

PAPER • OPEN ACCESS

Model-independent partial wave analysis using a massively-parallel fitting framework

To cite this article: L Sun *et al* 2017 *J. Phys.: Conf. Ser.* **898** 072025

View the [article online](#) for updates and enhancements.

Related content

- [Amplitude analysis of four-body decays using a massively-parallel fitting framework](#)

C Hasse, J Albrecht, A A Alves Jr. et al.

- [Partial Wave Analysis of \$J/\psi\$ Decays into](#)

Zhang Zhen-Xia and Wu Ning

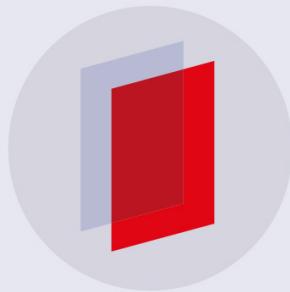
- [Performance studies of GooFit on GPUs vs RooFit on CPUs while estimating the statistical significance of a new physical signal](#)

Adriano Di Florio

Recent citations

- [GooStats: A GPU-based framework for multi-variate analysis in particle physics](#)

X.F. Ding



IOP | ebooks™

Bringing you innovative digital publishing with leading voices to create your essential collection of books in STEM research.

Start exploring the collection - download the first chapter of every title for free.

Model-independent partial wave analysis using a massively-parallel fitting framework

L Sun¹, R Aoude², A C dos Reis² and M Sokoloff³

¹ School of Physics and Technology, Wuhan University, Wuhan, Hubei Province, 430072 China

² Centro Brasileiro de Pesquisas Físicas (CBPF), Rio de Janeiro, 22290-180 Brazil

³ University of Cincinnati, Physics Department, ML0011, Cincinnati OH 45221-0011, USA

E-mail: sunl@whu.edu.cn

Abstract. The functionality of GooFit, a GPU-friendly framework for doing maximum-likelihood fits, has been extended to extract model-independent S -wave amplitudes in three-body decays such as $D^+ \rightarrow h^+ h^+ h^-$. A full amplitude analysis is done where the magnitudes and phases of the S -wave amplitudes are anchored at a finite number of $m^2(h^+ h^-)$ control points, and a cubic spline is used to interpolate between these points. The amplitudes for P -wave and D -wave intermediate states are modeled as spin-dependent Breit-Wigner resonances. GooFit uses the Thrust library, with a CUDA backend for NVIDIA GPUs and an OpenMP backend for threads with conventional CPUs. Performance on a variety of platforms is compared. Executing on systems with GPUs is typically a few hundred times faster than executing the same algorithm on a single CPU.

1. Introduction

The physics of scalar mesons below 2 GeV is a long-standing puzzle in light meson spectroscopy [1], this is partly due to the fact that these mesons often have large widths and overlap with each other, which make them hard to model. Contributions from Light scalars can be extracted by using Dalitz plot analyses [2] in three-body hadronic decays of charm mesons. First developed by the E791 Collaboration [3], the Model Independent Partial Wave Analysis (MIPWA) technique extracts S -wave amplitudes in the Dalitz plot analysis with no assumption about the nature of the S -wave. To pin down its large numbers of free parameters, the MIPWA technique requires the large samples of three-body decay events that have become increasingly available at the B -factories and the LHC. Notably, the LHCb experiment has recorded charm decays with sample sizes exceeding any previous experiment by more than an order of magnitude, offering a unique opportunity to study S -wave structures with unprecedented levels of precision. Analysing very large statistics samples requires disproportionately more computing power. Running all the calculations in a single Central Processing unit (CPU) thread would take a prohibitively long time.

Originally, a Graphics Processing Unit (GPU) was a specialised electronic circuit designed to rapidly manipulate and alter memory to accelerate the creation of images for output to a video display. The highly parallel structure makes GPUs more effective than CPUs for algorithms where large blocks of data are processed in parallel. Today, the functionality of GPUs has been extended to enable highly parallel computing for scientific and other more general applications. An open-source framework called GooFit [4] has been developed to exploit the processing power



Content from this work may be used under the terms of the [Creative Commons Attribution 3.0 licence](#). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

of these GPUs for parallel function evaluation, particularly in the context of maximum likelihood fits. This paper reports an extension of the GooFit framework to support MIPWA of a three-body decay with vastly reduced processing time.

2. MIPWA method

Essentially all studies of three-body hadronic $D_{(s)}$ decays employ the same technique: the unbinned maximum likelihood fit of the Dalitz plot, in which the quantum mechanical matrix element governing the decay is represented by a coherent sum of amplitudes [1]. These amplitudes correspond to the possible intermediate states in the decay chain $D_{(s)} \rightarrow Rh_3, R \rightarrow h_1 h_2$ ($h = K, \pi$). The amplitudes are grouped according to the orbital angular momentum L of R and h_3 in the rest frame of the $D_{(s)}$, and the total amplitude is

$$A(s_{12}, s_{13}) = \sum_L \sum_k c_k^L A_k^L(s_{12}, s_{13}), \quad (1)$$

where $s_{12(13)} \equiv m^2(h_1 h_{2(3)})$. The amplitudes A_k^L are weighted by constant complex coefficients c_k^L , and the series is truncated at $L = 2$. For a resonance with non-zero spin, the amplitude A_k is most often described using a relativistic Breit-Wigner function multiplied by a real spin-dependent angular factor [5]. For the \mathcal{S} -wave, two qualitatively different approaches exist. In the Isobar model, the \mathcal{S} -wave is treated as a sum of a constant non-resonant term and Breit-Wigner functions for the scalar resonances. While the Isobar model provides reasonably good descriptions for narrow resonances, it fails to describe the overlap of broad resonances. Additionally, the physical interpretation of the constant non-resonant term is problematic.

To address these issues, the MIPWA describes the ensemble of scalar components using a purely phenomenological set of parameters derived from the data. The $s_{12(13)}$ mass spectrum is divided into $N - 1$ slices with N boundary points separating the slices and at the two ends of the spectrum. The \mathcal{S} -wave is represented by a generic complex function $A_0(s) = a_0(s)e^{i\phi_0(s)}$. At each of the k boundary points, $A_0(s = s_k) = a_k e^{i\phi_k}$ where a_k and ϕ_k are real parameters. Between the N boundary points, the \mathcal{S} -wave is parametrised by a cubic spline [6] in the complex plane. The set of $\{a_k, \phi_k\}$ are free parameters, along with the coefficients c_l^L of the higher spin terms, are determined in the MIPWA fit. The large data sets studied, along with the large number of parameters ($2N$) required to describe the \mathcal{S} -wave, make maximising the likelihood a computationally intensive problem. Interpolating between the boundary points leads to correlations between the parameters, and hence to non-linear behavior.

3. MIPWA method with GooFit

GooFit provides an interface to allow probability density functions (PDFs) to be evaluated in parallel, using either GPUs or multicore CPUs as back-ends. While the original intention of GooFit was to utilise the massive computational power of NVIDIA GPUs based on the proprietary Compute Unified Device Architecture language (CUDA) [?], the Thrust parallel algorithms library [7] also supports an OpenMP backend for conventional CPUs.

GooFit has been used to perform a number of amplitude analyses, for example that of Ref. [8]. For a time-integrated Dalitz plot analysis, GooFit provides the `DalitzPlotPdf` class to model the Dalitz plot PDF. The `DalitzPlotPdf` object contains a list of `ResonancePdf` objects to describe resonant amplitudes as well as a constant nonresonant amplitude. The original GooFit package provided only the Isobar model to parameterise the \mathcal{S} -wave.

The `ResonancePdf` class has been extended to add support for the MIPWA method. The entire \mathcal{S} -wave is treated as a single resonance (`ResonancePdf` object), with a set of free parameters $\{a_0^k, \phi_0^k\}$ to be determined in the fit.

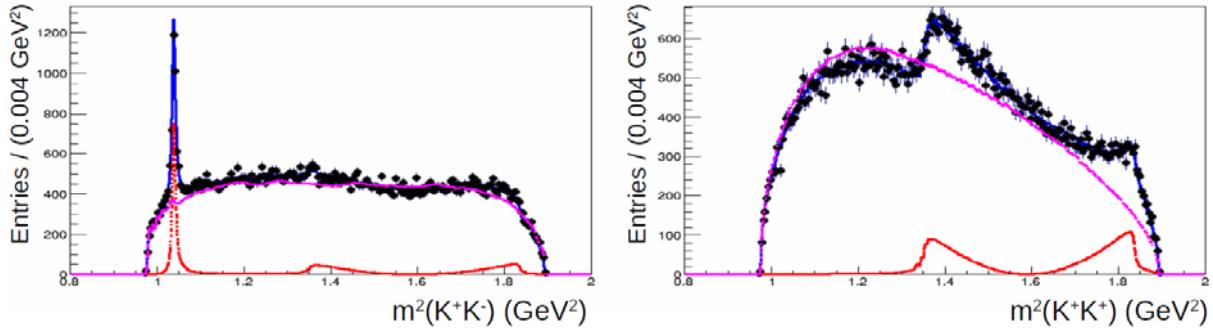


Figure 1. The projections of $m^2(K^+K^-)$ (left) and $m^2(K^+K^+)$ (right) from $D^+ \rightarrow K^+K^+K^-$. In each plot, the toy MC signal events (points with error bars) are shown together with the total fit (blue line), ϕ resonance (red line), and \mathcal{S} -wave (magenta line) determined from the MIPWA.

4. MIPWA in $D^+ \rightarrow K^+K^+K^-$ decay ¹

The $D^+ \rightarrow K^+K^+K^-$ decay offers a good opportunity to study the K^+K^- \mathcal{S} -wave amplitude directly. LHCb collected a large $D^+ \rightarrow K^+K^+K^-$ data sample from 2.1 fb^{-1} of $\sqrt{s} = 7 \text{ TeV}$ pp collisions recorded by the experiment during 2012. About 100K candidates were selected with a signal purity of 90%. The resonant structure of the K^+K^- \mathcal{S} -wave amplitude was studied for the first time using an Isobar model, and the results were first presented during the CHARM 2016 workshop [9]. This analysis indicates that the \mathcal{S} -wave component accounts for about 90% of the decay rate, while the \mathcal{P} -wave ($\phi(1020)$) resonance makes up the rest.

A specific goal of the work presented here is to provide a tool to extract the K^+K^- \mathcal{S} -wave by performing an MIPWA on the same LHCb data sample. In this case, the Dalitz plot is symmetrised along the two axes of $m^2(K^+K^-)$. The K^+K^- mass squared range is divided into 39 slices using 40 boundary points. Because the final state contains two indistinguishable K^+ mesons, the K^+K^- amplitudes interfere with themselves, and this sort of interference allows the \mathcal{S} -wave phase to be determined. To test the GooFit extension to support MIPWA fits, samples of 100K signal events are generated from a Dalitz plot PDF (“Toy MC”). The mass projections from the fit to a test sample can be seen in figure 1. First, the fit quality of the MIPWA method is tested by comparing the fitted $\{a_k, \phi_k\}$ values with the inputs. In each fit iteration, the Dalitz plot normalisation method (“`DalitzPlotPdf::normalise()`”) is called so that the integral of the total PDF is equal to one. The normalisation integral is calculated numerically based on evenly distributed grid points in the Dalitz plot plane. Figure 2 shows the improvement in the fit quality as the normalisation grid spacing is reduced. Although the finer granularity increases numbers of calculations, the high speed of the GPUs makes this problem tractable.

Table 1. GPU performance of the MIPWA fit.

Platform	GPU Model	Chip	CUDA cores	Run time (sec.)
tWorkstation	Tesla K40c	GK110BGL	2880	76
Desktop PC	GeForce GTX 980	2nd gen. Maxwell (GM204)	2048	67
Laptop ASUS N56V	GeForce GT 650M	GK107	384	179

¹ Charge conjugation is implied throughout.

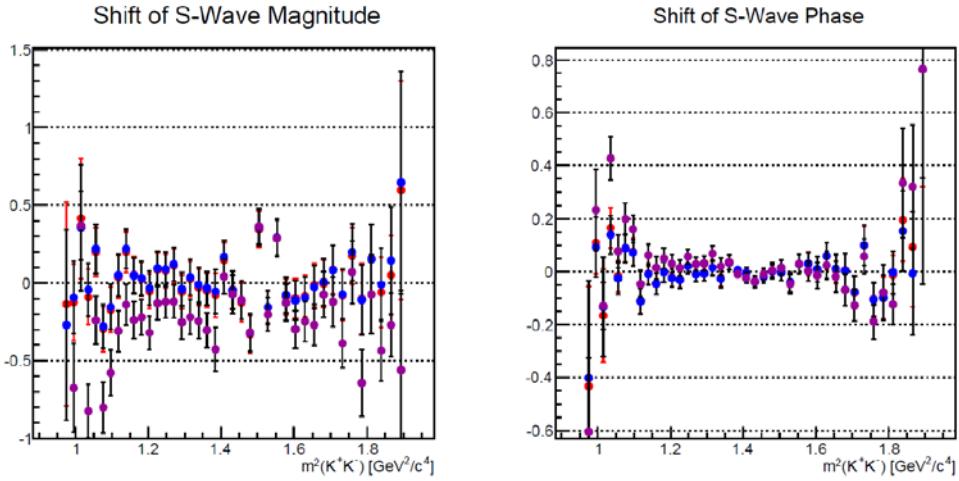


Figure 2. Differences between the fitted values and input ones for each boundary point. Shown are set of points with different grid spacings for the normalisation: 0.01 GeV^2 (purple), 0.004 GeV^2 (red), and 0.001 GeV^2 (blue).

As shown in table 1, the GooFit performance on three different GPU platforms has been measured by running the MIPWA fit over the same toy MC sample of 100 K events (see figure 1 for the fit projections). For comparison, the MIPWA fit using an older code with the same functionality that runs on one CPU core takes about eight hours to complete. Perhaps surprisingly, an older generation mobile GPU (a GeForce GT 650M with 384 cores) provides excellent performance; a newer HPC GPU board (a Tesla K40c with 2880 cores) provides better performance, but not in proportion to the number of cores; a high-end gamer board (a GeForce GTX 980) provides the best performance, albeit by a small margin. Based on other studies with GooFit, significantly better performance on the new P100 boards is anticipated which utilise NVIDIA's new Pascal GPU architecture.

Table 2. Specifications of the testing CPU platforms. Asterisks next to the number of cores indicate hyperthreading - two virtual processors per physical core.

Name	Chip type	# of Cores	Clock [GHz]	RAM [GiB]
Cerberus	Intel Xeon E5520	8*	2.27	24
Goofy	Intel Xeon CPU E5-2680 v3	24*	2.50	120

In addition to timing GooFit's new MIPWA performance on GPUs using Thrust's CUDA backend, its performance on two different CPU platforms is timed using Thrust's OpenMP backend, as shown in figure 3. With two Intel Xeon E5-2680 v3 CPUs, the fit uses 791 seconds for one OpenMP thread, 50 seconds for 24 threads. The speedup is almost linear with the increase of the number of OpenMP threads, up until the number of threads equals the numbers of physical cores.

5. Summary

This paper describes an extension of GooFit to support MIPWA fits for three-body decays, and have achieved speedups of a few hundred by using GPUs. The main branch of GooFit's source code is in a GooFit repository at <https://github.com/GooFit/GooFit>, while the updated code

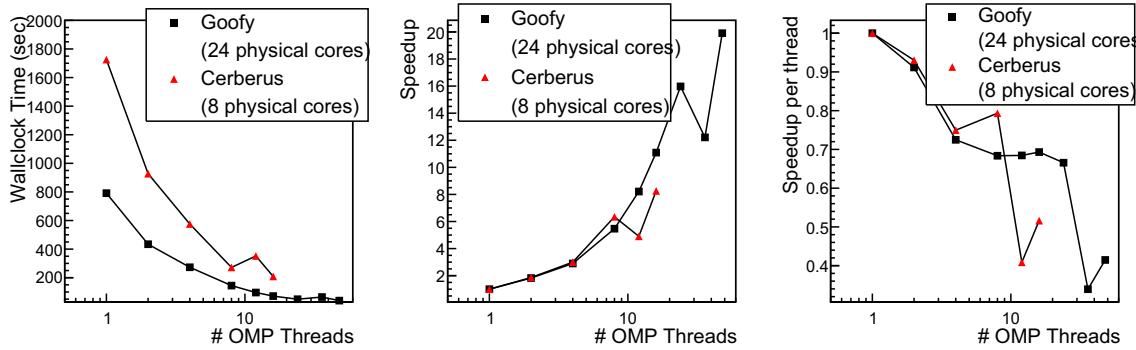


Figure 3. Timing results and speedups for the test MIPWA fit. Tested on two different CPU platforms as shown in table 2: Goofy (Intel Xeon CPU E5-2680 v3 x 2) and Cerberus (Intel Xeon CPU E5520 x 2).

with MIPWA support is in a personal GooFit branch at <https://github.com/liang-sun/GooFit>.

5.1. Acknowledgments

This work was performed with support from NSF Award PHY-1414736. NVidia provided K40 GPUs for our use through its University Partnership program. The Ohio Supercomputer Center made their “Oakley” computer farm available for development, for testing, and for GooFit outreach workshops.

References

- [1] Patrignani C *et al.* (Particle Data Group), *Chin. Phys. C* **40**, 100001 (2016).
- [2] Dalitz R H, 1953 *Phil. Mag.* **44** 1068.
- [3] Aitala E M *et al.* (E791 Collaboration), *Phys. Rev. D* **73**, 032004 (2006).
- [4] Andreassen R *et al.*, *IEEE Access* **2**, 160 (2014).
- [5] Kopp S *et al.* (CLEO Collaboration), *Phys. Rev. D* **63**, 092001 (2001).
- [6] Press W H *et al.*, *Numerical Recipes: The Art of Scientific Computing, 3rd Edition*, Cambridge University Press, 2007.
- [7] Bell N *et al.*, *Thrust: A Productivity-Oriented Library for CUDA*, available from the Documentation link at <https://thrust.github.io/>, accessed 29 January, 2017.
- [8] Lees J P *et al.* (BABAR Collaboration), *Phys. Rev. D* **93**, 112014 (2016).
- [9] LHCb Collaboration, LHCb-CONF-2016-008 (2016).