Heavy ion physics at the Large Hadron Collider: what is new ? What is next?

This content has been downloaded from IOPscience. Please scroll down to see the full text.

2013 Phys. Scr. 2013 014003

(<http://iopscience.iop.org/1402-4896/2013/T158/014003>)

View [the table of contents for this issue](#), or go to the [journal homepage](#) for more

Download details:

IP Address: 137.138.125.164

This content was downloaded on 21/01/2014 at 12:23

Please note that [terms and conditions apply](#).

Heavy ion physics at the Large Hadron Collider: what is new ? What is next?

J Schukraft

CERN, DIV. PH, CH-1211 Geneva 23, Switzerland

E-mail: schukraft@cern.ch

Received 1 October 2013

Accepted for publication 12 November 2013

Published 6 December 2013

Online at stacks.iop.org/PhysScr/T158/014003

Abstract

Towards the end of 2010, some 25 years after the very first collisions of ultra-relativistic heavy ions at fixed target energies, and some 10 years after the start of operation of the Relativistic Heavy Ion Collider, the Large Hadron Collider (LHC) opened a new era in heavy ion physics with lead on lead collisions at $\sqrt{s_{NN}} = 2.76$ TeV. After a short reminder of the main results from lower energies, this review highlights a few selected areas where significant progress has been made during the first three years of ion operation at the LHC.

PACS number: 25.75.-q

1. Introduction

The subject of ultra-relativistic heavy ion physics is the study of strongly interacting matter under extreme conditions of high temperature and/or high matter density. QCD predicts that at a sufficiently high energy density there will be a transition from ordinary nuclear or hadronic matter to a plasma of free (deconfined) quarks and gluons, the ‘quark–gluon plasma’ (QGP). The discovery and characterization of this plasma phase is thought to require a large volume of hot/dense matter and is therefore pursued in collisions of heavy nuclei at the highest energies.

This paper presents a subjective selection of results and interpretations from the first three years of the Large Hadron Collider (LHC) heavy ion program; it is based in parts on recent reviews [1, 2] in which most of the relevant LHC data and references to the primary literature can be found. Theoretical aspects and implications are covered and referenced in more detail in an accompanying paper [3].

1.1. Prior art: main results from fixed target experiments and RHIC

After pioneering experiments at relativistic energies (GeV per nucleon) in the 1970s in the USA (LBNL) and Russia (JINR), the quest for the QGP took off in 1986 at the fixed target machines at CERN (SPS) and BNL (AGS), first as an exploratory program with light ions (mass ≈ 30) and in the early 1990s with actual heavy ions (mass ≈ 200) [4–6]. An appraisal of the SPS program was made in 2000 [7], based on a ‘common assessment’ of the results from

half a dozen experiments collected and published over the preceding years. It concluded that *compelling evidence has been found for a new state of matter, featuring many of the characteristics expected for a QGP*¹. This conclusion was based primarily on three experimental observations: the copious production of hadrons containing a strange quark (strangeness enhancement), the yields of low mass lepton pairs (rho melting), and the reduced production of J/Ψ mesons (anomalous J/Ψ suppression).

The very high abundance of strange particles, in particular of hyperons (the omega-to-pion ratio increases by up to a factor 20 from pp to PbPb!), was predicted as a consequence of QGP formation. Today it is interpreted more generally as a manifestation of statistical hadronization from a thermalized medium, where most hadrons, not only those containing strange quarks, are created in thermal equilibrium ratios [8–11]. These ratios are governed essentially by a single scale parameter T , interpreted as a chemical freeze-out temperature.

The invariant mass distribution of prompt low mass lepton pairs showed an enhancement in the continuum yield just below the rho/omega resonance [13, 14]. This was qualitatively in agreement with signals expected from chiral symmetry restoration, where the mass and/or width of hadrons are modified in the vicinity of the QGP phase boundary (in this case the rho, observed inside the medium via its two lepton decay).

Finally, J/Ψ production was found significantly suppressed in central PbPb collisions relative to expectation

¹ Available at <http://press.web.cern.ch/press-releases/2000/02/new-state-matter-created-cern>.

(e.g. relative to a normalization process like Drell Yan) beyond what could be attributed to confounding cold nuclear matter effects like changes in the nuclear parton distribution functions or hadronic final state interactions [16–18]. Dubbed ‘anomalous’ J/Ψ suppression to distinguish it from the ‘normal’ one due to cold nuclear matter, the effect was within experimental and theoretical uncertainties fully consistent with the ‘smoking gun’ signal predicted for deconfinement in the QGP.

These experimental results have all stood the test of time, having been confirmed and refined subsequently at the SPS as well as at Relativistic Heavy Ion Collider (RHIC). The essence of the assessment, that there is a new state of matter in the SPS energy range, featuring some of the hallmarks of a QGP (thermalization, deconfinement and chiral symmetry restoration) seems however, in hindsight, today more compelling than in 2000, given for example much improved low mass lepton pair results from the SPS NA60 experiment [15] and new insights in some of the processes relevant for thermal particle production and quarkonia suppression from RHIC and LHC (see below).

The initial results from RHIC were summarized and assessed in 2005, based on a comprehensive (re)analysis of the first few years of RHIC running [19]. The experiments concluded that at RHIC *a new state of hot, dense matter was created out of the quarks and gluons .. but it is a state quite different and even more remarkable than had been predicted*². Unlike the expectation, with hindsight overly naive, that the QGP would resemble an almost ideal gas of weakly coupled quarks and gluons, the hot matter was found to behave like an extremely strongly interacting, almost perfect liquid, sometimes called the sQGP (where the ‘s’ stand for ‘strongly interacting’). It is almost opaque and absorbs much of the energy of any fast parton which travels through—a process referred to as jet quenching—and it reacts to pressure gradients by flowing almost unimpeded and with very little internal friction (i.e. has very small shear viscosity) [20, 21].

Also at RHIC, the crucial experimental results as well as the inferred characteristics of the QGP—a ‘hot, strongly interacting, nearly perfect liquid’—stood the test of time [1]. The temperature was inferred by measuring direct ‘thermal’ photons with an inverse slope of order 200 MeV, which leads to a (model dependent) estimate of an initial temperature of at least 300 MeV. The characterization of the QGP as almost opaque was based on the observation of a very significant suppression, up to a factor of five, of high p_T particles found in central, head-on nuclear collisions. This suppression of high p_T particles, which are typically leading jet fragments, is indicative of very strong final state interactions of the scattered partons with the medium, leading to significant energy loss via elastic scattering or enhanced gluon radiation. As essential control measurements, the suppression was neither seen in d-Au reactions, therefore excluding cold nuclear matter effects as the cause, nor with colour neutral probes like direct photons, which clearly establishes the effect as due to the strong (=QCD) interaction in the final state.

The ‘ideal liquid’ aspect of the QGP was based on the measurement of collective particle motions, the so called ‘elliptic flow’, which develops in response to initial

geometrical conditions and internal pressure gradients in the nuclear overlap zone. The elliptic flow magnitude at RHIC was found to essentially exhaust the maximal possible one predicted by hydrodynamics for the given initial deformation, equivalent to the response of an ideal liquid with vanishing shear viscosity. The shear viscosity over entropy ratio, η/s was found to be compatible with a conjectured lower bound of $\eta/s \geq 1/4\pi$ ($\hbar = k_B = 1$), a value reached in a very strongly interacting system when the mean free path approaches the quantum limit, the Compton wavelength.

1.2. What’s new from the LHC?

Prior to LHC, 25 years of heavy ion experimentation had already revealed a ‘QGP-like’ state at the SPS and not ‘the QGP’, but ‘a sQGP’ at RHIC. One could consider the discovery phase for the QGP as essentially over, the qualitative characterization as well under way, and quantitative precision measurements of QGP properties as just having started [22]. A main goal for the heavy ion program at LHC was therefore to measure with increased precision the parameters which characterize this new state of matter, making use of the particular strength of the LHC: a powerful new generation of large acceptance state-of-the-art experiments ATLAS, CMS, ALICE and LHCb³, and a huge increase in beam energy with the associated larger cross sections for hard probes and higher particle density, which makes for a QGP which will be ‘hotter, larger and longer living’. And indeed, LHC made significant progress towards increasing the precision on shear viscosity (see section 4) and plasma opacity (section 5) already during the first two years of ion running. However, when dealing with QCD in the nonperturbative regime, surprises should not come as a surprise, and a number of unexpected findings at LHC have helped shed new light on some old problems or issues. Two of those will be mentioned below, relating to particle production (i.e. thermalization, section 2) and quarkonia suppression (i.e. deconfinement, section 3). And finally the very first discovery made at LHC is discussed in section 5: the appearance of a mysterious long range ‘ridge’ correlation in high multiplicity pp reactions. It reappeared later—and much stronger—in the 2012 p-nucleus run, making it of great interest, and presumably of great relevance, to hot and dense matter physics, even if it is ultimate cause and connection to similar phenomena in nuclear collisions is as of today not finally settled.

2. Particle production

The production of different hadronic particle species in high energy collisions is a nonperturbative process and in general parameterized in phenomenological models and event generators, requiring usually a large number of parameters. A rather more economical approach is the thermal/statistical model of hadronization [8–12], which assumes that particles are created in thermal (or phase space) equilibrium; particles of mass m are essentially suppressed by a Boltzmann factor $e^{-m/T}$. In its simplest implementation, the scale parameter T , identified as the chemical freeze-out temperature, is the

² Available at www.bnl.gov/newsroom/news.php?a=1303.

³ LHCb participates in the p-nucleus part of the heavy ion program.

only free parameter of the model and defines the production of all hadronic species, together with some conservation laws (baryons, strangeness, etc) and an overall normalization constant proportional to the total particle multiplicity. For small systems (small number of final state hadrons), an additional parameter γ_s (or, equivalently, a reduced correlation volume) has to be introduced to describe the fact that strange hadrons are suppressed compared to the grand canonical thermal expectation. The temperature parameter T is found in all high energy collisions (pp , e^+e^- , AA) to be about 160 MeV, while γ_s increases from 0.5 – 0.7 in pp to 0.9 – 1 in AA . The fact that most hadrons containing light and strange quarks are described by the thermal model with typically better than 10–20% precision was considered an essential and well established fact in heavy ion collisions.

The origin of the success of the thermal model over a very large range of collision systems and beam energies is however not obvious. One of two qualitative mechanisms are usually invoked which can be summarized as (i) *born into (phase space) equilibrium* or (ii) *evolving into (thermal) equilibrium* [11, 23, 24]. The former (i) could arise for example from Fermi's golden rule, the fact that reaction rates are proportional in this case to the product of a QCD matrix element and a hadronic final state phase space factor. As many different channels (matrix elements) contribute to particle yields, phase space dominates in inclusive measurements and the most conspicuous QCD remnant is the observed suppression ($\gamma_s < 1$) of strange particles in elementary reactions (pp , e^+e^-), reflecting the higher mass of the strange quark and local strangeness conservation. This strangeness suppression is then 'somehow' lifted in large systems (AA), effectively replacing local strangeness conservation by average global (grand canonical) conservation. The alternative explanation (ii) postulates that inelastic reactions either in the *partonic phase* or in the final state *hadronic phase* drive initial abundances quickly towards thermal equilibrium via detailed balance. Reaction rates, in particular those involving more than two initial state hadrons, decrease drastically with temperature and therefore inelastic reactions seize abruptly when the system expands and cools, preserving chemical equilibrium ratios and a freeze-out temperature still very close to the hadronization transition.

Both explanations for the uncanny success of statistical models are conceivable, but difficult to underpin by quantitative dynamical calculations. And there is no fully satisfactory explanation for the strikingly similar, yet distinct, pattern of particle production in small (pp , e^+e^-) and large (AA) systems: in (i) the mechanism of strangeness enhancement (i.e. the change in γ_s or the correlation volume) remains qualitative and hand waving, whereas in (ii) particle production in small systems, which would not be expected to thermalize, is usually not even addressed. Some experimental observations are also counterintuitive: if hadronization is sudden, from an equilibrated partonic phase (the sQGP) with little inelastic hadronic final state interactions, the particle ratios should reflect parton equilibrium, which in general would result in particle fractions very different from hadron equilibrium. If, on the contrary, particle yields are established via hadronic reactions, a common chemical freeze-out of all particle species seems *a priori* unlikely. One should expect

to see at least *some* indications for sequential freeze-out, where hadrons with large inelastic cross section stay longer in equilibrium and freeze out later (at lower temperature) than those with significantly smaller reaction cross sections.

Measuring identified particles at LHC was nevertheless considered a somewhat boring exercise, as finding thermal particle ratios essentially identical to the ones measured at RHIC was thought to be one of the safest predictions [25]. It therefore came as quite a surprise when some particle fractions, in particular for the mundane proton, one of the most frequently produced hadrons, were found to differ considerably from expectations (and, to a lesser extent, from the ones measured at RHIC), while others, including those for multi-strange hyperons, were well in line with thermal predictions. Also the particle ratios measured in proton collisions seem less well described by thermal fits than data at lower energy, even when allowing for strangeness under-saturation ($\gamma_s < 1$) [26].

The increasing deviation of hadron ratios with energy in pp collisions, despite the overall increase in multiplicity which should bring results closer to statistical predictions, could fit naturally with the 'QCD plus phase space' interpretation. At the LHC, cross sections for hard processes are large and hard scattering is important even for the average minimum bias collision, so maybe semi-hard QCD processes are making their presence felt as stronger deviations from phase space dominance?

Converging on a potential reason for the measured nuclear particle ratios at LHC proved harder, and a lively discussion is taking place centring on three explanations: (i) Reduced thermal freeze-out temperature; (ii) sequential hadron freeze-out; and (iii) nonequilibrium parton freeze-out. Standard one parameter thermal fits (i) give a somewhat poor description of measured particle ratios and return a chemical freeze-out temperature T significantly below the one extracted previously from RHIC (by 6–10 MeV) [27]. Even if one accepts the LHC fits as tolerable and within the range of accuracy of the thermal model, the question why and how the chemical freeze-out temperature would come down with increasing energy remains open. Sequential models (ii) try to estimate abundance changes from inelastic hadronic final state interactions after initial chemical freeze-out [28]. In any thermally evolving system, sequential (time and temperature ordered) freeze-out of different degrees of freedom with different cross sections and mean free paths *must* exist at some level; the question is one of magnitude rather than of principle. Some nonequilibrium calculations, usually done with the help of event generators or kinetic transport models, give results remarkably close to the LHC data and also improve thermal model fits at lower energies, making this mechanism a plausible explanation. However, given the multitude of unknown hadronic reaction cross sections which are needed in the final state rescattering calculations, as well as the need to consider multiprong initial states to satisfy detailed balance requirements, make these calculations more of an art than an exact science and the arguments qualitative rather than quantitative ones. Finally (iii), extensions to the standard model [29], where nonequilibrium thermodynamics (including supercooling below the phase transition) is applied to the parton rather than to the hadron phase, can well fit all

currently available data, albeit at the expense of two additional free parameters.

The final resolution of the ‘proton puzzle’ is still outstanding, and will require a more complete set of particle ratio measurements at LHC as well as revisiting the RHIC results to confirm with better significance if particle ratios in central nuclear collisions indeed evolve with energy. Whichever explanation will finally prevail, the unexpected LHC results are a welcome fresh input likely to advance our understanding of the ‘unreasonable success’ of the statistical/thermal model of particle production.

3. Quarkonium suppression

Heavy flavour quarks (charm, bottom) have always been an important tool to probe the QGP [16–18]. They are created early in the collision, by (semi)hard processes amenable to QCD calculations, and their further dynamical evolution is then modified by the surrounding medium. In particular the J/ψ and Υ families should be suppressed in heavy ion collisions in comparison with pp , primarily as a consequence of deconfinement (‘melting’) in the QGP. The magnitude of the suppression for different quarkonium states should depend on their binding energy, with strongly bound states such as the Υ showing less or no modification.

While the ‘anomalous’ J/ψ suppression discovered at the SPS was considered one of the strongest indications for the QGP, the RHIC results showed essentially the same suppression at a much higher energy, contrary to most expectations and predictions from both QGP and non-QGP models. These initially very confusing results kept the interpretation of the smoking gun signal for deconfinement ambiguous for the last 10 years. A possible explanation was that the directly produced J/ψ is not suppressed at all, neither at SPS nor at RHIC, and only the high mass charmonium states ψ' and χ_c , which populate about 40% of the observed J/ψ yield, are affected by the medium. These weakly bound high mass states should not only dissociate very close to or even below the critical transition temperature but also they may be easily destroyed by hadronic final state interaction, without the need for invoking a QGP. Alternatively, it has been suggested that J/ψ suppression actually increases with energy but is more or less balanced by a new production mechanism, i.e. recombination at the phase boundary of two independently produced charm quarks [30].

LHC data seems to have resolved the J/ψ puzzle in favour of the coalescence picture [31]. As predicted by recombination, the large charm cross section at LHC leads to *less* J/ψ suppression at LHC, albeit not to J/ψ enhancement relative to pp , which was within the range of coalescence prediction and would have made the case clear cut. The suppression is also less strong at low p_T , where phase space favours recombination, in clear contrast to the opposite p_T dependence found at SPS and RHIC.

LHC results from the Υ family [32] are fully consistent with the expectation from a deconfining hot medium that quarkonia survival decreases with binding energy, i.e. in terms of suppression factors: $\Upsilon(3S) > \Upsilon(2S) > \Upsilon(1S)$. The $\Upsilon(1S)$ is suppressed by about a factor of two in central collisions, the $\Upsilon(2S)$ by almost an order of magnitude, and only upper

limits have been measured for the $\Upsilon(3S)$. As only about 50% of the observed $\Upsilon(1S)$ are directly produced, these results are actually compatible with almost complete melting of all high mass bottomonium states and survival of a lone, strongly bound $\Upsilon(1S)$, which according to lattice QCD may melt only at temperatures far above the critical temperature.

While at first sight charm quark coalescence may appear as yet another process complicating and masking quarkonium deconfinement, it is actually a respectable and important deconfinement signal in itself: only in a colour conducting, deconfining medium can quarks roam freely over large distances ($\gg 1$ fm), and this is exactly what two charm quarks—unlike light quarks produced independently at well separated locations and at early times—have to do in order to combine during hadronization.

Pending surprises from the LHC’s p-Pb run, which should quantify the role of cold nuclear matter effects (shadowing/saturation and final state interactions), heavy quark diffusion and recombination is well on the way of becoming a new and accepted ingredient needed to make sense of the different patterns seen in Υ and J/ψ production over a wide range of energies. Together with, and complementary to, the original advocated process of dissociation, it may even eventually deliver on the promise of quarkonia production as an unambiguous signal of deconfined partonic matter.

4. Elliptic flow

The observation of robust collective flow phenomena in heavy ion reactions [33] is arguably the most direct evidence for the creation of a strongly interacting, macroscopic (i.e. large compared to the mean free path) and dense matter system in nuclear collisions. Matter properties like the equation of state, sound velocity or shear viscosity, can be extracted by comparing measurements and hydrodynamic model calculations of elliptic (i.e. azimuth dependent) and radial (azimuthally averaged) flow. Flow depends not only on matter properties but also on initial conditions, in particular the geometrical distribution of energy density within the nuclear overlap zone and the resulting pressure gradients. In general the nuclear impact parameter, and therefore the reaction zone geometry relevant for initial conditions, can be measured rather well event-by-event for each individual collision, using global event observables like particle multiplicity or forward ‘zero degree’ energy. However, the remaining model dependence of the energy density profile is sufficiently large to dominate the systematic errors and limit the accuracy with which the matter parameters can be extracted from the data. Before LHC turn-on, the defining property of the ‘perfect liquid’, the shear viscosity-to-entropy ratio η/s , was only known to be within a factor of about five to the conjectured quantum limit $1/4\pi$.

Azimuthally dependent collective motions are usually analysed in terms of a Fourier expansion with respect to the reaction plane, with the first order component, $v_1 \propto \cos(\varphi)$, called directed and the second order component, $v_2 \propto \cos(2\varphi)$, called elliptic flow. Higher order components were thought to be small or vanish for symmetry reasons. Between 2005 and 2010, based on observations

at RHIC, suggestions were made that the geometrical overlap shape could fluctuate event-by-event, even at fixed impact parameter, because of the stochastic nature of nucleon–nucleon collisions. These fluctuations would generate ‘lumpy’ initial conditions which could give rise to higher harmonic Fourier components. However, these suggestions remained controversial.

When first azimuthal flow data from the LHC became available in early 2011, the evidence from all three experiments, as well as new results shown by the two RHIC collaborations, was overwhelming [34]: fluctuations, event-by-event, of the energy density in the initial state do give rise to complex collective flow patterns, which, when analysed in terms of Fourier coefficients, are measurable and significant up to at least sixth order ($v_1, v_2, \dots, v_6$)! The signal strength of different harmonics, their particle mass, centrality and momentum dependence were all in excellent agreement with expectations from hydrodynamics.

The correlation patterns induced by flow fluctuations had actually been strong and clearly visible since many years also in the RHIC data; however, before 2011, they were in general not recognized as hydrodynamic in origin but discussed in terms of fancy names (near side ridge, away side cone) and fancy explanations (gluon Cerenkov radiation, Mach cone, etc) [35]. At LHC, the large acceptance of the experiments, together with the high particle density (as a collective effect, the flow signal increases strongly with multiplicity) made the observation and interpretation straightforward and unambiguous.

The fact that energy density fluctuations on the scale of a fraction of the nuclear radius in the initial state are faithfully converted into measurable velocity fluctuations in the final state was a most amazing, and also most useful, discovery: one could not only identify the average, almond shaped ‘face of the collision zone’, but recognize much finer structures, the ‘warts and wrinkles’, of nuclear collisions. The analysis of flow has been invigorated and is advancing rapidly ever since, with direct measurements of the fluctuation spectrum [36], using event-by-event measurement and selection of flow as an analysis tool [37], and even finding nonlinear mode mixing between different harmonics [38]. Like temperature fluctuations in the cosmic microwave background radiation, which can be mapped to initial state density fluctuations in the early Universe, collective flow fluctuations strongly constrain the initial conditions (initial density distributions) and therefore allows a better measurement of fluid properties. Since 2011, the limit for the shear viscosity has come down by a factor of two (η/s smaller than two to three times $1/4\pi$). It is now precise enough to see a hint of a temperature dependence (slightly increasing from RHIC to LHC) [3], and future improvements in data accuracy and hydro modelling should either further improve the limit, or give a finite value for η/s . In either case, improved precision is relevant as the shear viscosity is directly related to the in-medium cross section and therefore contains information about the degrees of freedom relevant in the sQGP via the strength and temperature dependence of their interactions.

5. Jet quenching

High energy partons interact with the medium and lose energy, primarily through induced gluon radiation and, to a smaller extent, elastic scattering [39]. The amount of energy lost, ΔE , is expected to depend on medium properties, in particular the opacity (density, interaction strength) and the path length L inside the medium, with different models predicting a linear (elastic ΔE), quadratic (radiative ΔE), and even cubic (AdS/CFT) dependence on L . In addition, ΔE also depends on the parton type via the colour charge (quark versus gluon), the parton mass via formation time and interference effects (light versus heavy quarks), and finally somewhat on the jet energy. The total jet energy is of course conserved and the energy lost by the leading parton appears mostly in the radiated gluons, leading in effect to a medium modified softer fragmentation function. Jet quenching (i.e. measuring the medium modified fragmentation functions) is therefore a very rich observable which probes not only the properties of the medium but also properties of the strong interaction.

Jet quenching was discovered at RHIC not with jets, which are difficult to measure in the high multiplicity heavy ion background environment, but as a suppression of high p_T ‘leading’ jet-fragments. While the effect was experimentally very clean and significant with suppression factors up to five, the information was also very limited, impeding a quantitative and model independent determination of matter properties or energy loss mechanisms.

The high energy of LHC and the correspondingly large cross sections for hard processes make high energy jets easily stand out from the background even in central nuclear collisions. Jet quenching is therefore readily recognized and measured, with many unbalanced dijets or even monojets apparent in the data [40]. While the amount of energy lost in the medium can be of order tens of GeV and therefore even on average corresponds to a sizeable fraction of the total jet energy, it is nevertheless close to the one expected when extrapolating RHIC results to the higher density matter at LHC. The two jets remain essentially back-to-back (little or no angular broadening relative to pp) and the radiated energy (ΔE) is found in very low p_T particles ($< 2 \text{ GeV}/c$) and at large angles to the jet direction [41]. The latter two findings were initially a surprise, but are now incorporated naturally into models where the energy is lost in multiple, soft scatterings, and the radiated gluons are emitted at large angles. The parton then leaves the matter and undergoes normal vacuum fragmentation, i.e. looking like a normal pp jet but with a reduced energy.

Additional insight into the energy loss process has come from heavy flavours [42, 43]. Like at RHIC, the suppression of charm mesons is virtually identical to the one of inclusive charged particles, stubbornly refusing to show the difference expected from the stronger coupling (colour charge) of gluons, which are the source of the majority of charged particles, compared to quarks. The mass effect however seems to be as predicted: at intermediate p_T , beauty shows less suppression than charm, whereas at very high p_T ($E/m \gg 1$) b-jets and inclusive jets show similar modifications.

6. Discoveries

The first discovery made at LHC was announced in September 2010⁴ on a subject which was as unlikely as it was unfamiliar to most in the packed audience: the CMS experiment had found a mysterious ‘long range rapidity correlation’ in a tiny subset of extremely high multiplicity pp collisions at 7 TeV [44]. While in the meantime far eclipsed by the discovery of ‘a Higgs-like particle’, this ‘near side ridge’ is arguably still the most unexpected LHC discovery to date and spawned a large variety of different explanations, from mildly speculative to outright weird [45]. The most serious contenders are saturation physics, as formulated in the colour glass condensate model (CGC), and collective hydrodynamic flow. Hydrodynamics is of course a very successful framework to describe long range correlations in the macroscopic hot matter created in heavy ion reactions, but was not supposed to be applicable in small systems like pp collisions, where typically only a few, or at most a few tens, particles are produced per unit of rapidity. The CGC is a ‘first principles’ classical field theory approximation to QCD which is applicable to very dense (high occupation number) parton systems like those found at small- x and small Q^2 in the initial state wave function of hadrons. It has been successfully used to describe some regularities seen e.g. in ep collisions at HERA (geometric scaling) and to model the initial conditions in heavy ion physics.

Lacking further experimental input, no real progress was made to unravel the origin of these long range pp correlations until the ridge made a robust come-back with the first LHC proton–nucleus run some two years later (p-Pb at $\sqrt{s_{NN}} = 5$ TeV [46]). The correlation strength was actually significantly stronger than in pp at the same multiplicity, and in quick succession it was discovered that: the ridge was actually double-sided, showing correlations between particles both close by in azimuth as well as back-to-back; a Fourier analysis revealed both even (v_2) as well as odd (v_3) component; the correlation strength measured with four particles was almost identical to the one measured with two, indicating strongly a collective or at least multi-particle origin; and finally the dependence of the correlation strength on particle mass was virtually identical to the one expected from hydrodynamic flow.

All characteristics of the p-Pb ridge are very natural for and in good agreement with a hydrodynamic collective flow origin of the correlation. Even the strength of the signal and its multiplicity dependence are of the correct order of magnitude (within a factor of two) if one uses some reasonable geometrical initial conditions and a standard hydro model and just postulates that the system, some 1 fm in size and lifetime, behaves like a macroscopic ideal fluid. Just like the matter created in central Pb–Pb collisions, which is of order 5000 fm³ and therefore larger by orders of magnitude!

The question how such a tiny (few fm³) system could thermalize in essentially no time, maybe even become a small serving of sQGP, has kept the case open, despite what looks like considerable evidence. However, the idea of having small sQGP fireballs created even in pp collisions at the LHC may be less radical than it seems on first sight.

At high energy, a proton is qualitatively similar to a small nucleus in some respect; an extended, finite size collection of many (sea)partons which can undergo simultaneous and independent scatterings (called multi-parton interactions in pp , and N_{coll} in AA). Also the final state energy- or particle density (particles per unit volume) in high multiplicity pp is comparable to or even larger than in central Pb–Pb. The thermalization time scale commonly assumed for the sQGP created in nuclear collisions is significantly less than 1 fm c⁻¹ [3]; therefore the local volume required for thermalization is presumably of comparable size, i.e. at most a few fm³! Once equilibrated, the limit on η/s implies that the mean free path in the sQGP is compatible with zero or the Compton wavelength. We do find higher harmonic flow components which originate from density fluctuations with scales of order 1 fm or less. There is therefore actually quite some direct and indirect evidence that in matter with initial energy densities like those produced in high multiplicity pp and pA (and, of course, AA), the 1 fm scale is big enough (and long lived enough) to approach thermal equilibrium and exhibit collective phenomena. The applicability of hydro does not depend on absolute scales (fm or km), but on dimensionless numbers, and e.g. the ratio of system size over mean free path is a big number even in a small (but extremely dense) system like high multiplicity pp !

In any case, the ridge discovery in pp and pA at LHC is definitely more than a curiosity and likely to have profound implications for heavy ion physics, one way or another. If a sQGP (like) state can be created and studied in much smaller systems than anticipated, we can add an ‘extra dimension’, namely size, to our toolbox and compare pp , pA and AA to look for finite size effects, which may reveal information on correlation lengths and relaxation time scales not otherwise easily available⁵. If, on the contrary, initial state effects and saturation physics are the answer, we would have discovered at LHC yet another new state of matter, the CGC, opening a rich new field of activity for both experiment and theory.

7. What is next?

As shown in selected examples above, results have come fast and on a wide range of topics for heavy ion physics during the first three years of LHC Run-1: from subtle, as yet to be fully digested hints (particle ratios) to rather suggestive and clear messages (J/ψ recombination), the LHC has been shining a new and very instructive light on old problems. Some properties of the sQGP have been measured with significantly better precision (η/s , opacity), improving along the way substantially our understanding of the underlying mechanism (jet quenching) or even leading to a paradigm shift (higher harmonic flow) which opened a vast new range of observables to precision experiment and precision theory. And if the strong suspicion that the surprising long range structures in pp and pA collisions are of collective hydro origin turns out to be correct, the ‘new state of matter’ will have once more shown that ‘... it is a state quite different and even more remarkable than had been predicted.’

⁴ Available at <http://indico.cern.ch/conferenceDisplay.py?confId=107440>.

⁵ We also may have to consider changing the name ‘heavy ion physics’.

Some of the main experimental issues which can be addressed with nuclear beams in the upcoming runs at full LHC energy on the short to medium term are fairly clear based on the current results: one could hope to largely complete the measurements needed for understanding quarkonia production as a deconfinement signal and QGP thermometer. This includes quantifying ‘other effects’ from the past (and probably a future, high luminosity) pA run, and reducing the statistical error in particular for the higher mass members of the J/ψ and Υ families as well as midrapidity low p_T J/ψ . Precision and sophistication in the flow analysis should further improve over the coming years, on both experimental and theory fronts, hopefully reaching a precision in e.g. η/s of order 30% or better, at which point the result would be precise enough for quantum corrections to the AdS/CFT lower bound to become relevant. While this is definitely a long shot, its worth aiming at because there are not many alternatives on the horizon for experimental tests of quantum string theory. And last, not least, the ‘ridge puzzle’ may be solved quite soon—e.g. using multi-particle methods to rigorously (dis)prove collectivity—to decide between initial state (CGC) or final state (hydro) origin. If the correlation signal turns out to be the smoking gun for saturation physics, yet another new state of matter will have been discovered at the LHC; a dense, cold, quasi-classical state of gluon matter. And even if hydro is the answer, the search should go on in pA for any other sign of saturation physics. In particular in the forward direction at very low Feynman x (right in the LHCb acceptance), the conditions for saturation phenomena should be just right and proton–nucleus collisions will be the best place to look for them for some time, until an electron–ion collider comes into operation. If in fact a sensitive experimental search for the saturation physics will be carried out at the LHC, the results certainly can make, and possibly could break, the science case for such a future machine.

On a more than 10 year medium- to long-term time scale, which includes experimental and machine upgrades with better detectors and higher luminosity, a comprehensive program can take place to precisely measure and test the various aspects of jet quenching, including e.g. fragmentation functions with the golden γ -jet channel and precision heavy quark measurements down to zero momentum. On this time scale one hopefully can also address signals sensitive to chiral symmetry restoration (e.g. low mass lepton pairs); a defining property of the QGP which is experimentally extremely challenging and has therefore received comparatively less attention.

The exploration of the phases of strongly interacting matter is one of the four main pillars of contemporary nuclear physics, and one should see the LHC ion program in this broader context: the RHIC program is very active and competitive and continues to map the phase diagram at lower temperatures, in addition looking via an energy scan for the transition between normal matter and the sQGP and a ‘tri-critical’ point somewhere in the region at or below SPS fixed target energy. A continuation and strengthening of the SPS fixed target program is under discussion and two new low energy facilities (FAIR at GSI and NICA at JINR) are being built to study compressed matter, i.e. matter at high baryon

density and (comparatively) low temperature where the phase structure may be quite different (first order phase transition) and the matter is closer related to neutron stars than to the early universe. The LHC however is and will be the energy frontier facility not only of high energy physics but also of nuclear physics for the foreseeable future, with a well-defined and extensive program and wish list of measurements. And if the first three years are any guide, strong interaction physics, while firmly rooted in the Standard Model, shows no end to surprises and discoveries and promises to keep physics with heavy ions at the LHC interesting (and fun) for quite some time to come.

References

- As far as available, reviews or overview papers—rather than primary literature—are cited where the reader can find further details and references.
- [1] Muller B, Schukraft J and Wyslouch B 2012 First Results from Pb+Pb collisions at the LHC *Annu. Rev. Nucl. Part. Sci.* **62** 361
 - [2] Schukraft J 2012 Results from the first heavy ion run at the LHC *J. Phys. Conf. Ser.* **381** 012011
 - [3] Müller B 2013 Investigation of hot QCD matter: theoretical aspects *Phys. Scr.* **T158** 014004
 - [4] Schukraft J 2000 *AIP Conf. Proc.* 531, 3
 - [5] Antinori F, Billmeier A and Zaranek J 2003 Introduction to ultrarelativistic heavy-ion physics at the CERN-SPS *AIP Conf. Proc.* **631** 294
 - [6] Satz H 2004 The SPS heavy ion programme *Phys. Rep.* **403–404** 33
 - [7] Heinz U and Jacob M 2000 arXiv:nucl-th/0002042
 - [8] Braun-Munzinger P, Redlich K and Stachel J 2003 *Particle Production in Heavy Ion Collisions: Quark Gluon Plasma* ed R C Hwa *et al* pp 491–599
 - [9] Becattini F 2009 An Introduction to the statistical hadronization model arXiv:0901.3643 [hep-ph]
 - [10] Andronic A, Braun-Munzinger P and Stachel J 2009 Thermal hadron production in relativistic nuclear collisions *Acta Phys. Pol. B* **40** 1005
 - [11] Becattini F, Castorina P, Milov A and Satz H 2010 A comparative analysis of statistical hadron production *Eur. Phys. J. C* **66** 377
 - [12] Letessier J and Rafelski J 2008 *Eur. Phys. J. A* **35** 221
 - [13] Specht H J 2008 Lepton-pair production in nuclear collisions: past, present, future *Nucl. Phys. A* **805** 338
 - [14] Rapp R 2013 Dilepton production in heavy-ion collisions arXiv:1306.6394 [nucl-th]
 - [15] Specht H J (NA60 Collaboration) 2010 *AIP Conf. Proc.* **1322** 1
 - [16] Kluberg L and Satz H 2009 Color deconfinement and charmonium production in nuclear collisions arXiv:0901.3831 [hep-ph]
 - [17] Brambilla N, Eidelman S, Heltsley B K, Vogt R, Bodwin G T, Eichten E, Frawley A D and Meyer A B *et al* 2011 Heavy quarkonium: progress, puzzles and opportunities *Eur. Phys. J. C* **71** 1534
 - [18] Rapp R, Blaschke D and Crochet P 2010 Charmonium and bottomonium production in heavy-ion collisions *Prog. Part. Nucl. Phys.* **65** 209
 - [19] The Brahm's Phenix Phobos Star (Collaborations) 2005 *Nucl. Phys. A* **757** 1–283
 - [20] Jacobs P and Wang X-N 2005 Matter in extremis: ultrarelativistic nuclear collisions at RHIC *Prog. Part. Nucl. Phys.* **54** 443
 - [21] Muller B and Nagle J L 2006 Results from the relativistic heavy ion collider *Annu. Rev. Nucl. Part. Sci.* **56** 93
 - [22] Schukraft J 2012 *Phil. Trans. R. Soc. A* **370** 917

- [23] Stock R 2007 Hadron formation in high energy elementary and nuclear collisions *Int. J. Mod. Phys. E* **16** 687
- [24] Andronic A, Braun-Munzinger P and Stachel J 2006 Hadron production in central nucleus–nucleus collisions at chemical freeze-out *Nucl. Phys. A* **772** 167
- [25] Armesto *et al* 2008 Heavy ion collisions at the LHC—last call for predictions *J. Phys. G: Nucl. Part. Phys.* **35** 054001
- [26] Oeschler H (ALICE Collaboration) 2011 arXiv:1102.2745 [hep-ex]
- [27] Andronic A, Braun-Munzinger P, Redlich K and Stachel J 2013 *Nucl. Phys. A* **904–905** 535c
- [28] Becattini F, Bleicher M, Kollegger T, Schuster T, Steinheimer J and Stock R 2013 *Phys. Rev. Lett.* **111** 082302
- [29] Petran M, Letessier J, Petracek V and Rafelski J 2013 arXiv:1303.2098 [hep-ph]
- [30] Braun-Munzinger P and Stachel J 2009 arXiv:0901.2500 [nucl-th]
- [31] Abelev B *et al* (ALICE Collaboration) 2012 *Phys. Rev. Lett.* **109** 072301
- [32] Chatrchyan S *et al* (CMS Collaboration) 2011 *Phys. Rev. Lett.* **107** 052302
- [33] Voloshin S A, Poskanzer A M and Snellings R 2010 Collective phenomena in non-central nuclear collisions *Landolt-Boernstein: Relativistic Heavy Ion Physics* vol 1/23 (Berlin: Springer) pp 5–54 (arXiv:0809.2949 [nucl-ex])
- [34] Schutz Y and Wiedemann U A 2011 Quark matter *Proc. 22nd Int. Conf. on Ultra-Relativistic Nucleus–Nucleus Collisions, Quark Matter 2011* (Annecy, May 23–28, 2011); *J. Phys. G: Nucl. Part. Phys.* **38** 120301
- [35] Nagle J L 2009 Ridge, Bulk and Medium Response: How to kill models and learn something in the process *Nucl. Phys. A* **830** 147C
- [36] Aad G *et al* (ATLAS Collaboration) 2013 arXiv:1305.2942 [hep-ex]
- [37] Schukraft J, Timmins A and Voloshin S A 2013 *Phys. Lett. B* **719** 394
- [38] Qiu Z and Heinz U 2012 *Phys. Lett. B* **717** 261
- [39] Majumder A and Van Leeuwen M 2011 The theory and phenomenology of perturbative QCD based jet quenching *Prog. Part. Nucl. Phys. A* **66** 41
- [40] Aad G *et al* (ATLAS Collaboration) 2010 *Phys. Rev. Lett.* **105** 252303
- [41] Chatrchyan S *et al* (CMS Collaboration) 2011 *Phys. Rev. C* **84** 024906
- [42] Renk T 2013 Jet quenching and heavy quarks arXiv:1309.3059 [hep-ph]
- [43] Dainese A 2013 Heavy-quark production in heavy-ion collisions *J. Phys: Conf. Ser.* **446** 012034
- [44] Khachatryan V *et al* (CMS Collaboration) 2010 *J. High Energy Phys.* **JHEP09(2010)091**
- [45] Li W 2012 Observation of a ‘ridge’ correlation structure in high multiplicity proton–proton collisions: a brief review *Mod. Phys. Lett. A* **27** 1230018
- [46] Loizides C 2013 First results from p-Pb collisions at the LHC arXiv:1308.1377 [nucl-ex]