

Key Points:

- Correlations between space weather data can be artificially inflated by common cycles and trends unrelated to physical relationships
- Cycles and trends can be removed by spectral subtraction, differencing, or autoregressive moving average transfer function models
- Transfer function models with cycles removed show correlations between wave activity and solar wind velocity, but not with electron flux

Correspondence to:

L. E. Simms,
simmsl@augsborg.edu

Citation:

Simms, L. E., Engebretson, M. J., & Reeves, G. D. (2022). Removing diurnal signals and longer term trends from electron flux and ULF correlations: A comparison of spectral subtraction, simple differencing, and ARIMAX models. *Journal of Geophysical Research: Space Physics*, 127, e2021JA030021. <https://doi.org/10.1029/2021JA030021>

Received 7 OCT 2021
Accepted 18 JAN 2022

Author Contributions:

Conceptualization: Laura E. Simms
Data curation: Geoffrey D. Reeves
Formal analysis: Laura E. Simms
Funding acquisition: Mark J. Engebretson
Methodology: Laura E. Simms
Project Administration: Mark J. Engebretson
Software: Laura E. Simms
Validation: Laura E. Simms
Visualization: Laura E. Simms
Writing – original draft: Laura E. Simms
Writing – review & editing: Laura E. Simms, Mark J. Engebretson, Geoffrey D. Reeves

Removing Diurnal Signals and Longer Term Trends From Electron Flux and ULF Correlations: A Comparison of Spectral Subtraction, Simple Differencing, and ARIMAX Models

Laura E. Simms^{1,2} , Mark J. Engebretson¹ , and Geoffrey D. Reeves³ 

¹Department of Physics, Augsburg University, Minneapolis, MN, USA, ²Climate and Space Sciences and Engineering, University of Michigan, Ann Arbor, MI, USA, ³Space Science and Applications Group, Los Alamos National Laboratory, Los Alamos, NM, USA

Abstract Simultaneously cycling space weather parameters may show high correlations even if there is no immediate relationship between them. We successfully remove diurnal cycles using spectral subtraction, and remove both diurnal and longer cycles (e.g., the 27 days solar cycle) with a difference transformation. Other methods of diurnal cycle removal (daily averaging, moving averages [MAs], and simpler spectral subtraction using regression) are less successful at removing cycles. We apply spectral subtraction (a finite impulse response equiripple bandstop filter) to hourly electron flux (Los Alamos National Laboratory satellite data) and a ground-based ULF index to remove a 24 hr noise signal. This results in smoother time series appropriate for short-term (approximately < 1 week) correlation and observational studies. However, spectral subtraction may not remove longer cycles such as the 27 days and 11 yr solar cycles. A differencing transformation ($y_t - y_{t-24}$) removes not only the 24 hr noise signal but also the 27 days solar cycle, autocorrelation, and longer trends. This results in a low correlation between electron flux and the ULF index over long periods of time (maximum of 0.1). Correlations of electron flux and the ULF index with solar wind velocity (differenced at $y_t - y_{t-1}$) are also lower than previously reported (≤ 0.1). An autoregressive, MA transfer function model (ARIMAX) shows that there are significant cumulative effects of solar wind velocity on ULF activity over long periods, but correlations of velocity and ULF waves with flux are only seen over shorter time spans of more homogeneous geomagnetic activity levels.

Plain Language Summary Relationships between space physics processes are often based on correlations. However, variables following the same cycles or trends may show a spurious correlation that has nothing to do with their physical relationship. In space weather data, these common cycles may result from satellites orbiting the Earth daily, or the 27 days or 11 yr activity cycles of the Sun. The daily cycle can be removed using noise reduction techniques similar to that used to clean audio data. Differencing (subtracting the previous observation) can also remove both short-term cycles and longer trends. However, we find that an autoregressive-moving average time series model (ARIMAX) most successfully removes cycles and trends and allows the actual correlations between variables to be measured. Using ARIMAX models, we confirm that there are cumulative effects of solar wind velocity on the ULF index, but little correlation between either velocity or ULF waves with electron flux over long periods of time. This argues for limiting correlational studies to periods of constant geomagnetic activity.

1. Introduction

Periodicity is common in space weather data, with cycles ranging from diurnal (24 hr), the 27 days solar rotation, and to the 11 yr solar cycle. Parameters will be highly correlated if time series variables cycle in parallel, or if both follow a long-term trend, but this correlation may say nothing about the more immediate, causal relationship between them. As time series can be thought of as compositions of trends and cycles, it is possible to remove these by subtraction once they are identified (Hyndman & Athanasopoulos, 2018).

Both geosynchronous and ground-based measurements may depend on magnetic local time and thus exhibit diurnal variations. Although high energy electron fluxes remain mostly stable along any given drift shell, satellites at geosynchronous orbit, due to the asymmetric dipole of the Earth's magnetic field, do not stay within the same drift

shell or at constant geomagnetic latitude. For this reason, electron flux data collected at geosynchronous orbit show a diurnal cycle, typically with much higher levels on the dayside where the field is compressed and lower flux levels on the night side where the fields are stretched (Boynton et al., 2019; O'Brien & McPherron, 2003).

Various means of removing this diurnal cyclicity have been used, including daily averaging (Mathie & Mann, 2000; Reeves et al., 2011; Simms et al., 2016, 2018), or using only the maximum flux following a storm (Simms et al., 2014). However, fast injection and/or acceleration of electrons may occur in less than 24 hr (e.g., Reeves et al., 1998). Models using daily averaged data will miss the short-term processes that lead to these electron enhancements. This complicates the comparison of electron flux levels during and following geomagnetic storms. A moving average (MA), additionally, is also complicated by the fact that it will lag behind any trend that occurs in the time series. Even a small trend will result in a MA that is consistently above or below actual values in the region it seeks to represent (Hyndman & Athanasopoulos, 2018).

Other approaches have involved modeling each 24 hr MLT bin separately (Boynton et al., 2020), creating a noon proxy, using various magnetospheric inputs such as K_p , Dst, AE, and ULF wave power to probabilistically generate what a flux measurement would be if it occurred at noon (O'Brien & McPherron, 2003; Su et al., 2014), asynchronous regression (O'Brien et al., 2001), or applying the Kalman filter algorithm (Rigler et al., 2004). (Despite its name, the Kalman filter is not a filter in the signal processing sense. It does not remove a known nuisance frequency and allow others to pass through. Instead, it estimates signal within noisy data (Brown and Hwang, 2012). Thus, Kalman filtering should not be confused with the removal of a noise signal that we describe below. Similarly, although it depends on the state space variables of a previous time step, it is not a regression ARIMA model, which we present later.)

Another approach would remove just the nuisance cycle (such as the diurnal pattern), leaving the rest of the electron flux time series intact. This could be accomplished by subtracting an appropriate sine wave (at a 24 hr cadence). This method has been applied to a ULF ground index (Borovsky & Denton, 2014). However, as we show below, this introduces further cyclicity if the frequency needing to be removed is a strong signal.

Removing a single sine wave is a simple version of spectral subtraction or bandstop filter. Filters can be generated by calculating the Fourier transform of a time series, then subtracting the frequencies of the noise band from this frequency domain (Oppenheim & Schaffer, 2010). However, removing only a single sine function results in a sharp transition between filtered and non-filtered frequencies which in turn leads to unwanted oscillations in the resulting filtered series. To avoid this problem, a steep transition between filtered and unfiltered frequencies must be widened and smoothed to reduce the sharp switch between the two areas.

Differencing is another means of removing cyclical patterns from time series data (Makridakis et al., 1998; Simms et al., 2021). In the simplest case, subtracting the previous observation from the current observation ($y_t - y_{t-1}$) creates a new time series describing the change in the variable from one time step to the next. This effectively removes a linear trend from a time series. A difference of $y_t - y_{t-24}$ would remove a diurnal signal (from hourly data). In practice, differencing to remove one cycle ("seasonality"), such as the $y_t - y_{t-24}$ differencing, often removes trends and longer cycles as well. Thus, a single "seasonal" differencing often removes both the nuisance cycle and longer trends and cycles. (If needed, a second differencing could be done.) This removes not only the correlation between variables cycling at the same frequency, but also the correlation due to both deviating from a stationary mean or to a long-term cycle (Granger & Newbold, 1974; Hyndman & Athanasopoulos, 2018).

A disadvantage of simple differencing is that the change in a parameter may not be the best measure if actual values are needed. A differenced series does not give us a good visual picture of short-term events. If, for example, a superposed epoch analysis were desired, removal of the diurnal signal by spectral subtraction (using a bandstop filter) might be the better alternative.

Rather than using only differencing, more information may be obtained from a full ARIMAX model: autoregressive (AR), integrated (differenced), MA, transfer function. These models describe the cyclical nature of the time series of interest using autoregressive and MA terms, differencing the time series if necessary to produce a stationary time series in which the mean does not change over time. Short-term cycles may best be described by AR and MA terms, although differencing at this time step may also be helpful. However, for long-term cycles, such as the 27 days and 11 yr solar cycles, differencing may be the only viable way to remove them. (For these longer cycles, many more years of data would be necessary to model them strictly with AR and MA terms.) Input

variables (independent predictors) may then be added to study their effects on the dependent variable independent of the cycles and trends (Hyndman & Athanasopoulos, 2018; Pankratz, 1991; Simms & Engebretson, 2020).

In this study, we compare the output of several methods of removing the diurnal signal from two different time series: the high energy (1.8–3.5 MeV) electron flux from one Los Alamos National Laboratory (LANL) satellite at geosynchronous orbit and a ULF index (Kozyreva et al., 2007; Pilipenko et al., 2017). We also explore their correlation with differenced solar wind velocity, another strongly correlated parameter. We first identify the noise frequencies using spectrograms, then attempt to remove these using both bandstop filters and differencing. We also fit ARIMAX models to show the cumulative effect of one parameter on another. In all these cases, we find that removing the cycles and long-term trends by differencing greatly reduces the correlation between variables. Removal of only the diurnal cycle using a bandstop filter does not correct for correlations due to long-term trends, but it can aid in the study of shorter term events. Fitting ARIMAX models provide a better method for looking at the cumulative, lagged effects of one parameter on another.

2. Data

We analyze three kinds of observational data to compare the effects of these methods of removing unwanted or confounding signals. The first two are observables within the magnetosphere (electron flux and ULF waves), while the third is external (solar wind velocity) which contains solar periodicities but not diurnal variation.

We use hourly averaged \log_{10} electron fluxes ($\log(\text{electrons}/(\text{cm}^2/\text{s}/\text{sr}/\text{keV}))$) in the 1.8–3.5 MeV range from the LANL Energetic Spectrometer for Particles instrument located at geosynchronous orbit on the 1994-084 satellite. We limit the period to 25 October 1995 to 13 June 2002 when there was a minimum of missing data. For spectral analysis and bandstop filtering, short periods of missing values were interpolated from surrounding values.

We use the Kozyreva et al. (2007) \log_{10} (ULF Pc5 index) which provides an hourly measure of Pc5 (2–7 mHz) ULF power (in nT^2/Hz) observed in the local time range from 05:00 to 15:00 hr by ground-based magnetometers stationed between 60° and 70°N corrected geomagnetic latitude.

Hourly averages of the \log_{10} solar wind velocity (km/s), from the OMNIweb database are used as a comparison, as these provide a time series external to Earth's magnetosphere.

Analyses and filtering were performed in MATLAB. Spectrograms were obtained from the MATLAB spectrogram procedure. This computes the (Hanning) windowed discrete-time Fourier transform (FFT) of a signal using a sliding window of 256 points and a 50% overlap (128 points). We used a sampling frequency of 24 hr/day.

ARIMAX models are developed using the SPSS TSMODEL procedure. We chose a parsimonious model for each dependent variable, adding AR, MA, and differencing terms until the partial autocorrelation function contained no significant terms. Using the ARMA terms and differencing chosen by this procedure, we then added influence and decay terms for the independent variable until no more were statistically significant. Electron flux was fit with both autoregressive and MA terms (both at 1 hr) and daily autoregressive and MA terms (both at 1 day) with a 24 hr differencing transformation applied. When the ULF index was the dependent variable it was also fit to the same AR and MA terms with 24 hr differencing. This successfully removed the autocorrelation in both electron flux and the ULF index when either was the dependent variable.

When the ULF index was used as the independent variable instead (to predict flux) it was differenced at either 1 or 24 hr. Solar wind velocity, as a predictor, was differenced at 1 hr. While flux has a strong daily cycle that requires 24 hr differencing to remove, and while velocity is adequately differenced at 1 hr, the choice of differencing for the ULF index is not as obvious. We difference ULF at 1 hr when analyzing its response to V , for consistency between the two terms. However, we found, when studying the response of flux to the ULF, differencing ULF at 1 or 24 hr changed the influence pattern dramatically so we report both models.

When the effect of the 24 hr differenced ULF index on flux was tested, we found significant terms at lags 0, 1, 2, and 3 hr, along with decay terms that begin acting at lag 1 and 2 hr. We found no significant terms when the ULF was differenced at 1 hr. When solar wind velocity was used to predict the ULF index, we found that effects of velocity could best be modeled with statistically significant 0, 1, and 6 hr lags and a decay term at 1 hr. When solar wind velocity was used to predict flux, we found no significant influence terms.

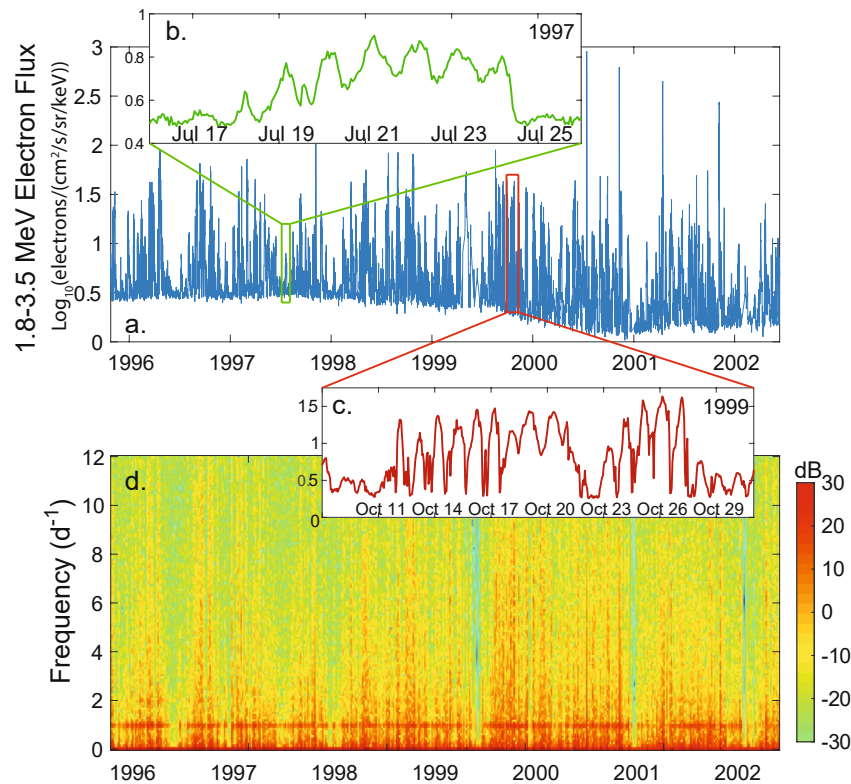


Figure 1. Electron flux (1.8–3.5 MeV) timeplot (a) over 1996–2002, (b) timeplot of quiet period. (16–25 July 1997), (c) timeplot of disturbed period (8–30 October 1999), and (d) spectrogram over 1996–2002 showing strong signal at 24 hr (1 day⁻¹) frequency with a weaker signal at 12 hr.

Although we could model the 24 hr cycle using seasonal ARMA terms, this was not possible for the 27 days cycle or changes due to the 11 yr solar cycle without a much longer time series. However, these are effectively removed using the differencing term.

In the ARIMAX models, we are able to calculate the partial correlation for adding a single predictor (both its influence and decay terms) by taking the square root of the difference of the R^2 between models with and without the predictor added, where the R^2 is the fraction of the variation in the data explained by each model (Neter et al., 1985).

3. Identifying the Signal

While it is difficult to discern the daily cycle of flux in a long time series over many years (Figure 1a), the diurnal pattern is clearly visible during undisturbed conditions over a period of a week (Figure 1b—inset panel), as well as during disturbed conditions (Figure 1c). In the frequency domain (using an FFT described above), the daily cycle of flux is clearly visible at a frequency of 1 day (Figure 1d). A less dramatic cycle is also seen at 12 hr (at a frequency of 2 times/day) in this spectrogram.

A daily pattern is less visible to the eye in the ULF ground index (Figure 2a, including inset panels 2b and 2c). There is, however, a weak diurnal signal in the frequency domain (Figure 2d). In addition, there appear to be cycles at 12, 8, 6 hr, etc. (respectively, 2, 3, and 4 cycles/day.)

4. Removing the Signal by Detrending

4.1. Subtracting Only the Diurnal Sine Wave

Subtracting off only the sine wave associated with the diurnal cycle in the LANL flux data does not accomplish the desired goal. Instead, this introduces an even stronger frequency, seen in the spectrogram near 1 day

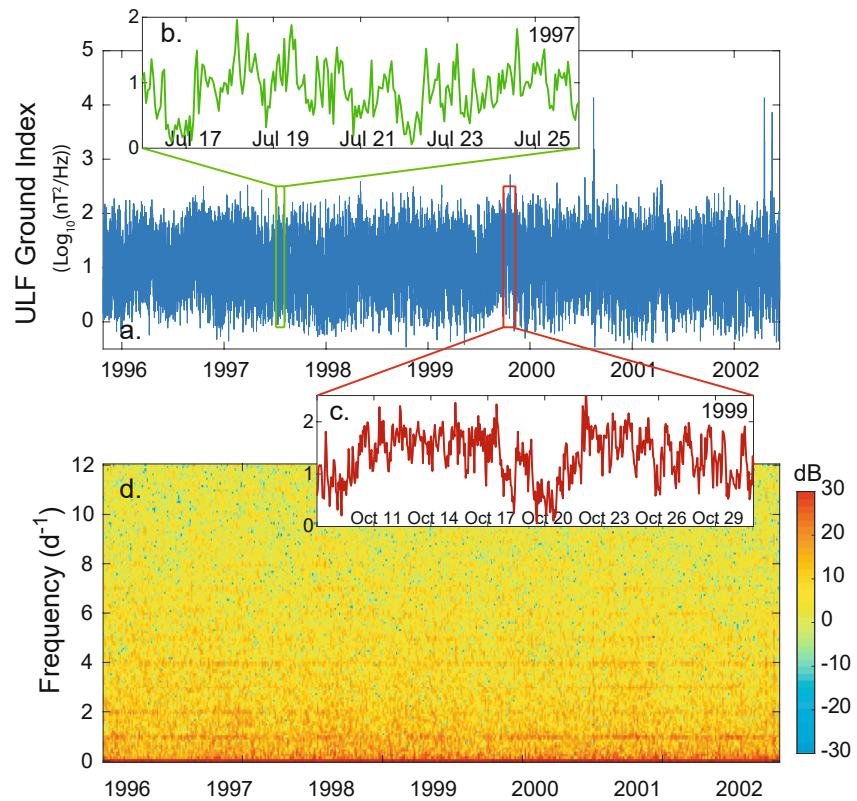


Figure 2. ULF index timeplot (a) over 1996–2002, (b) timeplot of quiet period (16–25 July 1997), (c) timeplot of disturbed period (8–30 October 1999), and (d) spectrogram over 1996–2002 showing weak signals at 24 hr (1 day⁻¹), and 12, 8, and 6 hr (1/2, 1/3, and 1/4 days⁻¹).

(Figure 3a). The introduced signal is stronger than the one we were attempting to remove, although slightly offset from the original (Figure 3b).

Subtracting only the 24 hr sine wave from the ULF index accomplishes little. The 1-day cycle is still clearly visible in the spectral frequency domain (Figure 4a). However, because it was a weak signal to begin with, this subtraction did not introduce a new signal of any great strength (Figure 4b).

4.2. Spectral Subtraction

The nuisance signal of the diurnal pattern can be removed from a time series with the use of a bandstop filter. The Fourier transform of the time series is calculated, the frequencies of the noise band are filtered from this frequency domain, and then the data is converted back to the time domain using an inverse Fourier transform. However, a sharp transition (“brick wall”) between stopband and passband frequencies in the filter can result in the introduction of oscillations (“ringing”) in the output. This explains the introduction of an oscillation (Figure 3) when only the single sine wave at 24 hr is subtracted from the signal. The solution is to smooth the “brick wall” sides of the filter by reducing the steepness of the transition from passband to stopband (Oppenheim & Schaffer, 2010).

We design an equiripple bandstop filter for both electron flux and the ULF index using MATLAB (Parks & McClellan, 1972). This is a finite impulse response filter. We placed the edges of the stopband where the diurnal signal was strongest, tapering the transition from the stopband edge up to the passband edge. For the daily signal in both electron flux and the ULF index, the stopband was between 23.95 and 24.05 hr⁻¹. The transition then tapered up to the full passbands on either side (21.82 and 26.67 hr⁻¹). For the half-day cycle, the stopband was between 11.99 and 12.01 hr⁻¹, rising to the passbands of 11.43 and 12.63 hr⁻¹. Further bands were removed at 8 and 6 hr⁻¹ from the ULF index (stopbands: 7.99–8.01, 5.99–6.01 hr⁻¹; passbands: 7.74, 8.28 and 5.85, 6.15 hr⁻¹, respectively).

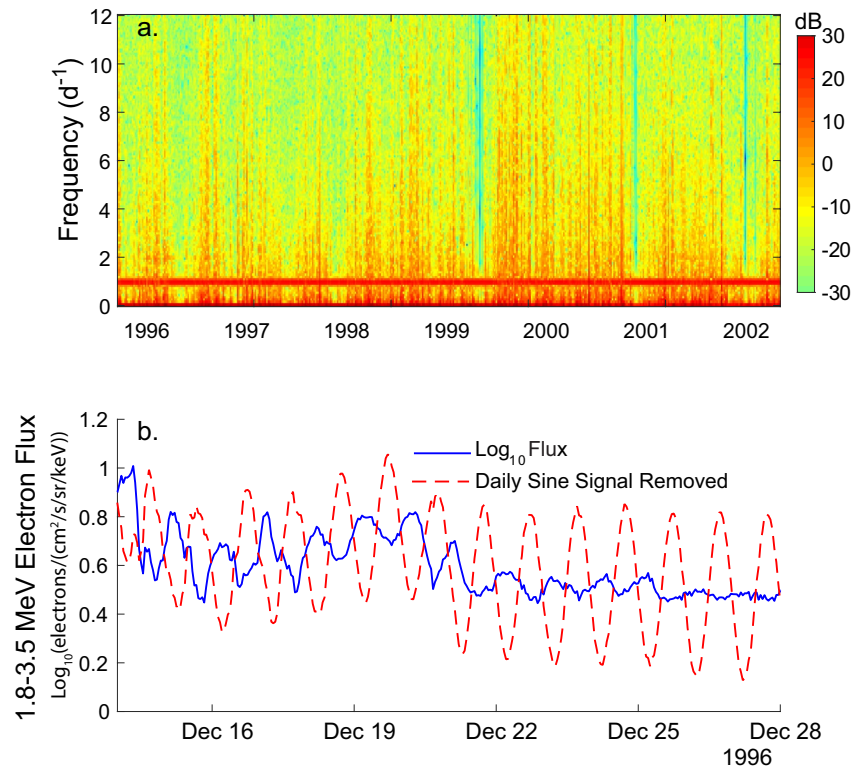


Figure 3. (a) Spectrogram of electron flux with daily sine signal removed showing the introduction of an even stronger diurnal signal. (b) Timeplot of 15–27 December 1996 period original flux (blue line) and series with daily sine wave subtracted (dashed red line).

When converted back to the time domain, the overall pattern of electron flux remains the same (Figure 5a). The two representative time periods (Figures 5b and 5c) still retain their respective shape. The diurnal pattern is effectively removed from the 16 July to 26 July 1997 quieter period. There are still peaks of flux in the 8 October to 20 October 1999 active period but they no longer follow a daily rise and fall. The spectrogram (5 days) shows that the daily and half-day signal have been removed.

Converting the ULF index back to the time domain results in a time series even more similar to the original data (Figure 6a). The July 1997 and October 1999 examples (Figures 6b and 6c) show only minor differences from the original. However, the spectrogram shows that these cycles have been successfully removed (Figure 6d).

It may seem troublesome that removing the diurnal cycle leaves a visible gap in the spectrograms (which plot in the frequency domain). However, while it may be supportable to interpolate in the time domain (otherwise the electrons would seemingly disappear), we do not have a similar common sense reason to interpolate in the frequency domain. While we can reasonably extrapolate that electrons must have been present during a gap in time, we have no reason to assume that there was still a cycle at the frequency we remove from the frequency domain. In fact, the gap left in the frequency domain is the interpolation in the time domain, and we are only able to admit that we cannot collect data at that frequency due to the strong diurnal signal due to the satellite's orbit. However, this is not a significant problem as we are mostly interested in the behavior of electrons over time and are removing frequencies only to see that time dependent behavior more clearly.

5. Removing the Signal by Differencing

An issue that arises when comparing time series is that either an overall trend over a long time period or predictable cyclical behavior can result in a high correlation between the two variables, even if there is no correlation at all on the smaller scales that are of actual interest (Hyndman & Athanopoulos, 2018).

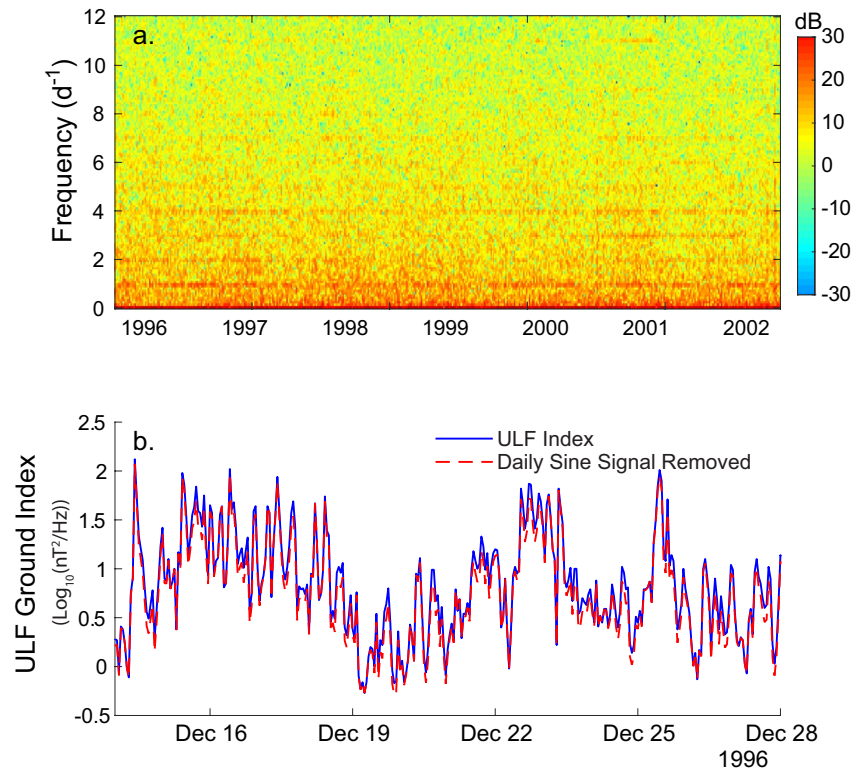


Figure 4. (a) Spectrogram of ULF index with daily sine signal removed showing the introduction of an even stronger diurnal signal, (b) timeplot of 15–27 December 1996 period original flux (blue line) and series with daily sine wave subtracted (dashed red line).

Electron flux, the ULF index, and solar wind velocity all track the changes in geomagnetic activity that occur over months and years, thus, we would expect a strong correlation between these variables solely on that basis. However, this correlation is of no interest when studying whether an increase in ULF wave activity or velocity drive electron acceleration in the short term. In addition, the strong 24 hr periodic cycle exhibited by the LANL satellite data, together with the weaker diurnal cycle in the ULF wave index will also result in a strong correlation between these two variables that is unrelated to any actual physical relationship.

An easy, if non-intuitive, solution to this problem is to detrend the data using differences. The simplest difference subtracts the previous observation from the current value ($y_t - y_{t-1}$). However, other differences are just as powerful at removing the overall trends and may have the added benefit of removing cyclicity at longer time lags (Hyndman & Athanasopoulos, 2018; Makridakis et al., 1998). The difference transformation over 24 hr ($y_t - y_{t-24}$) on the electron flux data removes both the overall trend (centering the time series at 0) as well as the 24 hr cyclicity (Figures 7a–7d). The disadvantage of this method is that we are no longer measuring the actual value of flux, but the change. While this will remove the unwanted correlation associated with the overall trend and daily cyclicity, over shorter time periods we may not find the resulting time series of “change” to be a satisfactory representation of “flux levels”. Additionally, the 24 hr differencing also removes information at all higher frequencies (Figure 7d: 1–12 hr⁻¹). The sharp rise at the end of the first example time period (Figure 7b) that is not present in either the original data (Figure 1b) or the spectral subtraction data (Figure 6b) suggests that we cannot interpret the differenced data in the same way we would interpret flux level.

However, there are further reasons for differencing time series data. Differencing can remove autocorrelation of variables, which can inflate correlation coefficients. It can also successfully remove cycles at periods other than the differencing period. Differencing of both the electron flux and the ULF index by 24 hr ($y_t - y_{t-24}$) removes not only the diurnal cycle but autocorrelation of each variable and the 27 days solar cycle (Figures 8a and 8b; note the difference in the right and left y- axis scales). Each of these undifferenced variables (blue line) shows the diurnal cycle, considerable autocorrelation (>0.75 out to 60 days), and a stronger autocorrelation at the 27 days

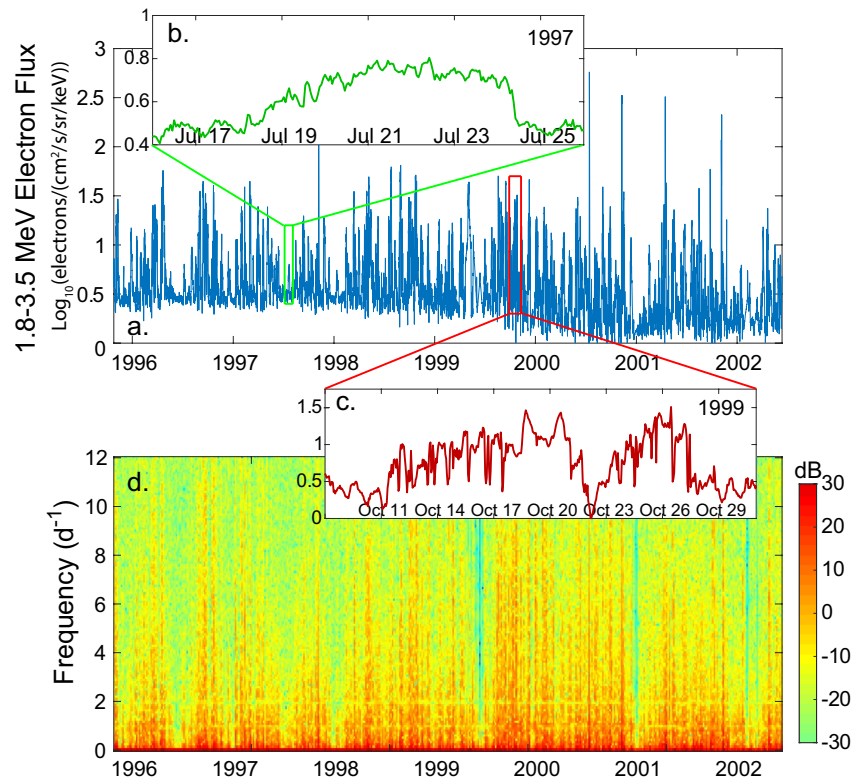


Figure 5. Electron flux (1.8–3.5 MeV) with daily and half-day cycles removed using a bandstop filter. (a) 1996–2002, (b) a quiet period (16–25 July 1997), (c) a disturbed period (8–30 October 1999), and (d) spectrogram over 1996–2002 showing strong signal at 24 hr (1 day^{-1}) frequency and weaker 12 hr signal have been removed.

solar cycle. When differenced at 24 hr (red line), both diurnal and 27 days cycle are removed, and autocorrelation is reduced to near 0. One hour differencing of these two variables (orange line) removes the 27 days cycle and the autocorrelation but leaves traces of the diurnal cycle (although still at < 0.2 correlation). As a comparison, solar wind velocity shows a high degree of autocorrelation (> 0.9) and the expected 27 days cycle, but, as expected, no diurnal cycle (blue line of Figure 8c). Differencing the velocity by either 1 or 24 hr removes both the autocorrelation and the 27 days solar cycle.

6. Using the Transformed Data

6.1. Time Series Simple Correlations

Log_{10} electron flux data (using all years) without the diurnal signal removed shows a varying simple correlation with the ULF index lagged over several days (Figure 9a). The peak of correlation with the unaveraged ULF index (red line; $r = 0.30$) occurs at 62 hr. The timing of this peak using the unaveraged ULF index could be wrong, because we are unable to see what occurs between the peaks due to the daily cycle. However, the bandstop filtered series show the same correlation (dashed line; $r = 0.30$ near 62 hr). We do note that the apparent correlation can be increased by using a 24 hr MA of both the ULF index and the electron flux (12 hr either side of the observation) showing a smooth pattern of correlation rising to an apparent peak when ULF power is measured 66 hr before flux (blue line; $r = 0.45$). This can be interpreted as 1 unit increase in the standardized ULF index leading to a 0.45 increase in the standardized flux. (Standardized variables can be obtained by subtracting the mean of the series from each observation and then dividing by the standard deviation of the series. In the special case of regression with a single predictor, the standardized regression coefficient will be the same as the correlation coefficient.)

However, this shows that removing only the diurnal cyclicity (by bandstop filtering) in the two time series is not enough to remove the overall trend and cyclicity of the time series. It is also possible that any association between

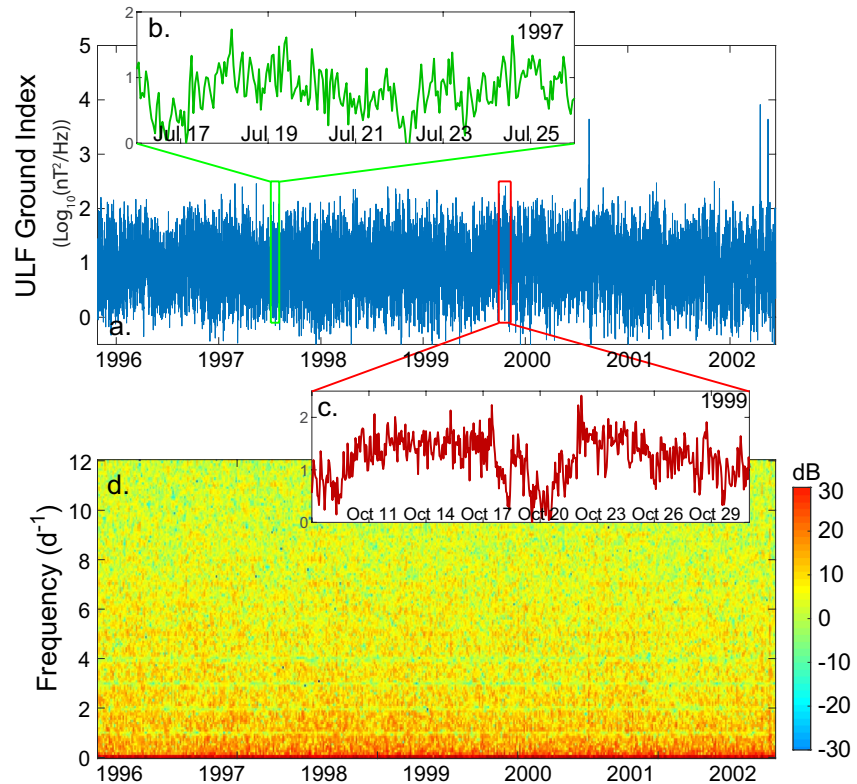


Figure 6. ULF index with daily and faster cycles removed using a bandstop filter: (a) 1996–2002, (b) a quiet period (16–25 July 1997), (c) a disturbed period (8–30 October 1999), and (d) spectrogram over 1996–2002 showing signal at 24 hr (1 day^{-1}) and weaker signals at 12, 8, and 6 hr ($1/2$, $1/3$, and $1/4 \text{ days}^{-1}$) have been removed.

ULF wave activity and electron flux only occurs at particular moments in time, not continually, and that correlating over long periods obscures this relationship. While there may be a noticeable pattern of concomitant rise in ULF activity followed by electron flux increases at particular moments, most likely geomagnetically disturbed periods, correlating the two variables over long periods of mostly low activity does not provide a detailed picture of that relationship. A long-term correlation such as this is likely only useful in the context of predictive models that also incorporate a measure of geomagnetic activity and correlational terms that change depending on that activity.

Although averaging may remove random variation and therefore result in a higher correlation, it comes at the expense of any information about the short-term relationship, and, more importantly, it does not remove the longer periodicity or trend that may contribute to inflated correlations. Differencing both time series (flux and the ULF index) by 24 hr results in a peak negative correlation in the first few lagged hours (black line; $r = -0.21$). Later lags show positive correlations (0–0.11 when lagged over 24–48 hr). For the differenced series, a 1 unit increase in standardized ULF change would lead to a 0.11 increase in change of the standardized flux. As cyclicity and long-term trends are removed, these correlations address the actual influence of ULF waves on flux and not merely the response of flux and ULF waves to the same external drivers.

As the solar wind is not influenced by Earth's rotation, its velocity does not show diurnal cycling and the correlation of its time series with that of electron flux does not show the peaks associated with this (Figure 9b). The highest correlation of the original time series of flux and V is 0.43 at 48 hr (red line). A moving 24 hr average of each series results in a slightly higher correlation (blue line; $r = 0.48$ at 51 hr). As there is no tandem diurnal cycling to correct, the increase in correlation when using 24 hr averaged flux data is smaller. However, the differenced series remove the longer 27 days solar cycle and result in a very low correlation (black line; maximum $r = 0.06$ at 41 hr).

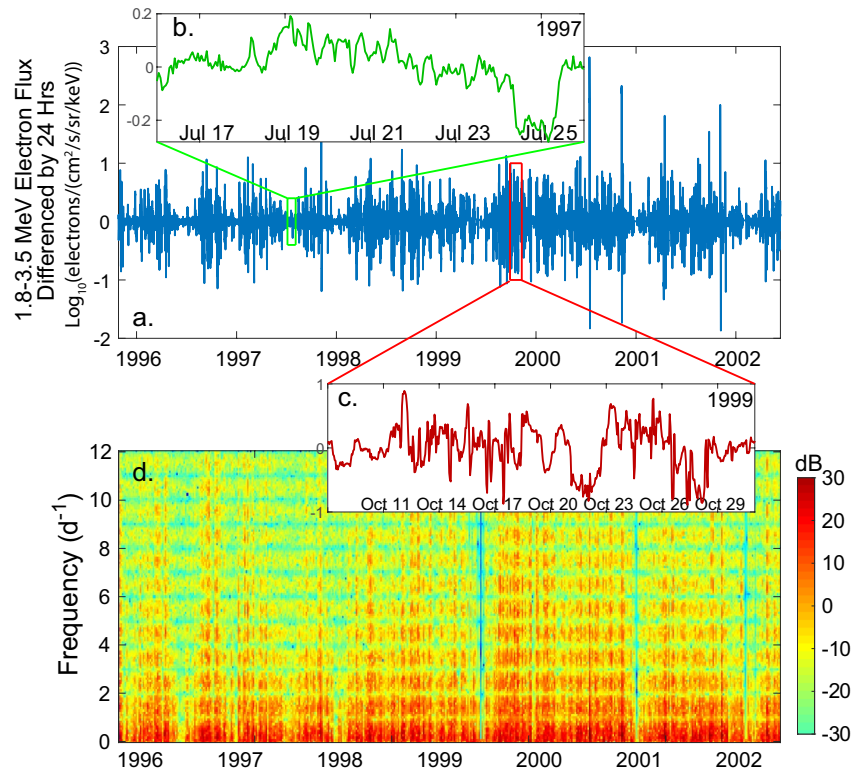


Figure 7. Electron flux (1.8–3.5 MeV) differenced by 24 hr ($y_i - y_{i-24}$): (a) 1996–2002, (b) a quiet period (16–25 July 1997), (c) a disturbed period (8–30 October 1999), and (d) spectrogram over 1996–2002 showing strong signal at 24 hr (1 day^{-1}) frequency and weaker 12 hr signal have been removed, as well as cyclicity at higher frequencies.

Although strong correlations have been found between ULF wave activity and solar wind velocity, differencing these two time series to remove the 27 days solar cycle and other trends results in a correlation of 0.10 when measured simultaneously (Figure 9c; black line), dipping to a negative, but also inconsequential, correlation at 24 hr ($r = -0.03$). This is in sharp contrast to the correlation between the uncorrected series (red line; $r = 0.45$). As might be expected, 24 hr averaging does not remove the correlation due to other trends and cycles and results in the higher correlation reported by other studies (blue line; $r = 0.65$).

The lack of relationship between these three variables when differenced is shown most dramatically in scatter plots (Figure 10). The explanatory variable (x -axis) is differenced by 24 hr (ULF index) or by 1 hr (V) then lagged by 1 hr. This lag shows the strongest correlation, however, to the eye, there appears to be little association between the variables when autocorrelation and both diurnal and 27 days cycles are removed by differencing. The correlations in the scatter plots of Figure 10 are equivalent to the solid black lines of Figure 9 at the single point where a 1 hr lag has been applied: (a) is slightly negative, (b) near zero, and (c) slightly positive. Not only have linear associations almost disappeared, but there is also no suggestion of a nonlinear component. Note that if the scatterplots were at a 24 hr lag these correlations would still be near zero for all three, but slightly positive for (a) and (b) and slightly negative for (c).

The most noticeable effect of the bandstop filter applied to electron flux and the ULF index is to smooth out the diurnal cycle (dashed line; Figure 9a). It does not correct for the 27 days solar cycle or longer trends, and, therefore, the correlation remains high between all three pairs of variables. Unfortunately, the 27 days solar cycle could only be successfully removed using a bandstop filter if the time series were significantly longer.

This drop in correlation between an original data series and its differenced data series is more noticeable when using hourly averaged data. We compare correlations of solar wind velocity with the ULF index between a more active period (the last half of 1993, as studied by Engebretson et al., 1998) and a quiet period (the last half of 1997). Daily or hourly averaged, undifferenced correlations differ only somewhat between these two periods, with that of the quiet period being slightly lower (Table 1a and 1c). Differencing daily averaged data results

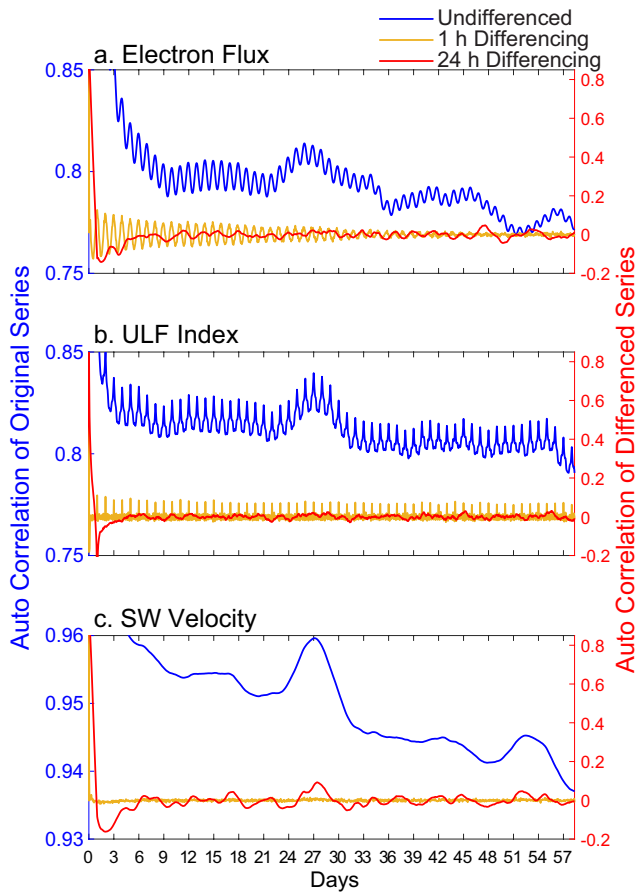


Figure 8. Autocorrelation of (a) electron flux (1.8–3.5 MeV), (b) ULF index, and (c) solar wind velocity. The high autocorrelation in each of the original time series (blue lines; left hand blue axis) can be removed by differencing at either 1 hr ($y_t - y_{t-1}$; orange line and right hand red axis) or 24 hr ($y_t - y_{t-24}$; red line and right hand red axis). Either differencing also removes the 27 days solar cycle. Differencing by 24 hr removes the diurnal cycle from both electron flux and the ULF index, but 1 hr differencing does not.

in lower correlations (Table 1b), but differencing the hourly averages effectively removes any correlation (Table 1d). This suggests that much of the correlation in hourly data is the result of common cycles and that the random variation otherwise is so high as to obscure any association. This might argue for using only daily averaged data, but if processes act more quickly than at a daily cadence, we must find ways to effectively study relationships at a shorter time step. One approach is to model these relationships as ARIMAX time series transfer function models (see Section 7).

6.2. Events

Bandstop filtering of the diurnal signal may be the best choice for analyzing individual events where actual values of flux and ULF activity are more pertinent to the questions being explored. Over short periods of time (less than the 27 days solar cycle), correlations can be performed between the filtered, but undifferenced, variables (diurnal signal removed) without the solar cycle or other long-term trends inflating the correlation coefficients.

In lineplots, the relationship between electron flux and the ULF index is somewhat obscured by the 24 hr cycling of both (Figure 11a; 26 October to 23 November 2000). The daily variation in flux is effectively smoothed by taking a 24 hr MA (Figure 11b). However, when filtered with a bandstop filter, the flux dips (at Dst minima) appear to be a more accurate representation than that seen in either the averaged or the unfiltered time series (Figure 11c). There are flux dips that do not show up in the unfiltered and averaged time series (e.g., on 29 October and 7 November), and flux highs may only appear as a spike in the unfiltered data when, in fact, they may last several hours (e.g., 9 November). Filtering has a more obvious smoothing effect on flux than on ULF activity. A shorter MA (<4 hr) may be helpful in further smoothing the ULF index.

We select three example periods to calculate short-term correlations between flux and ULF activity (Figure 12a). Over 16–24 July 1997 (quiet), 8–30 October 1999 (storm), and 26 October to 23 November 2000 (storm) periods, the highest flux-ULF power correlations using filtered data are 0.43 (November 2000 storm period at a 65 hr lag), 0.44 (October 1999 storm period at a 43 hr lag), and 0.39 (July 1997 quieter period at a 72 hr lag). When ULF power is measured 24 hr prior to flux, all correlations were <0.18, with those at less than a 10 hr time lag all <0. There is little difference in the pattern of correlation over time for these flux-ULF power correlations in these shorter time periods.

We also correlate both the filtered electron flux and filtered ULF index with solar wind velocity (unfiltered as there was no 24 hr periodicity in V). Both show differences between the storm and quieter periods. For the two storm periods, the flux-V correlation is <0.20 in the first 7 hr, rising to 0.58 (66 hr) or 0.66 (62 hr). However, the correlation in the quieter period rises to 0.71 at 25 hr, then falls off rapidly at longer lags (Figure 12b).

The pattern between storm and quieter period ULF index-V correlations is reversed (Figure 12c). The storm periods begin at a high of 0.60 and 0.50 when the variables are measured simultaneously, with correlations falling off as the lag increases. During the quieter period, the ULF-V correlation falls from 0.30 to negative values as the lag increases.

7. Cumulative Effects of Predictors: ARIMAX Models

Many phenomena are thought to act over longer periods than a single hour. A correlation at a single point in time is therefore less useful than describing the cumulative effects of a predictor over a number of hours. Transfer functions, which describe inputs and their decay, can be useful in modeling these cumulative effects. We show

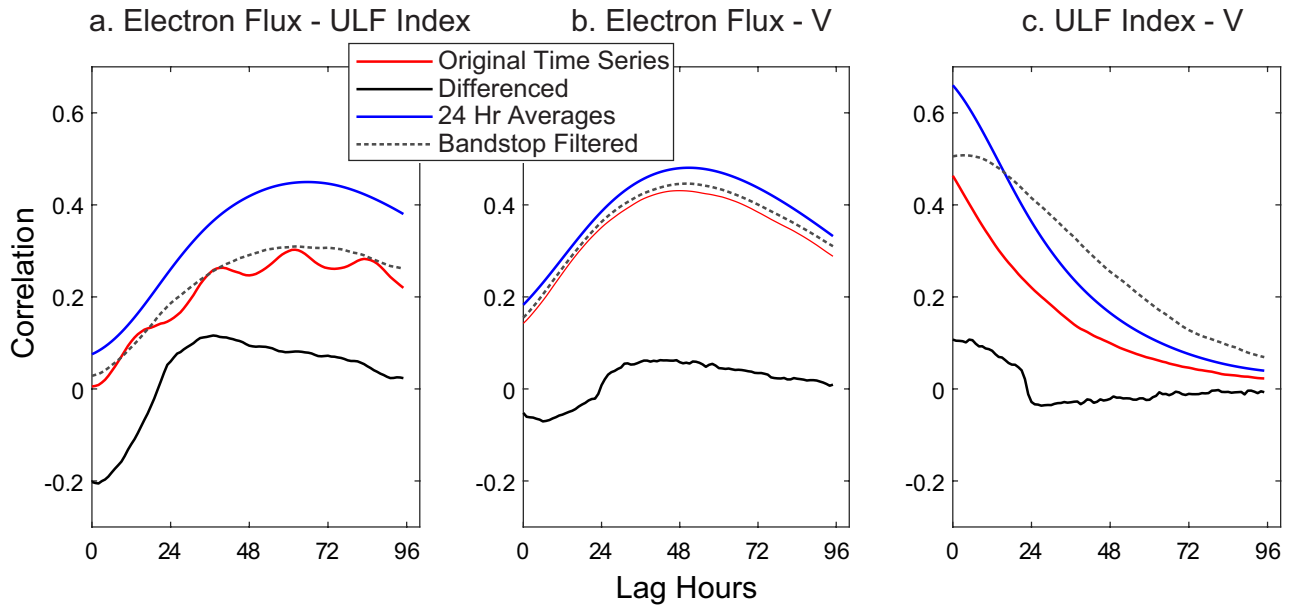


Figure 9. Correlations between electron flux, the ULF index, and solar wind velocity (V) at hourly lags over entire time series (1996–2002). The original time series (red), which may show a diurnal pattern, has lower correlations than the 24 hr averaged series (blue). Differencing by 24 hr (solid black) removes cycles and trends and results in lower correlations. Bandstop filtering (dashed line) successfully removes the diurnal cycle but not longer term trends. The bandstop filter is therefore not appropriate for long term correlation analysis.

that ARIMAX models, that remove common cycles and trends and describe the cumulative effects of predictors over many time steps, give a more complete picture of the relationships between time series variables than simple correlations. We develop (a) ARIMAX single predictor transfer function models with electron flux as the dependent variable, using either ULF wave activity or solar wind velocity (V) as the transfer function (predictor) variable, and (b) an ARIMAX model with the ULF index as the dependent variable and V as the predictor. See the Data section for details on the ARIMA terms used.

The correlation of each of these variables within the ARIMAX models can be modeled as a continual process that decays over time, testing the statistical significance of each influence and decay term within the regression model (Figure 13: note the different y-axis in each panel, and that the coefficients between them are not comparable because the units are different). As a predictor variable, the ULF index was entered at lags 0, 1, 2, and 3 hr, with decay terms at 1 and 2 hr. Attempts to enter longer lags or decay terms did not result in any further statistically significant terms. However, differencing the ULF index at 24 hr (Figure 13a: Model A) vs. 1 h (Figure 13b: Model B) resulted in a model with either all statistically significant ULF terms (A; all p -values <0.05) or none (B; all p -values >0.7). When both flux and ULF index are differenced at 24 hr (A), the ULF shows alternating positive and negative influences, with a maximum positive influence of 0.14 (at a lag of 9 hr). This predicts that a 1 unit change in the ULF index will result in a 0.14 increase in the change in flux. This alternation of influence could represent both positive and negative effects on electron flux attributable to ULF waves and might explain why the correlations with cycles removed merely by differencing results in low correlations (Figure 9). Without a more detailed model of the inputs attributable to the ULF waves, positive and negative influences sum to zero. However, the hypothesis suggested by Model A falls apart if the ULF index is differenced instead at 1 hr (Model B). In this second case, all significant effects of ULF waves disappear. In Model B, no terms in the ARIMAX model are significant and the summed influence is very low (note the very small scale on the y-axis). The choice of model, therefore, makes a sizable difference in how the data may be interpreted.

The influence of V on flux had no significant terms and therefore little influence (Figure 13c). As in Figure 13b, In order to show any effects at all, the y-axis scale must be considerably narrowed. However, there is a strong effect of V on the ULF index (Figure 13d).

All these models describe the behavior of the dependent variable reasonably well via the ARIMA terms alone ($R^2 = 0.97$ for Figures 13a–13c models, 0.64 for Figure 13d model). However, even when influence terms of

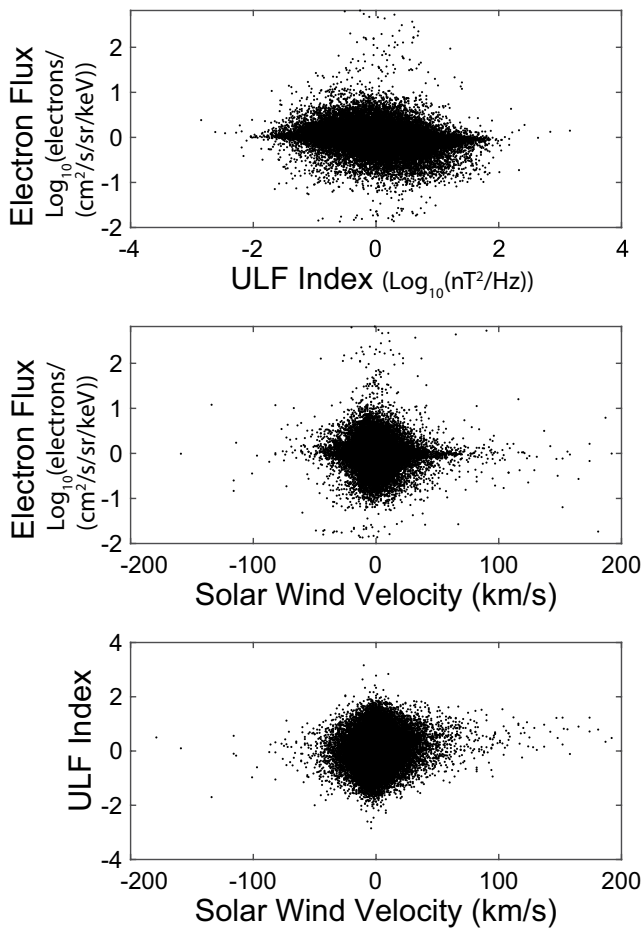


Figure 10. Lag scatterplots of the differenced time series. (a) Electron flux (differenced by 24 hr) vs. ULF index (differenced by 24 hr and measured 1 hr previous), (b) electron flux (differenced by 24 hr) vs. solar wind velocity (differenced by 1 hr and measured 1 hr previous), (c) ULF index (differenced by 24 hr) vs. solar wind velocity (differenced by 1 hr and measured 1 hr previous). There is little correlation between the differenced variables, showing that most of the correlation between the undifferenced variables is due to the common diurnal cycle and/or the 27 days solar cycle.

the predictors are significant, the additional correlation that can be attributed to the predictors is low. To test the overall effect of each parameter, rather than the individual influence and decay terms, we calculate a partial correlation coefficient by taking the square root of the difference of the R^2 between models with and without the predictor parameter. Only the V influence on the ULF index had a partial correlation above 0.1 (Figure 13d: $r = 0.104$). The greatest apparent influence of V on the ULF index occurs at a lag of 1 hr (regression coefficient of 4.12) after which the effect shows only a steady decay (Figure 13d). This is somewhat lower than the 6.52 coefficient previously found for daily averages (Engebretson et al., 1998), but is not strictly comparable due to the differencing of the two variables. Although the partial correlation due to the V influence on ULF is higher than for our other ARIMAX models, it is still much lower than that found in other studies that did not account for common cycles and trends (generally between 0.7 and 0.8).

A multivariate model that used both V and the ULF index to simultaneously predict electron flux showed only a statistically significant, and small, effect of V without any effect of ULF waves at all. This seemed an unlikely description of the system as the ULF wave processes are thought to be the physical drivers of electron flux changes while the solar wind velocity is merely a correlate. Consequently, we do not report results from this model.

8. Discussion

Cycles and trends are common in space weather data, where energy inputs may vary over the course of days (e.g., the 27 days solar cycle) or years (the 11 yr solar cycle). Measurements may also show marked daily cycles due to dependence on magnetic local time. Unfortunately, correlations will be inflated between two parameters that both show a common cycle or trend even if there is no real association between them. In the present study, we find simple correlations between electron flux, a ULF index, and solar wind velocity are much lower if cycles are removed. This would suggest that the presumed physical associations between these variables are not strong. However, an ARIMAX transfer function model, which corrects for cycles in both parameters, shows significant and cumulative effects of the solar wind velocity on the ULF index.

We explore several methods to remove the diurnal cycle from electron flux measured at geosynchronous orbit and from a ground-based ULF index, both hourly averaged. These methods include a 24 hr MA, spectral subtraction of the specific 24 hr frequency, a bandstop filter with smoothly sloping sides (reducing the steep transition from passband to stopband), a differencing transformation ($y_t - y_{t-24}$), and ARIMAX models. The bandstop filter, the differencing transformation, and the ARIMAX models successfully removed the diurnal cycle.

8.1. Moving Averages

We find that a MA increases the apparent correlation between these variables, likely by smoothing out the variation in both time series. However, it obviously does not remove the diurnal cycle (Figure 11b). Correlations using a 24 hr average will therefore still be artificially inflated due to flux and ULF power both cycling simultaneously at a 24 hr diurnal frequency. As both these parameters also rise and fall simultaneously over the 27 days and 11 yr solar cycles, correlations between them over long periods may mostly reflect these longer cycles. We find that using MAs results in a higher correlation

Table 1

Comparison of Solar Wind Velocity and ULF Index Correlations Between a Period of High Speed Solar Streams (Last Half of 1993) and a Quiet Period (Last Half of 1997), Showing That Differencing Hourly Averages Results in a Stronger Reduction in Correlation Than Differencing Daily Averages

	a. Daily average	b. Differenced daily average	c. Hourly average	d. Differenced hourly average
1993	0.675	0.472	0.487	0.024
1997	0.638	0.535	0.437	0.072

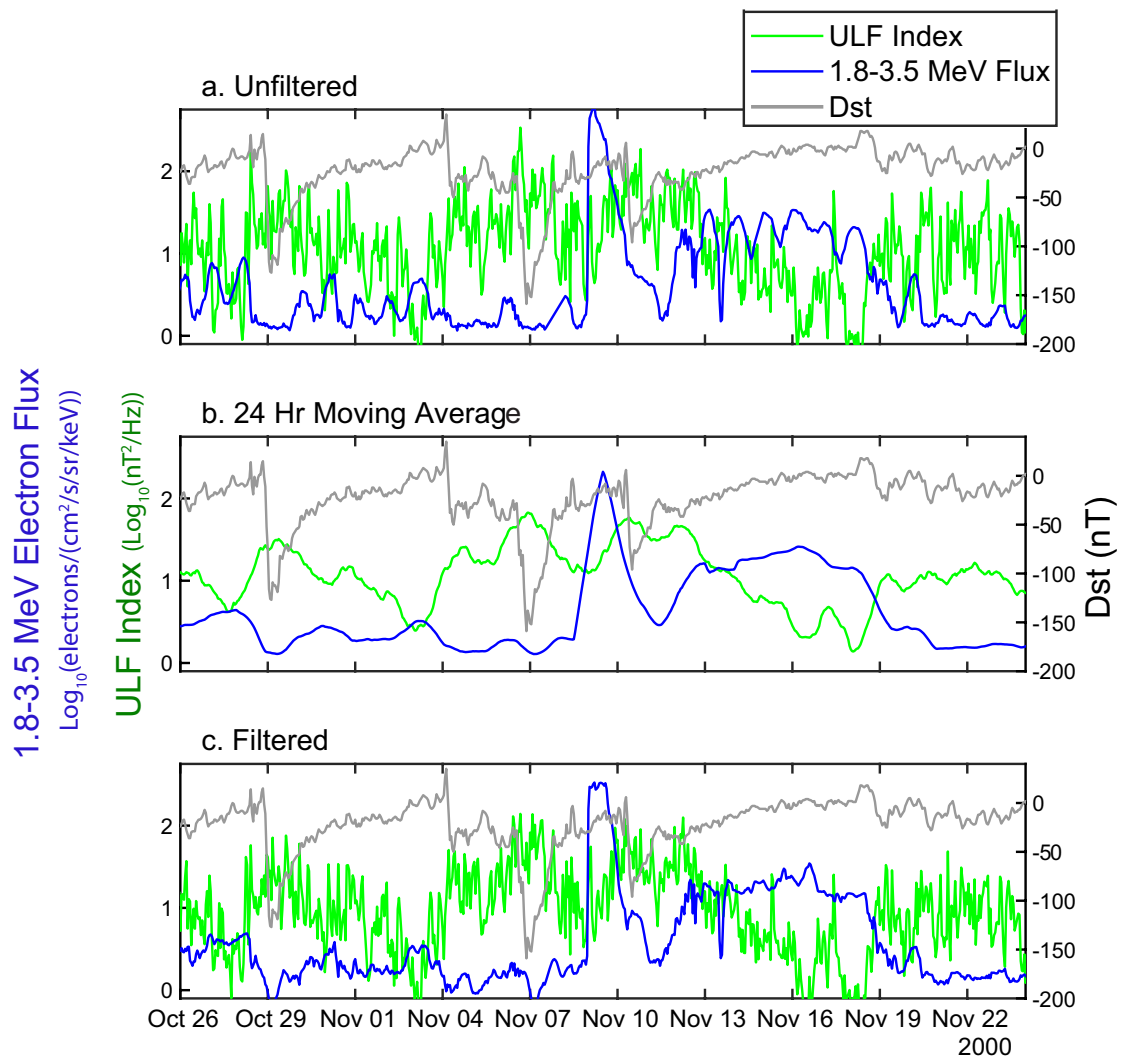


Figure 11. Timeplots of one month (26 October to 24 November 2000). (a) Original ULF index (green) and electron flux (blue) time series, with Dst (gray). (b) 24 h moving average of ULF and flux smooths the time series. (c) Bandstop filtering of ULF and flux removes the diurnal cycle.

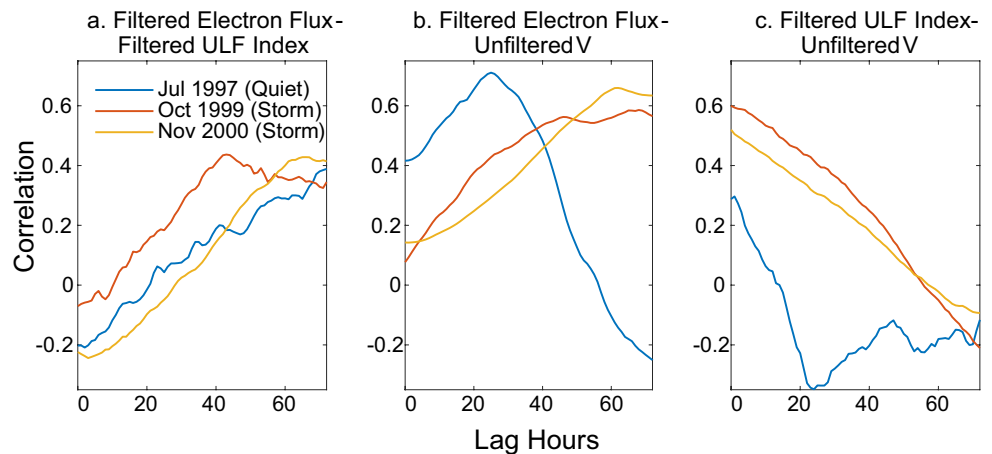


Figure 12. Correlations between electron flux, the ULF index (both with diurnal cycle removed using a bandstop filter), and unfiltered solar wind velocity (V) over 3 short time periods.

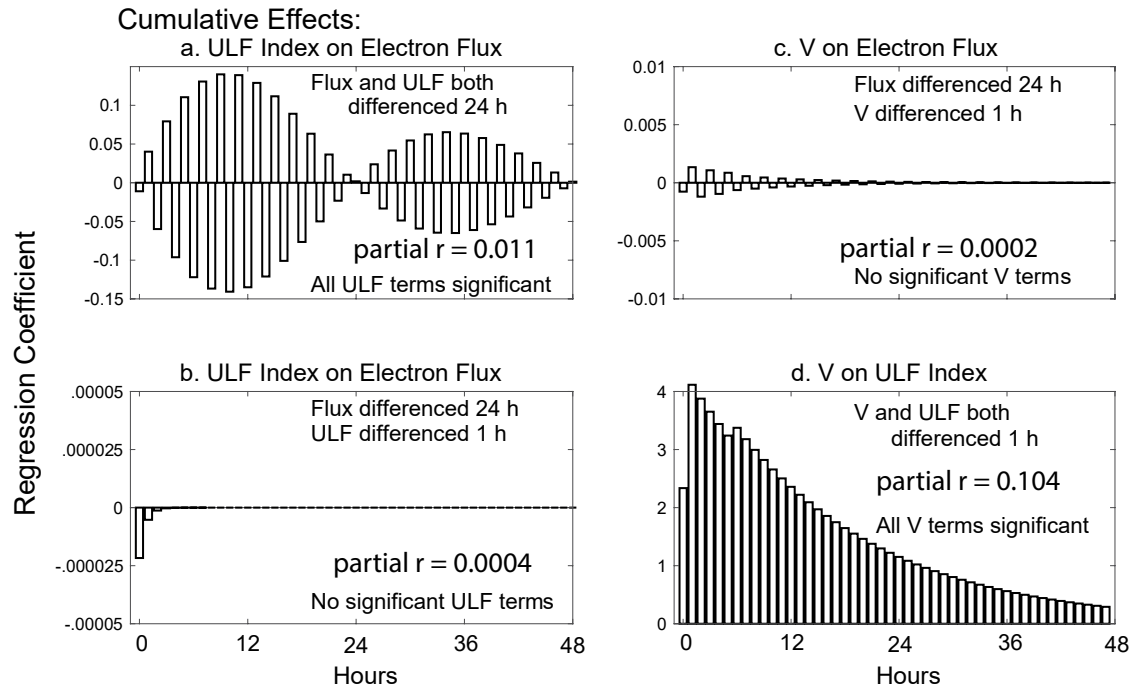


Figure 13. Cumulative lag effects from ARIMAX models: (a) the ULF index on electron flux (Model A: both differenced at 24 hr), (b) the ULF index on electron flux (Model B: flux differenced at 24 hr, ULF differenced at 1 hr), (c) solar wind velocity (differenced at 1 hr) on electron flux (differenced at 24 hr), and (d) V on the ULF index (both differenced at 1 hr). Partial correlation calculated by taking the square root of the difference in the R^2 between a base ARIMA model and the ARIMAX model with added predictor terms.

of 0.42. This is in line with correlations between ULF activity and electron flux found in other long-term studies (Mathie & Mann, 2000; 0.5–0.7; Kozyreva et al., 2007; 0.45; Simms et al., 2016), but all are artificially inflated by these cycles. Besides this problem, a lag will be introduced by using MAs if there are long cycles or trends (Hyndman & Athanasopoulos, 2018). Additionally, a 24 hr average may not be very useful if the response of these variables is at a shorter time scale.

8.2. Filtering

Removing just the sine wave associated with the daily cycle introduced severe nuisance oscillations into the electron flux data. While these nuisance oscillations were not as strong in the ULF index filtered in this way (Borovsky & Denton, 2014), we found that it also did not successfully remove the 24 hr signal from this index (Figure 4). Introduction of strong oscillations (ringing) into a transformed series is the result of a sharp transition between stopband and passband frequencies. These “ringing” oscillations will be more troublesome if the frequency being removed is a strong signal. A better filter can be designed if the stopband is entered more gradually (Oppenheim & Schaffer, 2010).

We were able to design bandstop filters with less steep sides that avoided this problem. A bandstop filter with a gradual transition in and out of the stopband successfully removed the diurnal cycle from both flux and the ULF index. These filtered time series could be used to correlate between flux and ULF power, but only over time periods much shorter than the 27 days solar cycle. Correlation between diurnally filtered flux and ULF power over the entire study period was still 0.30, the same as for unfiltered data, indicating that the longer cycles still inflate the correlation.

Using a bandstop filter to remove the diurnal cycling, we find that the two shorter disturbed periods we studied showed similar trends of correlation of the solar wind velocity with electron flux and the ULF index. Over longer lags, there was increased correlation between flux and solar wind velocity but decreased correlation between the ULF index and velocity. The example “quiet” period differs from these, with correlations falling as the

lag between the variable measurements is increased. Correlating solar wind velocity with either flux or ULF power over both types of conditions (quiet and storm) may therefore average the net correlation response to zero. However, in correlating electron flux and the ULF index, there is more similarity in the lag correlation patterns between quiet and storm conditions.

8.3. Differencing

Using differencing, we removed not only the diurnal cycle from flux and the ULF index, but also long-term trends and cycles from these two variables and from the solar wind velocity. We find that an appropriate difference transformation also removed autocorrelation (correlation of sequential observations) which can result in erroneous conclusions about the associations between variables. Neither spectral subtraction nor daily averaging removed autocorrelation. Differencing ($y_t - y_{t-24}$) reduced high autocorrelation (>0.85 in electron flux and the ULF index) to near 0 and successfully removed both the diurnal and the 27 days solar cycles. Autocorrelation of the solar wind velocity can be removed by a $y_t - y_{t-1}$ difference transformation. This reduces the autocorrelation of V (>0.95) to near zero and, for the most part, also removes the 27 days cycle. However, differencing results in a lower correlation between electron flux and ULF wave power than has been seen in previous studies ($r = -0.21$ – 0.11).

These advantages of differencing make it a more appropriate transformation for long-term correlations than band-stop filtering. However, differencing results in a measure of change, not value, and may therefore not be suitable for some types of studies. For the study of short-term events, filtered data (with only the diurnal cycle removed by a bandstop filter) may be more appropriate.

We also explored the differenced correlations of electron flux and the ULF index with solar wind velocity. With the 27 days cycle and longer-term trends removed from all three time series, correlations of flux and ULF activity with the solar wind velocity are more reliable, but much lower than previously reported. The peak correlations between the differenced electron flux and solar wind velocity (absolute magnitudes <0.08) and the ULF index and velocity (<0.11) are now much lower than those from the undifferenced time series (0.48 and 0.65, respectively). Previous work using long time series has found a high linear correlation between ULF activity and solar wind velocity (≈ 0.7 in a period dominated by persistent high-speed solar wind streams: Engebretson et al., 1998; ≈ 0.8 : Mathie & Mann, 2001; >0.70 : Potapov et al., 2012) and between electron flux and solar wind velocity (0.72: Baker et al., 1990; 0.5–0.7: Kozyreva et al., 2007; 0.54: Potapov et al., 2012; 0.37: Reeves et al., 2011; 0.5: Simms et al., 2016; Rigler et al., 2004 with no correlation reported). Reeves et al. (2011) noted that this linear measure may have missed the nonlinear component of the velocity-flux relationship. However, in our present study, with long-term trends and cycles removed from the time series by differencing, we note that not only is the linear correlation greatly reduced, but the nonlinear component is also removed (see Figure 10). Both Baker et al. (1990) and Rigler et al. (2004) employ linear prediction filters (Wiener and Kalman, respectively). However, despite the name, neither is a filter in the signal processing sense. They do not remove a known nuisance frequency and allow others to pass through. Instead, they estimate signal within noisy data (Brown & Hwang, 2012). Thus, Kalman filtering should not be confused with the removal of a noise signal as we have done above. Similarly, although the Kalman filter may depend on the state space variables at a previous time step, it is not a regression ARIMAX model and does not base its conclusions on a firmly established hypothesis testing methodology as regression does.

These low correlations after differencing might suggest an overall lack of a causal relationship of the solar wind velocity with ULF activity and flux. However, the association between these variables is likely to differ between disturbed and quieter, more normal conditions. It would not be surprising that long-term correlation studies, mixing these two types of conditions, might not be very high if the processes in these periods act opposite from each other or because there is only a connection between the variables during disturbed periods. A preponderance of “normal” (quieter) conditions in a data set may lead to the parameters overrepresenting more typical conditions at the expense of storm period activity. This could lead to models that do not model storm periods well and may be one reason why many predictive models lag behind the large changes in electron flux that occur during and after storms (Simms et al., 2016). The different patterns of electron flux and ULF power dependence on solar wind velocity between storms and periods of more typical conditions could both be included in a single model by adding terms to distinguish between conditions.

We note that while correlations between undifferenced variables may only be high due to common trends and not an actual relationship, this does not mean that predictive models built using undifferenced variables will necessarily be of poor predictive ability. Testing the hypothesis of a physical relationship between variables vs. creating predictive models are two very different goals. For hypothesis testing, it is important to know that the correlation is based on the physical relationship of interest and not due to common cycles shared by both variables. For predictive models, this is not as much of a concern if predictions are shown to be accurate. However, it is important to make the distinction between correlated variables that are exhibiting a physical or causal relationship (determined from hypothesis tests) and those that serve only as a proxy for each other, possibly due to sharing a common driving function (as may be the case in prediction models).

8.4. ARIMAX Models

We have found that merely differencing these time series results in lower correlations, with differencing of hourly data effectively removing all correlations. This suggests that much of the correlation in hourly data is the result of common cycles and that this variation obscures any other association between variables. However, we can do better with ARIMAX models which may difference to remove trends and enter cycles as AR and MA parameters. An ARIMAX model, with more precise removal of the noise of trends and cycles, could potentially reveal associations of the variables that are not the spurious correlations due to common cycling.

In general, ARIMAX models do a better job of forecasting relativistic electron flux than simpler regression models (Boynton et al., 2019, 2020; Simms & Engebretson, 2020). However, a transfer function model may not only result in a better fit, it can also show the cumulative effects of predictor variables. As many space weather parameters are thought to act over longer periods than a single hour, modeling the inputs and the decay can be a more effective way of describing their influence.

However, the specifics of the model chosen may seem to result in dramatic differences in a parameter's apparent influence, even if the part of the correlation attributable to this parameter remains the same. In Figure 13a, we found that a transfer function model may show alternating positive and negative influences of the ULF index on relativistic electron flux if both variables are differenced at 24 hr. This alternation of influence may represent the compilation of both positive and negative processes on electron flux that could be attributed to ULF waves. Possibilities include radial acceleration (e.g., Loto'aniu et al., 2006) or wave-particle interactions (e.g., Reeves et al., 2013) resulting in a higher electron population and outward radial diffusion due to wave-particle drift resonant interactions (Loto'aniu et al., 2010) resulting in electron loss. As these do not occur at the same time scale, it is possible that their opposing effects oscillate in influence. However, when viewed at a single time step, these oscillating influences may simply cancel out. This could explain why the simple correlations of the ULF index with electron flux (with cycles removed by differencing) results in low correlations. On the other hand, if flux is differenced at 24 hr and the ULF index at 1 hr, this alternating pattern disappears. No matter which model is used, the overall influence of ULF index (as measured by the partial correlation) is low (<0.02). With common cycles and trends removed, over long time periods, the influence of ULF waves on electron flux, therefore, appears minimal, at best.

The cumulative effect of V on ULF waves is superficially similar to that previously found (Engebretson et al., 1998) even though both variables are differenced in our present study. The maximum regression coefficient, at a 1 hr lag, is 4.12 using differenced data vs. the 6.52 found in the previous study using undifferenced data. However, the correlation that we find (0.104) is much lower than that found previously.

The ARIMAX models account for some cycles and trends by differencing the ULF wave index and the electron flux. This may be troubling, as there is a sense that this removes information about cycles that are actual physical events. Unfortunately, space weather data suffers from being purely observational and we cannot apply experimental treatments randomly (e.g., controlling the level of ULF waves applied before measuring the response of the electron flux). Without the ability to randomly set the independent predictor variables (treatments) in space physics studies, we are unable to infer causality and can only study the associations between variables. Random assignment of experimental treatments would allow us both to know a probability of how sure we are of associations between variables (a p -value) and to state that we have broken ties between possible hidden, confounding variables. Without randomization of treatments, no matter how many times we see flux changes following increasing ULF wave activity, we still have no assurance that there is a physical connection between these two

factors. However, we must do the best we can, and removing autocorrelation (often due to cycles and trends) is a necessary first step.

Granger and Newbold (1974) note three consequences of leaving uncorrected autocorrelation in regression models: (a) estimates of the coefficients may be wrong, (b) forecasts may not be optimal, and (c) significance tests of the coefficients will be invalid. Based on simulations, they point out that two independently autocorrelated time series often show a strong correlation even if there is no relation between them. Both seasonality and random walk processes can produce autocorrelation. A classic example of independent seasonal changes resulting in an apparent correlation is the association between ice cream sales and drownings. Presented in many introductory statistics textbooks, students are asked whether a high correlation between these two variables proves that eating ice cream causes drownings. Most students know intuitively that this is likely a spurious correlation representing the seasonal cycling of both parameters. Because both ice cream sales and swimming go up in the summer months, their increase or decrease is highly correlated with season, not with each other. We cannot, therefore, assume that ice cream causes drownings. While we can instantly see that the correlation in this ice cream example is likely ridiculous because it is really driven by seasonality, we have no independent way of making a similar common sense assessment of correlations in space weather data. We simply do not know, and cannot assess, if we are finding spurious correlations based solely on seasonality. Even worse, based on the simulations of Granger and Newbold, two time series of random walks can show as much correlation as series related by seasonality.

Therefore, we must have ways to at least model the seasonality and trends (even if a trend is only the result of random walk) so we can look at what remains of the correlation. In this way, we can assess how much of a correlation could be due to seasonality, how much to trends (random or otherwise), and how much remaining might be directly attributable to one variable driving another. (Although it must be remembered that even in the case of this “leftover” correlation, we cannot infer causality with certainty, as experimental treatments have not been randomly assigned to break the possible association with other possible hidden variables.)

It should be kept in mind that we are using one differencing transformation to accomplish two tasks at once: to remove a nuisance “seasonality” due to measurement difficulties, and to essentially model the longer-term solar cycles. Although the first cycle is artificial (due to satellite position) and the second due to real physical changes, both may obscure the actual relationship of ULF waves with flux. If one made the decision that only the diurnal cycle due to measurement errors should be removed, then one might accept the correlations from the bandstop filtered data. However, this approach makes the unsupportable assumption that correlations associated with further cycles and trends cannot be masking unmeasured variables.

However, differencing alone may not accurately describe the real association between variables, even though it removes cycles and trends. As we show above, by modeling the seasonality and trends solely with a difference transformation, we are led to believe that the “leftover” correlation between these variables is rather low. However, merely correlating differenced time series at specific hours is a somewhat crude method of determining the influence of one parameter on another. We may be more convinced if we can determine the statistically significant influences over time using ARIMAX models. By further describing the dependent time series with an ARIMAX model (which includes differencing), we find that the partial correlation due to adding the ULF wave index to an ARIMAX model of flux is, at best, $r = 0.011$ while that due to adding V to an ARIMAX model of the ULF wave index is $r = 0.104$. Only for the impact of solar wind velocity on ULF waves can we say that there appears to be some influence beyond the common cycling and trends, and this is much lower than previously found.

We conclude that correlation over long periods is not particularly useful in determining the association between electron flux, ULF waves, and solar wind velocity (unless terms are introduced into the model to account for geomagnetic conditions). Much (or all) of the correlation over long periods is due to common trends and cycles which do not describe the physical mechanisms between these processes. In addition, over long periods, including both active and quiet geomagnetic conditions, there may be many processes occurring, some possibly with opposing actions that sum to no change. There may be long periods of very little activity, which would obscure the correlations present during short periods of higher activity. The result is that we simply cannot see the action of V and ULF on flux if large periods of time are analyzed at once, and the influence of V on ULF waves is much lower than previously reported. However, by removing the diurnal cycle from electron flux and ULF data (with a

bandstop filter), we are more confident in our ability to understand the relationships between these variables (and the solar wind) within shorter periods defined by more homogeneous geomagnetic conditions.

9. Conclusions

1. We show that long period correlational analysis is not particularly useful: (a) because common cycles and trends result in inflated correlations unrelated to the physical associations between variables, and (b) because geomagnetically quiet and active periods may show opposing effects. Once common cycles and trends are removed by differencing, long period simple correlations between electron flux, the ULF index, and solar wind velocity are all much lower, as compared to previous reports. Removing the nuisance diurnal cycle and correlating over short periods showed often opposing influences between quiet and active periods. Future work should move toward modeling quiet and active time periods separately
2. Either bandstop filtering of the 24 hr frequency or differencing of the time series can remove the diurnal noise signal. Bandstop filtered data may be easier to interpret than differenced data. This may be the preferred transformation for the study of short term events. However, it does not remove longer trends which can result in spurious correlations and is therefore not an appropriate transformation for calculating correlations over long time series
3. For long period correlations, differencing of time series data (e.g., a $y_t - y_{t-1}$ transformation) easily removes long term trends and autocorrelation. Diurnal cyclicity in hourly data is best removed with a $y_t - y_{t-24}$ transformation, which also effectively removes autocorrelation and the longer cycles. However, an ARIMAX transfer function model may be a better choice for analyzing these associations as it can both account for the common cycles and describe the cumulative effects of one variable on another over time
4. Long period correlations of flux (with common cycles and trends removed) with either solar wind velocity or ULF waves are low to non-existent. The correlation between velocity and ULF waves is statistically significant, but much lower than that found with uncorrected time series
5. Over short periods, solar wind velocity shows different correlational patterns between quiet and active times with both electron flux and the ULF index (both with the diurnal cycles removed). The response of flux to ULF waves appears more similar between the three quiet and active periods we have chosen: ULF waves show a negative correlation with flux over the first day, rising to a positive correlation over the next two days

Data Availability Statement

Electron flux data were obtained from Los Alamos National Laboratory (LANL) geosynchronous energetic particle instruments (<https://zenodo.org/record/5834856>). The ULF index is available at http://ulf.gcras.ru/plot_ulf.html. Solar wind (V) data are available from Goddard Space Flight Center Space Physics Data Facility at the OMNI Web data website (http://omniweb.gsfc.nasa.gov/html/ow_data.html).

References

- Baker, D. N., McPherron, R. L., Cayton, T. E., & Klebesadel, R. W. (1990). Linear prediction filter analysis of relativistic electron properties at 6.6 RE. *Journal of Geophysical Research: Space Physics*, 95(A9), 15133–15140. <https://doi.org/10.1029/ja095ia09p15133>
- Borovsky, J. E., & Denton, M. H. (2014). Exploring the cross-correlations and autocorrelations of the ULF indices and incorporating the ULF indices into the systems science of the solar wind-driven magnetosphere. *Journal of Geophysical Research: Space Physics*, 119, 4307–4334. <https://doi.org/10.1002/2014JA019876>
- Boynton, R. J., Amariutei, O. A., Shprits, Y. Y., & Balikhin, M. A. (2019). The system science development of local time-dependent 40 keV electron flux models for geostationary orbit. *Space Weather*, 17, 894–906. <https://doi.org/10.1029/2018SW002128>
- Boynton, R. J., Aryan, H., Dimmock, A. P., & Balikhin, M. A. (2020). System identification of local time electron fluencies at geostationary orbit. *Journal of Geophysical Research: Space Physics*, 125, e2020JA028262. <https://doi.org/10.1029/2020JA028262>
- Brown, R. G., & Hwang, Y. C. (2012). *Introduction to random signals and applied Kalman filtering* (4th ed.). USA: John Wiley & Sons, Inc.
- Engebretson, M., Glassmeier, K. H., Stellmacher, M., Hughes, W. J., & Luhr, H. (1998). The dependence of high-latitude Pc5 wave power on solar wind velocity and on the phase of high-speed solar wind streams. *Journal of Geophysical Research*, 103, 26271–26283. <https://doi.org/10.1029/97JA03143>
- Granger, C. W. J., & Newbold, P. (1974). Spurious regression in econometrics. *Journal of Econometrics*, 2, 111–120. [https://doi.org/10.1016/0304-4076\(74\)90034-7](https://doi.org/10.1016/0304-4076(74)90034-7)
- Hyndman, R. J., & Athanasopoulos, G. (2018). *Forecasting: Principles and practice* (2nd ed, p. 291). Victoria, Australia: OTexts, Heathmont.
- Kozyreva, O., Pilipenko, V., Engebretson, M. J., Yumoto, K., Watermann, J., & Romanova, N. (2007). In search of a new ULF wave index: Comparison of Pc5 power with dynamics of geostationary relativistic electrons. *Planetary Space Sciences*, 55, 755–769. <https://doi.org/10.1016/j.pss.2006.03.013>

Acknowledgments

Work at Augsburg University was supported by NSF grants AGS-1651263 and AGS-2013648.

- Loto'aniu, T. M. (2006). Radial diffusion of relativistic electrons into the radiation belt slot region during the 2003 Halloween geomagnetic storms. *Journal of Geophysical Research: Space Physics*, *111*, A04218. <https://doi.org/10.1029/2005JA011355>
- Loto'aniu, T. M., Singer, H. J., Waters, C. L., Angelopoulos, V., Mann, I. R., Elkington, S. R., & Bonnell, J. W. (2010). Relativistic electron loss due to ultralow frequency waves and enhanced outward radial diffusion. *Journal of Geophysical Research: Space Physics*, *115*, A12245. <https://doi.org/10.1029/2010JA015755>
- Makridakis, S. G., Wheelwright, S. C., & Hyndman, R. J. (1998). *Forecasting: Methods and applications* (3rd ed., p. 652). New York, NY: John Wiley and Sons.
- Mathie, R. A., & Mann, I. R. (2000). A correlation between extended intervals of ULF wave power and storm-time geosynchronous relativistic electron flux enhancements. *Geophysical Research Letters*, *27*(20), 3261–3264. <https://doi.org/10.1029/2000GL003822>
- Mathie, R. A., & Mann, I. R. (2001). On the solar wind control of Pc5 ULF pulsation power at mid-latitudes—Implications for MeV electron acceleration in the outer radiation belt. *Journal of Geophysical Research: Space Physics*, *106*, 29783–29796. <https://doi.org/10.1029/2001JA000002>
- Neter, J., Wasserman, W., & Kutner, M. (1985). *Applied linear statistical models* (2nd ed., p. 112). Homewood, IL, USA: Richard D. Irwin, Inc.
- O'Brien, T. P., & McPherron, R. L. (2003). An empirical dynamic equation for energetic electrons at geosynchronous orbit. *Journal of Geophysical Research: Space Physics*, *108*(A3), 1137. <https://doi.org/10.1029/2002JA009324>
- O'Brien, T. P., Sornette, D., & McPherron, R. L. (2001). Statistical asynchronous regression: Determining the relationship between two quantities that are not measured simultaneously. *Journal of Geophysical Research: Space Physics*, *106*, 13247–13259. <https://doi.org/10.1029/2000JA900193>
- Oppenheim, A. V., & Schaffer, R. W. (2010). *Discrete-time signal processing* (3rd ed., p. 1132). Hoboken, NJ: Pearson Education, Inc.
- Pankratz, A. (1991). *Forecasting with dynamic regression models* (p. 386). New York, NY: John Wiley & Sons Inc.
- Parks, T., & McClellan, J. (1972). Chebyshev approximation for nonrecursive digital filters with linear phase. *IEEE Transactions on Circuit Theory*, *19*(2), 189–194. <https://doi.org/10.1109/TCT.1972.1083419>
- Pilipenko, V. A., Kozyreva, O. V., Engebretson, M. J., & Soloviev, A. A. (2017). ULF wave power index for space weather and geophysical applications: A review. *Russian Journal of Earth Sciences*, *17*, ES1004. <https://doi.org/10.2205/2017ES000597>
- Potapov, A. S., Tsegmed, B., & Ryzhakova, L. V. (2012). Relationship between the fluxes of relativistic electrons at geosynchronous orbit and the level of ULF activity on the Earth's surface and in the solar wind during the 23rd solar activity cycle. *Cosmic Research*, *50*(2), 124–140. <https://doi.org/10.1134/S0010952512020086>
- Reeves, G. D., Baker, D. N., Belian, R. D., Blake, J. B., Cayton, T. E., Fennell, J. F., et al. (1998). The global response of relativistic radiation belt electrons to the January 1997 magnetic cloud. *Geophysical Research Letters*, *25*(17), 3265–3268. <https://doi.org/10.1029/98gl02509>
- Reeves, G. D., Morley, S. K., Friedel, R. H. W., Henderson, M. G., Cayton, T. E., Cunningham, G., et al. (2011). On the relationship between relativistic electron flux and solar wind velocity: Paulikas and Blake revisited. *Journal of Geophysical Research: Space Physics*, *116*, A02213. <https://doi.org/10.1029/2010JA015735>
- Reeves, G. D., Spence, H. E., Henderson, M. G., Morley, S. K., Friedel, R. H. W., Funsten, H. O., et al. (2013). Electron acceleration in the heart of the Van Allen radiation belts. *Science*, *341*(6149), 991–994. <https://doi.org/10.1126/science.1237743>
- Rigler, E. J., Baker, D. N., Weigel, R. S., Vassiliadis, D., & Klimas, A. J. (2004). Adaptive linear prediction of radiation belt electrons using the Kalman filter. *Space Weather*, *2*, S03003. <https://doi.org/10.1029/2003SW000036>
- Simms, L. E., & Engebretson, M. J. (2020). Classifier neural network models predict relativistic electron events at geosynchronous orbit better than multiple regression or ARMAX models. *Journal of Geophysical Research: Space Physics*, *125*, e2019JA027357. <https://doi.org/10.1029/2019JA027357>
- Simms, L. E., Engebretson, M., Clilverd, M., Rodger, C., Lessard, M., Gjerloev, J., & Reeves, G. (2018). A distributed lag autoregressive model of geostationary relativistic electron fluxes: Comparing the influences of waves, seed and source electrons, and solar wind inputs. *Journal of Geophysical Research: Space Physics*, *123*. <https://doi.org/10.1029/2017JA025002>
- Simms, L. E., Engebretson, M. J., Pilipenko, V., Reeves, G. D., & Clilverd, M. (2016). Empirical predictive models of daily relativistic electron flux at geostationary orbit: Multiple regression analysis. *Journal of Geophysical Research: Space Physics*, *121*, 3181–3197. <https://doi.org/10.1002/2016JA022414>
- Simms, L. E., Engebretson, M. J., Rodger, C. J., Dimitrakoudis, S., Mann, I. R., & Chi, P. J. (2021). The Combined influence of lower band chorus and ULF waves on radiation belt electron fluxes at individual L-shells. *Journal of Geophysical Research: Space Physics*, *126*(5), e2020JA028755. <https://doi.org/10.1029/2020JA028755>
- Simms, L. E., Pilipenko, V. A., Engebretson, M. J., Reeves, G. D., Smith, A. J., & Clilverd, M. (2014). Prediction of relativistic electron flux following storms at geostationary orbit: Multiple regression analysis. *Journal of Geophysical Research: Space Physics*, *119*, 7297–7318. <https://doi.org/10.1002/2014JA019955>
- Su, Y.-J., Quinn, J. M., Johnston, W. R., McCollough, J. P., & Starks, M. J. (2014). Specification of >2 MeV electron flux as a function of local time and geomagnetic activity at geosynchronous orbit. *Space Weather*, *12*, 470–486. <https://doi.org/10.1002/2014SW001069>