ReDecay: a novel approach to speed up the simulation at LHCb

D. Müller^{1,a}, M. Clemencic¹, G. Corti¹, M. Gersabeck²

¹ European Organization for Nuclear Research (CERN), Geneva, Switzerland

² School of Physics and Astronomy, University of Manchester, Manchester, UK

Received: 26 October 2018 / Accepted: 20 November 2018 / Published online: 12 December 2018
© CERN for the benefit of the LHCb collaboration 2018

Abstract With the steady increase in the precision of flavour physics measurements collected during LHC Run 2, the LHCb experiment requires simulated data samples of larger and larger sizes to study the detector response in detail. The simulation of the detector response is the main contribution to the time needed to simulate full events. This time scales linearly with the particle multiplicity. Of the dozens of particles present in the simulation only the few participating in the signal decay under study are of interest, while all remaining particles mainly affect the resolutions and efficiencies of the detector. This paper presents a novel development for the LHCb simulation software which re-uses the rest of the event from previously simulated events. This approach achieves an order of magnitude increase in speed and the same quality compared to the nominal simulation.

1 Introduction

A common challenge in many measurements performed in high-energy physics is the necessity to understand the effects of the detector response on the physics parameters of interest. This response is driven by resolution effects that distort the true distribution of a quantity and by inefficiencies that are introduced by either an imperfect reconstruction in the detector or a deliberate event selection, resulting in an unknown number of true events for a given number of reconstructed events. The extraction of an unbiased estimate of the physics parameter of interest requires the correction of these effects. However, due to increasing complexity of the detectors employed and the challenging experimental conditions – especially at a hadron collider – the only feasible solution is the generation of Monte Carlo (MC) events and the simulation of their detector response to study the evolution from the generated to the reconstructed and selected objects. In the era of the Large Hadron Collider, the simulation of the

large event samples occupies a considerable fraction of the overall computing resources available. Hence, measures to decrease the time required to simulate an event are crucial to fully exploit the large datasets recorded by the detectors. In fact, recent measurements hinting at tensions with respect to the predictions of the Standard Model of elementary particle physics have systematic uncertainties that are dominated by the insufficient amount of simulated data [1]. Therefore, fast simulation options are necessary to further improve measurements such as these.

This paper presents a fast simulation approach that is applicable if the signal process is a decay of a heavy particle, e.g. $D^0 \rightarrow K^- \pi^+$ or $B^0 \rightarrow D^{*-} \tau^+ \nu_\tau$. It retains the same quality as the nominal full simulation while reaching increases in speed of an order of magnitude. In Sect. 2, the ReDecay approach is presented while Sect. 3 focuses on the correct treatment of correlations arising in this approach. Lastly, Sect. 4 discusses the experience with applying the approach within the LHCb collaboration. While the approach is applicable in all experiments which study the type of final state as defined above, the LHCb experiment is used as an example throughout this paper for illustrative purposes.

2 The ReDecay approach

In each simulated proton-proton collision, hundreds of particles are typically generated and tracked through a simulation of the detector, most of the time based on the GEANT4 toolkit [2,3]. Generally, the time needed to create an MC event for LHC experiments is dominated by the detector transport, accounting for 95–99% of the total time, which is proportional to the number of particles that need to be tracked. In the case that the intended measurement studies decays of heavy particles to exclusive final states with individual children being reconstructed for each particle, all long lived particles in the event can be split into two distinct groups: the particles that participate in the signal process

^a e-mail: dominik.muller@cern.ch

and all remaining particles where the latter group is henceforth referred to as the rest of the event (ROE). As the final state of the signal process is usually comprised of only a few particles, most particles in the event are part of the ROE and hence the majority of the computing time per event is spent on the simulation of particles that are never explicitly looked at. Ideally, this situation should be inverted and most of the computing resources should be spent on simulating the signal decays themselves. However, simply simulating the signal process without any contributions from the ROE results in a much lower occupancy of the detector. This lower occupancy in turn leads to a significant mis-modelling of the detector response compared to the real detector: resolution effects are underestimated while reconstruction efficiencies are overestimated.

In the ReDecay approach, problems with simulating only the signal process are mitigated by re-using the ROE multiple times instead of generating a new one for every event. The signal particle is kept identical in all sets of events that use the same ROE, to preserve correlations, but is independently decayed each time. To be more precise, in each redecded event, the origin vertex and kinematics of the signal particle are identical, while the decay time and thus decay vertex as well as the final state particle kinematics are different. The following procedure is applied:

1. A full MC event including the signal decay is generated.
2. Before the generated particles are passed through the detector simulation, the signal and its decay products are removed from the event and the origin-vertex position and momentum of the signal particle are stored.
3. The remaining ROE is simulated as usual and the entire output, the information on the true particles and the energy deposits in the detector, are kept.
4. A signal decay is generated and simulated using the stored origin vertex position and momentum.
5. The persisted ROE and the signal decay are merged and written out to disk as a full event.
6. The steps 4 and 5 are repeated N_{ReDecay} times, where N_{ReDecay} is a fixed number that is the same for all events.

An ROE and a signal decay associated to form a complete event according to the procedure above are digitised simultaneously. This ensures that simulated energy deposits that stem from particles in the ROE can interfere with the deposits produced by the signal decay products, as is the case in the standard method to simulate events. Different complete events, for example obtained combining the same ROE with different signal decays, are themselves digitised independently. This further implies that each event obtained could have been produced by chance in a nominal simulation as ReDecay just reorders the already factorised approach: hadrons are decayed independently and the quasi-

stable tracks are propagated through the detector individually. Therefore, the efficiencies and resolutions are, by construction, identical to those found in a nominal simulation. Furthermore, with increasing N_{ReDecay} , the average time per created event becomes more and more dominated by the time required to simulate the detector response for the signal particle and its decay products compared to the simulation of the particles from the ROE. However, an attentive reader will have noticed that different events stemming from the same original event are correlated, e.g. they all have – bar resolution effects – the same signal kinematics. The magnitude of these correlations will also depend on the studied observables and the following section attempts to quantify the strength of the correlations. Additionally, it presents a method to take them into account in an analysis using simulated samples that have been generated following the ReDecay approach.

3 The statistical treatment

Re-using the ROE multiple times and leaving the kinematics of the signal particle unchanged yields correlated events. Considering the signal decay $D^0 \rightarrow K^- \pi^+$ as an example large correlations are expected for the transverse momentum of the D^0 as the decaying particle's kinematics are invariant in all redecded events. On the other hand, quantities computed to evaluate, for example, the performance of the track reconstruction for the K^- meson tracks will be uncorrelated as those are based on the hits produced by the K^- meson traversing the detector that are regenerated in every redecded event. Additionally, many observables are constrained by the fixed D^0 kinematics but have a certain amount of freedom in each redecded event such as the transverse momentum of the K^- . In the following, a simple example is used to develop methods to quantify the degree of correlation between different events and to discuss methods to take the correlations into account.

Suppose x is a random variable, which is sampled in a two stage process: first, a random number x' is sampled from a normal distribution $\mathcal{N}(0, \sigma_1)$ with mean zero and width σ_1 . Subsequently, a value for x is then obtained by sampling one number from a normal distribution $\mathcal{N}(x', \sigma_2)$ with width σ_2 that is centered at x' . The resulting x is then distributed according to the convolution of the two normal distributions, $\mathcal{N}(0, \sqrt{\sigma_1^2 + \sigma_2^2})$. In the nominal case, a new x' is sampled for every x , while a procedure equivalent to the ReDecay approach is achieved by sampling N_{ReDecay} values for x from the same x' . Modifying the procedure for the generation of x in this way does not alter the final distribution and thus both approaches provide random numbers that can be used to obtain a histogram following the expected shape given by the convolution of the two normal distributions. However, the latter approach introduces correlations between different

Fig. 1 Illustration of the correlations between events produced with the ReDecay method, using $N_{\text{ReDecay}} = 100$: (left) autocorrelation function $R(\tau)$ for different values of the ratios of widths of the probability density functions, σ_2/σ_1 and (right) comparison of normalised distributions of a variable x generated with the ReDecay method and a standard generation method (nominal) for $\sigma_1 = 5$ and $\sigma_2 = 1$. The bottom part of the right plot shows the pull distribution of the difference Δ in each

bin of the distribution of the value obtained with the ReDecay method with the value from the nominal method, divided by the naive expectation for the statistical uncertainty, $\sigma = \sqrt{N}$ where N is the number of events in the bins. A unit Gaussian distribution is overlaid with the pull distribution to visualize their disagreement. The pull distribution has a width of 1.75 ± 0.09 indicating that uncertainties are underestimated in this illustration case

x values, which is visible in Fig. 1 as these correlations lead to an underestimate of the uncertainties when the values of x are filled into a histogram using uncertainties equal to $\sqrt{N}$ where N is the number of entries in the bin of the nominal distribution. One possibility to quantify the correlation between different entries in a sequence of random numbers is given by the autocorrelation of the random number x :

$$R(\tau) = \frac{1}{(N - \tau) \sigma^2} \sum_{t=1}^{N-\tau} (x_t - \mu)(x_{t+\tau} - \mu), \quad (1)$$

where N is the total number of entries, μ is the mean (zero in this example) and σ is the standard deviation ($\sqrt{\sigma_1^2 + \sigma_2^2}$ here) of the random numbers x . The autocorrelation itself is typically studied as a function of an integer offset τ , where $\tau = 0$ is the trivial case with $R(0) = 1$. Example autocorrelations as a function of τ are given in Fig. 1 for multiple combinations of values for σ_1 and σ_2 with different ratios of σ_2/σ_1 . As expected, the autocorrelation decreases both as a function of the offset τ as well as with increasing values for the ratio σ_2/σ_1 . While the former is simply caused by a decreasing overlap between the blocks of x from the same x' , the latter is a result of larger values of the ratio σ_2/σ_1 corresponding to a greater amount of freedom in the random process for the generation of x despite using the same x' .

In summary, events are correlated though the amount of correlation depends on the studied observable. This explains

the effect seen in the pulls where the approximation of $\sqrt{N}$ for the uncertainties breaks down in the presence of large correlations, when σ_2/σ_1 is small. In the following, an alternative approach will be discussed to obtain the statistical uncertainty in the presence of such correlations. The effect of the correlations highly depends on how the sample is used and which variables are of interest and hence deriving a general, analytical solution is difficult if not impossible. In use cases where the variables of interest are sufficiently independent in every redecayed event (corresponding to a large σ_2/σ_1 in the example above) even ignoring the correlations altogether can be a valid approach. Nonetheless, a general solution is provided by bootstrapping [4], which is a method for generating pseudo-samples by resampling. These pseudo-samples can then be used to estimate the variance of complex estimators, for which analytic solutions are impractical.

Starting from a sample with N entries, a pseudo-sample can be obtained according to the following procedure: first, a new total number of events N' is randomly drawn from a Poisson distribution of mean N . Then, entries are randomly drawn with replacement from the original sample N' times to fill the pseudo-sample. As sampling with replacement is employed, the obtained pseudo-sample can contain some entries of the original sample multiple times while others are not present in the pseudo-sample at all. Unfortunately, this cannot directly be applied to the ReDecay samples because it assumes statistically independent entries. In the study of time series, data

points ordered in time where each data point can depend on the previous points, different extensions of the bootstrapping algorithm have been developed to preserve the correlations in the pseudo-samples [5]. A common approach is the so-called block bootstrapping where the sample is divided into blocks. In order to capture the correlations arising in the ReDecay approach, a block is naturally given by all events using the same ROE (or the same x' in the example above). To obtain a pseudo-sample, the bootstrapping procedure above is slightly modified: instead of sampling N' individual entries from the original sample, entire blocks are drawn with replacement and all entries constituting a drawn block are filled into the pseudo-sample until the pseudo-sample has reached a size of N' entries. From these pseudo-samples, derived quantities, e.g. the covariance matrix for the bins in the histogram in Fig. 1 (right), can be obtained and utilised to include the correlations in an analysis.

4 Implementation and experience in LHCb

While the idea of ReDecay is applicable in different experiments, the actual implementation strongly depends on the simulation framework and no general solution can be provided. Nonetheless, this section discusses the experiences gained when introducing the ReDecay approach to the LHCb software and which can be transferred to other experiments.

In the most commonly used procedure to simulate events in LHCb, pp collisions are generated with the PYTHIA 8.1 [6, 7] generator using a specific LHCb configuration similar to the one described in Ref. [8]. Decays of particles are

described by EVTGEN [9] in which final-state radiation is generated with PHOTOS [10]. The implementation of the interaction of the generated particles with the detector, and its response, uses the GEANT4 toolkit. The steering of the different steps in the simulation of an event uses interfaces to external generators and to the GEANT4 toolkit. It is handled by GAUSS [11], the LHCb simulation software built on top of the GAUDI [12, 13] framework.

The ReDecay algorithm has been implemented as a package within the GAUSS framework, deployed in the LHCb production system and already used in several large Monte Carlo sample productions. In general, a speed-up by a factor of 10–20 is observed with the exact factor depending on the simulated conditions of the Large Hadron Collider and the complexity of the signal event. This factor is reached with a number of redecays per original event of $N_{\text{ReDecay}} = 100$. In this configuration, and depending on the simulated decay, around 95% of the total time required to create the sample is spent on the simulation of the detector response of the products of the studied decay. Hence, increasing the number of redecays further would have little to no impact on the speed. Due to the correlations between different simulated events using the same ROE as described above, ReDecay samples are typically used to obtain efficiency descriptions as a function of final state variables with high granularity. For example, the use of the ReDecay method is ideal to study variables describing multi-body decays, since these variables are largely independent between all ReDecay events. An example for this is given in Fig. 2, which compares some distributions for $D^0 \rightarrow K^- K^+ \pi^+ \pi^-$ decays between a ReDe-

Fig. 2 Comparison of a ReDecay sample with a nominally produced sample of $D^0 \rightarrow K^- K^+ \pi^+ \pi^-$ decays. Shown are the invariant mass squared of the kaon pair (left) and pion pair (right) which are both inde-

pendent of the kinematics of the decaying particle. Displayed uncertainties are computed assuming independent events. There is no sign of the effects caused by correlated events as seen in Fig. 1 (right)

cay and a nominally simulated sample. Despite ignoring the correlations and using uncertainties computed as $\sqrt{N}$, sensible pull distributions with a width compatible with one are observed.

In the case where the variables of interest are not obviously independent between different ReDecay events, the degree of correlation is difficult to predict and the only solution is the actual creation of a test sample. While the increase in speed when using ReDecay is substantial, the time to produce events is still not negligible and creating a test sample of sufficient size can quickly overwhelm the resources available for an individual analyst. To this end and in collaboration with the respective developers, the ReDecay approach has been added to the RapidSim [14] generator to enable the fast production of simplified simulation samples to judge the degree of correlation before committing the resources for the production of a large ReDecay sample.

5 Summary

The paper presents developments of a fast simulation option applicable to analyses of signal particle decays that occur after the hadronisation phase. An overall increase in speed by a factor of 10–20 is observed. Furthermore, this paper discusses procedures to handle the correlations that can arise between different events if necessary. ReDecay has already enabled the production of large simulated samples for current measurements at LHCb [15].

With the upgrade of the LHCb detector for the third run of the Large Hadron Collider, the detector will collect a substantially larger amount of data. ReDecay is becoming a crucial fast simulation option to extract high precision physics results from this data.

Acknowledgements We acknowledge all our LHCb collaborators who have contributed to the results presented in this paper. Specifically, we thank Patrick Robbe, Michal Kreps and Mark Whitehead for their editorial help preparing this document. We acknowledge support from the Science and Technology Facilities Council (United Kingdom) and the computing resources that are provided by CERN, IN2P3 (France), KIT and DESY (Germany), INFN (Italy), SURF (The Netherlands), PIC (Spain), GridPP (United Kingdom), RRCKI and Yandex LLC (Russia), CSCS (Switzerland), IFIN-HH (Romania), CBPF (Brazil), PL-GRID

(Poland) and OSC (USA). We are indebted to the communities behind the multiple open-source software packages on which we depend.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. Funded by SCOAP³.

References

1. LHCb Collaboration, R. Aaij et al., Measurement of the ratio of the $\mathcal{B}(B^0 \rightarrow D^{*-} \tau^+ \nu_\tau)$ and $\mathcal{B}(B^0 \rightarrow D^{*-} \mu^+ \nu_\mu)$ branching fractions using three-prong τ -lepton decays. *Phys. Rev. Lett.* **120**, 171802 (2018). [arXiv:1708.08856](https://arxiv.org/abs/1708.08856)
2. Geant4 Collaboration, J. Allison, et al., Geant4 developments and applications. *IEEE Trans. Nucl. Sci.* **53**, 270 (2006)
3. Geant4 Collaboration, S. Agostinelli et al., Geant4 a simulation toolkit. *Nucl. Instrum. Methods A* **506**, 250 (2003)
4. B. Efron, Bootstrap methods: another look at the jackknife. *Ann. Stat.* **7**, 1 (1979)
5. J.-P. Kreiss, S.N. Lahiri, 1—Bootstrap methods for time series, in *Time Series Analysis: Methods and Applications. Handbook of Statistics*, vol. 30 (Elsevier, New York, 2012), pp. 3–26. <https://doi.org/10.1016/B978-0-444-53858-1.00001-6>
6. T. Sjöstrand, S. Mrenna, P. Skands, PYTHIA, 6.4 physics and manual. *JHEP* **05**, 026 (2006). [arXiv:hep-ph/0603175](https://arxiv.org/abs/hep-ph/0603175)
7. T. Sjöstrand, S. Mrenna, P. Skands, A brief introduction to PYTHIA 8.1. *Comput. Phys. Commun.* **178**, 852 (2008). [arXiv:0710.3820](https://arxiv.org/abs/0710.3820)
8. I. Belyaev et al., Handling of the generation of primary events in Gauss, the LHCb simulation framework. *J. Phys. Conf. Ser.* **331**, 032047 (2011)
9. D.J. Lange, The EvtGen particle decay simulation package. *Nucl. Instrum. Methods A* **462**, 152 (2001)
10. P. Golonka, Z. Was, PHOTOS Monte Carlo: a precision tool for QED corrections in Z and W decays. *Eur. Phys. J.* **C45**, 97 (2006). [arXiv:hep-ph/0506026](https://arxiv.org/abs/hep-ph/0506026)
11. M. Clemencic, The LHCb simulation application, Gauss: design, evolution and experience. *J. Phys. Conf. Ser.* **331**, 032023 (2011)
12. G. Barrand, GAUDI—a software architecture and framework for building HEP data processing applications. *Comput. Phys. Commun.* **140**, 45 (2001)
13. M. Clemencic et al., Recent developments in the LHCb software framework Gaudi. *J. Phys. Conf. Ser.* **219**, 042006 (2010)
14. G.A. Cowan, D.C. Craik, M.D. Needham, RapidSim: an application for the fast simulation of heavy-quark hadron decays. *Comput. Phys. Commun.* **C214**, 239 (2017). [arXiv:1612.07489](https://arxiv.org/abs/1612.07489)
15. LHCb Collaboration, R. Aaij et al., *Measurement of the Charm-Mixing Parameter y_{CP}* . [arXiv:1810.06874](https://arxiv.org/abs/1810.06874) (submitted to PRL)