm 5.0$. There are interesting details to the distributions, as well: if one requires not only that dN/dy is flat, but also that the average momentum is constant, gives a central plateau which is even narrower, $\Delta y_{CP} \approx \pm 0.5$. Note that to answer the question of a central plateau, it is crucial to have identified particles. If one doesn't know the identity of the particle, one doesn't know the mass, and one can only compute the "pseudo"-rapidity, η; then it appears that the central plateau is relatively broad,
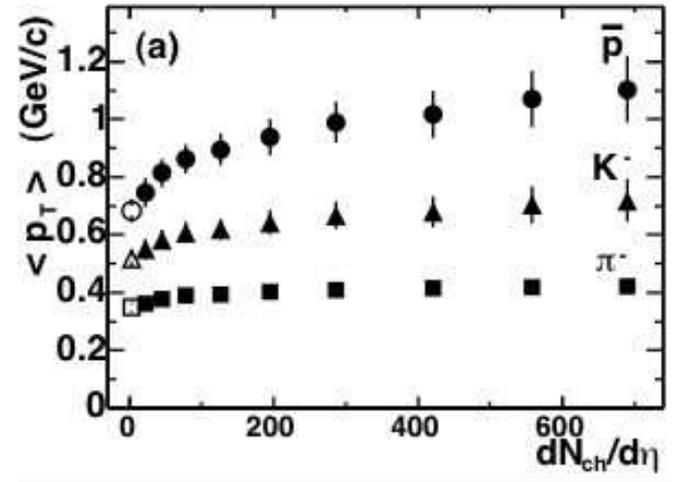

FIG. 3: Average momentum versus particle species

$\Delta \eta \approx \pm 2.0$.

At high energies, a nuclei becomes Lorentz contracted. For example, at energies of $\sqrt{s}/A : 45 \rightarrow 60$ GeV, a gold nucleus, which is about 15 fm in diameter in its rest frame, gets squashed into a pancake only $1/3 \rightarrow 1/4$ fm in width. As such, it represents an extremely strong color field. This gives rise to the picture of a Color Glass, proposed by McLerran and Venugopalan. As high energies, each pancake becomes a delta function, a gauge transformation on each side. This gives one a precise handle on the initial stages of an *AA* collision.

The Color Glass predicts that at high energies, *AA* scattering is semiclassical. As such, one can predict the dependence of the multiplicity and energy on the "saturation momentum" Q_{sat}. As a semiclassical process, one expects that the multi-

plicity grows as

$$\frac{dN}{d\eta} \approx \frac{Q_{sat}^2}{\alpha(Q_{sat})} \quad , \quad \frac{dE}{d\eta} \approx \frac{Q_{sat}^3}{\alpha(Q_{sat})} \quad , \qquad (5)$$

where N and E are the multiplicity and energy per unit area (this accounts for the dimensions in the above relations).

In going from SPS to RHIC energies, the increase in multiplicity is relatively modest; there are 600 particles per unit rapidity at the highest SPS energy, and about 1000 at the highest RHIC energy. This ruled out many cascade models of AA collisions, which had predicted large increases. For example, the increase in multiplicity from $\sqrt{s}/A = 130$ to 200 GeV is only 14%. This modest increase is rather natural in a Color Glass, where Q_{sat} grows slowly with energy.

In detail, however, the Color Glass does not describe particle production. This is especially true when one consider the average energy per particle. By the above equation, the increase of the average energy is greater than the increase in mulitplicity, by a factor of Q_{sat}. In contrast, the STAR experiment at RHIC claims that average transverse momentum, p_t, increases by only 2% between these energies, and not 21%, as predicted by the above.

The manner in which the energy is deposited is also wrong. A Color Glass produces gluons. Hadronization is modeled by invoking "Quark-Parton Duality", whereby one assumes that one gluon becomes one pion. This is an appealing picture; as Q_{sat} increases, one then expects that the average pion momentum will increase similarly. This is a logically consistent possibility: as one goes from pp to central AA, one would find a large increase in the pion momentum.

Fig. (3) shows the change in the average momentum, versus the particle species. In this Figure, the x-axis refers to centrality, so the left most points are for pp collisions, and the right most, central AA collisions. Contrary to expectations from the Color Glass, the average momentum of a pion increases by a rather small amount from pp to AA collisions. On the other hand, the average momentum of heavier particles, such as kaons and protons, increase dramatically.

This increase in momentum is interpreted naturally if the Color Glass only represents the initial, and not the final, state of AA collisions. Further, if a hydrodynamic picture can be invoked, then the increase of the average momentum of heavier particles can then be due to a large boost velocity for a medium in which these particles sit.

B. Signals at High Momentum

While the energies at which the nuclei are being collided is large, most of the particles are produced at small transverse momentum, $p_t \approx 170$ MeV or so. This is typical of hadronic collisions, from pp on up to AA. Since the Lattice indicates that the transition temperature is similar, one would naturally expect that the clearest signals of something new happening at RHIC are from soft momentum, on the order of the temperature.

Experimentally, this is not what happens. It turns out that the clearest signals are at high momentum, $p_t > 2GeV$ or so.

To quantify this, we define the ratio R_{AA}, as the ratio of the number of particles in an AA collision, to that in a pp collision. Of course it is necessary to normalize by the atomic number. In doing so, we expect that at hard momentum, the correct normalization factor is the number of collisions, which is given by the atomic number of one nucleus, $= A$, times the path length through which it travels, $= A^{1/3}$, for an overall factor of $A^{4/3}$. This factor is manifestly the number of hard collisions; soft collisions should only scale like A.

This ratio is plotted in Fig. (4), including all species of hadrons, for $\sqrt{s}/A = 200$ GeV. Since R_{AA} is constructed by dividing by the number of hard collisions, it is automatically below one for soft momenta, which should scale like A. (This explains why at zero momentum, R_{AA} goes to $1/A^{1/3} \sim 1/7$.)

There is a striking consistency of the measurements between the different experiments. At SPS energies, the ratio R_{AA} is about unity near $p_t \approx 2$ GeV. (Previous plots had found a value of $R_{AA} \approx 2 - 3$, but this was due to an inaccurate extrapolation of pp collisions.) Very recent measurements at $\sqrt{s}/A = 62$ GeV indicate that there is a smooth interpolation between SPS and RHIC energies, with R_{AA} decreasing smoothly, for a fixed value of p_t, as $\sqrt{s}/A$ increases. To simplify the discussion, I concentrate on the data at the highest energies.

The theoretical explanation of this effect is energy loss in a plasma. This is similar to energy loss in any medium: a fast particle scatters frequently, by many soft collisions, off the medium. There is a characteristic difference between energy loss in an abelian theory, such as QCD, and a non-abelian theory, such as QCD. In QED, energy loss is proportional to the total path length, L; in QCD, energy loss grows as L^2. Experimentalists are now extracting the path length dependence from the data, but preliminary analysis is confusing: a L^2 term is required, but the coefficient of an additional term, linear in L, is negative. This depends upon details of energy loss, also known as the Landau-Pomeranchuk-Migdal (LPM) effect. The L^2 behavior depends upon the ratio of the size of the system, the coherence length, and other details. Thus while a negative coefficient for a term linear in L is worrisome, it is not, a priori, nonsensical.

There is an important cross check to the interpretation of R_{AA} as due to scattering in the nucleus. The ratio was also computed in deuteron-gold collisions, which I refer to a pA (assuming dA is close to pA collisions). The ratio R_{pA} was then computed, normalizing now by the atomic number, A. One finds that at central rapidity, and $p_t \sim 2$ GeV, that R_{pA} is greater than one. The usual interpretation of this is initial scattering in the nucleus, termed the "Cronin" effect. Clearly one needs to understand the Cronin effect better, but the measurement certainly shows that was is going on in R_{AA} has to do with final state, and not initial state, interactions. An important test on consistency is that R_{dA} approaches unity as relatively modest momentum, $p_t \approx 8$ GeV and above. When this happens for R_{AA} is not clear; there is no evidence of it by momentum $p_t \approx 12$ GeV.

In this vein, the BRAHMS experiment also measured the ratio R_{dA} in the fragmentation regime of the deuteron, for pseudo-rapidity $\eta \approx 3.$. They find suppression of R_{dA} at for-
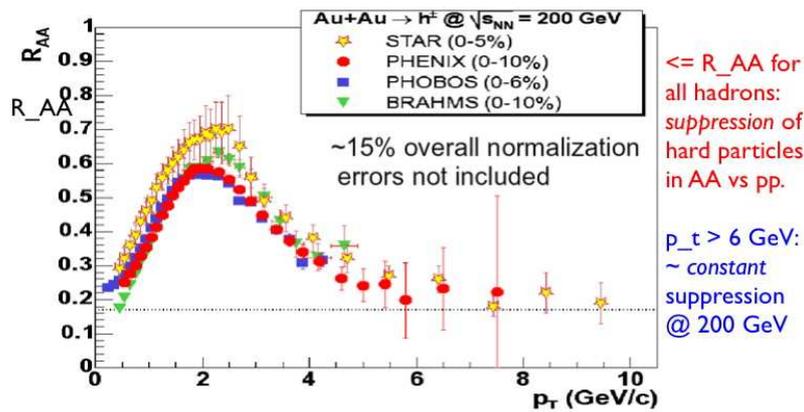
FIG. 4: The ratio R_{AA} versus transverse momentum

ward rapidity, $\approx .5$ by $p_t \approx 2$ GeV. This is expected from the Color Glass, as originally argued by Dumitru and Jalilian-Marian [10]. The fragmentation regime of the proton can be understood by sitting in its rest frame. In the proton rest frame, the nucleus is Lorentz contracted, and a large color charge. So purely on kinematic grounds, the best place to look for effects of the Color Glass is to look at the proton fragmentation regime in *pA* collisions. Similarly, final state effects from cold nuclei, such as energy loss in a cold nucleus, are best studied by looking at the nuclear fragmentation regime in *pA* collisions. Having said that, there are other models which predict the observed suppression of R_{pA} in the proton fragmentation regime; much more detailed measurements, such as of photon and dilepton production, will be able to definitively discriminate between the different models. As shown by the example of the overall multiplicity and energy, it is not enough to look at one quantity in isolation.

The behavior of R_{AA} also displays another interesting phenomenon, which was unexpected before the RHIC data. Note that R_{AA} has a maximum at $p_t \approx 2$ GeV, and then falls off. It appears to be constant for p_t greater than 6 GeV. This is due to a change in composition in the particle spectrum. If one looks at R_{AA} for neutral pions, then it is flat, $\approx .2$, from $p_t \approx 2$ GeV on up to the highest measured momenta, which are at present about 15 GeV.

In *pp* collisions, at $p_t \approx 2$ GeV the ratio of protons to pions is $\approx .1$. In contrast, at the same momentum in *AA* , this ratio is ≈ 1. That is, there is a "baryon bump" at intermediate momenta, for $p_t : 2 \rightarrow 6$ GeV, where the number of baryons is greatly enhanced over the number of mesons. This is not a mass effect, and is true for strange as well as non-strange mesons. An explanation of this is given in terms of a model of recombination of quarks into mesons and baryons: if there are primarily quarks about, then they are rather likely to coalesce and form baryons, as well as mesons. At higher momenta, the ratio is determined by perturbative QCD, and ordinary fragmentation functions. At lower momenta, one assumes that particle spectra are thermal; it is only in this window of momenta that recombination dominates. Recombination makes

one impressive prediction about the ratio of elliptic flow, described below. Even so, it does somewhat beg the question of why quarks dominate for these intermediate momenta.

There is an impressive test that the change in the spectrum is due to interactions with the nuclei. In a *pp* collision, one can directly look at jets: one sees a spray of particles in one direction, balanced (by momentum conservation) by a spray of particles in the other. In a *AA* collision, instead it is necessary to form a statistical measure of jets: one triggers on a hard particle, with momentum $p_t : 4 \rightarrow 6$ GeV, and then looks at the distribution of particles with $p_t : 2 \rightarrow 4$ GeV. In *pp* collisions, one sees a jet in the backward direction, as expected. In *AA* collisions, one does not: the backward jet is completely "eaten" by the nucleus in a central *AA* collision.

In Fig. 5 I show this jet–jet correlation for peripheral collisions. In peripheral collisions, by looking at the bulk of particles, which occur at soft momenta, one can unambiguously determine the reaction plane. This allows one to compute the jet–jet correlation for jets which occur in the plane, versus out of the plane. In a peripheral collision, the overlap region between the nuclei form an almond shaped region; a jet in the reaction plane transverses a smaller part of the almond than one perpendicular to the reaction plane. In Fig. 5, the forward, or trigger jet, is at $\Delta\phi \approx 0$ degrees; the backward jet is peaked about 180 degrees, or $\pi \approx 3.14...$ radians. In *pp* collisions, the backward jet is apparent. In *AA* collisions, what is striking is how the backward jet changes with the direction of the reaction plane. For jets in the plane, which go a short way through the almond, the backward jet is almost as large as that in *pp* collisions. In contrast, for jets out of the plane, which go through a long way through the almond, the backward jet is essentially absent. This provides a purely *geometrical* test that the change in jet behavior is due to the nucleus.

Further measurements have confirmed the picture that energy loss is occuring. Having identified the direction of the hard particle, one can look at the distribution of energy, going to lower and lower energies. Doing so, the data shows that in the direction backward to the jet, that the energy which disappears from high momentum goes into particles at low mo-
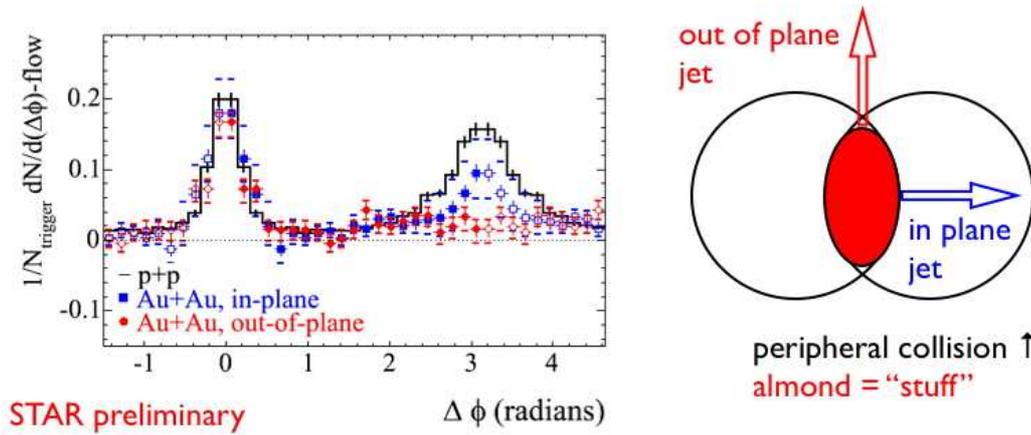
FIG. 5: Jet jet correlations for peripheral collisions

mentum. This is exactly what one expects from a fast particle slowing down, by numerous soft scatterings, in a medium.

C. Thermalization at soft momentum?

At hard momentum, there is unambiguous evidence that there is "stuff" formed in central *AA* collisions which dramatically affects the propagation of hard particles. This doesn't mean that the stuff is matter in thermal equilibrium; for that, rather detailed estimates of energy loss and the like are required.

But it certainly does show that *AA* collisions do not simply act like a superposition of *pp* collisions, and that qualitatively new behavior is occuring. As noted above, for a system as a temperature $T \approx T_c$, one would expect that the most obvious signals for possible thermalization are for momenta on the order of the temperature. In this section I summarize the results of experimental measurements at RHIC at soft momenta. While at first sight it appears to confirm the picture of a thermalized system, the details are such that there are rather significant questions open as to whether this has, in fact, occured. What cannot be avoided is that the interactions in the system are strong. I summarize several features.

Chemical equilibriation at $T_{chem} \approx 160$ *MeV*, with a small baryon chemical potential, $\mu_{baryon} \approx 24$ MeV. Consider first overall abundances, integrated over transverse momentum. Then one can fit literally dozens of particle abundances with two parameters: an overall temperature for chemical equilibriation, T_{chem}, and a small baryon chemical potential, μ_{baryon}. This includes strange mesons, the K's and ϕ's, and strange baryons, Λ's, Ξ's, and Ω's, along with all anti-particles.

It is known that similar fits can be done at lower energies, and smaller systems, even for *pp* collisions. For lower energies in *AA* collisions, however, it is necessary to add a "fugacity" for strangeness, which represents a departure from strict thermal equilibrium (that is, it is a fudge factor). Similarly, for *pp* collisions, corrections must be added for finite volume.

RHIC energies represent the first time that overall chemical ratios can be predicted from a textbook application of Bose-Einstein, or Fermi-Dirac, distribution functions. I also note that the parametrization does not work well for short lived reasonances, such as the ρ, Δ, K^*, and Λ^*. Note that this temperature is close to that for the transition temperature.

Of course this does *not* demonstrate that thermalization has occured; it is a necessary, but not a sufficient condition. Nevertheless, it is a remarkably good way of summarizing the data; it works even for Ω's, whose abundance is $\approx 0.1\%$ that of pions. As such, it cannot be ignored.

Kinetic equilibrium, $T_{kin} \approx 100$ *MeV,* $\beta \sim .7$. Besides looking at overall abundances, for each species, one can look at the momentum distribution, and compare to a thermal distribution. The basic feature was apparent from Fig. (3): the momentum of pions does not increase much, while kaons, and especially protons, have a strong increase in average momentum. The only way to incorporate this is to assume that particles are emitted from a local rest frame which has a large boost velocity with respect to the lab frame. Further, by symmetry, this boost velocity must vanish at the center of the nucleus, and have its maximum, β, at the surface of the nucleus. (The fits are not very sensitive to the dependence of this velocity with radius; typically, fits assume $\beta(r) \sim r^a$, where $a : 0.5 \to 2.0$.) Doing so, one finds good fits for pions, kaons, and protons with a single temperature of $T_{kin} \approx 100$ MeV, and a boost velocity of $\beta \approx .7$.

The usual explanation of the difference between the temperatures for chemical and kinetic equilibrium is the existence of a hadronic phase. Chemical equilibrium requires processes which change particle number, and are expected to decouple before processes which maintain kinetic equilibrium, which only require scattering which changes momenta in a collision, but not the number of particles.

However, while all particle ratios can be fit with a single temperature for chemical equilibrium, the momentum distributions cannot be fit with a single temperature and boost velocity. Instead, strange baryons cannot only be fit with a higher

temperature, and a lower boost velocity. This is usually described as due to "partonic equilibration", whereby strange baryons decouple earlier. This begs the question, though, of why overall abundances appear to be fit with a single temperature. If kinetic equilibrium is fit with a variety of temperatures and boost velocities, why isn't the same true for the overall ratios, which determines T_{chem}?

This suggests, at least to the author, that perhaps the distributions are not thermal, but something else.

(Ideal) hydrodynamics: works for elliptic flow. With about one thousand particles per unit rapidity, it is natural to try a hydrodynamic description. This assumes that locally, fluid elements are in thermal equilibrium, where the local element carries some large boost velocity. The particle distributions for pions, kaons, and protons, described above, can be fit with the above T_{kin} and boost velocity, β. The distributions of heavier particles, especially strange, cannot be described by hydrodynamics. Further, there is a peculiar feature: the initial time at which hydrodynamics begins is extremely short, $\tau \approx 0.6$ fm/c, and not several fm/c, as one might have expected.

There is a great success of a hydrodynamic analysis. For a peripheral collision, one can look at the distribution of particles with respect to the reaction plane. This allows one to compute "elliptic flow", which are moments with respect to the reaction plane. Doing so, one finds that there is a large asymmetry. For transverse momenta below ≈ 1 GeV, the elliptic flow of pions, kaons, and protons is well fit by ideal hydrodynamics.

There is another notable feature of elliptic flow. For momenta larger than ≈ 1 GeV, the elliptic flow is constant, for momenta up to the largest value measured, ≈ 5 GeV. There is no fundamental understanding of why elliptic flow is constant at these values of the momenta. Even so, the values do satisfy a prediction of recombination: the values scale with the numbers of quark in the hadron, so for $p_t : 1 \rightarrow 5$ GeV, the elliptic flow of baryons is $\approx 3/2$ that of mesons.

Indeed, the fit to elliptic flow works *only* if one assumes ideal hydrodynamics. If some reasonable value of viscosity is added, then the fit no longer works. This has led some to describe the theory at RHIC as a "sQGP", or strongly coupled QGP. As discussed previously, however, the Lattice data does not support large increases in the QCD coupling near the phase transition.

There are also problems in explaining the rapidity dependence of elliptic flow. Elliptic flow is constant over the narrow rapidity region in which both dN/dy and average momentum of identified particles are constant, ± 0.5 in rapidity. But at larger momenta, data from PHOBOS indicates that the rapidity dependence of elliptic flow falls off much quicker than particle multiplicity.

(Ideal) hydrodynamics: fails for HBT radii. Nearly ideal hydrodynamics appears to work well for single particle distributions of light hadrons (although again, it doesn't explain why heavier particles don't fit).

More detailed information about the space-time history of the collision can also be extracted. This applies the Hanbury-Brown-Twiss (HBT) effect to *AA* collisions. HBT is the standard method for determining the size of a star, and depends upon the interference which must occur between identical particles.

In words, one does the following. Consider a pair of, say, pions. Form the ratio of the two-particle correlation function, to the product of single particle correlation functions at the same momentum. This ratio is peaked and zero relative momentum, and then falls off with increasing (relative) pair momentum. The fall off is typically exponential, which allows one to determine a size. By kinematics, there are three possible directions: along the beam, which gives a longitudinal size R_{long}; along the direction of the pair, which gives a size R_{out}; perpendicular to the direction of the pair, which gives a size R_{out}.

Before the RHIC data, it was thought that a strongly first order transition would give large HBT radii. Then Coulomb corrections would be important, *etc.* Instead, the RHIC data showed that the HBT radii aren't large, but *small.* Indeed, from SPS to RHIC energies, they barely increase.

In fact, the HBT radii appear to be *much* too small. Comparing to the predictions of hydrodynamics, the longitudinal size R_{long} is too large by approximately a factor of two.

The other two distances also behave completely wrong. A nuclei isn't a sharp surface, but is represented by a Woods-Saxon form, with a smooth fall off of the nuclear density. Because of this, the nuclei tends to "burn" from the outside in. For a strongly first order transition, the analogy was to a burning log. The predicted behavior of the other two HBT radii was that R_{out}/R_{side} would be greater than one, say ≈ 1.5, and increase with increasing pair momentum. Instead, both PHENIX and STAR experiments show that after including Coulomb corrections, R_{out}/R_{side} is essentially one, and independent of momentum for pair momenta $100 \rightarrow 400$ MeV.

The HBT radii can be fit by a "blast wave" model. This abandons the connection between position space, and momentum space, which of necessity follows from ideal hydrodynamics. Emission from a sharp surface will also give R_{out}/R_{side} near one, since a ratio near one indicates that the emission surface is moving to the observer as fast as possible. To describe the longitudinal size, it is necessary to relax the assumption of boost invariance along the beam direction, and take a starting point closer to that of Landau hydrodynamics, where the two nuclei are assumed to stop, and then evolve.

One might ask, well, perhaps single particle correlation functions are all that we should hope to describe; perhaps two particle correlations are simply more than we have any right to expect. However, it is very difficult to accept that single particle distributions can be described by hydrodynamics, but that two particle distributions are such that the total space-time volume of the collision is off by a factor of *four* (R_{out} and R_{long} are each off by a factor of two; R_{side} is close to the data).

An alternate explanation is that we do not yet understand how hadronization occurs in *AA* collisions.

In summary, the data from RHIC is rather tantalizing. From hard momenta, such as R_{AA} and jet–jet correlations, it is clear that central *AA* collisions produce *some* sort of new matter. This matter slows down fast particles in a way which can be verified in a purely geometrical fashion.

However, when this matter is studied directly, by looking at

soft momenta, things behave in an unexpected fashion. Overall particle abundances are well described by a single temperature, but the momentum dependence appears to be described by temperatures and local boost velocities which depend on the identity of the particle. In looking at the angular distributions for peripheral collisions, the nuclei appear to be very "sticky", with large elliptical flow.

[1] BRAHMS collaboration, [arXiv:nucl-ex/0410020]; PHENIX collaboration, [arXiv:nucl-ex/0410003]; PHOBOS collaboration, [arXiv:nucl-ex/04100022]; STAR collaboration, [arXiv:nucl-ex/0501009].

[2] http://quark.phy.bnl.gov/~pisarski/talks/unicorn.pdf

[3] D. Scott, attributed by R. Stock, in Proceedings of Quark Matter '04; M. Gyulassy, numerous unpublished talks.

[4] A. Dumitru, Y. Hatta, J. Lenaghan, K. Orginos and R. D. Pisarski, Phys. Rev. D **70**, 034511 (2004) [arXiv:hep-th/0311223].

[5] O. Kaczmarek, F. Karsch, P. Petreczky, and F. Zantow, Phys. Lett. B **543**, 41 (2002) [arXiv:hep-lat/0207002]; S. Digal, S. Fortunato, and P. Petreczky, [arXiv:hep-lat/0211029]; O. Kaczmarek, F. Karsch, P. Petreczky and F. Zantow, Nucl. Phys. Proc. Suppl. B **129**, 560 (2004) [arXiv:hep-lat/0309121]; O. Kaczmarek, S. Ejiri, F. Karsch, E. Laermann and F. Zantow, [arXiv:hep-lat/0312015]; P. Petreczky and K. Petrov, [arXiv:hep-lat/0405009].

[6] R. D. Pisarski, Phys. Rev. D **62**, 111501 (2000) [arXiv:hep-ph/0006205]; A. Dumitru and R. D. Pisarski, Phys. Lett. B **504**, 282 (2001) [arXiv:hep-ph/0010083]; A. Dumitru and R. D. Pisarski, Phys. Lett. B **525**, 95 (2002) [arXiv:hep-ph/0106176]; Phys. Rev. D **66**, 096003 (2002) [arXiv:hep-ph/0204223]; O. Scavenius, A. Dumitru and J. T. Lenaghan, Phys. Rev. C **66**, 034903 (2002) [arXiv:hep-ph/0201079]; A. Dumitru, O. Scavenius, and A. D. Jackson, Phys. Rev. Lett. **87**, 182302 (2001) [arXiv:hep-ph/0103219]; A. Dumitru, D. Röder and J. Ruppert, [arXiv:hep-ph/0311119]; A. Dumitru, J. Lenaghan, and R. D. Pisarski, [arXiv:hep-ph/0410294].

[7] J. O. Andersen and M. Strickland, [arXiv:hep-ph/0404164]; J. P. Blaizot, E. Iancu and A. Rebhan, [arXiv:hep-ph/0303185]; U. Kraemmer and A. Rebhan, Rep. Prog. Phys. **67**, 351 (2004) [arXiv:hep-ph/0310337];

[8] O. Kaczmarek, F. Karsch, F. Zantow and P. Petreczky, Phys. Rev. D **70**, 074505 (2004) [arXiv:hep-lat/0406036].

[9] L. D. McLerran and R. Venugopalan, Phys. Rev. D **49**, 2233 (1994) [arXiv:hep-ph/9309289]; E. Iancu and R. Venugopalan, [arXiv:hep-ph/0303204]; R. Venugopalan, [arXiv:hep-ph/0412396].

[10] A. Dumitru and J. Jalilian-Marian, Phys. Rev. Lett. **89**, 022301 (2002) [arXiv:hep-ph/0204028].