Applicability of Parallel Computing to Partial Wave Analysis (PWA)

J. Ruger - Christopher Newport University
G.P. Gilfoyle - University of Richmond
D.P. Weygand - Jefferson Lab

Introduction

Jefferson Laboratory (JLAB) is a US national laboratory built to probe the atomic nucleus and to unravel the quark-gluon nature of matter. Our best theory of the strong interaction that binds nuclei together is Quantum Chromodynamics (QCD) which is not well understood in the nuclear environment. We are developing new technologies to analyze the large data sets that will be produced by one of the large detectors under construction in Hall B as part of the 12GeV Upgrade.

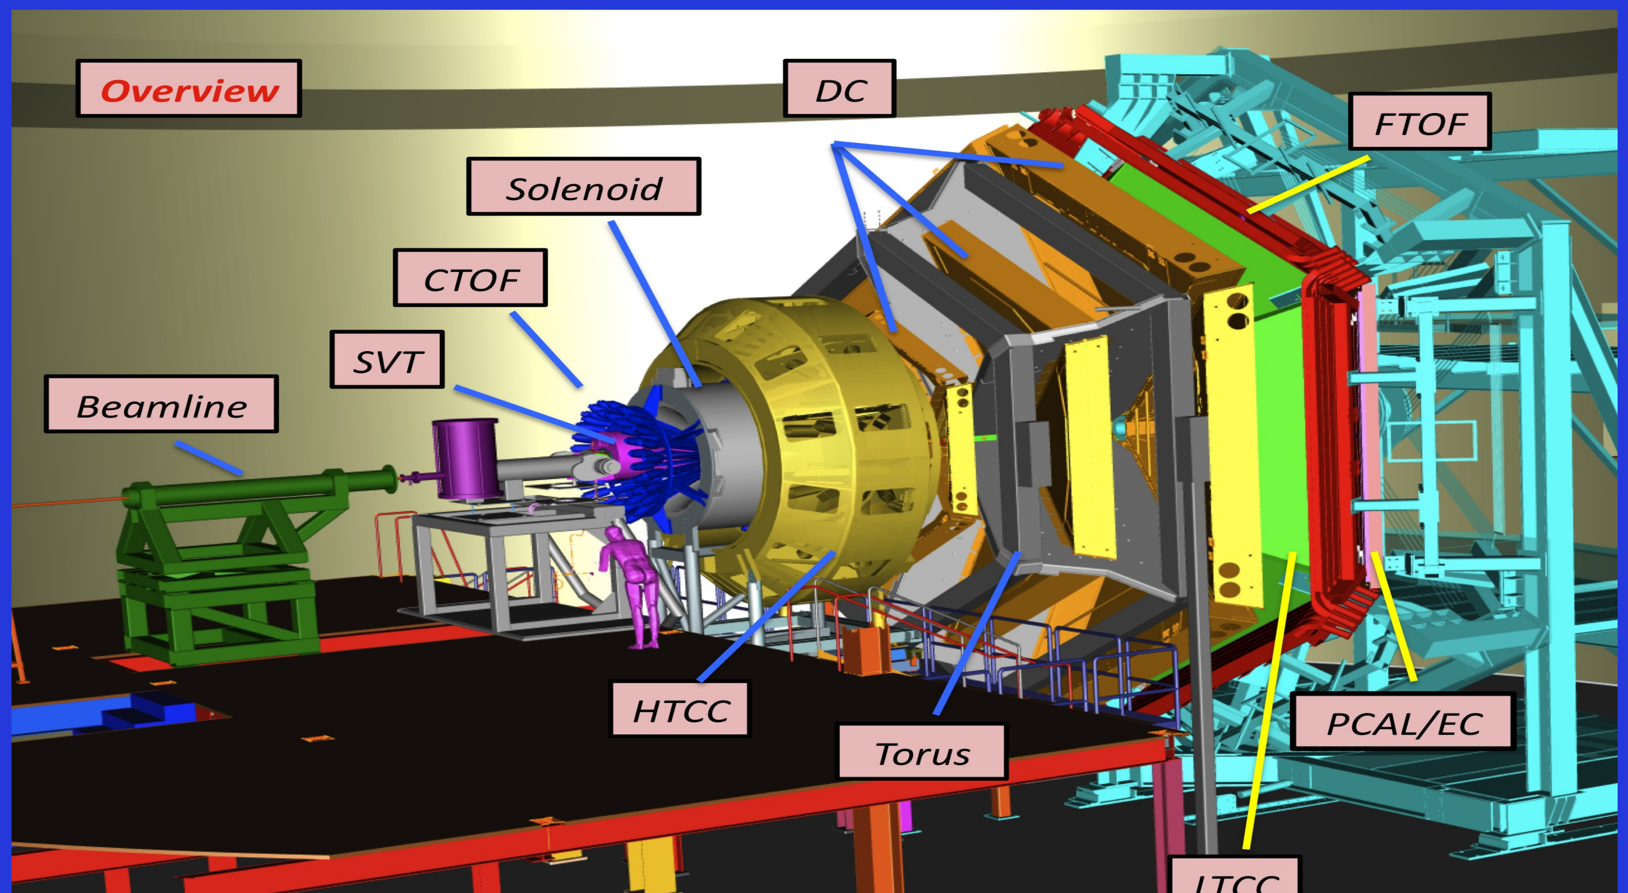

CLAS12

The primary goal of experiments using the CLAS12 detector at energies up to 12 GeV is the study of the internal nucleon dynamics. Towards this end, the detector has been tuned for studies of reactions in a wide kinematic range such as inclusive processes. The large acceptance of CLAS12 leads to high data rates since most of the debris from the collision of the electron beam with the target is detected. The computational requirements to acquire, reconstruct, analyze, and simulate these data are high. We expect to collect about 10,000 events per second where each event contains 10 kBytes of data. For typical running, this rate means we will gather about 5,000 GBytes per day and have to store about one million GBytes each year. It is the goal of our group to have the software in place on the day of first beams in CLAS12 to calibrate, reconstruct, and analyze these new data and to keep pace with the high rates.

Motivation

To understand the strong interaction, or Quantum Chromodynamics (QCD) it is necessary to understand the spectrum of bound states of the interaction, that is mesons and baryons. One of the best tools for the analysis of these reactions is Partial Wave Analysis (PWA). Using PWA we can extract the properties of individual particles from overlapping mass spectra as seen in Fig. 3 and Fig. 4. We, when analyzing CLAS12 data, will have to analyze millions of events that may require these computationally intensive calculations.

PWA requires fitting angular momentum states to large sets of data. At 700 ms per event per wave it will take roughly nine months to do one fit on the data scale that CLAS12 acquires. This timing analysis takes into account the number of waves, number of events and likelihood calculations.

The motivation for pursuing research with Xeon Phi's and its use in Partial Wave Analysis is to minimize the overall time required to perform calculations. By exploiting the vector and multi-threading capabilities of the MIC processor available from Intel, the likelihood function (Fig. 4) is multi-threaded to enhance performance of PWA calculations.

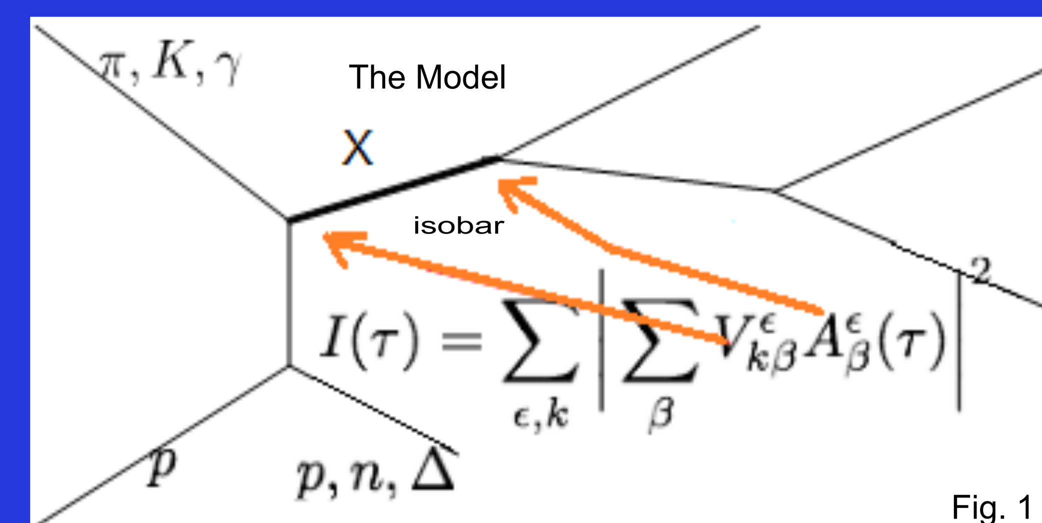

Fig. 1

Likelihood function:

$$\mathcal{L} = \frac{\bar{n}}{n!} e^{-\bar{n}} \prod_{i}^{n} \frac{I(\tau_i)}{\int I(\tau)\eta(\tau)d\tau}$$

Fig. 2

Our goal is to find a mathematical parametrization that explains the experimental observation that is differential cross sections (Fig. 1). A major portion of this parametrization is the calculation of the Likelihood function (L), which is defined as a product of probabilities. L is calculated using the quotient of the Intensity Distribution (I(τ)) and normalization integral. In L, the term outside the product in Fig. 2 is the Poisson probability of observing n events and reflects the fact that we use the extended maximum likelihood method. The integral in the denominator of the summed term contains the acceptance η(τ) and is the accepted normalization integral.

An independent calculation, I(τ), is a sum of amplitudes, squared to account for interference with the variable τ representing the set of variables necessary to define a configuration of the final state. Within I(τ), the variables V and A, represent the production and decay amplitudes respectively. The subscripts k and β for V and A, are the parameters that describe the partial wave decomposition we are using, α specifying properties of the different intermediate states that do not interfere, such as the spin states of the incoming or outgoing particles in the detector. The subscript β, on the other hand, represents the properties whose differing values do interfere, for instance the spin states of broad resonances produced as intermediate states in a sequential decay.

Using the calculations from above, PWA is able to extract the properties of the individual particles as seen below. Figure 3, represents the raw mass spectrum of $\pi^+ \pi^+ \pi^-$ from electron scattering on the proton. While, Figure 4 represents the extracted properties from the raw mass spectrum. As seen below, using PWA allows us to find the π2(1670) particle that was hidden behind the a2(1320) in unprocessed data.

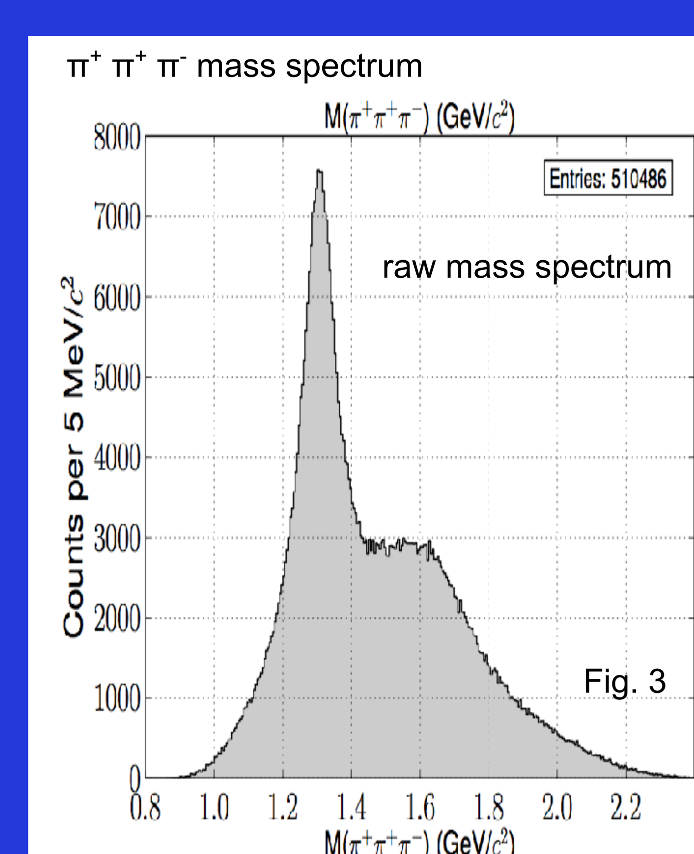

Fig. 3

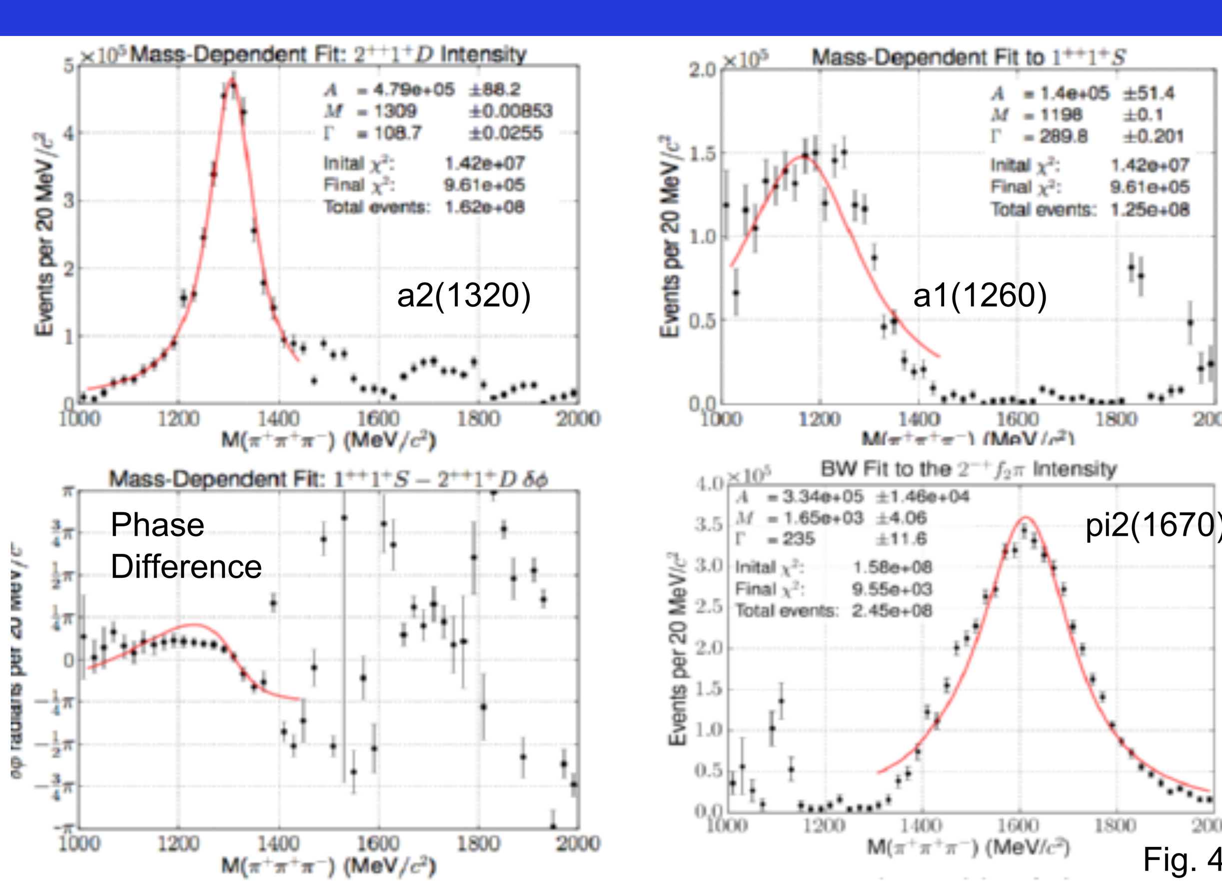

Fig. 4

Intel Xeon Phi

* Jefferson Lab has a 16-node Intel Xeon Phi (MIC) cluster
* Each node contains 4 Intel Xeon Phi 5110P Co-processors
* Intel Xeon Phi 5110P
  * 60 cores for a total of 240 threads at 1.053 GHz
  * Supports 16 single precision or 8 double precision calc.
* Jefferson Lab offers unique capability of running applications either in native mode on the CPU or through offloading from the host.
* Creating multi-threaded software packages is easy with the Phi's using OPEN-MP and OPEN-MPI.
* Code created for the Xeon Phi can be ran on single threaded platforms or GPU increasing value.

Computation Results

A representation of the processing power gained from the inclusion of a Xeon Phi is a comparison of the run time for PWA log-likelihood calculation versus the number of calculations deployed on the Xeon Phi and Host machine. Figure 5 shows that using the Xeon Phi decreased the processing time by a factor of 16. This will reduce the time need to do the 9 month PWA fit talked about in motivation to about two weeks.

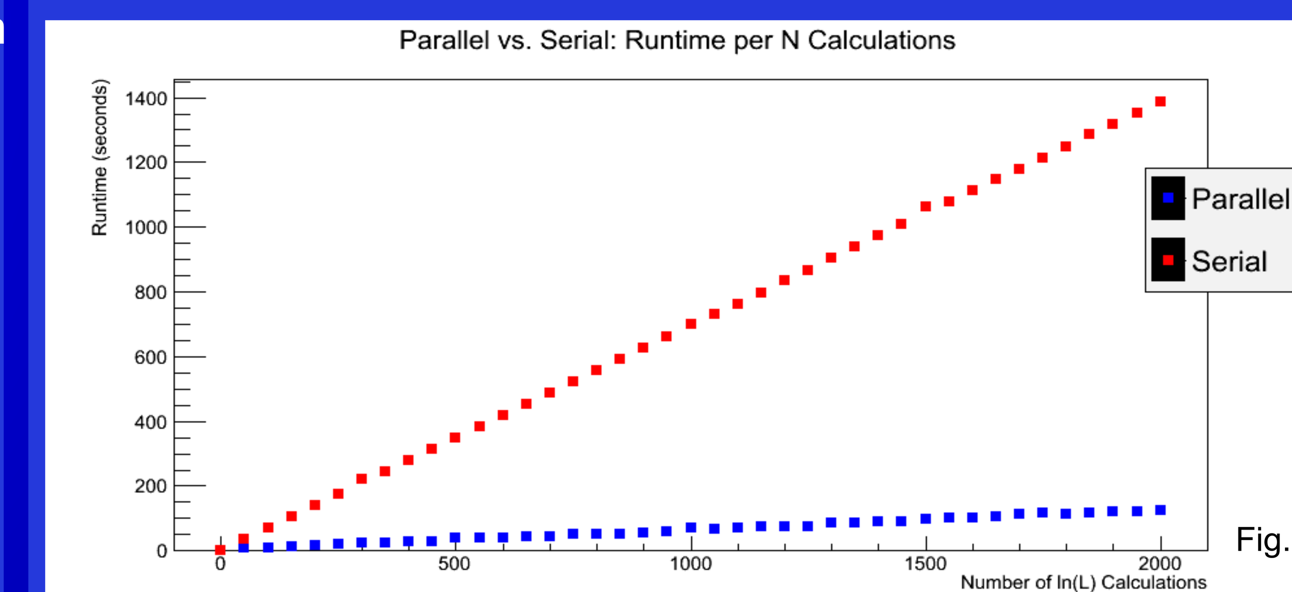

Fig. 5

To show the ability of the Xeon Phi to multi-thread log-likelihood calculations, Figure 6 represents the parallel scaling over available threads. The data showed that maximum efficiency was found using half the available threads. This will limit the amount of time spent in thread management having 120 threads always available.

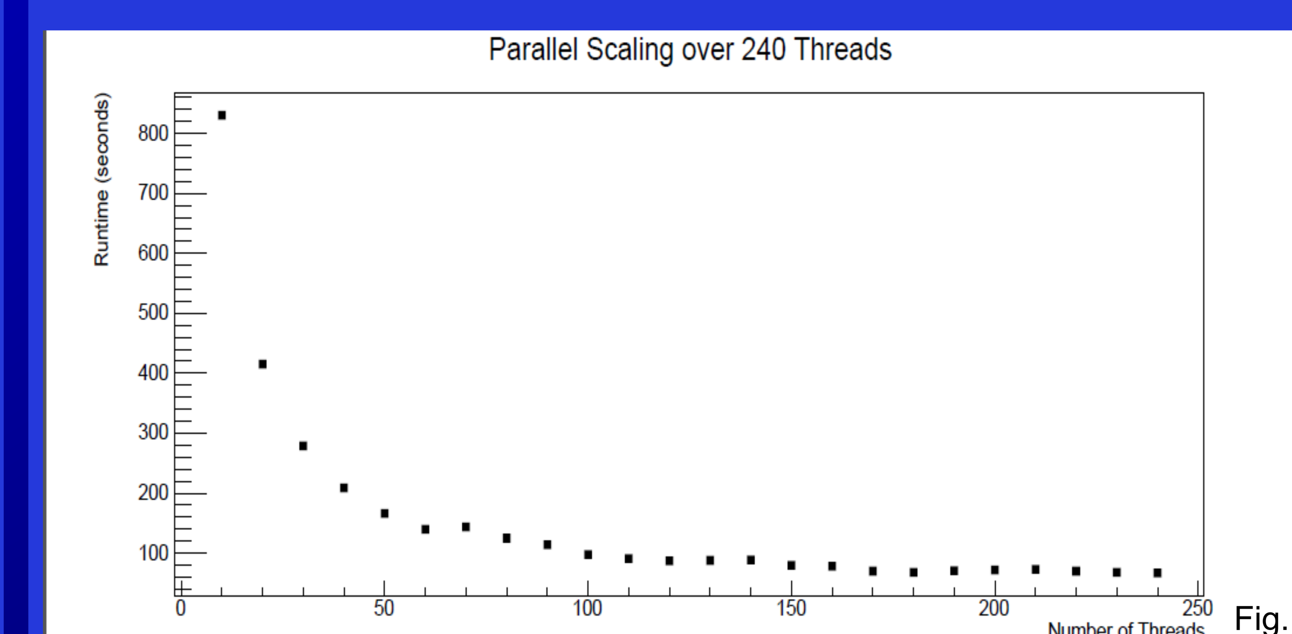

Fig. 6

The results show promise in using the Intel Xeon Phi for computationally intensive calculations. The ability to decrease run time by a factor of 16 allows data processing to happen quicker and minimizes the amount of time spent per fit.

Acknowledgments

I would like to thank Dr. Jerry Gilfoyle and the University of Richmond for the funding and opportunity to research PWA. Second, I would like to thank Dr. Dennis Weygand for the computational guidance received during the project. Last, I would like to thank Dr. Balint Joo for his knowledge and guidance when working with the Xeon Phi and Jefferson Lab's Xeon Phi Cluster.