Voyager Plasma Science at Jupiter
Error Analysis

Logan Dougherty, Kaleb Bodisch,
Rob Wilson, Fran Bagenal, Frank Crary
LASP University of Colorado Boulder

The purpose of this document is to address two issues related to fitting the Voyager Plasma Science (PLS) data obtained at Jupiter in March and July 1979:

(1) Evaluation of **errors in the measurement** of energy spectra. This is important for the fitting process since we need to weight the data with good vs. poor quality appropriately.
(2) Evaluation of **uncertainty in the fit parameters** derived from fitting the energy spectra with a sum of convected Maxwellian (or similar) functions.

To illustrate how we address these two issues we apply our analysis to two of our favorite "typical, best" cases – Voyager 1 Day Of Year (DOY) 64 1016 (we call Fred) and DOY63 15 35 (we call George).

Fred - Obtained at 5.269 RJ Fred is typical of the cold, inner torus with multiple well-resolved peaks plus a background of hot (presumably) pick-up ions. The flow is very close to strict corotation with Jupiter and comes directly into the main sensor with similar signatures in the A, B and C cups with little signal (only from the hot component) into the D cup.

Figure 1: Vertical axis is current (I) in femtoAmps (fA) and horizontal is channel number or energy-per-charge (or voltage, V in Volts).
Fred (64 1016):

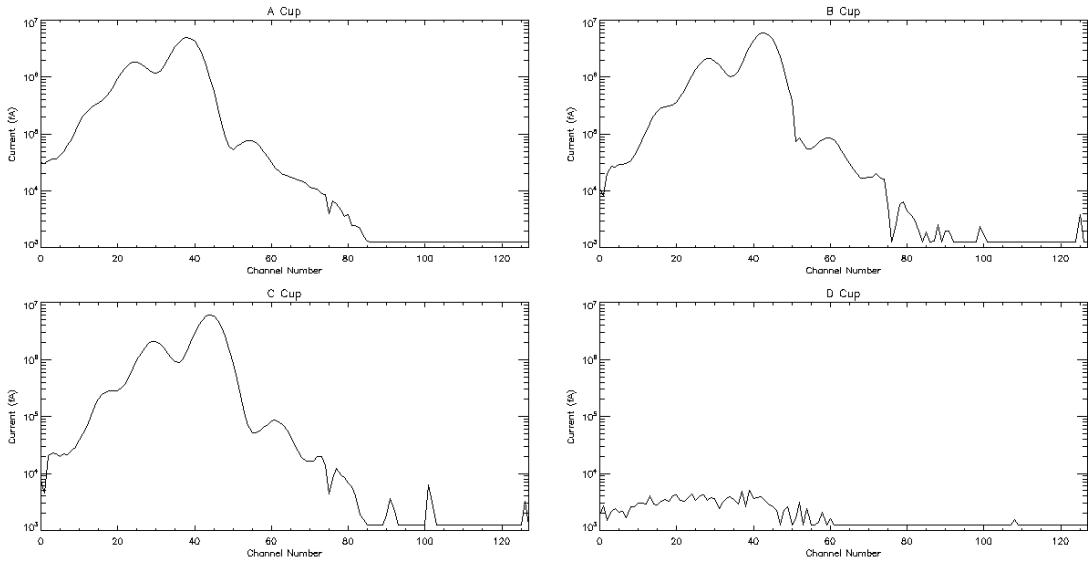

George - Obtained in the plasma sheet at 19.782 RJ, George also has well-resolved peaks for three heavy ions species plus a clear signal of protons. The flow is pretty much in the corotation direction but lagging by about 18%. There is minimal signal in the main sensor with the corotational flow directed into the D cup only.

George (63 1535):

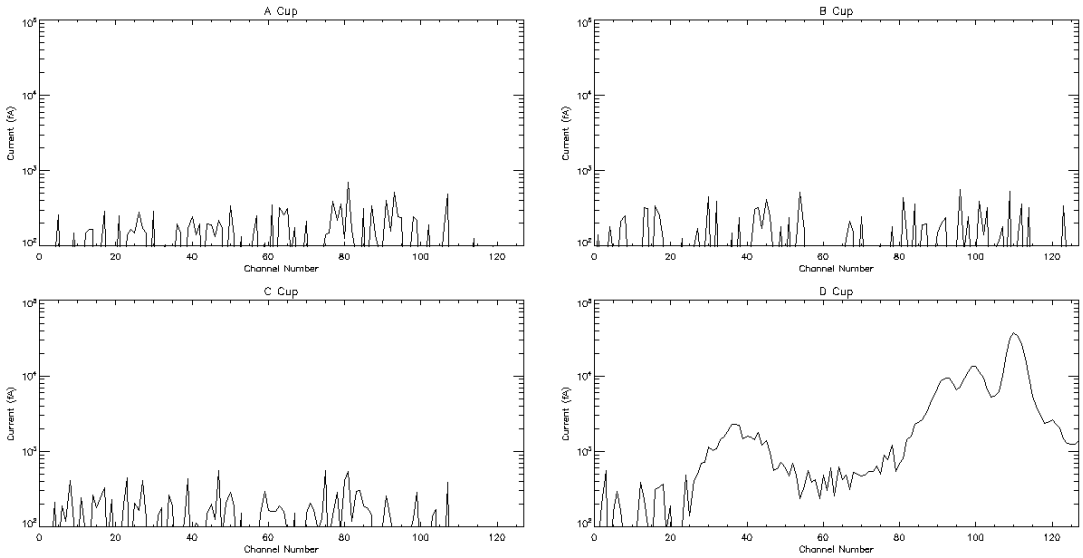

For all purposes in this memo, it will be assumed that currents are measured and displayed in femtoAmps (fA = 10^{-15} Amps). This is the natural range for the measured currents in the jovian system from PLS.

Errors in the PLS Ion Measurements

To get realistic estimates of the errors in the measurements, we first looked at a document from John Belcher:
http://lasp.colorado.edu/home/mop/files/2015/04/VoyagerDoc_2015.pdf
And then dug down into MJSANL to get this – as documented here:
http://lasp.colorado.edu/home/mop/files/2015/04/MeasurementError-raw.pdf

This provided the following estimate of the measurement error:

$$Error^2 = Error^2{}_{background} + Error^2{}_{Measurement}$$

where

$$Error^2{}_{background} = 54$$

$$Error^2{}_{Measurement} = 1.118x10^{-4}I^2$$

In our re-analysis of the Jupiter PLS data at LASP, this error seemed extremely small and did not even appear visible, as is obvious in Figure 1. Because of the range of these currents, a natural logarithmic scale is used. At some point in this scale, there should be a visual error, even if not at all points. So obviously, the ~7 fA background noise (square root of 54) does not do an adequate job of accounting for what is evidently noise in the data.

Figure 2: George and Fred with errors from the original Voyager Memo. Measurement errors are gold vertical bars (yes, hard to see. And ignore the fits).

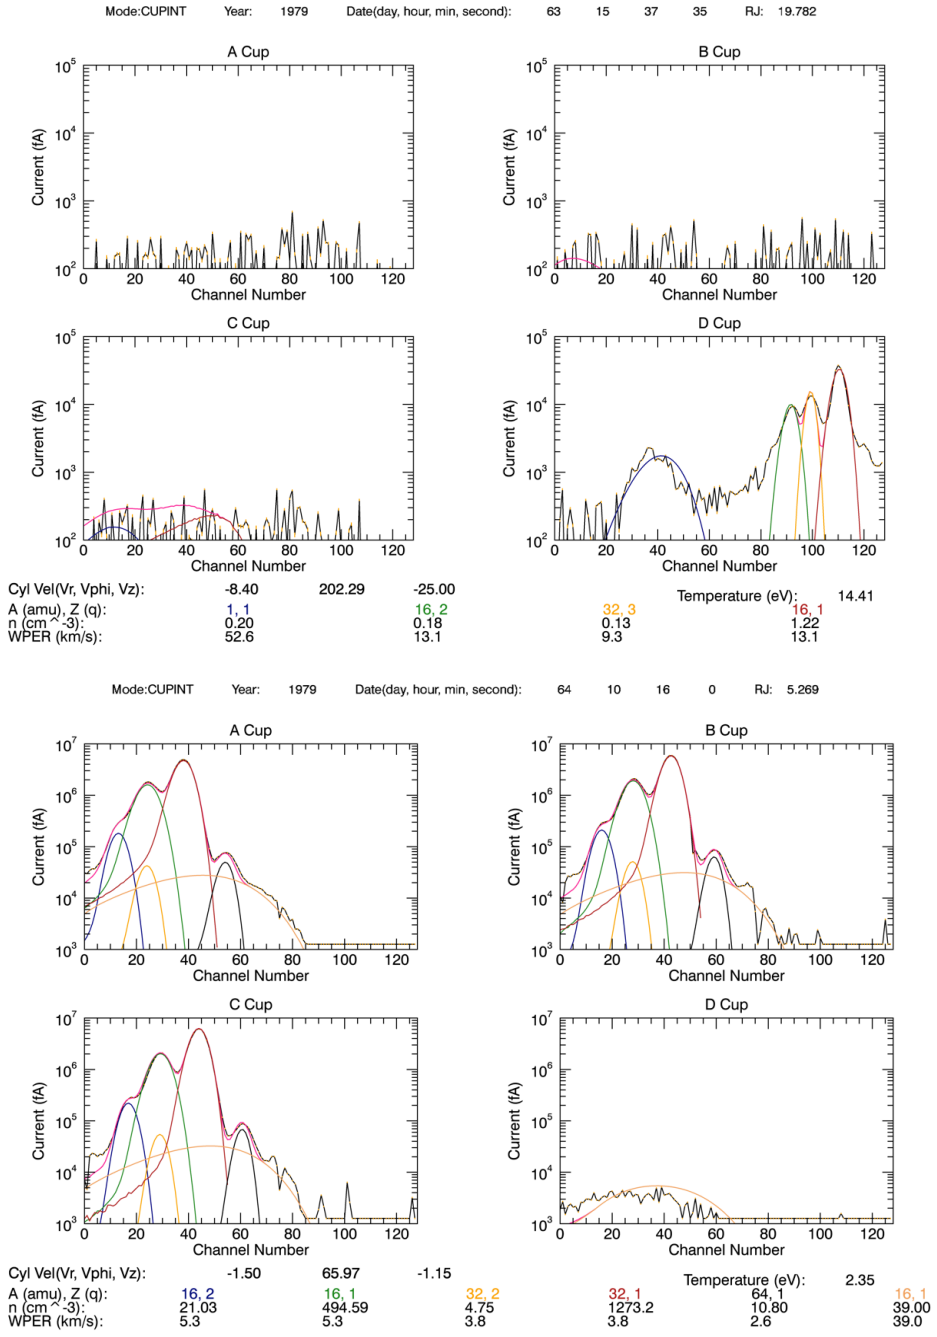

In order to solve this problem – the apparent lack of measurement error using the formal method from previous documentation – we took the residuals from a fit to the current-voltage (IV) shown in Figure 2 and plotted them in Figure 3 below. If the fit is a good fit, then the residuals should help to determine the true background noise. Similarly, we can determine a rough estimate for the background noise by eye. Residuals for the two fit spectra from are of the same order – about 100-1000 fA in the region of the spectrum where there is no clear signal of plasma flux.

Figure 3: Residual from fitting Fred (top) and George (bottom) – in fA.

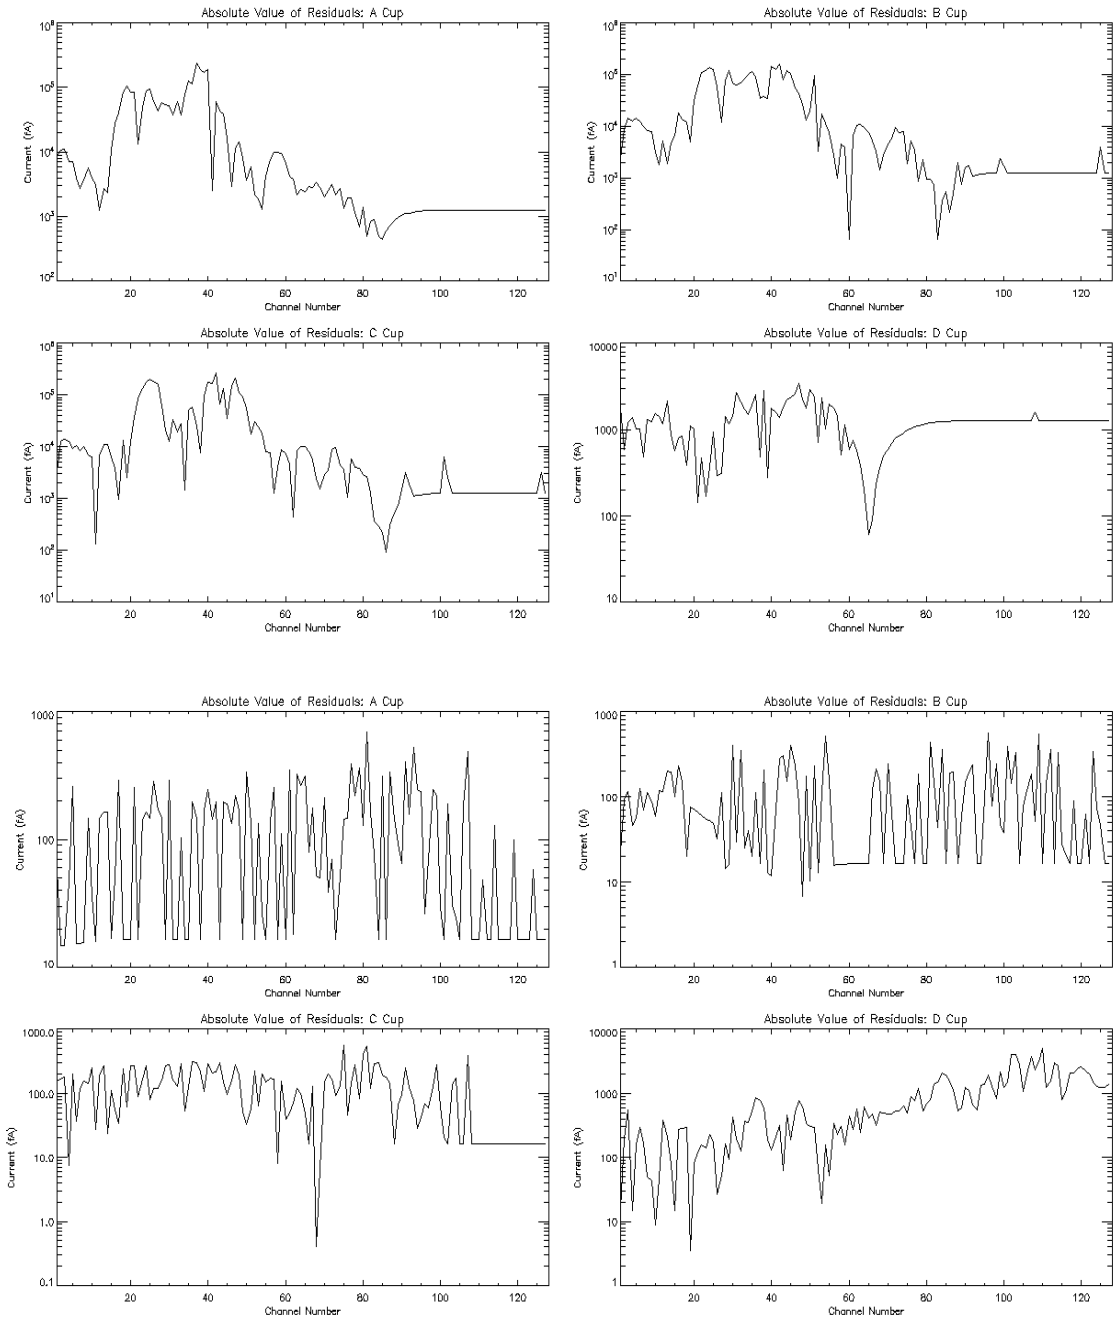

Frank Crary came up with another method – from a couple Soviet-era statisticians Savitzky and Golay. [Actually, they were at Perkin-Elmer Corp. at the time. But they may have been White Russians]

Simple, or box car, smoothing replaces each point with an average of the 2N+1 points. This is, in effect, fitting these points with a line and replacing the mid-point with the fit value. In the process, the shape of the smoothed spectrum is modified (to the extent that the linear fit can not quite match the shape of the spectrum over that interval of 2N+1 points. Peaks are broadened and their amplitude reduced, although the integral, or sum, over the peak is preserved.

Savitzky and Golay (1964) developed an improved form of smoothing. Using a specific weighted average of 2N+1, the midpoint is replaced value which results from fitting those points to a higher order polynomial. To the extent that the 2N+1 points are well-represented by a truncated Taylor series, this does not distort the shape of peaks. The Savitzky-Golay smoothed spectrum has a noise level reduced by $(2N+1)^{1/2}$ but a similar shape to the original spectrum. The difference between the two can be considered a reasonable measure of the noise level.

Figure 4: Savitzky-Golay method of estimating noise level. Plotted are difference between the raw and smoothed spectrum versus the raw values. The blue horizontal line is 300, the black slope is 1%.

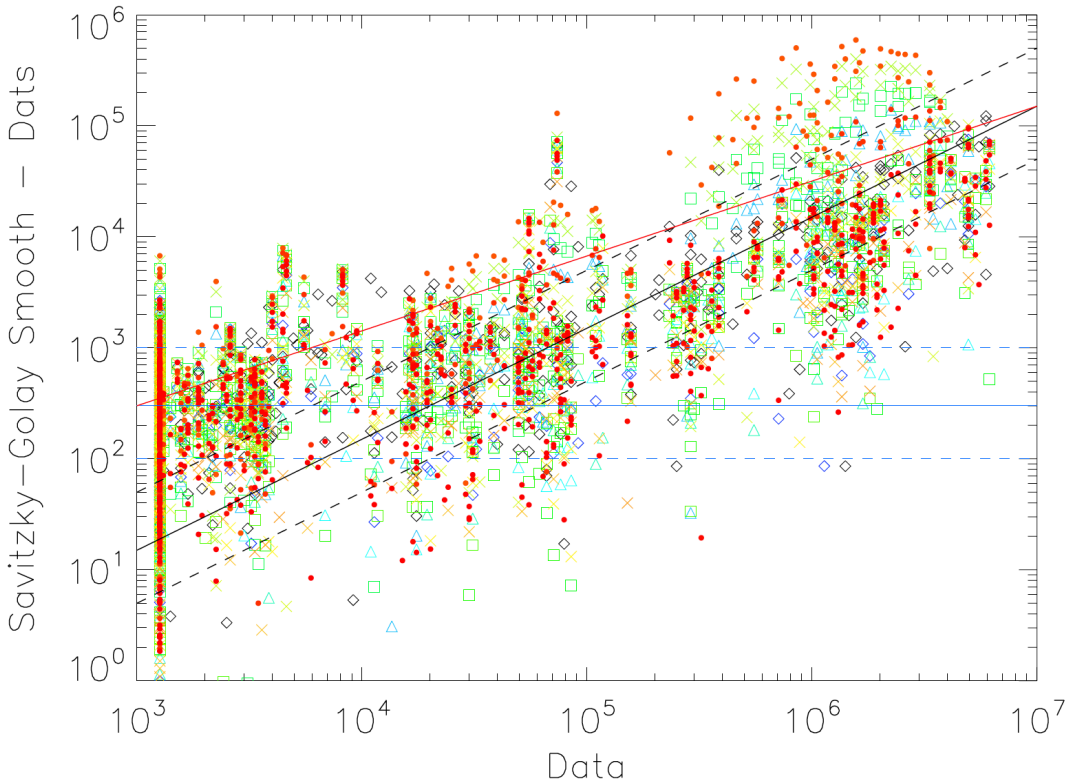

Based on these residuals, we made 2 estimates of the error. The larger estimate took an estimated background noise of 1,000 fA and a 1% error in the measurements. The smaller estimate assumed a background noise of 300 fA and a 0.5% error in the measurements. From these, the two errors below were derived:

$$Case\ A: \quad Error^2 = \sqrt{10^6 + 1.0x10^{-4}I^2}$$

$$Case\ B: \quad Error^2 = \sqrt{10^5 + 2.5x10^{-5}I^2}$$

Examples for spectra Fred and George are shown in Figure 5 for both errors and A and B. From the given spectra it would appear that error method B is a more accurate method for determining the errors in the measured data from the PLS instrument.

For completeness, we also evaluated a third method – Poisson counting statistics. Assuming a worse case scenario - all ions are doubly charged – then we calculate the number of ion impacts on a plate and the Poisson measurement error. Here, t is the accumulation period (in seconds) per channel of the PLS instrument. The accumulation time is typically 0.24 seconds for the Voyager PLS Faraday cup:

$$Case\ C: \quad Error^2 = \sqrt{10^5 + 2.5x10^{-5}I^2 + \frac{0.0179}{t}I}$$

Figure 5: Comparison between Cases A and B for estimating the measurement error with Fred (top) and George (bottom).

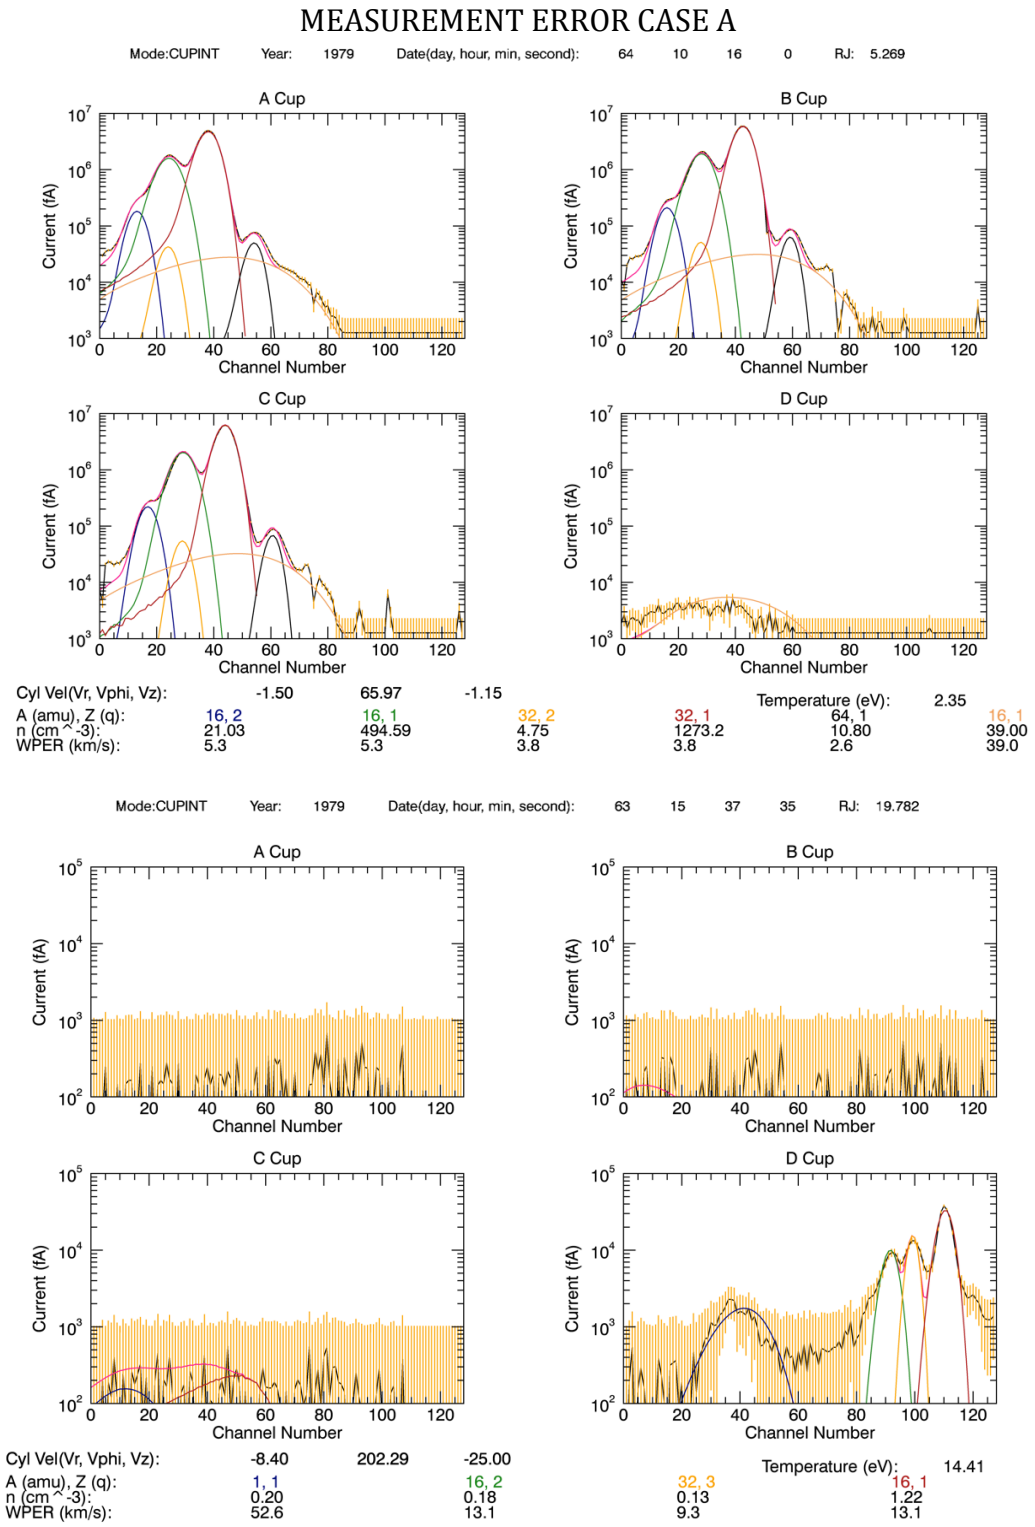

MEASUREMENT ERROR CASE B

Mode:CUPINT Year: 1979 Date(day, hour, min, second): 64 10 16 0 RJ: 5.269

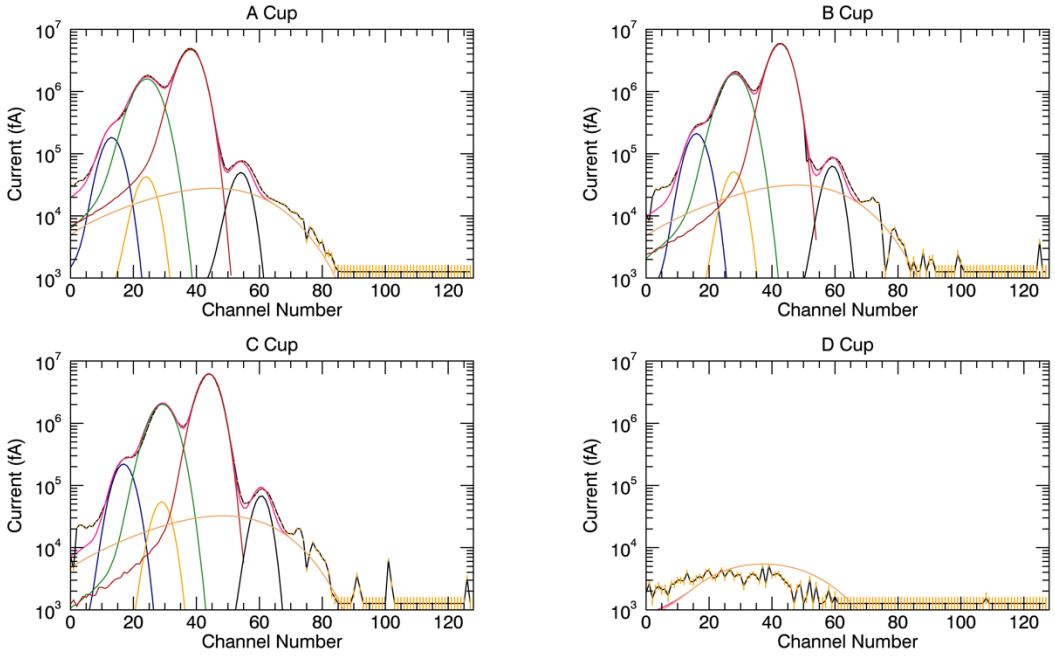

Cyl Vel(Vr, Vphi, Vz):	-1.50	65.97	-1.15			
					Temperature (eV):	2.35
A (amu), Z (q):	16, 2	16, 1	32, 2	32, 1	64, 1	16, 1
n (cm^-3):	21.03	494.59	4.75	1273.2	10.80	39.00
WPER (km/s):	5.3	5.3	3.8	3.8	2.6	39.0

Mode:CUPINT Year: 1979 Date(day, hour, min, second): 63 15 37 35 RJ: 19.782

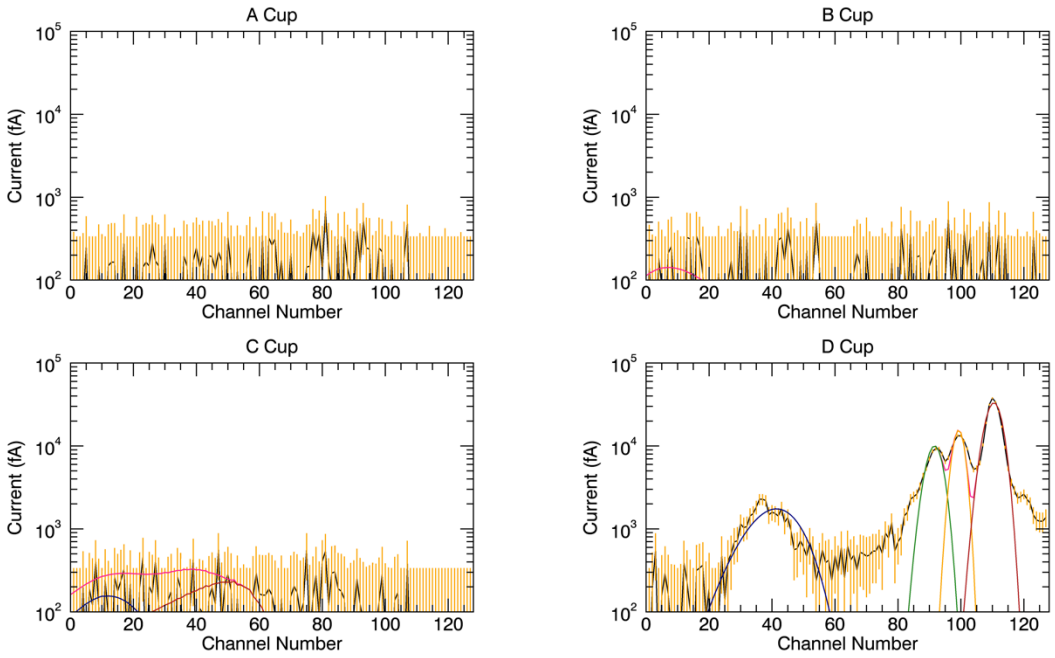

Cyl Vel(Vr, Vphi, Vz):	-8.40	202.29	-25.00			
					Temperature (eV):	14.41
A (amu), Z (q):	1, 1		16, 2	32, 3		16, 1
n (cm^-3):	0.20		0.18	0.13		1.22
WPER (km/s):	52.6		13.1	9.3		13.1

Figure 6 below shows the Case C error to illustrate that there is no visual difference between error B and C due to the addition of a Poisson-like term. Since the effect is minimal, this Poisson error was not used for our analysis of the Jupiter PLS data.

MEAUSUREMENT ERROR CASE C

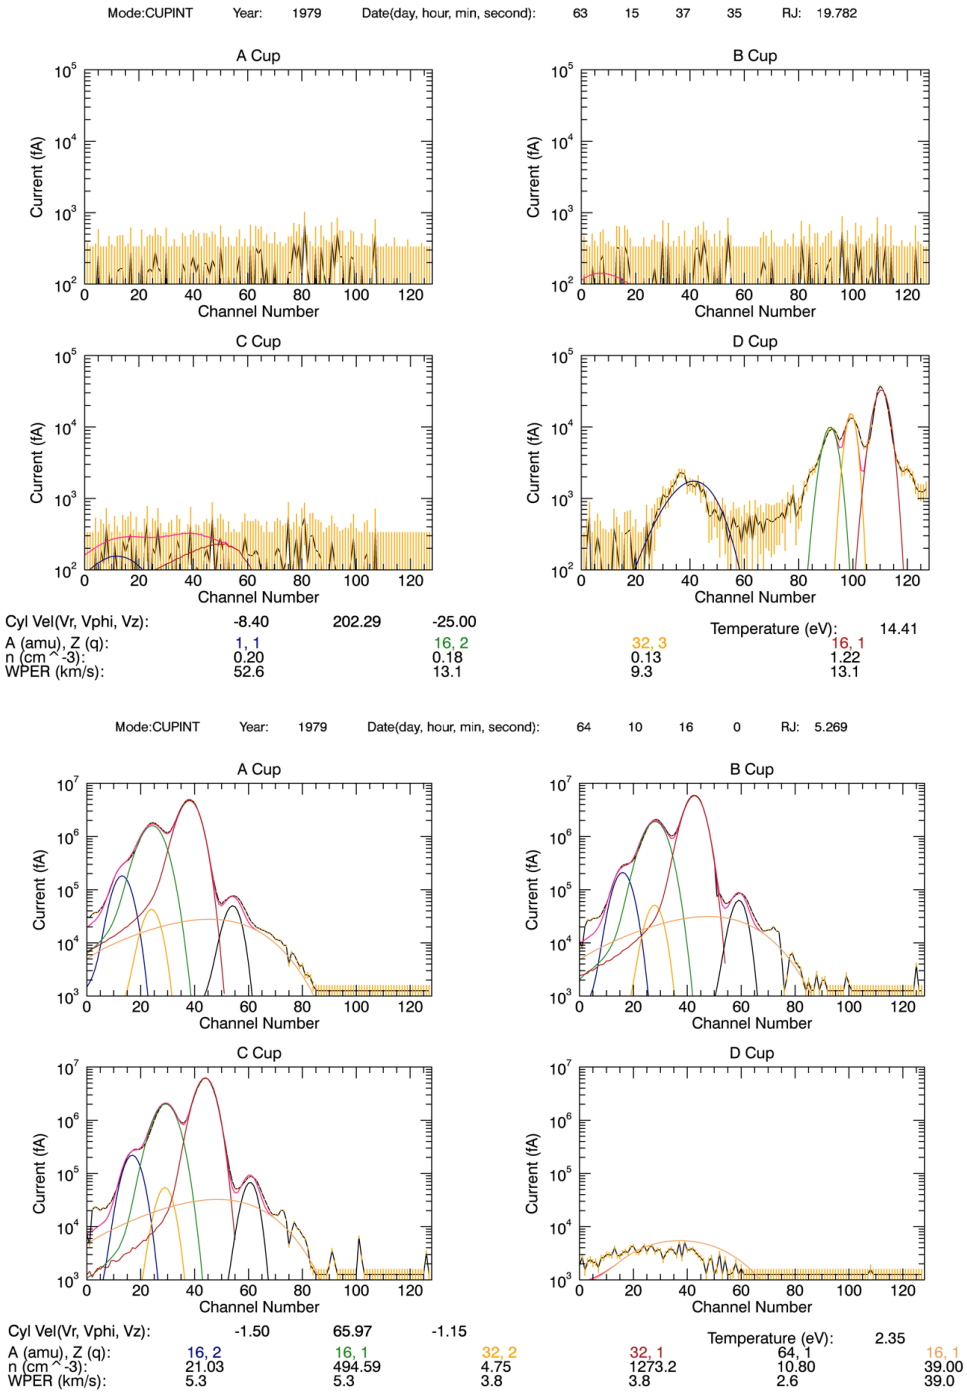

We decided that measurement error Case B looked about right and became the error method used for our analysis at Jupiter. We believe this accounts for the background noise for measurements in the jovian system and may be different for analysis at other planets or in the solar wind.

Uncertainties in the Parameters

In order to fit the PLS ion data, we need an instrument response to an input plasma velocity distribution. We have used Alan Barnett's response function developed for his 1983 (http://lasp.colorado.edu/home/mop/missions/voyager-2/references/). For the fitting process, we used the IDL procedure called MPFIT, written by Craig Markwardt. The above measurement error Case B provides appropriate weighting of the measurements and allows MPFIT to accurately minimize the chi-square value of the best fit parameters for any number of parameters. We have started by using the simplest function: convected isotropic Maxwellians. Usually, the fitting routine fits 3 components of the flow velocity, the temperature of the ion species, and the densities of 3 or 4 species for our model. We have explored letting all species have the same temperature vs. common thermal speed or varying each temperature separately. We have also constrained some of the densities either because limited information in the spectrum or because of the M/Q=16 ambiguity of O^+ and S^{++}.

In our experience, MPFIT does not converge to the true minimum, but gets extremely close. Therefore, after leaving the fitting routine, a check is made that the fit is at the true best fit location, making slight final adjustments to the parameters and getting them into the true best fit location as defined by minimizing the Chi-Squared value.

$$\chi^2 = \sum_{i=1}^{N} \frac{(Data_i - Model_i)^2}{Error_i^2}$$

Note that the value of χ^2 depends on the channels used – restricting the fit process to fewer channels (e.g. where the signals are strong) will decrease the value. We return to this issue later.

Once the true minimum is found then the uncertainties in the parameters can be derived. First, the Hessian matrix is calculated which describes the local curvature of a function of many variables. The size of the matrix is ν by ν, where ν is the number of free parameters. Each element of the Hessian matrix can be described below:

$$H_{i,j} = \frac{\partial^2 f}{\partial x_i \, \partial x_j}$$

From this, it is possible to construct the entire Hessian matrix. In this analysis, we specified the ∂x steps to be 1% of the best-fit parameters. This accommodates the wide range of absolute magnitude of the different parameters. If an absolute step size were used then a moderate change in the temperature in the cold Torus, about 0.02 of 2 eV, would be a very small change in the density of S^+, 0.02 of 1,200 n/cc. Therefore this percentage was used to reflect an accurate curvature for all of the

elements. It was also evident that the Chi-Squared space was noisy at small absolute step levels and the curvature matrix at that point would not reflect the overall curvature of the function itself; this is due to noise in the data and the co-dependence of the parameters in the function.

From the Hessian matrix, the curvature matrix is created, which is the same as dividing the Hessian matrix by a factor of 2. To get the covariance matrix, we take the inverse of the curvature matrix. The square roots of the diagonals of this matrix are the formal 1σ uncertainties in the best-fit parameters. This is a standard procedure of many off-the-shelf fitting routines.

Because this is in Chi-Squared space and not reduced Chi-Squared space, the curvature matrix does not represent the true curvature of the function, but rather a function scaled by the degrees of freedom (DOF). Therefore, from the channels that are fit in the fitting routine we can define the DOF, where N is the number of channels being fit and ν is the number of free parameters.

$$DOF = N - \nu$$

Since the code returns the un-scaled 1σ uncertainties from the covariance matrix, it just requires a scaling factor of the square root of the DOF to account for it. The code for calculating the 1σ uncertainties comes from Rob Wilson's 2015 paper. The only modification to the code is that the Δ in calculating the curvature matrix is not an absolute value but rather a relative value that depends on value of the best-fit parameter.

$$\sigma = 1\sigma_{un-scaled} * \sqrt{DOF}$$

$$parameters = \ parameter_{best} \pm \sigma$$

The best fit parameters and corresponding 1-sigma uncertainties from the above equation are returned directly from the fitting routine – and will be published in survey papers of the Voyager PLS data at Jupiter.

Note that MPFIT – and most standard fitting routines – assumes that each parameter is independent. In reality, solutions are not unique and similarly good fits could be obtained by varying different parameters. To explore the degree of correlation of the parameters we need to explore reduced chi-squared space.

Reduced Chi-Squared Space

In order to test whether or not the fit parameters are the best, and to tell how they are dependent on each other, it is important to look at the reduced chi-squared

space. This was done by taking the best-fit parameters and using the formal 1σ errors calculated above to create an array of reduced chi-square values. For this array, values of the reduced chi-squared are calculated within ± 3σ of the best-fit parameter.

$$\chi_R^2 = \frac{1}{N - v}\sum_{i=1}^{N} \frac{(Data_i - Model_i)^2}{Error_i^2}$$

N is the number of data points that the reduced chi-squared is being calculated for. For the formal analysis of these error points, data was selected carefully where the model matched the data. At points where the model goes to zero and there is just noise in the instrument, the chi-squared value was not calculated. The error value in the measurements that was used was the Case B measurement error described above. These values are calculated for variations in the parameter space (for just the parameters that are fit to the data) in a high resolution grid search between ± 3σ.

Note: The IDL code returns a CSV with the best-fit parameters. Make sure to use this unedited version of the CSV file to generate the chi-squared contour plots. If edits are made to this CSV file it appears that IDL encounters rounding errors and will displace the true minimum of the fit. Therefore if no changes are made to the actual CSV file for the parameters of interest, the chi-contours should produce the correct plots.

To illustrate, here is the following case with 2 best fit parameters. If the fit were to return the following values for T_i and V_φ:

$T_i = 4.3 \pm 0.3$ eV
$V_\varphi = 60 \pm 1$ km/s

The following grid represents the values at the corners of the plot:

$T_i = 3.4, V_\varphi = 57$	$T_i = 4.3, V_\varphi = 57$	$T_i = 5.2, V_\varphi = 57$
$T_i = 3.4, V_\varphi = 60$	$T_i = 4.3, V_\varphi = 60$	$T_i = 5.2, V_\varphi = 60$
$T_i = 3.4, V_\varphi = 63$	$T_i = 4.3, V_\varphi = 63$	$T_i = 5.2, V_\varphi = 63$

In between these values, many additional values are calculated to create a smooth contour plot of the reduced-chi squared values. For the plot, to put all of the parameters around the same value, we plot the delta of the reduced-chi squared value. Therefore, the actual value plotted is the reduced chi-squared value minus the minimum value of the reduced chi-squared value. Therefore, the minimum of every plot is by definition 0.

In addition, on the plots the 1σ uncertainties from MPFIT for each parameter are plotted in white error bars. To illustrate that the shape of the reduced chi-squared function, in reality, is not perfectly quadratic (in 2-D), the values of 1, 2, and $3\sigma_\Delta$ are

plotted for appropriate delta values. Assuming a model with one independent parameter, these values would respectively be 1, 4, and 9. For multiple parameters, Table 1 shows the appropriate values to use. This does not assume that the chi-squared values near the minimum form a perfect parabolic shape. On all of the contour plots, these values are plotted in red dashed lines. The equation used to solve for these contour levels of delta is:

$$\Delta = 2 * Inverse\ of\ incomplete\ gamma(p, \frac{\upsilon}{2})$$

Here, p is the confidence level, 0.6827 for $1\sigma_\Delta$, 0.954 for $2\sigma_\Delta$, and 0.997 for $3\sigma_\Delta$. These represent the 68%, 95%, and 99.7% confidence levels that the true value is with that value of σ_Δ. The number of parameters is denoted by the Greek letter υ.

Table 1: Delta values for different number of free parameters. Up to 10 free parameters are included in the table.

No. parameters	$1\sigma_\Delta$	$2\sigma_\Delta$	$3\sigma_\Delta$
1	1	4	9
2	2.29575	6.18007	11.8292
3	3.52674	8.02488	14.1564
4	4.71947	9.71563	16.2513
5	5.8876	11.3139	18.2053
6	7.0384	12.8488	20.0621
7	8.17624	14.3371	21.8466
8	9.30391	15.7891	23.5746
9	10.4234	17.2118	25.2569
10	11.536	18.6103	26.9011

Example contour plots for our 2 favorite spectra are included in Figure 7. All of these plots were created in the Matlab (better than IDL for plotting contours). All data was created in IDL.

If the two parameters (corresponding to the x and y axes) are independent of each other then the contours should be round and parabolic in curvature. If the calculation of $1\sigma_\Delta$ matches the MPFIT 1σ value of uncertainty then (a) the parameters are independent and (b) we have picked appropriate number of channels in our estimate of χ^2. We could adjust the channels to get a closer match between these 2 different methods of estimating sigma – but the point here is mainly to illustrate the degree of independence vs. correlation between pairs of parameters – and to show the different methods agree reasonably well.

Figure 7: Contours of Δ=Chi-Square-Minimum for spectra George and for Fred under two different fitting conditions.

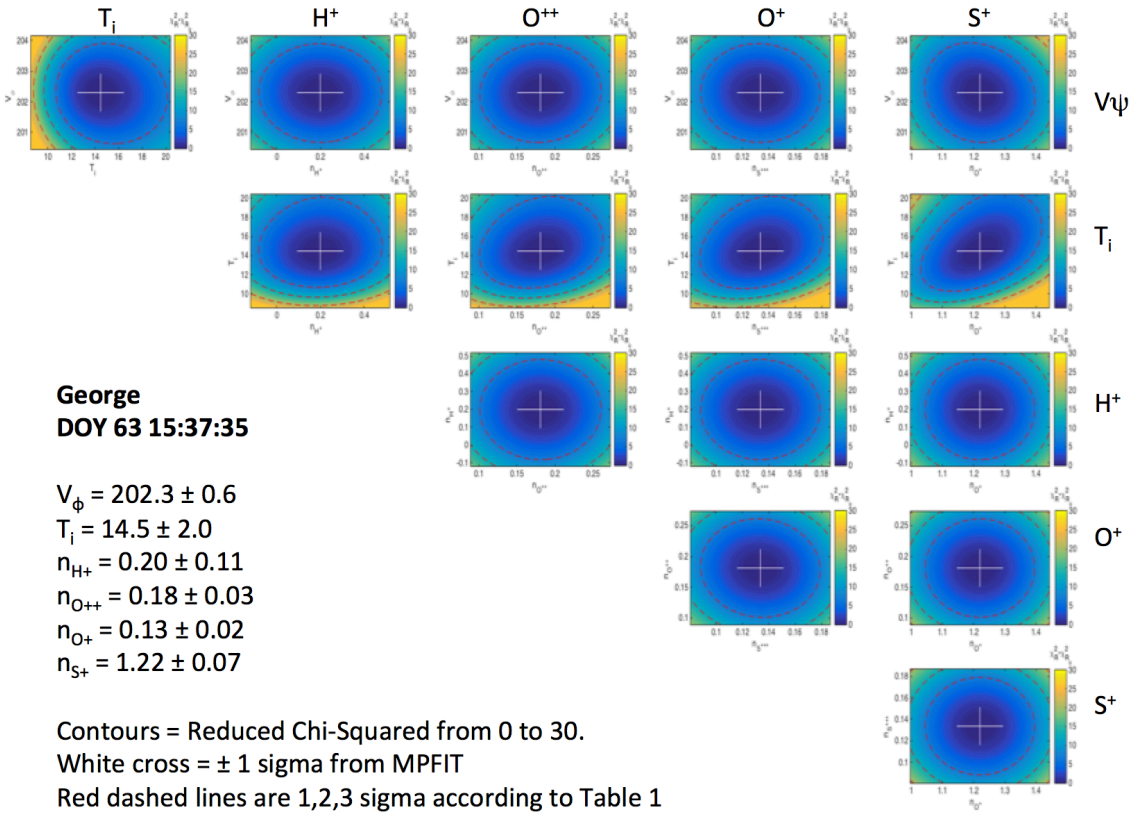

Note above that T_i seems to be asymmetric (chi-square values increasing more steeply towards lower temperatures rather than higher) and have the strongest correlation – particularly with S^+ density. Otherwise the contours are pretty circular, suggesting the parameters are fairly independent. The fact that the $1\sigma_\Delta$ contours are larger than the white cross from MPFIT suggests that we could have scaled the Δ values adjusting the number of channels in the calculation.

With significant signal only in the D-cup for George, we are only able to fit one component of the flow velocity – the azimuthal flow.

When we turning to Fred in the cold torus, we have signal in the 3 main sensors and can make a reasonable estimate of the radial and vertical flow components as well as azimuthal. As the second plot shows below, the flow is close to corotational.

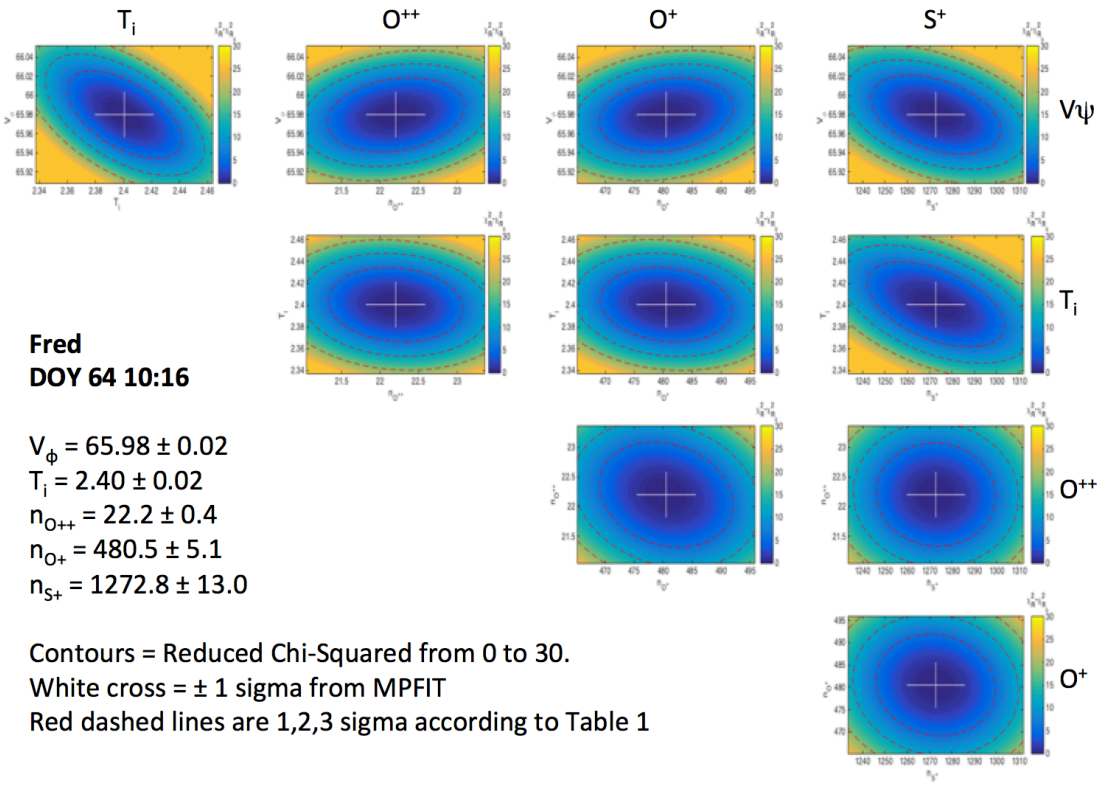

Fred
DOY 64 10:16

V_ϕ = 65.98 ± 0.02
T_i = 2.40 ± 0.02
n_{O++} = 22.2 ± 0.4
n_{O+} = 480.5 ± 5.1
n_{S+} = 1272.8 ± 13.0

Contours = Reduced Chi-Squared from 0 to 30.
White cross = ± 1 sigma from MPFIT
Red dashed lines are 1,2,3 sigma according to Table 1

Ti shows correlations with the other parameters while V seems anti-correlated.

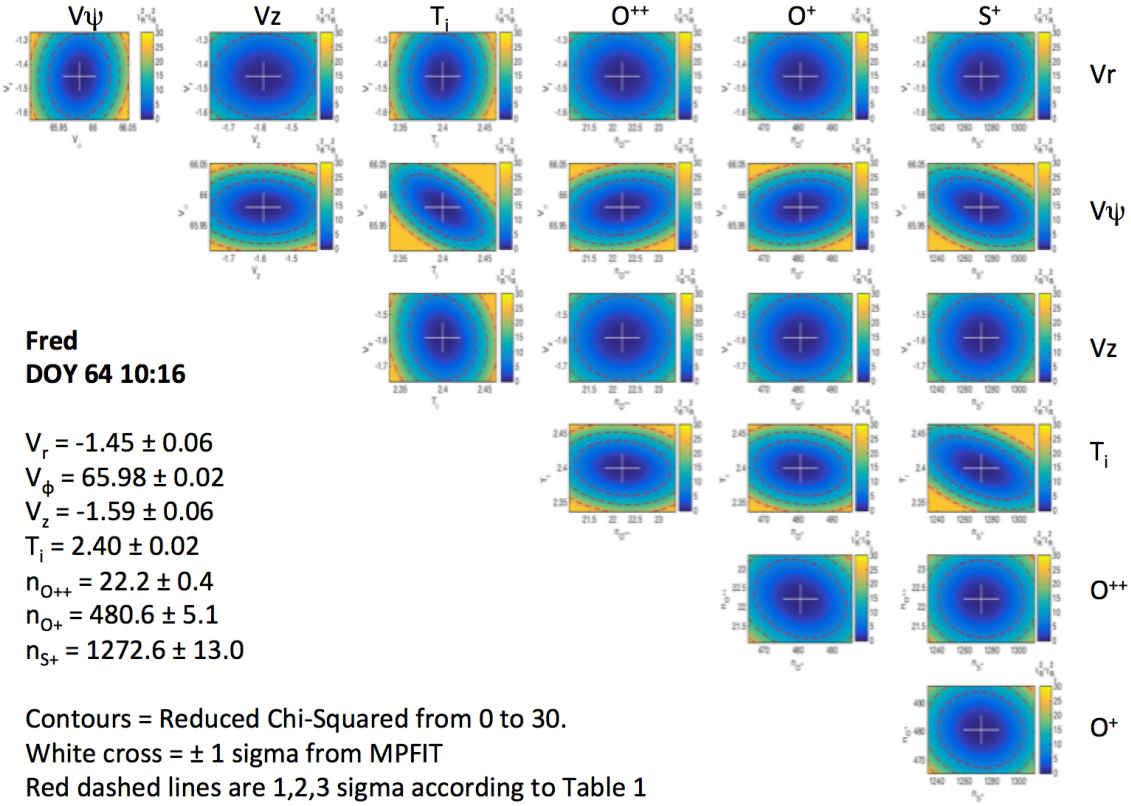

Fred
DOY 64 10:16

V_r = -1.45 ± 0.06
V_ϕ = 65.98 ± 0.02
V_z = -1.59 ± 0.06
T_i = 2.40 ± 0.02
n_{O++} = 22.2 ± 0.4
n_{O+} = 480.6 ± 5.1
n_{S+} = 1272.6 ± 13.0

Contours = Reduced Chi-Squared from 0 to 30.
White cross = ± 1 sigma from MPFIT
Red dashed lines are 1,2,3 sigma according to Table 1

With Vz and Vr we see relatively large uncertainties with no strong correlations.

In this examination of fitting procedures and calculation of uncertainties we are not so concerned about the absolute values of the output parameters. For example, we have fit the M/Q=16 peak as purely O^+ when we know that out in the plasma sheet (where the George spectra were obtained) there is about equal amounts of S^{++}. We know that we will have to specify the O^+/ S^{++} ratio (based on spectroscopic data and/or physical chemistry models) when we come to fit the rest of the PLS ion data.

References

Savitzky and Golay (1964), *Anal. Chem., 36,* 1627–1639

Wilson, R. J. (2015), Error analysis for numerical estimates of space plasma parameters, *Earth and Space Science, 2,* 201–222. doi:10.1002/2014EA000090.