

Predictions and postdictions for relativistic lead and oxygen collisions with the computational simulation code TRAJECTUM

Govert Nijs¹ and Wilke van der Schee²

¹*Center for Theoretical Physics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*

²*Theoretical Physics Department, CERN, CH-1211 Genève 23, Switzerland*



(Received 23 November 2021; revised 8 September 2022; accepted 12 September 2022; published 12 October 2022)

We introduce a global analysis of relativistic heavy ion collisions using TRAJECTUM of a significantly higher precision and including a new option to vary the normalization of the centrality estimator. We use the posterior distribution of our parameters to generate a set of high-statistics samples that allows us to make precise predictions including statistical and systematic uncertainties estimated from the model parameter distribution. The results are systematically compared with experiment whereby we also include many observables not included in the global analysis. This includes in particular (extremely) ultracentral anisotropic flow and mean transverse momentum, whereby we find satisfactory agreement with experiment where data are available. Lastly, we compute spectra and anisotropic flow for oxygen-oxygen collisions performed at the BNL Relativistic Heavy Ion Collider (RHIC) and to be performed at the the CERN Large Hadron Collider (LHC) and comment on how these collisions may inform us on properties of the quark-gluon plasma.

DOI: [10.1103/PhysRevC.106.044903](https://doi.org/10.1103/PhysRevC.106.044903)

I. INTRODUCTION

The collisions of heavy ions at the BNL Relativistic Heavy Ion Collider (RHIC) at Brookhaven, USA and the Large Hadron Collider (LHC) at CERN have led to an accepted picture of a short prehydrodynamic phase followed by a relativistic fluid composed of quark-gluon plasma (QGP) with remarkably small shear viscosity and finally a cascade of interacting hadrons [1–5]. An increasingly precise picture of these phases has recently emerged, including an accurate estimate of the temperature-dependent shear viscosity [6–12]. Nevertheless, due to the complicated interplay of the various stages and corresponding models it has remained challenging to advance our quantitative understanding of heavy ion collisions.

There are multiple outstanding grand questions in the field of heavy ion collisions [3], which range from the limits of hydrodynamics to the behavior of energetic quarks and gluons as a function of energy scale and coupling constant. Crucial questions are the nature of the prehydrodynamic phase (both its lifetime and dynamics [5]), the initial shape of the energy deposited by the colliding nuclei [13], the behavior of hydrodynamics when viscous corrections are (too?) large [14] or even if hydrodynamics can still make sense if only a few dozen particles per unit rapidity are produced. All these questions shed insight on a microscopic description based on QCD

and can hence, e.g., inform how a strongly coupled QGP can arise from a theory of quarks and gluons which are free at asymptotically high energies.

Given our current state of understanding it is often difficult to shed light on these often quantitative questions by studying a single aspect and linking it to a single measurement. Instead, for instance, both the viscosities and the initial shape characteristically alter the (anisotropic) momentum distribution of the final-state particles. The initial shape and initial momentum anisotropies furthermore depend on the prehydrodynamic far-from-equilibrium stage. It is hence clear that only a global analysis taking into account all these questions and a representative set of experimental observables can put strong quantitative constraints on individual parameters such as the shear viscosity over entropy density ratio.

The aim of the current work is twofold. First, we would like to assess our current understanding of heavy ion collisions, including the systematic uncertainties that necessarily arise from modeling its different aspects. Modeling uncertainties are especially well addressed by studying the robustness of the model under widening of the scope of the model, such that the model has the flexibility to be closer to reality. At first it may not be obvious why including more model parameters can help understanding systematic uncertainties or test the robustness of the model. Indeed, just adding model parameters would risk having as many parameters as (effective) data points, such that the model is basically a phenomenological fit to the data. On the other hand, it is clear that if a physical model parameter (say the slope of the specific viscosity η/s as a function of temperature) is artificially constrained to a specific value that does not correspond to physical reality this will worsen the model. In such a case the model is not complete and this fact has to be taken into account in a larger theoretical

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

systematic uncertainty. The easiest method to assess such a systematic uncertainty, however, is to include the parameter and vary it within the model. It is in this sense that adding parameters sheds light on the robustness and uncertainty of the model.

Recently, this approach of widening the scope of the model has been implemented for the initial shape [7], the initial stage [9], allowing the variation of second-order transport coefficients [9,12], widening the parametrization of the shear and bulk viscosities [12] and studying several different viscous particlization prescriptions [11,15]. In this work we also add one parameter (the centrality normalization), but instead we focus on the existing model and perform a high-statistics PbPb analysis of the model output at twenty randomly drawn model parameters from the posterior. This not only allows us to go beyond the experimental data used in the global analysis, but also allows an analysis of the systematic uncertainties due to our uncertainties in the model parameters. Discrepancies with the detailed experimental comparison then indicate a lacuna in our current understanding.

A second aim is to highlight smaller collision systems [4]. For this we do a similar analysis, but now for oxygen-oxygen (OO) collisions [16]. In May 2021, RHIC performed a 16 day run with 402M OO collisions at a nucleon-nucleon collision energy of $\sqrt{s_{NN}} = 0.2$ TeV and LHC has a special run planned likely in 2024 at $\sqrt{s_{NN}} = 6.8$ TeV. No analyzed data is yet available and hence, in contrast with the detailed PbPb analysis (mostly postdictions), the results of the OO analysis will include only predictions. Moreover, OO collisions are right at the threshold where a hydrodynamic QGP may well be present, but significant nonhydrodynamic effects are also expected. Predictions can hence inform us on which aspects are consistent with a hydrodynamic QGP.

Our analysis is performed using TRAJECTUM, which is a new computational code [9,10] to simulate heavy and light ion collisions that is entirely self-contained (with the exception of the hadronic gas phase) and includes one of the most versatile models for both the initial stage (based on the TRENTo model [6]) and a hydrodynamic phase including temperature-dependent first-order transport coefficients and all second-order coefficients. Importantly, it comes with an analysis routine that is fast and fully parallelizable, thereby allowing to analyze an unlimited number of events with methods that are designed to be as close to experimental measurements as possible.

In this work we present several improvements with regard to previous work. First of all, the simulation of TRAJECTUM in this work now has improved precision in its emulation of heavy ion collisions as a function of the parameter space whereby emulation uncertainties are now competitive or even significantly smaller than the corresponding experimental uncertainties. Second, we introduced a new parameter rescaling the anchor point (AP) for the 100% centrality point. This has direct consequences for all noncentral collisions and we show a significant change in the posterior of for instance the nucleon width.

Lastly, as stated we present a multitude of predictions (or postdictions) for lead-lead (PbPb) collisions at the LHC as well as oxygen-oxygen (OO) collisions at both RHIC and

the LHC. These predictions include a systematic uncertainty analysis from the full posterior distribution as well as statistical uncertainties for observables where this is relevant. Especially interesting are the (extremely) ultracentral results (down to 0.001% centrality), where measurements are mostly not available yet. For the OO results we present estimates how the upcoming experimental results can help in improving our understanding of the QGP.

II. MODEL

For the heavy ion simulations presented in this work, we use version 1.1 of the TRAJECTUM code [17]. As TRAJECTUM is described in some detail in Ref. [10], we will only give a brief outline here, highlighting the changes to the version of TRAJECTUM presented in Ref. [10]. TRAJECTUM allows the user to choose different models for the initial conditions, hydrodynamics, and transport coefficients, and also offers a variety of partial differential equation solvers for the hydrodynamical evolution. As in Ref. [10] we use the TRENTo model for the initial conditions, modified with a free streaming velocity that can differ from the speed of light. For hydrodynamics we use the 14-moment approximation, with transport coefficients that allow for temperature dependence in the shear and bulk viscosity over entropy density ratios, as in Ref. [8]. The hydrodynamical evolution is solved using the “fast” version of the MUSCL algorithm [10].

This leaves us with the following set of parameters that are varied in this work. From the boost invariant TRENTo initial stage, we have the minimal internucleon distance d_{\min} , the number of constituents n_c , and the nucleon width w . We also have the constituent width v , which is parametrized by the parameter χ_{struct} , in terms of which v is defined as follows:

$$v = v_{\min} + \chi_{\text{struct}}(w - v_{\min}),$$

where we fix $v_{\min} = 0.2$ fm. The constituents in each colliding nucleon each then source a thickness function with a Gaussian distribution, where for each individual constituent the norm of the Gaussian is given by $N\gamma/n_c$. Here γ is sampled from a gamma distribution with width $\sigma_{\text{fluct}}\sqrt{n_c}$, and N and σ_{fluct} are parameters. In this way, for each nucleus a thickness function is computed, \mathcal{T}_A and \mathcal{T}_B , respectively. The two thickness functions \mathcal{T}_A and \mathcal{T}_B so obtained are then combined to form the initial energy density e according to

$$e\tau = \left(\frac{\mathcal{T}_A^p + \mathcal{T}_B^p}{2} \right)^{1/p}, \quad (1)$$

with $\tau = 0^+$ the initial proper time and p another parameter.

The prehydrodynamic stage performs free streaming with a velocity v_{fs} until proper time τ_{fs} , which are the two parameters governing this stage. For the hydrodynamic stage the equation of state is a hybrid constructed from the HotQCD lattice equation of state, together with the hadron resonance gas (HRG) constructed from the chosen particle content [18–20]. We use a temperature dependent shear viscosity over entropy density η/s that is described by the parameters $(\eta/s)_{\min}$,

$(\eta/s)_{\text{slope}}$, and $(\eta/s)_{\text{crv}}$ through the formula

$$\eta/s = (\eta/s)_{\text{min}} + (\eta/s)_{\text{slope}}(T - T_c) \left(\frac{T}{T_c} \right)^{(\eta/s)_{\text{crv}}},$$

for $T > T_c$, and $\eta/s = (\eta/s)_{\text{min}}$ otherwise, with T_c fixed to 0.154 GeV. The analogous parameters for the bulk viscosity ζ are $(\zeta/s)_{\text{max}}$, $(\zeta/s)_{\text{width}}$, and $(\zeta/s)_{T_0}$, through the formula

$$\zeta/s = \frac{(\zeta/s)_{\text{max}}}{1 + \left(\frac{T - (\zeta/s)_{T_0}}{(\zeta/s)_{\text{width}}} \right)^2}.$$

In addition, we vary the shear relaxation time τ_π by setting $\tau_\pi s T / \eta$ as a temperature independent parameter, and we vary the second-order transport coefficient $\tau_{\pi\pi}$ by the parameter $\tau_{\pi\pi} / \tau_\pi$. The final parameter that we vary is the switching temperature T_{switch} , which enters in the particlization procedure (following Ref. [6]). In the following, we will describe the changes that we made to the model for this work.

A. Oxygen nuclei

To simulate a Pb nucleus, we use a Woods-Saxon distribution with parameters as given in Refs. [8,21], while enforcing a minimal distance d_{min} between nucleons. For a large spherical nucleus such a model is appropriate, but for smaller nuclei like O such an approach fails to capture important correlations between nucleon positions. Indeed, the 16 nucleons in oxygen tend to group together in four α particles, which potentially has important consequences for the shape of the geometry. For this reason, for O we randomly rotate a sampled oxygen nucleus from a list of 6000 O configurations computed using effective two- and three-nucleon potentials, which is able to capture these correlations [22–24]. The difference between a standard Woods-Saxon distribution and the α -clustered version has recently been studied in Refs. [24,25].

B. Continuous number of constituents

In the TRENTo model with substructure [7,26], the number of constituents n_c is a discrete parameter. In a Bayesian analysis, this can potentially cause problems for the emulator, since the emulator essentially interpolates between different design points assuming that the quantities it emulates depend continuously on the parameters. However, since n_c is discrete, in the TRENTo model the dependence on n_c is more like a stepwise defined function. Especially from $n_c = 1$ to $n_c = 2$ the effect on certain observables can be significant, which can cause the emulator to be less accurate.

In this work, we resolve this issue by changing the number of constituents to a continuous parameter. Each sampled nucleon either uses $[n_c]$ or $\lceil n_c \rceil$ constituents with probability such that the average number of constituents equals n_c . The definition agrees with the old definition for integer values of n_c , but for a large number of events this definition indeed leads to observables depending continuously on n_c .

C. URQMD and SMASH

Whereas [10] used SMASH [27–29], in this work we use URQMD [30,31] as a hadronic afterburner. These codes are

both state of the art, and are hence both well suited to be used in conjunction with TRAJECTUM. Both codes produce similar results [12], whereby URQMD leads to somewhat higher multiplicities (about 1%) and higher $\langle p_T \rangle$ for kaons and protons (about 1.5%). There is no $\langle p_T \rangle$ increase for pions, but low- p_T yields (< 0.4 GeV) are enhanced by about 1.5%, as are high- p_T (> 1.5 GeV) yields. The URQMD elliptic flow is also slightly higher (1%) across all centrality classes.

For the PbPb results presented in this work, it will turn out that the fitted free streaming velocity v_{fs} is substantially smaller than the previous result obtained in Refs. [9,10] (this can compensate for higher $\langle p_T \rangle$). The fitted minimum shear viscosity over entropy density ratio $(\eta/s)_{\text{min}}$ is slightly higher than the result obtained in Refs. [9,10] (this can compensate for the larger elliptic flow). Even though the differences between URQMD and SMASH seem mild they are indeed still the likely cause of the slightly changed posterior distributions, especially as the mean transverse momentum is a sensitive observable that is known experimentally to high precision (see also Fig. 5 later).. This probably indicates that there is some work remaining to be done to understand the correct way to describe the final stages of the collision, but at least between URQMD and SMASH the discrepancies are quite mild.

D. Centrality determination

Historically the determination of centrality classes has not always been clear. Theoretically often the impact parameter of the collision was used as a probe for centrality, with smaller impact parameter giving more central collisions. Fluctuations in geometry and multiplicity give this method a significant bias and it is therefore much better to order centrality by the final multiplicity (ordering by initial entropy also gives reasonable results [32]). TRAJECTUM has the option to specify several centrality definitions with different pseudorapidity η and transverse momentum p_T cuts to make these as close as possible to the actual experimental method. In this work we always work with the standard settings $|\eta| < 2.4$ and $p_T > 400$ MeV/c unless otherwise noted.

Experimentally the centrality ordering is usually defined by the number of tracks at forward rapidity (for ALICE the V0A and V0C detectors located at pseudorapidity $2.8 < \eta < 5.1$ and $-3.7 < \eta < -1.7$, respectively). This can lead to a problem comparing with theoretical models, which often work under the assumption of boost invariance and hence only compute multiplicities at midrapidity. Nevertheless, especially for PbPb collisions multiplicities at midrapidity are rather strongly correlated with the V0 amplitudes (see Fig. 15 in Ref. [33]) and unless one looks at rather extreme centralities the methods should match fairly well. An example of this is explored in Fig. 16. For pPb collisions this is, however, completely different and self-correlations in the centrality determination can be important [34].

A single crucial question remains though: when do we determine that a collision happened? Experimentally this is far from easy: first of all, the PbPb cross section is completely dominated by electromagnetic interactions of both charged nuclei. These interactions lead to the dissociation of a single (187 ± 12 b) or both (5.7 ± 0.4 b) Pb nuclei [35], which

is much larger than the estimated hadronic cross section of 7.7 ± 0.6 b. It is also necessary to correct for collisions that are hadronic but do not leave any tracks in the detector acceptance. Alternatively, it is possible to fit the multiplicity distribution with a MC Glauber model and take the 100% anchor point as given by the model, but this introduces a theoretical modeling uncertainty. Nevertheless, both methods agree well, and the final systematic uncertainty is less than 3% [33].

Also theoretically it is not quite clear when a hadronic collision actually happened. This is usually determined by a Monte Carlo Glauber model that puts nucleons at fluctuating positions and then decides in a collision whether two nucleons collide [36,37]. A nucleus-nucleus collision then occurs if at least a single nucleon-nucleon interaction happened and hence at least have two “participants.” A crucial ingredient herein is the nucleon-nucleon cross section σ_{NN} , which also features as a TRAJECTUM parameter and is usually fixed by the measured value in pp collisions (at high energies the difference in cross section between protons and neutrons is negligible). Nevertheless, several options are possible to satisfy the σ_{NN} cross section, and a black-disk approximation (as in the standard MC Glauber [23,38]) does not lead to the same result as the TRENTO model with nucleons of size w used by TRAJECTUM. Since here we work at fixed nucleons with smaller w can interact with high probability, but this depends strongly on the impact parameter. Nucleons with large w on the other hand often do not interact even if they collide at zero impact parameter (see Ref. [6] for full details).

In TRAJECTUM (as in TRENTO) it is possible that a nucleon-nucleon interaction happened, but that the switching temperature is not reached anywhere and no thermal particles are produced (since the particlization is a Poisson process this may even happen if T_{switch} is reached). Nevertheless, since a hadronic interaction happened particles realistically should be produced in a more complete model that takes the QGP corona into account and hence TRAJECTUM does include these events in its event list as having zero particles.

While both the theoretical and experimental choices and procedures are well motivated and studied still some significant uncertainty remains and also to test the robustness of the model we decided to introduce a new parameter $\text{cent}_{\text{norm}}$ that rescales the centrality classes. This parameter rescales the α - β percent centrality so that in the model we actually compute the $0.01\alpha\text{cent}_{\text{norm}}-0.01\beta\text{cent}_{\text{norm}}$ percent centrality bin. In other words, we perform the computation as if 100% experimental centrality corresponds to $\text{cent}_{\text{norm}}$ centrality in our model.

By including $\text{cent}_{\text{norm}}$ as a parameter, the uncertainty coming from this issue can be properly taken into account in the Bayesian analysis. That way, we can see which parameters correlate with $\text{cent}_{\text{norm}}$, and we are also able to fit $\text{cent}_{\text{norm}}$. It is however an interesting question whether in a way we may be double-counting uncertainties here, since the experimental measurements already contain a contribution in their uncertainty due to the centrality selection. Since indeed we will find $\text{cent}_{\text{norm}}$ to be consistent with 100% (see Sec. III B) this could motivate to reinstate $\text{cent}_{\text{norm}} = 100\%$ in future analyses.

Important for the discussion in Sec. II C is that $\text{cent}_{\text{norm}}$ is a parameter which enters purely in collecting the results of the (parallelized) analysis of all events, not the simulation or analysis. This means that it can easily be turned off at little extra computational cost, so we have been able to verify that the difference between this work and that obtained in Ref. [10] is not due to the inclusion of $\text{cent}_{\text{norm}}$ (see also Fig. 3 later).

III. PRECISE PbPb RESULTS AND OO PREDICTIONS

A. The design and emulator validation

In this work, we perform a Bayesian analysis of the PbPb system and subsequently make (predominantly) postdictions for PbPb and predictions for the OO system. In the end we also examine how much the precision of the model parameters can improve by the inclusion of future OO data. To make this as worthwhile as possible, we want to make the Bayesian analysis of both the PbPb and OO systems as precise as we can.

The experimental dataset that we fit to is the same as in Ref. [9], except that in this work we only fit to PbPb observables. In particular, we include charged particle multiplicity $dN_{\text{ch}}/d\eta$ at 2.76 [39] and 5.02 TeV [40] and transverse energy $dE_T/d\eta$ at 2.76 TeV [41]. We also include identified yields dN/dy and mean p_T for pions, kaons, and protons at 2.76 TeV [42], as well as p_T fluctuations at 2.76 TeV [43]. For integrated anisotropic flow, we include $v_2\{2\}$, $v_2\{4\}$, $v_3\{2\}$, and $v_4\{2\}$ at both 2.76 and 5.02 TeV [44]. We also include p_T -differential observables, in six coarse-grained p_T bins with bin boundaries at (0.5, 0.75, 1.0, 1.4, 1.8, 2.2, 3.0) GeV. With these p_T bins, we include identified transverse momentum spectra for pions, kaons, and protons at 2.76 TeV [42], as well as the p_T -differential identified anisotropic flow coefficient $v_2\{2\}(p_T)$ for pions, kaons, and protons at 2.76 TeV, and $v_3\{2\}(p_T)$ for pions [45]. For the p_T -differential anisotropic flow, we only use bins with good statistical uncertainties, which are the same as those used in Ref. [9] and indicated explicitly later in Fig. 2.

The starting point of any Bayesian analysis is Bayes posterior probability density for our parameters

$$\mathcal{P}(\mathbf{x}|\mathbf{y}_{\text{exp}}) = \frac{e^{-\Delta^2/2}}{\sqrt{(2\pi)^n \det(\Sigma(\mathbf{x}))}} \mathcal{P}(\mathbf{x}), \quad (2)$$

with $\mathcal{P}(\mathbf{x})$ the (flat) prior probability density and where

$$\Delta^2 = [\mathbf{y}(\mathbf{x}) - \mathbf{y}_{\text{expt}}] \cdot \Sigma(\mathbf{x})^{-1} \cdot [\mathbf{y}(\mathbf{x}) - \mathbf{y}_{\text{expt}}], \quad (3)$$

with $\mathbf{y}(\mathbf{x})$ being the predicted data for parameters \mathbf{x} , \mathbf{y}_{expt} the n experimental data points, and $\Sigma(\mathbf{x})$ is the sum of the experimental and theoretical covariance matrices. In principle the full experimental matrix should be provided by the relevant experiment, but as this is rarely available we follow the simple prescription from Ref. [20]. Obtaining the posterior distributions for our parameters involves a Markov chain Monte Carlo (MCMC) [46] that evaluates the model many times by use of a Gaussian Process emulator [20,47,48]. The uncertainty of the emulator is estimated by the emulator itself and included in Σ . For this analysis we use emulators only for the 50 most

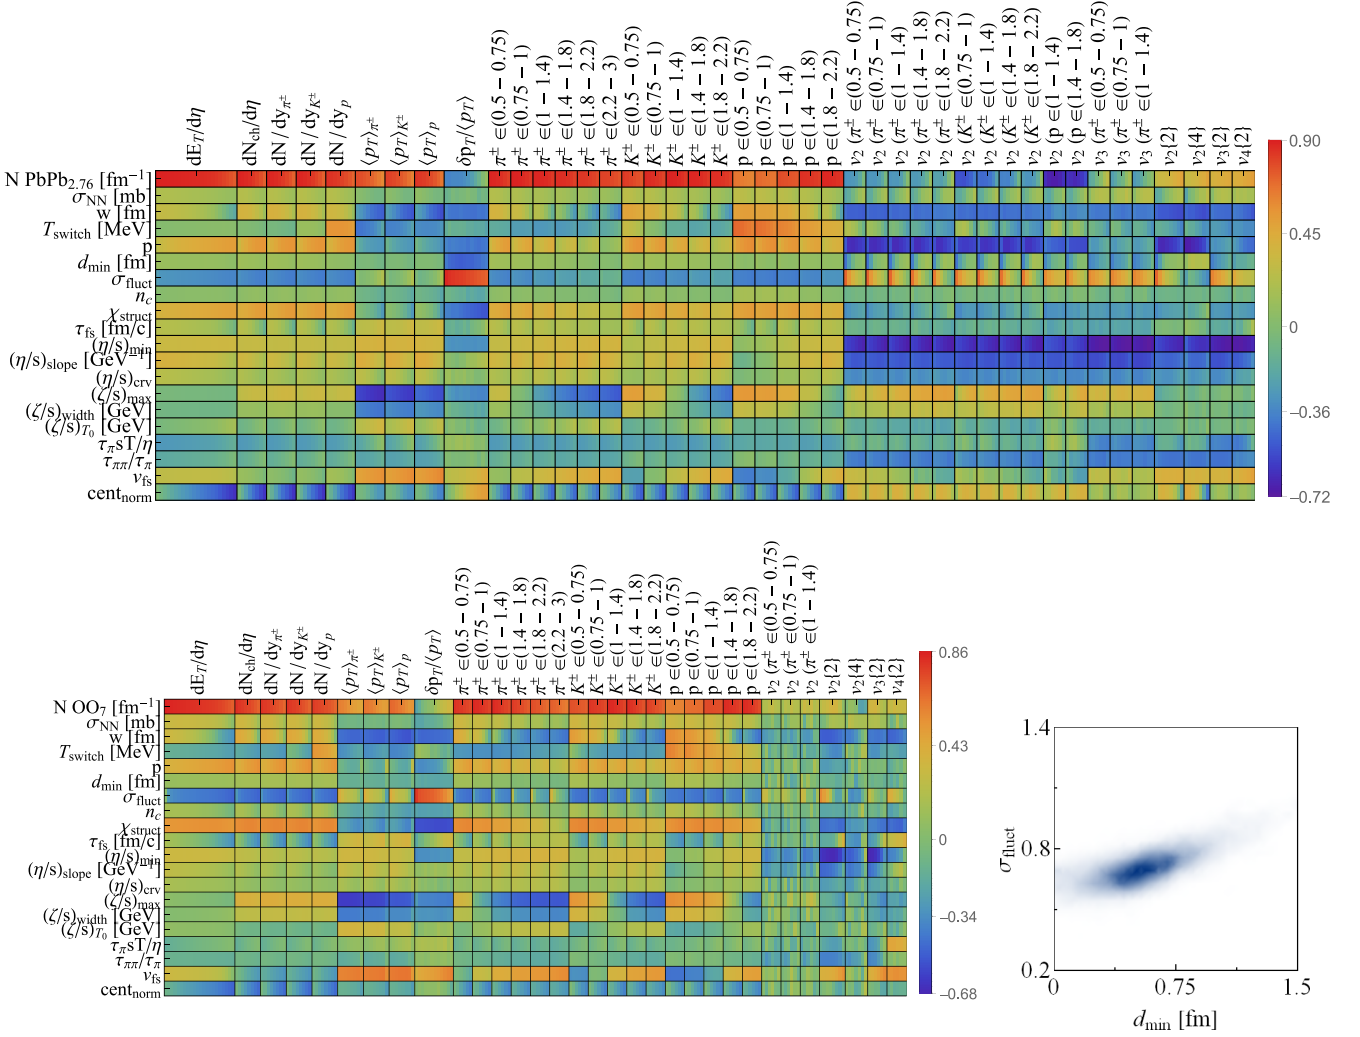


FIG. 1. We show the correlation between our parameters and all observables used for PbPb (top) and OO (bottom left) collisions. Inside each box different centralities are shown from most central on the left to most peripheral on the right. p_T bins are displayed in units of GeV/ c . (bottom right) The figure shows the posterior correlation between d_{min} and σ_{fluct} following from the posterior analysis done in Sec. III B.

significant principal components (PCs) of the experimental data.

A crucial element in the emulator of a particular PC is the characteristic length scale it assigns to a parameter: the emulator uses shorter length scales when the PC is more sensitive to the parameter. This means it has to rely on fewer design points and hence typically has a larger uncertainty. By restricting the prior ranges of our parameters it is possible to make PCs less sensitive (on the domain) and hence a more precise emulator is expected. The price to pay is of course the smaller domain of the prior range, but from the previous work [10] it was possible to tighten the prior range significantly without affecting our conclusions (our prior ranges can be seen later in Fig. 3). Compared with [10], we furthermore lowered the statistical uncertainty by increasing the number of events per design point from 6k to 15k (for clarity we note that this work does not oversample hydrodynamic events into several hadronic events). Given the increase in computation time needed for this number of events per design point, we lower the number of design points from 1000 to 750 so that the total computation

time remains manageable. Also, as discussed in Sec. II B, by making the number of constituents a continuous parameter, we no longer present the emulator with discontinuities. For OO collisions we ran the model at 1500 design points with 40k events each.¹

One of the first interesting outputs to look for in the result of all 750 design points is the correlation between the various parameters and the observables, which is shown in Fig. 1 for both PbPb (top) and OO (bottom). It can clearly be seen that total transverse energy, mean transverse momenta, and (identified) (p_T -differential) particle yields correlate positively with the norm. This makes sense, as the norm directly controls how much energy is deposited into the initial state.

¹The total computational time for the analysis was around 1M CPU hours, roughly equally divided over PbPb and OO collisions. For PbPb collisions half of the time was spent on UrQMD (the hadronic afterburner, we use an η cut of 2.5 for this analysis), whereas for smaller OO systems almost all time was used by TRAJETUM.

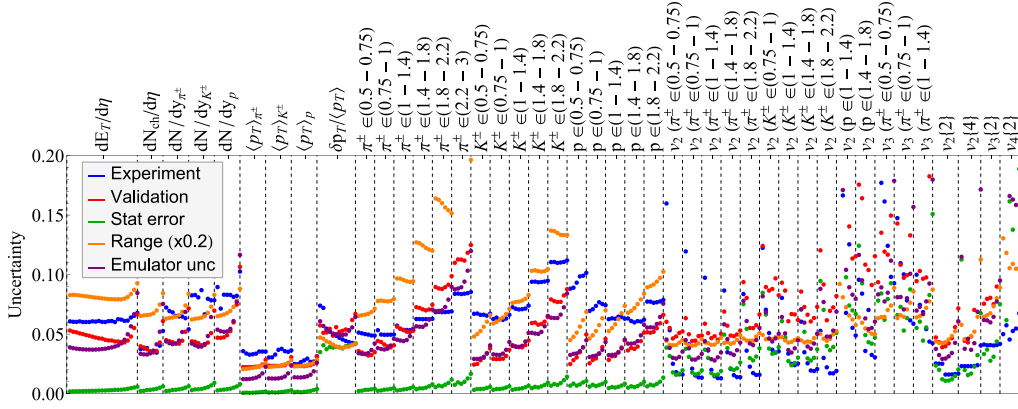


FIG. 2. For the same observables as in Fig. 1 we show as a fraction the experimental uncertainty, the validation of the emulator, the average statistical uncertainty of the model, the range of the model as well as the uncertainty reported by the emulator itself.

As is well known, the shear viscosity reduces all anisotropic flow coefficients $v_n\{k\}$ and the bulk viscosity suppresses the mean transverse momentum. As expected a higher switching temperature increases proton yields. Perhaps unsurprisingly, the per-nucleon entropy deposition fluctuations σ_{fluct} correlate strongly with fluctuations of mean transverse momentum $\delta p_T / \langle p_T \rangle$. Less obvious is that σ_{fluct} also correlates positively with the most central bin of the anisotropic flow coefficients, to which we will come back in Sec. IV. The free streaming velocity v_{fs} is also correlated with the mean transverse momentum, allowing both $(\zeta/s)_{\text{max}}$ and v_{fs} to adjust to fit this observable. However, whereas in PbPb v_{fs} does not have a correlation with p_T -differential anisotropic flow, $(\zeta/s)_{\text{max}}$ is moderately correlated, explaining why the inclusion of p_T -differential observables in Refs. [9,10] led to a substantially lower value for $(\zeta/s)_{\text{max}}$ than earlier work [6,8,49]. The new parameter $\text{cent}_{\text{norm}}$ is clearly more important for peripheral collisions, especially for the yields. Note however, that for especially v_2 the centrality normalization is (relatively) important across the entire centrality range, which is explained by the observation that v_2 is strongly centrality dependent towards central collisions.

A few parameters only have strong correlations with very few observables. A good example is d_{min} which strongly influences the transverse momentum fluctuations. This is however more heavily influenced by σ_{fluct} , so that we will indeed only find very mild constraints on d_{min} , albeit stronger constraints on the correlation between d_{min} and σ_{fluct} (see Fig. 1). Lastly, a few parameters have only mild correlations, such as τ_τ and n_c . This will mean that likely high precision will be needed to get constraints on those parameters. Note, however, that these figures can still hide important correlations when, e.g., two parameters interplay with each other.

The correlations for the OO figure show similar characteristics as the PbPb result, albeit likely showing higher statistical imprecision for the p_T -differential v_2 results. It is however clear that OO collisions due to their smaller size are more sensitive to the initial stage parameters such as v_{fs} and τ_{fs} .

A lower bound for the achievable precision in a Bayesian analysis is given by the experimental uncertainties on the data used in the analysis. Often, however, the theoretical

emulation uncertainty is considerably higher. In Fig. 2, we show the validation of the emulator in this work, which one can compare with Fig. 4 of Ref. [10]. One can see that both the statistical and emulator uncertainty have improved in this work, and for many observables the emulator uncertainty (which incorporates the statistical uncertainty as well) is now smaller than the experimental uncertainty, or comparable in size. This indicates that we are close to the most precise results which are achievable within the current model given current experimental uncertainties and the set of observables that we fit to. One potential issue is that the emulator's self-reported uncertainty (indicated in purple) tends to be lower than the error computed by comparing the emulator prediction to model computations (indicated in red). This indicates that the emulator may be slightly underestimating the true theoretical error. Compared with Ref. [10] the ranges of the observables are also decreased, which reflects the smaller prior ranges on our parameters. Nevertheless, given the previous analysis we were able to do so in a manner that still includes good coverage over the experimental data.

B. Posterior distribution

In Fig. 3, we show the posterior distributions for the parameters from (2) evaluated using the parallel tempered [50] emcee code [46] with three temperatures and 300 walkers applied to experimental data from PbPb collisions at 2.76 and 5.02 TeV. In particular, two different posteriors are shown. The solid blue lines correspond to a posterior in which $\text{cent}_{\text{norm}}$ was varied as a parameter, whereas in the dashed red lines $\text{cent}_{\text{norm}}$ is fixed to 100% at the level of the TRAJECTUM analysis. It can immediately be seen that the preferred value of $\text{cent}_{\text{norm}}$ is close to 100%, with a relatively small uncertainty. This indicates that, within the posterior uncertainties, 100% centrality in our model corresponds well with the experimental notion of 100% centrality. Note however, that the posterior distributions for some observables, such as w or $(\eta/s)_{\text{crv}}$, change substantially if $\text{cent}_{\text{norm}}$ is varied, indicating that small discrepancies between the model and experimental centrality determination can have a substantial effect on the posteriors of those parameters.

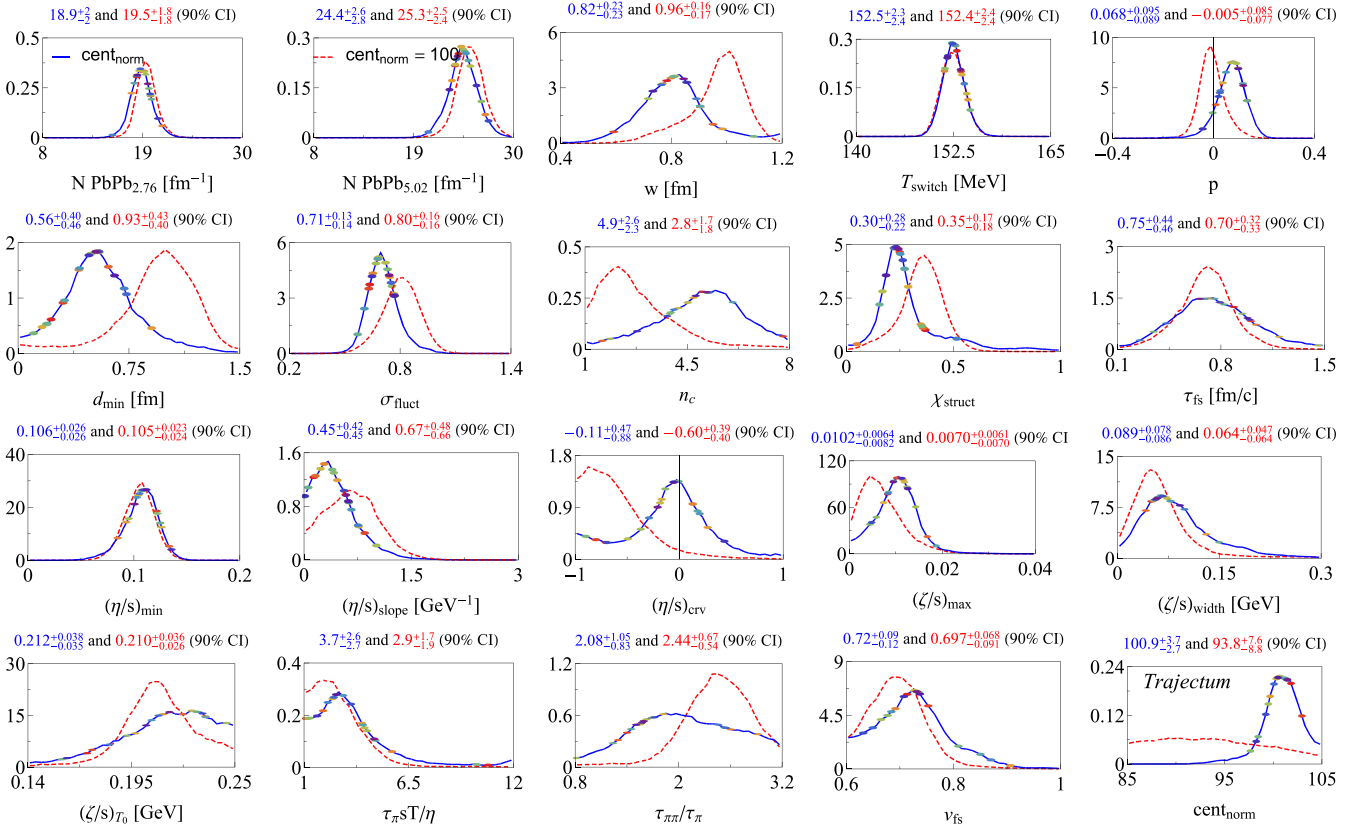


FIG. 3. We show MCMC posterior distributions for our parameters fit to PbPb collisions at collision energies of 2.76 and 5.02 TeV both with (solid blue) and without (dashed red) a varying centrality normalization $\text{cent}_{\text{norm}}$. From the distribution including variations in $\text{cent}_{\text{norm}}$ we randomly selected 20 sample points (indicated by colored dots) for which postdictions are shown in Sec. III C.

It is important to note that the red dashed line of $\text{cent}_{\text{norm}}$ is not completely flat, even though none of the observables depend on $\text{cent}_{\text{norm}}$ (it was fixed to 100%). Nevertheless, the emulator still associates the 750 design points to a given $\text{cent}_{\text{norm}}$ and attempts to find correlations between $\text{cent}_{\text{norm}}$ and the observables. Of course we could have explicitly told the MCMC analysis that our emulator does not depend on $\text{cent}_{\text{norm}}$ and that it hence could use a flat prior and posterior. In this work we chose not to do this to illustrate this effect and also to verify that indeed the MCMC does not exclude any of the $\text{cent}_{\text{norm}}$ prior distribution as should be expected. Also, for a flat prior and an ideal emulator we would expect to find the 90% confidence interval to be given by 95 ± 9 , which is indeed very close to the result of $93.8^{+7.6}_{-8.8}$.

Also, compared with Fig. 1 from [9], the posteriors of many of the parameters have become more constrained, due to the decreased statistical and emulator errors. This can be seen in some detail in Table I, where we compare the posterior including a varying $\text{cent}_{\text{norm}}$ from this work to the combined PbPb and p Pb posterior from Ref. [9]. Here we define the improvement I between the two works as

$$I = \frac{\sigma_{[9]}}{\sigma_{\text{this work}}} - 1,$$

with σ the 90% confidence interval from each respective work. One can see that the precision in the determination of both norms $N \text{ PbPb}_{2.76}$ and $N \text{ PbPb}_{5.02}$ improves substantially

in this work, as well as the switching temperature T_{switch} and the three parameters relating to the shear viscosity over entropy density ratio η/s . Of particular note is the parameter controlling the fluctuations in initial entropy deposition σ_{fluct} , which also has greatly improved precision, and which we will come back to in Sec. IV. Even though the improvement for $\tau_{\pi} s T / \eta$ is modest, it is worth noting that, with the present posteriors, the value obtained from holography [$4 - \ln(4) \approx 2.61$, [51]] is now clearly favored over the value obtained

TABLE I. Improvement in posterior uncertainty comparing the posterior including a varying $\text{cent}_{\text{norm}}$ from this work (solid blue in Fig. 3) to the combined PbPb and p Pb posterior from Ref. [9].

Parameter	I [%]	Parameter	I [%]
$N \text{ PbPb}_{2.76}$	50	$(\eta/s)_{\text{min}}$	50
$N \text{ PbPb}_{5.02}$	56	$(\eta/s)_{\text{slope}}$	97
w	-22	$(\eta/s)_{\text{curv}}$	21
T_{switch}	66	$(\zeta/s)_{\text{max}}$	-19
p	14	$(\zeta/s)_{\text{width}}$	40
d_{min}	22	$(\zeta/s)_{T_0}$	18
σ_{fluct}	44	$\tau_{\pi} s T / \eta$	21
n_c	0	$\tau_{\pi\pi} / \tau_{\pi}$	-15
χ_{struct}	22	ν_{fs}	33
τ_{fs}	-17		




FIG. 4. We show MCMC posterior distributions for the nucleon-nucleon cross section for both 2.76 (left) and 5.02 TeV (right) collision energies both with (solid blue) and without (dashed red) a varying centrality normalization $\text{cent}_{\text{norm}}$. All are consistent with the measured values of 61.8 ± 0.9 and 67.6 ± 0.6 mb for 2.76 and 5.02 TeV, respectively [53].

by a weakly coupled computation (5, [52]) for fixed $\text{cent}_{\text{norm}}$, and slightly favored with varying $\text{cent}_{\text{norm}}$. As noted before, the posteriors for v_{fs} and $(\eta/s)_{\text{min}}$ are substantially different from those obtained in Ref. [9], which is likely due to the difference between using UrQMD in this work, and using SMASH in Ref. [9].

In addition to the parameters shown in Fig. 3, in the model design we also vary the nucleon-nucleon cross-section σ_{NN} . The reason for this is mainly that this does away with the need to perform separate runs for both collision energies [10]. For a normal MCMC analysis we fix σ_{NN} to 63 or 70 mb (for 2.76 and 5.02 TeV), but it is also possible to attempt to fit σ_{NN} to the available data. This is shown in Fig. 4, where the analysis shown in Fig. 3 was repeated while allowing σ_{NN} to vary. Two interesting observations can be made. First, the previously known values of 63 mb ($\sqrt{s_{\text{NN}}} = 2.76$ TeV) and 70 mb ($\sqrt{s_{\text{NN}}} = 5.02$ TeV) are compatible with our fits, both with and without varying the centrality normalization $\text{cent}_{\text{norm}}$. Second, the fit is not very sensitive to the value of σ_{NN} at either collision energy. This indicates that one could probably produce reasonable predictions for both collision

energies by fixing σ_{NN} to a reasonable value compatible with the fits for both collision energies. It is interesting to see whether these fits can be made more precise in the future by including extra experimental data, such as the total hadronic PbPb cross section.

C. Postdictions with systematic uncertainties

In addition to showing posteriors for the parameters, we can also use the posterior distributions to compute observables that we did not use in the fit. To do this, we drew 20 random sets of parameters (indicated in Fig. 3) from the posterior and computed our observables for all of them using 0.5M events per parameter set for PbPb collisions and 1M (2M) events for OO collisions at $\sqrt{s_{\text{NN}}}$ of 7 (0.2) TeV. For 7 TeV collisions we extrapolated the norm N from the 2.76 and 5.02 norms by assuming a power-law scaling, whereas for 0.2 TeV we multiplied the 2.76 norm by 0.497 (this factor approximates the AuAu multiplicity at 0.2 TeV relatively well). We used 39.7 and 74.2 for σ_{NN} for, respectively, 0.2 and 7 TeV collisions. These 20 sets allow us to estimate a (one standard deviation) systematic uncertainty coming from our uncertainties in the model parameters in addition to the statistical uncertainty that is computed for all observables. Since we can use considerably more events as compared with the design points we can also make predictions or postdictions for more statistically demanding observables. Note that since we are drawing the parameters from the true posterior, and not just the one-dimensional (1D) projections shown in Fig. 3, all of the correlations between the posteriors of the parameters are automatically taken into account in the prediction.

The first such prediction is shown in Fig. 5, where we show identified particle yields and mean transverse momenta for PbPb collisions at $\sqrt{s_{\text{NN}}} = 5.02$ TeV together with experimental results from ALICE [54]. Note that, as in Ref. [9], at 5.02 TeV we only used charged particle multiplicity and integrated as well as p_T -differential $v_n\{k\}$ in the fit, so these particle identified multiplicities and $\langle p_T \rangle$ are true postdictions




FIG. 5. We show identified particle yields dN/dy (left) and mean p_T (right) for pions, kaons, and protons at midrapidity as a function of centrality for PbPb collisions at $\sqrt{s_{\text{NN}}} = 5.02$ TeV compared with colored ALICE data points [54] (also shown in the model/data ratio). The bands represent systematic uncertainties (1σ) in the TRAJECTUM parameters (see Fig. 3) and the grey points show statistical uncertainties.




FIG. 6. We show identified particle yields dN/dy (left) and mean p_T (middle) for pions, kaons, and protons at midrapidity as a function of centrality for OO collisions at $\sqrt{s_{NN}} = 7$ TeV and mean p_T at $\sqrt{s_{NN}} = 0.2$ TeV (right). The bands represent systematic uncertainties (1σ) around the median (shown dashed) in the TRAJECTUM parameters (see Fig. 3) and the gray points show statistical uncertainties.

(the equivalent 2.76 TeV data was used in the fit, however). The identified particle yields are consistent with the experimental data with the exception of the protons to which we will get back when we discuss the full spectra around Fig. 10. The mean proton transverse momentum is above the experimental data, but still consistent within two standard deviations when taking into account our systematic uncertainty.

It is also interesting to examine observables which are currently not measured, as this allows us to make novel predictions. In Fig. 6, we show the same observables as in Fig. 5, but for OO at $\sqrt{s_{NN}} = 7$ TeV and 0.2 TeV instead of PbPb at $\sqrt{s_{NN}} = 5.02$ TeV. This is the (approximate) collision energy at which RHIC and the LHC collide the OO system. We are able to predict the identified particle yields to around 10% precision, and the mean transverse momenta to around 2%–3%.

We use transverse momentum fluctuations of charged particles in PbPb collisions at $\sqrt{s_{NN}} = 2.76$ TeV in the fit, but the experimental data for this observable is not available at $\sqrt{s_{NN}} = 5.02$ TeV. In the left panel of Fig. 7, we show the

TRAJECTUM prediction not only for charged particles, but also for pions, kaons, and protons (see also Refs. [55–57] for other predictions or postdictions). For protons, the transverse momentum fluctuations are predicted with 5%–15% precision (statistics dominated, see the gray statistical uncertainties), depending on centrality, whereas the other three are predicted to less than 5%. Curiously, the identified transverse momentum fluctuations do not seem to be ordered according to mass, but rather the kaon transverse momentum fluctuations are the largest. The middle and right panels of Fig. 7 show the same observables, but this time for OO at $\sqrt{s_{NN}} = 7$ TeV (middle) and $\sqrt{s_{NN}} = 0.2$ TeV (right), which are the collision energies at which the LHC and RHIC, respectively, collide oxygen nuclei. One can see that the centrality dependence is weaker in OO from what it is in PbPb. Also, the uncertainties in OO are substantially larger than in PbPb, but still the precision for charged particles is less than 10% for most centralities.

An interesting family of observables which we did not use in the fit are the event plane angle correlations. In Fig. 8, we show $\langle \cos[4(\Phi_2 - \Phi_4)] \rangle_{SP}$, $\langle \cos[8(\Phi_2 - \Phi_4)] \rangle_{SP}$ and




FIG. 7. We show transverse momentum fluctuations of charged particles, as well as pions, kaons, and protons, all at midrapidity, as a function of centrality. This is done for PbPb collisions at $\sqrt{s_{NN}} = 5.02$ TeV (left), as well as OO collisions at $\sqrt{s_{NN}} = 7$ TeV (middle) and $\sqrt{s_{NN}} = 0.2$ TeV (right). The bands represent systematic uncertainties (1σ) in the TRAJECTUM parameters (see Fig. 3) and the gray points show statistical uncertainties.

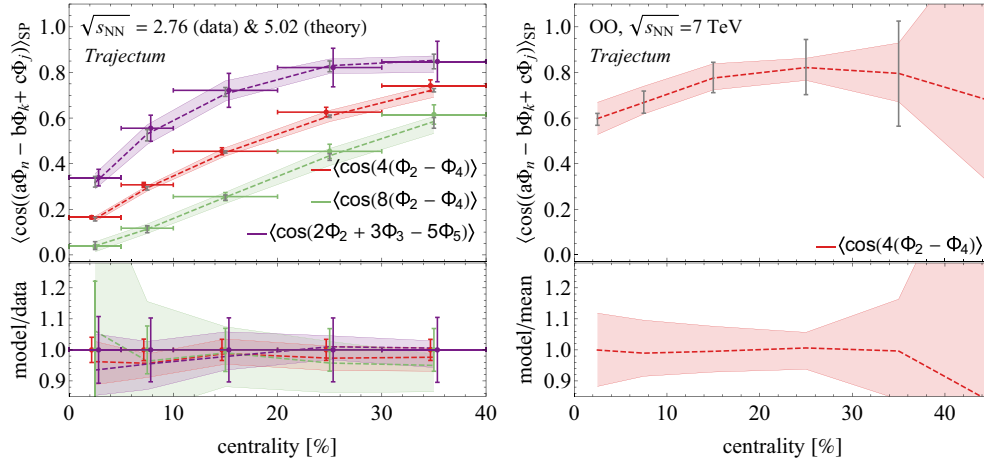


FIG. 8. We show event plane correlations $\langle \cos[4(\Phi_2 - \Phi_4)] \rangle_{\text{SP}}$, $\langle \cos[8(\Phi_2 - \Phi_4)] \rangle_{\text{SP}}$ and $\langle \cos(2\Phi_2 + 3\Phi_3 - 5\Phi_5) \rangle_{\text{SP}}$ of charged particles as a function of centrality for PbPb collisions at $\sqrt{s_{\text{NN}}} = 5.02$ TeV compared with colored ATLAS data points at $\sqrt{s_{\text{NN}}} = 2.76$ TeV (left) [58]. We also show predictions for $\langle \cos[4(\Phi_2 - \Phi_4)] \rangle_{\text{SP}}$ for OO at $\sqrt{s_{\text{NN}}} = 7$ TeV (right). The bands represent systematic uncertainties (1σ) in the TRAJECTUM parameters (see Fig. 3) and the gray points show statistical uncertainties.

$\langle \cos(2\Phi_2 + 3\Phi_3 - 5\Phi_5) \rangle_{\text{SP}}$ for PbPb at $\sqrt{s_{\text{NN}}} = 5.02$ TeV in the left panel. We compare the theory prediction to data at $\sqrt{s_{\text{NN}}} = 2.76$ TeV because this measurement is not available at $\sqrt{s_{\text{NN}}} = 5.02$ TeV. What makes these observables interesting is that they are not obviously related to the observables which we used in the fit, but we still describe them with an accuracy that can be even smaller than the experimental uncertainty. This means that, at least within our model, the event plane angle correlations are related to the observables that we do use in the fit in a nontrivial way, so that they can be completely described in terms of these simpler observables. In the right panel of Fig. 8, we make a prediction for $\langle \cos[4(\Phi_2 - \Phi_4)] \rangle_{\text{SP}}$ for OO at $\sqrt{s_{\text{NN}}} = 7$ TeV. For centrality bins up to about 30% centrality we are able to make a prediction with a precision of about 10%, with precision deteriorating for more peripheral collisions due to our limited statistics.

In the fit, we have used integrated anisotropic flow coefficients $v_n\{k\}$ at both $\sqrt{s_{\text{NN}}} = 2.76$ TeV and $\sqrt{s_{\text{NN}}} = 5.02$ TeV, with centrality bins bounded by (0, 5, 10, 20, 30, 40, 50, 60, 70)%, statistics allowing. In the left panel of Fig. 9, we show the TRAJECTUM results for $\sqrt{s_{\text{NN}}} = 5.02$ TeV using finer centrality bins. One can see that we obtain good precision, around 5%–10% for most centrality bins. For most centralities, we also have excellent agreement with ALICE data, with the largest deviations being around 2σ in $v_2\{2\}$ and $v_2\{4\}$ for centralities from 50% to 70%. In the middle and right panels, we again show predictions for OO collisions at $\sqrt{s_{\text{NN}}} = 7$ TeV (middle) and $\sqrt{s_{\text{NN}}} = 0.2$ TeV (right). We obtain predictions with a precision of less than 10% for most flow coefficients, with $v_3\{2\}$ and $v_4\{2\}$ being less precisely predicted for peripheral collisions. This can be understood, however, as $v_3\{2\}$ becomes very small for peripheral collisions, while $v_4\{2\}$ is very small for every




FIG. 9. We show integrated anisotropic flow coefficients $v_n\{k\}$ of charged particles at midrapidity as a function of centrality for PbPb collisions at $\sqrt{s_{\text{NN}}} = 5.02$ TeV compared with colored ALICE data points [60] (left), predictions for OO collisions at 7 TeV (middle) and 0.2 TeV (right). The bands represent systematic uncertainties (1σ) in the TRAJECTUM parameters (see Fig. 3) and the gray points show statistical uncertainties.




FIG. 10. We show the identified transverse momentum spectra for pions (left), kaons (middle), and protons (right) at midrapidity for different centralities for PbPb collisions at $\sqrt{s_{\text{NN}}} = 2.76$ TeV compared with colored ALICE data points [42]. The bands represent systematic uncertainties (1σ) in the TRAJECTUM parameters (see Fig. 3). We note as opposed to our 5.02 TeV results these spectra are generated using the emulator and hence also contain an uncertainty due to the emulator.

centrality class. Also note that the systematic uncertainties in all three panels include values where $v_n\{k\} < 0$, whereas from its definition $v_n\{k\}$ is either positive real or the k th root of a negative number. To make sense of the complex values, we redefine $v_n\{k\}$ as $\tilde{v}_n\{k\} = \text{sgn}(v_n\{k\}^k)|v_n\{k\}|$. This essentially rotates the complex values to the negative real axis while keeping the real ones unchanged, making $\tilde{v}_n\{k\}$ real in all cases. In Fig. 9, wherever $v_n\{k\}$ becomes complex, we therefore show $\tilde{v}_n\{k\}$ instead. We note that $v_2\{4\}$ in PbPb collisions also becomes complex for centralities smaller than 4% [59] (see also Ref. [10]), but our current simulation does not have sufficient statistics to fully probe this regime.

D. p_T -differential predictions

In addition to the various p_T -integrated observables discussed above, we also compute several p_T -differential observables. In Fig. 10, we show identified transverse momentum spectra for pions (left), kaons (middle), and protons (right), for PbPb at $\sqrt{s_{\text{NN}}} = 2.76$ TeV. Note that, in contrast

to the other figures described in this section, the contents of Fig. 10 were not made using a high-statistics run, but instead come from evaluating the emulator which was used to perform the Bayesian analysis. Whereas deviations from the model with the experimental data is typically small (around 10%), it is clear that especially for more off-central collisions we have too few high- p_T pions. This could indicate that perturbative QCD processes (which are only power-law suppressed at high p_T as opposed to the exponential Boltzmann suppression) become more important towards more peripheral collisions. Also mostly at peripheral collisions we have too many kaons as well as high- p_T protons.

Figure 11 shows the same quantities as Fig. 10, but now for PbPb at $\sqrt{s_{\text{NN}}} = 5.02$ TeV. In this figure, the high-statistics run was used to provide the TRAJECTUM prediction. Again the precision is around 10%, but here deviations from the experimental data are much larger, especially as the experimental uncertainty has decreased significantly with respect to 2.76 TeV data. Given that we did not fit to the spectra at $\sqrt{s_{\text{NN}}} = 5.02$ TeV, this is not too alarming, but it does

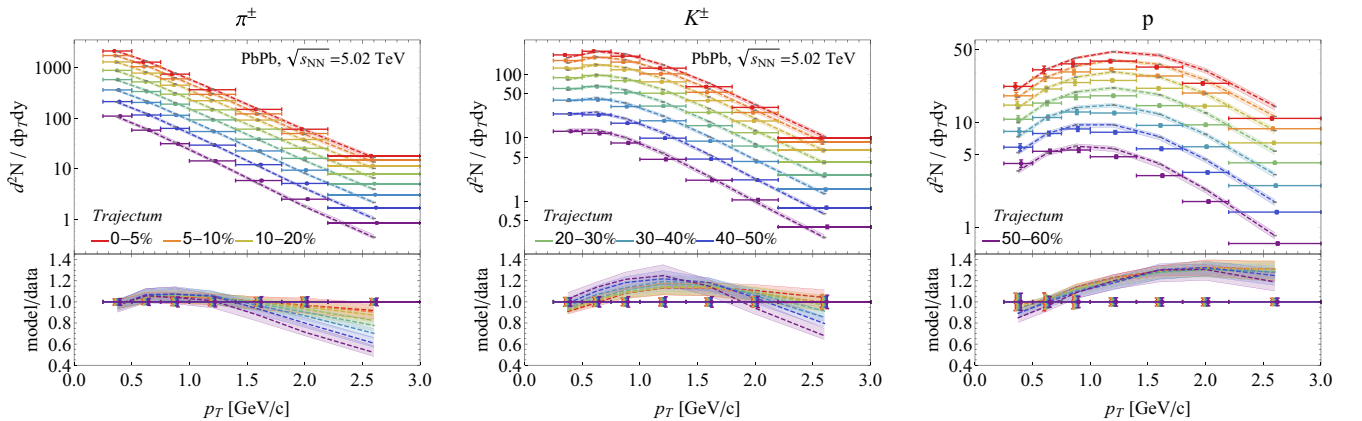


FIG. 11. We show the identified transverse momentum spectra for pions (left), kaons (middle), and protons (right) at midrapidity for different centralities for PbPb collisions at $\sqrt{s_{\text{NN}}} = 5.02$ TeV compared with colored ALICE data points [54]. The bands represent systematic uncertainties (1σ) in the TRAJECTUM parameters (see Fig. 3) and the grey points show statistical uncertainties.

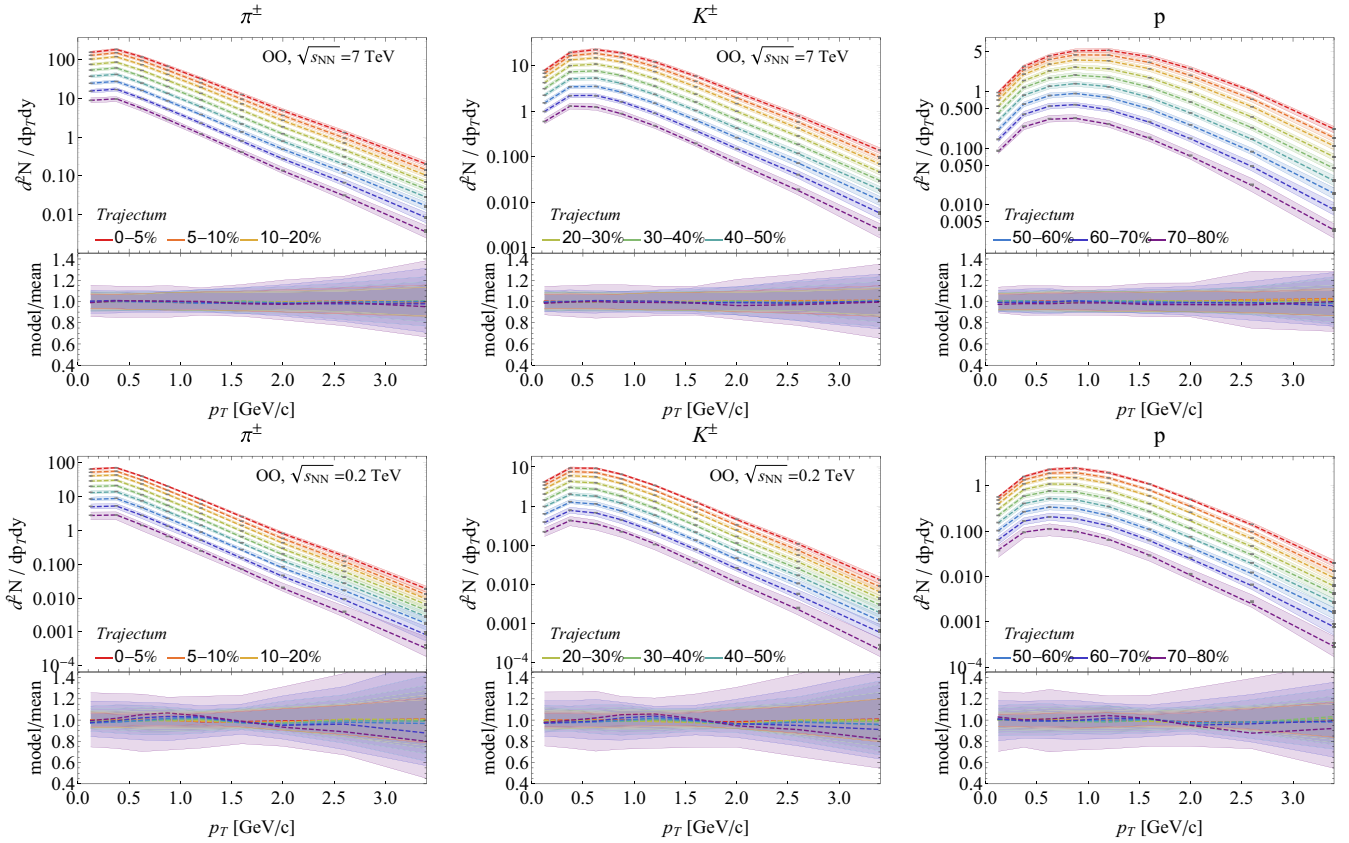


FIG. 12. We show the identified transverse momentum spectra for pions (left), kaons (middle), and protons (right) at midrapidity for different centralities for OO collisions at $\sqrt{s_{NN}} = 7$ TeV (top) and $\sqrt{s_{NN}} = 0.2$ TeV (bottom). The bands represent systematic uncertainties (1σ) in the TRAJECTUM parameters (see Fig. 3) and the gray points show statistical uncertainties.

indicate that including these measurements in a future fit is likely to improve constraints. We also show predictions for identified transverse momentum spectra for OO in Fig. 12, where we show pions (left), kaons (middle), and protons (right) at $\sqrt{s_{NN}} = 7$ TeV (top) and $\sqrt{s_{NN}} = 0.2$ TeV (bottom). The precision of the predictions depends on centrality and

collision energy, but is around 10%–30% for most centrality and p_T bins.

Lastly, Fig. 13 shows the identified p_T -differential anisotropic flow coefficient $v_2\{2\}(p_T)$ for pions (left), kaons (middle), and protons (right), for PbPb at $\sqrt{s_{NN}} = 5.02$ TeV. Depending on centrality and transverse momentum, the

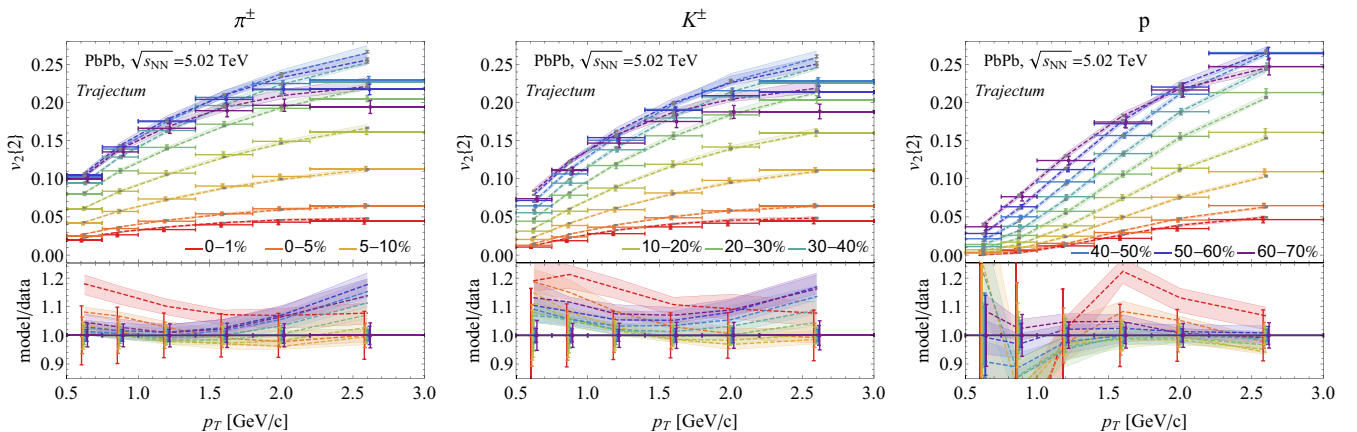


FIG. 13. We show the identified p_T -differential anisotropic flow coefficient $v_2\{2\}(p_T)$ for pions (left), kaons (middle), and protons (right) at midrapidity for different centralities for PbPb collisions at $\sqrt{s_{NN}} = 5.02$ TeV compared with colored ALICE data points [61]. The bands represent systematic uncertainties (1σ) in the TRAJECTUM parameters (see Fig. 3) and the gray points show statistical uncertainties.




FIG. 14. We show a zoom of Fig. 9 (left) together with data from ALICE [60].

precision ranges from around 5% to 10%. The largest relative deviations from experimental data can be seen in the 0%–1% centrality bin (this particular centrality bin was not used in the fit), but in absolute terms these still agree within two standard deviations. The only serious deviations are the overestimations of the elliptic flow for high- p_T pions and kaons for centralities bigger than 40%. This could again indicate the importance of pQCD contributions, as those contributions are expected to decrease the flow coefficients v_n . The agreement of the proton $v_2\{2\}$ is however impressively precise over the entire range of centralities (up to 70%) and transverse momentum probed.

IV. PREDICTIONS FOR ULTRACENTRAL COLLISIONS

Ultracentral collisions in the 0%–1% centrality class are rare events whereby the nuclei collide head-on and furthermore fluctuate specifically to have a high multiplicity. These collisions hence probe a special part of heavy ion collision physics, and have been historically difficult to describe accurately. A relatively famous example is the v_2 - v_3 -ultracentral puzzle, which states that it is difficult to simultaneously describe both v_2 and v_3 in the 0%–1% centrality class [13,62–64], but also for mean transverse momentum there are theoretical predictions [65]. The interesting observation in the latter work is that for extremely ultracentral collisions (0%–0.01% centrality) the size of the plasma is limited by the radius of the colliding nucleus and the increase in multiplicity is solely sourced by an upward fluctuation in the initial temperature. The increase in temperature increases the final radial flow and hence leads to an increase in $\langle p_T \rangle$.

In Fig. 14 we show TRAJECTUM postdictions for $v_n\{k\}$ as compared with ALICE measurements [60] [this figure is a zoom of Fig. 9 (left)]. This yields an excellent fit to the experimental data [60], even though these particular centrality classes were not used in the Bayesian model. The v_2 - v_3 puzzle is usually stated as a puzzle on the ratio, since it is often straightforward to reduce or increase the magnitude of both by modifying the shear viscosity. This, however, would

modify the $v_n\{k\}$ for higher centrality classes and we use our parameters within the posterior range that are constrained by such data. In that case it can be challenging to fit the absolute magnitude of both v_2 and v_3 (indeed, Ref. [64] is off by many standard deviations for $v_2\{2\}$, while the $v_2\{2\}/v_3\{2\}$ ratio is consistent with data within a single standard deviation). We remark that both Refs. [32] and [8,20] have an excellent description of the 0%–5% centrality class of both $v_2\{2\}$ and $v_3\{2\}$ so we regard it as likely that they will also describe the ultracentral centrality classes. The good fit of the Magma model [63], however, likely does not survive a full hydrodynamic simulation [13].

It is an interesting question why TRAJECTUM would obtain a better ultracentral description than e.g., v-USPhydro in Ref. [64], especially as the latter work uses the same initial-state model (TRENTo) and furthermore attempted some important changes to resolve the puzzle (albeit unsuccessfully). For this question we performed a simple TRENTo simulation [7,26] using our MAP settings from Ref. [9] to obtain the spatial anisotropies ϵ_2 and ϵ_3 , here defined by the initial energy density through [26]

$$\epsilon_n = \left| \frac{\int r dr d\phi r^n e^{in\phi} e}{\int r dr d\phi r^n e} \right|. \quad (4)$$

We furthermore show three variations obtained by setting the number of constituents to $n_c = 1$ as well as rescaling σ_{fluct} by a factor $\sqrt{3}$. We note that ϵ_2 and ϵ_3 correlate very well with $v_2\{2\}$ and $v_3\{2\}$, respectively [10,13,66]. We hence see in Fig. 15 that while the influence of σ_{fluct} is moderate over the complete centrality range (left panel) the contribution becomes relatively important in ultracentral collisions (middle and right panels) because of the small absolute magnitude of ϵ_2 and ϵ_3 in ultracentral collisions. The subnucleonic structure, varied by varying n_c is however only important for very peripheral collisions and more so for ϵ_3 than for ϵ_2 . To describe the ultracentral $v_2\{2\}$ it is hence essential to estimate σ_{fluct} accurately, which in TRAJECTUM is done by including transverse momentum fluctuations in the Bayesian analysis (see also Sec. III A and [20]). This can likely explain the difference with Ref. [64], which used an older Bayesian analysis [6] without transverse momentum fluctuations. Similar to the event plane angle correlations we hence see again the power of a full global analysis that links heavy ion physics across observables.

Several observables have not been measured in the ultracentral region, including the mean transverse momentum as advocated in Ref. [65]. In Fig. 16 we show the $\langle p_T \rangle$ for increasingly ultracentral collisions together with the multiplicity as well as the anisotropic flow coefficients $v_n\{2\}$. All these observables are statistically relatively cheap to compute, but we note that to obtain results for the 0%–0.001% centrality class it was still necessary to obtain 33M TRENTo events for each of the 20 parameter sets and subsequently sample the 0%–0.01% highest entropy events with a 100 times higher rate with respect to the other events (as stated, we used 0.5M hydro events in total). As expected, the multiplicity increases relatively steeply towards more central events. This increase can in principle be compared with the multiplicity beyond the




FIG. 15. We show a simple TRENTo simulation of the initial eccentricities ϵ_2 (left, middle) and ϵ_3 (right), where in the middle panel we show a zoom of the panel on the left together with ϵ_3 for comparison (black dot-dashed). In addition to the maximum a posteriori (MAP) values, we show results where σ_{fluct} was rescaled by a factor $\sqrt{3}$, as well as results where the number of constituents n_c was set to 1. As insets we also show ratios of these variations with respect to the MAP settings. One can see that σ_{fluct} has a large influence on both ϵ_2 and ϵ_3 , which is part of the resolution of the ultracentral v_2 to v_3 puzzle [13,62–64].

“knee” in the V0 amplitude in Ref. [33], which is also shown in the figure for 2.76 TeV collision energy [33] (for 5.02 TeV the V0 distribution suffers from pileup in this region and is harder to compare). Note that our model in this work is boost invariant and it is hence appropriate to compare our midrapidity multiplicity and midrapidity centrality selection ($|\eta| < 2.4$) with the ALICE V0 distribution (which uses the multiplicity and the amplitude for the centrality selection both at forward rapidity). Ultracentral multiplicity has also been studied at RHIC [67], although here it is problematic that the

STAR data are only available in raw form (the V0 data are also raw, but as a ratio with the 1%–2% class detector effects should be suppressed). In the right panel of Fig. 16 we find that there is no significant change in any $v_n\{2\}$ (with $n \leq 4$) downwards of 0.5% centrality.

The mean transverse momentum in extremely ultracentral collisions is especially interesting, since for these centrality classes the initial volume is approximately constant [65]. If one furthermore assumes that the total energy E and entropy S per unit rapidity are approximately conserved then the

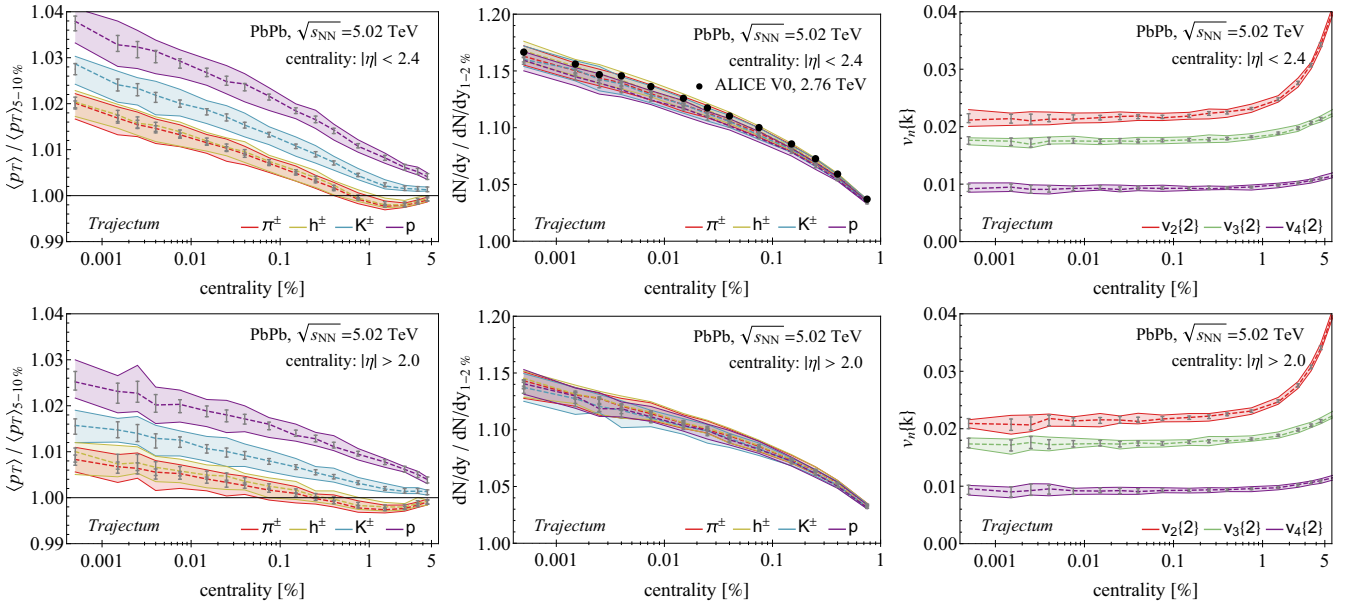


FIG. 16. (top) We show the mean transverse momentum for charged pions, hadrons, kaons and protons normalized to its value in the 5%–10% centrality class towards extremely ultracentral collisions (left, see also [65]) and similarly for the multiplicity at midrapidity (middle, we normalize with the 1%–2% class and uses dN/dy and $dN/d\eta$ for identified and unidentified hadrons respectively) and the $v_n\{2\}$ coefficients (right, not normalized). While the 3% difference for protons for the mean transverse momentum naively amounts to an effect that is within the systematic uncertainty the normalization with the 5%–10% class should allow much of the systematics to cancel, such that with the high luminosity delivered to ALICE this measurement is likely feasible. The $v_n\{2\}$ do not decrease further towards extremely ultracentral collisions, indicating that the QGP never becomes perfectly spherical. (bottom) Instead of our standard centrality measure (charged multiplicity with $p_T > 400$ MeV and $|\eta| < 2.4$ we show the same figures for a $|\eta| > 2.0$ cut, which is similar to the forward V0 centrality definition typically used by ALICE. By construction the effect on the multiplicity is now smaller (since we select in a different rapidity range), but an even greater reduction is found for the $\langle p_T \rangle$. Even bigger effects from the centrality measure can be expected for models that do not assume boost invariance.




FIG. 17. We show some 1D posterior distributions for the 20 Bayesian analyses including OO simulated data (solid, colored) together with the posterior distributions including only PbPb data (dashed, black). We show the nucleon substructure parameter χ_{struct} (left) and the minimum of the shear viscosity to entropy density ratio $(\eta/s)_{\text{min}}$ (right). It can be seen that the distributions for $(\eta/s)_{\text{min}}$ are all fairly close together, whereas those for χ_{struct} show more variation. For χ_{struct} , the distribution is usually narrower than the PbPb result, which can be easily seen because the peaks of the normalized distributions are taller than the PbPb result.

increase in $\langle p_T \rangle$ is directly related to the increase in e/s , with e and s being the initial energy and entropy densities that can potentially be interpreted as being directly related to the temperature. In reality the energy per rapidity decreases due to longitudinal work and entropy increases due to viscous corrections. Our results from the 1%–2% centrality class to the 0.1%–0.2% centrality class show an increase of the $\langle p_T \rangle$ for charged hadrons from approximately 715 to 724 MeV, which is in agreement with Ref. [65]. We note that, curiously, there is a small but significant decrease in the pion $\langle p_T \rangle$ going from centrality 5% to the 1%–2% class.

It is an interesting question how the extremely central observables are affected by a centrality selection which is different from the rapidity region where the observable is measured. This avoids autocorrelation effects from (rare) fluctuations and is typically what is done in experiment by, e.g., the ALICE forward V0 detector. To simulate this we redid our computation with a centrality measure given by the multiplicity at $|\eta| > 2.0$ (we cut of the particlization at $|\eta| = 2.5$ in our simulations, so we cannot use the “true” V0 ranges) and show this in the bottom row of Fig. 16. As expected, the increase in multiplicity is now weaker, since the multiplicity observable now decorrelates with the centrality selection. The reduction of the ultracentral effect on the $\langle p_T \rangle$ is even stronger, which can likely be explained because the dependence of $\langle p_T \rangle$ on centrality is much weaker than for the multiplicity. Note that our hydrodynamic model assumes longitudinal boost invariance, so any dependence on the centrality selection is due to thermal sampling during the particlization. Realistically the plasma itself also fluctuates (see, e.g., Refs. [68,69]), which would lead to stronger effects from the centrality selection than in our boost invariant model. To avoid this it is likely essential for extremely ultracentral observables to use the

same centrality rapidity interval as in which the measurement is made (as in the top row of Fig. 16).

V. BAYESIAN ANALYSIS USING OO SIMULATED DATA

In the previous section, we showed several predictions for the oxygen runs at the LHC and RHIC. While this is very interesting and important because it allows for a good test of the current model, one can answer an additional question. Assuming, as we have been so far, that the soft sector of OO collisions can be described by the same hydrodynamical model as for PbPb, one can wonder whether the addition of OO data can improve the constraints on the parameters such as those obtained from PbPb alone.

One might hope that OO yields constraints orthogonal to those already given by PbPb because OO is a much smaller system, but it is not *a priori* obvious that this is indeed the case. In this section, we obtain an answer to this question by using simulated data for OO at $\sqrt{s_{\text{NN}}} = 7$ TeV generated by the model itself (this is similar to a closure test [9–12]). We perform Bayesian analyses for 20 different sets of simulated data, where the simulated data come from high-statistics runs of the model itself using the parameters indicated in Fig. 3. Since we assume consistency between OO and PbPb, the “true” value should be one of the points in the posterior of our fit to PbPb data, so it makes sense to use parameter settings which are randomly drawn samples from the posterior (see Fig. 3). We have to make an additional assumption on the uncertainty of the simulated data. Estimates for the statistical and systematic uncertainties of the OO observables we would like to include in the fit have recently become available [16,70,71], but here we make the simplifying assumption that the relative experimental errors are the same as for the corresponding observables in PbPb. The observables we have used in the simulated data OO fits are the same as those used in the fit using only PbPb experimental data, except that for p_T -differential anisotropic flow we have only used the bins bounded by (0.5, 0.75, 1, 1.4) GeV/ c , and we only fit pion $v_2\{2\}(p_T)$ (as shown in Fig. 1).

In Fig. 17, we compare the results of the fit which only includes the PbPb experimental data to the 20 fits which also include simulated data for OO, where in both cases we use the fits that also vary $\text{cent}_{\text{norm}}$. In the left panel, we show the posterior distribution of the χ_{struct} parameter, whereas on the right we show $(\eta/s)_{\text{min}}$. There are two interesting observations. First, comparing the PbPb fit to the simulated data fits, one can see that the simulated data fits are narrower and shifted for χ_{struct} and virtually unchanged for $(\eta/s)_{\text{min}}$. This implies that the determination of χ_{struct} becomes more precise with the inclusion of simulated data, indicating that, under the assumptions we made, the inclusion of real OO data would also improve the precision of the fits. The second observation is that the χ_{struct} distributions do not all have their peak in the same place, whereas the $(\eta/s)_{\text{min}}$ distributions are almost on top of each other. Recall that the parameters for the simulated data are drawn randomly from the PbPb posterior. This means that, as one can see in Fig. 3, the true parameter values for the simulated data points are spread evenly according to the PbPb posterior distribution. Naively, one might think that the fits to

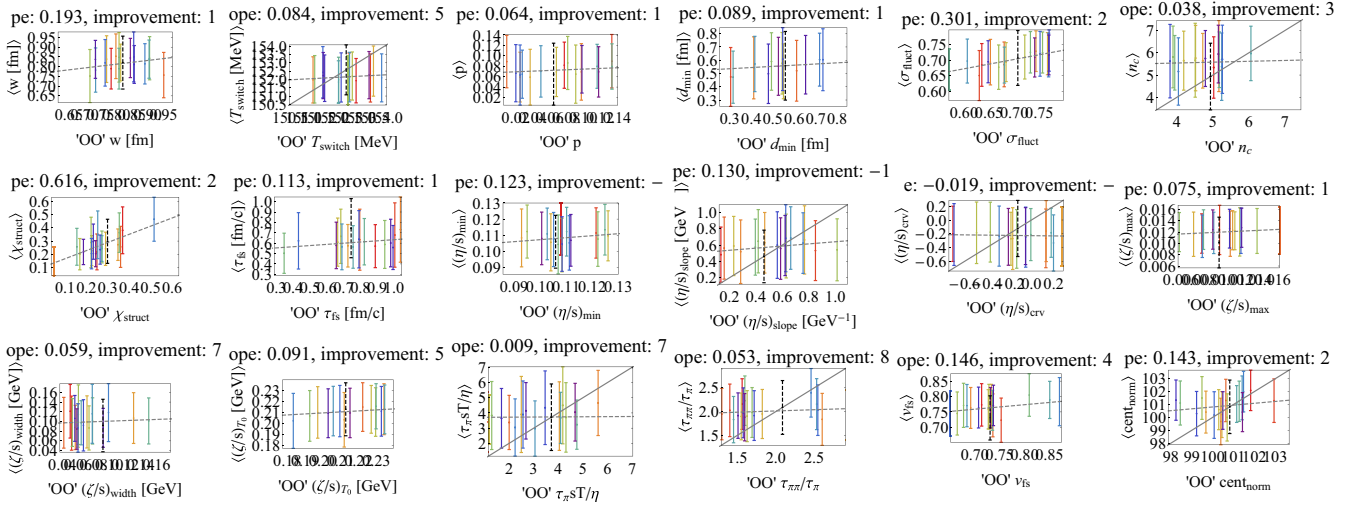


FIG. 18. For each parameter in our model, we show the original PbPb fit (black), as well as the results for the 20 fits which include the OO simulated data (colored). The horizontal position of the PbPb data point is given by its mean, whereas the horizontal positions of the data points which include the simulated data are given by their true values. We also show a linear function fit to the simulated data points (dotted) as well as a line with slope 1 (solid). The value of the slope of the linear function is also quoted, together with the improvement of the posterior uncertainty.

the simulated OO data should show reconstructed values centered around the true values, and therefore one may wonder why the simulated data posteriors for $(\eta/s)_{\min}$ are spaced so close together. The explanation for this is that, while for the simulated OO data we know the true parameters, this is not true for the PbPb data, which are still taken from the actual experiments. Both the PbPb data and the OO simulated data pull on the parameters, and the relative strengths of the pulls of each dataset determine the final posterior distribution. This leads us to the interpretation that the pull of OO on $(\eta/s)_{\min}$ is small relative to that of PbPb, whereas for χ_{struct} the pull of OO is larger.

Figure 18 explores these two observations for all the parameters in our model. The black points indicate the expectation value and precision of the parameter using just the PbPb experimental data. In color, we show the various fits using the 20 OO simulated data sets as well as PbPb experimental data, where the horizontal position corresponds to the true parameter value that was used to generate the particular simulated OO data set. Also shown is the slope of a linear function fit to the OO simulated data fits. If the parameter can be fully constrained by OO data with much higher precision than from PbPb data one would expect a slope of unity, whereas if PbPb data are more constraining one would expect a slope closer to zero. The slope hence quantifies the discussion above about the relative pull of the OO simulated data compared with the PbPb data. In addition, we show the improvement I in the size of the standard deviation of the simulated data posteriors with respect to the PbPb posterior, where we define I as

$$I = \sigma_{\text{PbPb}} / \langle \sigma_{\text{PbPb} + \text{OO}} \rangle_{\text{OO simulated data sets}} - 1,$$

where the average is taken over the posteriors of all OO simulated data fits, and σ denotes the standard deviation of the posterior. As we can see, the uncertainties in most parameters decrease modestly with the inclusion of OO simulated data,

as can be seen by the relatively small but mostly positive values for the improvement. This indicates that, under the assumptions mentioned in this section, the inclusion of real OO experimental data in future analyses could be beneficial to the precision of the analyses. We can also see that the correlation varies from parameter to parameter, indicating the different relative pulls between PbPb and OO data. Let us also note that correlations between parameters are important. The parameter d_{\min} is not used in OO collisions, however, its determination does benefit from the inclusion of OO simulated data. The reason for this is that it turns out that d_{\min} is strongly correlated with σ_{fluct} (see Fig. 1), which is also improved by the addition of OO simulated data. Through the correlation, an improvement in one of the two parameters then translates into an improvement in the other as well.

VI. DISCUSSION

We performed a Bayesian analysis where we fit PbPb experimental data to a sophisticated theoretical model, where we substantially reduced the statistical and emulator errors with respect to Refs. [9,10]. This results in reduced uncertainties in the posterior distributions for the varied parameters. By sampling 20 random parameter sets from the resulting posterior and performing high-statistics runs for each parameter set, we subsequently were able to make precise pre- and postdictions for both PbPb and OO collisions. Because the parameter sets with which these results were computed are randomly drawn, we are able to assign systematic errors to the results which arise because of the uncertainty in the parameters, whereby this also fully incorporates correlations between the parameters.

An interesting result from these high-statistics runs is a possible resolution of the ultracentral v_2 to v_3 puzzle. It turns out that, to correctly describe $v_2\{2\}$ and $v_3\{2\}$ for ultracentral

PbPb collisions, the magnitude of initial-state per-nucleon fluctuations in entropy deposition σ_{fluct} is important. Given that mean transverse momentum fluctuations $\delta p_T / \langle p_T \rangle$ are sensitive to σ_{fluct} (see Fig. 1), including $\delta p_T / \langle p_T \rangle$ in the set of observables used to fit the parameters improves the results for ultracentral $v_2\{2\}$ and $v_3\{2\}$ as well.

We then used the 20 randomly chosen parameter sets as simulated data to test what the effect of adding future OO data at $\sqrt{s_{\text{NN}}} = 7$ TeV would be on the precision of Bayesian analyses. This showed that such an addition would most often yield an improvement in the precision of the extracted parameters, albeit most often a moderate improvement and unfortunately no significant improvement for first- or second-order hydrodynamic transport coefficients. One should, however, keep in mind that this is true working under the assumption that OO can be described by the same hydrodynamical model that also describes PbPb. This assumption can be violated by several causes related to the smaller size of the OO system, such as the need to simulate hard processes in addition to hydrodynamics, or that OO should be simulated using $(3+1)$ D hydrodynamics instead of $(2+1)$ D. The presented precision of our OO predictions, following from simply assuming the TRAJECTUM model, will hence be an extremely interesting test. Even if it may not improve the precision of the parameters much we like to point out that the parameters are already rather precise and corroboration of the model using OO collisions will be a significant achievement which increases the trust in the precision of the model.

It was mentioned above that one reason one might expect that OO yields constraints orthogonal to those obtained from PbPb is that OO is a much smaller system. A valid point one might raise is why we are then comparing PbPb and OO to PbPb alone, and not PbPb, p Pb, and OO to PbPb and p Pb. After all, p Pb is also a small system, so one might expect that any constraints obtained from OO can also be obtained from p Pb. In Refs. [7,9,10], the analysis also included p Pb data, and there were actually extra constraints obtained by including p Pb in the analysis. In this work, however, with the substantial improvement in precision of the emulator, this is no longer the case, and inclusion of p Pb data actually makes the obtained fits less precise. Most likely the reason is the difficulty fitting the complete p Pb spectrum to high precision, which then pushes the MCMC algorithm to regions in parameter space where the emulator uncertainty is largest. According to Bayes theorem, this is indeed the region that is then most likely (since other regions would be ruled out). Nevertheless, from a modeling point of view this may not be wanted and the next necessary step would be to either improve this emulator artifact, or to improve the model itself (likely along the lines of Ref. [72]). Both options would be profoundly interesting, but are left for future work. This same issue could mean that the predictions for, e.g., the OO spectra from Fig. 12 are less reliable than they seem, although we stress that the OO system is significantly larger than p Pb and especially the centrality dependence is more robust to model than a p Pb collision [73].

Arguably the most interesting future avenue of research in the context of this work is the future measurements of

the OO observables we predict. RHIC performed OO collisions at $\sqrt{s_{\text{NN}}} = 0.2$ TeV in May 2021 and LHC has a special OO run planned in 2024 (likely at $\sqrt{s_{\text{NN}}} = 6.8$ TeV). It will be interesting to see how well our predictions compare with experimental data because this will shed light on how well small systems can be described using hydrodynamics.

In the context of the global analysis used in this work it is useful to comment on the strengths and weaknesses of such Bayesian analyses. Given the many unknowns in the current standard model of heavy ion collisions [3] and their complicated interplay in describing a wealth of experimental data, the advantage of a global analysis is quite straightforward. In particular, the emulator allows a direct modeling of the full nonlinear interplay on all observables and the Bayes formula can assess the parameter space that gives the best fit to the data. Weaknesses are also relatively clear. When using a model that is not entirely adequate to describe the complete relevant physics, the Bayes formula will push the posterior parameter range into regions that give the “best fit” but do not necessarily reflect the underlying physics. To avoid this situation we need a sufficiently versatile model that can realistically model the observables we wish to describe, and this can require many parameters, which degrades the predictiveness of the model.

Another potential weakness is the fact that the results of a global analysis depends on the scope of the data used. An example we saw here is that one of the first such analysis (Ref. [6]) did not use p_T fluctuations and consequently does not describe the ultracentral $v_2\{2\}$ well (see Sec. IV). A newer example is Ref. [74], where it was found that the nucleus-nucleus cross section as an observable has profound consequences for the nucleon size. As such it is important in global analyses to include a wide variety of observables and collision systems and study the effects of them on the analysis. The current work does not, for instance, include data taken at RHIC (see, however, Refs. [12,48]), which is an important open avenue for future inclusion.

There are also several ways in which we can further improve our model. One is the lowering of our theoretical uncertainties. From Fig. 2, it is clear that, while in the present analysis most observables have statistical and emulator uncertainties small enough that the experimental uncertainty dominates, this is not true for all observables, especially for the anisotropic flow observables. Since quite a few of our parameters influence anisotropic flow (see Fig. 1), it is reasonable to assume that improved statistical and emulator uncertainty in anisotropic flow observables would result in improved posterior estimates for these parameters as well.

Another improvement that will be done in future work is a more general initial stage. The free streaming picture of the initial stage used in this work assumes zero coupling, whereas models assuming strong coupling, such as AdS/CFT, yield a qualitatively different plasma just after the initial stage. Also, while the TRENTo formula interpolates between a wide range of model behavior, there are certain model features which it cannot describe, such as binary scaling. In future work we will attempt to address both these issues. Finally,

it could be interesting to add different particlization methods, as was explored in Refs. [11,12], as only including a single particlization method might bias the results.

ACKNOWLEDGMENTS

We thank Scott Moreland for the idea of making the number of constituents continuous. We thank Andrea Dainese,

Jasmine Brewer, Hannah Elfner, Derek Everett, Umut Gürsoy, Aleksas Mazeliauskas, Justin Mohs, Scott Moreland, Jamie Nagle, Jean-Yves Ollitrault, Krishna Rajagopal, Raimond Snellings, Mike Sas, You Zhou, and especially Giuliano Giacalone for very helpful discussions. G.N. is supported by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics under Grant Contract No. DE-SC0011090.

-
- [1] U. Heinz and R. Snellings, *Annu. Rev. Nucl. Part. Sci.* **63**, 123 (2013).
- [2] P. Romatschke and U. Romatschke, *Relativistic Fluid Dynamics In and Out of Equilibrium*, Cambridge Monographs on Mathematical Physics (Cambridge University Press, Cambridge, 2019).
- [3] W. Busza, K. Rajagopal, and W. van der Schee, *Annu. Rev. Nucl. Part. Sci.* **68**, 339 (2018).
- [4] J. L. Nagle and W. A. Zajc, *Annu. Rev. Nucl. Part. Sci.* **68**, 211 (2018).
- [5] J. Berges, M. P. Heller, A. Mazeliauskas, and R. Venugopalan, *Rev. Mod. Phys.* **93**, 035003 (2021).
- [6] J. E. Bernhard, J. S. Moreland, S. A. Bass, J. Liu, and U. Heinz, *Phys. Rev. C* **94**, 024907 (2016).
- [7] J. S. Moreland, J. E. Bernhard, and S. A. Bass, *Phys. Rev. C* **101**, 024911 (2020).
- [8] J. E. Bernhard, J. S. Moreland, and S. A. Bass, *Nat. Phys.* **15**, 1113 (2019).
- [9] G. Nijs, W. van der Schee, U. Gürsoy, and R. Snellings, *Phys. Rev. Lett.* **126**, 202301 (2021).
- [10] G. Nijs, W. van der Schee, U. Gürsoy, and R. Snellings, *Phys. Rev. C* **103**, 054909 (2021).
- [11] D. Everett *et al.* (JETSCAPE Collaboration), *Phys. Rev. Lett.* **126**, 242301 (2021).
- [12] D. Everett *et al.* (JETSCAPE Collaboration), *Phys. Rev. C* **103**, 054904 (2021).
- [13] R. Snyder, M. Byres, S. H. Lim, and J. L. Nagle, *Phys. Rev. C* **103**, 024906 (2021).
- [14] H. Niemi and G. S. Denicol, [arXiv:1404.7327](https://arxiv.org/abs/1404.7327).
- [15] D. Everett, C. Chattopadhyay, and U. Heinz, *Phys. Rev. C* **103**, 064902 (2021).
- [16] J. Brewer, A. Mazeliauskas, and W. van der Schee, in Opportunities of OO and pO collisions at the LHC, [arXiv:2103.01939](https://arxiv.org/abs/2103.01939).
- [17] <https://sites.google.com/view/govertnijs/trajectum>.
- [18] P. Huovinen and P. Petreczky, *Nucl. Phys. A* **837**, 26 (2010).
- [19] A. Bazavov *et al.* (HotQCD Collaboration), *Phys. Rev. D* **90**, 094503 (2014).
- [20] J. E. Bernhard, Ph.D. thesis, Duke University, 2018 (unpublished).
- [21] H. De Vries, C. W. De Jager, and C. De Vries, *At. Data Nucl. Data Tables* **36**, 495 (1987).
- [22] D. Lonardonì, A. Lovato, S. C. Pieper, and R. B. Wiringa, *Phys. Rev. C* **96**, 024326 (2017).
- [23] B. Alver, M. Baker, C. Loizides, and P. Steinberg, [arXiv:0805.4411](https://arxiv.org/abs/0805.4411).
- [24] M. Rybczyński and W. Broniowski, *Phys. Rev. C* **100**, 064912 (2019).
- [25] N. Summerfield, B.-N. Lu, C. Plumberg, D. Lee, J. Noronha-Hostler, and A. Timmins, *Phys. Rev. C* **104**, L041901 (2021).
- [26] J. S. Moreland, J. E. Bernhard, and S. A. Bass, *Phys. Rev. C* **92**, 011901(R) (2015).
- [27] J. Weil *et al.*, *Phys. Rev. C* **94**, 054905 (2016).
- [28] D. Oliinchenko, V. Steinberg, J. Weil, M. Kretz, J. Staudenmaier, S. Ryu, A. Schäfer, J. Rothermel, J. Mohs, F. Li, H. E. (Petersen), L. Pang, D. Mitrovic, A. Goldschmidt, L. Geiger, J.-B. Rose, J. Hammelmann, and L. Prinz, *smash-transport/smash: Smash-1.8* (2020).
- [29] T. Sjöstrand, S. Mrenna, and P. Z. Skands, *Comput. Phys. Commun.* **178**, 852 (2008).
- [30] S. A. Bass *et al.*, *Prog. Part. Nucl. Phys.* **41**, 255 (1998).
- [31] M. Bleicher *et al.*, *J. Phys. G* **25**, 1859 (1999).
- [32] B. Schenke, C. Shen, and P. Tribedy, *Phys. Rev. C* **102**, 044905 (2020).
- [33] B. Abelev *et al.* (ALICE Collaboration), *Phys. Rev. C* **88**, 044909 (2013).
- [34] J. Adam *et al.* (ALICE Collaboration), *Phys. Rev. C* **91**, 064905 (2015).
- [35] B. Abelev *et al.* (ALICE Collaboration), *Phys. Rev. Lett.* **109**, 252302 (2012).
- [36] M. L. Miller, K. Reygers, S. J. Sanders, and P. Steinberg, *Annu. Rev. Nucl. Part. Sci.* **57**, 205 (2007).
- [37] D. d’Enterria and C. Loizides, *Annu. Rev. Nucl. Part. Sci.* **71**, 315 (2021).
- [38] C. Loizides, J. Nagle, and P. Steinberg, *SoftwareX* **1-2**, 13 (2015).
- [39] K. Aamodt *et al.* (ALICE Collaboration), *Phys. Rev. Lett.* **106**, 032301 (2011).
- [40] J. Adam *et al.* (ALICE Collaboration), *Phys. Rev. Lett.* **116**, 222302 (2016).
- [41] J. Adam *et al.* (ALICE Collaboration), *Phys. Rev. C* **94**, 034903 (2016).
- [42] B. Abelev *et al.* (ALICE Collaboration), *Phys. Rev. C* **88**, 044910 (2013).
- [43] B. B. Abelev *et al.* (ALICE Collaboration), *Eur. Phys. J. C* **74**, 3077 (2014).
- [44] J. Adam *et al.* (ALICE Collaboration), *Phys. Rev. Lett.* **116**, 132302 (2016).
- [45] J. Adam *et al.* (ALICE Collaboration), *J. High Energy Phys.* **09** (2016) 164.
- [46] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, *Publ. Astron. Soc. Pac.* **125**, 306 (2013).
- [47] C. K. Williams and C. E. Rasmussen, *Gaussian Processes for Machine Learning* (MIT press Cambridge, 2006), Vol. 2.
- [48] J. S. Moreland, Ph.D. thesis, Duke University, 2019.
- [49] S. Ryu, J. F. Paquet, C. Shen, G. S. Denicol, B. Schenke, S. Jeon, and C. Gale, *Phys. Rev. Lett.* **115**, 132301 (2015).
- [50] D. J. Earl and M. W. Deem, *Phys. Chem. Chem. Phys.* **7**, 3910 (2005).

- [51] R. Baier, P. Romatschke, D. T. Son, A. O. Starinets, and M. A. Stephanov, *J. High Energy Phys.* **04** (2008) 100.
- [52] G. S. Denicol, S. Jeon, and C. Gale, *Phys. Rev. C* **90**, 024912 (2014).
- [53] C. Loizides, J. Kamin, and D. d'Enterria, *Phys. Rev. C* **97**, 054910 (2018); **99**, 019901(E) (2019).
- [54] S. Acharya *et al.* (ALICE Collaboration), *Phys. Rev. C* **101**, 044907 (2020).
- [55] W. Broniowski, M. Chojnacki, and L. Obara, *Phys. Rev. C* **80**, 051902(R) (2009).
- [56] G. Giacalone, F. G. Gardim, J. Noronha-Hostler, and J.-Y. Ollitrault, *Phys. Rev. C* **103**, 024910 (2021).
- [57] B. Schenke, C. Shen, and D. Teaney, *Phys. Rev. C* **102**, 034905 (2020).
- [58] G. Aad *et al.* (ATLAS Collaboration), *Phys. Rev. C* **90**, 024905 (2014).
- [59] M. Aaboud *et al.* (ATLAS Collaboration), *J. High Energy Phys.* **01** (2020) 051.
- [60] S. Acharya *et al.* (ALICE Collaboration), *J. High Energy Phys.* **07** (2018) 103.
- [61] S. Acharya *et al.* (ALICE Collaboration), *J. High Energy Phys.* **09** (2018) 006.
- [62] CMS Collaboration, CMS-PAS-HIN-12-011 (2012).
- [63] F. Gelis, G. Giacalone, P. Guerrero-Rodríguez, C. Marquet, and J.-Y. Ollitrault, [arXiv:1907.10948](https://arxiv.org/abs/1907.10948).
- [64] P. Carzon, S. Rao, M. Luzum, M. Sievert, and J. Noronha-Hostler, *Phys. Rev. C* **102**, 054905 (2020).
- [65] F. G. Gardim, G. Giacalone, and J.-Y. Ollitrault, *Phys. Lett. B* **809**, 135749 (2020).
- [66] H. Niemi, G. S. Denicol, H. Holopainen, and P. Huovinen, *Phys. Rev. C* **87**, 054901 (2013).
- [67] P. Carzon, M. D. Sievert, and J. Noronha-Hostler, *Phys. Rev. C* **105**, 014913 (2022).
- [68] P. Bożek, W. Broniowski, and J. Moreira, *Phys. Rev. C* **83**, 034911 (2011).
- [69] W. Ke, J. S. Moreland, J. E. Bernhard, and S. A. Bass, *Phys. Rev. C* **96**, 044912 (2017).
- [70] I. Altsybeev, <https://cds.cern.ch/record/2751545>.
- [71] ALICE Collaboration, ALICE-PUBLIC-2021-004 (2021).
- [72] W. Zhao, C. M. Ko, Y.-X. Liu, G.-Y. Qin, and H. Song, *Phys. Rev. Lett.* **125**, 072301 (2020).
- [73] A. Huss, A. Kurkela, A. Mazeliauskas, R. Paatelainen, W. van der Schee, and U. A. Wiedemann, *Phys. Rev. Lett.* **126**, 192301 (2021).
- [74] G. Nijs and W. van der Schee, [arXiv:2206.13522](https://arxiv.org/abs/2206.13522).