C-A/AP/68
January 2002

Experimental Study of Beam Parameters Extracted from a Flag in the U Line – an attempt to quantify errors

Leif Ahrens

Collider-Accelerator Department
Brookhaven National Laboratory
Upton, NY 11973

C-A AP Note (A Beam Study Note)

Experimental Study of Beam Parameters Extracted from a Flag in the U Line – an attempt to quantify errors

Leif Ahrens January 2002 (7Feb02)

UF3 (a thin 1 mil aluminum + 2 mils phosphor flag in the U transport line) sending light into a CCD camera and from there into Frame Grabber # 4. A low intensity (1e11) proton beam creates the light. Details on the flag available on the RHIC Instrumentation group web page, link from C-A Acc Div page.

Objective: to quantify the appropriate error to associate with the beam profile widths extracted from the flag measurements in AtR. Why is this important? These beam widths are used to deduce the transverse emittances for both AGS beam and RHIC beam. They provide what would seem to be a very robust measurement. The density distribution of the beam in transverse coordinates at the end of the acceleration cycle in AGS is expected to follow a nearly Gaussian distribution. The tails (i.e. beam at 3 or more sigma from the center) no doubt deviate from a Gaussian, but for most beam size discussions the tails are not the point of the exercise since they represent only a tiny fraction of the beam. Falsely generated tails and also falsely supressed distribution peaks leading to distortions in the Gaussian fit would be a problem.

A second objective is to validate a simple procedure to allow identification of "good" width measurements. What is conjectured is that if the two rather different fitting procedures available in the AtR flag profile program (see below) both fit the data well and both agree on the profile widths then we have the right answer.

The responses given to the question of beam width errors fall into two groups: 1) assume 10% based on undocumented experience, or 2) make many identical beam width measurements and extract the statistical error associated with these, the standard deviation of the resulting distribution of widths. Then the appropriate estimate of the error in the mean of this distribution, assuming random errors in the many measurements is that the error goes down as root N where N is the number of measurements. The second approach is well defined and justified to the extent that systematic effects are insignificant. The first approach is unfortunately more appropriate if one can show that the systematic error is indeed of order 10% and dominates the statistical errors. If the subtle details of how the data is collected and processed from a flag result in a change in the extracted beam width that varies by say 10% then simply taking lots of measurements doesn't help. We want to try to characterize at least some potential systematic effects and so perhaps make some headway on this question.

Plan: take beam profiles with nominally unchanging beam from AGS, with one flag, CCD camera, analogue light link, and one frame grabber. Vary the attenuation (neutral density filters) between the flag and the camera. For a given attenuation, fit a given beam

picture (the projections in the x and y directions) using two algorithms which have been developed to cope with this data - the “first pass” (1P) fit and the “least squares” (l.s.) fit.

These fit algorithms are built into the flag acquisition application program – which belongs to Steve Tepikian. The fits are optimized in different ways, which we do not claim to fully understand. The 1P fit starts from a pure Gaussian, invokes some relative weighting of the center region relative to the tails to get reasonably performance, and does not involve a least squares fitting optimization. The l.s. fit does involve a least squares fitting not unexpectedly, but to get reasonable performance allows the distribution to have non-gaussian tails.

The plan then is to apply both fittings to a given data set for several beam extractions and then repeat the entire sequence for several attenuations. (This was done on 23 Jan.) If the extracted beam widths depend strongly on the attenuation used, the whole idea of extracting a beam size from the flag is seriously damaged. What is the right answer? If there is a range of attenuations over which the beam width is nearly independent of the attenuation, and if we can ascertain that we are comfortably within this window, perhaps systematics at least from this source are not important.

Analysis: Enter the ‘raw data’ into a spread sheet (i.e. the fit results – sigmas - 1P and l.s. and other measurements provided by the program for each picture). For each attenuation, extract the average sigma and its standard deviation over the data taken. The 23 Jan data stepped through six attenuations, and each attenuation had at least two pictures – two data sets. For convenience, in order to simultaneously compare the behavior of horizontal and vertical profiles, the sigmas of each type (horizontal and vertical) is normalized to its average. This average is taken to be just the average sigma over the attenuations (transmissions now) of 3.3%, 6.8%, and 8.8%.

On 24 January (the last day of the polarized proton run), having had a quick look at the first set of data, we took a little more data with the same goal, and (or but) with a (surprisingly) different beam. This was a smaller set with only transmissions 1.3, 3.13, and 6.8; and not much repetition – e.g. only one measurement at 1.3.

“Typical” good profiles are shown in figures 1 and 2. Figure 1 is taken with less attenuation (more transmission) than figure 2, and the results of the two types of fits are provided by parts a) and b). The largest deviations between fit and data is in figure 1a) with the fit (the red curve if you have color) having a smaller width than the data in the horizontal, and a larger width in the vertical. The l.s. fit matches the data much closer. The widths from the fits are given in millimeters. I apologize for the scale change between the two attenuations - between fig 1 and fig2.

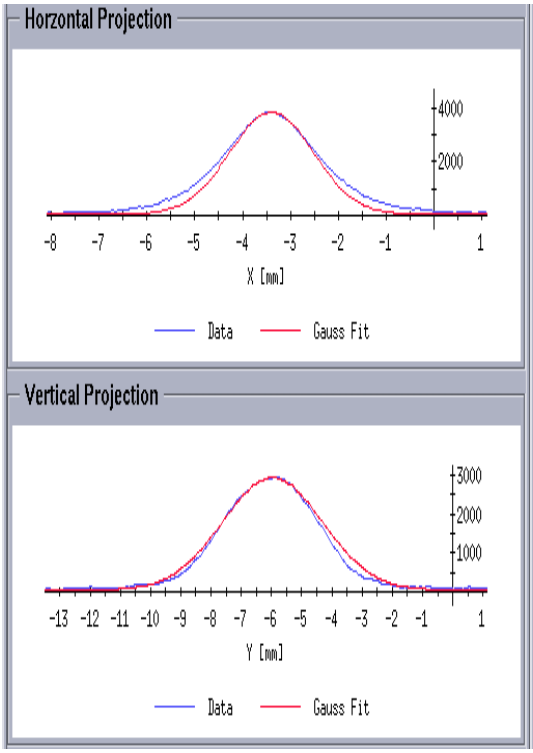

Figure 1a IP fit 6.7% trans (σ .87,1.65)

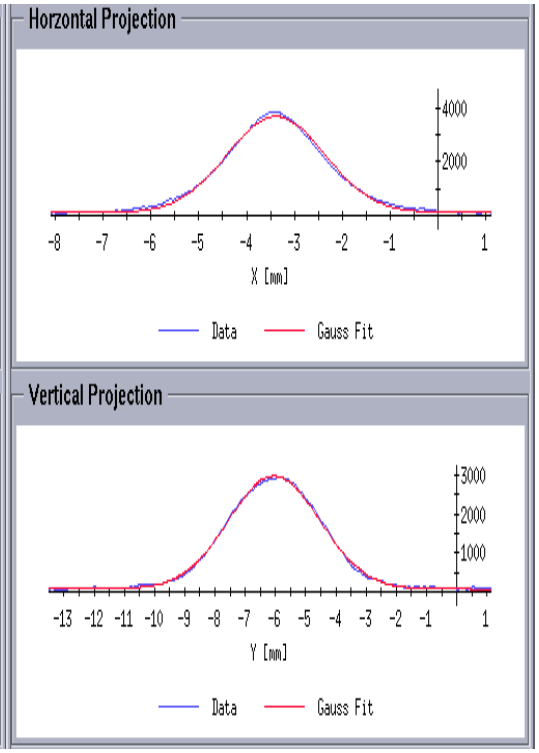

Figure 1b l.s. fit 6.7% trans(σ 1.01,1.51)

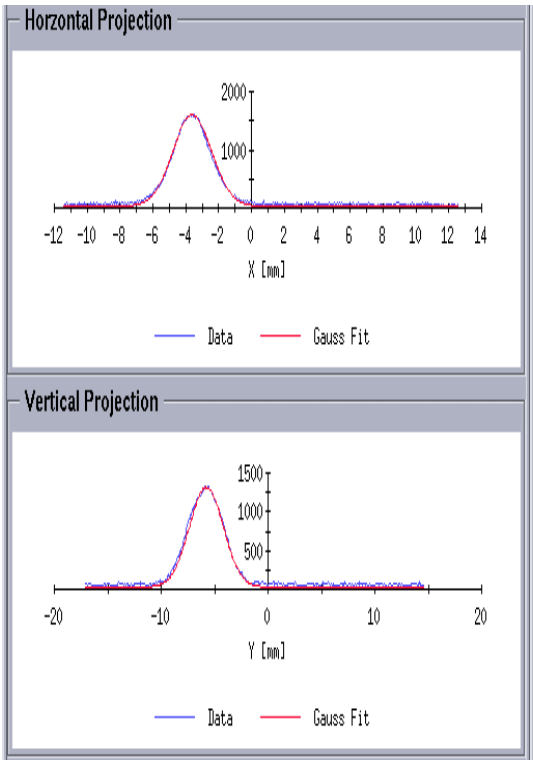

Figure 2a IP fit 3% trans(σ 1.19,1.59)

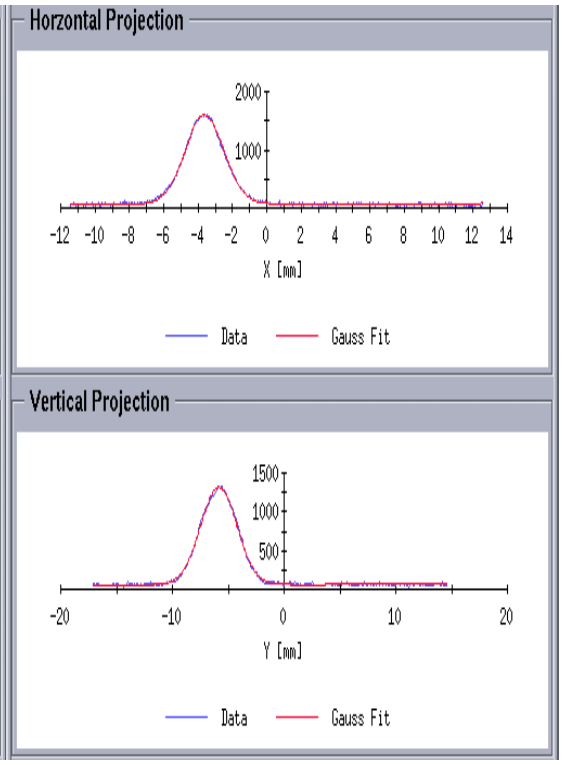

Figure 2b l.s. fit 3% trans(σ 1.14,1.62)

What sort of systematic error might one predict to see in an attenuation scan? As the image supplied to the camera gets stronger, surely the tails of the distribution become more visible and hence more measurable. If there is any saturation of the image near the peak then the extracted width would be expected to grow in a Gaussian fit. On the other end, as the image intensity gets very faint information from the tails must disappear into the noise. The core should determine the fit, the quality of the fit must ultimately suffer. Figure 3 gives the result from the first day of scanning.

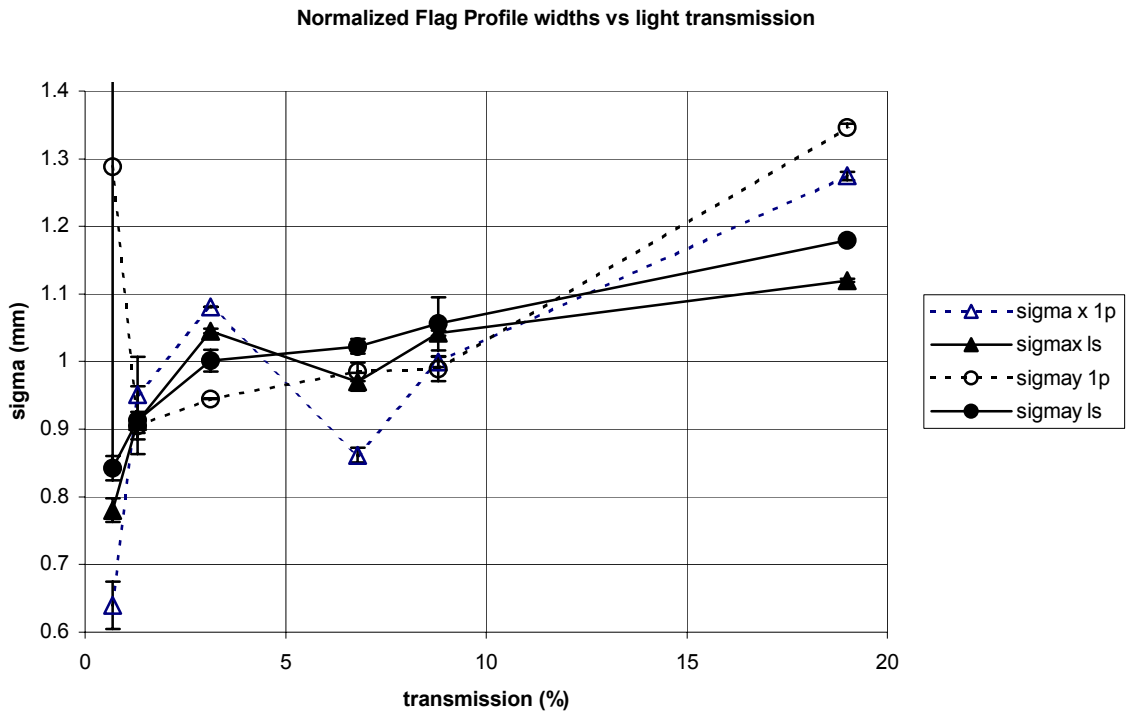

Figure 3 Beam Widths vs Attenuations for the two fitting procedures

Figure 3 is a bit confusing. The triangles identify the horizontal profile fits, the circles give the vertical profile fits. Horizontal and vertical (x and y) curves overlay on average "by construction". The two types of fits are shown as the open (1P) and closed (l.s.) markers. The error bars measure the standard deviation of sigmas reported at the same setting. The statistics for this are usually small and not equal for different attenuations. The fits have in general the expected behaviors at the extremes. There is apparently something about the least squares fit that produces a narrower width at low camera intensity. The desired "attenuation independent" region is perhaps evident, more clearly in the least square fit. To further investigate whether the "dip" in the reported horizontal width near the 6% transmission is a systematic effect or a statistical fluke was one motivation for the second round of data collection.

Not intentionally, the second scan was done with a significantly lower intensity beam, shown in figure 4. However the beam was also significantly smaller in width. Up till now we have been plotting against transmission setting. From figure 4, it is clear that we could equally well have been plotting against the reported flag intensity –

transmission and intensity are linearly related. If the systematics we are searching for depend on the peak intensity that the camera or anything downstream has to deal with, then plotting against intensity divided by the product of the beam widths in x and y is appropriate, as the resulting number should be appropriate to the central intensity. The sigmas used for this scale change are the averages from the reasonable part of the scan. Figure 5 shows what happens for the flag pictures taken when this calculated "peak" intensity is plotted against the transmission setting. For the new run, with smaller intensity but also smaller beam widths the effects nearly cancel. Nevertheless we move to plotting against this calculated peak intensity rather than transmission setting.

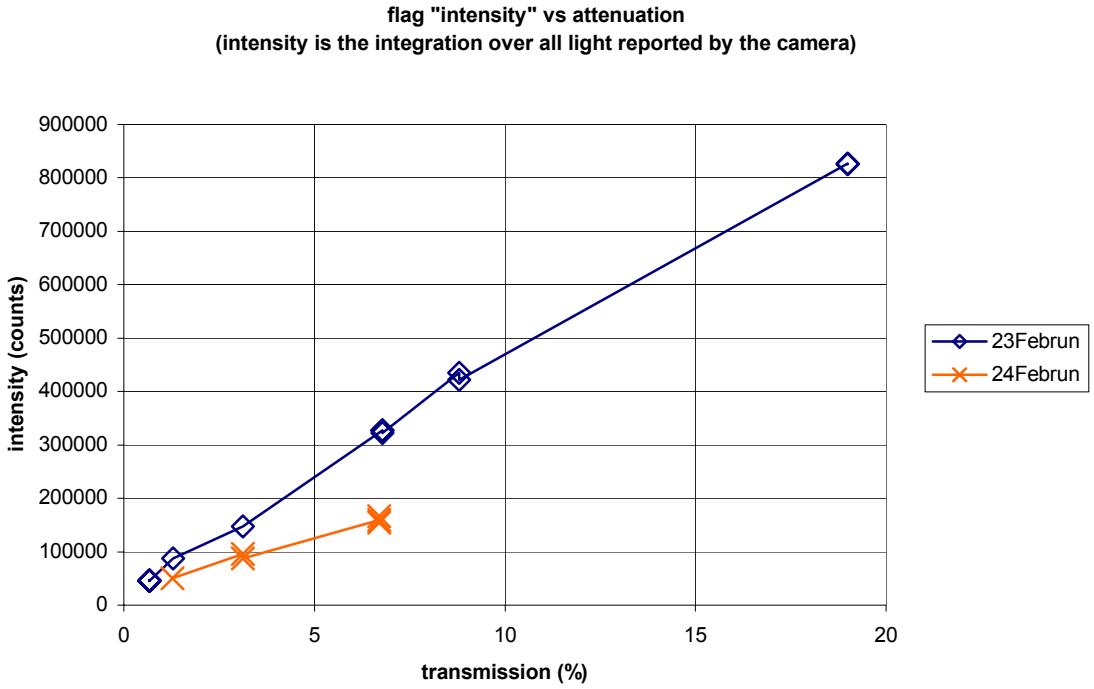

Figure 4 Total Flag Intensity for the scans on the 23rd and 24th

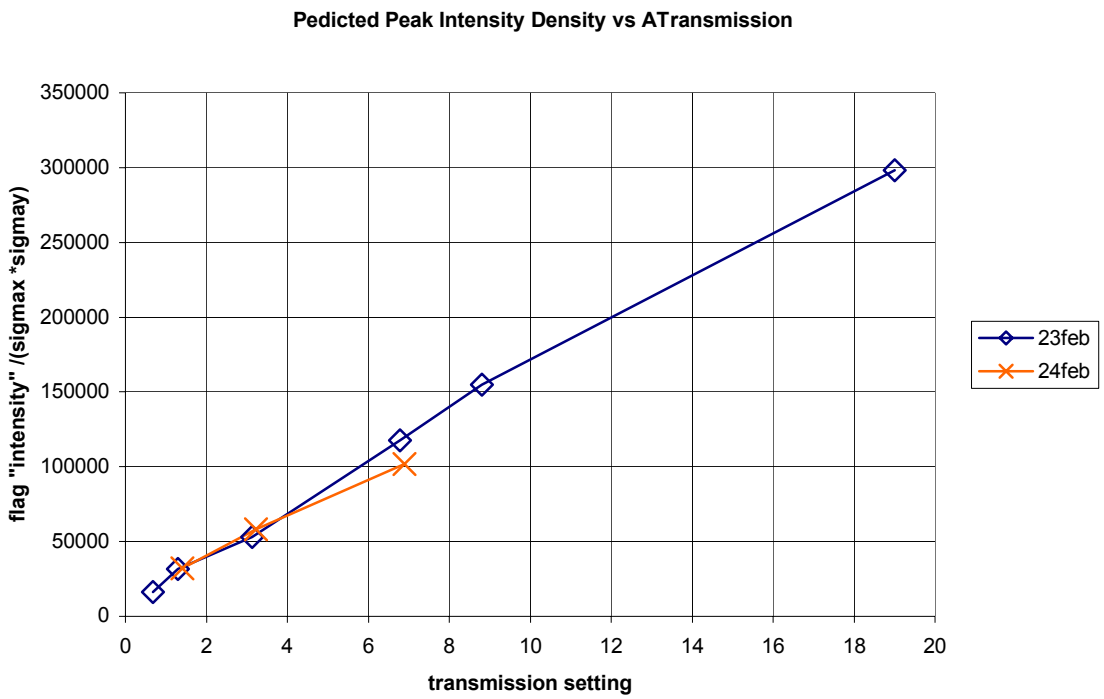

Figure 5 "Peak" flag "intensity" for the scans on the 23rd and 24th

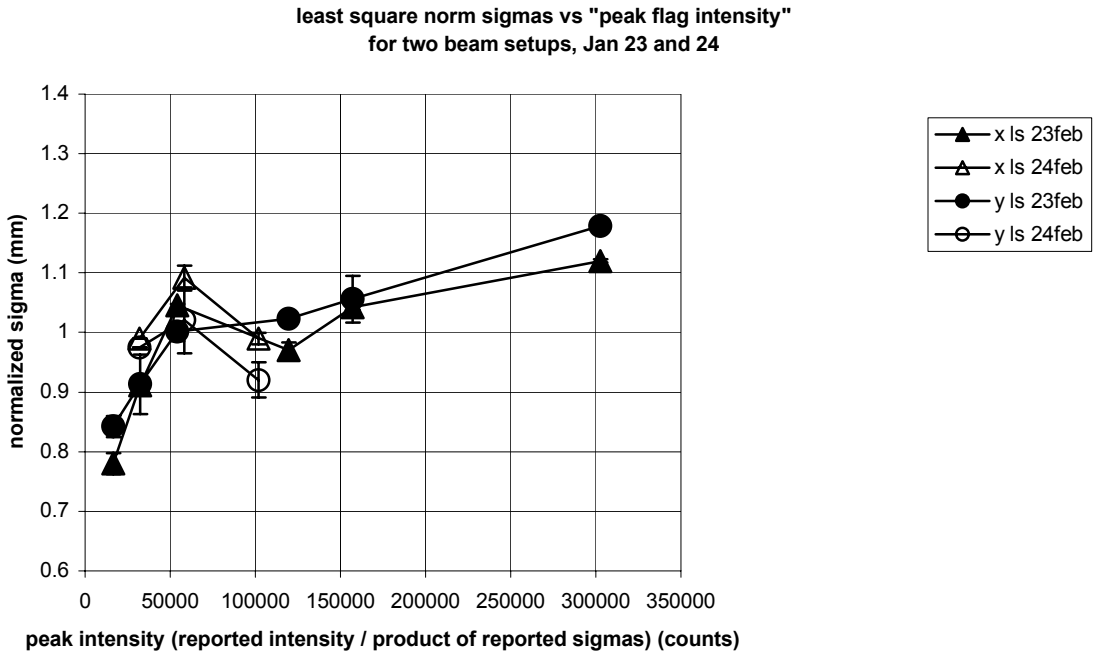

Figure 6 Profile Widths vs Peak Intensity, now combining both run days, but (for clarity) only results from the l.s. fit procedure

Some tentative conclusions:

Figure 6 is then the same presentation as figure 3, except looking only at the 'l.s.' fitting, and including another day of data, taken under very different beam conditions, which nevertheless yielded very similar flag local intensities. We probably got what we deserved. The speculation that there might be a "dip" in the extracting of a width is certainly not ruled out – which was what we had desired. More generally the conclusion to be drawn here – is there some sort of systematic effect with peak intensity – is likely though not rigorously proven. We don't at the moment have a mechanism for a "dip" in the middle of the attenuation scan, and for the horizontal data there may indeed be no need to postulate it, but **to set a systematic error at +/- 5% would be a reasonable somewhat informed "guess" – for data acquired in the ideal range** - away from the clear low and high deteriorations. Taking repeated measurements till the cows come home won't get rid of this uncertainty.

If we stray into the situation associated with the extremes in this data collection then we clearly get hit by systematic effects of large magnitudes - 15 or 20% for the least squared analysis, worst for the 1P. To avoid this we return to a suggested simple procedure to judge how good a flag setup is for a given beam situation. It appears consistent with this data that if the two procedures , 1P and l.s. disagree substantially then we do not get a good width determination. Can we then use a measurement showing that they agree as indication that the fit is good? Figure 7 gives this analysis for the above data. The error bars are misleading as there is too little data in some bins – namely only one measurement – which show up here as perfect points.

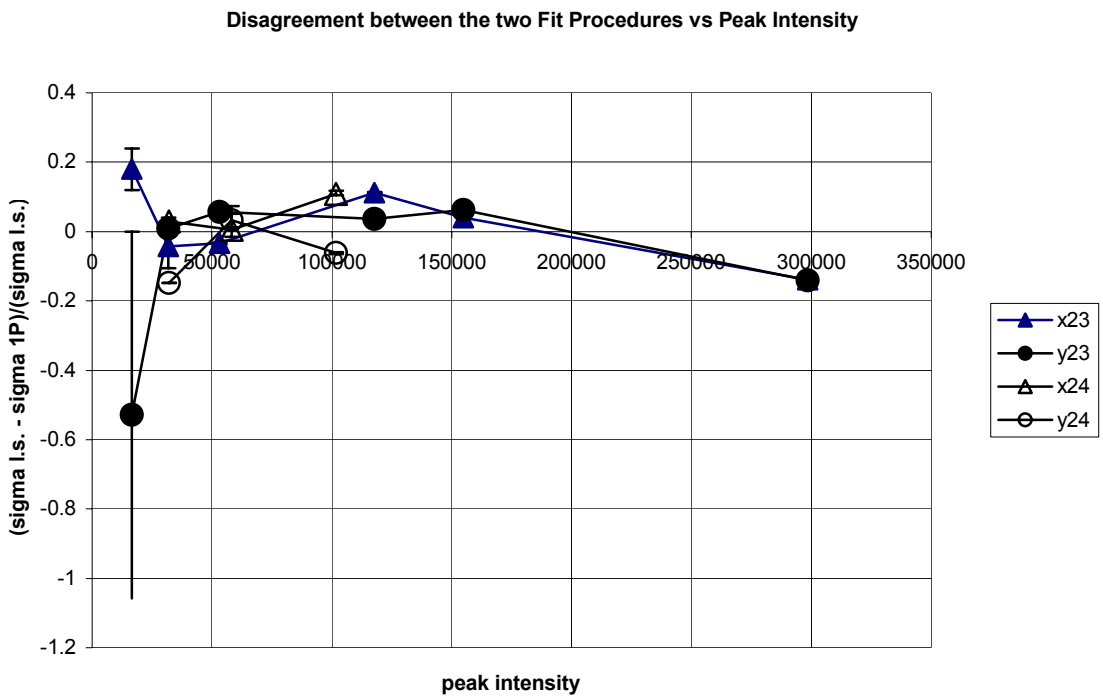

Figure 7 The normalized width disagreement between the two fitting procedures

Apparently **insistence, perhaps by an automatic procedure of attenuation adjustment, on agreement between the two fitting procedures – down to +/- 10% -** would move the setup to the peak intensity region of about 75000 (Intensity_counts/((mmσx)*(mmσy))), and this **might not be a bad place to be**. There is good indication of a systematic differing between the two fitting procedures at high peak intensity (1P reporting a wider beam than l.s., and both larger than what we believe is the true width) and with less certainty and in the other direction (l.s. wider than 1P) at intermediate peak intensity – another way of seeing the "dip" noted before.

We leave this subject with several questions that need experimental verification. Is our particular flag and frame grabber special, or can a similar rule be applied to the others? Was our beam special? Does such a rule apply to the high intensity proton beam in the U-V line? In fact, other people who have worked more with these flags may already know these answers as well as better questions to ask.