Response to referee no. 2

We thank the referee for their constructive comments, and believe that addressing the points presented in the comments will greatly improve the manuscript. Responses to each comment can be found below.

First on a general note, in order to understand the foreshock conditions surrounding the transients, we have compared the plasma properties of the transients to those of the surrounding ULF wave field. We do this by finding the troughs in the ULF wave field, since they can be directly compared to the transients consisting of decreases in density/magnetic field magnitude. This method is illustrated in Figure 1 below. We define the troughs as local minima in the proton number density below the input solar wind value. Unlike cavitons and SHFAs, we do not track the motion of the troughs, but use them only to calculate statistics of various plasma properties (e.g., density, temperature and bulk speed), which are compared to the tracked transients. Only troughs in the relevant region are selected for these statistics (e.g., within 1/4/10 RE from the bow shock). We will refer to these results in the responses below, and also add them to the revised manuscript.

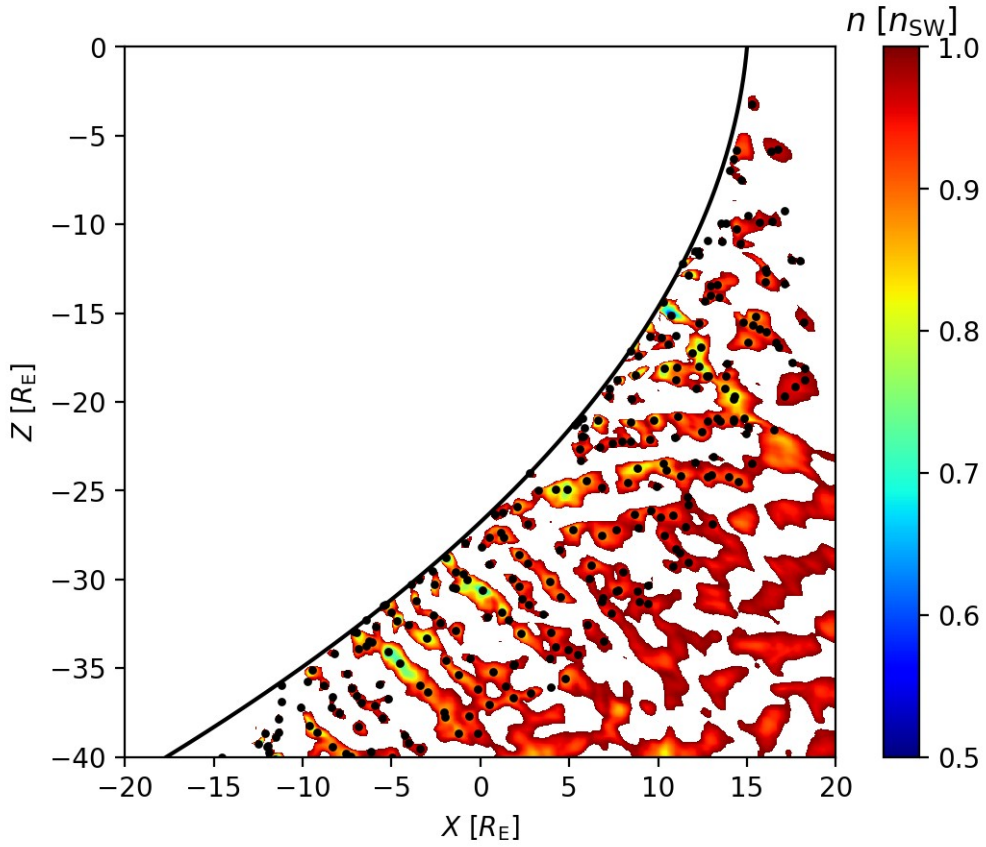

Above: A plot showing ULF wave troughs / local minima as black dots within 10 RE from the bow shock at time t=900.0 s. The colormap shows values of proton number density below the input solar wind density nSW. The bow shock is modelled with the 4th order polynomial described in the manuscript.

MAJOR:

Throughout there is very little explicit comparison of the properties of the transients compared to the foreshock in general, let alone the ambient foreshock at the transient's location. Instead mostly only values in the pristine solar wind are used for comparison. However, understanding how the structures differ from their surroundings is of vial importance and needs to be incorporated into the work throughout. This affects numerous aspects of the work, including:

*** Are the choices of properties and thresholds for detection of the transients suitable? How does a 20% decrease in density compare to the variability in density associated with the foreshock ULF wave field? Is plasma beta a sensible parameter to use to distinguish between cavitons and SHFAs (I would have thought a temperature criterion would have been more appropriate) and how does a value of 10 compare to the typical foreshock and its variability?**

The choice of a 20% limit is the same as in earlier spacecraft studies by Kajdic et al. (2013, 2017). However, in these studies, the events had to fulfill a subsequent criterion based on a function defined as χ(t) = (n(t) -<n>) * (B(t) - <B>) (where n(t) and B(t) are the density and magnetic field magnitude at time t and <> denotes a time average). The criterion requires that the value of χ inside cavitons must be at least 5 standard deviations larger than the temporal average of χ over the observation interval. We have omitted this subsequent criterion in order to be able to detect small transients and study the temporal evolution of the transients.

In general, the density and magnetic field magnitude fluctuate ~5-10% from their solar wind values in the foreshock, and the amplitude of the fluctuations is below our caviton detection criteria. This is demonstrated below, where the temporal averages of proton number density np and magnetic field magnitude B are shown over a 120 s interval.

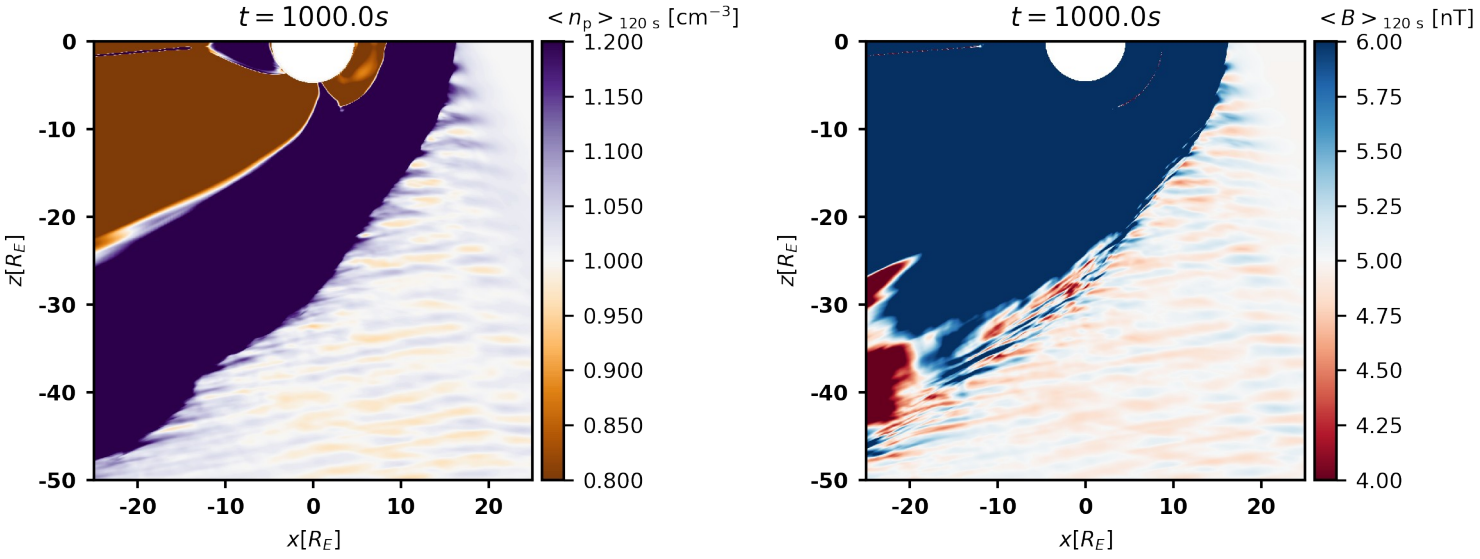

<u>Above: Temporal averages of proton number density (left) and magnetic field magnitude (right) in the foreshock over a 120 s period.</u>

More specifically, the depths of the transients can also be compared to those of the surrounding ULF waves by taking into account each trough (i.e., local minima) in the

foreshock below the input solar wind density. In the region in which cavitons are present (< ~10 RE from the bow shock), the mean depth of a trough in the wave field is ~12%, a bit over half of the caviton detection criterion. Structures below the 20% limit represent ~17.4% of all troughs in this region. We will add these results concerning the general wave field in the revised manuscript to better motivate the chosen detection criteria.

Beta was chosen as the SHFA criterion due to the large variation of temperature in the foreshock, which makes choosing an explicit temperature condition challenging. The physical motivation behind the choice of beta is that a large beta indicates that the transients are dominated by the plasma instead of the magnetic field. In the end, we retained the beta criterion as it appears to pick the differences between cavitons and SHFAs well, and in order to keep our results comparable with the earlier Vlasiator caviton/ SHFA study by Blanco-Cano et al. (2018), where the beta-criterion was originally used. A value of 10 was chosen visually. In the region where cavitons are found (< ~10 RE), the beta in a ULF wave trough has a mean value of ~4.5, well below our criterion. Values below 10 represent 92.3% of all troughs. Finally, we acknowledge that other SHFA criteria could be also applied. If a temperature criterion was used instead, we would obtain a similar classification, as panel d) in Figure 2 shows.

*** In Figure 2, how do the density of suprathermals and temperatures of cavitons and SHFAs compare to typical foreshock conditions? Are the velocities in these structures significantly different from the ambient?**

Using the same method as above, the suprathermal densities and temperatures inside cavitons/SHFAs and ULF waves can be compared by considering the values at the troughs in the wave field. However, both the suprathermal density and the temperature are sensitive to the location in the foreshock, with both increasing rapidly near the shock. We will supplement the revised manuscript with the below results.

In the region where most cavitons are found (< ~10 RE from the shock), the suprathermal density in a trough of a wave has a mean value of ~0.03 nSW. Panel g) in Figure 6 shows that the suprathermal density inside cavitons far from the shock is similar to this value. Near the shock, where both cavitons and SHFAs are present (< ~4 RE), the mean suprathermal density in a wave trough is ~0.05 nSW, and a comparable increase is found inside cavitons.

For the temperature, the corresponding mean values at <10 RE and <4 RE are 2.8 MK and 4.3 MK, respectively. Comparing these values to panel i) of Figure 6 shows that the temperature inside cavitons is similar for both ranges. As stated in the manuscript, the suprathermal densities and temperatures are larger inside SHFAs than inside cavitons, and thus by extension, larger compared to the surrounding ULF wave field.

Since the bulk velocity is also dependent on the suprathermal density, it shows a general decrease towards the bow shock. Within (10, 4, 1) RE from the shock, a mean bulk speed of (~728 km/s, ~714 km/s, ~680.2 km/s) is found in a wave trough. The speeds inside cavitons are similar to these values.

*** In Figure 3, are the correlations presented simply extensions of the overall foreshock or do they constitute distinct populations?**

For the proton number density and magnetic field magnitude, there exists a continuity across fluctuations within 20% of the ambient solar wind values to cavitons/SHFAs. Since the transients form from the ULF wave field, there is no clear limit between a ULF wave and a transient. However, the density and magnetic field magnitude are relatively well

correlated inside structures surpassing the 20% limit, as shown in panel a) of Figure 3. When all troughs are taken into account, it is seen that the values of the density and magnetic field magnitude appear spread for small-magnitude fluctuations. For larger depressions, the parameters become correlated. This is illustrated in the figure attached below.

For the suprathermal density, temperature and bulk flow speed, the distributions for density fluctuations within 20% of the solar wind density exhibit similar shapes as shown in panels b) and c) of Figure 3. The values are concentrated near their respective solar wind values, having similar ranges as cavitons. The 90th (10th) percentile values for the suprathermal density and temperature (bulk flow speed) are 0.07 nSW and 5.6 MK (690.6 km/s), respectively.

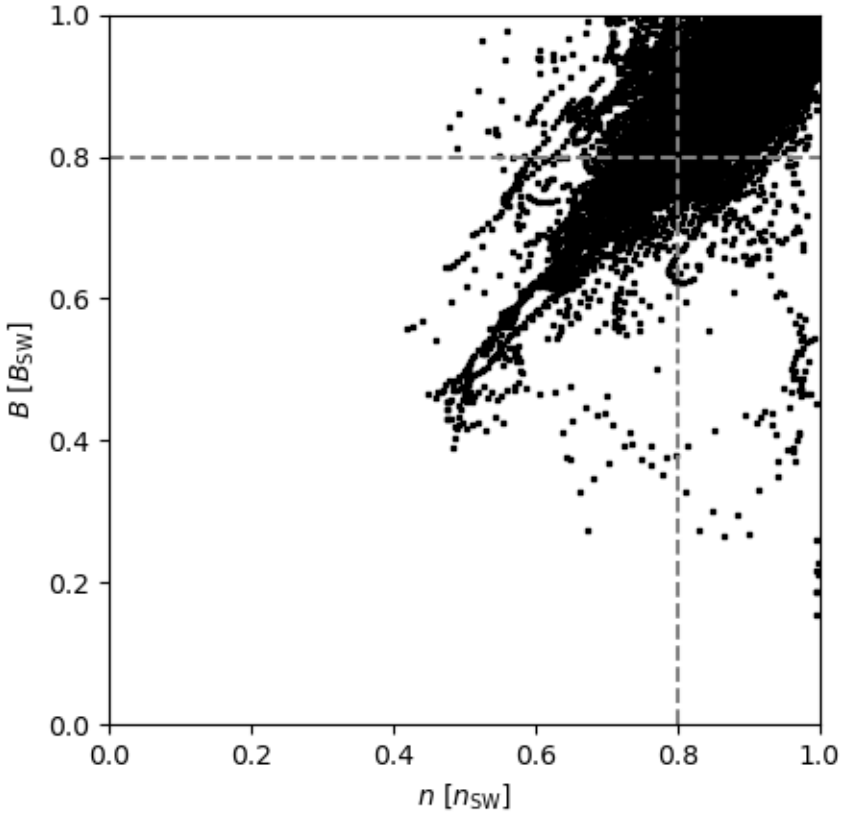

Above: A scatterplot of proton number density (n) and magnetic field magnitude (B) for all ULF wave troughs below the input solar wind number density.

*** I also have concerns over the results surrounding the suprathermal ions. The method employed of distinguishing between core and suprathermals uses the velocity and temperature of the pristine solar wind. This seems unsuitable for transients associated with flow anomalies, as the authors concede on line 200, and thus many of the results are likely micharacterising the solar wind and suprathermal ions in these structures. I would suggest the authors reprocess the data separating out regions in phase space using a distance condition in velocity space (based on the temperature in the pristine solar wind) either from the bulk or peak phase space density.**

Unfortunately, the recategorisation of the core and suprathermal populations is not possible post-run. The suprathermal population is resolved from the velocity distribution function as the simulation is running, and the distribution function is saved afterwards only in selected simulation cells to keep the sizes of the simulation bulk files tractable (of the order of GB as opposed to TB for a file with the velocity distribution available everywhere). Input solar wind values are utilised in the categorisation since the process is automated during the simulation run.

We also note that the method is valid in the majority of the foreshock (excluding major hot flow anomalies), as the deflections seen in SHFAs are a result of the velocity moments, as opposed to deceleration of the solar wind core. This effect is demonstrated for HFAs by Parks et al. (2013).

*** Related to the above, many conjectures around how the solar wind beam vs. the suprathermals are affecting the moments of the distribution are made, however, no velocity distributions are presented within the manuscript. It is known that the distrbutions within foreshock transients can evolve from multicomponent to single component plasmas, whereas the authors posit only the former.**

Since the manuscript already contains quite a lot of material with several large figures, we feel that the scope of this study should be limited to presenting statistics of the general properties and evolution of foreshock transients. In this regard, we propose to reduce the emphasis on the velocity distributions in the revised manuscript, and leave detailed study of their evolution as a topic for future work. The structure of foreshock velocity distributions is however briefly discussed in the previous Vlasiator study of cavitons/SHFAs by Blanco-Cano et al. (2018), and more extensively in a similar simulation performed by Battarbee et al. (2020), who studied the helium foreshock using Vlasiator.

*** Finally, the results with relation to the "nose angle" (which may be better described in the manuscript as meridional angle or solar zenith angle throughout) need to be understood in terms of the theta_Bn angle that the transient is magnetically connected to, since this largely controls the physics of the foreshock. This may aid in the interpretation of the results.**

ThetaBn is not used in this study since most of the analysed transients are located at the flank of the bow shock, where thetaBn has a narrow value range. Thus, we chose to use the nose angle instead, as it is better suited for analysing the spatial variation of transient formation and properties. We will add discussion on the range of thetaBn in the revised manuscript.

MINOR:

Lines 20-21: "before it is deflected by the magnetopause" This could do with rewording, since the bow shock also deflects the solar wind and the pressure gradients present throughout the magnetosheath (between bow shock and magnetopause) act to deflect the plasma around the boundary.

The sentence will be reworded followingly:

“The bow shock slows the solar wind down to submagnetosonic speeds while compressing and heating it. This allows the solar wind to flow around the magnetopause that separates the solar wind from the magnetosphere.”

Line 23: "far back into the upstream." This is not true for the entire region of the shock connected to the IMF, as the sentence suggests, only in the quasi-parallel case. Please reword this sentence, for example, removing the word "far".

The sentence containing this phrase will be rephrased as follows:

“At the quasi-parallel part of the bow shock (defined as the region where the shock normal and the IMF have an angle theta_Bn < 45), ions reflected off the shock can propagate several hundred ion inertial lengths upstream, forming the foreshock region. At the quasi-perpendicular shock (theta_Bn < 45), the upstream motion of the reflected ions is limited to an order of ion gyroradius, and a more abrupt shock crossing is found.”

Line 59: "SHFAs evolve" I would say they are "thought to evolve" since this is point requires further evidence in general and the results of the manuscript show it be the case only for some SHFAs.

This will be rephrased as requested.

Line 188: "SHFAs tend to be more depleted than cavitons" This could simply be an effect of the plasma beta condition so needs further comment.

For the large majority of SHFAs in our study, the relative increase in the temperature (T) is much greater than the relative decrease in the magnetic field strength (B). While beta has a stronger dependence on B than T, the decrease in B is countered to some extent by the comparable decrease in plasma number density (n), so that beta = nT/B^2 behaves as ~T/B. Since the change in T is much larger than the change in B, B should not impact the classification of the transients.

Figure 4: PDFs would be more helpful to readers than CDFs to see the regions where the transients actually form, rather than cumulatively from the bow shock up to some region where a certain proportion form. Some of the cumulative numbers can remain in the text, however.

We will change Figure 4 to the figure attached below, so both PDFs and CDFs are shown.

Table 1: Minimina and maxima of probability distributions are not robust statistics, the 25th and 75th percentile would be more appropriate columns to use. This would also remove potential confusion between the minimum and maximum value for each a particular transient used in the left column, which is appropriate.

This will be changed.

Figure 4: The label states these are counts, but they are proportions

This will be fixed.

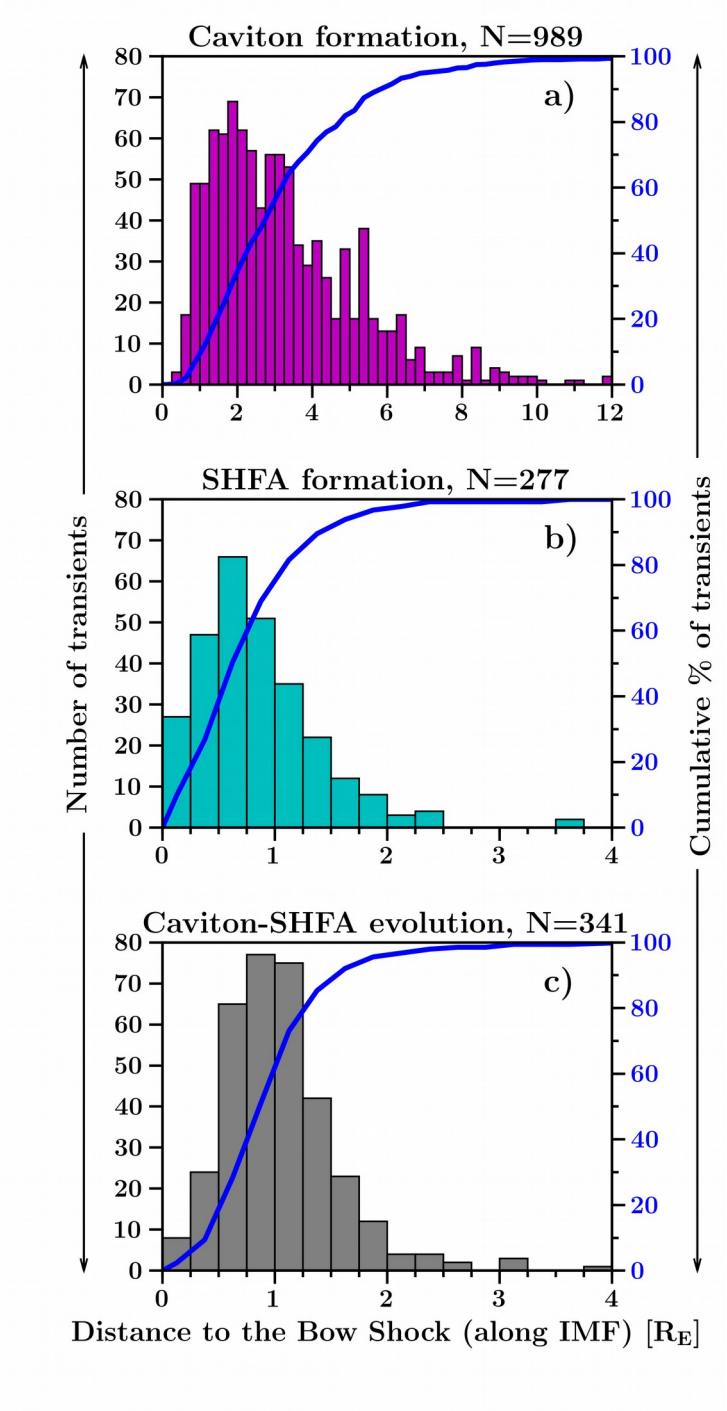

<u>Above: Updated Figure 4 displaying both PDFs and CDFs.</u>