Communications Research Centre

AN INVESTIGATION OF HF DIRECTION-FINDING ACCURACY ON THE CHURCHILL-OTTAWA AURORAL ZONE PATH

by

G.O. Venier

**DEPARTMENT OF COMMUNICATIONS**
MINISTÈRE DES COMMUNICATIONS

CRC TECHNICAL NOTE NO. 702

CANADA

OTTAWA, JULY 1980

COMMUNICATIONS RESEARCH CENTRE

DEPARTMENT OF COMMUNICATIONS
CANADA

AN INVESTIGATION OF HF DIRECTION-FINDING ACCURACY
ON THE CHURCHILL—OTTAWA AURORAL ZONE PATH

by

G.O. Venier

(Radio and Radar Research Branch)

CRC TECHNICAL NOTE NO. 702

July 1980

OTTAWA

TABLE OF CONTENTS

AN INVESTIGATION OF HF DIRECTION-FINDING ACCURACY

ON THE CHURCHILL-OTTAWA AURORAL ZONE PATH

by

G.O. Venier

ABSTRACT

The CRC HF Direction-Finding Array, consisting of a 1181 meter by 236 meter cross, was used to make angle-of-arrival measurements on HF waves propagated over a 1900 km path from Churchill, Manitoba to Ottawa, in May, 1976. This path is of interest because much of it is in the auroral zone where conditions are generally more disturbed than at lower latitudes.

Swept-frequency transmissions allowed the investigation of the variations in angle of arrival of individual propagation modes, while fixed-frequency transmissions permitted a test of two azimuth-estimation schemes. One of these used a phase-front planarity test and averaging of azimuths over relatively frequent samples, while the other was based on the separation of modes by Doppler processing. The former technique was found to provide slightly better accuracy than the latter on this path, and both compared well in accuracy with estimates from separated modes using swept-frequency data.

1. INTRODUCTION

The CRC high-frequency direction-finding array has been used in a number of experiments intended to determine the characteristics of errors in estimates of the bearing of a transmitter, caused by propagation through the ionosphere. In the experiments described here, the transmitter was situated at Churchill, Manitoba, at about 59 degrees N latitude, and at a distance of about 1900 km from the receiving system near Ottawa. Thus, the northern part of the path was in the auroral zone, and the mid-point conditions were probably strongly influenced by auroral effects.

This experiment complements one[1] in which the transmitter was situated at Frobisher Bay, at a slightly higher latitude. That experiment was performed in early winter of 1976 while this one was carried out in the spring of 1976.

A ranging capability in the transmitter-receiver system allowed the separation and identification of various propagation modes, so that each could be studied without interference from the others, and a high level of coherency in the system permitted the measurement of Doppler shift with high accuracy and resolution. The large aperture of the array provided the capability of sampling phase fronts over a large spatial extent.

The results to be presented here include the following:

a)   the azimuth-error statistics of the various propagation modes,

b)   the azimuth-error statistics of a scan-by-scan (a scan is the single sampling of all antenna elements) azimuth measurement technique in which averages of azimuth are taken over all scans in a sampling interval which pass a phase-front planarity test, and

c)   the azimuth error statistics of an estimation scheme based on the computation of a weighted mean of azimuths obtained from Doppler-separated modes. The weights depend on phase-front linearity, elevation angle, and power, all of which correlate with the accuracy of the azimuth estimate[2].

2.   DESCRIPTION OF EXPERIMENT

A description of the CRC High-Frequency Direction-Finding system can be found in reference 3, but some of the important characteristics are included here.

The antenna configuration used is shown in Figure 1. The direction of the transmitter was not far from that of the axis of array 1. Thus, the azimuth measurement depended mainly on the shorter array, while the elevation-angle measurement depended mainly on the longer array. This means that, the azimuth resolution was not as good as that for the Frobisher Bay experiment[1], but the elevation-angle resolution was better.

There were 42 elements in array 1 and 16 in array 2, a total of 58, each feeding its own receiver. These were scanned sequentially, and in-phase and quadrature (I and Q) channel samples, with respect to a reference receiver, were digitized to 12 bits and recorded for each antenna for each scan. The AGC level of each receiver was also recorded to allow amplitude correction. The signal was heterodyned down to essentially zero frequency before sampling, with frequency components both positive and negative in the tens-of-Hertz range. Phase errors caused by the time required to sequentially scan the receivers were corrected when the recorded signal was later processed.

319.72° REL. TO TRUE NORTH

CHURCHILL 327.38° E. OF N.

7.66°

12 SPACES OF 38.1 METERS

1181
METERS

267
METERS

8 SPACES OF 15.24 METERS

CENTRAL SPACE = 22.9 METERS, BOTH ARRAYS

ARRAY 2 16 ELEMENTS

7 SPACES OF 15.24 METERS

236 METERS

EACH ARRAY SYMMETRICAL
ABOUT CENTRE

ARRAY 1 42 ELEMENTS

Figure 1. Antenna Configuration

4

Two types of transmission were used in this experiment, a fixed-frequency signal with a duration of ten seconds, and a swept-frequency signal sweeping through 100 KHz in two seconds. These signals were transmitted once per minute, the fixed frequency signal starting 20 seconds after the beginning of the swept-frequency signal. The sampling rates were 128 antenna scans per second for the swept frequency, and 25.6 scans per second for the fixed frequency. The time required for each scan of the array was about three milliseconds. Rubidium-vapour frequency standards were used in both the transmitter and receiver to provide accurate timing and coherence.

Every 20 minutes, recording was stopped for about three minutes during which time the transmitter sent a long frequency sweep, and the receiving and data processing system processed it to produce an oblique ionogram. These were used to help select frequencies for transmission, and also, in later analysis, to identify the propagation modes present. Larger gaps in the data (8 minutes) occurred once per hour when it was necessary to change tapes.

The experiment was carried out over the eleven-day period from May 11 to May 21, 1976, as shown in Figure 2. Data were processed only for the times indicated with cross-hatching. At other times the received signals were too weak to provide useful data, or problems such as incorrect frequency settings occurred. As can be seen, data from all times of day were analyzed.
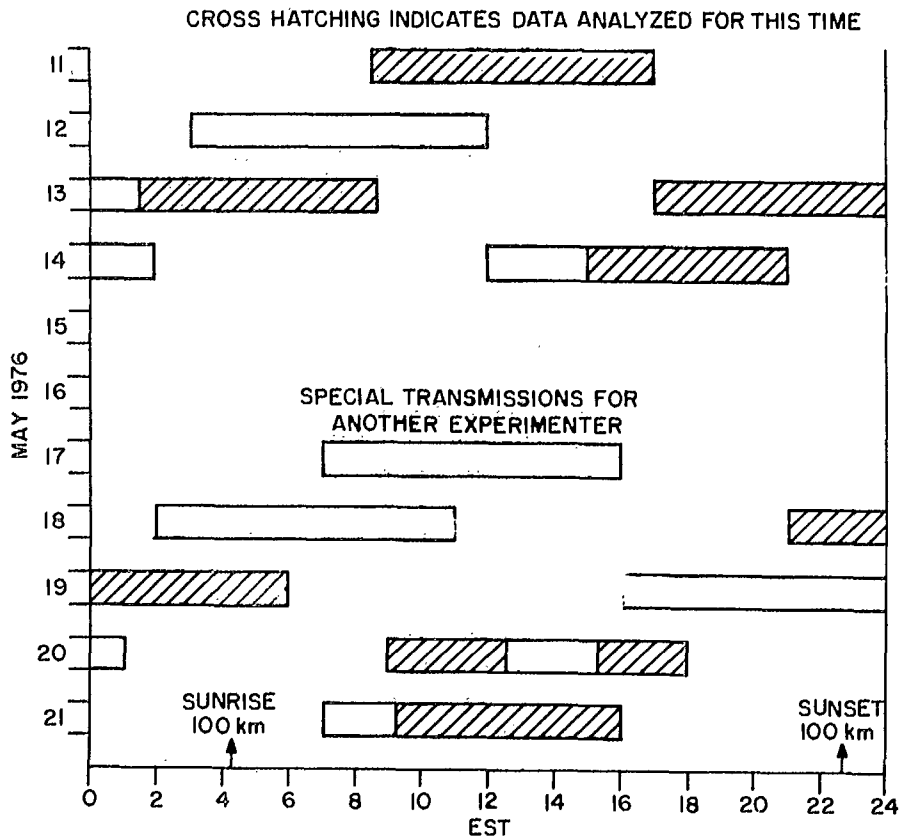
Figure 2. Schedule of Operation

3. METHOD OF ANALYSIS

3.1 FIXED-FREQUENCY DATA

The fixed-frequency operation was included in the experiment to allow the investigation of the Doppler characteristics imparted to the signals by the ionospheric path, and also to allow a study of the short-term characteristics which can be found from a scan-by-scan analysis.

For Doppler processing, the recorded data for each receiver were cosine-squared preweighted, then Fourier transformed to provide frequency spectra, one for each receiver. A complex transform was performed on the I-and Q-channel data over the 256 samples (25.6 samples per second over 10 seconds). An offset of 4 Hz was used between the transmitter and receiver frequencies to prevent the zero of Doppler shift occurring at zero frequency where the receiver introduces additional noise. The resulting 256 complex numbers for each receiver, representing the amplitude and phase of the signal for 256 Doppler shifts uniformly spaced over a 25.6 Hz frequency window, were recorded on a secondary magnetic tape for input to later stages of the analysis. Since the phase of the signal was retained, the direction of arrival could be determined from the relative phases across the array of receivers. Further, this direction-finding process could be carried out independently for each of the 256 Doppler frequency cells. Therefore, if two or more modes were received with different Dopplers, the direction of arrival of each could be determined.

3.1.1 Analysis of Doppler-Processed Tapes

The Doppler-processed tape may be used as input to a number of programs which make use of the Doppler-processed information. These programs compute cone angles for each of the arrays by means of a linear least-squares fit to the measured phases along the array. Cone angle is defined here as the angle between the line of the array and the normal to the phase front, since for a linear array, the possible directions of arrival corresponding to a measurement fall on the surface of a cone. Azimuth and elevation are determined by solving for the intersection of the cones for the two arrays. For each of the least-squares lines computed, the RMS deviation of the phases at each element from this straight line is also computed. The larger this number, which will be referred to here as the RMS phase deviation, the more distorted (from planar) the phase front must be. Since distortions in the phase front result from wave interference[4,5] and generally result in bearing errors, the RMS phase deviation may be used as an indication of the quality of the measurement.

One method for improving azimuth estimates by making use of the Doppler-processed data is referred to here as the weighted-mean method. This method forms a weighted mean of azimuths determined for each of a number of Doppler cells with the greatest amplitudes. The method is described in more detail in reference 2 but a general description is included here.

For each of the 6 Doppler cells with largest amplitudes (this number is a parameter in the program which may be changed), the azimuth and elevation are computed and the azimuth is given a weight which depends on RMS phase deviation across the array, elevation angle, and power. The weight increases

with decreasing RMS deviation, decreasing elevation angle and increasing power. A weighted mean is then computed from the weighted azimuths and is taken as the azimuth estimate. An overall weight, indicating the quality of the azimuth estimate is also computed (this was not indicated in the above reference). It can be used in a decision as to whether or not to accept the result. The overall weight is computed from the sum of the cell weights multiplied by a Doppler-spread weight and a bearing-spread weight. These last two are computed from the standard deviation of Doppler shift and azimuth across the cells used. Large spreads on Doppler and azimuth are an indication of poor conditions caused by either the propagation medium or interference, and result in a low overall weight. The formulas and parameters used in computing the weights for this experiment are given in the Appendix. The parameters were chosen from experience with a shorter, lower-latitude path. The selection was based on consideration of the geometry of the propagation path, the statistics of the received data, and some trial and error. No formal optimization procedure was used. It is likely that a small improvement in accuracy could be obtained for the path in this experiment by a selection of parameters better suited to the conditions.

3.1.2 Scan-by-Scan Analysis of Fixed-Frequency Data

In the Doppler processed data, only one phase is available from each Doppler cell for each antenna element. This phase represents an average over the full ten-second recording time. In this Section we consider the analysis of the signal on a scan-by-scan basis; that is, we look at the phase of each element for each scan of the array.

The antenna elements are sampled sequentially within a scan, and the resulting time delay between element samples produces phase errors. Such errors are frequency dependent, and the corrections are applied in the frequency domain. For this reason, the scan-by-scan data were regenerated by an inverse Fourier transform on the corrected Doppler shift data. (Sample-and-hold amplifiers have since been installed in the system to make the sampling simultaneous, and these corrections will not be necessary in the future).

Since it was assumed that the correlation time of the phase-front distortions was much longer than 0.04 seconds (approximately the inverse of the scan rate), and that Doppler frequencies did not exist outside the Doppler frequency of ± 3.2 Hz, only one quarter of the Doppler cells, centered about zero Doppler, were transformed back to the time domain. This resulted in only 64 uniformly spaced time samples over the 10-second period. That is, the effect was to reduce the scan rate by a factor of four. These data allowed us to look at the phase, and hence angle of arrival, every 0.156 seconds.

Actually only 63 samples were used since the cosine-squared pre-weighting on the original data had the effect of reducing the resultant values at each end of the time window to the levels where quantization errors were significant; the first value was generally below the quantization level and was discarded. Quantization noise on some of the remaining samples may have had a small effect on the results.

The linear least-squares-fit method was employed, as with the Doppler-processed data, to determine azimuth, elevation angle, and RMS phase deviations for each of the 63 scans. These were used in a simple azimuth-estimation

scheme based on averaging azimuth results over those of the 63 scans that had an RMS phase deviation below a threshold. It was assumed that a high RMS phase deviation is an indication of a probable large error.

In the choice of threshold levels for RMS phase deviation, we wish to use an optimum threshold which will keep the single-scan error low, and yet allow enough scans to pass the test to provide a good average. The RMS phase-deviation thresholds used in this analysis were 80 degrees for array 1 and 15 degrees for array 2. These were chosen on the basis of a statistical study of the data, which will be reported elsewhere.

It is of course recognized that for scans taken a short time apart, there will likely be considerable correlation in the phase linearity and also in the angle errors. That is, the angle samples are not all independent, and averaging of a given number of scans will not reduce the error by the factor expected for independent samples. The correlation time of the wave-front shape has been investigated in reference 5. It will be sufficient here to note that the correlation time is on the order of one second although it is usually less in the presence of spread F. When modulated signals are received, this correlation time can be considerably less, as a result of the effect of the modulation on the interference between two or more modes[6]. Further work is required to determine whether this will actually allow more independent samples to be averaged in a given time.

A program called AZPLT6 computes azimuth estimates as described above, and plots the results. One estimate is produced each minute for the ten seconds of data.

AZPLT6 also computes the RMS error from the true azimuth, of the 10-second averaged results over any desired interval (roughly 40 minute intervals were used in this study). Results can be witheld from the RMS-error calculation if a specified number of scans within the 10 seconds did not pass the RMS phase deviation test and were not included in the ten-second average. A criterion of 20 out of 63 scans, for inclusion in the statistics, was used in analyzing the results presented here.

3.2 SWEPT-FREQUENCY DATA

Swept-frequency signals were transmitted to allow separation of propagation modes on the basis of path length, or range. The received signal was heterodyned with a replica of the transmitted signal, giving an output whose frequency was proportional to range. The 100 KHz sweep used provided a range resolution of about 10 microseconds. The scan rate was 128 scans per second in this experiment, providing 256 samples per channel in the two-second transmission time. This resulted in a range window of 2.56 milliseconds, which could be centered at any desired range by control of the delay between the transmitted waveform and the replica used for heterodyning.

Analysis of the swept-frequency data is essentially the same as that of the fixed-frequency data. Whereas a spectral analysis of the latter provides Doppler information, a spectral analysis of the swept-frequency data provides range information, as a result of the correspondence between range and frequency after heterodyning. Thus the first stage of processing is a discrete Fourier

transform as before. The result is now 256 range cells for each receiver, again with complex data, allowing the phase for each receiver at each range to be determined. These are recorded on a secondary magnetic tape which is used as input for other analysis routines.

One such routine provides a display of the amplitude, averaged over all receivers, for each range cell, as a function of time (one per minute). Range and time form the two spatial dimensions of the plot and the amplitude is coded as a character, providing a kind of contour plot. Modes can be identified on these plots by comparison with the ionograms, and the plots provide a history of the ranges for the various modes between ionograms, and are useful in choosing range windows to select individual modes for statistical analysis.

Program AZPLT5 produces an output of azimuth and elevation as a function of time for a selected range window. It computes the azimuth and elevation for the range cell with the greatest amplitude within this window, using the linear least-squares-fit method. If the elevation falls within specified limits and the RMS phase deviation is below a specified values, the computed azimuth is included in a long term average of RMS error which is computed at specified intervals, usually about 40 minutes. If the elevation and RMS-phase-deviation criteria are not satisfied, the range cell with the next highest amplitude is tried until the criteria are satisfied or ten ranges have been tried.

The range window is used to select a particular mode and the RMS-phase-deviation and elevation-angle criteria help to discriminate against interference and ambiguities. Thus, the result should be a good measure of the best azimuth estimate which can be made for the two-second signal when the propagation occurs in only one mode. This is intended, first, as a means of comparing the potential of different modes for azimuth estimates, and secondly, as a reference against which to test possible operational techniques which cannot make use of range to separate modes.

4. EXPERIMENTAL RESULTS

4.1 AZIMUTH-ERROR STATISTICS

4.1.1 Range-Separated Modes

The swept-frequency data were analyzed by the AZPLT5 routine, for each of the modes present, to produce plots of azimuth as a function of time, and to compute RMS errors in azimuth measured over approximately 40 minute periods. Data where interference was evident, or where the signal was very weak, or near the MUF, were discarded. RMS phase deviation thresholds were set at 40 degrees for array 1 and 20 degrees for array 2 to eliminate phase fronts of poor quality. These values were chosen on the basis of experience with the system; they have been found to give good rejection of phase fronts with large errors while not rejecting very many accurate ones.

Figure 3 shows histograms of the RMS azimuth errors, computed over the approximately 40 minute periods, for each mode, over the entire experiment.
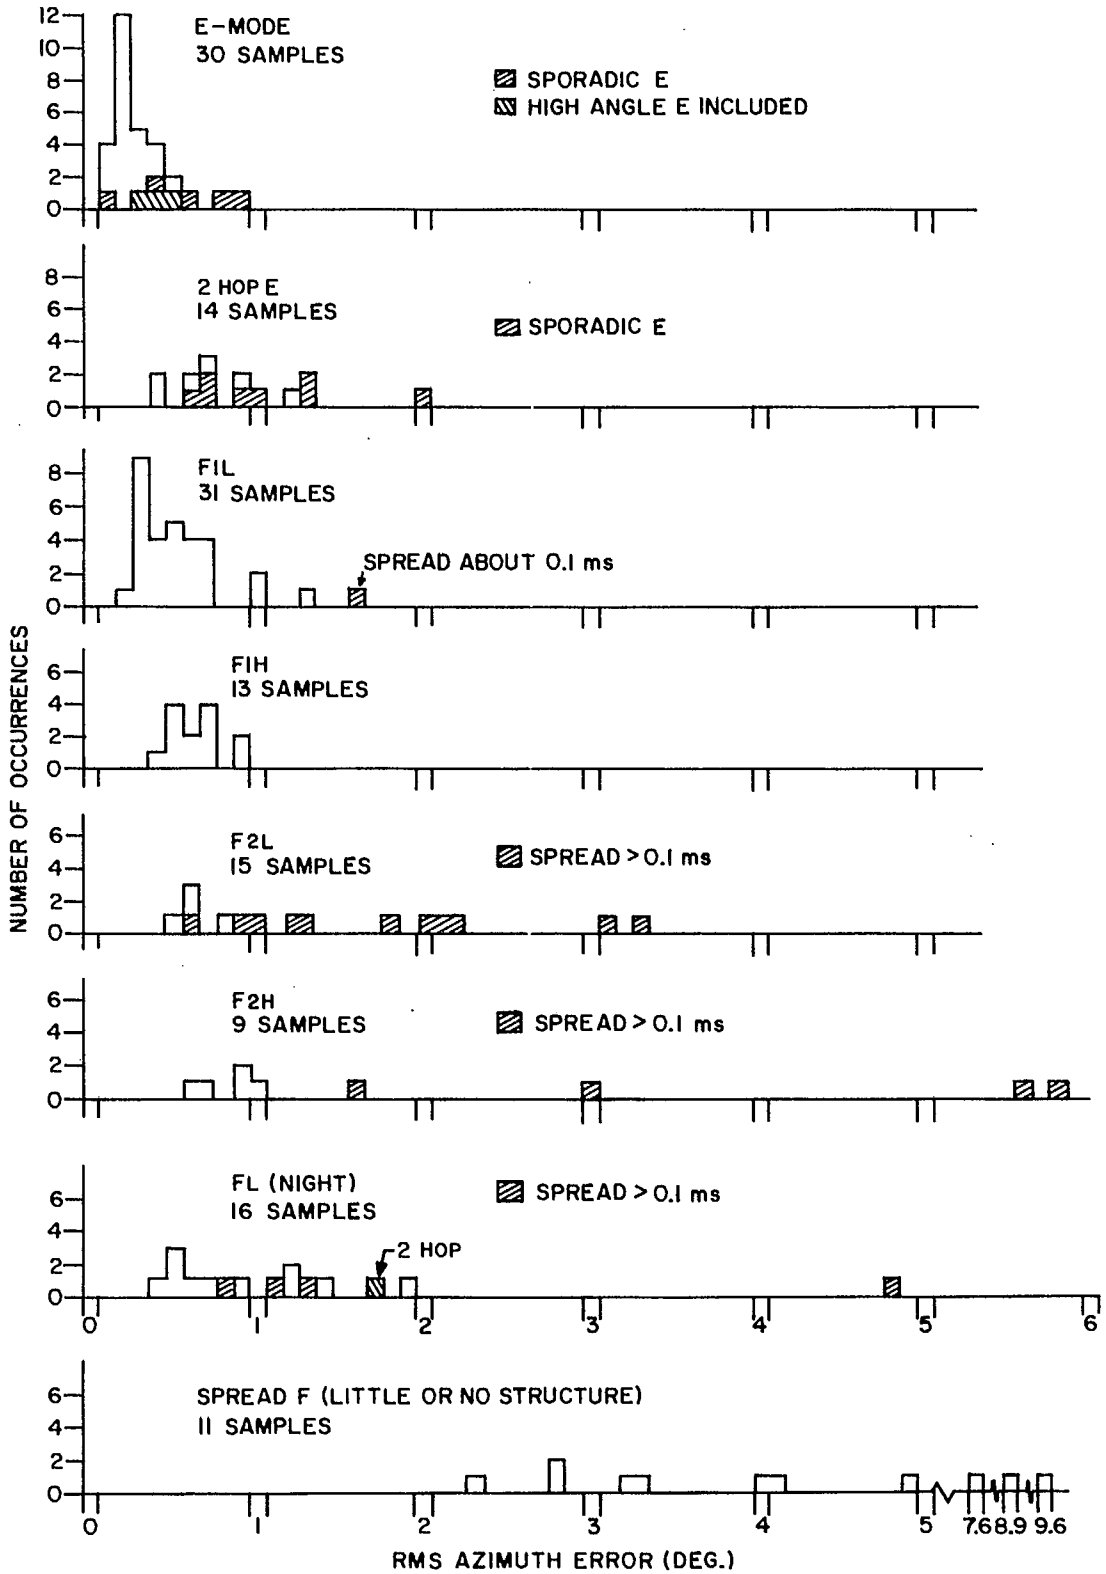
Figure 3. Histograms of RMS Azimuth Errors for Separated Modes

The number of samples shown for each histogram is a rough indication of the relative occurrence of that mode at the frequency of operation. Since this frequency was chosen (and changed from time to time) in order to accumulate data on the various modes, the number of samples may not be a good measure of the likelihood of those modes. Also, some data were discarded because of poor quality, and the measurement periods varied considerably about the nominal 40 minute value.

The E-mode histogram shows extremely good azimuth accuracy, particularly if the samples identified as sporadic E are removed. In three of the samples some returns thought to be high-angle E were received. These samples were near the high end of the histogram, but represent relatively small errors in comparison with other modes.

The 2-hop E mode had somewhat higher errors, and again the sporadic case gave the highest errors.

The low angle F1 mode (F1L) histogram shows reasonably low errors. One of the samples corresponded to a time when there was a range spread of about 0.1 milliseconds on this mode, and this sample was found to have the highest RMS azimuth error of all the samples for that mode.

The high-angle F1 mode (F1H) histogram also shows fairly low errors, although this mode appears to be not quite as good as the low-angle mode.

The F2 modes, both high and low angle, show relatively high errors, but if the samples identified as having a range spread in excess of 0.1 milliseconds are ignored, the errors are not much worse than those for the F1 modes. However, spread in range appears to be a common occurrence on the F2 modes.

Towards sunset the F1 modes disappeared and the F2 modes moved down until there was only a single F layer. We have identified the resulting modes as FL(night) and FH(night).

The histogram for the FL(night) case indicates relatively poor accuracy even when the range-spread samples are ignored. A single two-hop sample is included in this histogram.

When the F layer had such spreading in range that the ionograms showed little evidence of high-and low-angle structure, the condition was called spread F, and the samples from these times are included in the lowest histogram. Extremely large errors are apparent, and it is clear that this is a very poor mode for direction finding. This occurred mainly at night, and another mode was rarely present. At the mid-point of the path at this time of year the sun is below the horizon only for a few minutes each day at a height of 300 km, and for about 6½ hours at 100 km. It was during this latter period that most of the spread F conditions occurred.

These histograms show how the RMS azimuth error measured over some fairly long time interval (about 40 minutes) was distributed. The distribution of individual azimuth estimates (one per minute) is also of interest, and some examples of histograms showing those distributions are presented in the next Section.

4.1.2 Test of Azimuth-Estimation Schemes

Both the weighted-mean method, using Doppler processing, and the scan-by-scan method of azimuth estimation, were tested on the fixed-frequency data from this experiment. As well, the scan-by-scan method was run with very high RMS-phase-deviation threshold, to simulate the results which would be obtained with simple averaging of the scan-by-scan-measurements without the phase-front linearity test. This sometimes resulted in wild points far from the true azimuth. Such points were removed from the average if they were clearly not representative. One might expect an operator to perform such a task on the basis of his observations over some reasonable time period even though he did not know the true azimuth.

In comparing these results with those from the range-separated modes, the difference in the intervals over which measurements were made should be considered. The range-separated results were computed from swept-frequency data which had a duration of two seconds, while the fixed-frequency data had a duration of ten seconds. Therefore if the errors were correlated over a period of less than two seconds, and if they were stationary and unbiased, we should reduce the swept-frequency errors by a factor of $\sqrt{5}$ or 2.24 before comparing them with the fixed-frequency ones. In fact, this correction factor was probably too high since the above conditions are not likely to have been satisfied. It is thought that the factor was near unity in most cases and depended to some extent on conditions. No correction factor was actually applied to the data, but the existence of this effect should be kept in mind, as it offers an explanation of why, in some cases, the fixed-frequency estimation schemes gave smaller errors than the best single mode.

It is assumed that, in most cases, the errors in azimuth resulted from the medium and not from the effects of noise. Thus we do not consider the differences in the coherent integration times of the different methods to have been a significant factor.

Figure 4 shows histograms of azimuth error for a relatively simple situation where only two modes were present, sporadic E and two-hop sporadic E, with the latter a little stronger. The two-hop mode has almost twice the RMS error of the single-hop mode. The scan by scan method gives about the same RMS error as the single-hop mode, although the histogram for the latter looks a little better. After some allowance for the correction factor mentioned above, the scan by scan method probably gives an accuracy somewhere between the two single modes. The weighted-mean method shows larger errors than the scan-by-scan method but has smaller errors than the worst of the single modes, at least if the required correction factor is not too great.

The simple-average method gives relatively poor results, even though four wild points were removed. It is no doubt influenced by the interference between the two modes.

The phase-linearity test appears to have worked quite well in the scan-by-scan method. It probably would have been even better if the more accurate mode had been the stronger. The weighted-mean method, while better than the simple average, was not quite so good. This may be a result of the lack of large Doppler on the E-mode signals; if there is not much Doppler on
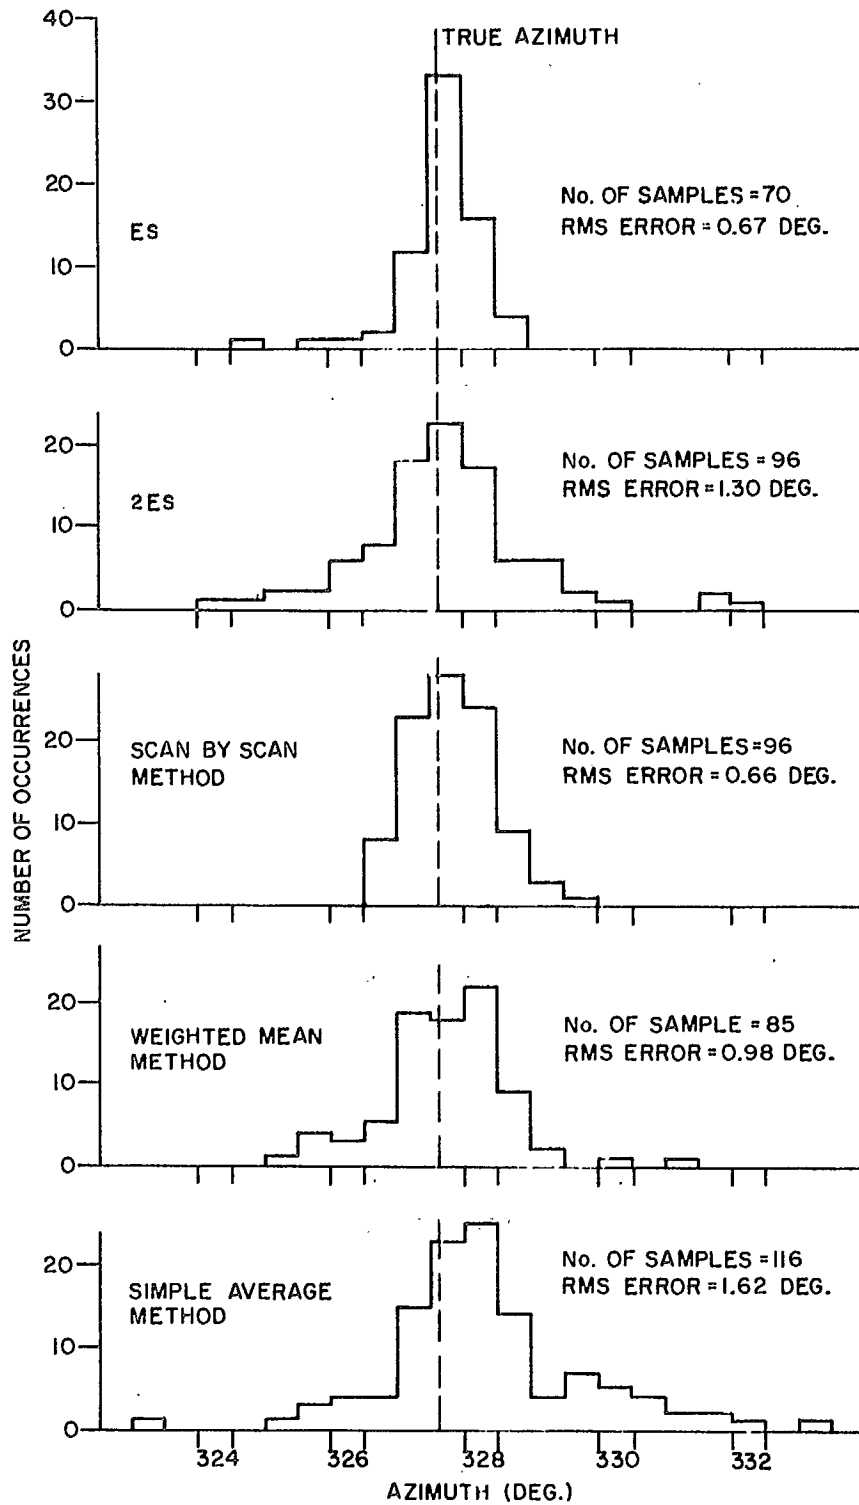
12

Figure 4. Azimuth Histograms for Day 134, 01:24 – 04:19
Es and 2Es Modes Present
No. of Scans Available = 120

the single-hop mode, there will not likely be much on the two-hop, and separation of the modes by Doppler difference is not likely.

Another two-mode case is shown in Figure 5. Here the modes were high- and low-angle night F, the low-angle mode having the greater amplitude. The scan-by-scan method provided an accuracy comparable to that of the stronger and more accurate single mode. The same is true for the weighted-mean method, which shows slightly greater errors than the scan by scan method. Both show considerable improvement over the simple average.

Three modes were present in Figure 6, a very accurate E mode, and high-and low-angle F1 modes. The low-angle F1 mode was very weak, with the other two modes about equal in amplitude. This time the weighted-mean method gave slightly better results than the scan-by-scan method, and both were much better than the simple-average. However none approached the accuracy of the E mode. These conditions should have been relatively good for the weighted-mean method since the accurate E mode generally has very low Doppler, and if there is significant Doppler on the other modes they should be separable. It appears that they did not have large Dopplers however during most of this period. Still, the errors for both the weighted-mean and scan-by-scan methods were rather low.

In Figure 7 four modes were present, sporadic E, low angle F1, and high-and low-angle F2, with the last mode having both o and x paths. These were not separated in the histogram however; the strongest return from a range window encompassing both o and x modes was used. All modes were rather weak with the E_s mode probably the weakest. The scan-by-scan and weighted-mean methods gave about the same accuracy, again considerably better than the simple average. This accuracy was comparable to that of the single modes, with the exception of the very accurate E_s mode. However, this mode was so weak that azimuths could not always be computed from it; only 26 out of 43 of the one-minute intervals yielded an azimuth estimate.

The above histograms show a few typical cases. A summary of the accuracy over the entire experiment is shown in Figure 8, which presents histograms of RMS error measured over approximately 40 minute intervals. These intervals covered all conditions over the entire experiment. Some data were rejected from these histograms when conditions were so poor that very few acceptable estimates were produced. When data from any one method were rejected, they were also rejected for the other methods for the same period. This may have biased the results slightly in favour of better conditions, but the same conditions applied for each of the methods in the histogram results.

It is clear that both the scan-by-scan and the weighted-mean methods give much better accuracy than the simple average. The scan-by-scan and weighted-mean histograms were not greatly different. The latter has a few more very large errors but the medians of these histograms are almost equal. As explained earlier, it should be possible to select better values for the parameters of the weighted-mean method for this path. This should result in a small improvement in accuracy.

The number of estimates shown on each histogram is the total number (one per minute, each representing ten seconds of data) that were used in computing the RMS errors. These values for the scan-by-scan and weighted-mean
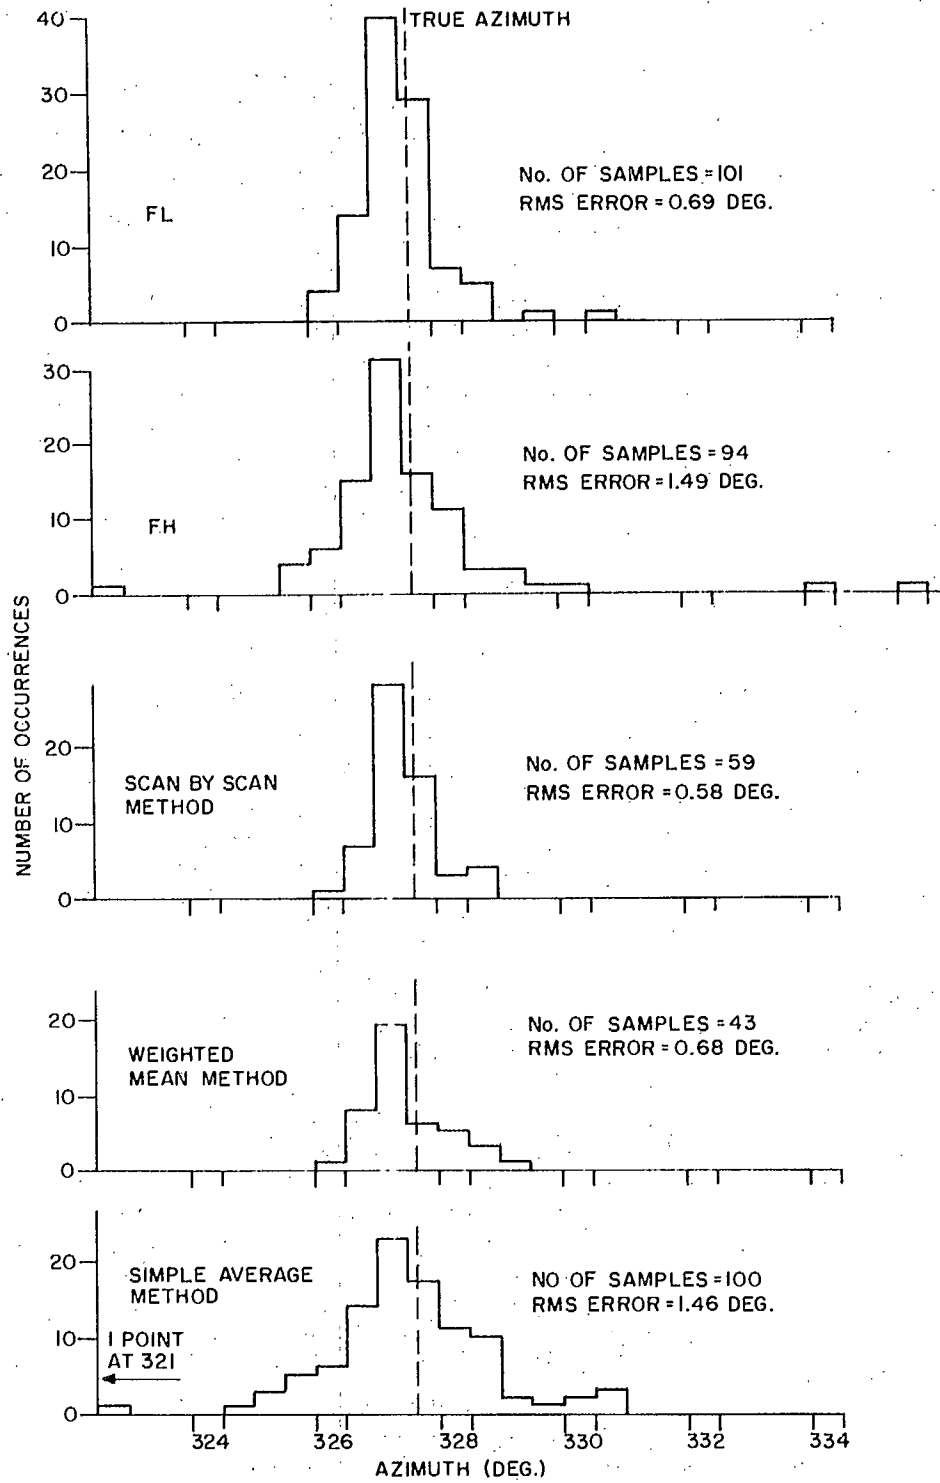
14

Figure 5. Azimuth Histograms for Day 135, 18:04 — 20:29
*FL and FH (Night) Modes*
No. of Scans Available = 106

Figure 6. Azimuth Histograms for Day 141, 09:05 – 11:15
E, E1L, F1H Modes Present
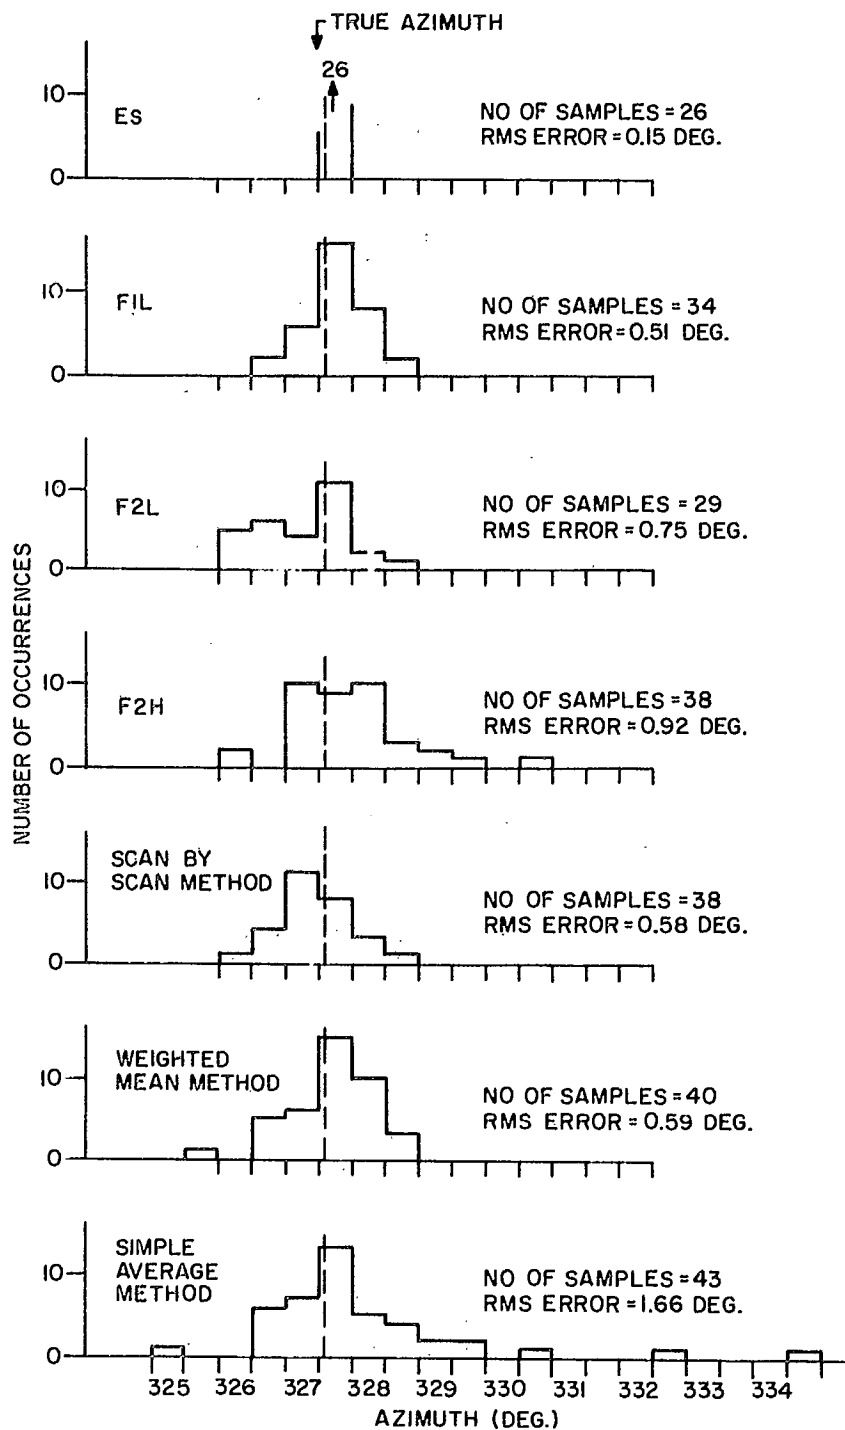No. of Scans Available = 98

Figure 7. Azimuth Histograms for Day 134, 17:04 — 17:55
Es, F1L, F2L F2H (0+X) Present
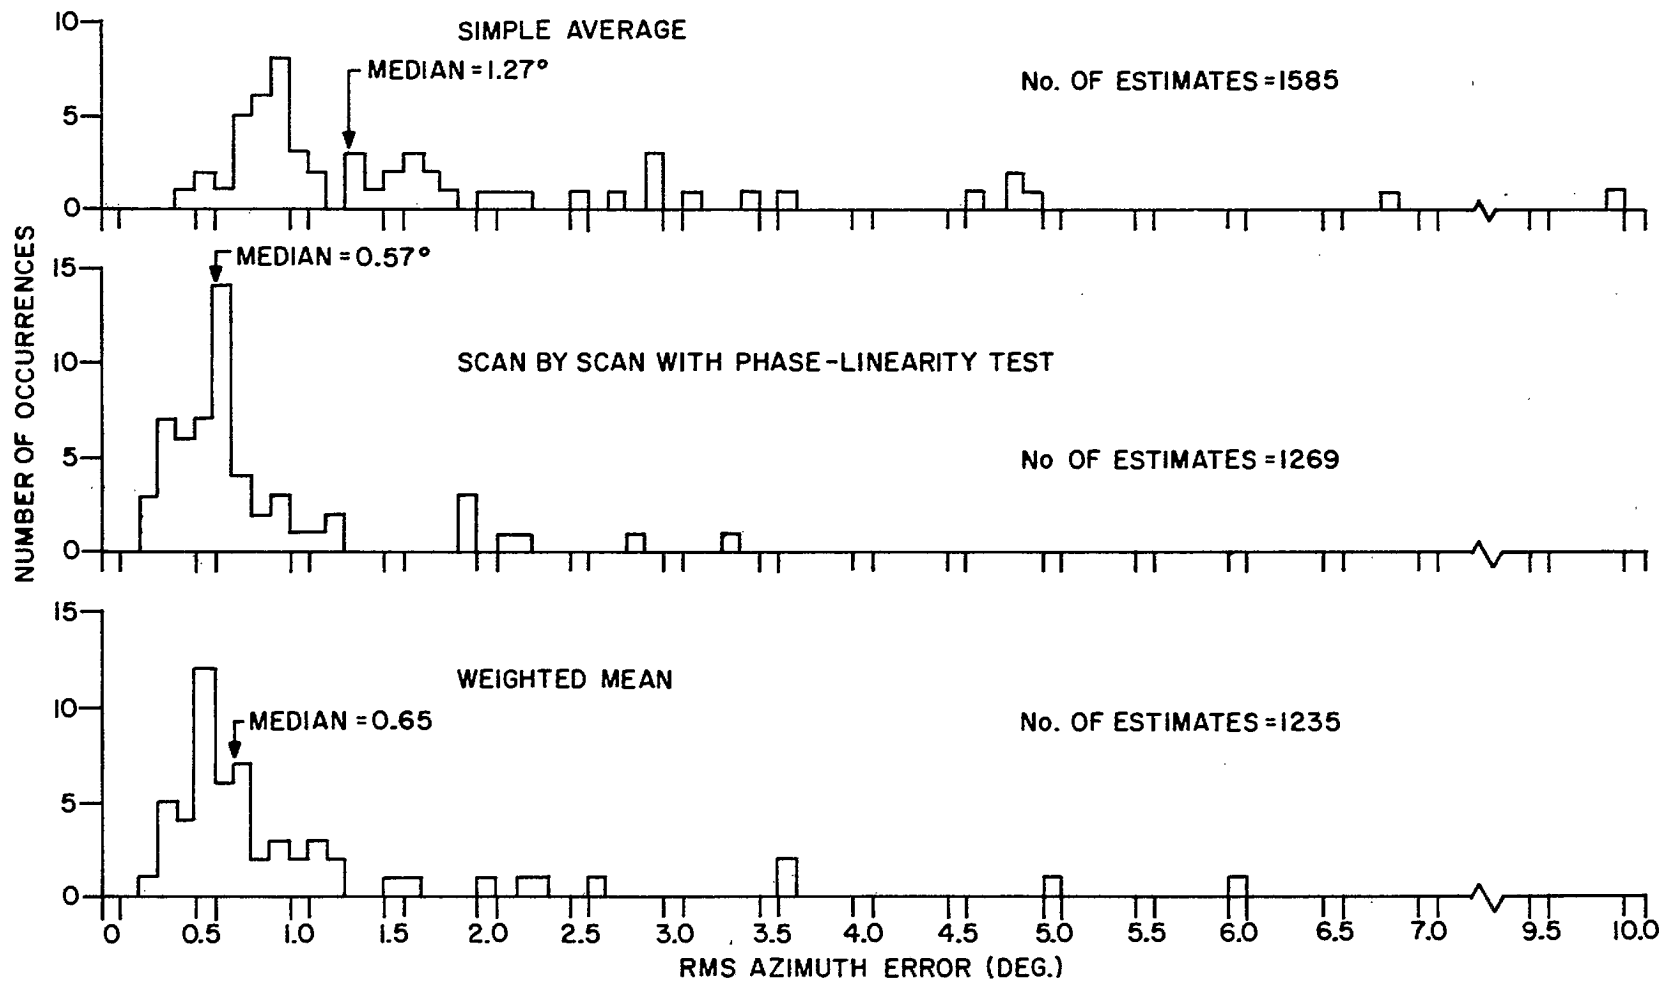No. of Scans Available = 43

Figure 8. Histograms of RMS Azimuth Error Measured on 40 Minute Intervals over Entire Experiment

methods, when compared against that for the simple average, give some indication of the loss of data resulting from the testing. They show that there was a loss of 20% of the scan-by-scan estimates and 22% of the weighted-mean estimates. This was an average loss and was higher for poor conditions, and lower for good conditions. The loss can be reduced by a relaxation of the acceptance criterion, but this will cause an increase in the RMS error; there is a trade-off between the two.

Figure 9 illustrates the loss of data for the scan-by-scan method when conditions were poor. Propagation was by spread F with the deterioration of conditions towards the end of the time period shown. The estimates from the simple average, indicated by small circles, deviated considerably from the true azimuth. The estimates from the scan-by-scan method are separated into two types; those that pass the test (20 out of 63 scans have low enough RMS phase deviations) are identified by an x and those that do not pass the test are identified by a dot. We see that, for the time before the ionogram, when estimates are relatively accurate, only one of the scan-by-scan estimates was rejected, while after the ionogram, when the estimates became much worse, 12 out of 16 estimates were rejected. The rejected estimates were usually very inaccurate. It may be argued that a poor estimate is better than none at all; that will depend on the situation. In any case, the poor estimates could be provided to the user, but identified as such, so that he could decide how to treat them.

5. SUMMARY AND DISCUSSION

The measurements reported here were made on a 1937 km path, a large section of which is under the auroral zone. In the mid-May time-period of the experiment, both F1 and F2 layers were present during the day, along with the E layer, and during the relatively short nights, spread-F and sporadic-E conditions occurred fairly often. Propagation conditions were generally more disturbed than those typically observed on lower latitude paths.

Separation of modes, by range discrimination with swept-frequency transmissions, provided information on the azimuth errors to be expected from the various modes. E-mode propagation provided the best accuracy, with a median RMS error of less than 0.3 degrees. However, when the propagation was identified as sporadic E, the errors were somewhat larger, with a median RMS error about twice that for normal E. When propagation was by two hops there was considerable degradation in accuracy, with the sporadic cases again giving larger errors.

Low-angle F1 propagation during the day showed fairly good accuracy, with a median RMS error of less than 0.5 degrees. The high-angle errors were only slightly greater for the F1 layer, but the F2 layer showed rather large errors for both the low-and high-angle cases. This mode was usually spread in range by more than 0.1 milliseconds. For the few cases where it was not, the accuracy was just about as good as for the F1 layer.

The single F layer which occurred at night provided relatively poor accuracy, even when spread-F conditions were not included. When this layer was heavily spread in range, so that the ionogram showed little high-and low-angle
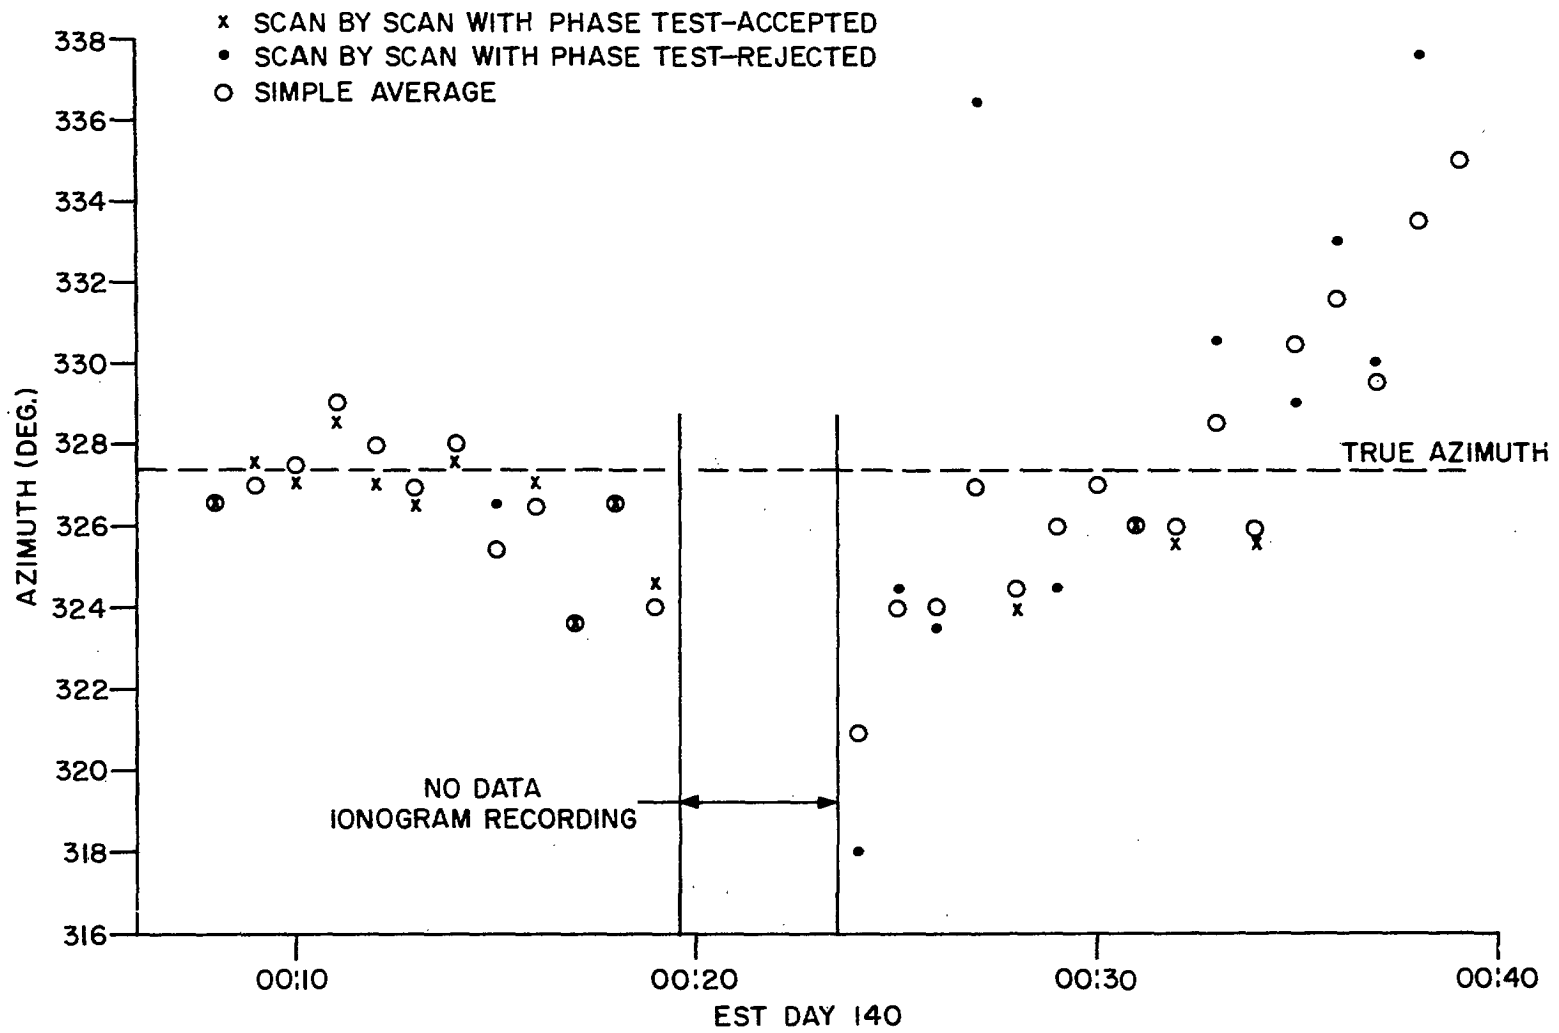
Figure 9. Rejection of Poor Azimuths by Phase Test — Spread-F Conditions

structure, the RMS azimuth errors were extremely large, and approached ten degrees for the worst periods.

One method of estimating azimuth on non-cooperative signals, for which range discrimination is not possible, has been referred to here as the scan-by-scan method. This method computes a mean value of azimuth based on a large number of samples (scans across the arrays). Samples are used in the averaging only if the RMS phase deviations across each array are below specified values. This is equivalent to a test for planarity of the phase front. Tests of this scan-by-scan method on the data from this experiment, showed that it gave an accuracy better than that of the worst mode in a multiple mode situation, but not always as good as that of the best mode. A simple average of scans without the phase-planarity test, on the other hand, gave an accuracy which was generally worse than that of any single mode. This is a clear indication of the value of the phase-planarity test in the scan-by-scan method.

Another method for azimuth estimation, proposed in an earlier report[2] and referred to here as the weighted-mean method, provided an accuracy not quite as good as that of the scan-by-scan method. The weighted-mean method uses Doppler processing in an attempt to separate modes, and is a more complex method than the scan-by-scan one. Thus it appears that the scan-by-scan method is the best one to use on a path of this type. However, the parameters used in the weighted-mean method probably were not optimum for this path, and some improvement may be possible.

The weighted-mean method is likely the better method for shorter, lower-latitude paths. On such paths, the elevation-angle differences between modes are greater, and the propagation conditions should be more stationary. The former condition will tend to give greater weight to the more accurate, lower-elevation-angle modes, and the latter will reduce the smearing in Doppler which can be caused by variations during the Doppler processing time.

6. ACKNOWLEDGEMENTS

7. REFERENCES

1. Venier, G.O., *An Investigation of H.F. Direction-Finding Accuracy on the Frobisher Bay - Ottawa Auroral Zone Path*. CRC Technical Note to be published, Communications Research Centre, Department of Communications, Ottawa.

2. Venier, G.O., *An Experiment to Investigate the Value of Doppler Shift Measurement as a Means of Improving Wide Aperture HFDF Azimuth Estimates* (U), CRC Technical Note No. 688, Communications Research Centre, Department of Communications, Ottawa, November 1977. RESTRICTED

3. Rice, D.W. and E.L. Winacott, *A Sampling Array for HF Direction-Finding Research*, CRC Report No. 1310, Communications Research Centre, Department of Communications, Ottawa, November 1977.

4. Hayden, E.C., *Propagation Studies Using Direction-Finding Techniques*, Journal of Research of the National Bureau of Standards, D, Radio Propagation, Volume 65D, No. 3, pp. 197-212, May-June, 1961.

5. Rook, B.J., *Study of the Behaviour and Stability of Phase Fronts on Short Time Scales*, CRC Report No. 1312, Communications Research Centre, Department of Communications, Ottawa, February 1978.

6. Rice, D.W., G.O. Venier, and G. Atkinson, *The Effect of Signal Modulation in the Application of a Wave-Front Linearity Test in HF Direction-Finding*, Proceedings of the 1978 NRL-ONR Symposium of the Effect of the Ionosphere on Space and Terrestrial Systems, Arlington, VA, January 24-26, 1978, paper 3-12.

A P P E N D I X   A

Formulas for Computation of Weights in the Weighted-Mean Method

The weight given to the azimuth computed for ith Doppler cell is:

$$W_i = Ge_i \cdot Gd_i \cdot Ga_i,$$

where

Ge_i, Gd_i and Ga_i are "quality" factors depending on elevation angle, RMS phase deviation across the arrays and received power, respectively in the ith cell. They are described below.

$$Ge_i = \exp[(2.0-\theta e_i)/5.0] \qquad \text{for } \theta e_i > 2.0$$
$$= 1 \qquad \qquad \theta e_i \leq 2.0,$$

where

θe_i is the elevation angle in degrees in the ith cell.

$$Gd_i = \exp\{[20.0-(d1_i+2.0d2_i)]/7.0\} \qquad \text{for } d1_i+2.0d2_i > 20.0$$
$$= 1 \qquad \qquad \text{for } d1_i+2.0d2_i \leq 20.0,$$

where

$d1_i$ and $d2_i$ are the RMS phase deviations in degrees across arrays 1 and 2 respectively in the ith Doppler cell.

$$Ga_i = V_i/Vs,$$

where

V$_i$ is the received voltage magnitude in the ith cell, and V$_s$ is the sum of the voltage magnitudes over all Doppler cells used in the computation.

The overall weight for the weighted mean azimuth is given by:

$$W_o = W_s \cdot \exp(\sigma_d/0.2) \cdot \exp(\sigma_b/0.50)$$

where

W$_s$ is the sum of weights for all the Doppler cells used in the computation,

σ_d is the standard deviation of Doppler frequency in Hz across the cells, and

σ_b is the standard deviation of bearing in degrees across the cells.

For the results presented here an overall weight threshold of 10^{-5} was used. Estimates whose weights were less than this were discarded.