Detection and Characterisation of Travelling Ionospheric Disturbances Using a compact GPS network

Joseph Reid[1], Richard Penney[1], Natasha Jackson-Booth[1] and Luke Selzer[1]

[1]QinetiQ
Malvern Technology Centre
St. Andrew's Road
Great Malvern
Worcestershire
WR14 3PS
United Kingdom

ABSTRACT

Computationally efficient methods for detecting and characterising Travelling Ionospheric Disturbances (TIDs) and other ionospheric phenomena have been developed. Innovative techniques to use GPS data to estimate TID motion, including: building a background fit of complex trends in the behaviour of the ionosphere, preserving wave-like behaviour and eliminating noise; inverting waveform changes to mitigate the effect of satellite motion, and cross correlating receiver sites to correct for specific satellite ephemeris; and pre-processing and intelligently down-sampling data to preserve trends while allowing for smaller computational cost.

The GPS data used to compute the motion estimates is collected by 3 GPS receivers, located in a triangle over a 3km baseline. The novel concept came from using the aforementioned data to produce results that preserve the subtle spatio-temporal features of TIDs. While cross correlating data from the three receiver sites, mitigation had to be made for satellite motion. Each site has a corresponding "pierce point" on the ionosphere, which is the point through which the transmission passes on its way from the satellite to the receiver. These pierce points move with respect to one another as the satellite traverses its ephemeris, and this effect must be allowed for as a time dependant variable, and the received waveform has to be correctly modulated to allow for correct ionospheric measurement.

These techniques allow confidence in the accuracy of the measurement of ionospheric behaviour. Extracting the TID waveform for cross-correlation requires fitting of the background ionosphere, identifying and modelling the complex diurnal trends. This allows preservation of oscillatory features in observed data, divorcing these from the usual but complex behaviour of the ionosphere.

This has been used in two ways; the development of a TID Identification Index, and in TID motion estimation. The Index is a computationally inexpensive tool for indicating the presence of TID like behaviour. By working from complex modelling and analysis described above, it can reliably suggest the presence of TID behaviour, and is an intuitive indication of unexpected phenomena.

TID motion estimation is a more complex field, but progress has been made. The clarity offered by the background fitting and ephemeris-aware cross-correlation means that oscillatory behaviour in the ionosphere can be reliably observed, and as such TID motion can be estimated for comparatively little computation power. Using this, TIDs moving over the UK on a daily basis have been characterised, with regular TIDs travelling in a south-east direction with a speed usually around 100 m/s. Due to the computational efficiency of this innovative solution, these observations have been shown to be relatively constant over many months of observed data, from multiple satellites.

1. INTRODUCTION

In 2013, we launched an investigation into Travelling Ionospheric Disturbances (TIDs). The aim of this investigation was: to identify general characteristics of TIDs, such as their form, magnitude, velocity and regularity; to better understand the impact of TIDs on radio systems; and to develop systems to detect and characterise TIDs from observed GPS data.

A novel element of this investigation was whether reliable and useful TID ionospheric data could be collected from a small network of GPS receivers over a short geographic baseline. To investigate this the TEMPLAR network was deployed.

2. GPS OBSERVATION: THE TEMPLAR NETWORK

The TEMPLAR Network is a small volume, small footprint GPS observation network. Over a 3 km baseline, the team has deployed 3 GPS recording stations (TEMPLARs). All of the analysis has been performed on data collected from these stations.

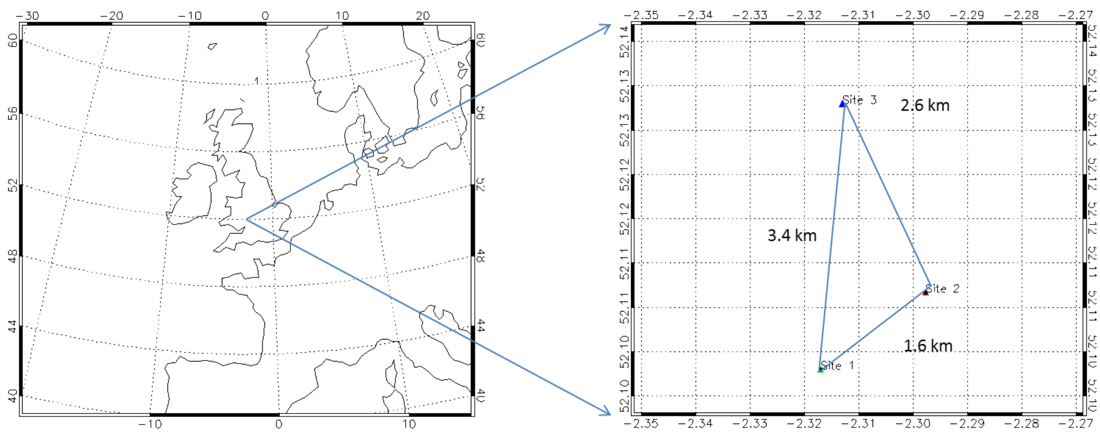

Figure 1: The TEMPLAR GPS receiver network

Each TEMPLAR has a number of components in order to facilitate useful data observation, smooth data collection, and system reliability. Each TEMPLAR comprises: a commodity dual frequency GPS receiver, with an antenna; a control laptop, remotely accessible to the research team via 3G WiFi dongle; an easily accessible external hard drive for data storage; a Rubidium (Rb) standard atomic clock for independent verification of data readings; and a software suite to integrate the parts, including remote access, networking, logging, and pre-processing tools.

There are two configurations developed by the team. The primary configuration, constructed and distributed at the start of 2013, used NovAtel GPS receivers, and did not use the Rb Standard. In July 2013, the GPS receivers were upgraded to Septentrio PolaRx4 models, and the Rb standard was incorporated.

The upgrade of the receivers was due to a number of factors, including the fact that the NovAtel receivers were no longer supported. The Septentrio receivers were chosen because they provided a less noisy GPS signal, more facilities for monitoring other GNSS constellations (such as GLONASS), and facilities for synchronisation with an external 10MHz frequency (the Rb standard)

The three TEMPLAR systems are disposed across a small footprint, in order to understand the capability of such a network. This layout is a triangle, with “sides” (inter-TEMPLAR distances) of 3.4km, 2.6km, and 1.6km. This network is pictured in Figure 1.

3. PRE-PROCESSING CHAIN

Once the GPS data has been collected, steps must be taken to obtain useful ionospheric measurements from the dataset. Effort has been applied to develop a bespoke pre-processing chain that performs this function, as well as to mitigate the errors introduced during observation, keep the data to a useable size, and push it into a format easily useable by the analysis team.

The primary function of the pre-processing suite is to transform the dual-band GPS readings into Total Electron Content. By combining the L1 and L2 carrier-phase measurements [e.g. *Ya'acob et al.,* 2000], we obtain the Total Electron Content between the receiver and the satellite at that point in time; this is also called the Slant TEC.

As with the collection of any real world data, the datasets contain errors. These errors fall mostly into two categories; "dropouts" and Cycle-slips. In order to reliably and correctly obtain the Slant TEC readings, these errors must be understood and corrected before the aforementioned transformation from dual phase GPS readings.

Cycle-slips (carrier phase cycle-slips) are a well-documented phenomena in recorded GPS. When measuring carrier phase, receivers maintain a count of the number of whole wavelengths between themselves and the satellites they are observing. If lock on a satellite is lost, this count is lost, and when the satellite lock is re-established, this count may have erroneously changed. This causes "jumps" in the data, and has to be accounted for.

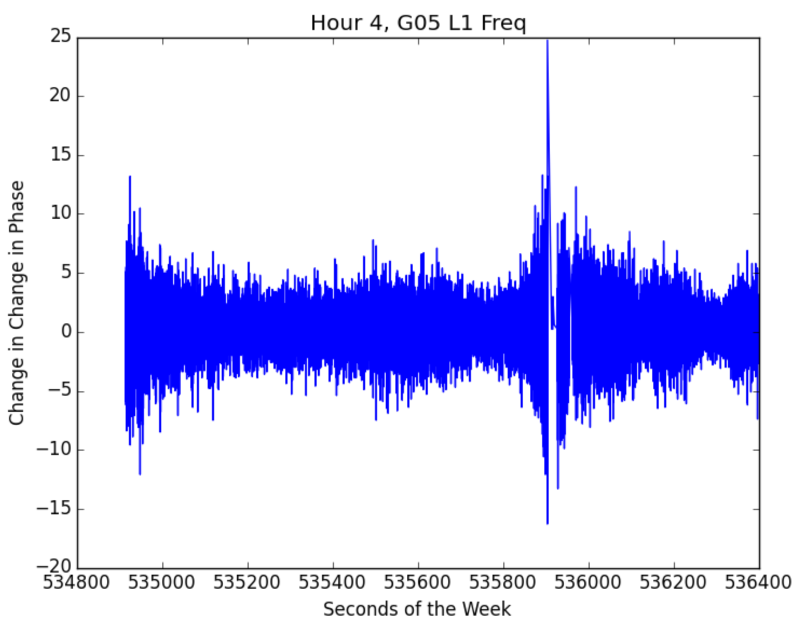

Figure 2: The double differential of a GPS time series, highlighting a dropout

"Dropouts" are periods of no readings between a satellite and a receiver, generally between 0.1 and 60 seconds in length (Figure 2). These occur randomly due to environmental factors and receiver noise, with increased likelihood when a satellite nears the horizon. If improperly accounted for, these can cause highly erroneous readings and incorrect analysis.

For the mitigation of these errors, we have incorporated GPSTk [*Tolman et al.,* 2004] into our pre-processing chain. An open source GPS tool suite maintained by the University of Texas at Austin, GPSTk has a number of routines to detect and correct errors in the readings, implementable through executables or through well documented C++ libraries. These highly reliable fixing routines allow faith in the data from which the Slant TEC readings are generated, and are an integral part of the pre-processing

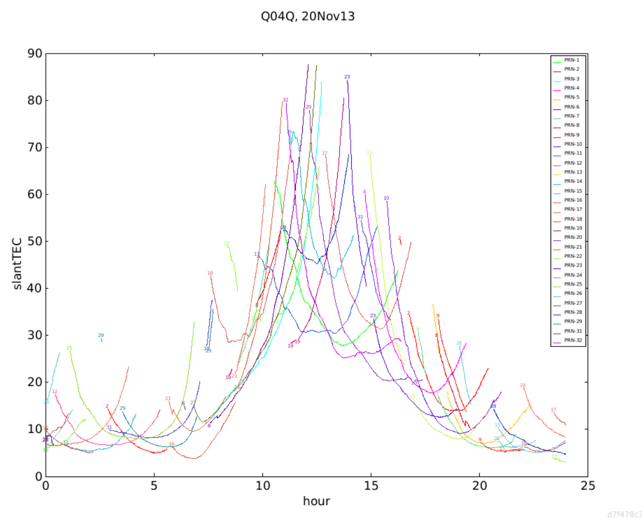

Figure 3: Sample Slant TEC time series, following pre-processing

Once the data has been corrected for discontinuities, it is down sampled from 10Hz to 1/30Hz, taking account of dropouts and unpredictable track start times across the network. It is from this down sampling of data that we can obtain Slant TEC time series, as illustrated in Figure 3. This down sampling is an essential part of the pre-processing that ensures computational efficiency.

Once the data has been down sampled, it is saved into a bespoke HDF5 file format [*Penney and Jackson-Booth*, 2015]. This system allows for easy and intuitive access to the data, and provides an implementation-agnostic format on which to build analysis and systems as part of the ongoing TID investigation. The pre-processing, and the analysis of data at every stage from observation to output, increases the reliability of the data and confidence in the systems tools developed.

4. TID INDEX

The TID Index is the first analytical system built on the observed data. The intention behind the Index is to efficiently identify the occurrence of anomalous ionospheric conditions. While it gives very limited information about the TID waveform, the Index will use provided data to very quickly detect TID-like behaviour in the ionosphere, allowing warning of erratic RF behaviour, and to allow more stringent analysis between system malfunction and ionosphere-induced behaviour change.

The TID Index is designed to efficiently process large amounts of data, meaning that it can quickly analyse months, or years, of data. It does this using a variable size window; this window has a default length of 2 hours, and a "window step size" of 1 hour, though the window length can be configured to any length, and the step size will usually be half of that value. This window-themed analysis works in the following way, starting at the first time point of the dataset:

1. For each satellite, the window identifies whether continuous STEC readings for all three TEMPLAR systems exist for the length of the window (which, given their almost colocation and the reliable pre-processing, is usually the case).
2. In the case that there exists a continuous data segment, a 4th order polynomial fit is applied to each of the three series (one per TEMPLAR). This fit is then subtracted from the respective series, to yield a residual of anomalous behaviour. These residuals are averaged if they are sufficiently similar.
3. The first and last 10% of the averaged series are discarded, and the maximum value of the remaining readings is calculated and stored.
4. The window moves to the next analysis point in the data (as governed by the window step size) and the process repeats.

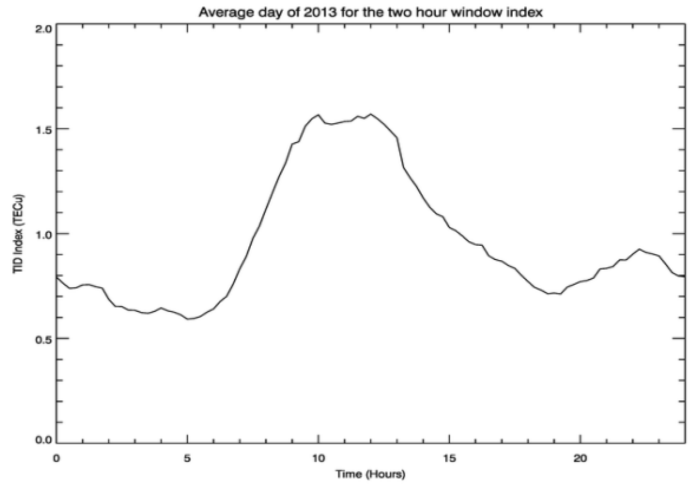

Figure 4: TID Index output for a single day in April 2013

Once all of these maximum data points have been calculated and stored, the variation of maximum observed Slant TEC values can be easily analysed against time, and given the computational efficiency this analysis can easily stretch to yearly comparisons.

Figure 4 shows the TID Index's output over the course of a day. It clearly represents the diurnal behaviour of the ionosphere, with a TEC peak at around midday, and shows maximum TEC values for the day. In this case, no unusual ionospheric activity for the day is detected.

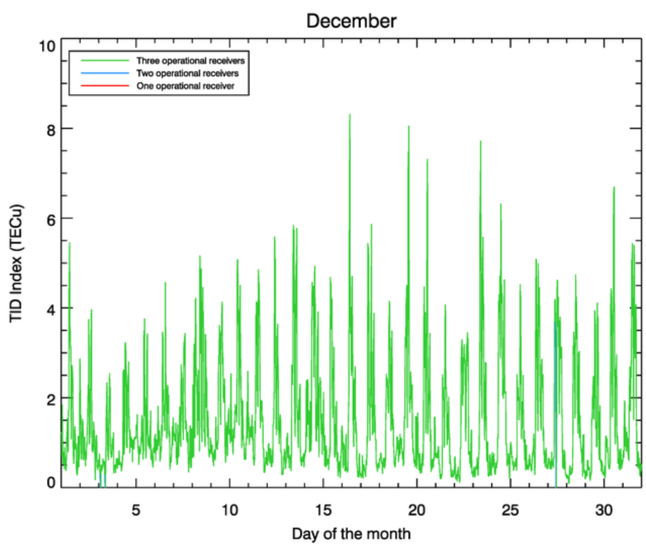

Figure 5: TID Index output for the month of December 2013

Figure 5 shows the Index's output for all of December of 2013. The daily peaks are easily observable, and we see a slight increase in ionospheric activity between the 17^{th} and 24^{th} of the month.

The two hour windowing system accurately profiles the day-to-day ionospheric activity as well as registering small scale TIDs (Figure 6a). For larger scale TIDs, with a period of 2 hours or greater, the window size must be increased to ensure the 4^{th} order polynomial fit doesn't hide the wave-like structure of the TID. By increasing the window to 6 hours, we can observe these larger scale TIDs. Using both the 2 hour window and the 6 hour window data, we get a feel for the overall behaviour of the ionosphere, and the occurrence of both small and larger scale TIDs. Figure 6a shows the Index using a 2 hour window for 2013, and Figure 6b shows the 6 hour window. Notice that the average reading of the 6 hour window is higher, and it shows more readings caused by larger scale TIDs.

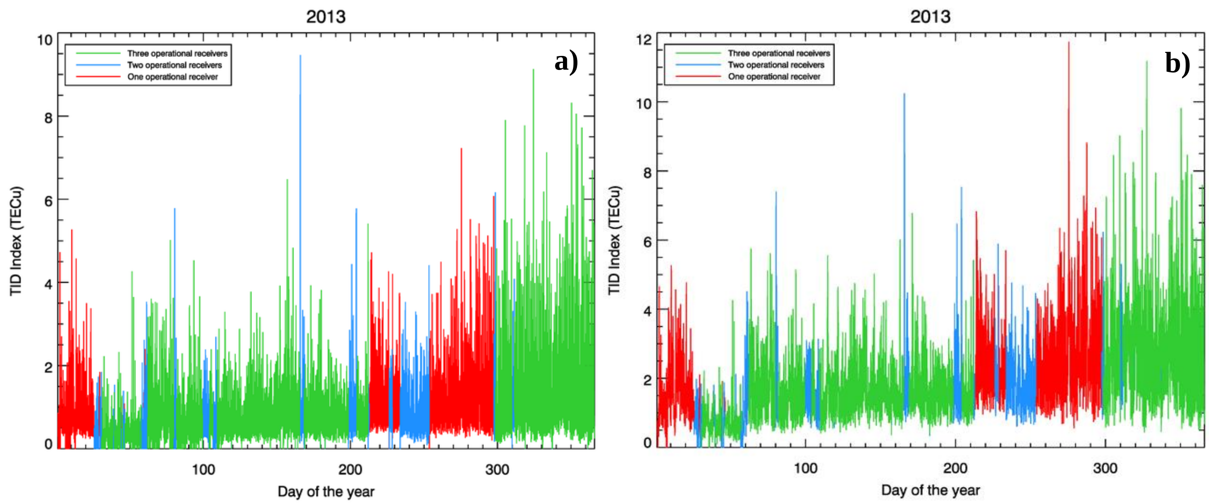

Figure 6: a) TID index output using a 2 hour window for the year 2013. b) TID index output using a 6 hour window for the year 2013.

5. MOTION ESTIMATION

While the TID Index offers us a computationally efficient method to detect the presence of TIDs, it offers no forecasting capability. In order to provide mitigation of TID impact on system behaviour, it becomes necessary to understand the motion of detected TIDs; this allows prediction of their location, and as such their impact, in the future. Estimating TID velocity is much more subtly influenced by satellite motion than the amplitude as measured by our TID index.

In order to characterise TIDs, as with their detection, the background ionospheric behaviour must be separated from the anomalies caused by the TIDs. This immediately gives rise to two major challenges to be addressed: picking the TID oscillations out from those of the background ionosphere, given that these resemble each other closely; and mitigating the time-varying Doppler shifts introduced onto the TID oscillations by the motion of the GPS satellites.

The first of these challenges is addressed by augmenting the usual approach of using polynomial fits [e.g. *Valladares and Hei*, 2012] on the background ionosphere by using minimum-curvature polynomial fits. This is a Bayesian approach to curve fitting, which introduces a cost function that incorporates observed scale for the rate of change of the TEC with time, typical TEC acceleration, and expected TID amplitude. This cost function strikes a balance between the closeness of the fit to the observed time series, and the degree of oscillation to the fitting function (with too much oscillation being punished, in order to preserve TID-like behaviour). An example of this fitting function can be seen in Figure 7, where both 8^{th} and 16^{th} order minimum-curvature fits preserve similar levels of detail of the TID waveform and show less over-fitting than a conventional un-weighted 8^{th}-order polynomial fit.

Once the TID oscillations have been extracted from the ionosphere, estimating the motion presents several new and interesting challenges. Estimating TID velocity from individual satellites gives very little information about the spatial extent of the TID, or of the wavefront shape. Combining this with the errors that can be introduced due to satellite motion and it becomes necessary to cross-correlate between satellites to more accurately estimate TID velocity. However, this cross-correlation must take account of the change in spacing between pierce-points due to satellite motion; an effect that occurs on a similar timescale to the TID period itself.

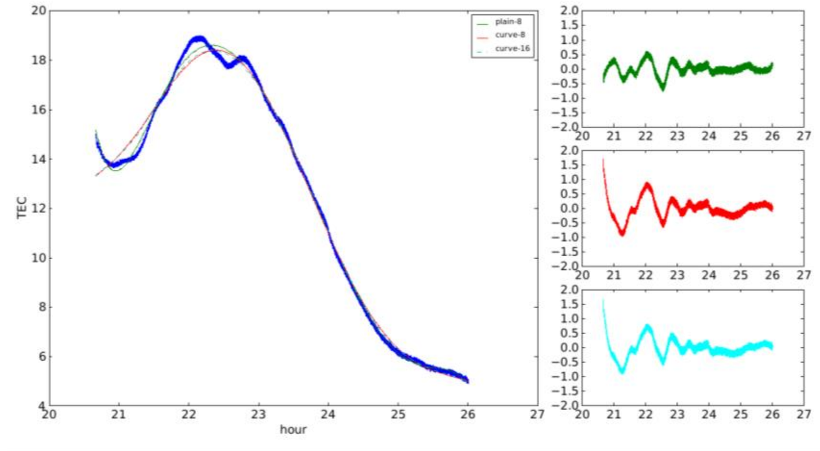

Figure 7: A typical TEC time series (dark blue) with three polynomial fittings. The green is 8th order, which over fits; the cyan is 16th order, and accurately preserves oscillatory behavior; the red is our Bayesian fit to the 8th order.

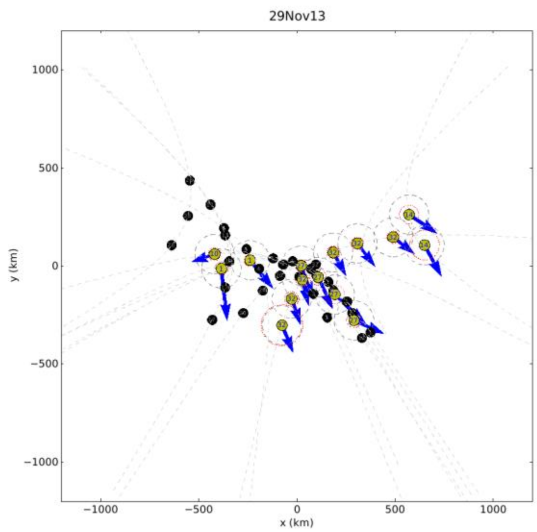

Figure 8: Estimation of the spatio-temporal distribution of TID velocity observed by the TEMPLAR network

There are two approaches to this multi-satellite TID characterisation: Combining measurements made from individual satellites across the three ground stations to give a more complete picture; and combining raw data collected from individual satellites from one receiver then using that to build a picture of the TID. Both of these avenues are being explored.

The combination of individually measured TID velocities across multiple satellites, together with each satellite's ephemeris data, allows the construction of a TID model incorporating spatial and temporal change. Figure 8 shows a plot of observed TID velocity from independent satellites, with dotted grey circles indicating a velocity of 100 m/s. This overall consistency, across a number of satellites and spanning a period of several hours, lends confidence to the existence of a south-easterly moving TID with a velocity of approximately 100 m/s.

The combination of raw TEC readings from multiple satellites to identify and characterise TIDs is a more complex task. It involves identifying TID like behaviour from observations often made hundreds of km apart, which becomes even more difficult given the small baseline of the TEMPLAR network. We are pioneering techniques that involve fitting a single TEC background trend to the observations of all satellites, and identifying TID like behaviour where the satellites deviate from the fit [*Penney and Jackson-Booth*, 2015]. This fitting must include satellite motion, as satellites necessarily have different look directions that sweep through the ever-changing ionosphere at hundreds of kilometres per hour.

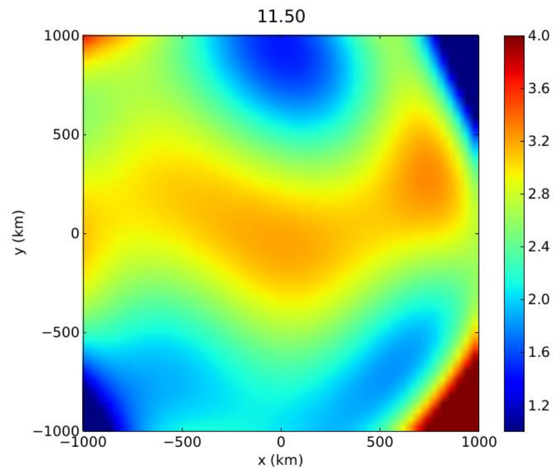

Figure 9: A snapshot of the multi-satellite TEC background trend

Our prototype solution generalises the already developed 1-dimensional trend estimation into 3-dimensions, including a 3-dimensional polynomial fit with an awareness of meaningful cut off values in spatial wavelength and wave-period. This prototype is providing promising results, and Figure 9 shows its TEC background trend for a day in December 2013; the outputs have proven themselves close to actual measurement.

6. CONCLUSIONS

The ongoing investigation into the detection and characterisation of TIDs continues to produce novel results. The use of the TEMPLAR network shows that meaningful ionospheric observation through GPS does not require large geographical baselines. We have also managed to develop new techniques for ionospheric analysis on the observed data from our novel network. While focussing on computational efficiency, we have: developed bespoke pre-processing that automatically ensures reliability and correctness in the data we observe, and reduces the data set to an efficient size while preserving all important aspects of the dataset; innovated the TID Index, an extremely efficient, customisable tool that highlights TID like behaviour, and can do so for data captured over massive lengths of time with speed; developed methods of cross-analysing observed TID like behaviour from many GPS satellites, lending useful insight into TID shape, magnitude, and velocity; and prototyped a 3 dimensional fit that incorporates raw TEC readings from many satellites, together with those satellites' ephemeris data, to intelligently map the movement of TEC in 3 dimension through time.

ACKNOWLEDGEMENTS

This work has been funded by the United Kingdom Ministry of Defence Science and Technology program.
We are grateful for many interesting and informative discussions with Prof. Matthew Angling, Dr. Gemma Attrill, Robert Moore and Nigel Taylor.

REFERENCES

Penney, R. W. & Jackson - Booth, N. K. (2015), Mitigating satellite motion in GPS monitoring of travelling ionospheric disturbances. Radio Sci., In Prep.

Tolman, B. R. B. Harris, T. Gaussiran, D. Munton, J. Little, R. Mach, S. Nelsen, B. Renfro, D. Schlossberg (2004), The GPS Toolkit – Open Source GPS Software., Proc. 17th Int. Tech. Mtg. Satellite Division of the Institute of Navigation(ION GNSS 2004), Long Beach, California.

Valladares, C. E., and M. A. Hei (2012), Measurement of the characteristics of TIDs using small and regional networks of GPS receivers during the campaign of 1730 july of 2008, *International Journal of Geophysics*, *2012.*

Ya'acob N., Abdullah M., and Ismail M (2000), GPS total electron content (TEC) prediction at ionpshere layer over the equatorial region, *Trends in Telecomms. Tech*. pp485-508.