Computer exercise about elliptic flow

P. Christiansen (Lund University)

December 16, 2015

Abstract

The goal of this exercise is to develop a toy simulation that can test various ways of measuring the elliptic flow. The material has been prepared for the course MNXB01 where the compulsory part of the exercises is covered in two tutorials.
The material can be found at: http://www.hep.lu.se/courses/MNXB01/.

1 The quark gluon plasma and elliptic flow

Elliptic flow is a special phenomenon that occurs in collisions of heavy ions, e.g., at the Large Hadron Collider (LHC) at CERN on the boarder between Switzerland and France. In high energy heavy ion collisions a phase of deconfined quarks and gluons called the Quark Gluon Plasma (QGP) is produced, see Fig. 1. After production the QGP expands and cools off before it finally hadronizes. It is by studying the produced hadrons that we attempt to understand the properties of the QGP. Surprisingly it has been observed that the expansion of the QGP medium is the same as expected for an almost ideal fluid. This means that the ratio of the shear viscosity to entropy density is so small that the expansion is nearly reversible and indicates that there are very strong correlations in the QGP.

In a heavy ion collision the reaction plane is spanned by the impact parameter vector and the beam axis, see Fig. 1 right. When the impact parameter is small we talk about central collisions and when it is large we have peripheral collisions. In most collisions the volume of nucleons that are actually interacting in the two nuclei is not symmetric, see Fig. 1 right. The expansion is therefore asymmetric, being faster in the reaction plane than in the plane perpendicular to this. This gives rise to so-called elliptic flow: the distribution of particles in the final state is asymmetric in the azimuthal plane. The elliptic flow will clearly have a dependence on the impact parameter of the collisions, and therefore in practice one always analyzes it for narrow regions of centrality. In general the elliptic flow is largest in mid-central collisions where the impact parameter is of order the nuclear radius.

In hydrodynamical flow calculations the QGP is treated as a macroscopic medium and the result is the 3 dimensional matter density and associated flow velocity field. Because of this velocity boost
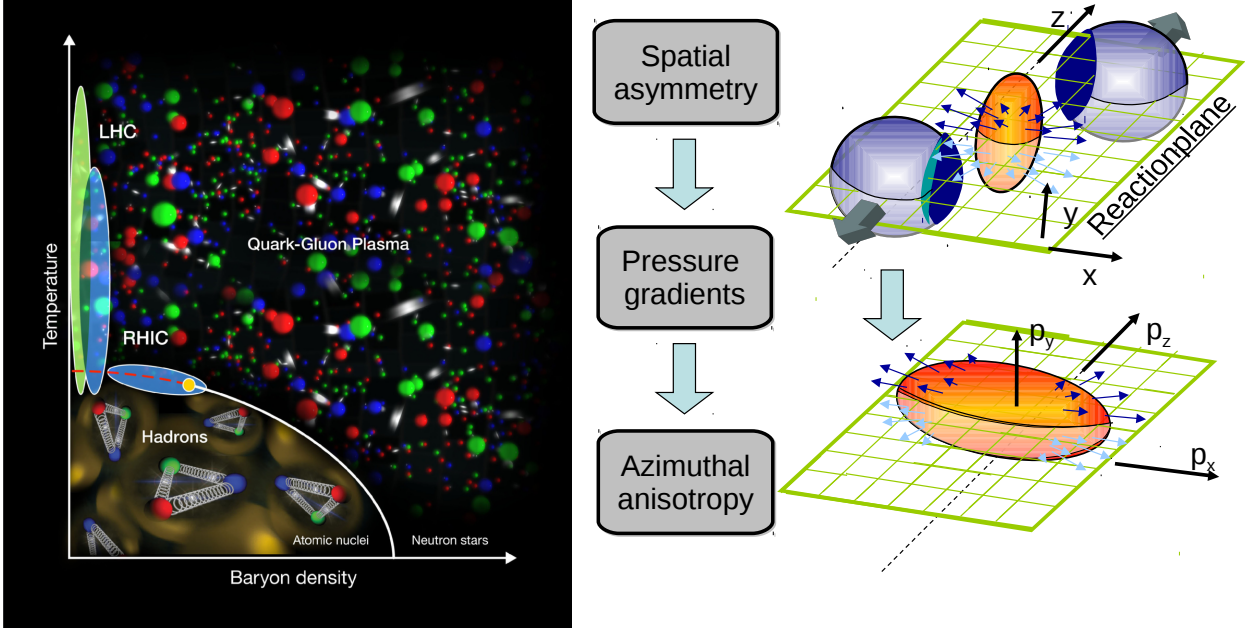
Figure 1: Left: QCD phase diagram [1]. Right: illustration of how elliptic flow arises in non-central heavy ion collisions.

the elliptic flow also depends on the transverse momentum of the particles and peaks at $p_T \approx 3$ GeV/c, see Figure 2 right. There is also a mass dependence such that heavier particles are more affected by the flow. This can easily be understood if one considers a particle with mass m produced at rest $(p = 0)$ and then boosts it with the transverse fluid flow velocity β where we find the laboratory $p_T = \beta \gamma m$.

2 How to measure elliptic flow

There are several ways to measure the elliptic flow. In this exercise we will test some of the different ways.

2.1 Event Plane method

In the event plane method we estimate the reaction plane in each collision and then align the particles relative to this plane [4]. In this way we can average over many events.

If the particles in an event is distributed according to:

$$f(\varphi) \propto 1 + 2v_2 \cos[2(\varphi - \Psi_2)], \tag{1}$$

then the event plane we want to determine is Ψ_2.

Figure 2: Left: Data from the STAR experiment used to estimate v_2 using the event plane method ($2 < p_T < 6\,\text{GeV}/c$). Note that Ψ_{plane} is what we call Ψ_2 in the text. The elliptic flow v_2 is ≈ 0.1 for the most central collisions (0-10%) and ≈ 0.2 for the most peripheral collisions (31-77%) [2]. Right: Results from the ALICE detector at LHC showing v_2 as a function of p_T extracted using 4-particle correlations [3]. The grey bands show comparable results from the STAR experiment at a beam energy that is ≈ 10 times smaller indicating that indeed we have a Quark Gluon Plasma with similar properties at both energies.

It can be shown (good optional exercise) that Ψ_2 can be estimated as:

$$\Psi_2 = \frac{1}{2}\tan^{-1}\left(\frac{\langle\sin(2\varphi)\rangle}{\langle\cos(2\varphi)\rangle}\right), \tag{2}$$

where it is smart to use the ROOT function `TMath::ATan2(y, x)`. The brackets $\langle x \rangle$ means that one should take the average of x over all tracks in a single event.

One then histograms all tracks relative to this plane in each event. After measuring over many events one fits the final histogram using a constant (normalization) times Eq. 1, see Fig. 2 left for an example.

The elliptic flow measured this way is often denoted $v_2\{EP\}$. It is in this authors opinion the most easy method to understand and one can easily visually check that one did not make a grave mistake, but is not as precise as the other methods.

2.2 2-Particle correlations

The Event Plane method is criticized for the need to first determine the event plane and then measure v_2 since the statistical precision with which we can determine the event plane event-by-event will affect the result (even one typically corrects for this using a resolution function). One can avoid this

by studying 2-particle correlations [1]:

$$
\begin{aligned}
\langle \cos[2(\varphi_1 - \varphi_2)]\rangle &= Re\langle e^{i2(\varphi_1 - \varphi_2)}\rangle \\
&= Re\langle e^{i2(\varphi_1 - \Psi_2 - \varphi_2 + \Psi_2)}\rangle \\
&\approx Re\left[\langle e^{i2(\varphi_1 - \Psi_2)}\rangle\langle e^{i2(\varphi_2 - \Psi_2)}\rangle\right] \\
&= \langle \cos[2(\varphi_1 - \Psi_2)]\rangle\langle \cos[2(\varphi_2 - \Psi_2)]\rangle \\
&= v_2^2,
\end{aligned}
\tag{3}
$$

where the average is over all pairs, and the assumption in line 3 is that there are no direct correlations between particle 1 and 2, but only indirect correlations through the common event plane Ψ_2.

In this way we determine the 2-particle correlation factor $\langle 2\rangle = \langle \cos[2(\varphi_1 - \varphi_2)]\rangle$ and we write $v_2\{2\} = \sqrt{\langle 2\rangle}$.

As there is no event plane determination the average indicated by the brackets are over all events! In the simulations we will have the same number of tracks in each event and so one can just average first in an event and then average over events. The quantity one wants to average is $\langle 2\rangle$.

2.3   2-Particle correlations using the Q-vector

To determine the 2-particle correlations we need a nested double loop, e.g, if we have 1000 tracks we loop for track 1 over the remaining 999 particles, for track 2 over the remaining 998 particles, and so on. Especially when one goes to 4-particle correlations and higher this becomes impossible due to the time it takes.

It turns out that one can determine $Q_n = \sum_{\text{tracks}} e^{in\varphi}$ with which it is easy to show that:

$$
\langle 2\rangle = \frac{|Q_2|^2 - M}{M(M-1)}
\tag{4}
$$

In this way one just have to loop one time over all tracks to calculate $\langle 2\rangle$.

3   Different Ways to Generate Random Numbers According to a Distribution

To be able to test the different methods we want to be able to generate random numbers that are distributed like Eq. 1. Here we shall not discuss how one can generate a so-called flat distribution of random numbers between 0 and 1, but show how one can use these to generate any distribution.

In the following example we always wants to generate numbers distributed according to $\sin x$ and we will show plots from the macro `generate_sinx.C` that is also made available to you.
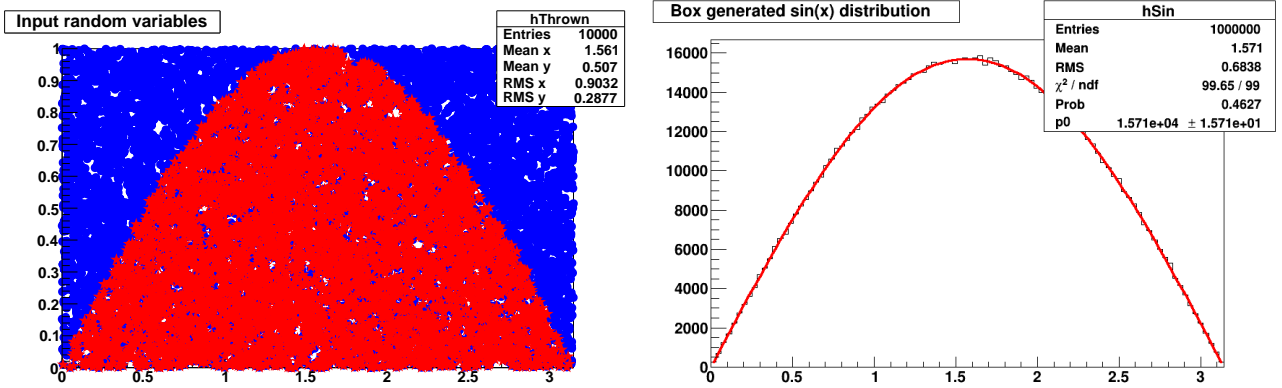
Figure 3: The box method. Left: 2-dimensional histograms showing all attempted (x, y) blue, and those accepted in red. Right: The distribution of accepted x values. The red function is a fit to the generated data.

3.1 The box method

This is a simple method that works as long as one generates random numbers in a restricted range, but is not always very efficient. We generate two random numbers x and y. We scale x so that it gives a random point in the restricted range we want to generate random numbers in. Now we scale y so it matches the range from 0 to the maximum value of the function $f(x)$ we want to generate. Now we accept x if $y < f(x)$ and reject x otherwise.

Figure 3 gives an example of this method.

3.2 The analytical

In some cases one can integrate and invert the probability density distribution. In this case, e.g.:

$$
\begin{aligned}
P(x) &= \frac{1}{2}\sin x \\
\int_0^\pi P(x')dx' &= 1 \\
y = \int_0^x P('x)dx' &= \frac{1}{2}[-\cos x']_0^x \\
&= \frac{1}{2}(1 - \cos x), \qquad \text{so that} \\
x &= \cos^{-1}(1 - 2y)
\end{aligned}
\tag{5}
$$

This means that if we generate a random number y between 0 and 1, then we should just apply the transformation $x = \cos^{-1}(1 - 2y)$ to get random numbers distributed according to $\sin x$.

From this method one will get a similar distribution as Figure 3 right, but one has a 100% efficiency.

---

[1] The biggest gain is when we go to higher order correlations, see Advanced Topic 2.

This method cannot always be used, because it is not always possible to invert a function. Sometimes it can be used even when x can take unbounded values, e.g., if we want to distribute according to an exponential distribution $\exp^{-x/k}$. There are also unbounded cases where one can combine this with the box method.

4   The histogram method

Figure 4: The histogram method. Left: Cumulative probability for 100 bins. Right: The distribution of generated x values. The red function is a fit to the generated data.

This method can be used to approximate any function in a restricted range and can also be used in the case where one only has a histogram, e.g., with real data and want to generate random numbers according to this.

We first choose a binning and then fill in each bin the cumulative function value in this bin. Then we normalize it to the total integral so that in the last bin we have the value 1.

Now it works in some sense exactly like a numerical approximation of the analytical method. We generate a random value z and we find the corresponding lowest bin where z is smaller than the histogram value. Now we use as a random number the x that corresponds to this bin. Figure 4 gives an example of this method.

This is the least elegant method, but it is quite simple to implement and can easily be generalized to 2d and 3d distributions, but one will of course have binning effects even if one can smooth out this.

Here it is meant to illustrate how we understand the basic principles of ROOT's generic function as explained in the ROOT class description [2]

```
Double_t TF1::GetRandom(Double_t xmin, Double_t xmax)

Return a random number following this function shape in [xmin,xmax]
```

---

[2]See http://root.cern.ch/root/html/TF1.html.

The distribution contained in the function fname (TF1) is integrated
over the channel contents.
It is normalized to 1.
For each bin the integral is approximated by a parabola.
The parabola coefficients are stored as non persistent data members
Getting one random number implies:
- Generating a random number between 0 and 1 (say r1)
- Look in which bin in the normalized integral r1 corresponds to
- Evaluate the parabolic curve in the selected bin to find
the corresponding X value.
The parabolic approximation is very good as soon as the number
of bins is greater than 50.

IMPORTANT NOTE
The integral of the function is computed at fNpx points. If the function
has sharp peaks, you should increase the number of points (SetNpx)
such that the peak is correctly tabulated at several points.

5 Exercises

In the following one can find the list of exercises for day 1 and for day 2.

The goal of day 1 is as a minimum to make a code that can generate a certain number of events, each with a certain amount of tracks, where each track is represented by a φ angle generated according to Eq. 1 using the box method. The events will be stored in a text file. Some optional exercises are suggested.

The goal of day 2 is as a minimum to make a code that can read in the tracks and analyze them using the 2-particle correlation method. Some optional exercises are suggested.

For the interested student, several advanced topics are included in the back that are difficult to solve, but shows important advanced aspects of flow analyses.

5.1 Day 1 (tutorial 7a)

The goal of the day 1 program is to make a simple event generator that can generate a track distribution with a certain elliptic flow v_2 and store them in a text file.

5.1.1 Exercise 1.1: study the example (compulsory)

Familiarize yourself with the example `generate_sinx.C`. Run the 3 different methods (box, analytical, histogram), and understand the differences and advantages. You can get the macro using the Linux command:
`wget http://www.hep.lu.se/staff/christiansen/generate_sinx.C`
Note that in the top of the macro there is information about how to load it and how to call the function.

5.1.2 Exercise 1.2: make a v_2/φ box generator (compulsory)

Extend the box method to be able to generate random numbers according to Eq. 1. In all cases we set $\Psi_2 = 0$. Figure 5 gives an example of how results from the final working program could look.
Technical help:

1. Copy `generate_sinx.C` to a new file: `EG_v2.C`

2. Remove all functions except the box function

3. Rename all places sinx to v2

4. Change the code to work for v_2. Be careful to also change ranges to match the new situation. The function should also take the new argument `Double_t v2`.

Make sure after each step that the code still works before going to the next step.

Figure 5: Generating φ according to $v_2 = 0.05$. Left: 2-dimensional histograms showing all attempted (φ, y) blue, and those accepted in red. Right: The distribution of accepted φ values. The red function is a fit to the generated data. Note that in the left we only used 10,000 events while in the right we used 100,000 events.

Note that in real data the total transverse momentum will have to be 0. This constrain we do not have here and so the data we generate does not necessary have to match even idealized collision data. For larger number of tracks this is likely not an issue.

As an optional exercise one can extend the code to use all 3 generation methods: box, analytical, histogram.

5.2   Exercise 1.3: make an event generator (compulsory)

First we need to convert this to a generator that now should take 3 arguments: `Int_t nEvents`, `Int_t nTracks` (in each event), and `Double_t v2`.

1. To speed up visualization only draw hAccepted and change the draw option to COLZ (you can even remove hThrown).

2. Add the `Int_t nTracks` argument and make a temporary storage for the φ angle, e.g.:

```
Double_t phi[nTracks]; // array to store phi angles
Int_t nt = 0; // track counter

while (nt < nTracks) {

  ....
  if(....) {
    // for each accepted track
    phi[nt] = x;
    nt++;
```

```
        }
    }
```

To test this we fill a histogram, `hPhi`, with the φ values in another loop

```cpp
    for(Int_t i = 0; i < nTracks; i++) {

        hPhi->Fill(phi[i]);
    }
```

3. Add the `Int_t nEvents` argument and make a loop over `nEvents`, e.g.:

```cpp
    for (Int_t ne = 0; ne < nEvents; ne++) {

        .... // here we generate nTracks
    }
```

Make sure after each step that the code still works before going to the next step.

5.2.1   Exercise 1.4: store the result in a text file (compulsory)

Here we want to output the generated data to a normal text file.

1. First just write the information to the standard output using the normal `cout` operator like this:

```
Event 0
nTracks 10
0 : 0.969496
1 : 4.96748
2 : 1.21711
3 : 4.87544
4 : 5.41607
5 : 3.08065
6 : 4.88619
7 : 5.34159
8 : 0.18106
9 : 5.36895
Event 1
nTracks 10
0 : 4.43049
1 : 5.82934
2 : 0.939175
3 : 4.88806
```

```
            ....
```

2. Now we want to open a file and redirect the output to the file. We need to include the `fstream` (input/output file stream class) header file.

```cpp
// open output file
ofstream file("phi_dist.dat");
....
// write to file
file << "Event " << ne << endl
....
// close file
file.close();
```

5.2.2 Solution to day 1

The solution to the day 1 exercises that should be mailed in is

- a macro called `EG_v2.C`

- the macro have to be compilable in ROOT without any warnings:
  `.L EG_v2.C+`

- the macro must have a function called:
  `void EG(Int_t nEvents, Int_t nTracks, Double_t v2)`

- the function should produce a histogram called hPhi that is drawn and shows the generated φ distribution

- the function should produce a text file called
  `phi_dist.dat`
  with the generated events.

5.3 Day 2 (tutorial 7b)

The goal of the day 2 program is to read back in the data from the output file and to analyze the data using the 2-particle correlation method described in Sec. 2.3.

5.3.1 Exercise 2.1: read back in the file (compulsory)

Make a new ROOT macro called `Analyze_v2.C` with a function:
`void analyze_v2_2particle()`.

1. Now we want to open a file and read back the input in a formatted way. We need to include the `fstream` (input/output file stream class) header file.

```
    Int_t nEvents = 0; // event counter
    string helpString; // help variable
    // open input file
    ifstream file("phi_dist.dat");
    Int_t eventNo = -1;

    while(file >> helpString >> eventNo) {

      cout << "Reading event : " << eventNo << endl;
      nEvents++;

      Int_t nTracks = -1;
      file >> helpString >> nTracks;
      cout << "Event contains " << nTracks << " tracks" << endl;

      Double_t phi[nTracks];
      Int_t    trackNo;

      for(Int_t nt = 0; nt < nTracks; nt++) {

        // Read back the phi values
        ....
      }

      // Here we want to analyze the phi values in the next exercise
    }
```

5.3.2 Exercise 2.2: extract v_2 (compulsory)

Now you have to implement the 2-particle correlation analysis.
To help with the 2-particle correlation function part I give here an example of how the loop can be done:

```
// Q vector (Sec. 2.3)
Double_t sum_cos2 = 0;
Double_t sum_sin2 = 0;
// 2-particle (Sec. 2.2)
Double_t sum_cos2_diff = 0;
for(Int_t i = 0; i < nTracks; i++) {

  // Compute Q vector
  sum_cos2 += TMath::Cos(2*phi[i]);
```

```
    sum_sin2 += TMath::Sin(2*phi[i]);

    // Compute 2-particle
    for(Int_t j = i+1; j < nTracks; j++) {

      sum_cos2_diff += 2*TMath::Cos(2*(phi[i]-phi[j]));
    }
  }

  // Average <2> for this event for Q-vector and 2-particle
  ....

  // Update <2> event average for Q-vector and 2-particle
  ....
```

And then one can compare the methods event-by-event and see they give exactly the same, e.g., by printing out:

```
v2 (average of 100 events) = 0.0493544
Q: v2 (average of 100 events) = 0.0493544
```

Note that one can only calculate the $v_2\{2\}$ for an event when $\langle 2 \rangle > 0$. Also note that the 2-particle method (without the Q vector trick) is extremely slow for large `nTracks`.

5.3.3 Exercise 2.3: statistical uncertainty and fitting (optional)

The idea here is to go through Advanced topic 3 with step 1 and step 2 done in a much simpler way.

Instead of step 1 and 2 make 10–40 different data files using the code developed on day 1. Analyze them using the code just made and make a histogram of $v_2\{2\}$. What does this histogram tell about the statistical precision of the measurement?

Likely one has to generate one file and analyze it, since we overwrite the files. Note that if the random number generator is initialized with the same `seed` the random values will follow the same pattern. To initialize the random generator with a random seed one has to call it like this:

`TRandom* randGen = new TRandom(0);`

Finally one can implement the fitter in step 3 following the instruction in Sec. 8.3. But first one have to implement the event plane method, see Sec. 2.1. To do the event plane method you need to loop two times: firstly, to estimate the event-plane, Ψ_2 and secondly to fill the histogram with $\varphi - \Psi_2$. Make sure that $\varphi - \Psi_2$ is in the actual range of your histogram (add or subtract 2π if it is outside this range).

5.3.4 Exercise 2.4: compare the two methods (optional)

Which method does the best job of calculating v_2?

How does this depend on `nTracks` and `Nevents`?

Bonus question: in the event plane method one has trivial correlations between each particle and the event plane. Since the event plane in each event just tries to maximize the v_2 one in general overestimates v_2. One can calculate the event plane for each track where one ignores the track itself. If one does this one instead always finds a too small v_2. Can you understand why? (Hint: study the event plane resolution in your code and think about what it does to v_2).

5.3.5   Solution to day 2

The solution to the day 2 exercises that should be mailed in is

- a macro called `Analyze_v2.C`

- the macro have to be compilable in ROOT without any warnings:
  `.L Analyze_v2.C+`

- the macro must have a function called:
  `void analyze_v2_2particle()`

- the function should produce a histogram called hPhi that is drawn and shows the φ distribution read in from the file `phi_dist.dat`

- the function should calculate $v_2\{2\}$ using the Q-vector method and write the result to the standard output
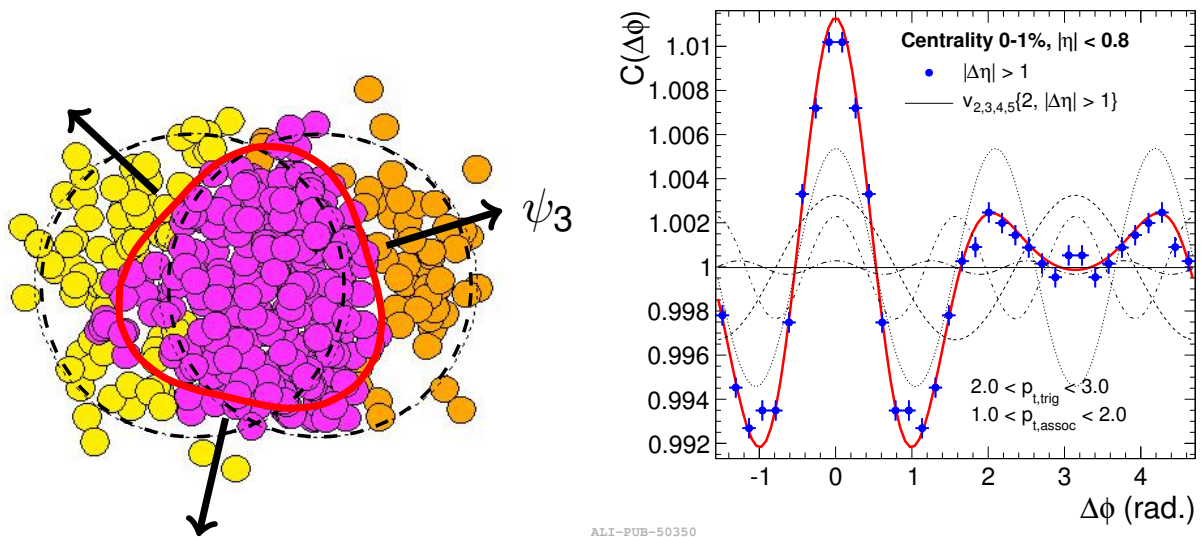
6 Advanced Topic 1: Triangular flow

Figure 6: Triangular flow. Left: fluctuations can give rise to a triangularity in the created medium. Right: by going to the 1% most central collisions one can suppress v_2 and directly observe the triangular correlations using 2-particle correlation techniques as maxima at 0, $2/3\pi$ and $4/3\pi$ radians (0, 120, and 240 degrees). Notice that the red curve is *calculated* from the 2-particle flow coefficients v_2, v_3, v_4, v_5 [6].

In 2005 the PHOBOS experiment at RHIC have shown that to understand the flow fluctuations observed in Au-Au and Cu-Cu collisions one needs to take into account fluctuations in the positions of nucleon-nucleon collisions. This means that the impact parameter plane is not always the best symmetry plane to use. In 2010 two experimental physicists, Alver and Roland [5], showed that the same fluctuations could give rise to triangular flow, v_3, see Fig. 6 Left:

$$f(\varphi) \propto 1 + 2v_3 \cos[3(\varphi - \Psi_3)], \tag{6}$$

The ALICE experiment at LHC confirmed this prediction by going to the most central collisions where the geometric elliptic flow is suppressed, see Fig. 6.

6.1 Step 1: extend your code to handle v_3 instead of v_2

In this part assume that $\Psi_3 = 0$.

This should be straight forward, but recall that v_3 has 3 peaks and in general one therefore needs $\cos 3\varphi$ instead of $\cos 2\varphi$ and so on.

15

6.2 Step 2: extend your code to handle v_2 and v_3 at the same time

In detailed simulations and confirmed by measurements the symmetry angles Ψ_2 and Ψ_3 are not correlated. Ψ_2 is geometrical and close to the impact parameter angle, while Ψ_3 is related to fluctuations. This might be necessary to implement to get the methods to work.

7 Advanced topic 2: 4-particle correlations

It is known that in addition to correlations through the event plane Ψ_2 there are also correlations from decays, e.g.: $\phi \to K^+ + K^-$. These correlations are called non-flow and can be removed by studying higher order correlation functions which is the goal of this advanced exercise.

7.1 Step 1: implement non-flow

An easy way to implement non-flow is for each particle (or some fraction of particles) to add one extra particle with exactly the same φ angle [3].

Do this in your code and check that you now estimate a wrong $v_2\{2\}$.

7.2 Step 2: implement the 4-particle correlation function

In the same way as we defined the 2-particle correlation factor we can define the 4-particle correlation factor:

$$\langle 4 \rangle = \langle \cos[2(\varphi_1 + \varphi_2 - \varphi_3 - \varphi_4)] \rangle. \tag{7}$$

It turns out that by using the 4-particle cumulant one can subtract the 2-particle correlations to v_2 so one only are left with genuine 4-particle correlations:

$$v_2\{4\} = \sqrt[4]{-\langle 4 \rangle + 2\langle 2 \rangle^2}. \tag{8}$$

In this way we expect to be insensitive to the non-flow we just added before.

The problem with $\langle 4 \rangle$ is that if you have e.g. 1000 tracks in each event then it will take forever to calculate the correlation as a nested loop (over i, j, k, l). The trick is to use the Q-vector. It can be shown that:

$$\begin{aligned}
\langle 4 \rangle = &\frac{|Q_2|^4 + |Q_4|^2 - 2 \cdot \mathrm{Re}[Q_4 Q_2^* Q_2^*]}{M(M-1)(M-2)(M-3)} \\
&- 2\frac{2(M-2) \cdot |Q_2|^2 - M(M-3)}{M(M-1)(M-2)(M-3)}.
\end{aligned} \tag{9}$$

This analytic result and many more can be found here [7].

Implement this estimate and show that now you obtain the correct v_2 independent of non-flow. While $v_2\{2\}$ converges fast and works well even for a few particles in general $v_2\{4\}$ converges much slower and one therefore needs to run with a lot of tracks in each event [4].

---

[3] The advantage of this method is not only that it is easy, but one can in fact calculate the exact effect on the 2-particle correlation analytic. If you do this, then try to understand the nTrack dependence (Hint: consider the number of pairs).

[4] The number of 4-particle associations grows approximately as (Ntracks)^4 so when calculating higher order correlations one therefore also has to be careful that this can impose a significant bias towards the subsample of the events with the most tracks

7.3 Final remarks about flow fluctuations

Due to fluctuations the v_2 varies event-by-event. As the 2- and 4-particle correlation functions measures v_2 to some power they are biased by these fluctuations. The easiest case to understand is the 2-particle correlation function. If $\sigma_{v2}^2 = \langle v_2^2 \rangle - \langle v_2 \rangle^2$, then it is clear that one in general with $v_2\{2\}$ measures (neglecting non-flow):

$$v_2\{2\}^2 = \langle v_2 \rangle^2 + \sigma_{v2}^2, \tag{10}$$

so that $v_2\{2\} \geq \langle v_2 \rangle$.

For 4-particle correlations one finds that when $\sigma_{v2} \ll \langle v_2 \rangle$ then:

$$v_2\{4\}^2 = \langle v_2 \rangle^2 - \sigma_{v2}^2, \tag{11}$$

so that $v_2\{4\} \leq \langle v_2 \rangle$.

This means that we can estimate the fluctuations as:

$$2\sigma_{v2}^2 = v_2\{2\}^2 - v_2\{4\}^2 \tag{12}$$

Use your code to study this relation (remove non-flow).

8 Advanced topic 3: the statistical uncertainty

It can sometimes be difficult to calculate the statistical uncertainty on a quantity that is measured in a complex way like $v_2\{2\}$. Here we will illustrate how it can be done by subdividing the total sample into smaller samples and then from the spread of those assign an uncertainty for the average.

8.1 Step 1: extend your code to make 1000s of analyzes

In this part you want to know the answer for a certain configuration of Ntracks and Nevents. To find this you can simply just change your code so it will make $N_{\text{experiments}}$. Then you can make a histogram with the results of e.g. $v_2\{2\}$ for each experiment and then in the end you will have a histogram which is approximately Gaussian and where the mean should be very close to the input v_2 and where the σ is the statistical uncertainty.

8.2 Step 2: implement an unbiased estimator for the mean

In this part we now want to divide the original Nevents sample into 10 smaller samples (Nevents/10) that we then analyze independently ($N_{\text{experiments}} = 10$). In this way we can obtain:

$$\sigma_{1/10}^2 = \frac{1}{N_{\text{experiments}} - 1} \sum_{i=1}^{N_{\text{experiments}}} (v_{2,i} - \langle v_2 \rangle)^2, \tag{13}$$

and then we estimate the statistical uncertainty for the full sample to be:

$$\sigma = \frac{\sigma_{1/10}}{\sqrt{N_{\text{experiments}}}}. \tag{14}$$

The important thing to notice is that one should divide by $N_{\text{experiments}} - 1$ and not $N_{\text{experiments}}$ in Eq. 13. This estimator for $\sigma_{1/10}$ is called the unbiased estimator because it takes into account that we use 1 degree of freedom to estimate also $\langle v_2 \rangle$. If you would use the generator v_2 in Eq. 13 then you should in fact use $N_{\text{experiments}}$ and not $N_{\text{experiments}} - 1$.

Test the performance of this method, i.e., compare to the result found in step 1.

Note that since we simulate only few events here (vs millions in real experiments) then the method is not so good. This is because $v_2\{2\}$ depends non-linearly on $\langle 2 \rangle$. The fluctuations in $\langle 2 \rangle$ must be rather symmetric and so if they are large this gives an asymmetric influence on $v_2\{2\}$, e.g. if it fluctuates down to 0 then $v_2\{2\}$ decreases by 100% while if it fluctuates up a factor 2 $v_2\{2\}$ only increases by 40% [5]. Due to the 10 times smaller statistics one also easier gets negative values for $\langle 2 \rangle$ that one has to decide how to handle. For this reason it is recommended to have a large number of events and a large v_2, e.g., 1000 events with 100 tracks and $v_2 = 0.1$.

Bonus question: what is the motivation to subdivide into 10 subsamples and not 3 or 50?

---

[5]One can actually do the σ analysis for $\langle 2 \rangle$ instead to reduce these problems.

8.3 Step 3: implement fitter

In the event plane method one actually obtains also a statistical uncertainty from the ROOT fitter.

In this step we want to implement a simple fit program and use that to fit the event plane data and estimate the statistical uncertainty.

If we normalize the event plane histogram properly we only need to fit the v_2:

$$f(\varphi) = 1 + 2v_2 \cos[2\varphi]. \tag{15}$$

The idea is to make a 1 dimensional histogram where on the x-axis we vary v_2 and on the y-axis we calculate χ^2:

$$\chi^2(v_2) = \sum_{\text{bins}} \frac{(f(\varphi) - (\text{hist value in bin}))^2}{(\text{hist uncertainty in bin})^2}. \tag{16}$$

ROOT will keep track of the statistical uncertainty for a histogram when you call the method `Sumw2()` just after creating it. The statistical uncertainty in each of the event plane histogram bins is in this case just: $\sqrt{N_{\text{entries}}}/\text{Normalization}$.

If you select a reasonable range of parameters for the guessed v_2 then you should now see a minimum in your histogram. This is the best estimate of v_2 given the current binning. Write a loop into your code that finds this minimum. Now it turns out that if the fit has good quality then the so-called reduced χ^2 should be close to 1. The reduced χ^2 is defines as χ^2/N_{dof}, where the N_{dof} is the number of degrees of freedoms given as the number of bins minus the number of fit parameters, so in this case it is number of bins minus 1.

The reduced χ^2 quantifies the actual deviation between the fitted curve and the data points and gives a measure of how similar they are compared to the actual statistical uncertainty of your measurements.

If the reduced χ^2 is much less than 1 then it means that your statistical uncertainty is too large or you have too many parameters in your fit. If the reduced χ^2 is much larger than 1 it means that your statistical uncertainty is too small and likely there are systematic uncertainties that you need to take into account.

Test if the reduced χ^2 is close to 1 in your case when you vary the number of histogram bins, the number of events, and/or the number of tracks.

The statistical uncertainty on the fitted v_2 is approximately the absolute distance in v_2 you have to step away from the minimum for the χ^2 to be larger by 1 (this step is in χ^2 and not in the reduced χ^2). Estimate this in your program and compare to what ROOT gets for the fit.

In more advanced fitters the fitting routine estimates a local derivative in χ^2 space around the current value and then uses that to find where to go next, in this way stepping towards the minimum. For situations where there are more than one minimum more advanced methods are needed.

In general things gets even more complicated (unstable and time consuming) the more parameters one has in the fit.

References

[1] B. V. Jacak and B. Muller, Science **337**, 310 (2012).

[2] C. Adler *et al.* [STAR Collaboration], Phys. Rev. Lett. **90**, 032301 (2003).

[3] K. Aamodt *et al.* [ALICE Collaboration], Phys. Rev. Lett. **105**, 252302 (2010).

[4] A. M. Poskanzer and S. A. Voloshin, Phys. Rev. C **58** (1998) 1671 [nucl-ex/9805001].

[5] B. Alver and G. Roland, Phys. Rev. C **81**, 054905 (2010).

[6] K. Aamodt *et al.* [ALICE Collaboration], Phys. Rev. Lett. **107**, 032301 (2011).

[7] A. Bilandzic, R. Snellings and S. Voloshin, Phys. Rev. C **83**, 044913 (2011).