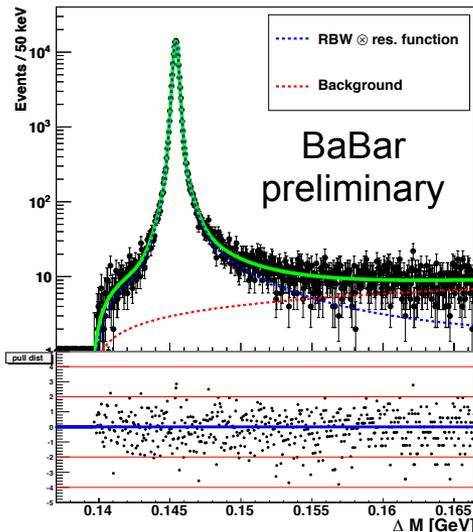


GooFit for HEP Analysis

Karen Tomko, Ohio Supercomputer Center
Rolf Andreassen, Univ. of Cincinnati
Michael Sokoloff, Univ. of Cincinnati

- Exploring hybrid (MPI and CUDA/OpenCL) parallel implementations for fitting large datasets and complex models
 - Developing a package, GooFit, similar to CERN's RooFit using GPUs to parallelize PDF evaluations
 - Will be shared within High Energy Physics (HEP) community and others with similar data analysis problems such as Cosmology
 - Current implementation uses Thrust Template Library for CUDA

- Why GPUs for parallel fitting?
 - Data sets from next generation accelerators are expected to be 100x larger than prior experiments
 - Goal of reducing complex analysis time to overnight run
 - Massive parallelism well suited to SIMD execution in maximum-likelihood fit
 - Many PDFs applied to a set of events can amortize data transfer costs



Fit with 16 parameters describing 154k events. The variable of interest is the underlying width of the mass distribution, which is obscured by the unknown detector resolution. The fit extracts both quantities. The lower part shows the deviation of the fit from the data points in units of the error.

- Single GPU results, multi GPU implementation underway

		Minutes	Speedup
Toy Example*	RooFit (on 1 core)	240.97	1
	MPI (8 processes)	11.53	21
	MPI (16 processes)	5.65	43
	GPU (Nvidia C2050)	0.50	484
Physics Example**	RooFit (on 1 core)	1632.00	1
	GPU (Nvidia C2050)	8.30	197

* Single Gaussian

** Mass distribution convolved with sum of Gaussians, plus background component