Polar Science 28 (2021) 100638

Contents lists available at ScienceDirect

Polar Science

journal homepage: www.elsevier.com/locate/polar

Comparison of interferometer calibration techniques for improved SuperDARN elevation angles

Gareth Chisham [a,*], Angeline G. Burrell [b], Aurélie Marchaudon [c], Simon G. Shepherd [d], Evan G. Thomas [d], Pasha Ponomarenko [e]

[a] British Antarctic Survey, Cambridge, UK
[b] Naval Research Laboratory, 4555 Overlook Ave SW, Washington, DC, USA
[c] Institut de Recherche en Astrophysique et Planétologie, University of Toulouse, CNRS, CNES, Toulouse, France
[d] Thayer School of Engineering, Dartmouth College, Hanover, NH, USA
[e] University of Saskatchewan, Saskatoon, SK, Canada

ARTICLE INFO

Keywords:
SuperDARN
Ionosphere
HF radar
Interferometer calibration
Geolocation

ABSTRACT

The high frequency radars in the Super Dual Auroral Radar Network (SuperDARN) estimate the elevation angles of returned backscatter using interferometric techniques. These elevation angles allow the ground range to the scattering point to be estimated, which is crucial for the accurate geolocation of ionospheric measurements. For elevation angles to be accurately estimated, it is important to calibrate the interferometer measurements by determining the difference in the signal time delays caused by the difference in the electrical path lengths from the main array and the interferometer array to the point at which the signals are correlated. This time delay is known as t_{diff}. Several methods have been proposed to estimate t_{diff} using historical observations; these methods are summarised in this paper. Comparisons of the t_{diff} estimates from the different calibration methods are presented and sources of uncertainty discussed. The effect of errors in the estimated t_{diff} value on the accuracy of geolocation is evaluated and discussed. The paper concludes with a series of recommendations for both scientific SuperDARN data users and SuperDARN radar operators.

1. Introduction

The Super Dual Auroral Radar Network (SuperDARN) is a major tool for studying ionospheric and magnetospheric dynamics in both the polar regions (Chisham et al., 2007) and at mid-latitudes (Nishitani et al., 2019). One of the most important measurements that the high frequency (HF) radars make is the line-of-sight Doppler velocity of ionospheric F-region plasma that moves as a result of **E**×**B** drift. Combining this velocity data from the extensive fields-of-view (FOVs) of multiple SuperDARN radars allows the production of polar maps of ionospheric convection (Ruohoniemi and Baker, 1998).

Combining velocity measurements from multiple radars requires a high level of accuracy in the geolocation of the radar backscatter targets. To achieve this accuracy, the propagation paths of the HF radio signals to and from the scattering locations need to be well estimated (Greenwald et al., 2017). The SuperDARN radars are equipped with interferometers that make it possible to determine the elevation angle of arrival of returning radio signals, and hence, to estimate the most likely propagation path to the scattering volume (Milan et al., 1997; André et al., 1998; McDonald et al., 2013; Burrell et al., 2015; Shepherd, 2017; Chisham, 2018). However, attempts to calibrate these elevation angle measurements to ensure their accuracy have been hindered by a variety of issues, and they have been rarely used over the lifetime of SuperDARN. Instead, most studies have relied on virtual height models (VHMs) to estimate the locations of scattering volumes (Chisham et al., 2008; Greenwald et al., 2017). These VHMs have limited accuracy, and at far ranges their use can result in errors in the estimation of the ground range to the backscatter location of 100s of km (Yeoman et al., 2008).

The calibration problem exists because of differences in the electrical path lengths from the main and interferometer antenna arrays to the point at which the return signals from the two arrays are correlated with each other. This electrical path length difference leads to a difference in the signal travel times along the two different electrical paths, which is colloquially known in the SuperDARN community as t_{diff} (also referred to as TDIFF and δt_c in different SuperDARN publications). This difference in the signal travel times is responsible for a systematic offset

* Corresponding author.
E-mail addresses: gchi@bas.ac.uk (G. Chisham), angeline.burrell@nrl.navy.mil (A.G. Burrell), aurelie.marchaudon@irap.omp.eu (A. Marchaudon), simon.g.shepherd@dartmouth.edu (S.G. Shepherd), evan.g.thomas@dartmouth.edu (E.G. Thomas), pbp672@usask.ca (P. Ponomarenko).

https://doi.org/10.1016/j.polar.2021.100638
Received 8 July 2020; Received in revised form 30 October 2020; Accepted 4 January 2021
Available online 10 January 2021

in the measured phase difference between the signals from the two arrays, which must be corrected for accurate geolocation to be possible.

Until recently, the value of t_{diff} had only been estimated using engineering methods applied directly to the radar hardware, such as looking at the different delays of test signals in the radar cabling and electronics. However, for some SuperDARN radars, their remote location or the lack of local engineering resources has resulted in the absence of any engineering t_{diff} estimation. In such cases, a default value of zero has been assumed (equivalent to no systematic phase offset). Although it is known that different radar hardware components, such as the cabling and electronics, can have frequency and temperature dependent properties, no conclusive evidence has yet been presented in the literature that shows that the value of t_{diff} varies definitively with operational frequency, radar beam direction, or with season. Generally, the hardware dependencies of different radars have not been fully assessed. However, it is known that changes to the radar hardware can result in a significant change in t_{diff}, and so t_{diff} should be re-estimated in such cases. As an alternative to on-site engineering calibration, a number of methods have been proposed for calibrating elevation angle measurements by identifying t_{diff} using historical data (Chisham and Freeman, 2013; Ponomarenko et al., 2015; Burrell et al., 2016; Chisham, 2018; Ponomarenko et al., 2018).

During the 2017 annual SuperDARN workshop, poorly-estimated t_{diff} values were identified as the main issue preventing accurate elevation angle estimations. In response, an 'Elevation Angle Task Force' was formed to address this issue using data-based calibration methods. This paper presents the results and conclusions of the activities of this task force. Section 2 presents a brief summary of these data-based methods, including their strengths and shortcomings. Section 3 presents a comparison of t_{diff} values determined by the different analyses for three different SuperDARN radars located at different latitudes. Section 4 investigates the sensitivity of elevation angles and geolocation to variations in t_{diff}. Section 5 evaluates and discusses the factors that may affect the estimated value of t_{diff}. Section 6 summarises the paper and presents recommendations and an outline for future work.

2. Method summaries

Each of the methods presented here provides an independent way of estimating t_{diff} in order to calibrate the SuperDARN elevation angle data. All of the techniques are indirect methods (non-engineering) that use the measured range and either the elevation angle measurements or the phase difference measurements from which the elevation angle is derived. They also make certain assumptions about how these quantities are related. It is important to remember that all of these methods produce only estimates of t_{diff}, and the true value of t_{diff} is unknown.

All the methods have the following aspects in common:

- In contrast to direct engineering measurements of t_{diff}, no access to the radar hardware is necessary.
- They are all based on the analysis of SuperDARN data a posteriori. Hence, these methods can be applied to all historic SuperDARN data, including data from radars that are no longer operational, or for which the hardware has been significantly changed over the operating lifetime (as long as interferometer data have been recorded by that radar at that time).
- As the elevation angle determination involves a phase difference between two signals, estimations of t_{diff} by all the methods may be affected by a 2π phase ambiguity. That is, the same phase difference between the main and interferometer array signals (which relates to the elevation angle) will occur even if one of the phase values is shifted by a multiple of 2π. This means that the same calibration adjustment will occur for a t_{diff} value that is shifted by $\frac{n}{f}$, where n is an integer, and f is the operational frequency of the radar. This ambiguity can be eliminated by identifying the t_{diff} value for which measurements at all frequencies give the same (or very similar) values (typically within the uncertainty of the solutions), as presented in Chisham (2018). This frequency-independent solution is only strictly possible if t_{diff} is truly independent of frequency; this assumption is discussed in Section 5.1. In the absence of multiple frequency measurements, the best t_{diff} estimate is typically that identified closest to the hardware t_{diff} determined using engineering methods (if available).
- Systematic errors in the radar range determination will affect the accuracy of all the methods, but to different degrees. Those methods that rely on data from ranges near to the radar (such as meteor and E-region scatter) will be more strongly affected as the errors represent a larger fraction of their propagation path. This issue is a particular problem when determining the height of meteor trail scattering points (Section 2.2); the effect of these errors on t_{diff} estimates is discussed in Section 5.4.
- Many radars, particularly those with log-periodic antennae, have a strong back lobe in the radar antenna pattern (i.e., in the rear FOV opposite the intended look direction of the radar). Most types of scatter (but particularly ground and meteor scatter) can return from this back lobe, especially when the radar is operating at lower frequencies. Backscatter from the back lobe is problematic; the phase lags in these signals have the wrong sign as the standard elevation angle calculation assumes that the return direction of backscatter is the front FOV. This assumption leads to a population of backscatter echoes with dramatically incorrect elevation angles (Milan et al., 1997). Determining the return direction of backscatter echoes is possible, although the localised nature of meteor echoes pushes the boundaries of existing return direction identification methods (Burrell et al., 2015).

The different methods are summarised in the following subsections, with examples of applications for methods that have not been fully discussed in previous papers.

2.1. Virtual height model comparison method

The premise behind this technique is derived from the Chisham et al. (2008) VHM study, where a statistical analysis was performed on elevation angle measurements of ionospheric scatter from the Saskatoon (SAS) SuperDARN radar using a 5-year interval of data. They produced an empirical VHM for the SuperDARN radars by studying the distribution of measurements in range-elevation angle and range-virtual height space.

This method similarly determines the distribution of measurements from a particular radar and time interval in either range-elevation angle space or range-virtual height space. The expectation is that, for a properly calibrated interferometer, the elevation angle or virtual height variations with range would match the characteristic variations observed by Chisham et al. (2008). Viewing distributions in range-elevation space or in range-virtual height space (as in Figures 3–5 of Chisham et al., 2008) clearly shows when the t_{diff} value being used is severely in error (i.e., the distribution is significantly different from that expected from typical propagation modes). The profiles of these distributions can be adjusted by changing the value of t_{diff} used in the elevation angle determination until the distributions match closely with the variations predicted by the VHM.

Fig. 1 presents an example of this analysis method. Panel (a) shows the range-virtual height distribution for $\frac{1}{2}$-hop E- and F-region scatter at the Inuvik SuperDARN radar (INV) for the month of February 2017 for the frequency band around 12.4 MHz, where the elevation angles (and hence, the virtual heights) have been determined using the engineering t_{diff} value of 0.0 ns from the INV hardware file. The shaded regions show the probability of a particular virtual height at each range, where the distribution has been normalised for each slant range. The solid blue

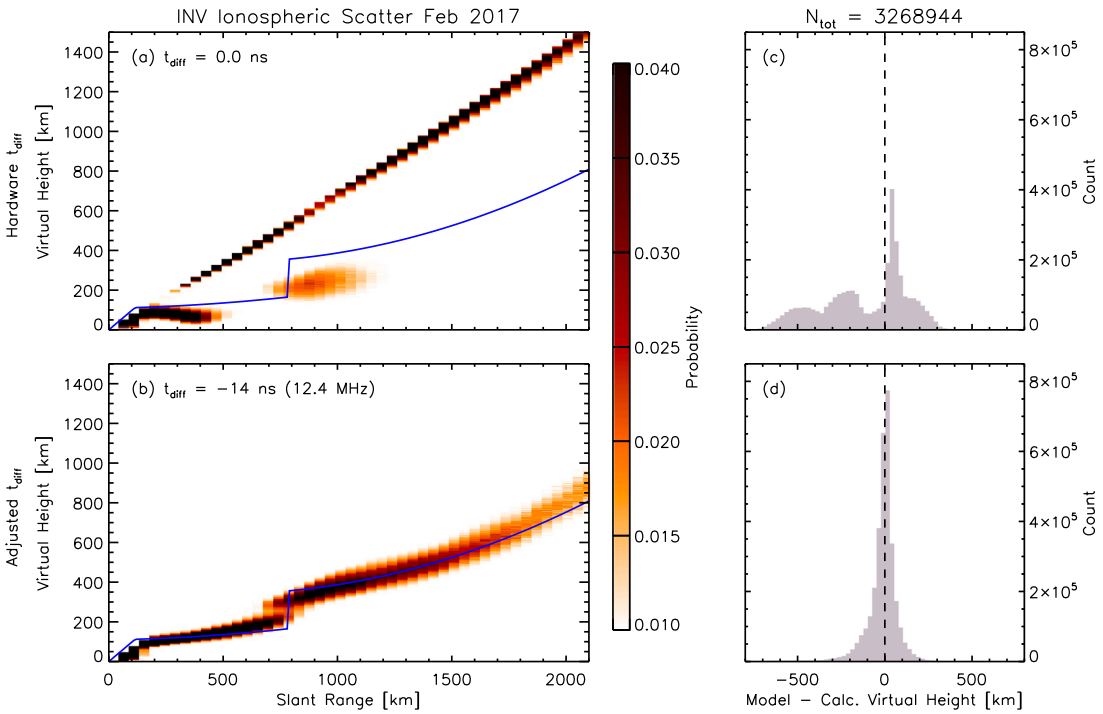

Fig. 1. Comparison of virtual height distributions when using the hardware t_{diff} (0.0 ns) and the t_{diff} found by adjusting the $\frac{1}{2}$ hop ionospheric backscatter to match the Chisham VHM (-14 ns at 12.4 MHz). Panels (a) and (b) show two-dimensional histograms of the slant range and virtual height distribution of the INV data from February 2017 for ranges 0–2100 km, with the VHM overlaid in blue. Panels (c) and (d) show histograms of the difference between the modelled and measured virtual heights at these same ranges. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

line presents the variation of virtual height from the Chisham et al. (2008) model across these ranges. It is immediately clear that there is very little agreement between the measurements and the model, indicating that the hardware t_{diff} is severely erroneous for this time and frequency. Most of the distribution appears as aliased values near the maximum observable elevation angle, giving the observed linear variation in the distribution (see Chisham, 2018 for a description of aliasing and the maximum observable elevation angle). The histogram in panel (c) shows the distribution of the differences between the model and the observations, confirming that they match very poorly. The peak in the histogram closest to zero is predominantly comprised of E-region data where the altitude variation is more restricted, minimising the magnitude of the differences between the measured and model virtual heights. This close agreement does not necessarily indicate that the hardware t_{diff} is suitable for the analysis of data from this region.

Panel (b) of Fig. 1 shows the range-virtual height distribution after adjusting t_{diff} to −14 ns to provide the best comparison with the Chisham VHM across both the E and F regions. The corresponding histogram in panel (d) shows that the agreement between the observations and the model is now excellent, strongly peaking around a difference of 0 km.

There are some limitations to the methodology:

- The method assumes that the Chisham et al. (2008) VHM variations are the same for all SuperDARN radars and for all times and conditions. This is not the case. The radar ray path depends heavily on the geographic and geomagnetic variations in the radar FOV, particularly the ionospheric electron density variation and the magnetic field inclination. These factors mean that the actual virtual height profile will vary with time of day, season, solar cycle, as well as the beam direction, as the electron density conditions are spatially variable. In addition, the transition region between scatter that is predominantly from the E-region and that which is predominantly from the F-region also changes with these factors.
- This method requires the removal of ground scatter from the data set (these are echoes returning from the ground or sea) to leave purely ionospheric backscatter, as the two types of scatter have different propagation modes and hence, different range-elevation angle profiles. The separation of ionospheric and ground scatter echoes using the default SuperDARN ground scatter criteria can be challenging for some radars, particularly those at mid-latitudes where low-velocity ionospheric scatter is often misinterpreted as ground scatter (e.g., Ribeiro et al., 2011; Burrell et al., 2018).
- This method assumes t_{diff} was properly calibrated for the SAS radar when the Chisham VHM was developed.

2.2. Meteor scatter method

This method to estimate t_{diff} was first proposed by Chisham and Freeman (2013) and is presented in detail by Chisham (2018). It uses SuperDARN meteor echoes from near ranges (typically between 200 and 350 km). Meteor echoes occur when radio waves backscatter from ionised trails left by meteoroids (meteor trails) as they enter the Earth's atmosphere (Ceplecha et al., 1998). These echoes represent the largest population of SuperDARN backscatter at near ranges (Hall et al., 1997; Chisham and Freeman, 2013).

The method involves taking meteor scatter data measured at a minimum of three ranges over a set time interval, and determining the elevation angle distributions for each range for this time interval using the zero-lag phase from the SuperDARN cross correlation analysis (e.g., Shepherd, 2017). These elevation angle distributions are then converted to meteor echo height distributions, and the distributions for the different ranges are compared. The value of t_{diff} used when determining the elevation angles is then incrementally varied to identify the t_{diff} value for which the height distributions of the meteor echoes for the different ranges show the minimum difference in the heights of the distribution peaks. This value is taken as the best estimate of t_{diff}.

There are limitations to this methodology that may cause errors and require careful consideration:

- As discussed above, meteor echoes can arise in the back lobe of the radar antenna pattern. These echoes can create a second erroneous population in the meteor height distribution, that (if not removed) can complicate the analysis process.
- Contamination of the meteor echoes by E-region echoes is a potential source of error. These echoes return from a higher altitude, and hence would create a second unwanted population in the echo height distribution. Significant effort is made to remove E-region scatter before the analysis (see Chisham and Freeman, 2013).
- Systematic errors in the range determination are a potential source of error as discussed above. These errors shift the peak height of the adjusted meteor height distributions with consequences for the estimated t_{diff} values (see discussion in Section 5.4). These systematic errors can be estimated and mitigated by fine-tuning the analysis through a comparison with other t_{diff} estimation methods or a comparison between the final meteor echo peak heights and theoretical or otherwise observed meteor echo peak heights.

2.3. Known location method

The Known Location (KL) method for estimating t_{diff} was first proposed by Burrell et al. (2016). When the location of a scattering volume can be determined independently of the SuperDARN data set, the elevation angle equation can be solved iteratively for t_{diff}. Additionally, if the uncertainty in the independent location measurement is known, it is possible to estimate its contribution to the uncertainty in t_{diff}.

Some 'known' location coordinates are more reliable than others. Backscatter from irregularities created by ionospheric heaters (e.g., Stubbe et al., 1982; Papadopoulos et al., 1990; Robinson et al., 2006) is excellent for calibration purposes, because the irregularity location can usually be determined with relatively good accuracy and the backscatter signatures are easily identified. These irregularities provide a spatially-localised target for HF radio waves and hence, very localised backscatter. They have been used for some time to estimate the uncertainties in ground range determination using VHMs (Yeoman et al., 2001, 2008). This method was used to successfully estimate t_{diff} for the Hankasalmi SuperDARN radar (HAN) by Burrell et al. (2016) using the ionospheric heaters at Svalbard and Tromsø, Norway.

In some instances, ground or sea scatter may be used successfully with this method, as this type of scatter often returns consistently from specific targets (such as mountain ranges, bodies of water and coastlines). The main caveat when using this type of scatter as a target is that its origin location must be obvious and unambiguous.

For example, INV is a questionable radar for the application of this method because the number of locations from which ground or sea scatter could return from are numerous. The wide availability of water, mountains, and coastline in the INV FOV provides a wealth of targets when sea ice is not present (Ponomarenko et al., 2010). As an example, Fig. 2 presents three possible return locations for ground scatter in beam 15 of the INV FOV. Fig. 2(d) shows the postulated 'known' locations of the most distant ground scatter observed along INV beam 15 between 1600 UT on 1 June 2016 and 0400 UT on 2 June 2016. This far range ground scatter was chosen as its origin would most likely occur near the coastline of Greenland, away from the numerous islands in the near FOV.

Fig. 2(a)–(c) shows histograms of the location results when using the selected ground scatter with the hardware t_{diff} (grey), the estimated t_{diff} from this method (different shades of blue that correspond to the location stars in Fig. 2(d)), and the 'known' location in geographic latitude marked as a black dashed line for the 3 postulated return locations. The histograms in these figures are shown as a function of latitude rather than slant range, as latitude was used in the optimisation process. The closest and furthest locations, panels (a) and (c) respectively, have extremely large differences between the median location of the solution and the expected 'known' location. The difference is smallest for the central location (panel (b)), but the difference between the 'known' and median locations is still on the order of a degree of latitude. It is therefore likely that this backscatter is originating from the mountain range near the western coast of Greenland rather than the western coastline.

The failure of this attempt to accurately identify t_{diff} ironically reveals a strength of the known location method: if a clearly incorrect location is specified for the backscatter, the method will not converge onto that location. Thus, it may be possible to adapt the known location method to iteratively identify t_{diff} using backscatter coming from a *single* location by optimising two parameters: the known location function and the known location itself, assuming a characteristic distribution about the known location.

The limitations of this methodology are:

- The method is heavily reliant on the existence of known backscatter targets that have a localised extent in range, meaning that in most instances their location can be accurately identified. Although using artificially generated backscatter is the most reliable method, there are currently only two heating facilities (EISCAT, the European Incoherent SCATter radar, and HAARP, the High-frequency Active Auroral Research Program) within the region observed by SuperDARN. Successfully applying this method to other radars requires the identification of small and distinct regions of sea or ground scatter. At present the use of this type of scatter as a known location is in its infancy.
- This method works best when the known location is either a $\frac{1}{2}$-hop or 1-hop scattering location. If the scattering location is either a $1\frac{1}{2}$-hop or 2-hop scattering location (or greater) then the assumption needs to be made that the ionospheric reflection heights on multiple hops are identical. Although these heights may be similar in many cases, for long propagation paths where rays are passing through different physical regions on each hop (e.g., the auroral zone and then the polar cap), the reflection heights may be very different due to the different electron density profiles in these regions. Hence, a cautious approach is required with scatter from $1\frac{1}{2}$-hop or greater propagation paths.
- When using either meteor or ground scatter as the known location, this method is susceptible to problems with scatter originating from the back lobe of the radar, as discussed above.

2.4. Ground scatter method

This method to estimate t_{diff} was first proposed by Ponomarenko et al. (2015). The premise behind this technique is that the reflection height for a spatially-extended region of ground scatter echoes at a given operational frequency is approximately constant, meaning that properly calibrated ground scatter elevation angles will decrease towards zero with increasing range. The technique involves consecutive adjustments of t_{diff} until the elevation angle data show the expected pattern, i.e., a general decrease of elevation angle with range which shows at large distances a near-zero background accompanied by sporadic isolated jumps.

There are a number of limitations to this methodology:

- Ground scatter with an elevation angle near 0° is not always observed for several reasons. Firstly, due to the limited spatial extent of ground scatter targets (e.g., finite-sized mountain ranges and bodies of water). Secondly, because in the typical SuperDARN vertical antenna pattern the gain drops off rapidly below elevation angles of 10°, suggesting that the scattered signal will be much weaker at very low elevation angles. Thirdly, due to features that obscure the horizon (e.g., nearby mountain ranges). Hence, an asymptotic approach is generally required.

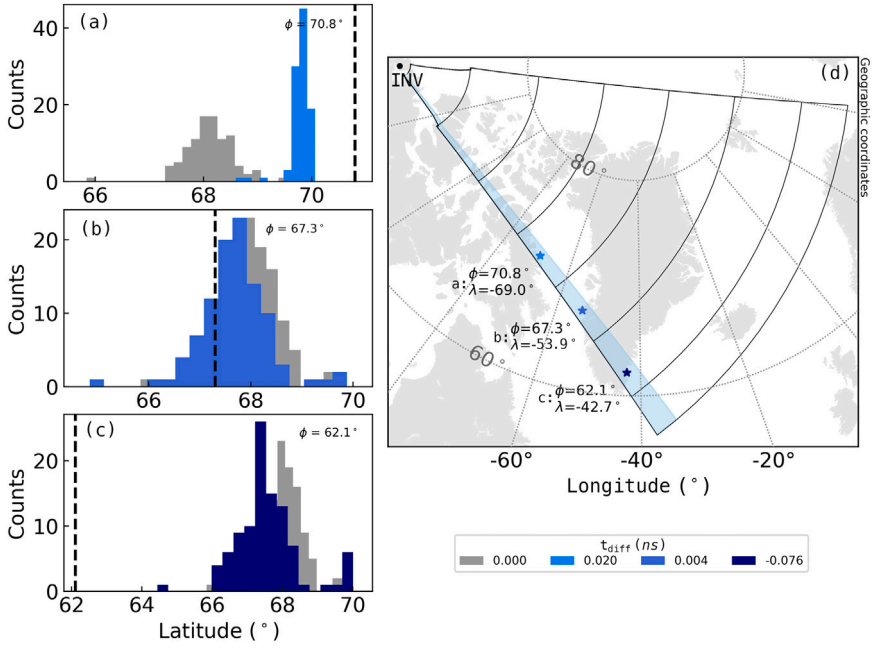

Fig. 2. INV t_{diff} estimations for the 12.2–12.5 MHz frequency band made along beam 15 using ground scatter with power at or above 1 dB, slant ranges between 3555 km and 4230 km, and times between 1600 and 0400 UT (from the 1–2 June 2016). Panels (a)–(c) show the geolocation results of the t_{diff} estimation assuming the three different 'known locations' from panel (d). In each of the first three panels the black dashed line marks the 'known' latitude, the grey histogram shows the geolocation of the data using the hardware t_{diff}, and the blue histograms show the geolocation of the data using the estimated t_{diff}. The different shades of blue correspond to different starred locations in panel (d), which are also marked with the corresponding panel letter. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

- As discussed above, many radars observe strong ground scatter from the back lobe of the radar (Burrell et al., 2018), which is problematic for this method.
- Statistical fluctuations will cause aliasing in the phase difference measurements near 0°, so that the elevation angles determined lie near the maximum observable value. In addition, the elevation angle near 0° is highly sensitive to small variations in the phase difference. Uncertainty in the measured phase difference, therefore, has a larger impact on elevation angles near 0°.
- The ionosphere is regularly perturbed by magnetospheric waves and atmospheric gravity waves that result in perturbations in the ground scatter. These waves cause spatiotemporal fluctuations in the signal reflection height through their impact on the electron density height profile and the magnetic field inclination, which results in a broadening of the ground scatter location profile.
- Not all radars observe ground scatter. For those that do, its occurrence rate and location can vary with season due to the presence of sea ice, which is a poor reflector compared to the rough ocean surface. This fact would restrict the estimation of t_{diff} using this method at many polar radars to particular times of year.
- The high variability in F-region propagation conditions means that this method is not easily automated. Application of the method requires visual analysis of data records.

2.5. E-region method

This method to estimate t_{diff} was first proposed by Ponomarenko et al. (2018). It is based on the observation that E-region backscatter echoes often dominate ranges between 300 and 800 km, and that they always exhibit a gradual decrease in elevation angle with range. This fact suggests that the E-region echoes propagate along low-angle rays that correspond to scatter from the bottom part of the layer (in contrast to Pedersen rays that travel along the maximum of the E layer, which is frequently the case with the F-region echoes). Due to the well-defined structure of the propagation and the comparatively low variability of the E-region altitude and width, it is possible to predict with high probability that starting from a certain range (beyond 700 km) these echoes will be coming from near-zero elevation. Therefore, it is in principle possible to align the observed phase difference distributions at these ranges with their expected values by adjusting t_{diff}.

This method is similar to the ground scatter method described above, but the fact that the E region is a more altitudinally confined ionospheric layer than the F region means that the range and altitude coverage are much more constrained, and the method is much easier to automate.

There are, however, some limitations to the methodology:

- The method does not account for phenomena such as sporadic E, which alters the heights of E-region backscatter.
- As the method relies on near-range backscatter, it is sensitive to range errors (Section 5.4), as discussed above for the meteor scatter method (Section 2.2).
- In the summer months E-region scatter at certain ranges is almost completely masked by F-region scatter, whose effective scattering volume moves closer to the radar site. However, this issue is often easily mitigated by removing data from the ranges where the F-region backscatter start to dominate.

3. Method comparisons

One of the main objectives of the SuperDARN Elevation Angle Task Force was to validate indirect estimates of t_{diff} through the comparison of results from different calibration methods. In the first instance, a small number of SuperDARN radars were chosen for the method comparison, covering a range of latitudes from the polar regions to mid-latitudes. The radars chosen for the initial study were INV and Rankin Inlet (RKN) from the polar radars in northern Canada, HAN from the auroral zone in Finland, and Christmas Valley East (CVE) from mid-latitudes in the USA. The time interval chosen for this initial analysis was the years 2016–2017. The results of these initial analyses were

compiled in the task force annual report for the SuperDARN executive committee, and are briefly discussed below.

In this initial analysis, the estimated values of t_{diff} for each particular radar were generally consistent between the methods within a range of about 5 ns. The results at HAN were the exception to this, where problems arose as a result of the large distance between the main and interferometer arrays (185 m), which limits the measurements of elevation angles to below 30° and particularly inhibits the accurate determination of the elevation angle of meteor scatter at near ranges (<300 km slant range). In addition, the lack of measurements at widely separated frequencies made it difficult to deal adequately with the 2π ambiguity discussed in Section 2.

Similarly, this paper presents comparisons of the derived values of t_{diff} for three selected SuperDARN radars, one from polar latitudes (INV), one from the auroral zone (HAN), and one from mid-latitudes (Christmas Valley West - CVW). Some of the comparisons presented here represent an extension of the initial analyses to longer time intervals. The comparisons employ all possible methods for those intervals, except the ground-scatter method which is deemed too subjective and labour intensive. The ground-scatter method is used for some validation analysis, however.

The results for the different methods are presented at different temporal resolutions due to the restrictions and limitations inherent to each method. The two methods that provide temporal variations of t_{diff} at relatively high resolution are the meteor scatter and E-region scatter methods. Chisham (2018) presented estimates of t_{diff} using the meteor scatter method at 3-month, 10-day, and daily resolution. They concluded that the estimates at 10-day resolution provided the right balance between precision and resolution, and the results here are presented at that resolution. Ponomarenko et al. (2018) presented estimates of t_{diff} using the E-region scatter method at monthly and daily resolution. Their results at different temporal resolution are broadly consistent, but the point-to-point variability of the t_{diff} estimates is increased at the daily resolution. Here, we present the results from this method at daily resolution. The results from the VHM comparison method are presented as monthly averages, whereas those from the known location method are effectively from a fixed instant in time.

3.1. Polar-latitude SuperDARN radar example (INV)

Fig. 3 presents a comparison of the t_{diff} estimates determined by three of the methods at INV for the 10-year epoch from 2008 to 2017 inclusive, and for the frequency band centred around 12.4 MHz (one of two major operational frequency bands used at INV). The black, orange, and red symbols in panel (a) represent estimates from the E-region, meteor scatter, and VHM comparison methods, respectively. There are no reliable known location estimates for this radar over this interval (see Section 2.3). The horizontal purple and white dotted line represents the hardware value of 0 ns, which is the only hardware value for this whole 10-year interval.

There are long intervals during this 10-year epoch where the estimates of t_{diff} appear relatively invariant with time, although a degree of statistical variability is evident in the higher-resolution estimates. Although the temporal variation of t_{diff} is steady for long intervals, there are several very sharp jumps in the estimated t_{diff} value across the epoch. As the different methods identify these sharp jumps at identical times, it seems likely that these jumps are not a result of the failing of a particular method. Indeed, the engineering team responsible for the INV radar have confirmed that all these jumps in the estimated t_{diff} are the result of either hardware problems, software calibration, or hardware calibration. These problems and calibrations will all have resulted in changes in the electrical path lengths from the two antenna arrays, and subsequently the true value of t_{diff}.

There are also indications of seasonal variations in the t_{diff} estimates at times, particularly in the VHM comparison method estimates from 2009 to 2013, and in the meteor scatter method estimates from 2011 to 2013, and 2015 to 2017. Note that the seasonal variations appear to be in the opposite sense for these methods and the possible origins of these temporal variations in t_{diff} will be discussed later in Section 5.3.

Panel (b) of Fig. 3 presents histograms of the different t_{diff} estimates for an interval during which the methods showed approximately consistent estimates of t_{diff} for a long period of time; the interval from January 2009 to June 2013 inclusive, marked by the grey region in panel (a). The black, orange, and red histograms correspond to the E-region, meteor scatter, and VHM comparison methods, respectively. The horizontal purple and white dotted line represents the hardware value of 0 ns. Throughout this interval, the E-region and meteor scatter methods result in relatively consistent estimates of t_{diff} covering a range of a few ns close to the hardware value; the meteor scatter method estimates of t_{diff} are approximately offset 1 ns on average in the positive direction. However, the VHM comparison method results are split into two distributions. The winter estimates are offset approximately 2 ns in the negative direction, whereas the summer estimates are offset approximately 6 ns in the negative direction. The reason for the observed differences between the method estimates remains to be fully understood, and will be a subject of further work.

3.2. Auroral-latitude SuperDARN radar example (HAN)

Fig. 4 presents a comparison of t_{diff} estimates at HAN for the last 3 months of 2006, using the known location method with EISCAT Tromsø heater irregularities (Stubbe et al., 1982), the E-region method, and the VHM comparison method. It was not possible to produce estimates of t_{diff} from the meteor scatter method for this time due to the unreliability of the meteor scatter elevation angles (as discussed earlier in this section). The results were determined for the Stereo channel A, and the frequency band 11.075–11.275 MHz. The lack of multiple frequency observations at this time makes it difficult to assess the effect of the 2π ambiguity at HAN, and so the t_{diff} estimates are shown as their estimated value closest to the hardware value.

The daily E-region method estimates are shown as black symbols in panel (a), and a black histogram in panel (b). These are characterised by a fair amount of scatter ($\sigma = 1.58$ ns) across the 3-month interval, with a mean and median of 154 ns. Only one known location t_{diff} estimate from a single heater experiment was available during this period of time (indicated by the dashed, blue line in both panels at 156 ns). The duration of the validity of this known location estimate (highlighted by the extent of the dashed blue line in panel (a)) was established using comparisons with the ground scatter method discussed above. A similar comparison has previously been presented by Burrell et al. (2016) in Fig. 10 of that study. The VHM comparison method estimation was evaluated for 1-month intervals, and remained constant over the 3-month interval at 144 ns, shown as the solid red line in both panels (a) and (b). The median difference between the E-region and known location method estimates is -1.9 ns, while the difference between these two methods and the VHM comparison method estimates is ~10 ns. The difference between the E-region and known location methods is similar to the uncertainty in the daily estimates of the E-region method (assuming the scatter in the E-region method to be a result of random errors in the estimation process).

The hardware value of 181 ns is indicated in both panels (a) and (b) by the purple dotted line. Although the difference between the hardware value and that of the other estimated values seems large, it should be stated that analysis of data from the Stereo channel B using the VHM comparison method estimated t_{diff} to be consistent with the hardware t_{diff}. The multi-channel analysis using the VHM comparison method is highly suggestive that the different Stereo channels of this radar are characterised by different t_{diff} values. As a consequence, the present format of the SuperDARN hardware files, which allow for only one t_{diff} value at a particular time, restricts the determination of elevation angles to data from only one channel. The reason for the difference in t_{diff} between channels is likely due to differences in various electrical components along the respective paths of the two separate channels.

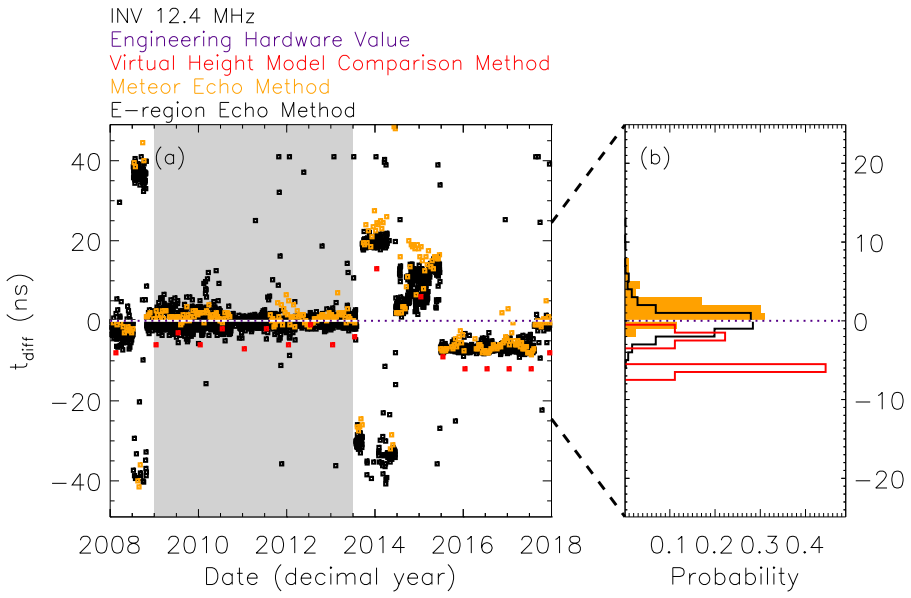

Fig. 3. Comparison between the results of different t_{diff} estimation methods for the INV radar for the interval 2008 to 2017 inclusive. The estimates have been determined for a frequency band centred around 12.4 MHz, with a 2π integer ambiguity added (where necessary) to obtain the t_{diff} estimate closest to the hardware value of zero. (a) The black, orange, and red symbols represent the estimates from the E-region, meteor scatter, and VHM comparison methods, respectively. The purple and white horizontal dotted line at 0 ns represents the hardware value. (b) t_{diff} estimate histograms for the interval January 2009 to June 2013 inclusive. This interval is marked by the grey region in panel (a). The black, orange, and red histograms represent the t_{diff} estimate distributions from the E-region, meteor scatter, and VHM comparison methods, respectively. The purple and white horizontal dotted line again represents the hardware value of 0 ns. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

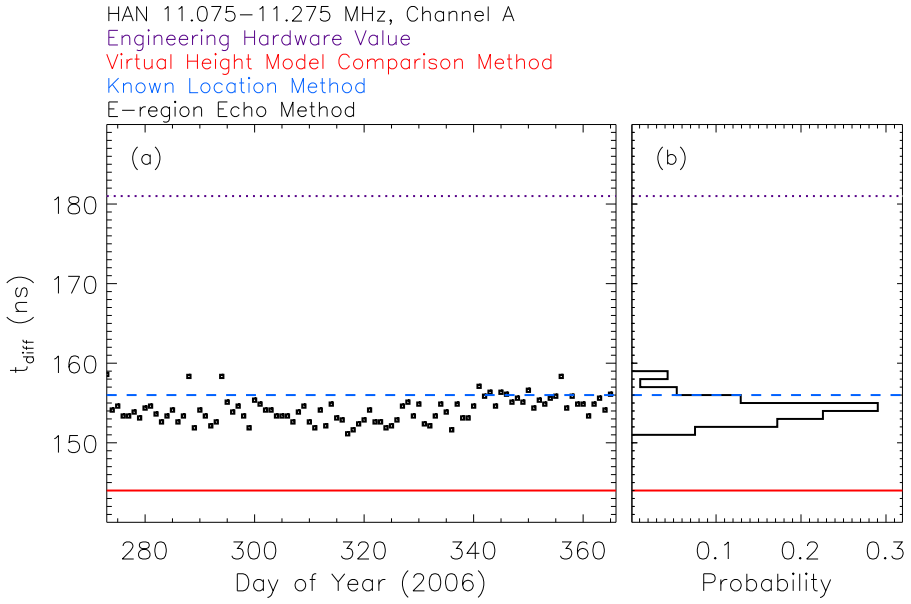

Fig. 4. Comparison between the results of different t_{diff} estimation methods for the HAN radar for the last 3 months of 2006. The estimates have been determined for the Channel A 11.075–11.275 frequency band, with a 2π integer ambiguity added (where necessary) to obtain the t_{diff} estimate closest to the hardware value of 181 ns. (a) The black symbols represent the estimates from the E-region method. The blue dashed line and the solid red line represent the estimates from the Known Location and VHM comparison methods, respectively, over the interval that they were valid. The purple and white horizontal dotted line at 181 ns represents the hardware value. (b) The t_{diff} estimate histogram for the E-region method for the same interval. The blue and red lines represent the t_{diff} estimates from the Known Location and VHM comparison methods, respectively, which were constant over this evaluation period. The purple and white horizontal dotted line again represents the hardware value of 181 ns. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

3.3. Mid-latitude SuperDARN radar example (CVW)

Fig. 5 presents a comparison of the t_{diff} estimates for CVW between 2014 and 2018 inclusive for the frequency band 10.3–10.9 MHz. In panel (a), the orange, red, and blue symbols represent estimates from the meteor scatter, VHM comparison, and known location methods, respectively. Panel (b) shows the distributions of the t_{diff} estimates for the meteor scatter (orange) and VHM comparison method (red) estimates. The single horizontal blue dashed line in panel (b) represents the known location estimate. The median of the t_{diff} estimates from the meteor scatter method is −348.8 ns, whereas the median of the estimates from the VHM comparison method is −347.0 ns. Although the difference in the medians for these two methods is 1.8 ns, the differences vary between ∼1 ns and ∼4 ns across the 5-year interval.

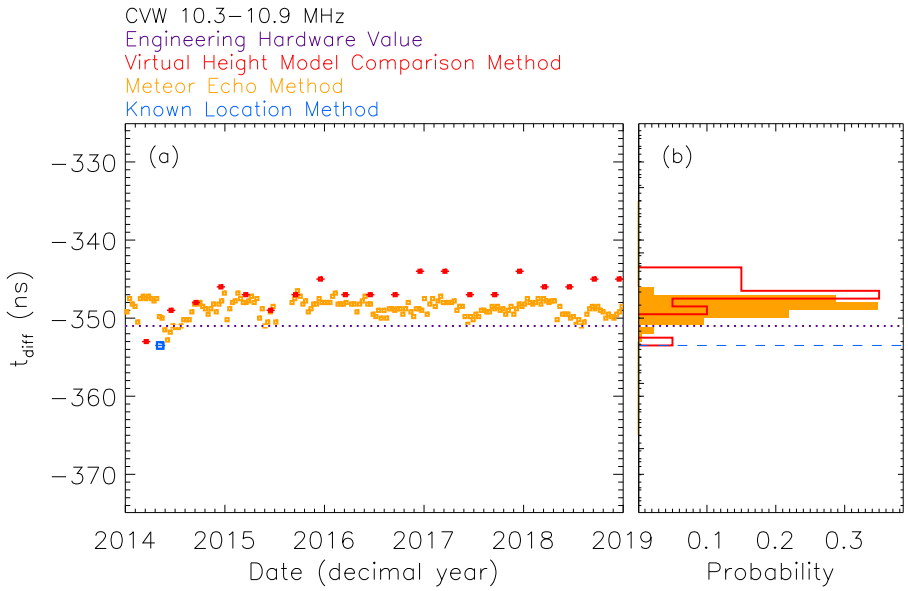

Fig. 5. Comparison between the results of different t_{diff} estimation methods for the CVW radar for the interval 2014 to 2018 inclusive. The estimates have been determined for a frequency band centred around 10.6 MHz, with a 2π integer ambiguity added (where necessary) to obtain the t_{diff} estimate closest to the hardware value. (a) The orange, red, and blue symbols represent the estimates from the meteor scatter, VHM comparison, and known location methods, respectively. The purple and white horizontal dotted line at −351 ns represents the hardware value. (b) The orange and red histograms represent the distributions of the meteor scatter and VHM comparison method estimates. The horizontal blue dashed line represents the value of t_{diff} from the known location method. The purple and white horizontal dotted line again represents the hardware value of 0 ns. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

The known location method estimate (of −353.5 ns) was obtained using irregularities generated by the HAARP ionospheric heater and observed along a $2\frac{1}{2}$-hop F-region propagation path. The difference between this method and the other two is ∼1–2 ns at the time of the measurement (23 April 2014). The horizontal purple and white dotted line represents the CVW hardware value of t_{diff} (−351.0 ns), determined using engineering methods. For 2015 onward, the results from the meteor scatter method appear to closely match the hardware value during the summer months, but differ by ∼3 ns in the winter months. The seasonal variations in the different method estimations will be discussed in Section 5.3. The differences between the estimates from different methods will be further examined in Section 4.

Not included in this figure are other known location estimates made using irregularities generated by the HAARP ionospheric heater, because these t_{diff} estimates were made while the radar was operating on the 14.7–15 MHz frequency band. The known location method returned t_{diff} estimates of −348.7 ± 1 ns for 21 Feb 2017 and −349.9 ± 1 ns for 26 Sep 2017 (both along $1\frac{1}{2}$-hop F-region propagation paths). These estimates are consistent over the course of the year and differ by ∼1–2 ns from the hardware value.

4. Sensitivity of elevation angles and geolocation to t_{diff} variations

In the three example comparisons presented above, the differences in the estimates of t_{diff} between the different methods were often as large as 5 ns. However, without knowing how big an influence differences of this size make to the determination of the elevation angle it is difficult to assess their significance and whether uncertainties in t_{diff} of this magnitude are acceptable. It is also unknown at this time which of the methods provides the most reliable estimates of the true unknown t_{diff} value. Hence, it is important that future analysis improves the understanding of the differences between the methods. Given this present uncertainty, it is important to evaluate the effect that the observed level of variations in t_{diff} have on the elevation angle distribution, and consequently on the geolocation of backscatter. The level of disagreement between different methods complicates the decision of the choice of t_{diff} values for a particular radar, and makes it difficult to know if we can average the estimates from different methods and different times, or whether we need to choose the estimates for one particular method at high precision.

One way of comparing the effect of the different t_{diff} estimates from different methods is to examine the different elevation angle and virtual height variations of intervals of ground scatter determined using the different t_{diff} values. Fig. 6 uses an adaptation of the ground scatter method as a validation check for the different t_{diff} estimates obtained for CVW in early 2014 from the hardware file, the known location method, the meteor scatter method, and the VHM comparison method. The panels on the left (a, c, e, and g) show the elevation angle calculated using each t_{diff} value (specified on the vertical axis), as a function of range and universal time on 23 April 2014. This data is limited to the ground scatter returned from the 10.3–10.8 MHz band along beam 9. The main band of single-hop F-region ground scatter between ∼0700 UT and ∼1000 UT is characterised by gradually decreasing elevation angle with increasing range, as is expected. However, as the elevation angle reaches zero, the band is suddenly interspersed with aliased elevation angles close to 40°. The amount of aliasing varies with t_{diff}. We discuss the problem of aliasing later in this section, as adjusting t_{diff} to remove all aliasing does not necessarily result in a better estimate of t_{diff}.

The ground scatter in the black box, covering 13:00–14:15 UT and slant ranges between 600–1680 km (combined for all beams) is used to determine the virtual height distributions in the right-hand panels (b, d, f, and h). The time and range selection of this box was made to include a large span of distances for backscatter with a 1-hop F-region propagation path. Assuming that the returning ground scatter in this frequency band has a similar ionospheric refraction height, it is assumed that the best t_{diff} estimate will result in a virtual height distribution that is constant as a function of range. These panels show that the known location method (panel (b)) has the flattest virtual height variation with range. However, this method also has the greatest amount of aliased data. The other estimates more quickly develop a tail that shows increasing virtual height with increasing slant range. Notwithstanding this, the virtual height distributions for all 4 methods are similar.

Another way of assessing the effect of different errors in t_{diff} is to examine the effect of changing t_{diff} on the range-elevation angle

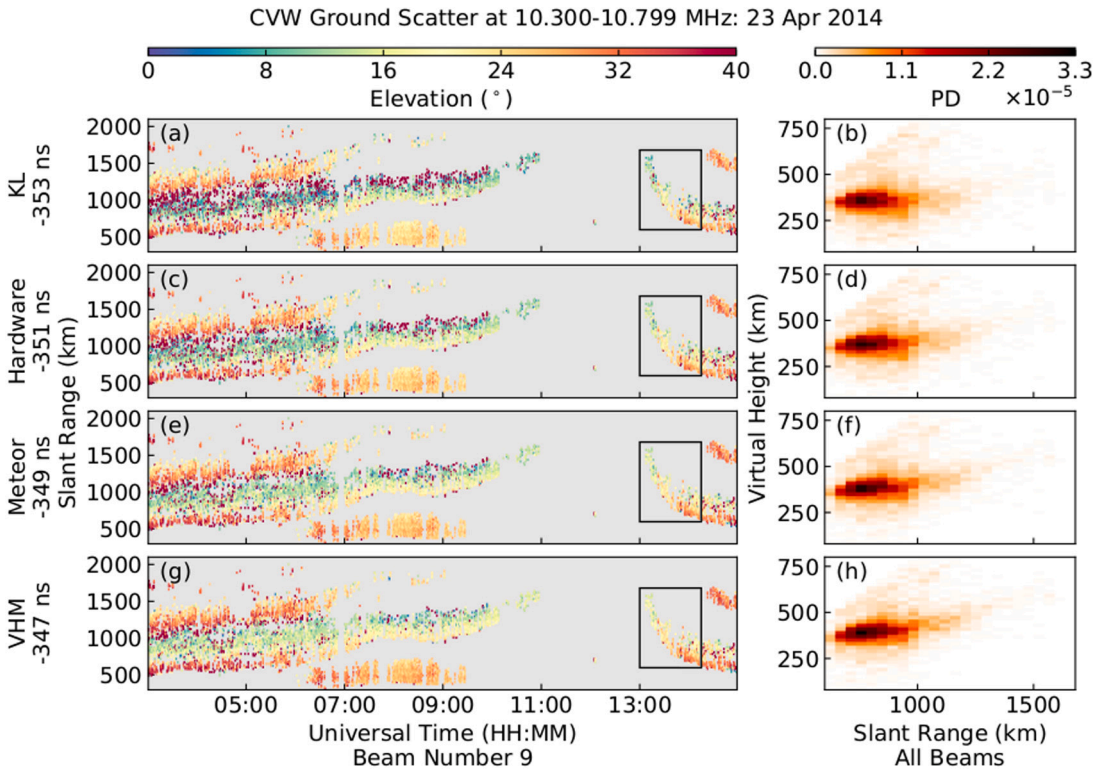

Fig. 6. Comparison of elevation angles and virtual heights determined using the t_{diff} estimates from the known location method (KL), the hardware value, the meteor scatter method, and the virtual height model comparison method (VHM) for the CVW radar. The left column (panels a, c, e, and g) shows the elevation angles along beam 9 for the 10.3–10.8 MHz frequency band as a function of slant range and universal time. The right column (panels b, d, f, and h) contains 2D histograms of the probability density of ground scatter in 45 km slant range bins by 20 km virtual height bins for the times and locations marked by the black boxes in the left-hand column (only for all radar beams).

distribution for a particular radar at a particular time. Fig. 7 presents nine range-elevation angle probability distributions for the INV radar for the interval 2009 to 2013 inclusive. This relates to the grey shaded region highlighted in Fig. 3(a). The nine panels show the distributions for nine different t_{diff} values ranging from −8 ns to 8 ns, in 2 ns intervals. The black and white dashed line illustrates the variation from the Chisham VHM, for reference.

Comparing the different panels in Fig. 7 shows how the elevation angle distribution changes in 2 ns t_{diff} steps, with the bulk of the distribution decreasing in elevation angle as t_{diff} increases. In addition, as t_{diff} increases the size of the aliased elevation angle population (that which appears around 40° close to the maximum measurable elevation angle), increases significantly. As an example of the effect of changing t_{diff} on this distribution, the high probability region at a range of 1000 km varies from ~14° to ~27° for a t_{diff} of −8 ns, to ~5° to ~19° (with a large aliased distribution between 36° and 40°) for a t_{diff} of 8 ns. This difference is approximately a change of ~ 0.5° for a 1 ns change in t_{diff}.

The Chisham VHM variation presented in Fig. 7 should not always be assumed to be an accurate representation of the expected average elevation angle variation as it was determined using a single auroral zone radar (SAS) for which the reliability of the t_{diff} calibration at the time was uncertain. Notwithstanding this point, it represents a consistent marker when comparing the distributions for the different t_{diff} values. The distribution matches closest with the VHM variation for a t_{diff} of −6 ns (Fig. 7(b)), as expected given the red histogram in Fig. 3(b). Fig. 3(b) shows that the other methods determine the value of t_{diff} to be ~0 ns. Fig. 7(e) shows that this t_{diff} estimate has an elevation angle distribution that is shifted slightly from the VHM variation and a small but significant population of aliased elevation angles.

To assess the effect that these changes in elevation angle have on the geolocation of backscatter targets, the data shown in Fig. 7 are plotted as distributions in slant range-ground range space (as in Chisham et al., 2008) in Fig. 8. In all the panels, the main part of the distribution is almost linear in slant range-ground range space, closely following the Chisham VHM variation shown by the black and white dashed lines. The spurs diverging from this main distribution, that occur for t_{diff} values of −2 ns and above, relate to the aliased part of the distribution and are clearly erroneous. The variations in the main distribution are very subtle on the scale shown. For the E and nearer F-region scatter (up to slant ranges ~1000 km) there is little difference in the distributions with changing t_{diff} (except for the aliased distribution). The differences between the distributions increase with slant range, being ~10 km ground range for each ns of t_{diff} error at slant ranges around 2000 km. At this range, a t_{diff} error of ~5 ns would lead to a geolocation error of 50 km. However, although the ground range distribution varies only gradually with t_{diff}, estimates of virtual height will show a larger relative change.

Many of the distributions in Figs. 7 and 8 are characterised by aliased values near the maximum measurable elevation angle, which can be an indication of an error in t_{diff}. Adjusting t_{diff} to remove all aliasing has at times been suggested as a method to estimate t_{diff}. However, there are other factors that also lead to aliasing in these distributions. Statistical random errors in the measured phase differences between the return signals measured by the main and interferometer antenna arrays can also lead to a significant aliased distribution.

Fig. 9 presents the results of a simple simulation that demonstrates how random statistical errors in the measured phase difference lead to aliasing. The black line shows a potential elevation angle distribution of echoes from a particular range gate. This is equivalent to a horizontal cut through a theoretical version of the panels shown in Fig. 7 at a fixed range gate. In this example, the elevation angle population has been simply modelled as a Gaussian probability distribution with a mean of 17° and a standard deviation of 6°. To illustrate the effect of random phase errors, this elevation angle distribution was randomly sampled one million times. Each randomly-sampled elevation angle was

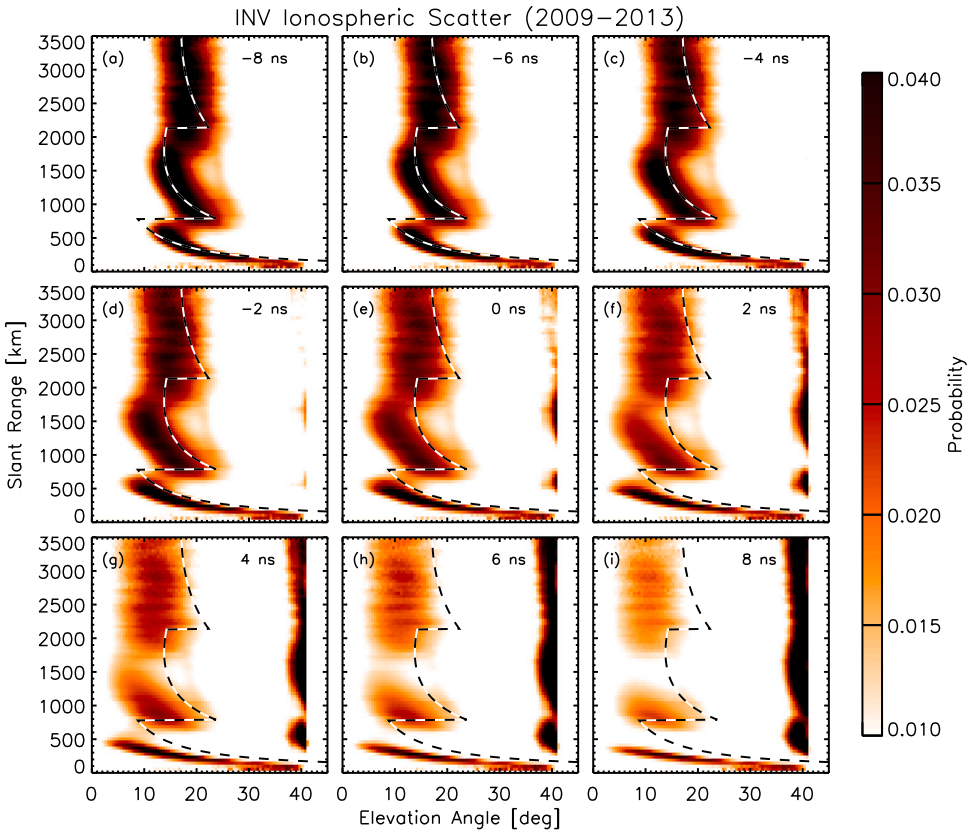

Fig. 7. Two-dimensional probability distributions of elevation angle and slant range observed by the INV radar for ionospheric scatter from the years 2009–2013, for a frequency band centred around 12.4 MHz using the fixed t_{diff} values given in the upper right of each panel. The black and white dashed line shows the variation of the Chisham VHM. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

converted to the phase difference value that would have given rise to that elevation angle. A random phase error value was then added to that phase difference, and the phase difference was converted back to an elevation angle value. The phase error values were randomly selected from a Gaussian probability distribution with a mean of 0° and a standard deviation of 30°. The orange line in Fig. 9 shows the resulting elevation angle distribution after the addition of the phase errors. This new elevation angle distribution is characterised by a strong aliased peak between 36° and 41°. This simple example shows how aliasing can occur in these elevation angle distributions without there being any error in t_{diff}, and implies that adjusting t_{diff} to completely remove aliasing in these distributions is an unreliable way to estimate t_{diff}.

In summary, the analysis in this section shows that the typical uncertainties in t_{diff} that are seen between methods (~5 ns) result in an acceptable geolocation error (~50 km), as long as the aliased elevation angle population is removed. With future fine tuning of these methods, reducing this uncertainty to ~1 ns will improve the geolocation accuracy. The analysis also shows that the major issue affecting geolocation is the aliased population of elevation angles that leads to severe errors. It is likely that no elevation angle measurement within 5° of the maximum measurable elevation angle can be wholly trusted. The maximum measurable elevation angle for a particular radar, beam, and frequency can be determined using the formula presented in Chisham (2018), but only for radars for which the interferometer antenna array is offset from the main antenna array solely in the boresight direction. For more complex antenna layouts, the formula can be deduced from Shepherd (2017).

5. Discussion

This section presents a discussion of issues that may affect the empirical estimation, interpretation, and implementation of t_{diff}, and whether it is possible to resolve the problems they present.

5.1. Does t_{diff} vary with operational frequency?

The first application of the meteor scatter method by Chisham (2018) included the estimation of t_{diff} at a 3-month resolution over a 9-year interval for the SAS auroral-region radar for four different 2 MHz frequency bands stretching from 10 to 18 MHz. Their analysis suggested that the estimated value of t_{diff} varied with frequency, sometimes being different by up to 5 ns for a 2 MHz frequency difference, although more typically being around 2 ns. Subsequent (unpublished) analyses of data from multiple radars (not shown), using this and the other methods described above have shown that the estimates of t_{diff} often show differences of this size when evaluated for different frequency bands. However, these differences are not always consistent, and for some radars and time intervals there is little or no evidence of any variation of t_{diff} estimates with frequency.

If t_{diff} truly varies with frequency it will have a significant effect on the determination of elevation angles within the SuperDARN data analysis software. Consequently, it is important to understand the origin of t_{diff} frequency differences at different radars, and to determine if they require a change to the SuperDARN hardware files and to the SuperDARN Radar Software Toolkit.

Potential origins for the observed variations in estimates of t_{diff} with frequency are:

1. time averaging in the t_{diff} estimation processes.
2. ray path propagation differences.
3. different signal travel properties in the radar cabling and electronics for different frequency signals.

Possibility (1) could occur because averaging measurements over intervals of a day or longer would mask any variations of t_{diff} with time of day. These differences could exist due to either changes in

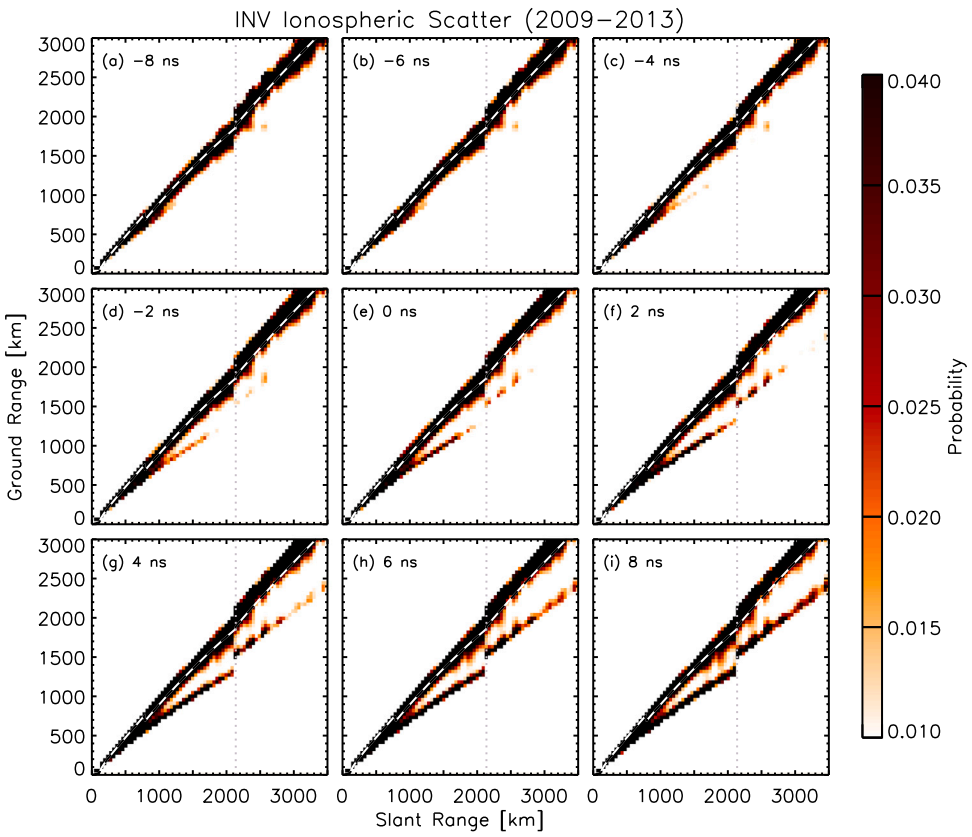

Fig. 8. Two-dimensional probability distributions of slant range and ground range observed by the INV radar for ionospheric scatter from the years 2009–2013, for a frequency band centred around 12.4 MHz using the fixed t_{diff} values given in the upper left of each panel. The transition between $\frac{1}{2}$-hop and $1\frac{1}{2}$-hop F-region ionospheric scatter is assumed to occur at 2137.5 km slant range (vertical dotted line) for consistency with the Chisham VHM (black and white dashed line). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

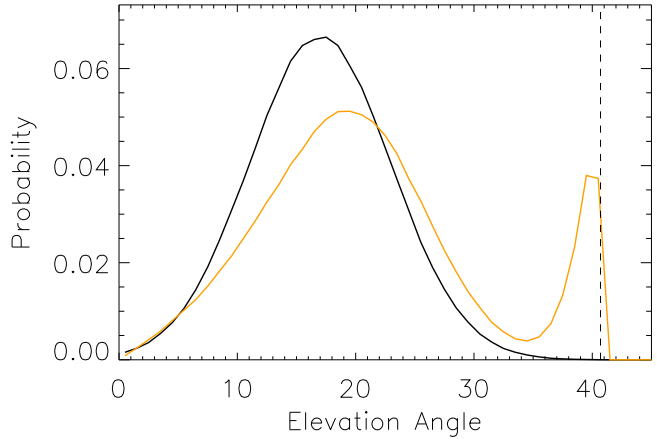

Fig. 9. Simulated elevation angle distributions showing the effect of measurement errors in the phase difference between the signals from the main and interferometer antenna arrays. The black line represents a typical elevation angle distribution, modelled as a Gaussian probability distribution with a mean of 17° and a standard deviation of 6°. The orange line represents the probability distribution that results from adding random errors to the phase difference values. The error values are randomly selected from a Gaussian error distribution with mean 0° and standard deviation 30°. The vertical dashed line represents the value of the maximum measurable elevation angle. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

ionospheric propagation between day and night having a different effect on the t_{diff} determination method assumptions, or be the effect of diurnal temperature differences on the transmission properties of the radar cabling. As the operational frequency variations in many radars are organised so that one frequency operates during the day, and another during the night (to optimise ionospheric backscatter by adapting to the changing ionospheric conditions), then it may be that the t_{diff} at one frequency represents the 'daytime' t_{diff} and the other represents the 'nighttime' t_{diff}, and that there is actually no variation in t_{diff} that is solely the result of a frequency change.

A simple analysis of INV data from 2017 (not shown) shows that this radar operates in two frequency bands through both daytime and nighttime (10.3–10.8 MHz and 12.2–12.8 MHz). Hence, any t_{diff} variation with frequency at this radar for this time cannot be due solely to the different frequencies being used at different times of day. Calculating t_{diff} for the different frequency bands using data from all local times could demonstrate that the frequency dependence is not caused by the diurnal variations. However, to investigate if this is the origin for these variations with frequency, t_{diff} estimates would need to be made separately for daytime and nighttime intervals for a constant frequency of operation. Although even then, these variations may still be a result of the estimation methodology.

Possibility (2) has parallels with the above possibility, but the differences would be solely associated with the changes in ionospheric propagation resulting from changing frequency and not due to diurnal differences. The VHM comparison method could be affected in this way, as it involves a comparison with a 'one fits all' VHM, which in reality will vary with frequency, as well as time of day and season.

For possibility (3), there are several potential sources for a frequency-dependent t_{diff} value, although they will typically be radar site specific as there is no standard design for the radar electronics and the physical layout of SuperDARN antenna arrays. Because t_{diff} results from a difference in the electrical paths that signals travel along from the two respective antenna arrays to the digitising electronics, any

frequency dependent components that appear in only one of the paths will result in a frequency dependent t_{diff} value. One likely source are RF components (switches, amplifiers, filters, etc.), which have delays that can vary by several ns over the ~8–20 MHz frequency range of SuperDARN radars. Another possible source comes from the very large difference in cable lengths that communicate signals from the respective antenna arrays to the radar electronics. These RF cables can exceed 300 m in length and differences of up to nearly 150 m are possible. Any frequency dependence that exists in the velocity factor of these cables would then lead to a frequency-dependent t_{diff} value.

The big question remaining is whether or not to produce averaged t_{diff} values using measurements from all frequencies, or different t_{diff} values for each frequency band. At present, the SuperDARN data analysis software only allows a single t_{diff} for each radar at a particular time; this would restrict the validity of any elevation angle determination solely to that frequency range. But presenting a single value that is averaged over many frequencies may reduce the accuracy of the elevation angles that are determined, although this is dependent on the size of the frequency variation. Further analysis of this issue is required for a full understanding of the significance of the differences with frequency. However, the issue can only be fully resolved by significant changes to the radar analysis software.

5.2. Does t_{diff} vary with beam direction?

For the SuperDARN radars, adjusting the relative phases of the signals transmitted by the different antennae allows the focusing of the radiated power into narrow beams in the horizontal plane, the direction of which can be changed by varying the phasing. One result of the geometry of the antenna layout is that the antenna radiation pattern for beams aligned close to the boresight direction is different from those beams closer to the edge of the radar fields-of-view. Hence, it has been suggested that t_{diff} may vary with beam direction.

Fig. 10 presents the variation of the distribution of elevation angle with beam for two SuperDARN radars (CLY and CVE), using t_{diff} values estimated from the meteor scatter method. The elevation angle distribution variation at both of the radars is characterised by the same features:

1. A broad band that is relatively constant with changing beam number, that extends from approximately 10° to 30° in elevation angle. This represents the expected peak in the elevation angle distribution given the vertical profile of the SuperDARN antenna radiation pattern.
2. Curved bands at high elevation angle. These are the result of aliased elevation angle measurements and are located immediately below the maximum measurable elevation angle (see Section 4 and Figs. 7, 8, and 9). The maximum measurable elevation angle varies with frequency (hence the multiple bands for the multiple frequencies used at CVE), and with look direction (hence the curved variation with beam number) (see Chisham, 2018, equation 8).

The variation for the CLY radar (Fig. 10(a)) shows a consistent variation of the core elevation angle variation with beam. The only exception is a faint reduction in probability in an arc between 10° and 20° elevation angle that peaks in the boresight direction (like the high elevation angle arc). The reason for this reduction, and its variation with beam, is unclear. The variation for the CVE radar (Fig. 10(b)) is generally consistent with beam but with slightly more beam-to-beam variability. The same faint reduction in probability is clearer in this panel but extends over a wider range of elevation angles from 10° to 30°.

Given the lack of any major variations in the core elevation angle distribution with beam, we can conclude that any effect is not large enough to be a major problem and that data from all beams can be combined in any t_{diff} estimation process, and that any t_{diff} value is valid for all beams. However, the variation of elevation angle distribution with beam should be checked for each radar to be sure that this result can be consistently assumed.

5.3. What are the origins of temporal variations in the t_{diff} estimates?

It is clear from looking at time series of t_{diff} estimates, such as those presented in this paper, and in previous method papers (e.g., Chisham, 2018; Ponomarenko et al., 2018), that the estimated value of t_{diff} does not often remain constant with time as expected, but appears to vary in different ways on different timescales. It is important to understand the origins of these temporal changes, and which are the result of methodological uncertainty and which represent an accurate reflection of changes to the hardware t_{diff}.

The observed temporal variations in the t_{diff} estimates can be grouped into three general categories:

1. Small point-to-point fluctuations.
2. Sharp (and occasionally large) changes.
3. Seasonal variations.

Longer-term trends and variations in t_{diff} may exist, such as solar cycle variations, or changes due to the degradation of the radar hardware over time. However, it is difficult to identify such variations clearly without longer datasets.

The small point-to-point fluctuations are clearly evident in the comparison plots presented in Figs. 3, 4, and 5. These fluctuations are likely to be the result of random errors due to the varying amounts of data being used to determine each of the estimates. In some cases, such as the 1-minute E-region method estimates presented in Fig. 3, there are occasionally estimates that are significantly in error. This is likely due to only small amounts of E-region echoes being available on these days. These point-to-point fluctuations are unlikely to represent true changes in the actual t_{diff} but are likely a result of uncertainties in the t_{diff} estimation techniques. Hence, averaging these estimates over longer intervals of time, or increasing the size of the time interval used to determine each estimate, will most likely increase their reliability and consistency. Hence, we conclude that these fluctuations are the result of methodological uncertainty.

Sharp changes in t_{diff} are clearly seen in the t_{diff} variation for the INV radar, presented in Fig. 3. This figure shows several large jumps in t_{diff} that are consistently identified at identical times by both the meteor scatter and E-region methods. Time series of t_{diff} estimates from some other radars have shown similar features. The origin of these jumps can be explained by hardware and software changes to the radar, as discussed in Section 3.1. Hence, we conclude that such changes represent an accurate reflection of changes to the hardware t_{diff}.

Seasonal variations in the t_{diff} estimates can be seen in the CVW comparison results (Fig. 5) and for some years in the INV comparison results (Fig. 3). In addition, Chisham (2018) observed an apparent annual variability in the t_{diff} estimates for the SAS radar using the meteor scatter method, with the value of t_{diff} varying by up to ~4 ns between summer and winter. The major question is whether these gradual changes do actually represent a change in the hardware t_{diff} or whether they reflect a changing systematic error in the estimated value that is method-dependent.

If the seasonal variation is a true variation in t_{diff} then it would imply that there was a temperature effect on the transmission properties of the radar cabling and electronics. If this temperature dependency was the only reason for seasonal variation, the t_{diff} determined by every method would show the same seasonal variation. The seasonal variations in t_{diff} evident in the results presented in this paper are in the opposite sense for different methods (as can be seen for INV in Fig. 3 when comparing the meteor scatter and VHM comparison methods). Hence, although there may be systematic seasonal variations in t_{diff},

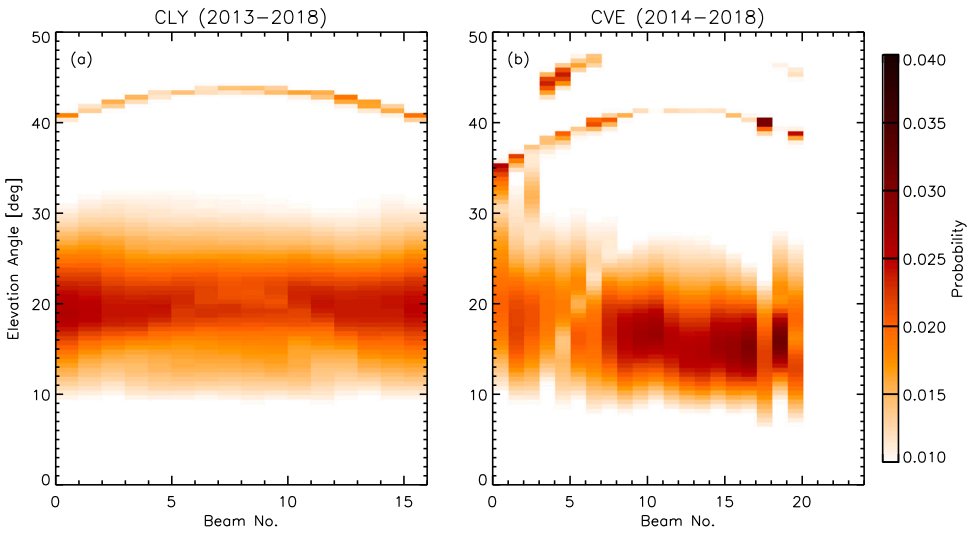

Fig. 10. Two-dimensional histograms of the beam number and elevation angle distributions observed by the CLY and CVE radars for all scatter using common mode control programs from the years given at the top of each panel, using t_{diff} values from the meteor scatter method. Note that CLY is a 16-beam radar while CVE typically operates on only 20 of its 24 beams for common mode control programs.

there are also seasonal dependencies in the different estimation techniques that prevent the clear determination of any systematic changes in t_{diff} that may be associated with ambient temperature variations.

The consequence of this interpretation is that the seasonal variation for each method needs to be fully understood to determine if the estimated value is more reliable at a particular time of year, or whether an averaged value across the year provides the most reliable estimate. For the meteor scatter method, it is well known that the peak heights of meteor occurrence distributions vary throughout the year in response to seasonal variations in atmospheric density (Clemesha and Batista, 2006; Liu et al., 2017). This may result in the level of applicability of the assumptions in the meteor scatter method (such as that of straight line propagation) changing with season. For the VHM comparison method, the measured range-elevation angle profile is shifted by changing t_{diff} until it matches best with the Chisham VHM. The Chisham model was developed using a combination of elevation angle data from all times of year. However, the ionospheric electron density profile at a particular location, and hence the likely variation of virtual height with range, varies with time of year. In addition, the temporal distribution of data available in the original development of the VHM was variable throughout the year. Hence, the seasonal variations in t_{diff} estimated by this method are very likely to be a result of seasonal variations in virtual height that are not captured in the Chisham VHM. Hence, we conclude that the observed seasonal variations in the different estimation methods are partially the result of methodological uncertainty. Understanding the seasonal variations, and developing methods to remove them should allow for more accurate estimates of t_{diff}.

5.4. What is the effect of range errors on the estimation of t_{diff} ?

As discussed in Section 2, systematic errors in the radar range determination will affect the accuracy of all the methods, but to different degrees. All the methods use the range measurement as part of the estimation methodology. However, for those methods that typically use echoes from ranges beyond the E-region (such as the known location, VHM comparison, and ground scatter methods), any systematic error in the range determination is likely to be very small in comparison to the range values being used. However, for those methods using echoes from the nearer ranges (such as the meteor scatter and E-region methods), any systematic error in the range determination may be a significant source of error. However, the origin of any large systematic errors in the range determination is presently unknown, and requires further investigation.

As the meteor scatter method is likely to see the largest effect of any range errors, this method is used to discuss the potential size and significance of these range errors on the t_{diff} estimates, and how it may be possible to estimate and mitigate the effect of this error. As discussed in Section 2.2 and in Chisham (2018), the meteor scatter method estimates t_{diff} by adjusting height distributions of meteor echoes from multiple ranges in response to a changing t_{diff} and determining the t_{diff} value for which there is the least variation in the peak heights of these distributions. As discussed by Chisham and Freeman (2013) (and references therein), the peak height of the echo distribution for the HF operational frequencies typically used by the SuperDARN radars is expected to be in the range of ~95 to ~105 km. Hence, if the peak height is determined to be significantly outside of this range, there is very likely a systematic range error for that radar. Knowledge and estimation of range errors can be used to improve the calibration of other t_{diff} determination methods, and of geolocation in general.

Fig. 11 presents the effect that systematic range errors would have on the t_{diff} value estimated by the meteor scatter method for measurements in two frequency bands (10.3–10.9 MHz and 14.6–15.0 MHz) for the CVW radar, averaging over all the estimates from 2013 to 2018 inclusive. Fig. 11(a) presents the changing estimate of t_{diff} as the range error changes, whereas Fig. 11(b) presents the changing value of the peak height of the meteor echo distributions. The black symbols and error bars represent the variations for the 10.3–10.9 MHz frequency band, whereas the orange symbols and error bars represent the variations for the 14.6–15.0 MHz frequency band. The straight lines represent linear fits to the variation with range error. The vertical dotted line represents the value for zero assumed range error, as was the case for the meteor scatter method estimates shown in Fig. 5. Here, the average estimated t_{diff} value of −349 ns for the 10.3–10.9 MHz band matches the results presented in Fig. 5.

In Fig. 11(b), the peak height value in the case of no range error (indicated by the vertical dotted line) is ~113 km for the 10.3–10.9 MHz frequency band, and ~109 km for the 14.6–15.0 MHz band. Such a difference between frequency bands is expected due to the meteor echo height ceiling effect that results in the shifting of the peak of the meteor echo distribution to lower altitudes for increasing operational frequency. However, both these peak height measurements are higher than those typically expected for these frequency bands (highlighted by the grey shaded region in Fig. 11(b)). For the interferometer layout

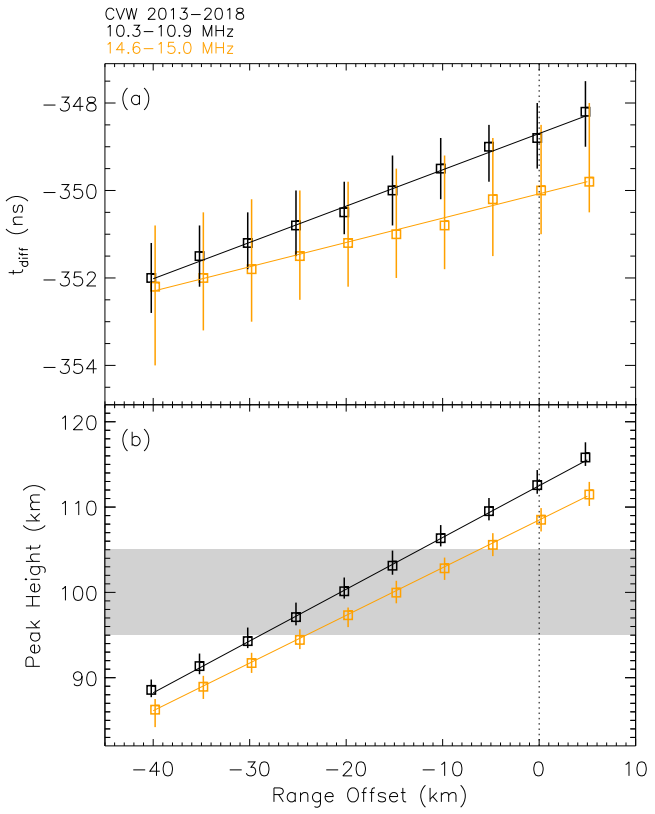

Fig. 11. The effect of errors in radar range determination on (a) the estimated value of t_{diff} using the meteor scatter method, and (b) on the peak height of meteor height distributions. The estimates shown were determined for the CVW radar averaging data covering the years 2013 to 2018 inclusive. Results are presented as symbols and error bars for two frequency bands, 10.3–10.9 MHz (black) and 14.6–15.0 MHz (orange). The straight lines represent linear fits to the observed variations. The grey shaded region in panel (b) highlights the meteor peak height region typically expected for HF operational frequencies. The vertical dashed line represents the case of zero range error that is typically assumed. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

at CVW (where the interferometer antenna array is behind the main antenna array), introducing a negative range offset to compensate for any systematic error results in a reduction in both peak height and the value of t_{diff}. Fig. 11 presents the variation of both peak height and t_{diff} for a series of range errors 5 km apart. It is clear that both vary linearly as a function of this range offset, with the peak height for both frequency bands decreasing by ~0.6 km for each km change in range, and the t_{diff} estimate decreasing by ~0.06–0.09 ns for each km change (depending on frequency). For the case presented here, assuming a systematic range error of approximately -20 km results in peak height values close to those expected, as well as a change of ~1–2 ns in t_{diff}.

This discussion highlights how systematic errors in the range determination shift the peak height of meteor echo height distributions, and as a consequence, the estimate of t_{diff}. However, it also shows how these systematic errors can be estimated by fine-tuning the analysis by assuming that the peak heights should match theoretical or otherwise-observed meteor peak heights.

6. Summary and recommendations

Improvements to the calibration of SuperDARN interferometers can significantly improve the accuracy of the determined elevation angles, and hence, of the geolocation of SuperDARN scattering locations. This paper has shown how this calibration can be satisfactorily performed purely using existing historical SuperDARN data sets, without the need for engineering measurements or access to the radars. The paper can be summarised as follows:

- Five different complementary methods for determining t_{diff} have been presented.
- Comparisons of t_{diff} estimates using the different methods are roughly consistent (mostly within ~5 ns). However, it must be stressed that there is presently no consensus as to which of the methods is the 'best' or 'most accurate'.
- Errors in t_{diff} of ~5 ns lead to only small ground range errors at near ranges and ~50 km error at ranges around 2000 km.
- Aliasing in the elevation angle distribution is a major problem. Elevation angles within 5° of the maximum elevation angle should not be trusted or used.
- t_{diff} varies with the operational frequency of the radar.
- t_{diff} does not vary significantly with beam direction.
- Large, abrupt changes in t_{diff} are associated with changes in the radar hardware or operating software.
- Seasonal changes in t_{diff} are partially related to the shortcomings of the determination methods, as they are method dependent. It is unknown if the true value of t_{diff} changes significantly between hardware changes as a result of seasonal temperature variations.
- Different channels on Stereo radars may have different t_{diff} values.
- Systematic errors in the radar range determination result in small errors in t_{diff} determined using the meteor scatter method.

As a result of the findings presented in the paper, the authors make the following recommendations for scientists using SuperDARN elevation angle measurements:

- For high-spatial resolution analysis, the use of elevation angle measurements is recommended over virtual height models.
- Do not always assume the t_{diff} values in the radar hardware files are appropriate for an interval of data. The above methods can be used to estimate t_{diff} for smaller intervals of time, to provide more accurate elevation angles.
- Always assess the uncertainty in elevation angle determinations that is introduced by the uncertainty in t_{diff}.
- Do not use elevation angles within 5° of the maximum elevation angle for a given radar, beam and frequency.
- Always include values of t_{diff} in any publications that use elevation angles, as well as the details of the calibration methods and any associated software used to estimate this t_{diff} value.

As a result of the findings presented in the paper, the authors make the following recommendations for SuperDARN radar operators and engineers:

- Any changes to radar hardware should be followed by a re-estimation of the t_{diff} value, through either engineering methods or the post-processing methods presented in this paper. If this is not immediately possible, the end time for the old value should be specified and elevation angle data after this time should not be used until a new value of t_{diff} has been estimated.
- Any of the different t_{diff} methods presented here can be used to monitor the health of the radar system. Monthly checks are recommended to look for sharp deviations from the previous value that may imply that there are hardware problems at the radar site. Yearly checks are recommended to check for any longer-term drift in the t_{diff} value.
- If analysis shows a t_{diff} variability with season, care should be taken to not over-specify the variability. In the absence of hardware changes or failures yearly averages are likely to suffice.

The work in this paper has highlighted the following areas for future work:

- Complete the validation of t_{diff} estimates using all methodologies.

- Understand the error sources for the different t_{diff} estimation methods, and the reasons for the different estimates between methods. This will help to identify which of the methods provides the most accurate estimate of t_{diff}.
- Determine the optimum way of combining t_{diff} estimates from different t_{diff} estimation methods, and of averaging these estimates over time.
- Determine how best to deal with t_{diff} variations with frequency, and whether this requires a change to the radar hardware file and the SuperDARN data analysis software.

Acknowledgements

The authors acknowledge the use of SuperDARN data. SuperDARN is a collection of radars funded by the national scientific funding agencies of Australia, Canada, China, France, Italy, Japan, Norway, South Africa, UK, and United States. The raw SuperDARN data are available from the British Antarctic Survey (BAS) SuperDARN data mirror (https://www.bas.ac.uk/project/superdarn). The operations of the SuperDARN Canada radars are supported by the Canada Foundation for Innovation, Innovation Saskatchewan, Canada, and the Canadian Space Agency. GC is supported as part of the BAS Polar Science for Planet Earth Programme, funded by the UK Natural Environment Research Council (NERC) as part of United Kingdom Research and Innovation (UKRI). AGB is supported by the Chief of Naval Research. EGT thanks the National Science Foundation, USA (NSF) for support under grant OPP-1836426. AM thanks the Programme National Soleil-Terre de l'Institut des Sciences de l'Univers (PNST/INSU) funded by CNRS, France, CNES, France and CEA, France and the Institut Paul Emile Victor (IPEV) for supporting French SuperDARN activities. SGS acknowledges support from NSF, USA grants AGS-1341925 and AGS-1934997.

References

André, D., Sofko, G.J., Baker, K., MacDougall, J., 1998. SuperDARN interferometry: Meteor echoes and electron densities from groundscatter. J. Geophys. Res. 103, 7003–7015.

Burrell, A.G., Milan, S.E., Perry, G.W., Yeoman, T.K., Lester, M., 2015. Automatically determining the origin direction and propagation mode of high-frequency radar backscatter. Radio Sci. 50, 1225–1245. http://dx.doi.org/10.1002/2015RS005808.

Burrell, A.G., Perry, G.W., Yeoman, T.K., Milan, S.E., Stoneback, R.A., 2018. Solar influences on the return direction of high–frequency radar backscatter. Radio Sci. 53, 577–597. http://dx.doi.org/10.1002/2017RS006512.

Burrell, A.G., Yeoman, T.K., Milan, S.E., Lester, M., 2016. Phase calibration of interferometer arrays at high-frequency radars. Radio Sci. 51, 1445–1456. http://dx.doi.org/10.1002/2016RS006089.

Ceplecha, Z., Borovicka, J., Elford, W.G., Revelle, P.O., Hawkes, R.L., Porubcan, V., Simek, M., 1998. Meteor phenomena and bodies. Space Sci. Rev. 84, 327–471.

Chisham, G., 2018. Calibrating SuperDARN interferometers using meteor backscatter. Radio Sci. 53, 761–774. http://dx.doi.org/10.1002/2017RS006492.

Chisham, G., Freeman, M.P., 2013. A reassessment of SuperDARN meteor echoes from the upper mesosphere and lower thermosphere. J. Atmos. Sol.-Terr. Phys. 102, 207–221.

Chisham, G., Lester, M., Milan, S.E., Freeman, M.P., Bristow, W.A., Grocott, A., McWilliams, K.A., Ruohoniemi, J.M., Yeoman, T.K., Dyson, P.L., Greenwald, R.A., Kikuchi, T., Pinnock, M., Rash, J.P.S., Sato, N., Sofko, G.J., Villain, J.P., Walker, A.D.M., 2007. A decade of the Super Dual Auroral Radar Network (SuperDARN): scientific achievements, new techniques and future directions. Surv. Geophys. 28, 33–109. http://dx.doi.org/10.1007/s10712-007-9017-8.

Chisham, G., Yeoman, T.K., Sofko, G.J., 2008. Mapping ionospheric backscatter measured by the SuperDARN HF radars - Part 1: A new empirical virtual height model. Ann. Geophys. 26, 823–841.

Clemesha, B., Batista, P., 2006. The quantification of long-term atmospheric change via meteor ablation height measurements. J. Atmos. Sol.-Terr. Phys. 68, 1934–1939.

Greenwald, R.A., Frissell, N., Larquier, S., 2017. The importance of elevation angle measurements in HF radar investigations of the ionosphere. Radio Sci. 52, 305–320. http://dx.doi.org/10.1002/2016RS006186.

Hall, G.E., MacDougall, J.W., Moorcroft, J.P., Manson, A.H., Meek, C.E., 1997. Super dual auroral radar network observations of meteor echoes. J. Geophys. Res. 102, 14603–14614.

Liu, L., Liu, H., Chen, Y., Le, H., Sun, Y., Ning, B., Hu, L., Wan, W., 2017. Variations of the meteor echo heights at beijing and mohe, china. J. Geophys. Res. Space Phys. 122, 1117–1127. http://dx.doi.org/10.1002/2016JA023448.

McDonald, A.J., Whittington, J., d. Larquier, S., Custovic, E., Kane, T.A., Devlin, J.C., 2013. Elevation angle-of-arrival determination for a standard and a modified SuperDARN HF radar layout. Radio Sci. 48, 709–721. http://dx.doi.org/10.1002/2013RS005157.

Milan, S.E., Jones, T.B., Robinson, T.R., Thomas, E.C., Yeoman, T.K., 1997. Interferometric evidence for the observation of ground backscatter originating behind the CUTLASS coherent HF radars. Ann. Geophys. 15, 29–39.

Nishitani, N., Ruohoniemi, J.M., Lester, M., Baker, J.B.H., Koustov, A.V., Shepherd, S.G., Chisham, G., Hori, T., Thomas, E.G., Makarevich, R.A., Marchaudon, A., Ponomarenko, P., Wild, J.A., Milan, S.E., Bristow, W.A., Devlin, J., Miller, E., Greenwald, R.A., Ogawa, T., Kikuchi, T., 2019. Review of the accomplishments of mid-latitude Super Dual Auroral Radar Network (SuperDARN) HF radars. Prog. Earth Planetary Sci. 6, 27. http://dx.doi.org/10.1186/s40645-019-0270-5.

Papadopoulos, D., Bernhardt, P.A., Carlson, H.C., Gordon, W.E., Gurevich, A.V., 1990. Haarp, Research and Applications. Report No. ADA355641, Naval Research Lab., Washington, DC.

Ponomarenko, P., Nishitani, N., Oinats, A.V., Tsuya, T., St.-Maurice, J.P., 2015. Application of ground scatter returns for calibration of HF interferometry data. Earth Planets Space 67, 138. http://dx.doi.org/10.1186/s40623-015-0310-3.

Ponomarenko, P.V., S. Maurice, J.P., Hussey, G.C., Koustov, A.V., 2010. HF ground scatter from the polar cap: Ionospheric propagation and ground surface effects. J. Geophys. Res. 115, A10310. http://dx.doi.org/10.1029/2010JA015828.

Ponomarenko, P., St.-Maurice, J.P., McWilliams, K.A., 2018. Calibrating HF radar elevation angle measurements using e layer backscatter echoes. Radio Sci. 53, 1438–1449. http://dx.doi.org/10.1002/2018RS006638.

Ribeiro, A.J., Ruohoniemi, J.M., Baker, J.B.H., Clausen, L.B.N., d. Larquier, S., Greenwald, R.A., 2011. A new approach for identifying ionospheric backscatter in midlatitude SuperDARN HF radar observations. Radio Sci. 46, RS4011. http://dx.doi.org/10.1029/2011RS004676.

Robinson, T.R., Yeoman, T.K., Dhillon, R.S., Lester, M., 2006. First observations of SPEAR-induced artificial backscatter from CUTLASS and the EISCAT svalbard radars. Ann. Geophys. 24, 291–309.

Ruohoniemi, J.M., Baker, K.B., 1998. Large-scale imaging of high-latitude convection with Super Dual Auroral Radar Network HF radar observations. J. Geophys. Res. 103, 20797–20811.

Shepherd, S.G., 2017. Elevation angle determination for SuperDARN HF radar layouts. Radio Sci. 52, 938–950. http://dx.doi.org/10.1002/2017RS006348.

Stubbe, P.H., Kopka, H., Lauche, H., Rietveld, M.T., 1982. Ionospheric modification experiments in northern Scandinavia. J. Atmos. Terr. Phys. 44, 1025–1041.

Yeoman, T.K., Chisham, G., Baddeley, L.J., Dhillon, R.S., Karhunen, T.J.T., Robinson, T.R., Senior, A., Wright, D.M., 2008. Mapping ionospheric backscatter measured by the SuperDARN HF radars - Part 2: Assessing SuperDARN virtual height models. Ann. Geophys. 26, 843–852.

Yeoman, T.K., Wright, D.M., Stocker, A.J., Jones, T.B., 2001. An evaluation of range accuracy in the Super Dual Auroral Radar Network over-the-horizon HF radar systems. Radio Sci. 36, 801–813.