A Short Introduction to the DMSP SSIES-3 Quality Flags and How to Use Them

A hypothesis or theory is clear, decisive, and positive, but it is believed by no one except the man who created it.
Experimental findings, on the other hand, are messy, inexact things, which are believed by everyone except the man who did that work.
— Harlow Shapley (US astronomer, *Through Rugged Ways to the Stars,* 1969)

The single biggest mistake a researcher can make in using a dataset they did not create is to assume that the values in these data represent a simple and precise measurement of the physical parameter being sampled. They assume any measurement must be straightforward, like a thermometer measuring the temperature of a bowl of water. Just put the thermometer in the water, wait until it reaches equilibrium, then read the result. If they repeat this (and the water does not change) they will get exactly the same result. Simple.

The SSIES-3 plasma measurements are neither that simple nor straightforward.

For example, the ion temperature, the ram drift velocity of the ions, the fraction of the plasma that is composed of H+, the fraction that is He+, the fraction that is O+, and the electrostatic potential difference between the instrument versus the plasma ground are all calculated from the Retarding Potential Analyzer (RPA) instrument. The values of these six parameters are based on a Gaussian fit to the curve of the measured ion currents as a function of the retarding potential applied to the instrument (referred to here as the "I-V curve"). We are doing an analysis that searches through a six-dimensional phase space to find the true local minimum that simultaneously gives the most accurate values of all six parameters. Since there are a lot of local minima in this six-dimensional phase space, this is not a trivial task. Despite our best efforts we sometimes end up settling on the wrong minimum. Or we find we can get reliable results for some of the parameters, but not the rest. Or the original current-voltage curve is so noisy that we cannot settle on a single unique solution; this results in multiple solutions, all appearing equally valid, but giving significantly different values for one or more of the parameters. And sometimes the current-voltage measurements are too noisy to produce any valid solutions. Given enough time and effort we could take almost any single reasonable current-voltage curve and keep searching until we found an acceptable result that is at least equal or somewhat better than the result we are publishing here. But since one DMSP SSIES-3 RPA generates over 30 million current-voltage curves in a year, and that currently there are over 40 satellite-years of SSIES-3 data available, we will never have that much time and effort.

This is why we include quality flags in these data sets; we want to give the users an initial guide to which data are of reasonable quality and which data are bad and should not be considered. ***But it is important that the user understand this*: *These data and quality flags are only a level-2 quality. It is ultimately the end user's responsibility to double-check that these data are***

reasonable before they use them in their own research. Blindly dumping a year's worth of data into your analysis program without first checking it will bring you grief.

Plot of a single orbit

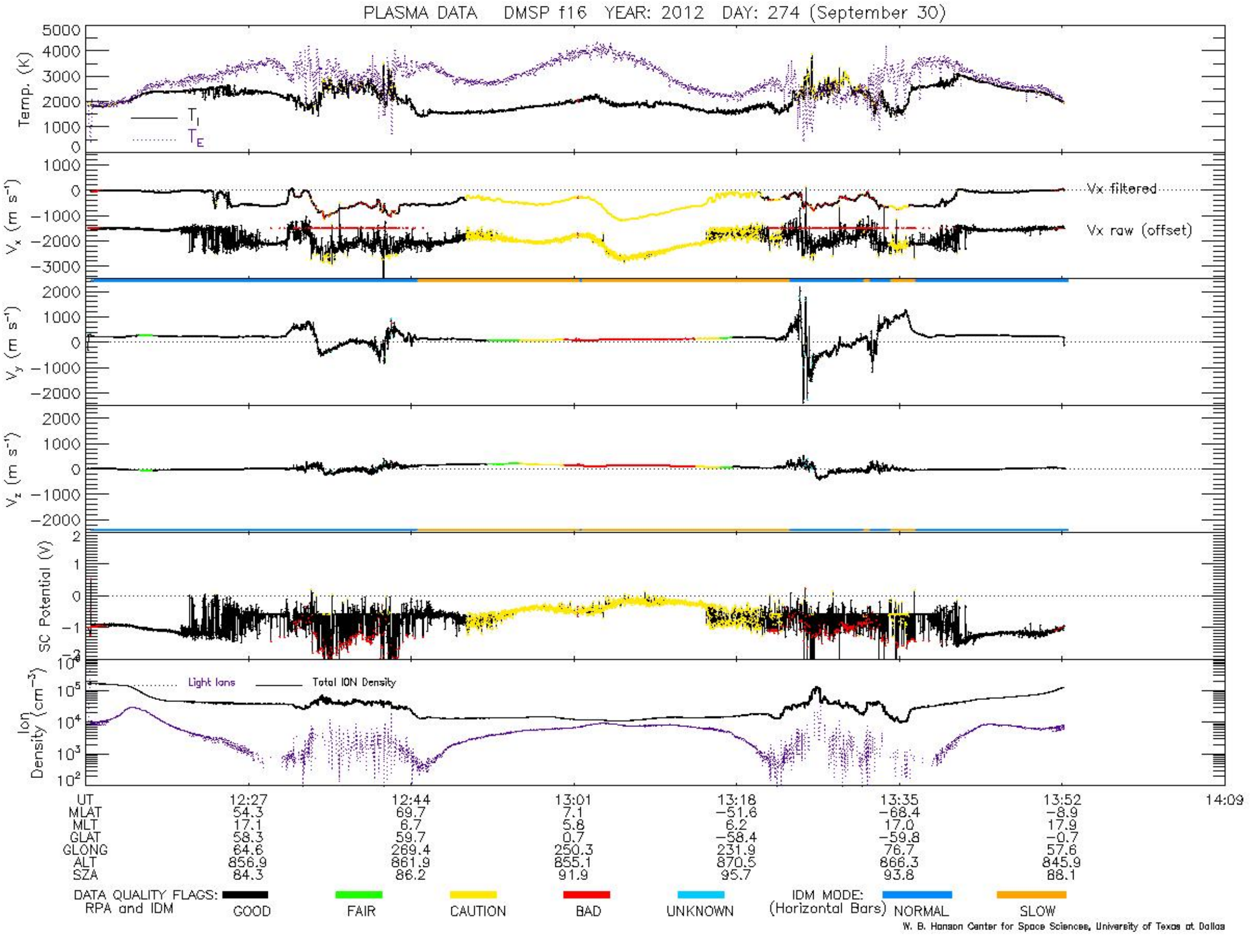

Figure 1

Figure 1 shows a typical plot of F16 data from a single orbit starting at 12:10 UT on 30 September 2012. The orbit starts at the northbound geographic equatorial crossing on the duskside. The spacecraft goes over the northern hemisphere midlatitude on the duskside, over the north polar region (left side of plot), heads southbound on the dawnside recrossing the equator at about halfway across the plot, then over the southern polar region (right side of plot), and finally going northward on the southern hemisphere duskside to end back at the equatorial crossing.

The top panel shows the electron temperature (Te) plotted in purple and the ion temperature (Ti). The quality of the Ti data are shown in the color of the plot, and for this pass all the Ti data are good and plotted as black. The second panel shows the ram flow data (Vx) in two formats. The top line is the Vx after it has been run through a median filter and smoothed and all the Vx = 0 points have been removed. The data are color coded based on the quality flag of the Vx data as shown by the color code at the bottom of the plot. The lower line shows the raw Vx data (offset by -1500 m/s for clarity here) before the median filtering and smoothing have been applied. The red horizontal lines are the Vx = 0 points in the raw data that are flagged as 4 (bad, do not use). The third and fourth panel show the two crosstrack ion flow data from the IDM with the horizontal crosstrack flow (Vy) in the third panel and the vertical crosstrack flow (Vz) in the fourth panel. As with the Vx data the Vy and Vz data are color coded based on the quality of these data. Furthermore, the top edge of the Vy panel and bottom edge of the Vz panel are color coded to show whether the IDM was operating in normal mode (blue) or slow mode (orange). The fifth panel shows the calculated potential difference between the plasma ground and the sensor plane of the SSIES-3 instruments. The color-coding of the data here uses the Vx quality flag. Finally, the bottom panel shows the total ion density from the scintillation meter (black) along with the light ion (H+ and He+) density (purple). There is no quality flag for the scintillation meter density so the line is always black. Although there are quality flags for fracH and fracHe values, we do not color code the light ion density plot.

Here are the quality flags for various parameters followed by examples.

Retarding Potential Analyzer

For the retarding potential analyzer (RPA) we have separate quality flags for the overall RPA fit (RPAqual), the ram drift ion velocity (Vx), the ion temperature (Ti), the RPA measured ion density (RPA-Ni), and the three fractional compositions (fracO+, fracH+, and fracHe+). All of these flags can have different settings for the same one-second set of data.

The RPA flags for each of these parameters (and color code on the plots) are:

1 = data are **good** and can be used with high confidence **(black)**
2 = data are **fair** and can probably be used with confidence **(green)**
3 = data should only be used with **caution (yellow)**
4 = data are **bad** and should not be used **(red)**
5 = it is **uncertain** what the quality of the data are **(blue)**

To start with the easy cases, if the ion density (measured by the scintillation meter) is less the 10^2 ions/cc then there are not enough ions for the RPA to make a valid measurement, so for these cases all the flags are set to 4 (bad) and the flag for RPA-Ni (densqual) is set to 3 (caution).

The ion density of the topside ionosphere at 850 km never goes above the upper levels of 10^6 ions/cc so any measurement of densities is greater than 10^7 ions/cc is an instrument error from the scintillation meter and we see these occasionally. For these cases all the flags are set to 4 (bad) and the flag for RPA-Ni (densqual) is set to 3 (caution).

The RPA only works properly when its ground potential is near the ambient plasma potential so the ions' motions are not affected by the spacecraft's potential (generally around -15 to -20 V relative to the plasma). The senpot circuit on SSIES holds all the instruments (RPA, IDM, scintillation meter, and Langmuir probe) at the same potential, generally between -2 V and 0 V relative to the ambient plasma potential. This potential difference between the plasma and the SSIES-3 instruments (referred to as "SC potential" on the plot) is shown in the fourth panel. If the potential difference is more negative than -2 V or greater than +0.5 V then the results from the RPA are problematic. For these cases all the flags are set to 3 (caution).

The curve fitting algorithm attempts to calculate the fractional amounts of the H+, He+, and O+ ions in the plasma. However, a somewhat common error is where the analysis results in no H+ but there is a fractional amount of He+. This is generally an unphysical condition so in this case we set the quality flags of all three components to 3 (caution).

If the fitting of the curve results in an ion temperature (Ti) of less than 500 K or greater than 10,000 K, we consider that result to be erroneous so the quality flag for Ti is set to 3 (caution) and the overall RPA quality flag (RPAqual) is set to 3 (caution).

In the ionosphere the electron temperature (Te) is generally hotter than the ion temperature (Ti) by a few hundred to 1000 K, so we assume that any time the calculated Ti exceeds Te then the Ti calculation must be in error. In those cases we set the Ti flag at 3 (caution) and the overall RPA quality flag (RPAqual) is set to 3 (caution).

While the Vx is generally centered around zero, and we assume that Vx is close to zero outside of the polar regions, our curve fitting routine almost never results in a Vx value of *exactly zero*. However, there are times where we cannot produce a valid fit for all the parameters. In some of these cases the problem is we only have a single component in the plasma and we have too many unknowns for solve for valid solutions for all the parameters. For these cases we reduce the number of unknowns by setting Vx ***exactly to*** zero and then solving for the rest of the parameters. In general we have found that the solution for Ti is not sensitive to the variations in Vx, so setting Vx to zero still allows us to obtain reasonable Ti values. Thus for cases where Vx is identically zero the quality flag for Vx is set to 4 (bad, do not use) as this zero value is an assumption and not a measured value.

In principle the SSIES-3 (both RPA and IDM) can measure large ion flows up to 3500 km/s. In reality we almost never see valid values over 2000 m/s and values higher than that we view with caution. For Vx we set the limit at 3000 m/s so any values of Vx greater than 3000 or less than -3000 m/s have the Vx flag set to 4 (bad, do not use). Because most of the DMSP orbits are roughly dawn-dusk we tend to see the largest flows in the Vy component. Thus we tend to view any Vx flows greater then 1000 m/s or less than -1000 m/s with caution. Thus with any Vx with an absolute value between 1000 and 3000 m/s we set the Vx flag to 3 (caution).

The RPA-derived parameters are based on the fit of a Gaussian distribution to the measured currents seen as a function of the retarding potential, or the "I-V curve". Obviously the better the fit to the data, the more trustworthy the results. We measure the "goodness of fit" to the curve by calculating the root-mean-square (RMS) of the fit with smaller values of the RMS indicating better fits. If the RMS is greater than 0.12 then the Vx flag is set to 2 (fair) unless it has already been set to 3 or 4 by one of the criteria above, and the overall RPA flag (RPAqual) is set to 2 (fair). If the RMS is less than zero an error in the analysis routine has occurred and the Vx quality flag and the overall RPA flag (RPAqual) are both set to 4 (bad, do not use).

As stated above the RPA analysis routine has difficulty solving for Vx when there is only a single species of ion in the plasma. In the mid-latitude and equatorial regions we expect Vx to be near zero, but we frequently see large (say |Vx| > 500 m/s) flows in Vx on the dawnside of the orbit where the plasma is nearly 100% H+. See Figure 2 below where Vx is close to zero after leaving the subauroral region in northern hemisphere on the dawnside, but partway towards the equator Vx jumps to about -900 m/s and these unrealistic flows continue to the equator. To prevent the use of such erroneous data we flag as 3 (caution) any Vx data between +40° and -40° geographic latitude where the O+ density (calculated from fracO times the scintillation meter density) is less than 10^4 ion/cc.

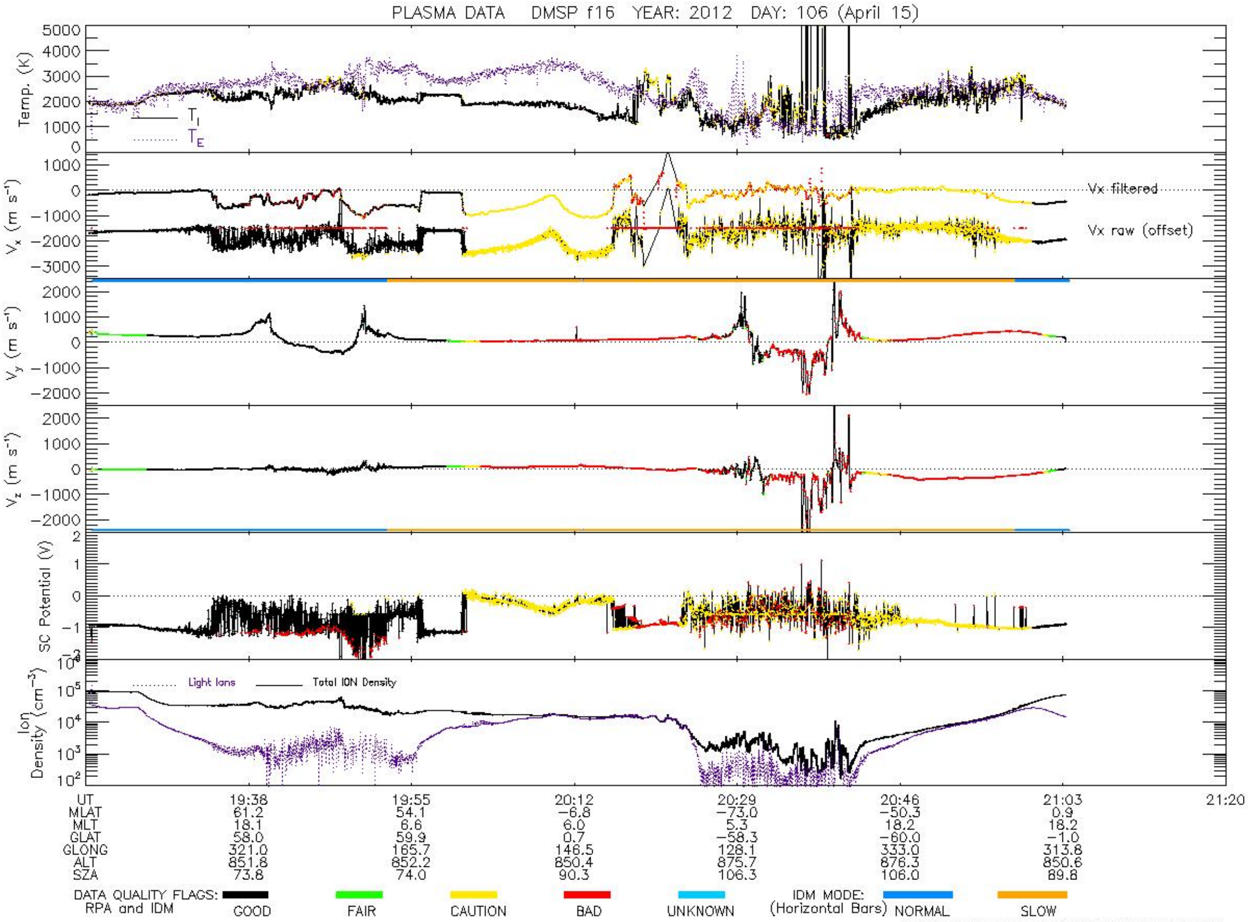

W. B. Hanson Center for Space Sciences, University of Texas at Dallas

Figure 2

If the quality flags for the overall RPA (rpaqual), Vx (vxqual), the ion temperature (tiqual), and RPA density (densqual) have not yet been set by one of the previous filters listed above, then we assume these surviving data are good and these four RPA quality flags are set to 1 (good).

Next are the quality flags for the fractional H+, fractional He+, and fractional O+. For each component, if the fraction is negative (which can happen in the analysis or the analysis failed and a fill data value of -99999 was set), then the fraction is invalid and the quality flag for that component is set to 4 (bad, don't use). For each component if the fraction is between 0.00 and 0.05, then the quality flag for that component is set to 2 (fair). If the fraction is greater than 1.05 (which can happen in the analysis) then the quality flag is set to 3 (caution). All the cases that are left therefore have a fraction between 0.05 and 1.05, and the quality flag for that component is set to 1 (good). We allow the fraction to be rated good even if it is between 1.0 and 1.05 because the error for the fractional composition is estimated to be about 5%, so these are most likely cases where the composition of that ion is essentially 100%. We set the quality flag to caution (rather than bad) for cases where the fractional composition is over 1.05 because the results for the other parameters (Ti, Vx, etc.) can be valid even if the analysis produces these unphysical fractional composition values. We leave it to the user to decide whether to use these composition data or not.

Last, for any quality flag that reaches the end without having been changed (i.e. it fell into none of the categories listed above) then that flag is set to 5 indicating that we are uncertain about the quality of these data. In principle this condition should never be reached in the code, but this is set up as the last resort backup.

Ion Density Meter

For the ion drift meter (IDM) we have separate quality flags for the horizontal crosstrack ion velocity (Vy), the vertical crosstrack ion velocity (Vz), and the IDM itself, but all of these will always have the same value for a given one-second set of data. The main factors that affect the quality of the IDM data are the total ion density and the percentage of O+ ions. In general, high density (> 3 x 10^4 ion/cc) plasma with a majority (> 75%) O+ produce good quality results. The details will be explained below.

The IDM flags (and color code on the plots) are:

1 = data are **good** and can be used with high confidence **(black)**
2 = data are **fair** and can probably be used confidence **(green)**
3 = data should only be used with **caution (yellow)**
4 = data are **bad** and should not be used **(red)**
5 = it is **uncertain** what the quality of the data are and it is the end user's responsibility to decide whether to use these data or not **(blue)**
The following flags are for F17 only
6 = data are judged to be from **good** conditions but some caution should still exercised by the user because of baseline issues **(black) (F17 only)**
7 = data are judged to be from **fair** conditions but some caution should still exercised by the user because of baseline issues **(green) (F17 only)**

8 = data are judged to be from conditions that warrant **caution**, and since these are F17 data the user should be doubly cautious in using them **(yellow) (F17 only)**
(The special flags for the IDM for F17 will be explained later in this section.)

The IDM for SSIES-3 differs from the earlier SSIES IDMs in a major way. The earlier IDMs sampled each component six times a second and since the public data were synced to the four-second cycle of the earlier SSIES RPAs, the published Vy and Vz data were nominally a four-second average of 24 samples. However, under low density conditions there was some indication of electronic "ringing" in the IDM as it alternated sampling between the two components twelve times a second. To correct for this the SSIES-3 IDM operates in two modes. If the total ion density is above 3.0 x 10^4 ions/cc then it samples each component six times per second just as the earlier IDMs did and this is referred to as the "normal mode". Since the RPA cycle for SSIES-3 is one second, then the published Vy and Vz data in these files are at a one-second cadence and are the averages of these six samples of each component.

If the total ion density is below 2.1 x 10^4 ions/cc then the IDM switches to what is referred to as "slow mode". Here it samples a single component six times during the one-second period, but only reports the final sample in the telemetry. The idea here is to allow the electrometers in the IDM to settle down and stop "ringing" before the final sample is taken. In the next second the IDM repeats this process for the other component. This results in the Vy and Vz being sampled only once every other second using only a single sample for each. Since the SSIES-3 data set here has a one-second cadence, the Vy and Vz values are repeated twice in the datafiles and the standard deviation for the Vy and Vz are fill data values of 99999.0. (For the times when the total ion density is between the two values, 2.1 to 3.0 x 10^4 ions/cc, the IDM remains in whichever mode it was in previously, so IDM data taken when the density is in this range can be in either normal or slow mode.) In the plots the top edge of the panel for Vy and the bottom edge of the panel for Vz on the plots are color coded to denote in which mode the IDM is operating with blue denoting normal mode and orange denoting the slow mode.

To start with the easy case, if the total density is less than 10^3 ions/cc the IDM is not sampling enough ions to produce a valid result no matter what the composition is. For those cases we set all the IDM quality flags to 4 (bad, do not use). There are also special cases that we have identified from within the telemetry (IDM density reported as negative) where the electrometers in the IDM get into a state where they produce invalid data. These points have the IDM quality flags set to 4 (bad, do not use) and the Vy and Vz data in the file are set to fill data values (-9999.0).

While the IDM can, in principle, measure crosstrack flows with absolute values as high as 3500 m/s we rarely see absolute flows get as high as 2000 m/s and we do not trust any absolute flows higher than 2500 m/s. Thus if Vy or Vz has an absolute value above 2500 m/s we set the IDM quality flags to 4 (bad, do not use).

The rest of the quality flags require knowing the composition of the plasma, but there are cases where the RPA analysis failed for that second. In these cases the IDM results may be valid or not, but there is no way to determine which without a user examining the data. So for cases where the fraction composition of O+ is negative (usually fill data of -99999.) or greater than 1.03 (unphysical) then we set the IDM flags to 5 (unknown). We leave it to the user to determine if these Vy and Vz data are valid. In some cases the densities are low and both the RPA and the IDM are producing bad data, but the RPA is failing before the IDM is so that there are no usable RPA data while there are still IDM data. In such cases the Vy and Vz flows should be obviously bad to the user. On the other hand we have cases such as seen in figure 3 where the Vy and Vz data in the northern polar region are flagged as unknown (blue on the plot). In most cases the RPA data will fail for only a few seconds resulting in short bursts of Vy and Vz data being flagged as unknown. For some unknown reason in the figure 3 orbit, there is a period of several minutes where the RPA failed (as seen in the data gaps in Ti, Vx, and the light ion density trace in the bottom plot. Even through the RPA analysis failed here, the Vy and Vz flow data here seem realistic and follow the patterns we would expect under these conditions. Thus we would personally rate these “unknown” Vy and Vz data as good and use them with high confidence.

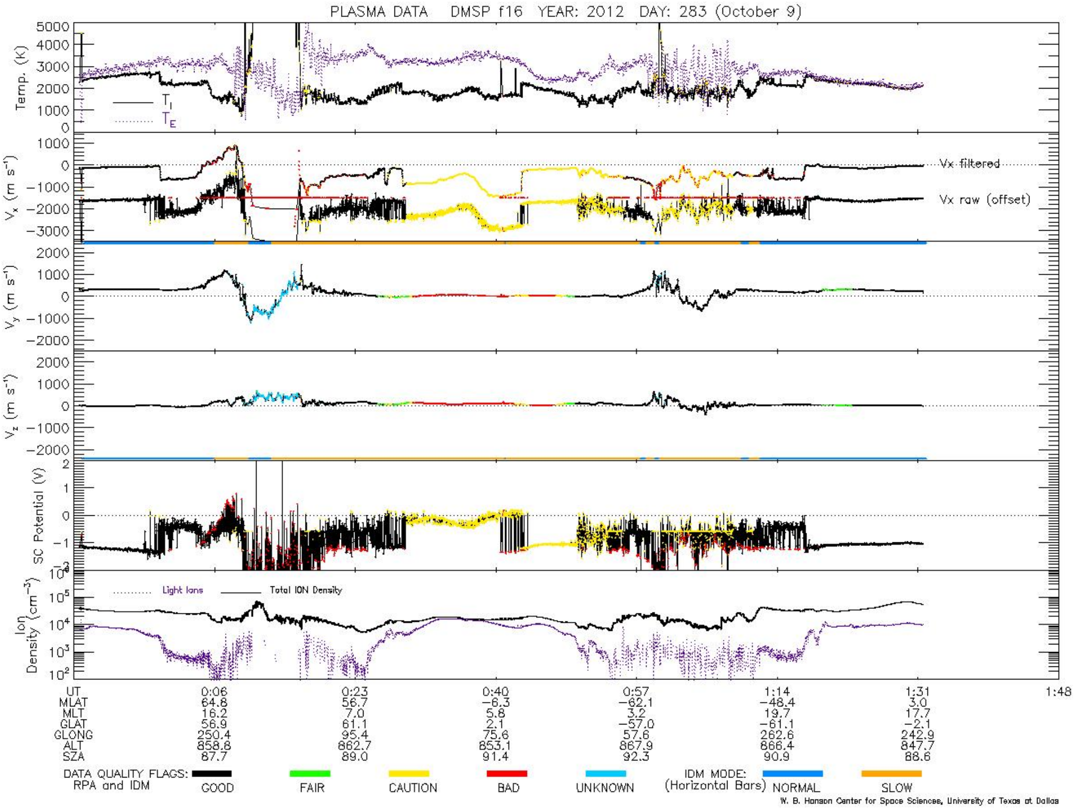

Figure 3

The rest of the IDM quality flags are set based on the fractional amount of O+ in the plasma. The UT Dallas-built IDMs were originally designed to function at lower altitudes with higher plasma densities, and the topside ionospheric conditions are frequently outside the envelope of their optimal operating conditions. In general the higher the density and the higher the percentage of O+ in the plasma, the better the quality of the data. While we have tried to determine the exact density/composition boundaries where the data transition from good to fair to caution to bad, we have not been completely successful. We have discovered that for some ranges of a given density and composition of the plasma, there are multiple examples of a variety of data quality. So we cannot do a simple one-to-one mapping where we say for a certain density and percentage of O+ the results from the IDM will *always* be of a certain quality. Instead we have come up with what are admittedly "rules of thumb" criteria. We have been slightly conservative in setting these boundaries, but we have been less strict than we were in setting the SSIES-2 quality flags. Our hope is that we have increased the overall fraction of the IDM data flagged as good or fair compared to the earlier SSIES-2 data. In general we feel all the Vy and Vz flagged here as good or fair can be trusted with only cursory checking by the user. The IDM data flagged as caution we leave to the user to judge. And in many cases the IDM data flagged as bad should be discarded, but not always. There are cases where "bad" IDM data may actually be usable as will be discussed below.

Our rule of thumb settings of the quality flags are as follows:

We start here assuming that the plasma density is already above 10^3 ions/cc because those data have already had the IDM quality flags set to 4 (bad) by the performed check above. After that:

If the fractional amount of O+ is greater than 0.75 and less than 1.03 then the IDM quality flags are set to 1 (good).
If the fractional amount of O+ is greater than 0.65 and less than or equal to 0.75 then the IDM quality flags are set to 2 (fair).
If the fractional amount of O+ is greater than 0.55 and less than or equal to 0.65 then the IDM quality flags are set to 3 (caution).
If the fractional amount of O+ is less than or equal to 0.55 then the IDM quality flags are set to 4 (bad, do not use).
Note that these boundaries of fractional O+ are based on the F16 and F18 data from 2012. As the ionosphere changes throughout the solar cycle we will adapt and change these boundaries for the data from other years.

As stated above, these flags are intended as a first guide, not necessarily as an absolute final judgment. For example, figure 4 is a pass from May 2012 where the Vy and Vz data in the northern polar regions (left side) are obviously good while the Vy and Vz data in the southern polar region (right side) are obviously bad. This pass occurred in late spring and the bottom panel displaying the total density (black) and the light ion density (purple) show that in the northern polar region the total density is in the upper range of 10^4 ion/cc while the light ions are in the low to midrange of 10^3 ion/cc. The color blue at the upper edge of the Vy panel and lower edge of the Vz panel indicate that the IDM was in normal mode throughout most of the northern hemisphere. Thus the Vy and Vz data in the northern polar cap are all flagged as good (black) and the changes in the flows as the spacecraft flies through the polar convection pattern are smooth and regular. In the southern polar region this pass occurs in late autumn where the total ion density drops to the mid 10^3 to the low 10^2 ion/cc ranges and the light ions comprise

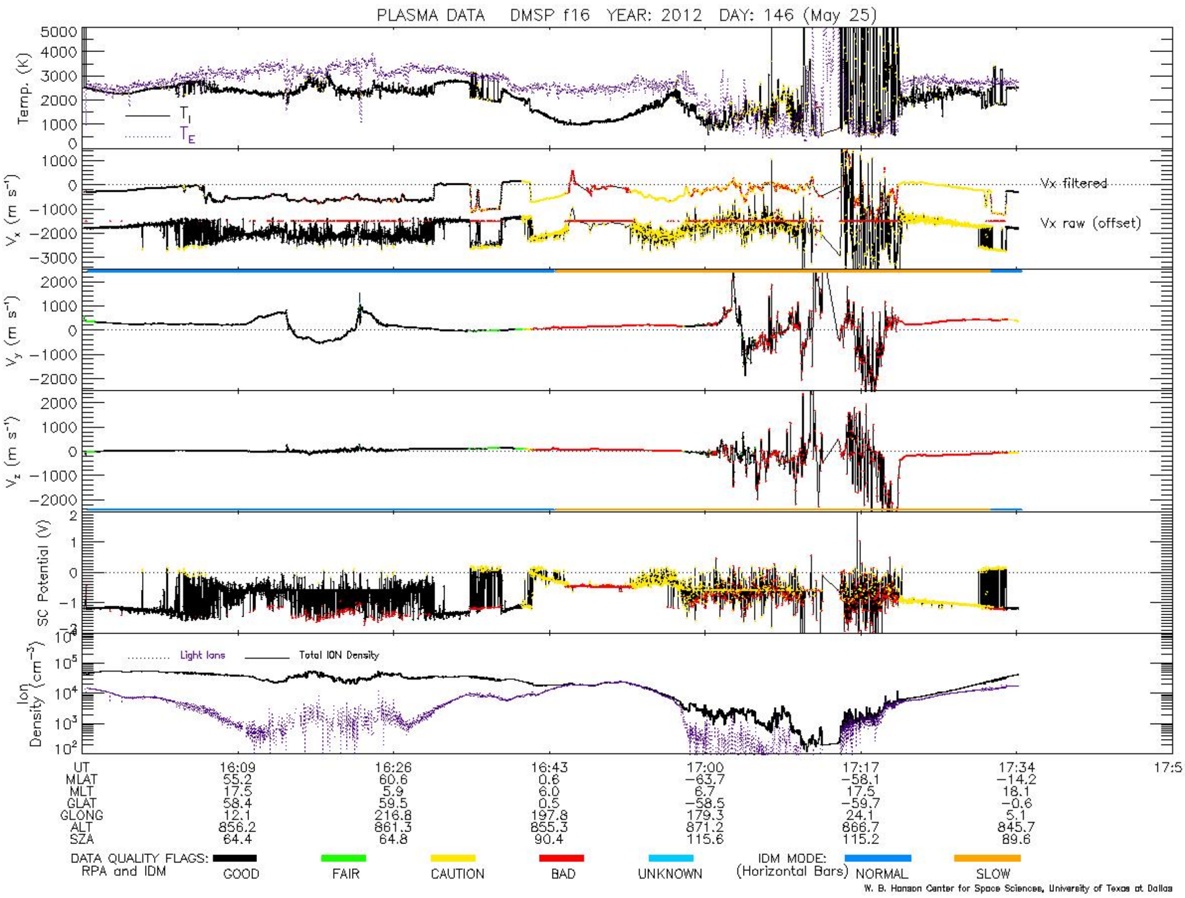

Figure 4

more than 50% of the plasma. The color orange at the upper edge of the Vy panel and lower edge of the Vz panel indicate that the IDM was in slow mode throughout most of the southern hemisphere. Thus the Vy and Vz data in the southern polar cap are all flagged as bad (plotted as red dots though the lines connecting them are plotted in black). Here the Vy and Vz flows are extremely noisy with frequently unphysically large magnitudes. In general we would discard these bad Vy and Vz data, though a user could run these data through a median filter to get a qualitative (not quantitative) sense of the southern polar regions convection flows.

While these flags work well in distinguishing good from bad IDM data in the polar regions, it is not as successful in the midlatitude and equatorial regions. Notice in figure 4 the dawnside Vy and Vz data between the subauroral region in the northern hemisphere and the subauroral region in the southern hemisphere. This makes up the middle third (horizontally) of the plot. Both the Vy and the Vz flows in this region are near zero for the entire time with little variation or noise. This is what we would normally expect of the topside ionosphere flows in these regions. But note that on the left side of this segment (northern hemisphere) the data are all colored black to indicate they are flagged as good, while on the right side of this segment (southern hemisphere) the data are all colored red to indicate bad data that should be discarded. And yet these "bad" flow data appear to be just as valid and reliable as the good flow data in the northern hemisphere. These are flagged as bad because, as can be seen below them in the bottom panel, these data are taken in a region that is essentially 100% light ions (the purple light ion density trace overplots the black total ion density trace). Thus we are faced with a dilemma. We cannot flag these bad points as “good” without also flagging the obviously bad flow data in the southern polar region as also “good”. So our final compromise is to do this: stick with our relatively safe and conservative algorithm of flagging IDM data comprised of less than 55% O+ ions as bad knowing that this will catch the vast majority of the truly bad data, and leave the final decision to the user. If the user is in a situation where, say, they want to use southern midlatitude dawnside IDM flow data that are initially flagged as bad, but they feel based on other criteria these appear to be valid data, then they can override our quality flags and use these data.

The sun glint

There is another issue of questionable data that appears in the SSIES-3 IDM data and that is the occurrence of what we term the “sun glint”. The IDM is designed such that there should be no contamination of the crosstrack drift measurements from photoemissive electrons. If sunlight shines directly into the aperture and onto the measuring collectors, any photoemissive

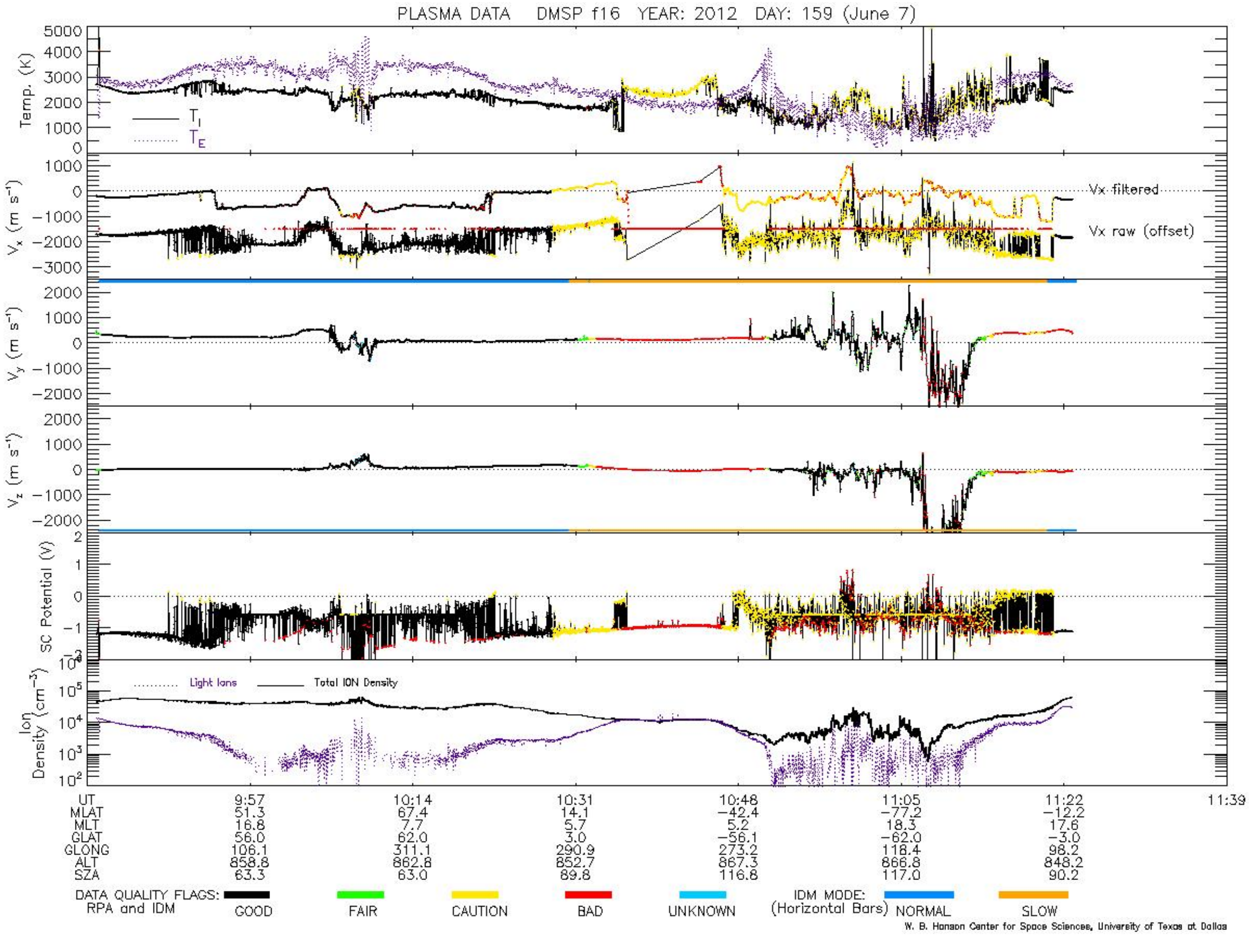

Figure 5

electrons coming off the collectors would be repelled back by the negatively charged grid just above the collectors so the electrons would return to the collectors giving a net current of zero. However, if the sunlight is falling almost perpendicular to, and slightly in front of, the aperture, then it is possible that sides of the IDM cup behind the negative repellor grid and in front of the collectors could be illuminated and produce photoelectrons. These photoelectrons would reach the collectors and create a spurious current that has nothing to do with current from the crosstrack ions. These electrons would fall disproportionately on one side of the collector producing a large and spurious flow in the measured data. Figure 5 shows an example of this where the region of large negative (about -2000 m/s) Vy flow and even larger negative (< -2500 m/s) Vz flow are seen in the duskside subauroral region on the duskside southern hemisphere around 11:10 UT. This pass is from 7 June 2012 were the spacecraft is just coming out of the darkness near sunset in the winter (southern) hemisphere and the sun is on the horizon at nearly right angles to the spacecraft's velocity vector. Obviously, these large flows are not real but they appear in every orbit at roughly the same location in the orbit for about a month. Thus we conclude this is an effect of the sun angle causing a "sun glint" in the IDM that causes these spurious crosstrack flows. This region of bad data slowly appears in the data set and then leaves the data set over a period of several weeks. For F16 in 2012 the effect begins about May 20 and disappears after the end of July. This is an effect caused by the change in the sun's angle relative to the spacecraft's orbit over the course of the seasons. Thus the problem occurs only during a limited time of the year. It is not observed in the IDM data the rest of the time because during the other times the orbital path is such that the spacecraft never reaches the proper orientation for the sun to shine into the IDM producing the "sun glint".

Figure 6 shows an example of the "sun glint" in a F18 orbit on 14 September 2012. The bad data (highlighted by the vertical red lines on the left) occur between roughly 13° to 17° north latitude on the duskside at about 20 hours SLT. The region of bad data in Vy and Vz is not as large in magnitude and extent as the bad data in the F16 orbit seen in figure 5 above. However this run of "sun glint" bad data is longer lived than the F16 "sun glint" bad data. Because it occurs near the equator the location of the "sun glint" bad data does not move and persists for almost the entire year. In January 2012 the bad data is in the midlatitude northern hemisphere and migrates down to just north of the equator by the start of March 2012. It remains in that area (as seen in figure 6) through about mid-November 2012 when it disappears. It could be argued that these bad data were caused by conditions of low density / high percentage of light ions as seen between the two red lines on the left side in the bottom panel during the bad data period. However, F18 experiences essentially the same conditions later in the orbit on the duskside in the darkness of the southern midlatitudes which are denoted by the two vertical red lines on the right side of the plot. The overall ion densities are essentially the same in both regions and the percent of light ions is even higher in the southern midlatitudes compared to the northern midlatitudes, yet in the southern region the Vy and Vz flow are near zero and noise-free. The only significant difference between these two regions is the presence of the sun on the horizon in the northern midlatitude compared to the darkness in the southern midlatitude.

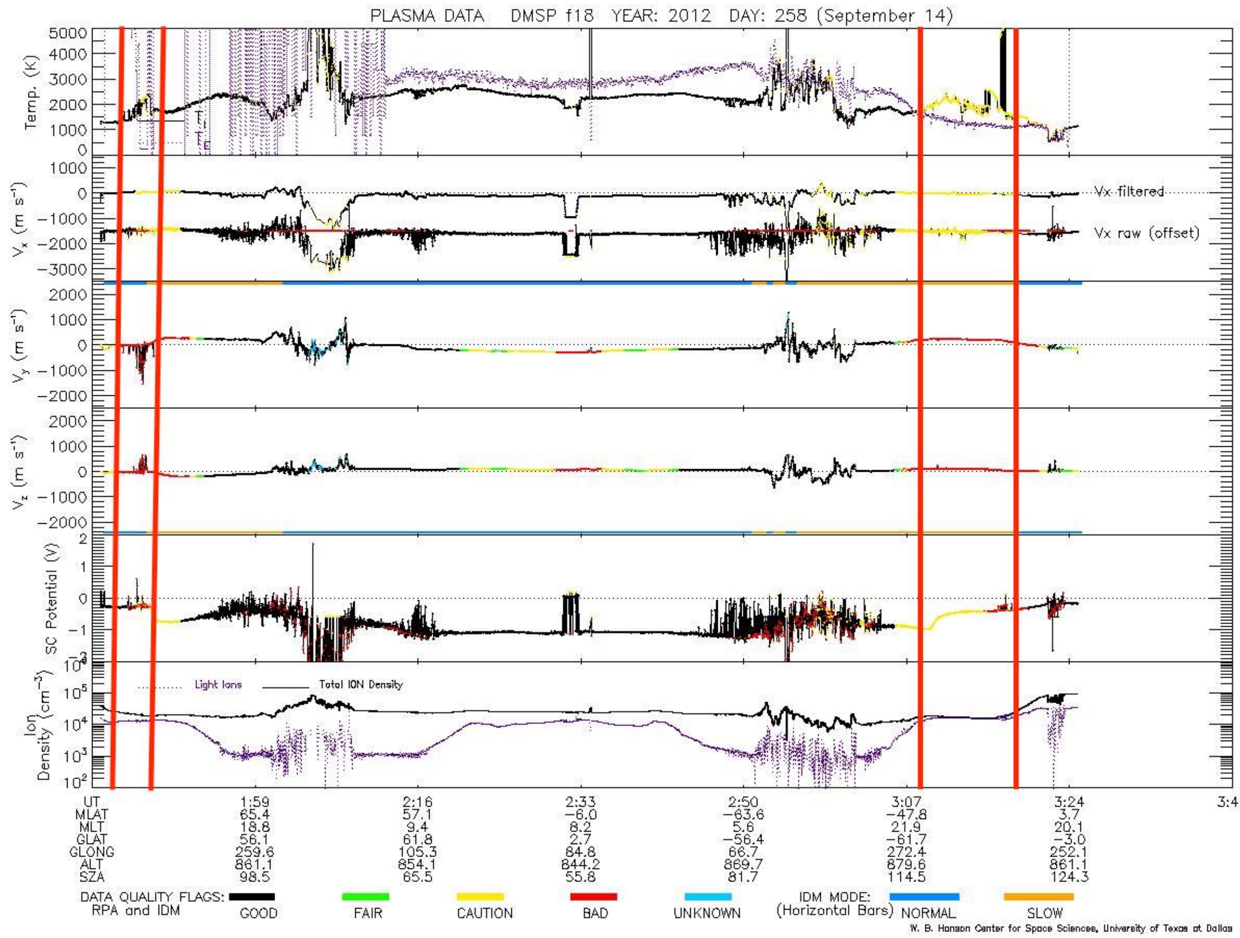

Figure 6

In the F17 data for 2012 there are places where obviously bad data occur in the IDM data. Figure 7 shows a F17 orbit from 4 January 2012 where the vertical red lines in the left side mark a region of large negative (< -2500 m/s) Vy flows and large positive (> +2500 m/s) Vz flow in the dawnside auroral region in the northern hemisphere. These erroneous flows are not associated with any "sun glint" but rather are the result of extremely low ion densities (< 10^3 ion/cc) in the polar cap as seen in the bottom panel between the two red lines. This extremely low density region is also associated with greater than normal noise in the Vx (second panel) and erroneously high Ti values (top panel). The quality flags for the Vy and Vz flows in this region are already flagged as bad based on the low density, so they are already removed from consideration as valid data.

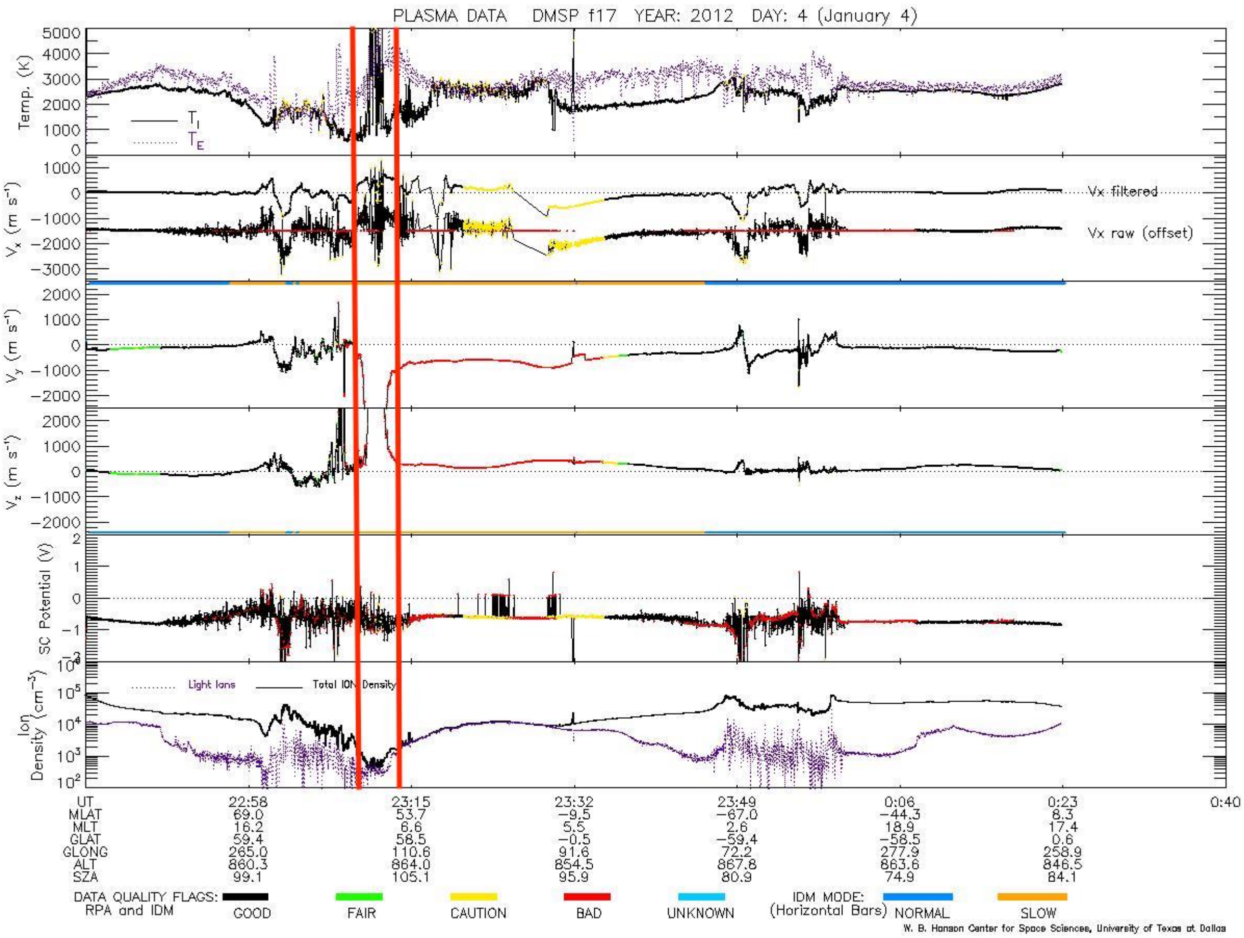

Figure 7

While spotting these bad data is easy by eye, we have not had any success so far in developing an algorithm that will correctly flag these bad data for ***all*** possible cases. We are still working on this problem and hope a future version of this code will incorporate such an algorithm. Until then we can only warn the users to be on the lookout for these short-lived "sun glints" in their data that may contain data that have been incorrectly flagged as good or fair.

Skimmers

All the DMSP spacecraft are in sun-synchronous polar orbits with an inclination of 98°. With this inclination the orbital plane precesses counter-clockwise with respect to the celestial sphere about once per year, thus keeping the orbit tracking with the local solar time on the Earth below throughout the year. While the spacecraft's orbit covers the same path in local time and latitude for every orbit, the Earth is turning beneath the orbit allowing the spacecraft to cover most of the Earth twice a day. Although the spacecraft never goes above ±82° geographic latitude, the tilt of the Earth's magnetic dipole causes it to "rock" back and forth under the orbital plane once per day. This means that there are times during the day when the magnetic poles are close to the spacecraft's orbital plane and other times when the poles are distant from the orbital plane. For all of the SSIES-3 DMSP (F16 through F19) the early part of the day from roughly 2:00 UT to 9:00 UT are when the magnetic poles are most distant from the orbital plane. Figure 8 presents a typical F16 orbital plot starting at 16:37 UT. In the third

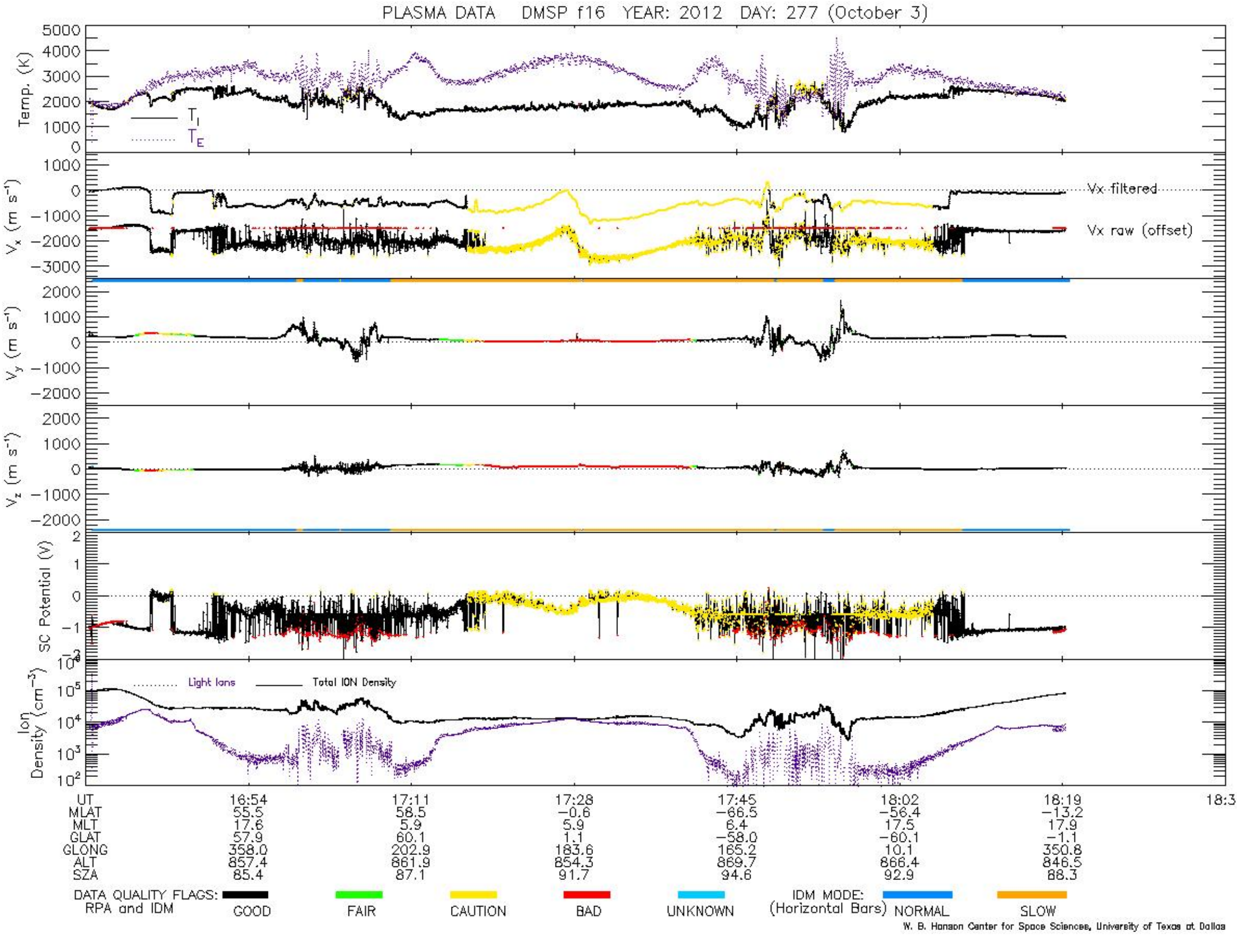

Figure 8

and fourth panels the Vy and Vz flows show a clear two-cell convection flow in both polar caps and the ion temperature (black trace in the top panel) clearly shows structure and variation in the polar cap region. In this orbit the northern polar path passed almost over the north magnetic pole while the southern polar path got within 5° of the south magnetic pole. Meanwhile the F16 pass from nearly 12 hours before is shown in figure 9 below. Here the northern polar pass is on the dayside and never gets closer than 15° to the north magnetic pole and the southern polar pass is on the nightside and never reaches any closer than 20° of the south magnetic pole. Note that the Vy and the Vz flows here in figure 9 are essentially flat and that the ion temperatures in both polar regions here are almost constant. Both these passes are so far away from the magnetic poles that the spacecraft either misses or skims the polar ion convection pattern in the ionosphere. Because of this we refer to polar passes like this as “skimmers”. We point this out so that the user will recognize these “skimmer” passes as normal and not as times when the instruments were malfunctioning. While these “skimmer” passes are regular during quiet geomagnetic periods, during active geomagnetic periods or

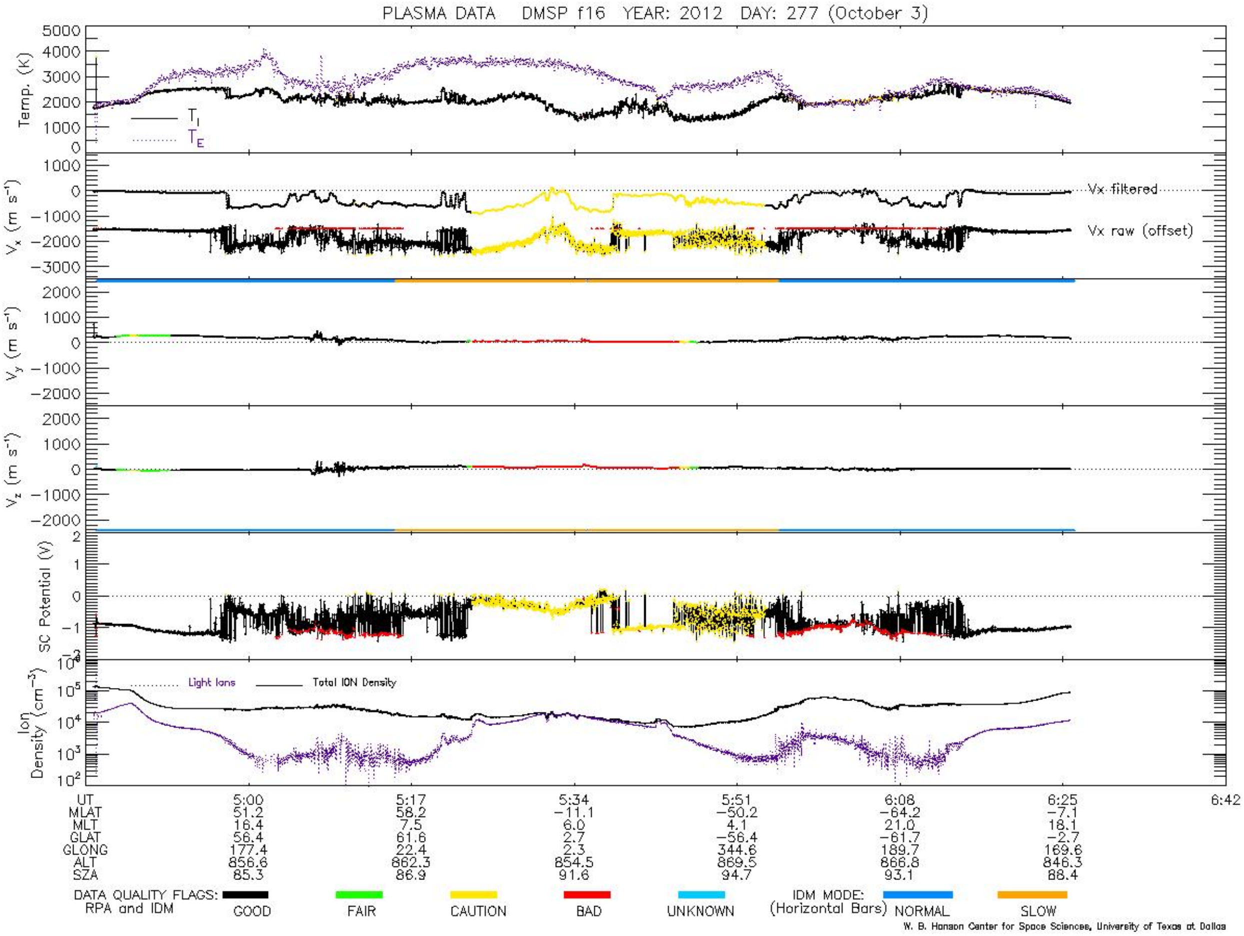

Figure 9

geomagnetic storms the convection patterns expand equatorwards so polar passes during these times so then the instruments do observe significant Vy and Vz flows.

SEU and Equatorial Resets

Another common type of missing data in the SSIES-3 dataset are missing "half-orbits" of data caused by a single event upset or "SEU". The electronics on SSIES-3 proved to be more susceptible to upsets caused by high energy particles interfering with the processing electronics than the previous SSIES-1 and SSIES-2. Thus every few weeks or so the electronics suffers an upset and puts SSIES into safe mode. This occurs most often near the equator in the South Atlantic Anomaly. Figure 10 shows an example of an orbit where this occurs. The spacecraft completes half an orbit and near the dawnside equatorial crossing an SEU occurs and there are no more data for the rest of the orbit. Initially we thought such SEUs would be rare and we would just radio up a reset command if we saw the SSIES-3 had entered safe mode. During the early phase of F16 we realized this was happening more often than expected and as it required a ground common to reset we were losing too much data (orbits worth of it) as the SSIES-3

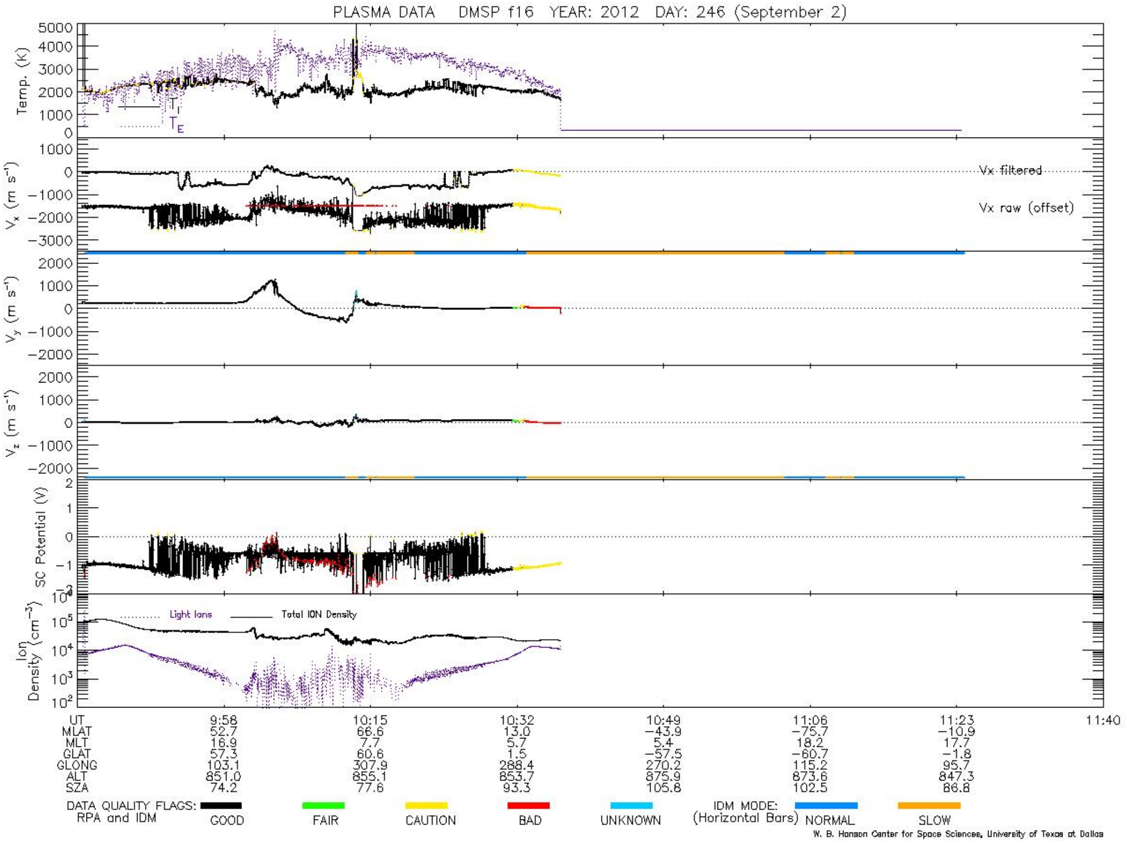

Figure 10

went into safe mode too often. To correct for this we reprogrammed the SSIES-3 to do an automatic reset twice each orbit near the equatorial crossing. This way we would only lose at most half an orbit of data.

This leads to another source of bad data. When SSIES-3 is operating nominally and then resets, there can be several seconds of bad data as the instruments reset themselves. Most of the instruments just return fill data during the reset, but the IDM can return bad data during this period. Usually the bad data are flagged as bad because the values are outside the norm, but there can be times when the values are close enough to nominal that they may be flagged as good. So anytime near the equator when there is a block of four to eight seconds that seem “off” or show a short spike or step function in their values, the user should view those data with caution. An example of this can clearly be seen in figure 11 where there is spike in the Vy (third panel) at the center at about 20:18 UT. A smaller discontinuity appears in the Vz (fourth panel) at the same time, but is only barely visible in the plot here. These points are already flagged as bad because of the high percentage of H+ in this region, but there are times when these small spikes appear in the middle of good/fair data and are also incorrectly flagged as good/fair.

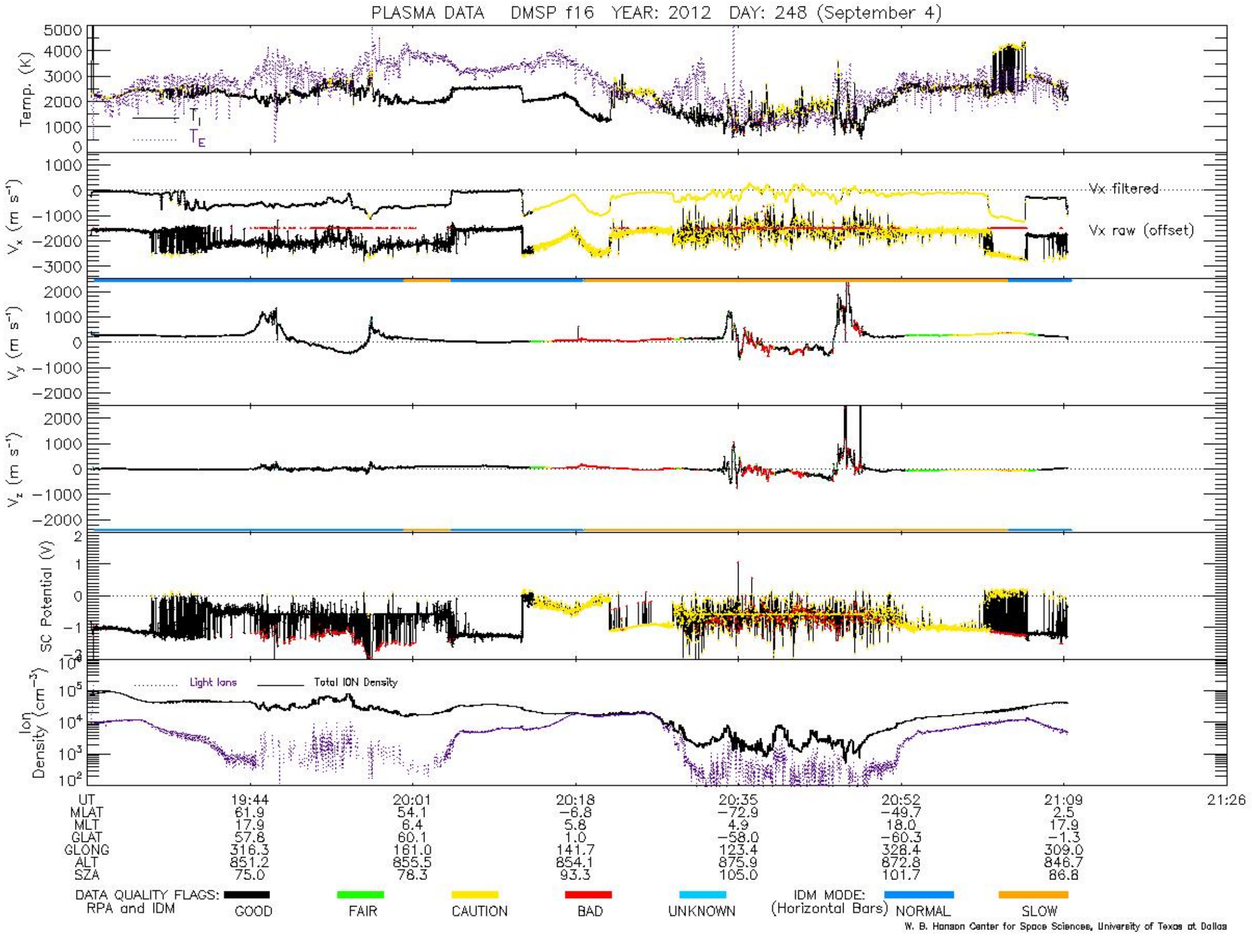

Figure 11

Special IDM quality flags for F17

Prior to 23 June 2009 the IDM on F17 operated nominally and matched the output of the other SSIES-3. But on 23 June 2009 around 6:30 UT the IDM began returning anomalous data. Figure 12 show a before and after examples of the IDM data. (Note, these are plots from our earlier plotting routines so the format does not exactly match the plots seen above.) Here the third and fourth panels show the Vy and Vz flows respectively. The orbit on the left was from two days (21 June) before the anomaly when there was enough geomagnetic activity to show a standard two-cell convection pattern in the northern hemisphere. Because of the low density in the southern (winter) hemisphere the IDM data there are very noisy and unusable (right side of the plot). The plot on the right shows an orbit from a few hours after the anomaly on 23 June. At this point the IDM was returning flows of identically zero values for both Vy and Vz data except for the northern polar region where both were returning large negative flows close to the instrumental limit. (The cosine wave seen in the Vy data in the plot results from the corotation correction that is applied to the all-zero Vy flow data.) Despite intense investigation into the cause of this anomaly, we were never able to make a final determination of what caused the anomaly nor could we come up with any way to mitigate this problem. After several

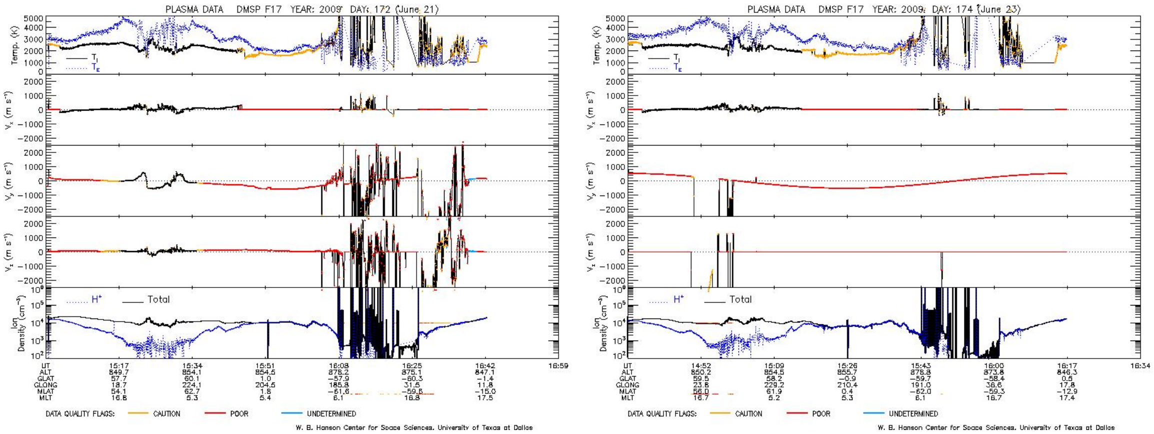

Figure 12

months the F17 IDM was declared to be inoperable and we assumed that all future IDM data from F17 would be lost.

The anomaly occurred during the extended “deep solar minimum” from 2005 to 2011 and as the overall density of the ionosphere at the 850 km altitude of F17 began to increase, we started to see indications of the IDM beginning to work again. Figure 13 below shows two examples of this partial return of data signal. On the left is an orbit from 23 December 2010 where the data in the southern (summer) polar cap (right side of the plot) shows structure that

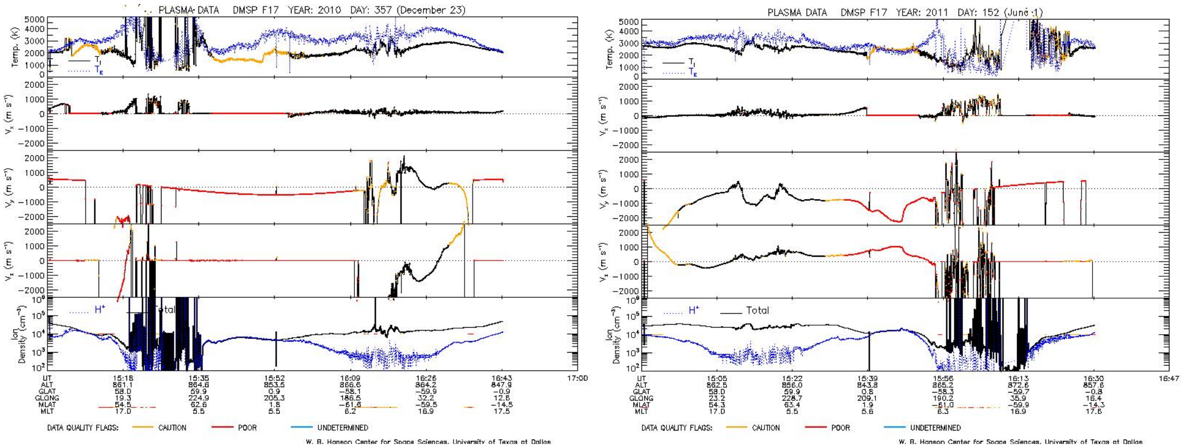

Figure 13

resembles what we would see in a polar convection flow, but the magnitudes and the baseline are very wrong. By 1 April 2011, the time of the orbit on the right, the Vy flows in the northern hemisphere (left side of the plot) clearly shows a clean two-cell convection along with upward flows in the Vz in the auroral oval. However the baseline of both the Vy and the Vz here are obviously wrong. And in both cases the ionospheric conditions in the opposite hemisphere (north on the 23 December 2010 plot and south on the 1 April 2011 plot) are low density, high percentage of light ion conditions, so the IDM data there are so noisy as to be garbage or else they revert to the "all zero" conditions.

By the beginning of 2012 the IDM data had returned to nominal conditions. Figure 14 shows a F17 orbit from 1 April 2012. We never did solve the mystery of what caused the anomaly in the first place nor what caused the IDM's recovery. All we know for certain is that the quality of the IDM roughly correlates with the solar cycle although the we could not find a direct correlation between the data quality and either the total density or the O+ density.

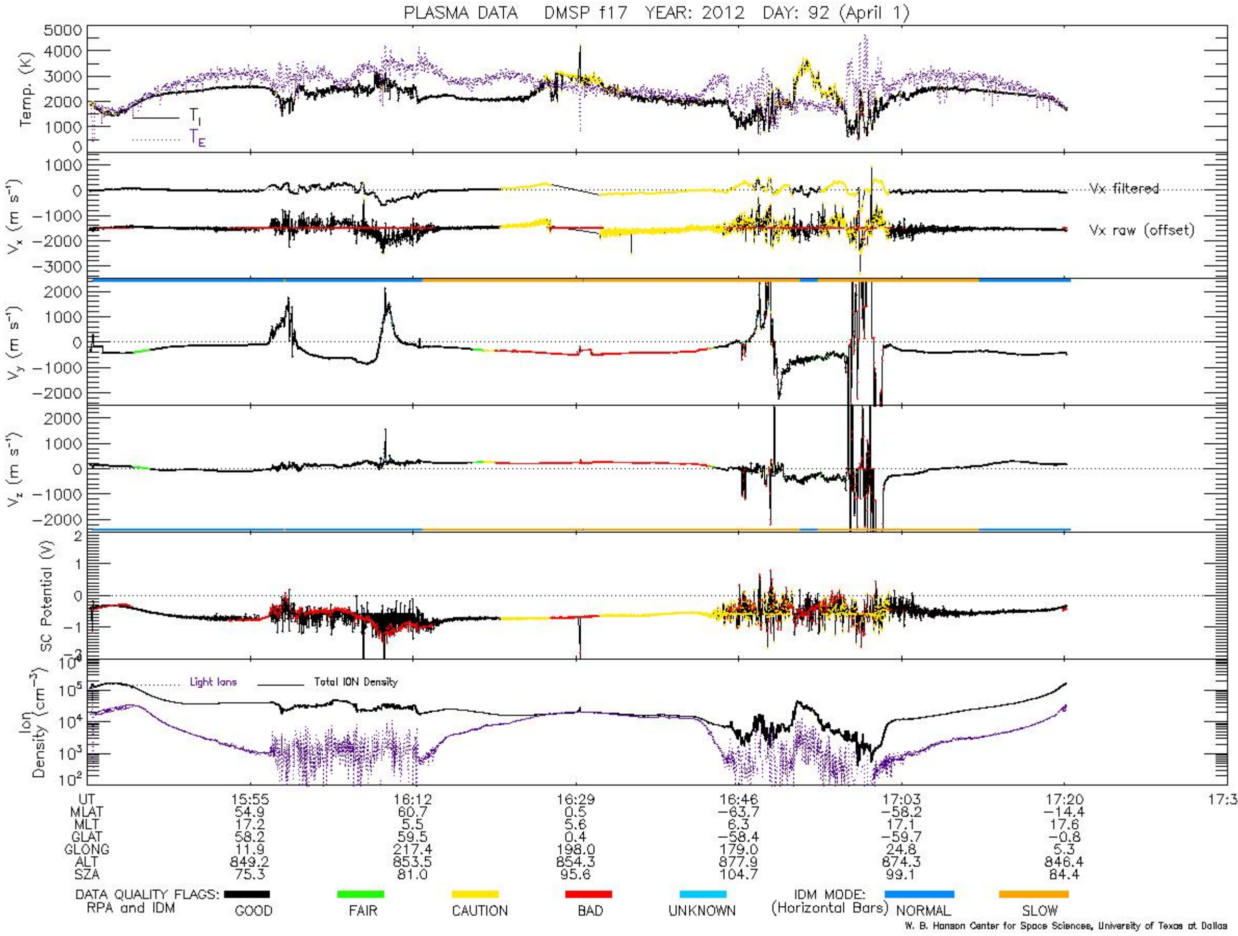

Figure 14

Because of this anomaly we knew that the IDM quality flags for F17 would have to be separate from the other SSIES-3 during the solar minimum period, but we had hoped that during the solar maximum condition we would be able to use the same quality flags as F16 and F18. Unfortunately, that did not work out. Examining the 2012 and 2013 data (which corresponds to the peak of the solar maximum) showed that the baseline flows of Vy and Vz in the midlatitudes and the equatorial regions are still suspect. In figure 14 above the Vy flows outside of the polar regions show a significant non-zero negative flow of between -200 and -400 m/s. The Vz flow offset is not as large but is more complicated. On the dawnside leg (middle of the plot) Vz is slightly positive. On the duskside (the two outside segments flanking the polar regions) Vz is slightly negative near the subauroral regions and slightly positive in the midlatitudes and equatorial regions.

Thus, even under the best solar maximum conditions, we do not have as much confidence in the IDM data from F17 as we do from the other SSIES-3 IDMs. We cannot say for certain that the absolute values of the F17 Vy and Vz data are reliable because we do not know for certain what their baselines are. To compensate for this all the F17 IDM quality flags after 23 July 2009 use a second set of quality flags. Repeating what was stated above, the flag codes and the colors codes for F17 after 23 July 2009 are:

6 = data are judged to be from **good** conditions but some caution should still exercised by the user because of baseline issues **(black) (F17 only)**
7 = data are judged to be from **fair** conditions but some caution should still exercised by the user because of baseline issues **(green) (F17 only)**
8 = data are judged to be from conditions that warrant **caution**, and since these are F17 data the user should be doubly cautious in using them **(yellow) (F17 only)**

The bad IDM data from F17 after 23 July 2009 will be flagged as 4 (bad, do not use) just as with the rest of the SSIES-3 data.

For the 2012 and 2013 F17 data we are using the same boundary conditions on densities and O+ fraction that we use for the F16 and F18 data. Despite the fact that the algorithm is the same we felt that we had to use the separate quality flags to remind the user that the F17 IDM data values are not as certain and reliable as the other IDM data values. We expect that for other years we will use different boundary conditions to set the flags for F17 relative to the flags for the other spacecraft, but that will be determined on a year by year, case by case basis.

Marc Hairston
September 2019