22 September 2024
andrei.lupasco@cern.ch

RF manipulations to improve SPS spill quality in the MHz range

Andrei Lupasco
CERN, CH-1211 Geneva, Switzerland

Keywords: SPS, SHiP, longitudinal phase-space, de-bunching, phase displacement blow-up, optimization, evolutionary algorithms, Pareto fronts

Summary

The Super Proton Synchrotron (SPS) at CERN will be used to deliver beams of protons for the new Search for Hidden Particles (SHiP) experiment. SHiP is a proposed fixed target experiment that aims to search for new particles beyond the Standard Model of particle physics, as predicted by various models. During its nominal operation, SHiP will require a $\sim$1s spill of 4×10^{13} 400 GeV protons from the SPS via slow extraction. Moreover, there are also requirements for the quality of the spill, e.g., it has to be as uniform as possible at all time scales. This report will present two methods for manipulating the beam's longitudinal phase-space distribution prior to performing slow extraction. The aim is to achieve a distribution that maximizes spill quality for the SHiP experiment. First, we will formally define the objectives the methods are trying to achieve. Then, we proceed to optimize their parameters using different optimization techniques and compare their performance.

Contents

1 Introduction **3**

2 Spill Quality **3**
2.1 Slow Extraction . . . 3
2.2 Momentum Distribution . . . 4
2.3 Time-Offset Distribution . . . 4
2.4 Turns taken . . . 6

3 RF Manipulations **6**
3.1 De-bunching . . . 6

3.2 Phase displacement blow-up 6

4 Optimization **6**
4.1 Optimization goal 6
4.2 Optimization technique 8

5 Results **9**
5.1 De-bunching 9
5.2 Phase displacement blow-up 10
5.3 Comparison 11

6 Conclusion **12**

1 Introduction

The Super Proton Synchrotron (SPS) at CERN is a powerful circular accelerator that plays a key role in the chain of accelerators leading up to the Large Hadron Collider (LHC). Operational since 1976, the SPS can accelerate protons to energies of up to 450 GeV (giga-electron volts). Furthermore, the SPS also delivers beams to the CERN North Area (NA) for various fixed target experiments.

One new planned experiment that is expected to take place in the CERN North Area is the Search for Hidden Particles (SHiP) experiment [1]. SHiP is a proposed fixed target experiment that aims to search for new particles beyond the Standard Model of particle physics, as predicted by various models.

During its nominal operation, SHiP will require a $\sim$1s spill of 4×10^{13} 400 GeV protons from the SPS via slow extraction. Moreover, there are also requirements for the quality of the spill, e.g., it has to be as uniform as possible at all time scales.

This report will present two methods for manipulating the beam's longitudinal phase-space distribution prior to performing slow extraction. The aim is to achieve a distribution that maximizes spil quality for the SHiP experiment. First, we will formally define the objectives the methods are trying to achieve. Then, we proceed to optimize their parameters using different optimization techniques and compare their performance.

2 Spill Quality

2.1 Slow Extraction

The SPS provides beams to the North Area by employing the process known as slow extraction. By exciting a third-integer transverse resonance in a controlled manner, particles are slowly "peeled off" from the circulating beam and extracted towards the experimental area.

The theoretical details on slow extraction can be found in [2]. In this report, a simplified approach is employed to model the process. In the SPS, particles are pushed towards the resonance by ramping all magnets in the lattice simultaneously, which is equivalent to sweeping the machine's reference momentum. This is illustrated in the Steinbach diagram in Figure 1. It can be seen that, for the nominal SPS emittance, the resonance has a characteristic stop-band width of $\hat{\delta}_{stopband,rms} = 5.4 \times 10^{-5}$. In our model, we will assume that particles within such a width are extracted simultaneously and we will refer to such a set of particles as a "momentum slice". Additionally, we will assume that the reference momentum is ramped linearly over 1 second. This leads to the following equation describing the time-offset coordinate of a particle at the moment of extraction:

$$\tau_{extracted} = \tau_0 + \eta\delta_0 \frac{N_{max}}{\delta_{max} - \delta_{min}} (\delta_{mean} - \delta_{min}), \tag{1}$$

where τ_0 is the time-offset coordinate of the particle at the start of the extraction, δ_0 is the particle's momentum offset, δ_{min} and δ_{max} are the minimum and maximum momentum offsets in the beam, δ_{mean} is the midpoint of the momentum slice the particle belongs to, $N_{max} = \frac{1}{23\times10^{-6}}$ is the number of turns taken by the slow extraction process.

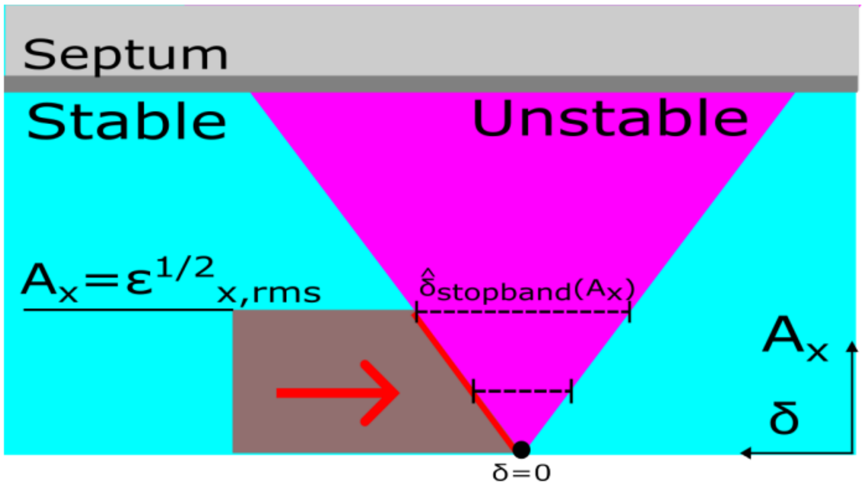

Figure 1: Steinbach-diagram illustration of COSE [4]. The horizontal arrow indicates that the beam (brown) is pushed from the stable region (blue) into the unstable region (fuchsia). The red band indicates the strip of particles that is resonant at a given instant. The momentum spread of such a strip is known as the r.m.s. momentum stop-band. Image source: [2].

2.2 Momentum Distribution

The SPS magnetic ramp is modulated by an undesired ripple. This introduces variations in the flux that can lead to the loss of valuable data for the experiments' acquisition systems.

The impact of ripple can be minimized by starting with a wide momentum distribution. Moreover, it is important for such a distribution to be as uniform as possible so that the flux is not modulated as the reference momentum is linearly ramped. In reality, corrections can be performed by using a non-linear ramp computed via a feed-forward approach. Still, this correction is easier the more uniform the initial distribution is.

To express the uniformity and the width of the momentum offset distribution as a metric we want to optimize, we will use the Kullback-Leibler divergence (which we want to minimize) between the resultant momentum distribution and a uniform distribution covering the whole SPS bucket. The KL divergence is defined as follows:

$$D_{KL}(P||Q) = \sum_i P(i) \ln \frac{P(i)}{Q(i)}, \tag{2}$$

where P is the resultant momentum distribution and Q is the desired uniform distribution.

2.3 Time-Offset Distribution

The perturbations described in the previous section affect the flux on timescales longer than the revolution period (23μs for the SPS). However, the flux is also modulated on shorter timescales, which can lead to problems for the experiments, such as increased combinatorial background.

Therefore, it is also important to uniformize the time-offset distribution inside the ring. More precisely, the distribution needs to be uniformised individually for all momentum slices of width $\sim \delta_{stopband,rms}$, since this represents the slice of particles that will be extracted from the machine simultaneously.

Therefore, it is possible to informally define spill quality as the level of uniformity of the intensity profile delivered to the user. In an ideal world, we would like to extract a perfectly rectangular spill, whose particle rate $I(t)$ would be described by:

$$I(t) = \begin{cases} I_0/T_s, & \text{if } t \in [0, T_s] \\ 0, & \text{otherwise,} \end{cases} \tag{3}$$

where T_s is the spill time and $I_0 = \int_0^{T_s} I(t)dt$ is the total integrated intensity.

Unfortunately, this is not physically attainable, since it would require a 'crystalline' beam, i.e., a particle lattice with predictable spacing. In a realistic 'hot' beam, a better model to describe the spill is a Poisson process of mean rate λ:

$$I(t) = \begin{cases} Pois(\lambda = I_0 \Delta t/T_s), & \text{if } t \in [0, T_s] \\ 0, & \text{otherwise,} \end{cases} \tag{4}$$

where δt is the chosen time-binning. We chose $\Delta t = 1ns$ for our simulations.

In order to quantify the proximity of the spill as a metric we could optimize, we will perform statistical tests on the time-offset distributions across momentum slices to determine if they follow a Poisson process.

For a given momentum slice, let $X_1, X_2, ..., X_N$ be discrete random variables denoting the number of particles in each bin after binning the slice into bins of width $\Delta t = 1ns$. The null hypothesis that we will test against states that $X_i \sim Pois(\lambda = I_0 \Delta t/T_s) \forall i$, where I_0 is the number of particles and T_s is the time-offset width of the momentum slice. If this holds, then $E[X_1] = E[X_2] = ... = E[X_N]$, i.e., in expectation each member of the family of random variables should have he same value. Thus, for a single member of the family, we compute the p-value that it follows the above-mentioned Poisson distribution by using the "Poisson means test", also known as the E-test [5]. Finally, for a given simulation, the time-offset metric T that we want to maximize is computed as follows:

$$T = \sum_{i \in \text{momentum slices}} \frac{1}{\text{\# of bins}} \sum_{j \in \Delta t-\text{width bin}} \ln(\text{Poisson means test}(n_{j,i}, N_i, n_i, N_i)), \tag{5}$$

where N_i is the number of particles in the momentum slice i, $n_i = N_i/\#$ of bins is the average number of particles across the time-offset bins of the slice and $n_{j,i}$ is the number of particles in the j-th time-offset bin of the i-th momentum slice. We take the natural logarithm of the p-values for better numerical stability, as maximizing the sum of logarithms is equivalent to maximizing the product of p-values, but the former is more numerically stable. We also take the average of logs of p-values across the bins for each momentum slice to better compare the methods, since the phase displacement blow-up method produces momentum slices with a much wider time-offset distribution, and thus significantly more bins. To experienced readers, this formula may be familiar to Fisher's combined probability test for a family of random variables.

2.4 Turns taken

Last but not least, we want to minimize the number of turns taken to perform the RF gymnastics, in order to maximize the number of spills delivered to the experiment and the overall cost efficiency of the operation.

3 RF Manipulations

3.1 De-bunching

The first method we will consider for improving the spill quality is the de-bunching method. This is also the method currently employed in operation for other experiments at the SPS. It aims to increase the momentum spread of the bunches via a RF phase jump and the time offset spread via de-bunching (waiting). More specifically, the method consists of the following steps:

1. Arrival at flattop: The beam is accelerated to the desired energy. Particles are performing stable oscillations around the stable fixed point.

2. Phase jump: The RF-waveform phase is shifted almost instantaneously by π, causing the particles to follow the bucket separatrix. This stretches the momentum distribution of the beam.

3. Re-phasing & bunch rotation: The RF-waveform phase is shifted back, bringing the particles back to the stable fixed point. This slightly increases the time offset spread of the beam, but also modifies the momentum distribution.

4. De-bunching: The RF voltage is turned off and the particles are left to drift for a certain amount of time. This increases the time offset spread of the beam.

3.2 Phase displacement blow-up

In the phase displacement blow-up method [8], we sweep the empty bucket through the beam by ramping f_{RF}. If we seep fast enough to eliminate the empty bucket area, the momentum distribution is blown-up with no mean displacement, as can be seen in Figure 2.

4 Optimization

4.1 Optimization goal

Let P be a set of parameters that define a configuration of the SPS machine, agnostic of the method used to extract the beam. When focusing on a specific method, we will denote the set of parameters as P_D for the de-bunching method and P_B for the phase displacement blow-up method. For a given set of parameters P, we can compute a vector of quality metrics $\vec{q}$ that describe the quality of the spill. This is achieved by simulating the manipulations and

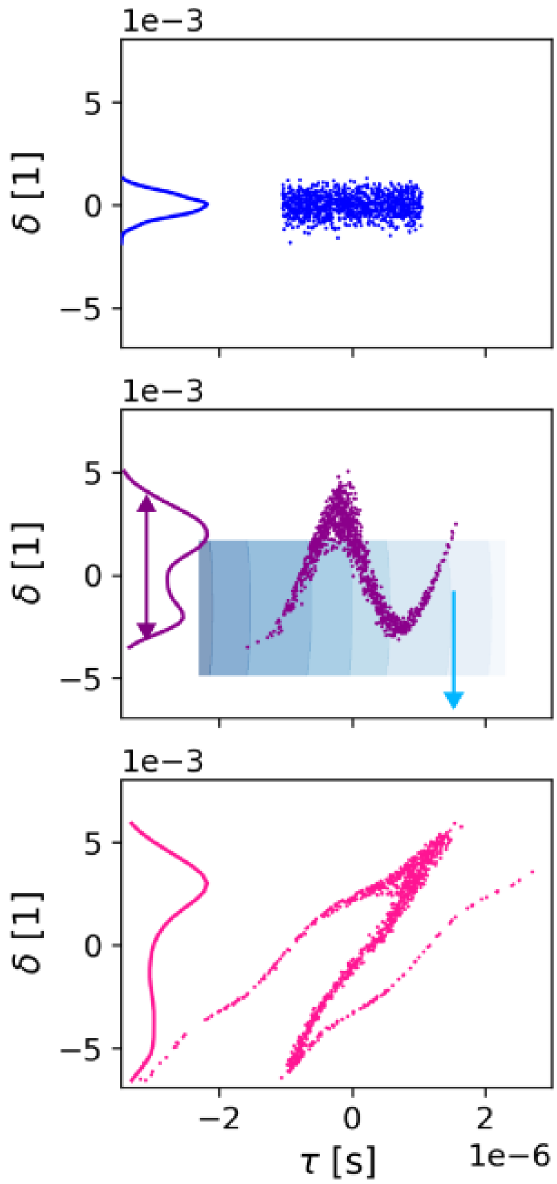

Figure 2: Simulation of phase displacement blow-up in the Proton Synchrotron at CERN. he momentum distribution is shown before (top), during (centre) and after (bottom) the manipulation, with the blue arrow and colour map (centre) showing the direction and location of the zero-area 'empty bucket', re- spectively. Image source: [8].

the extraction of the beam using the *henontrack* [2] simulation software, which is a Python package.

Let $\vec{q}$ be a vector of quality metrics assigned to some parameters P via simulating the configuration in henontrack. For the SHiP experiment, we will consider the following quality metrics:

- q_1: The KL divergence between the resultant momentum distribution of the simulation and the desired uniform distribution covering the whole SPS bucket.
- q_2: The p-value of observing the resultant time offset distribution of the simulation, under the null hypothesis that the spill follows a Poisson process.
- q_3: The number of turns taken to perform the RF gymnastics. This includes all the operations performed on the beam, such as the de-bunching or the phase displacement blow-up. However, this does not include the time taken by the slow extraction process itself, as this is a fixed parameter.

Since $\vec{q}$ is a vector of quality metrics and thus there does not exist a total order on the set of all possible $\vec{q}$, we will use a Pareto order to compare the quality of different configurations. More formally, we define a partial order $\preceq$ on the set of all possible $\vec{q}$ as follows:

$$\vec{q} \preceq \vec{r} \iff q_1 \leq r_1 \land q_2 \geq r_2 \land q_3 \leq r_3. \tag{6}$$

We say that $\vec{q}$ Pareto dominates $\vec{r}$ if $\vec{q} \preceq \vec{r}$. Extending this definition to the sets of parameters, we say that a set of parameters P Pareto dominates a set of parameters R if the vector of quality metrics $\vec{q}$ assigned to P Pareto dominates the vector of quality metrics $\vec{r}$ assigned to R.

Finally, we are interested in finding the set of configurations that are extremes in the Pareto order. These are the configurations that are not Pareto-dominated by any other configuration. We will refer to these as the Pareto front. Obviously, these configurations are the candidates to be used in the SPS machine for the SHiP experiment.

4.2 Optimization technique

We have employed several optimization techniques to find the Pareto front of the set of parameters P for the de-bunching and phase displacement blow-up methods. This includes Bayesian optimization with the qNEHVI acquisition function specifically designed for finding Pareto fronts, as well as adversarial machine learning. However, we have found that by far the best results are obtained by using genetics algorithms. More specifically, the algorithm SPEA2 [9] has been used to find the Pareto front of the set of parameters P for both methods.

The SPEA2 algorithm is a genetic algorithm constructed for finding Pareto fronts. It achieves this by maintaining a small archive of non-dominated individuals, which are used to guide the evolution of the population. On each generation, the archive is updated with the best individuals from the current population. If there are not enough individuals to fill in the archive, the algorithm will also add dominated candidates based on how many individuals they dominate. If, however, there are too many non-dominated non-dominated individuals, the algorithm will keep only a subset of them, the one with the best crowding distance. Since we have not found any Python implementation of the SPEA2 algorithm, we have implemented it ourselves [7]. It can evaluate the quality metrics of a set of parameters P in parallel, which is crucial for the optimization process.

The whole optimization loop can be summarized as follows:

Algorithm 1: Optimization loop

Data: the size of the population, the size of the archive, the number of generations
Result: the Pareto front

Initialize the population with random individuals;
for *each generation* **do**
  for *each individual in the population* **do**
    Evaluate the quality metrics of the individual using the henontrack package;
  end
  Update the archive with the best individuals from the population;
  Generate the new population by mutating an individual from the archive selected by a binary tournament;
end
Return the archive;

---

5 Results

5.1 De-bunching

The de-bunching method has been optimized using the SPEA2 algorithm. The genomes of the individuals were composed of the following parameters:

- Amount of turns in the phase jump stage, range: [0, 150]
- Amount of turns in the re-phasing (rotation) stage, range: [0, 50]
- The voltage of the 200 MHz RF cavity, range: $[0,\ 7.3 \times 10^{-3}$ GV]
- The voltage of the 800 MHz RF cavity, range: $[0,\ 1.4 \times 10^{-3}$ GV]

As the first step of the simulation, the beam arrives at the flattop in a fixed manner independent of the above parameters. The parametrized voltages are applied only during the phase jump and re-phasing stages. We have also investigated the effect of allowing different voltages for the cavities during these stages, but the results were not significantly different from the fixed voltage case.

The amount of turns taken for the de-bunching stage was not a free parameter. Instead, it was computed as the number of turns required for the bunches to overlap in their time offset in all momentum slices after the slow extraction procedure. This number was computed using the following formula:

$$N_{\text{debunch\_turns}} = \frac{f_0}{\eta\bar{\delta}}(\tau^2_{\text{min}} - \tau^1_{\text{max}}), \tag{7}$$

where f_0 is the revolution frequency of the SPS, η is the momentum slippage factor of the SPS, $\overline{\delta}$ is the width of a momentum slice equal to 5.4×10^{-5}, $\tau^1_{\max}$ is the maximum time deviation of a particle from the first bunch and momentum slice, and $\tau^2_{\min}$ is the minimum time offset of a particle from the second bunch and first momentum slice, both after the slow extraction. The values of $\tau^1_{\max}$ and $\tau^2_{\min}$ were computed using the henontrack package.

The optimization process was run for 50 generations with a population size of 2000 and an archive size of 500. The mutation of an individual selected by a binary tournament was performed by modifying each free parameter with a probability of 0.5 by adding a random value from a uniform distribution in the range $[-0.1 * (\text{max value} - \text{min value}), 0.1 * (\text{max value} - \text{min value})]$, where the min and max values are the bounds of the parameter. If the new value of the parameter was outside the bounds, it was clipped to the closest bound.

After the optimization process, the de-bunching method is showing promising results. One example can be seen in Figure 3. The Pareto fronts are shown in Figures 4 and 5. The results show that the momentum and time metrics are positively correlated, meaning that improvements in one metric generally lead to improvements in the other. On the other hand, the time metric is an increasing function of the number of turns taken to perform the manipulations. While this result does not come as a surprise, it is still important to note that we want to minimize the number of turns taken to maximize the number of spills delivered to the experiment, therefore one has to choose a configuration that balances the two metrics. The same plot also indicates that the time metric saturates at around 1500 turns (this excludes the arrival at the flattop), meaning that this would be a good choice for the number of turns taken to perform the manipulations. This is also the number of turns usually required by the bunches to de-bunch in order for them to overlap in time in all momentum slices.

5.2 Phase displacement blow-up

The phase displacement blow-up method has been optimized in a similar environment as the de-bunching method. The genomes of the individuals were composed of the following parameters:

- Amount of turns in the phase jump stage, range: [0, 150]
- Amount of turns in the re-phasing (rotation) stage, range: [0, 50]
- The voltage of the 200 MHz RF cavity during the phase jump and re-phasing, range: [0, 7.3×10^{-3} GV]
- The voltage of the 800 MHz RF cavity during the phase jump and re-phasing, range: [0, 1.4×10^{-3} GV]
- Γ, called "sine of the stable phase". For a given RF voltage, it quantifies the speed at which the RF frequency is varied. Larger Gamma corresponds to a faster frequency sweep [3], range: [1, 6]
- The voltage of the 800 MHz RF cavity during the phase displacement blow-up stage, range: [0.5×10^{-3}, 1.4×10^{-3}]

- The number of sweeps, range: [1, 20]
- The amplitude of the sweeps, in terms of the bucket height, range: [0.5, 2.0]

Similarly to the de-bunching method, the beam first arrives at the flattop, after which we perform a phase jump and re-phasing for a few turns each. Only after this did we perform the phase displacement blow-up manipulations, during which the voltage of the 800MHz cavity was parametrized, while the voltage of the 200MHz cavity was fixed to 0.

The ranges for the parameters were initially much larger, but after a few optimization trials, we have narrowed them down to the values presented above, for the reason that the good individuals were always found clustered in the presented ranges.

It is worth noting that we have also experimented with non-linear RF schedules for the phase displacement blow-up stage (as can be seen in Figure 8), but the results were not significantly different from the linear schedules (which can be seen in the presented examples). This is why we have decided to use linear schedules for the optimization process.

The optimization process was the same as for the de-bunching method, namely 50 generations with a population size of 2000 and an archive size of 500. A non-dominated individual can be seen in Figure 6.

The results of the optimization process for the phase displacement blow-up method are shown in Figures 9 and 10. The method exhibits similar behavior to the de-bunching method. The metrics tend to be in the same ranges. However, for this method the time metric gets saturated slower, at around 2500 turns. This way, according to the simulations, the de-bunching method should be preferred over the phase displacement blow-up method, as it requires fewer turns to perform the manipulations.

5.3 Comparison

The optimization process has yielded promising results for both methods. Their respective momentum and time metrics have landed in the same ranges, indicating that the methods are comparable in terms of spill quality. For the non-dominated individuals from both methods, the resultant $\frac{\Delta p}{p}$ distribution covers the whole bucket and is close to uniform. The same can be said about their time-offset distributions. However, the number of turns taken to perform the manipulations is significantly lower for the de-bunching method, which makes it the preferred choice for the SHiP experiment.

It is important to note that from the experiments in the SPS, it has been observed that the beam suffers from re-bunching due to collective effects [6], which are not accounted for in our single-particle simulations. This may lead to the beams not overlapping in time offset in all momentum slices, which would require more turns to perform the manipulations. In order to address this issue, we have tried to imitate this behavior by lowering the momentum slippage factor η in the simulations. The results of these experiments can be seen in Figures 11, 12 and 13. The results indicate that depending on the value of η, and thus the bunching factor unaccounted for in the simulations, the number of turns taken to perform the manipulations can increase significantly. Thus, if the bunching factor turns out to be significant, the phase displacement blow-up method may be the preferred choice.

6 Conclusion

In this report, we have presented two methods for manipulating the beam's longitudinal phase-space distribution prior to performing slow extraction. The aim was to achieve a distribution that maximizes spill quality for the SHiP experiment. We have defined the objectives the methods are trying to achieve and optimized their parameters using different optimization techniques. The results have shown that the de-bunching method is the preferred choice for the SHiP experiment, as it requires fewer turns to perform the manipulations. However, if the bunching factor turns out to be significant, the phase displacement blow-up method may be the preferred choice.

References

[1] Sergey Alekhin, Wolfgang Altmannshofer, Takehiko Asaka, Brian Batell, Fedor Bezrukov, Kyrylo Bondarenko, Alexey Boyarsky, Ki-Young Choi, Cristóbal Corral, Nathaniel Craig, David Curtin, Sacha Davidson, André de Gouvêa, Stefano Dell'Oro, Patrick deNiverville, P S Bhupal Dev, Herbi Dreiner, Marco Drewes, Shintaro Eijima, Rouven Essig, Anthony Fradette, Björn Garbrecht, Belen Gavela, Gian F Giudice, Mark D Goodsell, Dmitry Gorbunov, Stefania Gori, Christophe Grojean, Alberto Guffanti, Thomas Hambye, Steen H Hansen, Juan Carlos Helo, Pilar Hernandez, Alejandro Ibarra, Artem Ivashko, Eder Izaguirre, Joerg Jaeckel, Yu Seon Jeong, Felix Kahlhoefer, Yonatan Kahn, Andrey Katz, Choong Sun Kim, Sergey Kovalenko, Gordan Krnjaic, Valery E Lyubovitskij, Simone Marcocci, Matthew Mccullough, David McKeen, Guenakh Mitselmakher, Sven-Olaf Moch, Rabindra N Mohapatra, David E Morrissey, Maksym Ovchynnikov, Emmanuel Paschos, Apostolos Pilaftsis, Maxim Pospelov, Mary Hall Reno, Andreas Ringwald, Adam Ritz, Leszek Roszkowski, Valery Rubakov, Oleg Ruchayskiy, Ingo Schienbein, Daniel Schmeier, Kai Schmidt-Hoberg, Pedro Schwaller, Goran Senjanovic, Osamu Seto, Mikhail Shaposhnikov, Lesya Shchutska, Jessie Shelton, Robert Shrock, Brian Shuve, Michael Spannowsky, Andy Spray, Florian Staub, Daniel Stolarski, Matt Strassler, Vladimir Tello, Francesco Tramontano, Anurag Tripathi, Sean Tulin, Francesco Vissani, Martin W Winkler, and Kathryn M Zurek. A facility to search for hidden particles at the cern sps: the ship physics case. *Reports on Progress in Physics*, 79(12):124201, October 2016.

[2] Pablo Arrutia. Advanced RF techniques for CERN's future slow-extracted beams, 2024. Presented 06 Mar 2024.

[3] Pablo A. Arrutia Sota, Philip N. Burrows, Matthew A. Fraser, and Francesco M. Velotti. Millisecond burst extractions from synchrotrons using rf phase displacement acceleration. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 1039:167007, 2022.

[4] V Kain, F M Velotti, M A Fraser, B Goddard, J Prieto, L S Stoel, and M Pari. Resonant slow extraction with constant optics for improved separatrix control at the extraction septum. *Phys. Rev. Accel. Beams*, 22(10):101001, 2019.

[5] K. Krishnamoorthy. A more powerful test for comparing two poisson means. *Journal of Statistical Planning and Inference*, 119:23–35, 01 2004.

[6] A. Lasheen, T. Argyropoulos, T. Bohl, J. F. Esteban Müller, H. Timko, and E. Shaposhnikova. Beam measurement of the high frequency impedance sources with long bunches in the cern super proton synchrotron. *Phys. Rev. Accel. Beams*, 21:034401, Mar 2018.

[7] Andrei Lupasco. A Python implementation of the SPEA2 algorithm available at: https://github.com/JauntyKnight/spea2.

[8] P A Arrutia Sota, M A Fraser, F M Velotti, P N Burrows, V Kain, G Papotti, R Piandani, F Roncarolo, A Spierer, and M Vadai. RF techniques for spill quality improvement in the SPS. *JACoW IPAC*, 2023:MOPA116, 2023.

[9] Eckart Zitzler, Marco Laumanns, and Lothar Thiele. Spea2: Improving the strength pareto evolutionary algorithm. Report, Zurich, 2001-05.

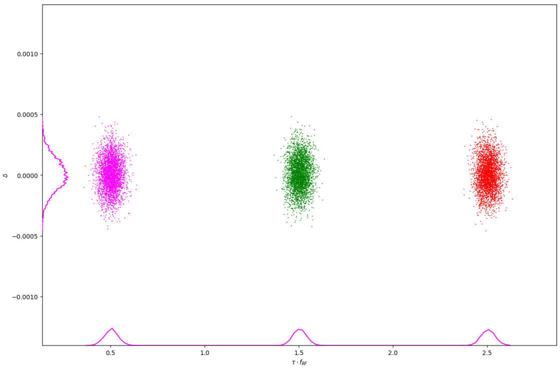

(a) Arrival at flattop. Turns 0 to 300.

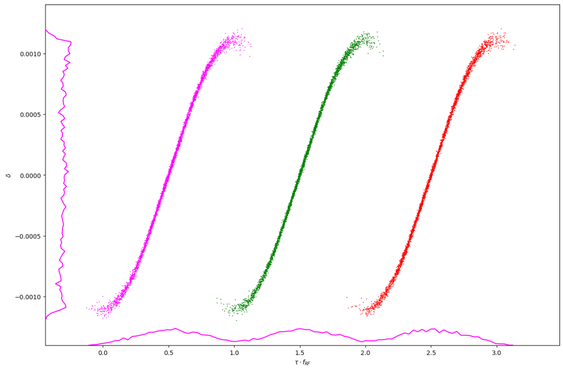

(b) Phase jump. Turns 300 to 388.

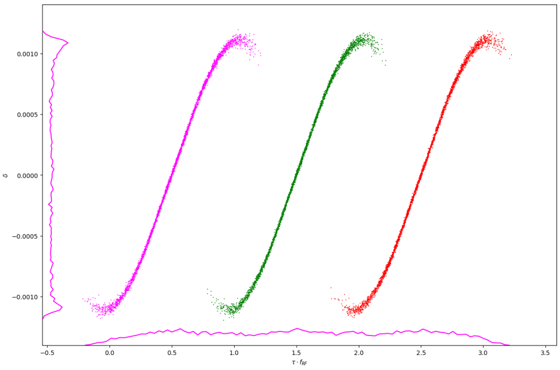

(c) Re-phasing. Turns 388 to 396.

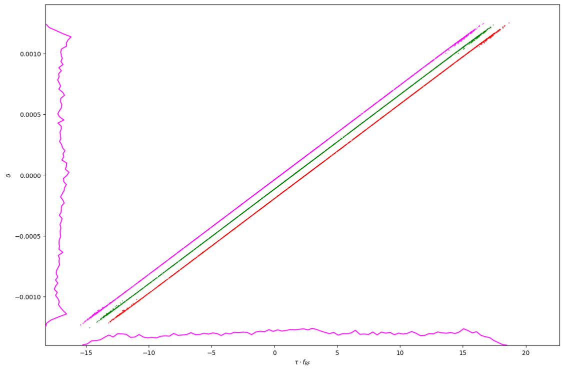

(d) De-bunching. Turns 396 to 1860.

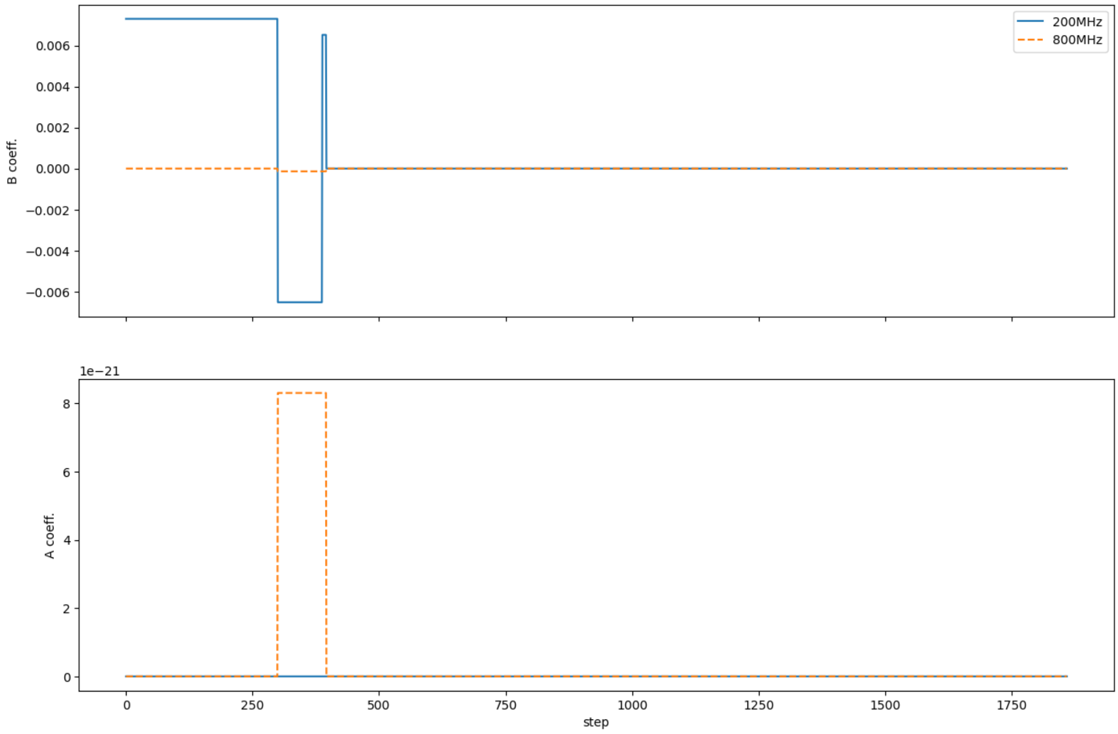

(e) The RF schedule of the example.

Figure 3: The longitudinal phase space of the beam during the de-bunching method throughout the stages. Different bunches have different colors. At the bottom of the figure, the RF schedule is shown.

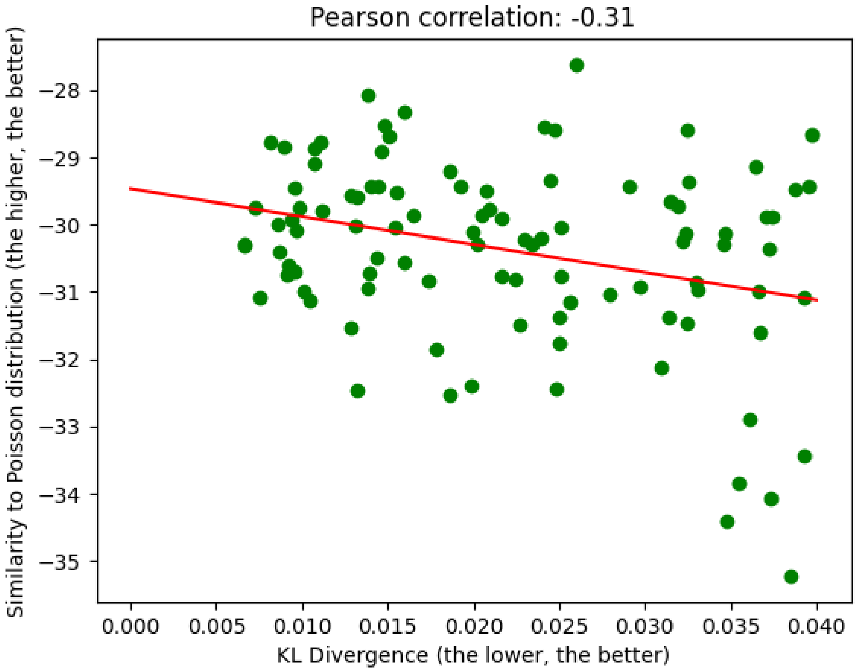

Figure 4: The Pareto front of the de-bunch method between the momentum and time metrics.

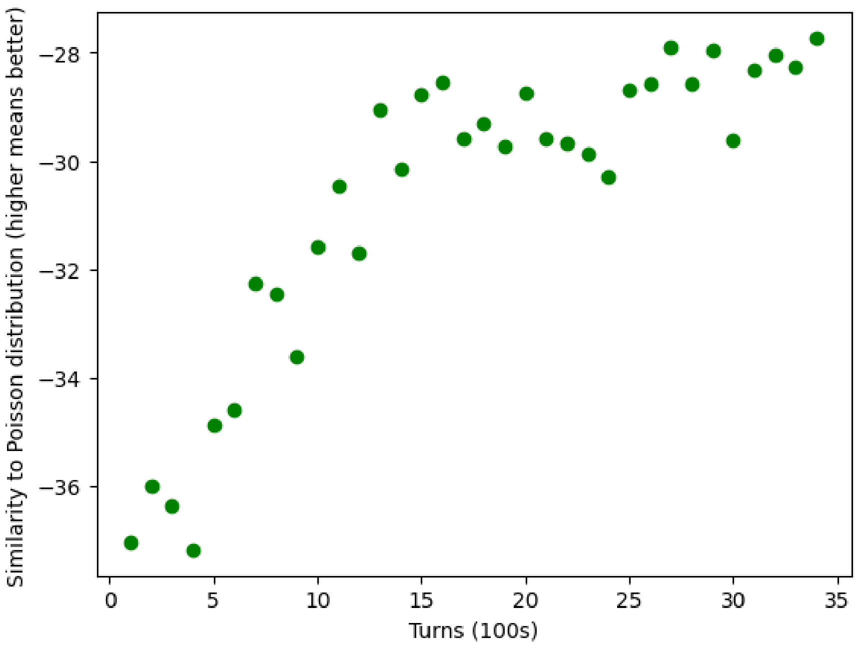

Figure 5: The Pareto front of the de-bunch method between the time metric and number of turns taken.

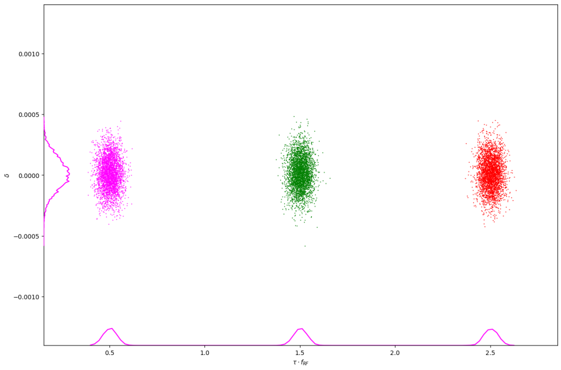

(a) Arrival at flattop. Turns 0 to 300.

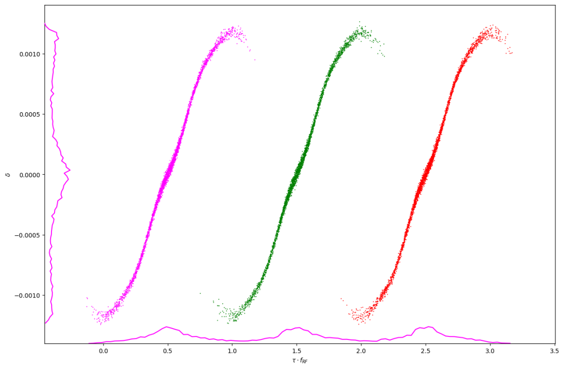

(b) Phase jump. Turns 300 to 394.

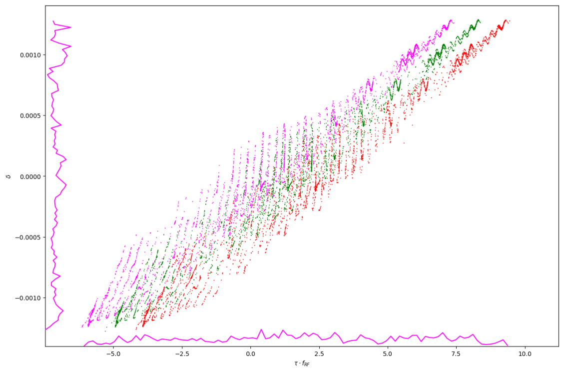

(c) Phase displacement blow-up. Turns 394 to 984.

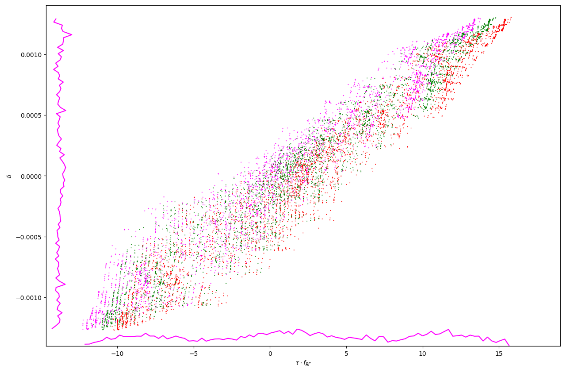

(d) Phase displacement blow-up. Turns 984 to 1574.

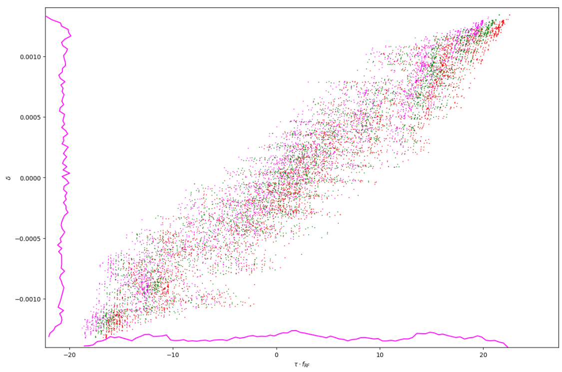

(e) Phase displacement blow-up. Turns 1574 to 2164.

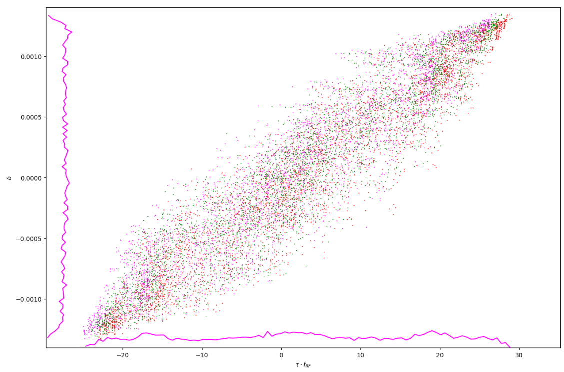

(f) Phase displacement blow-up. Turns 2164 to 2754.

Figure 6: The longitudinal phase space of the beam during the phase displacement blow-up method throughout the stages. Different bunches have different colors.

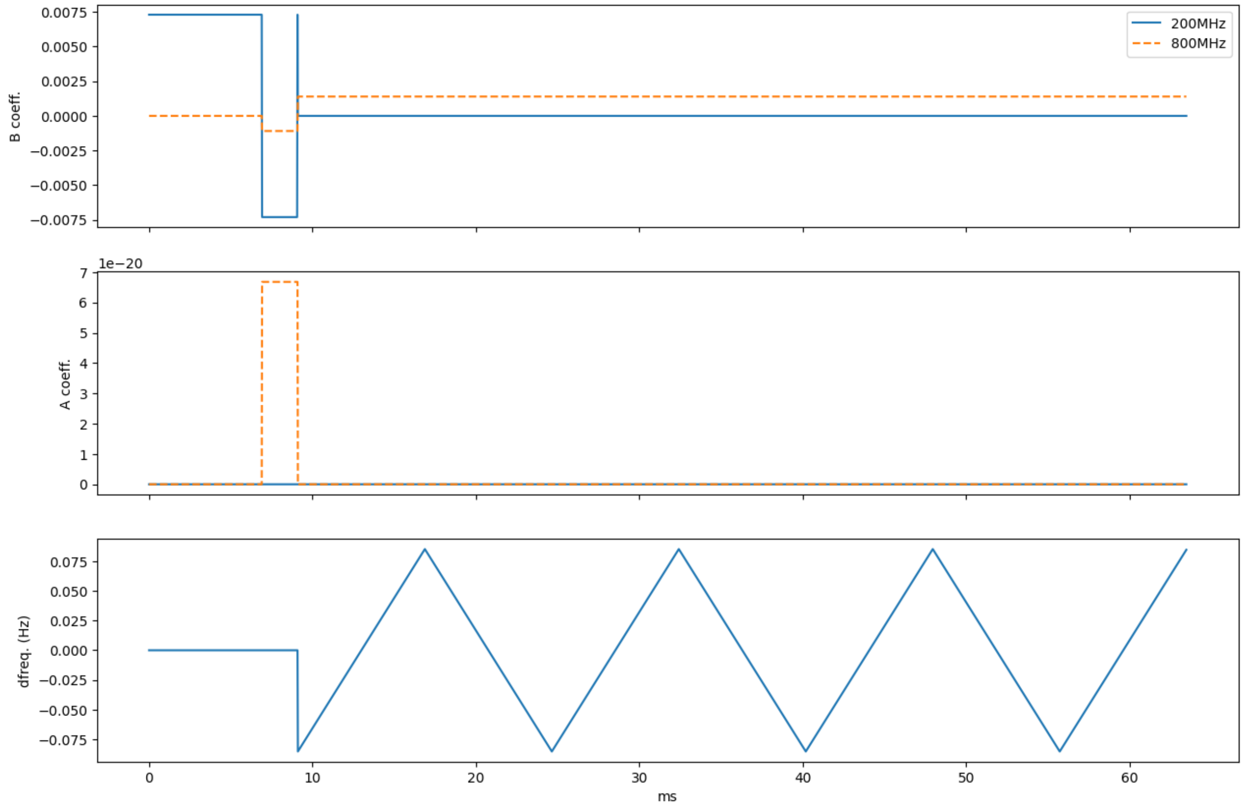

Figure 7: The linear RF schedule for Figure 6.

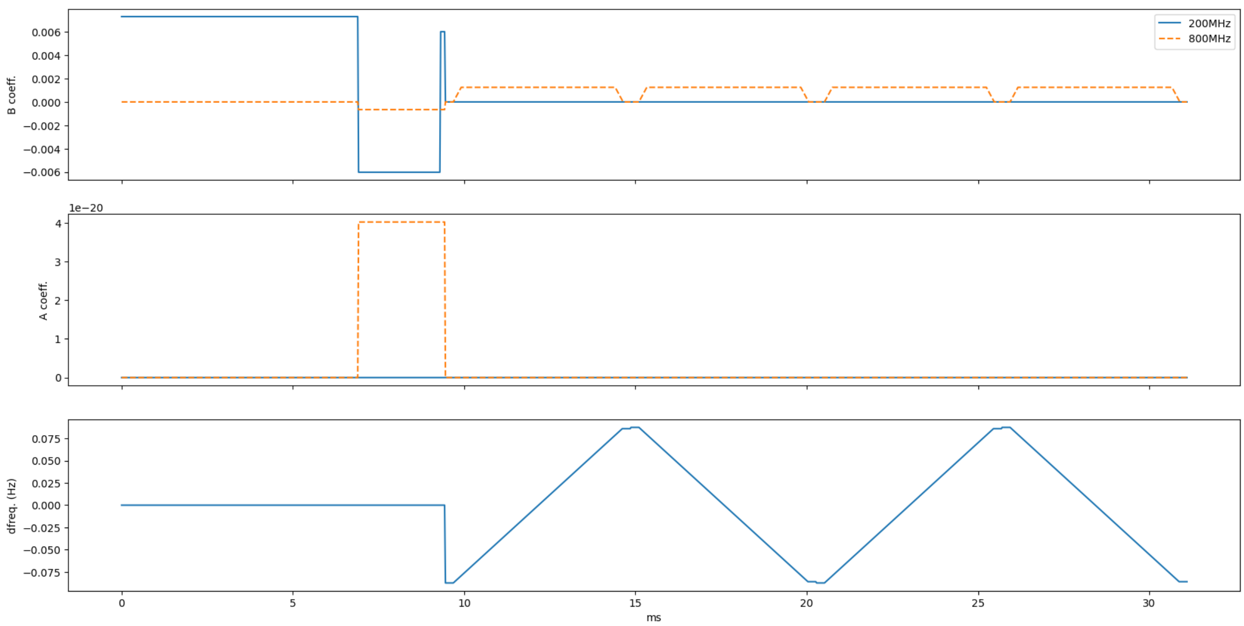

Figure 8: A non-linear RF schedule of the phase displacement blow-up method.

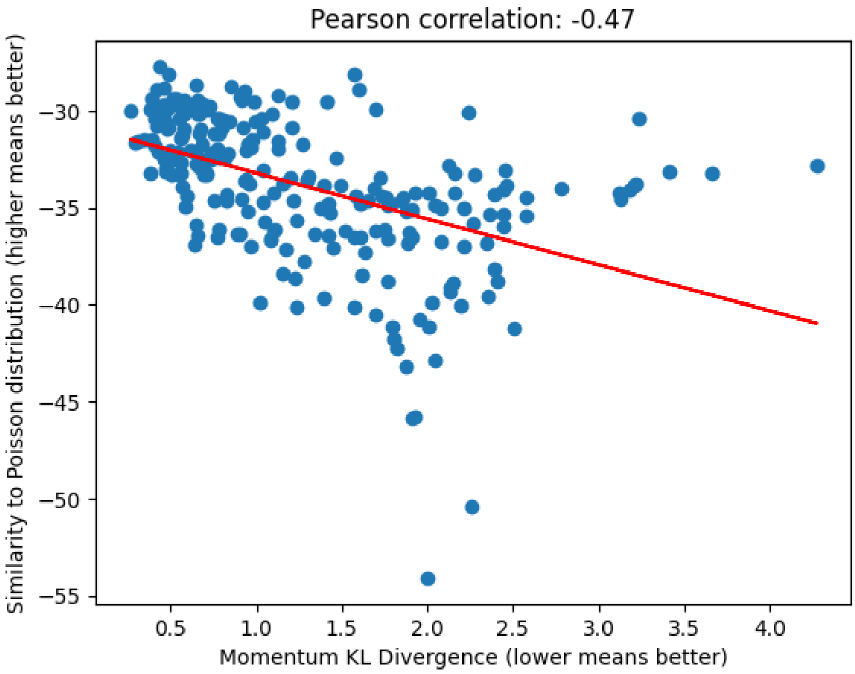

Figure 9: The Pareto front of the phase displacement blow-up method between the momentum and time metrics.

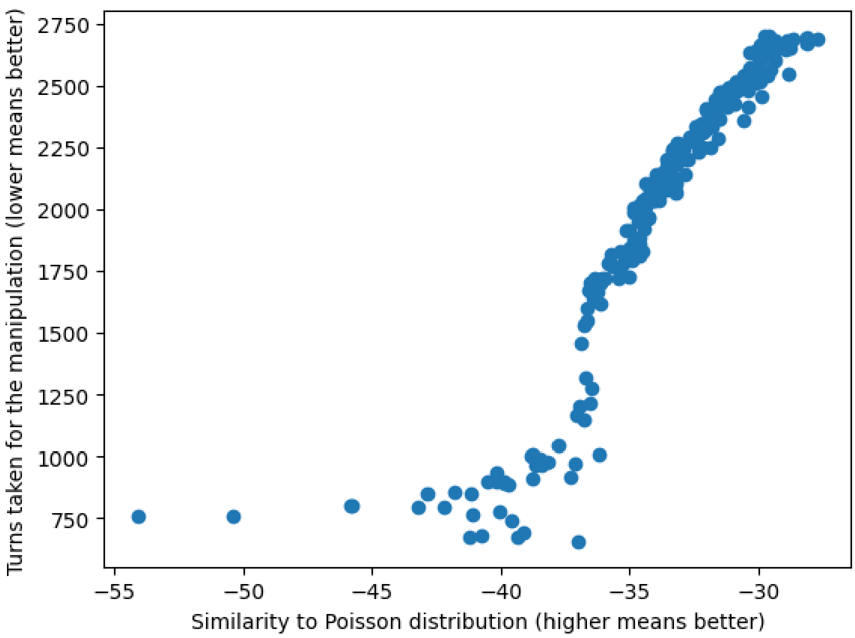

Figure 10: The Pareto front of the phase displacement blow-up method between the time metric and number of turns taken.

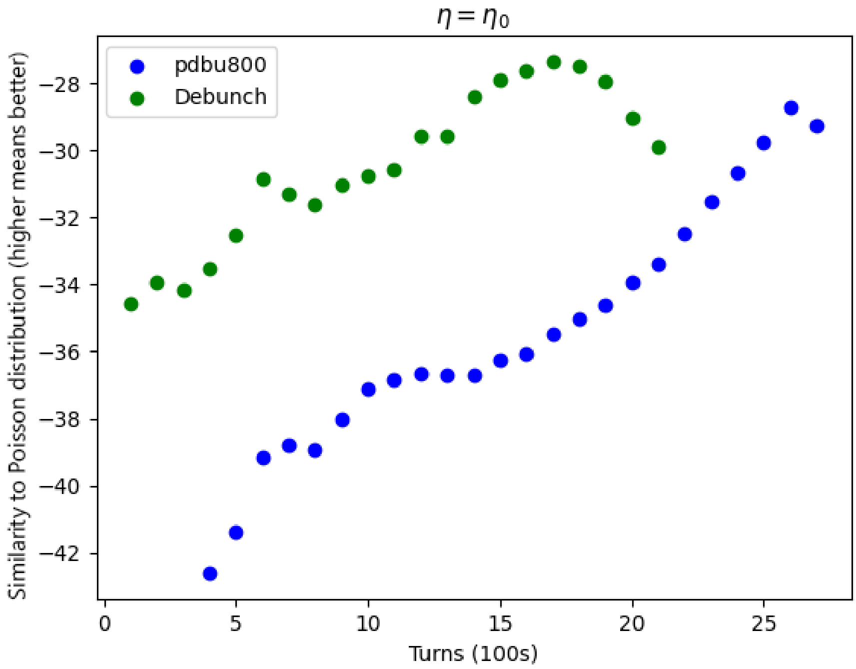

Figure 11: The Pareto front between the time metric and the number of turns taken of the methods with the nominal momentum slippage factor η.

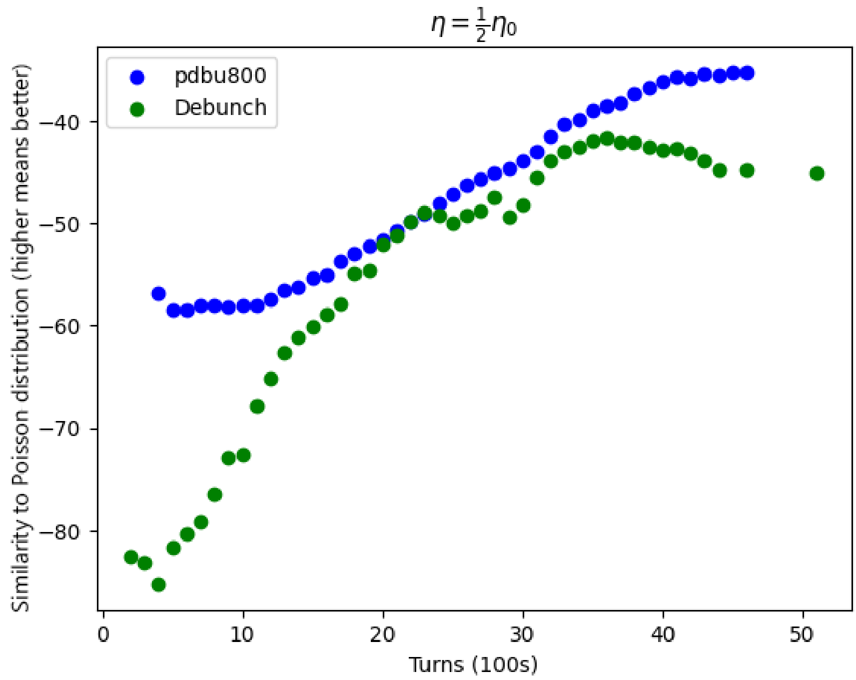

Figure 12: The Pareto front between the time metric and the number of turns taken of the methods with η set to 0.5 of the nominal value.

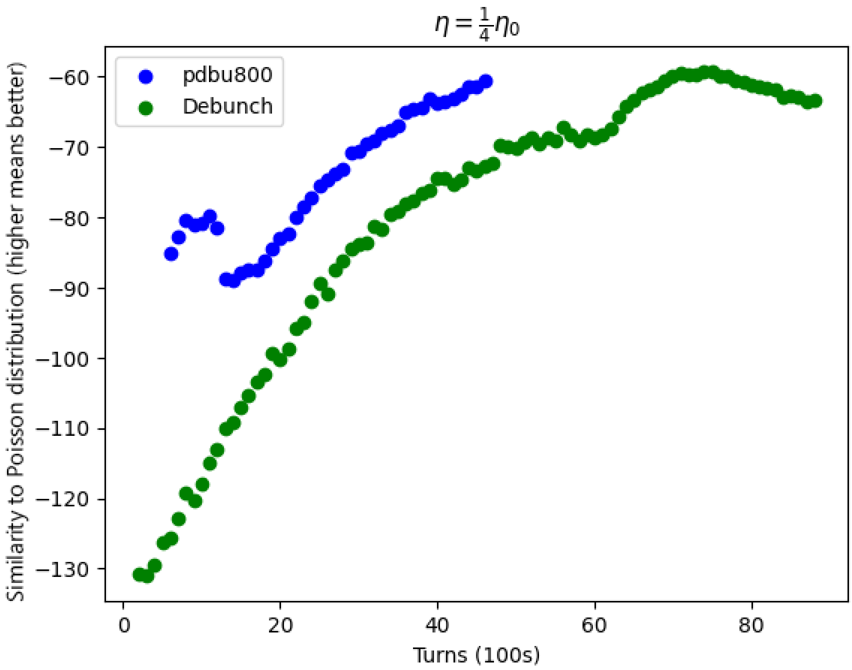

Figure 13: The Pareto front between the time metric and the number of turns taken of the methods with η set to 0.25 of the nominal value.