Towards an Accuracy Bound for OTHR

Samuel J. Davey†§ and Robert Gardiner-Garden†
† National Security and ISR Division, Defence Science and Technology, AUSTRALIA
§School of Electrical and Electronic Engineering, The University of Adelaide, AUSTRALIA

***Abstract*—Skywave over-the-horizon radar is a sensor modality that uses refraction of radio waves through the Earth's ionosphere to observe at very long range. An OTHR system uses secondary sensors to probe the ionosphere in order to characterise propagation. The accuracy of estimated positions from OTHR depends on both the radar measurements and the system's ability to map the ionosphere using these secondary sensors. We describe a method to build a Cramer Rao lower bound on localisation accuracy for a realistic OTHR network.**

I. Introduction

The ionosphere is a region of the Earth's atmosphere that contains charged particles due to ionising radiation from the sun. The density of these particles varies as a function of altitude as determined by the physics and chemistry of the atmosphere and usually has multiple local maxima. This region refracts radio waves in the high frequency band and transmissions can "bounce" off the bottom side of the ionosphere, creating a propagation path much farther than line-of-sight. This effect is used by skywave over-the-horizon radar (OTHR) to achieve coverage at ranges beyond 1 000 km [5], [6].

The simplest model of the ionosphere is a spherical mirror at a fixed altitude but in reality propagation is much more complex. In particular, there are multiple *layers* within the ionosphere that refract radiated signals. The layers commonly involved with skywave OTHR are referred to as the E layer, the F1 layer and the F2 layer. These different layers facilitate multiple propagation paths between the sensor and the targets and the range illuminated by each path is a function of the radar carrier frequency. At some combinations of range and frequency a single target can appear as a cluster of returns at the receiver [6]. Figure 1 shows an example of ray tracing through a representative ionosphere. In the figure, the E layer is at an altitude of around 100 km and there are two refractions from the F2 layer: one at around 200 km and another at around 250 km. These are referred to as the high and low paths. Rays from both the E and F2-low paths reach ground ranges between around 1250 km and 1750 km. This would result in four two-way paths, referred to as *modes*. Geometry means that paths over higher layers are longer in delay and have a greater elevation angle. The Doppler frequency depends on the projection of the target velocity vector onto the ray path, so an increase in elevation corresponds to a decreased magnitude Doppler shift for a horizontal velocity vector. The elevation angle also affects the difference between the true target azimuth and the steer angle for a linear receiver array due to the coning effect [6]. Figure 2 shows an example of how four modes can appear to the sensor. The figure uses the sensor parameters described later for our simulations: the grid shows the delay and steer angle bins; the start of each mode is marked with a circle and the end with a square.

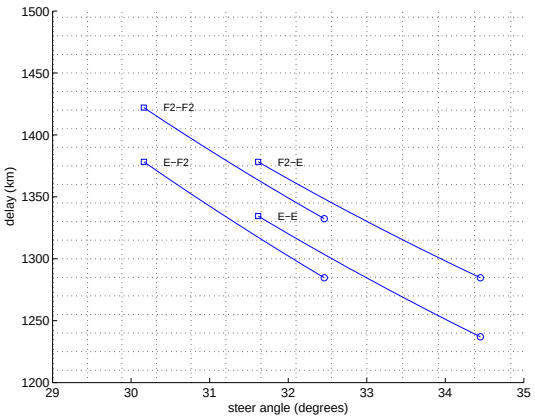

Fig. 2. Delay-steer plot for two ionospheric layers

The interaction of the radar signal with the ionosphere means that the measured delay and azimuth are dependent on the state of the ionosphere, in the simple case, the height of the mirror. In practice the ionospheric state is complex and is part of the weather of the upper atmosphere with structure on various spatial and temporal scales. Transforming measured delay and steer angle into a geographic reference, such as latitude and longitude, requires knowledge of the current ionospheric state. OTHR gathers this knowledge through secondary sensors that probe the ionosphere, referred to as sounders.

The purpose of surveillance is to detect the number of objects in the field of view and to localise them; the quality of achieving these dual tasks can be described in terms of sensitivity and accuracy. Previous research has developed a detailed model of OTHR network sensitivity [7] and our aim here is to develop a similar model of accuracy. The model needs to account for a number of factors: radar-target geometry; sounder geometry; radar measurement accuracy; sounder accuracy; ionospheric variability in space and time; seasonal and diurnal effects. The intention is to derive a Cramer Rao lower bound on accuracy, which is independent from the signal processing algorithms used to analyse the measurement data. Such a bound provides the best possible performance of the system. The bound can be used to evaluate actual achieved system performance and also as a design tool to guide sounder placement or prioritise investment in network enhancements.

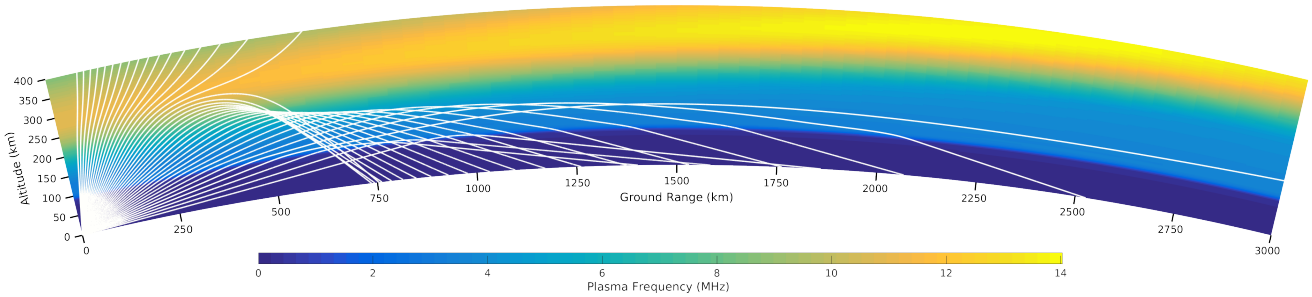

Fig. 1. Ray tracing through an example ionosphere

II. Related Work

The existing literature defines accuracy bounds for OTHR in a few places. These tend to use very simplified models for the ionospheric data and processing.

Rutten [15] used a spherical mirror ionosphere with height varying in a linear deterministic way over time. Direct measurements of height were assumed available. In that context, the ionospheric height and height-rate can be appended to the target state vector and the bound can be derived for a non-linear radar measurement function using the augmented state.

Romeo et al. [14] used numerical ray tracing to derive an accuracy bound but used the international reference ionosphere [1] and assumed that the ionospheric state was perfectly known. Similarly, Tang et al. [16] derive a bound assuming that the set of path transformations are known.

Our aim is to model the ionospheric measurement and modelling in much greater detail because the uncertainty in these components of the OTHR system makes a significant contribution to the overall accuracy.

III. Radar Network

An OTHR network consists of one or more transmit and receive sites along with ionospheric sounders that provide awareness of the propagation medium. We now describe the components of the network and the sources of error in each of these components. Our overall goal is to understand how these errors combine to limit the achievable localisation accuracy.

Figure 3 shows a block diagram illustrating the main data processing components of the radar network. There are two sources of measurement data. The left side of the diagram shows the processing applied to radar data. The initial form of this data is time series of the returned signals across an array of antennae. Transmitted energy is propagated through the ionosphere to the surveillance region, back-scattered, and then propagated again through the ionosphere to the receiver site. The radar data has a non-linear coupled dependence on the physical states of targets and the ionosphere. The right side of the diagram shows the processing applied to ionospheric sounder data. These are a collection of sensors that probe the ionosphere and give measurements that are independent of the targets. Broadly, the network data flow aims to use the sounder data to compensate for the influence of the ionosphere on the radar data and thus estimate target states.

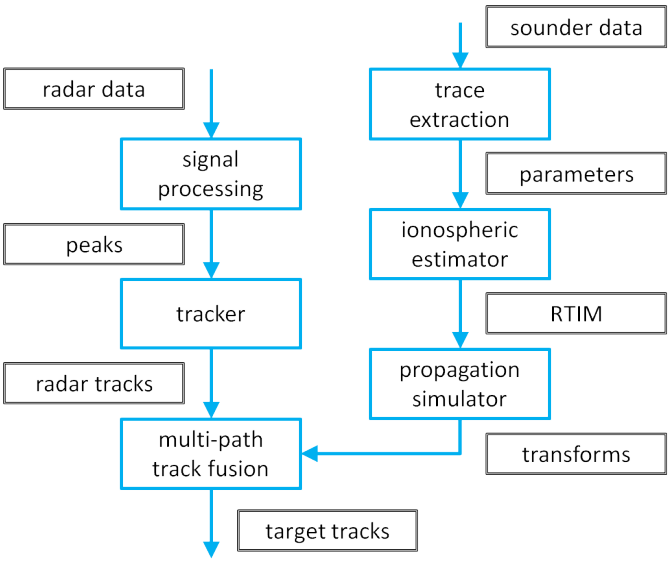

Fig. 3. OTHR network data processing

The first processing block applied to the radar data is labelled *signal processing* and is actually a sequence of numerous algorithms that transform the data into an azimuth-range-Doppler data cube and mitigate clutter, interference and noise. The final stage in this block is a single-dwell detector that condenses the date cube into a collection of point-measurements, often referred to as peaks [3]. Some of these peaks are noisy observations of the true steer angle, group range, and Doppler of a targets, i.e.

$$\boldsymbol{y}_{t,j} = f\left(\boldsymbol{x}_{t,j}, \boldsymbol{i}_t\right) + \boldsymbol{n}_{t,j}^{\text{peak}}, \tag{1}$$

where $\boldsymbol{y}_{t,j}$ is a peak, $\boldsymbol{x}_{t,j}$ is the state of target j at time t, $\boldsymbol{i}_t$ is the ionospheric state, and $\boldsymbol{n}_{t,j}^{\text{peak}}$ is measurement noise. OTHR tend to operate in a regime where the false alarm rate is high, so most of the peaks are not from targets. Of course the data does not identify measurement source. It is tempting to think of the function $f\left(\boldsymbol{x}_{t,j}, \boldsymbol{i}_t\right)$ as a non-linear projection from a point in the target state space to another point in the measurement space, with the noise $\boldsymbol{n}_{t,j}^{\text{peak}}$ a function of waveform parameters. However, in many cases there can be multiple closely separated propagation paths that are unresolved by the radar. In this case a single peak

will combine contributions from each of the unresolved paths and the variability of the peak will depend on the complex interaction of these returns as well as the accuracy implied by the waveform.

The peaks are then processed by a tracking algorithm [2] that aims to identify which peak sequences belong together, how many targets are present, and the states of those targets. The architecture in figure 3 shows a processing stream where the tracker ignores multipath and treats each collection of peaks as an independent target. This is a robust approach although it is possible to integrate the modes into the tracker [4], [8]. The tracks at this point are effectively smoothed peak sequences with low false alarm rate. Assuming that the tracker is unbiased,

$$\hat{\boldsymbol{y}}_{t,j} = f\left(\boldsymbol{x}_{t,j}, \boldsymbol{i}_t\right) + \boldsymbol{n}_{t,j}^{\text{track}}. \tag{2}$$

The noise at the output of the tracker $\boldsymbol{n}_{t,j}^{\text{track}}$ can be notionally decomposed into smoothed measurement noise and tracker introduced error. As discussed above, the measurement noise contains contributions from the interaction of unresolved paths. The two main sources of error potentially introduced by the tracker are mismatch between the target's true dynamics and the tracker's model, and degradation caused by data association errors. Mismatch in the dynamics does not affect the accuracy bound because the tracking model is an algorithmic choice, not an inherent property of the data. In contrast, data association is inherent and existing literature includes this as an information reduction factor on the measurements [9], [11]. This will be explained in more detail in the numerical example.

Accuracy bounds for tracking have been well-studied in recent years. Subject to assumptions about the measurement error the CRLB can be determined for a prescribed target trajectory. For an ensemble of targets the posterior CRLB forms an expectation over possible trajectories. These are known methods in the tracking literature, e.g. [13].

The right-hand path through figure 3 shows the processing applied to ionospheric sounder measurements. The native data product from each sounder is an image of received power as a function of frequency and group delay. Figure 4 shows an example oblique ionogram: the curves of higher power show the frequencies that support propagation to a particular range. The image in figure 4 has markers overlaid that are features from an automatic *trace extraction* algorithm. These features are constrained by a parametric model of electron density as a function of altitude. The output of the sounder processing is a measurement of these parameters. There are two types of error at this point: first there is measurement noise, that is the automatic extraction algorithm can produce a trace that doesn't exactly go where it should; and second is representation modelling error, which is the difference between the true arbitrary electron density profile and the parametric approximation.

The time series of sounder measurements from across a spatial network feeds into a real-time ionospheric model (RTIM), which is a statistical estimation engine that uses the measurements to form a real-time map of the ionosphere.

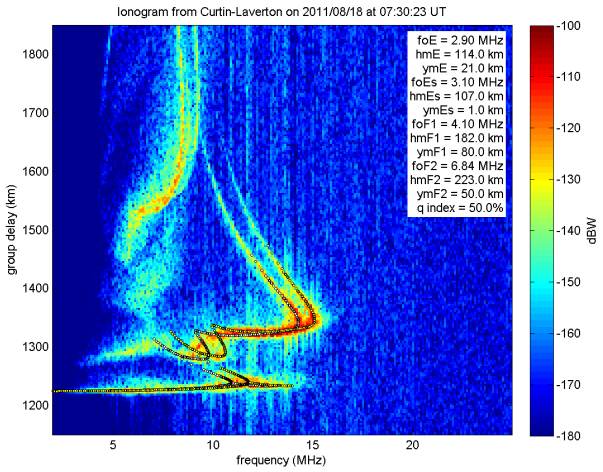

Fig. 4. Example oblique ionogram

This model makes assumptions about the spatial and temporal correlation structure of the ionosphere in order to extrapolate measurements to the position and time where the radar signals propagated and to average out measurement noise. Just like the tracker, the ionospheric model can introduce mismatch error if the assumed spatial and temporal scales are not the same as the true physics of the ionosphere. As before, this is of no concern to a performance bound, which is premised on a valid statistical model. The RTIM operates in the same parameter space as the sounder measurements and so its output error can be notionally divided into estimation error and modelling error. The first is the difference between the best possible modelled ionosphere and the actual reported value, and the second is again the difference between the true profile and the best parametric approximation.

The next processing stage on the right-hand data path is propagation simulation, which uses numerical ray tracing through the ionospheric map to generate transformations between the measurement space and physical geographic space. The ray-tracing can account for magnetic field effects and cross-range gradients. The parameters of the ray-tracing can be selected so that the error introduced by this numerical process is small compared with other system errors. The propagation simulator can also cluster together paths that are unresolvable.

The sounders do not run synchronously with the radar or even at a comparable tempo. Whereas the radar integration time can be of the order of seconds, it takes minutes for a sounder to sweep through the whole HF band. It is possible to update the RTIM on the radar clock and redo the ray-tracing for every dwell but this is computationally impractical and much of the time there would be no new sounder data. Instead the RTIM and transforms are generated at a slower rate, which means there is an additional forecast error.

Finally the radar and sounder data are combined in the form of radar tracks and transforms. Here another layer of association and estimation is required. The multi-path track fusion (MPTF) block in figure 3 is a track-oriented multiple

hypothesis association engine that assigns propagation paths to each radar track, applies the corresponding transformation and then fuses tracks on a common target. The main potential error introduced by MPTF is association error, in this context referred to as miss-moding where the incorrect propagation path is assigned to a track and a large localisation error results from applying the wrong transformation.

IV. Registration Accuracy

The intention of this study is to predict the achievable geo-registration error performance, which is the RMS error between a registered track and the corresponding true target location. The approach is to construct a budget of all the error sources and then analyse how these errors progress through the system to impact track registration. The previous section described our catalogue of errors, we now discuss how to quantify the track registration error. This builds by understanding the accuracy of radar tracks and the ionospheric model before we combine these to form ground registered tracks.

A. Radar Track Accuracy

The progression through the left-hand chain in figure 3 is relatively straightforward. As mentioned earlier, the bound on track error can be derived as a function of the measurement error under assumptions about the target dynamics. In practice this can vary as a function of geometry because unresolved modes are a major contributor to the measurement error. The overall accuracy figure could be determined for a specific trajectory or it could be averaged over an ensemble, depending on the application. [13] describes how to derive the tracking CRLB for known data association. The bound is a covariance matrix that can be calculated using the recursions of the extended Kalman filter where the required Jacobians are evaluated using the true state. For a randomly evolving state, the posterior CRLB is the ensemble average CRLB determined by sampling realisations of the process and taking the mean of their CRLBs.

In he numerical example to follow, we consider the simplest case where the target motion can be modelled as an almost constant rate of change of delay and steer angle. In this case, the tracking problem is linear and the steady state covariance of a radar track is the solution to the Ricatti equation, which can be found numerically.

B. Ionospheric Model

We choose to model the ionosphere at a parametric level rather than numerically approximate a three dimensional electron density map. The numerical example to follow uses a virtual height model, which is the simplest parameterisation and replaces the ionosphere with a mirror for each layer. The virtual height needs to be variable in time, space, and frequency to approach a useful level of accuracy because the ionosphere is neither uniform nor stationary. A more detailed model would be used in practice to capture effects such as gradients in the ionosphere.

The trace extraction algorithm generates parameter fits to the sounder data that use the same parameterisation as the real time model: the role of the ionospheric estimator is to interpolate and extrapolate in space and time. We assume a Gaussian process model over the parameter map with a squared exponential spatial covariance and an exponential decay in time. These models lead to a low-pass filter in space and time, which reduces the noise variance but can also attenuate real small scale ionospheric variability. A higher fidelity description of the physics would need to consider wavelike structures in the ionosphere but this level of detail is beyond our scope.

We assume a large-scale prior distribution on the ionospheric parameters that understands seasonal and solar variation based on historical data. The Gaussian process requires a mean and variance on the parameters: in practice these will not be spatially or temporally uniform. The remaining model parameters are the assumed sounder fit accuracy and the correlation scales of the ionosphere in space and time. Given all of these, standard smoothing results exist for such a Gaussian process [12] and we will use these in the numerical example to follow. The result is a map of variance in ionospheric parameter estimates given the assumed system model.

The next stage on the right-hand data path is the propagation simulator that generates transforms. In the simplest virtual height case, these transforms are analytic non-linear functions and the virtual height accuracy feeds directly into the final registration stage below. In a more general model there is no analytic relationship between the ionospheric model parameter accuracy and the transform accuracy in which case numerical methods are required. An example is the unscented transform, described in more detail below.

C. Mode Association

The MPTF block at the bottom of figure 3 associates together transforms from the high-hand ionospheric model and radar tracks from the left-hand tracker. When the mode structure is complicated, this task can be ambiguous and degrades statistical achievable accuracy. Mode association is a significant source of localisation error because it can contribute large step-changes in position rather than small perturbations. In the tracking-level problem, the effect of association error is handled by an approximation referred to as *information-reduction* [11]. This method inflates the variance of the measurement process to compensate for associate error. It is not clear that such an approach is suitable for the mode association case and this is an important area of further study. The numerical example to follow assumes a single mode and therefore avoids this issue.

D. Registration Accuracy

The final task is to combine the accuracy values for the two processing chains, that is radar tracks and transforms, to obtain the target track accuracy. Even in the simplest case, this is a non-linear transformation problem and cannot be solved

analytically. There are fundamentally two suitable approaches: sigma-point methods, or Monte Carlo sampling.

The unscented transform is an example sigma point method [10] and works by finding a small number of sample points based on a decomposition of the state covariance. It places samples in the state space according to the Eigen structure of the covariance and transforms each sample individually. The output is formed by a weighted sample covariance. This is suitable for the simpler analytic case, such as the virtual height model used in the numerical example. For example, the relationship between height, group delay and ground range (great circle distance) can geometrically be derived as

$$r_{\mathsf{gnd}} = 2r_e \cos^{-1}\left(\frac{r_e^2 + (r_e+h)^2 - (r_{\mathsf{group}}/2)^2}{2r_e(r_e+h)}\right) \quad (3)$$

where r_e is the Earth's radius. In the example we use the unscented transform to compute ground range and azimuth accuracy from delay, steer and height accuracy.

The unscented transform requires us to be able to specify the accuracy of the transform in terms of a small covariance matrix. In the case of a distributed numerical parameter map, this will not be possible and the alternative is to randomly sample errors for an ensemble of ray paths and compute covariances over these paths.

V. Numerical Example

We now present a numerical example using a simplified model. We consider a single radar located at Alice Springs aligned with a boresight North-West. Suppose that the radar system has a single downrange vertical sounder located 700 km from the radar along boresight. We look at the accuracy effect of inserting a second sounder in two potential locations: at 700 km and ± 35 degrees azimuth. The simulation uses the nominal air target detection parameters described in table 1.3 of [6]; a subset of parameters relevant to the example is presented in table I. The table also presents the statistical parameters we assume for the ionospheric prior and measurement accuracy, and the tracker model parameters.

A. Tracker Model

We assume that the standard deviation of the peak location is half of a resolution cell, for example 7.5 km in range. The tracker uses an almost constant velocity model with an assumed 1 g process noise acceleration.

B. Ionospheric Model

The time and space correlation are assumed to be separable

$$\mathrm{cov}\Big[h_m(\mathbf{p},t), h_m(\mathbf{p}',t')\Big] = \mathrm{cov}_m(\mathbf{p},\mathbf{p}')\ \mathrm{cov}_m(t,t'). \quad (4)$$

where height is denoted by $h_m(\mathbf{p},t)$ with m a mode index, $\mathbf{p}$ a 2D position on the earth, and t time. The temporal term is assumed to be exponentially decaying

$$\mathrm{cov}_m(t,t') = \exp\left\{-\frac{1}{\tau}|t-t'|\right\}, \quad (5)$$

where τ is a time constant. The advantage of an exponential model is that is satisfies the decomposition $\mathrm{cov}_m(t_1,t_3) = \mathrm{cov}_m(t_1,t_2)\mathrm{cov}_m(t_2,t_3)$ for $t_1 < t_2 < t_3$. This property allows us to construct recursive estimators that are computationally efficient. However, as discussed earlier, it does not allow for wavelike disturbances that are physically realistic. We adopt a squared exponential covariance kernel for the spatial term

$$\mathrm{cov}_m(\mathbf{p},\mathbf{p}') = \sigma_p^2 \exp\left\{-\frac{1}{2}(\mathbf{p}-\mathbf{p}')^{\mathsf{T}}\mathrm{P}^{-1}(\mathbf{p}-\mathbf{p}')\right\} + \sigma_m^I(\mathbf{p},t)\delta(\mathbf{p}-\mathbf{p}'), \quad (6)$$

where $\sigma_p^2 = 30^2$ km^2 is the prior accuracy and P is a matrix that describes the spatial structure of the ionospheric correlation. The characteristics of the ionosphere vary differently with latitude and longitude: here we use a standard deviation of 10 degrees in longitude and 5 degrees in latitude.

The parameter measurements from the sounders are assumed to be unbiased with a fit error standard deviation of 5 km. We acknowledge that the physics means that height fluctuations are not in reality symmetrical.

TABLE I
Simulation parameters

Radar parameters	
azimuth resolution	0.38 deg
range resolution	15 km
velocity resolution	5 ms^{-1}
revisit period	30 s
Ionospheric parameters	
height prior variance	$\sigma_m^I = 30^2$ km^2
spatial coefficient	$\sigma_m^P = \sigma_m^I$
sounder measurement noise	$\sigma_n^S = 5^2$ km^2
spatial correlation scale	$\mathrm{P} = \begin{bmatrix} 10^2 & 0 \\ 0 & 5^2 \end{bmatrix}$ degrees^2
time constant	$\tau =$
sounder period	5 minutes

C. Localisation Error

Consider a surveillance region spanning ± 45 degrees in azimuth and 1000 km$-$2000 km in range. The accuracy bound at any point in the region is determined by the following steps: transform the geographic location into delay and steer using a known virtual height; determine the mid-point of the great-circle between the radar and the ground location; evaluate the ionospheric parameter accuracy at this midpoint; combine the delay, steer, radar track accuracy and ionospheric parameter accuracy to derive the target accuracy in ground range and azimuth; combine the covariance in ground range and azimuth as an error distance using $d^2 = \sigma_R^2 + R^2\sigma_\theta^2$.

Figure 5 shows the spatial distribution of the error distance in km. The top map shows error for the case of a single-sounder at 700 km on boresight; the middle map places a second sounder at +35 degrees azimuth and the bottom map places a second sounder at -35 degrees azimuth. In the plots, the diamond represents the radar location, a black plus shows the sounder locations and a white plus shows the position of an oblique path passing exactly over each sounder.

The influence of each sounder spreads more at ranges longer than the path directly over the sounder, which is why the first sounder location is short of the mid-point in range. We also note that the worst case localisation accuracy, shown in white, is around 20 km, which is significantly less than the uncertainty in height for these paths. This is a consequence of the geometry. The boresight sounder provides better accuracy in the Northern sector rather than the Western sector because the spatial correlation of the ionosphere is broader in an East-West direction.

VI. Conclusion

Radar performance modelling is useful as a network design tool and to diagnose system performance. For OTHR, the localisation accuracy is dependent on both the tracking performance and on the systems used to monitor the ionosphere. This paper is a start towards a Cramer Rao lower bound on accuracy for OTHR. Future efforts will focus on identifying where this simplistic model is too naive and on developing more realistic representations for these areas.

References

[1] D. Bilitza, "35 years of international reference ionosphere–karl rawers legacy," *Advances in Radio Science*, vol. 2, no. G. 2, pp. 283–287, 2005.

[2] S. B. Colegrove and S. J. Davey, "PDAF with multiple clutter regions and target models," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 39, no. 1, pp. 110–124, Jan. 2003.

[3] S. B. Colegrove, S. J. Davey, and B. Cheung, "Clutter rejection using peak curvature," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 42, no. 4, pp. 1492–1496, Oct. 2006.

[4] S. J. Davey, G. A. Fabrizio, and M. G. Rutten, "Detection and tracking of multi-path targets in over-the-horizon-radar using a group tracking model," *IEEE Transactions on Aerospace and Electronic Systems*, submitted to.

[5] K. Davies, *Ionospheric Radio*. London, UK: Peter Peregrinus Ltd., 1990.

[6] G. A. Fabrizio, *High Frequency Over-the-Horizon Radar*. McGraw Hill, 2013.

[7] D. B. Francis, M. A. Cervera, and G. J. Frazer, "Performance prediction for design of a network of skywave over-the-horizon radars," *IEEE Aerospace and Electronic Systems Magazine*, vol. 32, no. 12, pp. 18–28, 2017.

[8] S. V. Fridman and L. J. Nickisch, "Sifter: Signal inversion for target extraction and regis," *Radio Science*, vol. 39, no. 1, feb 2004.

[9] C. Jauffret and Y. Bar-Shalom, "Track formation with bearing and frequency measurements in clutter," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 26, no. 6, pp. 999–1010, 1990.

[10] S. J. Julier and J. K. Uhlmann, "New extension of the kalman filter to nonlinear systems," in *Signal processing, sensor fusion, and target recognition VI*, vol. 3068. International Society for Optics and Photonics, 1997, pp. 182–194.

[11] T. Kirubarajan, H. Chen, and Y. Bar-Shalom, "Parameter estimation and the crlb with uncertain origin measurements," *Methodology and Computing in Applied Probability*, vol. 3, no. 4, pp. 387–410, 2001.

[12] C. E. Rasmussen, "Gaussian processes in machine learning," in *Advanced lectures on machine learning*. Springer, 2004, pp. 63–71.

[13] B. Ristic, S. Arulampalam, and N. J. Gordon, *Beyond the Kalman Filter: Particle Filters for Tracking Applications*. Artech House, 2004.

[14] K. Romeo, Y. Bar-Shalom, and P. Willett, "Detecting low SNR tracks with OTHR using a refraction model," *IEEE Transactions on Aerospace and Electronic Systems*, vol. preprint, 2017.

[15] M. G. Rutten, "Multipath tracking and track fusion for over-the-horizon radar," Ph.D. dissertation, The University of Melbourne, Aug. 2005.

[16] X. Tang, Q. Wu, R. Tharmarasa, and T. Kirubarajan, "Multipath ML-PMHT for low observable targets," *IEEE Transactions on Aerospace and Electronic Systems*, vol. preprint, 2017.

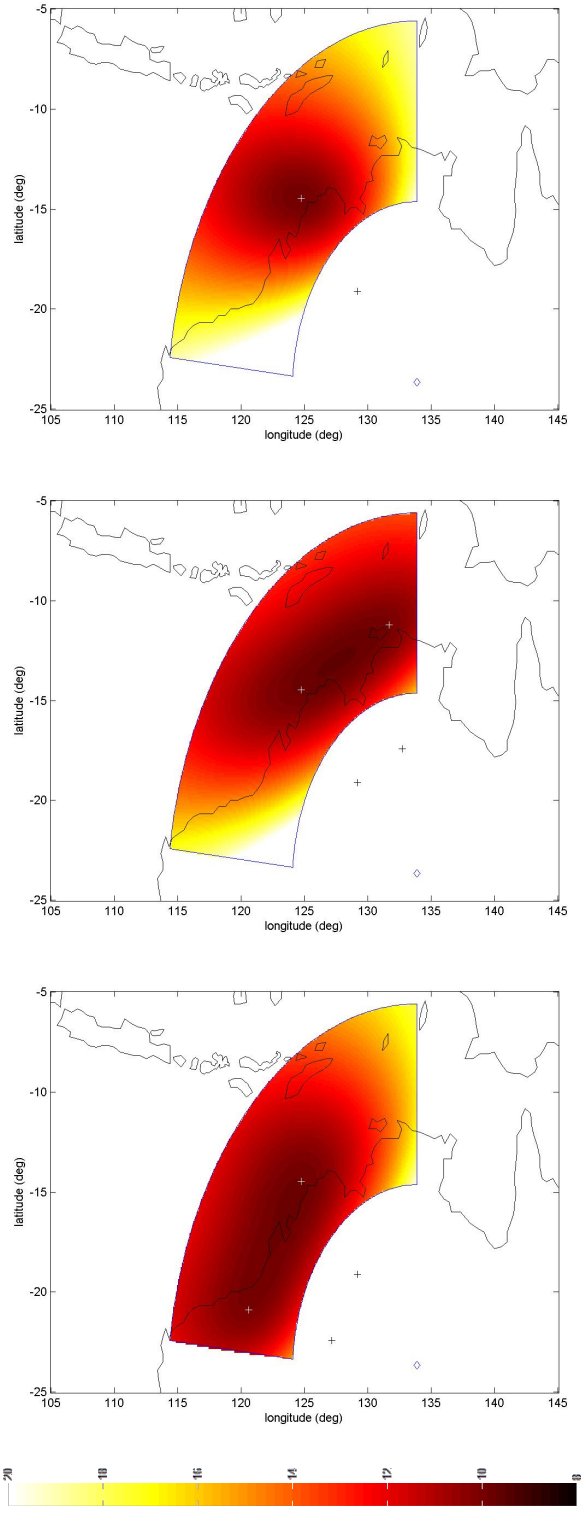

Fig. 5. Example error distance (km) maps