Medusa, a multithread 4-body decay fitting and simulation software

Alessandro Maria Ricci[1,2,*], *Antonio Augusto* Alves Junior[3], *Davide* Brundu[4], *Andrea* Contu[4], *Francesca* Dordei[4], and *Piera* Muzzetto[4]

[1]Università di Pisa
[2]INFN Sezione di Pisa
[3]Institute for Astroparticle Physics – Karlsruhe Institute of Technology (IAP/KIT)
[4]INFN Sezione di Cagliari

Abstract. We present a new C++14 compliant application to perform physics data analyses of generic 4-body decays in massively parallel platforms. Medusa is highly based on Hydra, a header-only library which hides most of the complexities of writing parallel code for different architectures. Medusa has been tested through the measurement of the CP-violating phase ϕ_s in b-hadron decays exploiting the data collected by the LHCb experiment. Medusa executes the optimization of the full model, running over 500000 events, until 330 times faster than a non-parallelized program. Medusa is freely available on GitHub under GPL v.3.0 license.

1 Introduction

Among the biggest computational challenges for High Energy Physics (HEP) experiments, there are the increasingly larger datasets that are being collected, which often require correspondingly complex data analyses. In particular, the Probability Density Functions (PDFs) used for modeling the experimental data can have hundreds of free parameters [1]. The optimization of such models involves a significant computational effort and a considerable amount of time, of the order of days, before reaching a result.

Medusa is a C++14 compliant application designed to perform physics data analyses of generic 4-body decays deploying massively parallel platforms (TBB, OpenMP and CUDA). It relies on Hydra [2], a C++14 compliant and header-only library that provides a high-level and user-friendly interface for common algorithms used in HEP, abstracting away the complexities associated with the implementation of code for different massively parallel architectures. Thanks to Hydra, the code is independent from the architecture. Indeed, changing only a flag at the compilation time, it is possible to compile the same code for different architectures (TBB, OpenMP, CUDA).

Medusa has been tested through the measurement of the CP-violating phase ϕ_s in b-hadron decays exploiting the data collected by the LHCb experiment [3] between 2015 and 2016, the results of which have been published in Ref. [4, 5] and which we refer to test the implementation success. By deploying such technologies as CUDA, TBB and OpenMP,

*e-mail: alessandro.ricci@df.unipi.it

Medusa accelerates the optimization of the full model, running over 500000 events, by factors 73 (multicore CPU) and 330 (GPU) in comparison with a non-parallelized program.

Besides measuring the CP-violating phase in $B_s^0\left(\bar{B}_s^0\right)$ decays, Medusa is willing to be a set of ready-made models to perform many diverse data analyses, as e.g. the Fadeeva functions [6], in order to accelerate the develop of new models.

2 Use case: $B_s^0\left(\bar{B}_s^0\right)$ decay at LHCb

Medusa has been tested through the measurement of the CP-violating phase ϕ_s in $B_s^0\left(\bar{B}_s^0\right) \rightarrow J/\psi\left(\mu^+\mu^-\right)\phi\left(K^+K^-\right)$ decay, one of the golden channels for this type of research at the LHCb experiment. A such complex model is an excellent example to test Medusa's performance. Moreover, it also contains all the principal ingredients of a typical flavour analysis at colliders.

Between 2015 and 2016, the LHCb Collaboration collected a sample of about 209000 events. To extract the ϕ_s phase, it is necessary to perform a maximum-likelihood fit with 32 free parameters, by using a model which includes both the signal and the background. Moreover, this model must include the modeling of the distribution of the B_s^0-decay time t, of the decay angles $\theta_\mu, \theta_K, \phi_h$ (see Fig. 1), of the so-called flavour-tagging needed to distinguish between B_s^0 and $\bar{B}_s^0$ mesons at production and the modeling of other experimental artifacts.

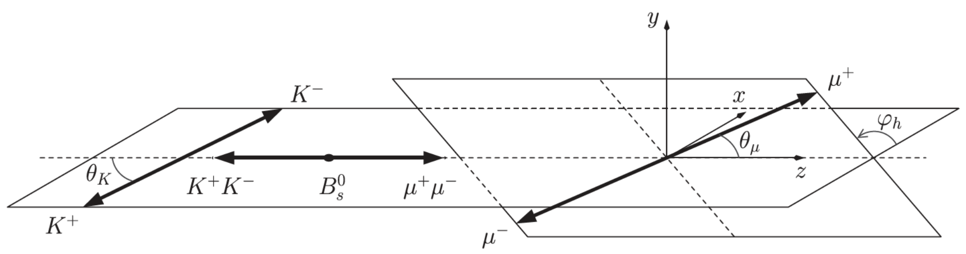

Figure 1. Definition of the angles in the helicity basis, which is used to describe the decay geometry [7].

The distribution of the decay time and angles in the helicity basis for the B_s^0 and $\bar{B}_s^0$ mesons produced at time $t = 0$ is described by the sum of 10 terms, corresponding to the four squared polarization amplitudes N_k (see Fig. 2) and their interference terms. Each of these is given by the product of the decay time- and angle-dependent functions. So, the full time- and angle-dependent decay rate is [7]:

$$\frac{d^4\Gamma}{dt\, d\theta_\mu\, d\theta_K\, d\phi} \propto \sum_{k=1}^{10} N_k h_{k,q}(t) f_k\left(\theta_\mu, \theta_K, \phi\right), \tag{1}$$

with

$$h_{k,1}(t) = \frac{3}{4\pi} e^{-\Gamma t}\left(a_k \cosh\left(\frac{\Delta\Gamma t}{2}\right) + b_k \sinh\left(\frac{\Delta\Gamma t}{2}\right) + c_k \cos\left(\Delta m t\right) + d_k \sin\left(\Delta m t\right)\right) \tag{2}$$

for the B_s^0 meson and

$$h_{k,-1}(t) = \frac{3}{4\pi} e^{-\Gamma t}\left(a_k \cosh\left(\frac{\Delta\Gamma t}{2}\right) + b_k \sinh\left(\frac{\Delta\Gamma t}{2}\right) - c_k \cos\left(\Delta m t\right) - d_k \sin\left(\Delta m t\right)\right) \tag{3}$$

for the $\bar{B}_s^0$ meson. Here, Γ is the B_s^0 decay width, $\Delta\Gamma$ and Δm are respectively the width and mass differences between the light and heavy mass eigenstates of the B_s^0 meson. The squared polarization amplitudes N_k must satisfy the condition $A_0^2 + A_\perp^2 + A_\parallel^2 = 1$. So, only A_0^2 and $A_\perp^2$ are used as parameters for the fit. Moreover, each of the three amplitudes comes with a different phase, labeled as $\delta_0, \delta_\perp$ and $\delta_\parallel$. Since only the phase differences are observable, the three phases are reduced to two phase differences $\delta_\perp - \delta_0$ and $\delta_\parallel - \delta_0$ which are used for the fit. The formulas for the angular functions $f_k\left(\theta_\mu, \theta_K, \phi\right)$, the squared polarization amplitudes N_k and the coefficients a_k, b_k, c_k, d_k, containing all phases, can be found in Table 1.

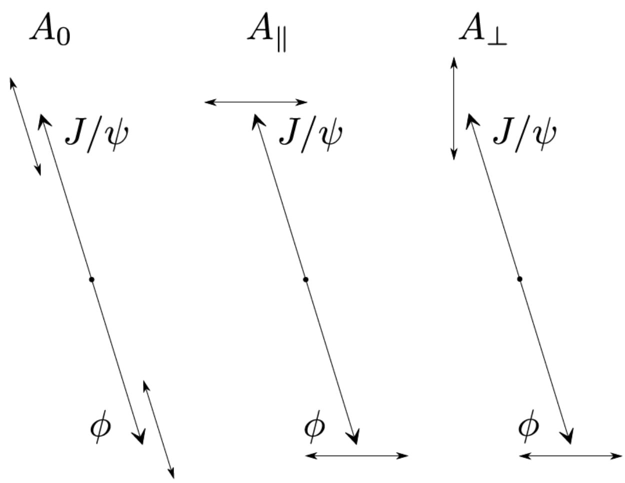

Figure 2. Polarization amplitudes of the $J/\psi\,\phi$ system. The short arrows indicate the spin orientation of the two vector mesons [7].

The signal-only model uses 18 free parameters for the fit. However, we need to consider also the flavour tagging, i.e. the identification of the initial b-hadron in the decay, which adds other 4 free parameters to the model. Then, we have to consider other experimental artifacts like the decay-time resolution, the decay-time and angular acceptances and the S-wave. Decay-time resolution, decay-time and angular acceptances increase the model complexity without raising the number of free parameters. The S-wave is composed of K^+K^--couples, which originate from the direct decay $B_s^0\left(\bar{B}_s^0\right) \to J/\psi K^+K^-$. This contribution has a different angular distribution and then must be separated from the signal-only model. The S-wave is split in six bins of the invariant mass of the two kaons, m_{KK}, namely: [990 – 1008], [1008 – 1016], [1016 – 1020], [1020 – 1024], [1024 – 1032], [1032 – 1050] MeV/c^2. The two S-wave associated parameters are left free to vary in each bin. This brings to a simultaneous fit on six m_{KK}-bins with 20 free parameters in common between the bins and two specific for each bin. Hence, totally we have 32 free parameters.

Finally, the background is statistically removed by associating a weight to each event [8], which indicates the probability that the event is a signal or background.

Table 1. Angular and time-dependent functions with the polarization dependent CP violation [4]. Abbreviations used include $c_K = \cos\theta_K$, $s_K = \sin\theta_K$, $c_\mu = \cos\theta_\mu$, $s_\mu = \sin\theta_\mu$, $c_\phi = \cos\phi$ and $s_\phi = \sin\phi$.

f_k	N_k	a_k	b_k	c_k	d_k
$c_K^2 s_\mu^2$	$\|A_0\|^2$	$\frac{1}{2}(1+\|\lambda_0\|^2)$	$-\|\lambda_0\|\cos(\phi_0)$	$\frac{1}{2}(1-\|\lambda_0\|^2)$	$\|\lambda_0\|\sin(\phi_0)$
$\frac{1}{2}s_K^2(1-c_\phi^2 s_\mu^2)$	$\|A_\|\|^2$	$\frac{1}{2}(1+\|\lambda_\|\|^2)$	$-\|\lambda_\|\|\cos(\phi_\|)$	$\frac{1}{2}(1-\|\lambda_\|\|^2)$	$\|\lambda_\|\|\sin(\phi_\|)$
$\frac{1}{2}s_K^2(1-s_\phi^2 s_\mu^2)$	$\|A_\perp\|^2$	$\frac{1}{2}(1+\|\lambda_\perp\|^2)$	$\|\lambda_\perp\|\cos(\phi_\perp)$	$\frac{1}{2}(1-\|\lambda_\perp\|^2)$	$-\|\lambda_\perp\|\sin(\phi_\perp)$
$s_K^2 s_\mu^2 s_\phi c_\phi$	$\|A_\perp A_\|\|$	$\frac{1}{2}\Big[\sin(\delta_\perp-\delta_\|)-\|\lambda_\perp\lambda_\|\|$ $\sin(\delta_\perp-\delta_\|-\phi_\perp+\phi_\|)\Big]$	$\frac{1}{2}\Big[\|\lambda_\perp\|\sin(\delta_\perp-\delta_\|-\phi_\perp)$ $+\|\lambda_\|\|\sin(\delta_\|-\delta_\perp-\phi_\|)\Big]$	$\frac{1}{2}\Big[\sin(\delta_\perp-\delta_\|)+\|\lambda_\perp\lambda_\|\|$ $\sin(\delta_\perp-\delta_\|-\phi_\perp+\phi_\|)\Big]$	$-\frac{1}{2}\Big[\|\lambda_\perp\|\cos(\delta_\perp-\delta_\|-\phi_\perp)$ $+\|\lambda_\|\|\cos(\delta_\|-\delta_\perp-\phi_\|)\Big]$
$\sqrt{2}s_K c_K s_\mu c_\mu c_\phi$	$\|A_0 A_\|\|$	$\frac{1}{2}\Big[\cos(\delta_0-\delta_\|)+\|\lambda_0\lambda_\|\|$ $\cos(\delta_0-\delta_\|-\phi_0+\phi_\|)\Big]$	$-\frac{1}{2}\Big[\|\lambda_0\|\cos(\delta_0-\delta_\|-\phi_0)$ $+\|\lambda_\|\|\cos(\delta_\|-\delta_0-\phi_\|)\Big]$	$\frac{1}{2}\Big[\cos(\delta_0-\delta_\|)-\|\lambda_0\lambda_\|\|$ $\cos(\delta_0-\delta_\|-\phi_0+\phi_\|)\Big]$	$-\frac{1}{2}\Big[\|\lambda_0\|\sin(\delta_0-\delta_\|-\phi_0)$ $+\|\lambda_\|\|\sin(\delta_\|-\delta_0-\phi_\|)\Big]$
$-\sqrt{2}s_K c_K s_\mu c_\mu s_\phi$	$\|A_0 A_\perp\|$	$-\frac{1}{2}\Big[\sin(\delta_0-\delta_\perp)-\|\lambda_0\lambda_\perp\|$ $\sin(\delta_0-\delta_\perp-\phi_0+\phi_\perp)\Big]$	$\frac{1}{2}\Big[\|\lambda_0\|\sin(\delta_0-\delta_\perp-\phi_0)$ $+\|\lambda_\perp\|\sin(\delta_\perp-\delta_0-\phi_\perp)\Big]$	$-\frac{1}{2}\Big[\sin(\delta_0-\delta_\perp)+\|\lambda_0\lambda_\perp\|$ $\sin(\delta_0-\delta_\perp-\phi_0+\phi_\perp)\Big]$	$-\frac{1}{2}\Big[\|\lambda_0\|\cos(\delta_0-\delta_\perp-\phi_0)$ $+\|\lambda_\perp\|\cos(\delta_\perp-\delta_0-\phi_\perp)\Big]$
$\frac{1}{3}s_\mu^2$	$\|A_S\|^2$	$\frac{1}{2}(1+\|\lambda_S\|^2)$	$\|\lambda_S\|\cos(\phi_S)$	$\frac{1}{2}(1-\|\lambda_S\|^2)$	$-\|\lambda_S\|\sin(\phi_S)$
$\frac{2}{\sqrt{6}}s_K s_\mu c_\mu c_\phi$	$\|A_S A_\|\|$	$\frac{1}{2}\Big[\cos(\delta_S-\delta_\|)-\|\lambda_S\lambda_\|\|$ $\cos(\delta_S-\delta_\|-\phi_S+\phi_\|)\Big]$	$\frac{1}{2}\Big[\|\lambda_S\|\cos(\delta_S-\delta_\|-\phi_S)$ $-\|\lambda_\|\|\cos(\delta_\|-\delta_S-\phi_\|)\Big]$	$\frac{1}{2}\Big[\cos(\delta_S-\delta_\|)+\|\lambda_S\lambda_\|\|$ $\cos(\delta_S-\delta_\|-\phi_S+\phi_\|)\Big]$	$\frac{1}{2}\Big[\|\lambda_S\|\sin(\delta_S-\delta_\|-\phi_S)$ $-\|\lambda_\|\|\sin(\delta_\|-\delta_S-\phi_\|)\Big]$
$-\frac{2}{\sqrt{6}}s_K s_\mu c_\mu s_\phi$	$\|A_S A_\perp\|$	$-\frac{1}{2}\Big[\sin(\delta_S-\delta_\perp)+\|\lambda_S\lambda_\perp\|$ $\sin(\delta_S-\delta_\perp-\phi_S+\phi_\perp)\Big]$	$-\frac{1}{2}\Big[\|\lambda_S\|\sin(\delta_S-\delta_\perp-\phi_S)$ $-\|\lambda_\perp\|\sin(\delta_\perp-\delta_S-\phi_\perp)\Big]$	$-\frac{1}{2}\Big[\sin(\delta_S-\delta_\perp)-\|\lambda_S\lambda_\perp\|$ $\sin(\delta_S-\delta_\perp-\phi_S+\phi_\perp)\Big]$	$-\frac{1}{2}\Big[-\|\lambda_S\|\cos(\delta_S-\delta_\perp-\phi_S)$ $+\|\lambda_\perp\|\cos(\delta_\perp-\delta_S-\phi_\perp)\Big]$
$\frac{2}{\sqrt{3}}c_K s_\mu^2$	$\|A_S A_0\|$	$\frac{1}{2}\Big[\cos(\delta_S-\delta_0)-\|\lambda_S\lambda_0\|$ $\cos(\delta_S-\delta_0-\phi_S+\phi_0)\Big]$	$\frac{1}{2}\Big[\|\lambda_S\|\cos(\delta_S-\delta_0-\phi_S)$ $-\|\lambda_0\|\cos(\delta_0-\delta_S-\phi_0)\Big]$	$\frac{1}{2}\Big[\cos(\delta_S-\delta_0)+\|\lambda_S\lambda_0\|$ $\cos(\delta_S-\delta_0-\phi_S+\phi_0)\Big]$	$\frac{1}{2}\Big[\|\lambda_S\|\sin(\delta_S-\delta_0-\phi_S)$ $-\|\lambda_0\|\sin(\delta_0-\delta_S-\phi_0)\Big]$

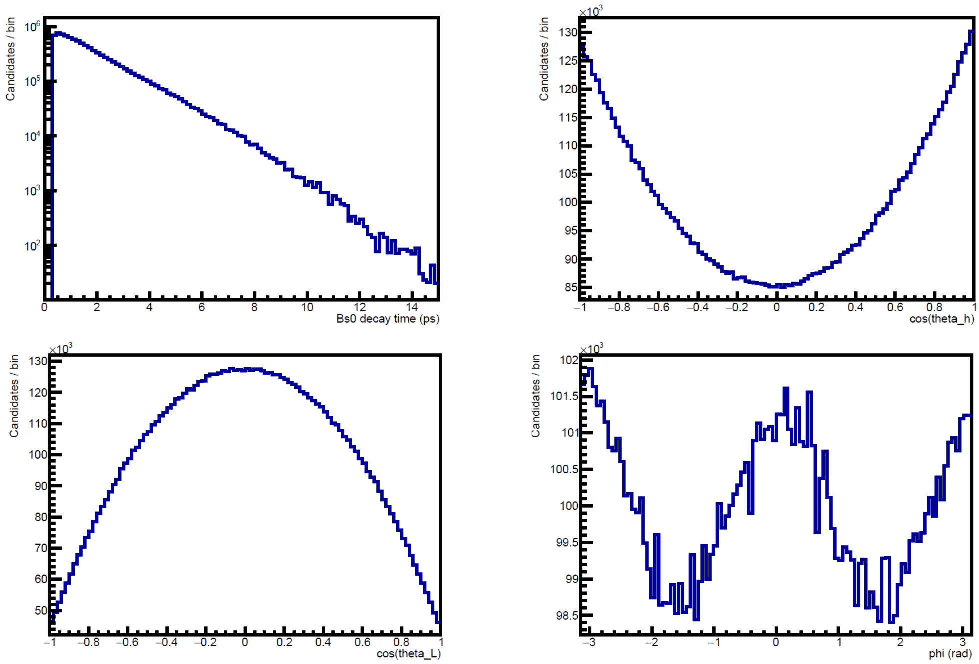

Figure 3. Distributions of the B_s^0 decay time and angles in the helicity basis with 10 millions of simulated events in the first m_{KK}-bin and using a biased trigger (h means hadron and L means lepton).

The Probability Density Function (PDF) used to fit the dataset is

$$PDF_{y,c}^{j}\left(t,\Omega\,|\,q^{OS},q^{SS},\eta^{OS},\eta^{SS},\delta_t\right)=\frac{1}{N_{q^{OS},q^{SS},y,c}^{\eta^{OS},\eta^{SS},\delta_t,j}}\sum_{k=1}^{10}\tilde{N}_k^j f_k(\Omega)\,\epsilon_{y,c}(t)$$
$$\Big\{\Big[\left(1+q^{OS}\left(1-2\omega^{OS}(\eta_{OS})\right)\right)\left(1+q^{SS}\left(1-2\omega^{SS}(\eta_{SS})\right)\right)\cdot h_{k,1}(t)+$$
$$\left(1-q^{OS}\left(1-2\bar{\omega}^{OS}(\eta_{OS})\right)\right)\left(1-q^{SS}\left(1-2\bar{\omega}^{SS}(\eta_{SS})\right)\right)\cdot h_{k,-1}(t)\Big]\otimes G\left(t\,|\,\sigma_{eff}(\delta_t)\right)\Big\}, \tag{4}$$

where $\Omega=\left(\theta_\mu,\theta_K,\phi\right)$, $\left(q^{OS},q^{SS},\eta^{OS},\eta^{SS}\right)$ are the tagging variables, δ_t is the estimated per event decay-time uncertainty and $G\left(t\,|\,\sigma_{eff}(\delta_t)\right)$ is the effective Gaussian used to model it, $\epsilon_{y,c}(t)$ is the cubic spline used for the decay-time acceptance, with the indexes y and c which run on the year and the two trigger categories, biased and unbiased, respectively. Instead, the index j runs on the six m_{kk}-bins. Finally, $N^{\eta^{OS},\eta^{SS},\delta_t,j}_{q^{OS},q^{SS},y,c}$ is a normalization factor given by

$$N^{\eta^{OS},\eta^{SS},\delta_t,j}_{q^{OS},q^{SS},y,c} = \int_{0.3\,\mathrm{ps}}^{15\,\mathrm{ps}} \sum_{k=1}^{10} \tilde{N}^j_k \epsilon_{y,c}(t)\, \omega^k_{y,c}$$

$$\Big\{\Big[\left(1 + q^{OS}\left(1 - 2\omega^{OS}(\eta_{OS})\right)\right)\left(1 + q^{SS}\left(1 - 2\omega^{SS}(\eta_{SS})\right)\right)\cdot h_{k,1}(t) +$$

$$\left(1 - q^{OS}\left(1 - 2\bar{\omega}^{OS}(\eta_{OS})\right)\right)\left(1 - q^{SS}\left(1 - 2\bar{\omega}^{SS}(\eta_{SS})\right)\right)\cdot h_{k,-1}(t)\Big] \otimes G\left(t\,|\,\sigma_{eff}(\delta_t)\right)\Big\}\, dt\,, \tag{5}$$

which can be computed analytically [6], but the calculation is expensive. The interested reader can find more detains in Ref. [1, 4, 5, 7].

3 Validation and performance

We validated the model comparing the results of Medusa with those published by the LHCb Collaboration in Ref. [9]. Table 2 shows the difference between the numerical values found by Medusa and the LHCb Collaboration for main parameters of interest as extracted from the fit.

As a further check, we generated a 10 million events Monte Carlo dataset of the decay times and angles in the helicity basis using a uniform distribution. Then, we shape it according to Eq. (4). Fig. 3 shows the distributions of the decay times and angles which we obtained. They are compatible with those reported in Ref. [7, 9].

Table 2. Difference between the results of Medusa (M) and the LHCb Collaboration [9] (left) and, for comparison, the statistical uncertainties of the LHCb results (right).

Differences	Statistical uncertainties
$\phi_s^{LHCb} - \phi_s^M = 0.0009$ rad	$\delta\phi_s^{LHCb} = 0.043$ rad
$\lambda_0^{LHCb} - \lambda_0^M = -0.013$	$\delta\lambda_0^{LHCb} = 0.045$
$\Gamma^{LHCb} - \Gamma^M = 0.00055$ ps^{-1}	$\delta\Gamma^{LHCb} = 0.0024$ ps^{-1}
$\Delta\Gamma^{LHCb} - \Delta\Gamma^M = -0.0007$ ps^{-1}	$\delta\Delta\Gamma^{LHCb} = 0.0077$ ps^{-1}
$\Delta m^{LHCb} - \Delta m^M = 0.029$ ps^{-1}	$\delta\Delta m^{LHCb} = 0.060$ ps^{-1}
$A_\perp^{2,LHCb} - A_\perp^{2,M} = -0.0035$	$\delta A_\perp^{2,LHCb} = 0.025$
$A_0^{2,LHCb} - A_0^{2,M} = 0.0011$	$\delta A_0^{2,LHCb} = 0.018$
$(\delta_\perp - \delta_0)^{LHCb} - (\delta_\perp - \delta_0)^M = 0.031$ rad	$\delta(\delta_\perp - \delta_0)^{LHCb} = 0.15$ rad
$(\delta_\parallel - \delta_0)^{LHCb} - (\delta_\parallel - \delta_0)^M = 0.029$ rad	$\delta(\delta_\parallel - \delta_0)^{LHCb} = 0.083$ rad

We tested the performance of Medusa measuring the time spent to evaluate the objective function, the log-likelihood of Eq. (4), with an increasing number of events. Table 3 shows the results for different architectures. Comparing the obtained values, we see that, running over 500000 events, Medusa can accelerate the model resolution about 73 times (multicore CPU) and 330 times (GPU) faster than a non-parallelized software. Fig. 4 shows how the evaluation time scales with the number of events.

Finally, also the compilation time also has been strongly optimized. namely, the GCC compiler spends about 1 minute to create the executable for TBB and OpenMP and NVCC about 4 minutes for CUDA.

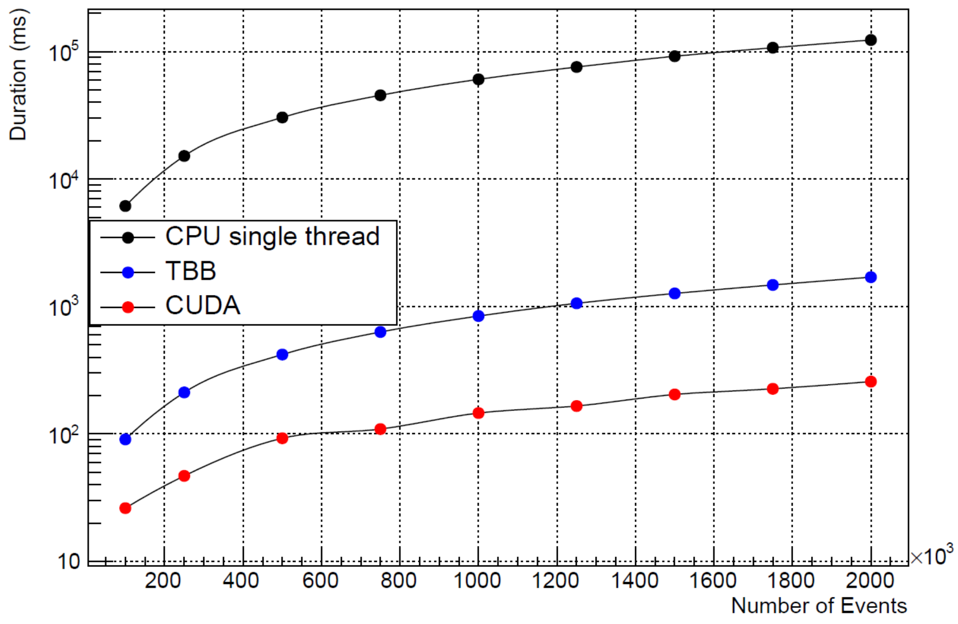

Figure 4. Evaluation time of the objective function per number of events.

Table 3. Time spent by Medusa to evaluate the objective function with 500000 events.

System	Time/call (ms)
AMD EPYC 7452 @ 1.5 GHz (1 Thread)	30424
AMD EPYC 7452 @ 1.5 GHz (128 Threads)	419
NVIDIA A100	92

4 Conclusions

HEP experiments collect ever-larger datasets and their analyses get more and more complex. Moreover, the PDFs used for data modelling become increasingly more complicated, with hundred of free parameters. Not rarely, a computation spends hours if not days to reach a result, which very often needs to be re-tuned.

Medusa is a multithread 4-body decay fitting and simulation software created to speed up the physics data analysis. Medusa is a set of ready-made models for the analysis and ready-made multithread functions to accelerate the develop of new models. Thanks to the library Hydra, all software is independent from the architecture. Indeed, changing only a flag at the compilation time, it is possible to compile the same code for different architectures (TBB, OpenMP, CUDA).

Medusa has been validated through the measurement of the CP-violating phase in $B_s^0\left(\bar{B}_s^0\right) \to J/\psi\,(\mu^+\mu^-)\,\phi\,(K^+K^-)$ decay, one of the golden channel for this type of research at the LHCb experiment. We compared the results of Medusa with those obtained by the LHCb experiment using data collected between 2015 and 2016 [9]. We demonstrated that, using massively parallel platforms as TBB, OpenMP and CUDA, Medusa can accelerate the optimization of the full model, running over 500000 events, by about 73 times on multicore CPU and about 330 times on GPU.

Medusa is freely available on GitHub under GPL v.3.0 license [10].

References

[1] R. Aaij et al. (LHCb Collaboration), arXiv:2308.01468 (2023).

[2] A. A. Alves Junior, *MultithreadCorner/Hydra*, doi:10.5281/zenodo.1206261 (2018).

[3] R. Aaij et al. (LHCb Collaboration), International Journal of Modern Physics A **30**, 1530022 (2015).

[4] R. Aaij et al. (LHCb Collaboration), European Physical Journal C **79**, 106 (2019).

[5] R. Aaij et al. (LHCb Collaboration), European Physical Journal C **80**, 601 (2020).

[6] T. M. Karbach et al., arXiv:1407.0748v1 (2014).

[7] S. Stemmle, *CP violation and lifetime measurements in the decay* $B_s^0 \to J/\psi\, \phi$ *with the LHCb experiment* (Ph.D. Thesis, Heidelberg, 2019) 199 pages.

[8] H. Dembinski et al., Nuclear Instruments and Methods A *1040*, 167270 (2022).

[9] S. Benson et al., *Measurement of the CP-violating phase* ϕ_s *using* $B_s^0 \to J/\psi K^+ K^-$ *decays in Run* 2 (LHCb-ANA-2017-028, CERN, 2020) 389 pages.

[10] A. M. Ricci et al., *Medusa, a multithread 4-body decay fitting and simulation software*, https://github.com/dbrundu/Medusa/tree/develop-aricci.