PERGAMON

Journal of Atmospheric and Solar-Terrestrial Physics 63 (2001) 155–169

Journal of ATMOSPHERIC AND SOLAR-TERRESTRIAL PHYSICS

www.elsevier.nl/locate/jastp

Real-time determination of meteor-related parameters utilizing modern digital technology

W.K. Hocking[a,*], B. Fuller[b], B. Vandepeer[b]

[a] *Mardoc Inc., London, Ont., Canada N6G 4N4*
[b] *Genesis Software Pty. Ltd., Adelaide, Australia*

Received 7 September 1999; accepted 30 January 2000

Abstract

Modern fast digitization techniques and computer methods have been combined with both new and old theoretical approaches to permit construction of a new class of meteor radar. This radar can simultaneously stream data into memory, detect occurrences of meteors, and determine the location of meteor trails (range and angle), as well as find their radial drift speeds and decay times. The meteor entrance speeds as the meteors enter the atmosphere can also be determined. These parameters may then be used to deduce meteor fluxes, as well as winds, temperatures and diffusion coefficients at altitudes of 80–100 km. This information can also be used to deduce source positions of meteor shower radiants. Some of these capabilities are very new, especially the ability to measure atmospheric temperatures at $\sim$ 90 km altitude.

Keywords: Meteors; Atmosphere; Astronomy; Interferometry

1. Introduction

Meteor studies have a long history, and the field was particularly active prior to the mid-1970s. These studies were made to determine both information about the meteors themselves, and also to learn about the atmosphere in which they burn up. Many of the earlier astronomical results have been described by McKinley (1961), and some of the atmospheric results have been described in a special MAP handbook (Roper, 1987). An excellent, and more recent review, is given by Ceplecha et al. (1998).

During the 1980s and early 1990s, the incidence of new meteor research, especially for atmospheric applications, diminished. This was in part due to the retirement of some active researchers, and also because, to some extent, new possible advances were limited by inadequate technology. Many of the previous studies used photographic film, or somewhat primitive computer detection algorithms, and some of the research was manpower intensive. Computer algorithms were restricted in their capabilities and often "non-meteor" echoes could not be adequately discriminated from true meteor echoes. There was also, in the atmospheric community, a perception that middle atmosphere wind measurements could best be made with other techniques such as spaced antenna studies (e.g. Briggs, 1980, 1984; Roettger, 1981; Hocking et al., 1989) and VHF Doppler radars (Woodman, 1985; Zrnic, 1979; Hocking, 1997b). One notable exception was the work of Avery et al. (1983, 1989), who used narrow beam VHF radars to undertake meteors studies, but because the radars were often optimized for MST VHF studies, the meteor detection rates were often low. Tsuda et al. (1987, 1995) also carried out meteor studies in this time frame.

The availability of fast digitization devices, and high-speed personal computers with large memory buffers, changed this situation. It was now possible to use multitasking operating systems like UNIX on personal computers, and simultaneously stream data to memory and analyze it "on the fly". Furthermore, very detailed meteor selection and analysis algorithms were now possible, permitting very high rejection probabilities for lightning, E-region echoes, impulsive RF interference, auroral echoes, and so forth.

* Corresponding author. Tel.: +1-519-657-7822; fax: +1-519-657-7822.

E-mail address: mardoc-inc@home.com (W.K. Hocking).

1364-6826/01/$ - see front matter
PII: S1364-6826(00)00138-3

Procedures which once required manual intervention and took hours or days to perform on large data sets, could now be performed in real time. Newer, more sensitive detection algorithms were also possible.

Meteor radars have generally used either narrow beams (Avery et al., 1983) or, more commonly, interferometric techniques. The interferometric techniques generally utilize wider beams. In the limit, beams which are almost isotropic can be used, and these are capable of seeing meteors over the whole sky (the so-called "all-sky" systems). Interferometric techniques use phase information recorded at the receiving antennas to determine meteor locations, whereas narrow-beam systems rely on the assumption that most meteors detected occurred somewhere in the main beam of the radar. In regard to the interferometric method, one other limitation of earlier meteor studies was that these systems often used receiving antennas spaced only a half-wavelength apart, and the coupling between these antennas could be severe (up to ~ -10 dB). This therefore produced biases in the phases measured by the antennas, and so produced errors in location of the meteors. Improvements to this design were only forthcoming in the 1990s. Hocking and Thayaparan (1997) used four antennas spaced by typically 1.5–3 wavelengths, whilst Jones and Webster (1992), and subsequently Hocking (1997a) used a 5-antenna system with minimum antenna spacing of 2 wavelengths. The latter system has been analyzed for coupling effects by Jones et al. (1998). The availability of these new antenna arrangements improved the reliability of meteor systems.

Finally, two other phenomena caused a regeneration of interest in meteor studies. The first was the imminent arrival of the Leonids meteor storm in 1998/1999 (e.g. see Brown et al., 1997, 1998b; Brown, 1999), and the second was the possibility of meteor damage to satellites. These events produced a renewed interest from the astronomical and military communities.

It was within this environment of renewed interest, and significantly improved technology, that the SKiYMET radar was developed.

2. Radar objectives

The purposes behind the construction of the SKiYMET radar were multiple. It was intended that a system be developed which could acquire data at the highest possible rate, and simultaneously perform many different analyses. The ability to run continuously and unattended was also an important requirement. The system was designed to employ interferometry, using optimal receiver antenna spacings, and was also designed to be an all-sky system. This is in contrast to many earlier radars which transmitted their power into moderately narrow beams aligned at low elevation angles. The system was also designed to operate at very high pulse repetition frequencies (PRFs) — up to 2000 Hz and higher. This was a new development for routine meteor studies, since systems in the past often used PRFs of 600 Hz and less. These higher PRFs would allow the system to be used to determine parameters previously not amenable to standard meteor radar studies, such as meteor entrance speeds. Use of such a high PRF, however, suffered from one limitation — the aliasing range was small. For example, if the PRF is 2143 Hz, the aliasing range is 70 km. Thus, in principle, it would be impossible to resolve whether a meteor had a range of say 65, 135, 205 km, etc. This limitation, however, is easily resolved. Because it is known that the vast majority of meteors burn up in the altitude range between 70 and 110 km, and because the angular location (azimuth and elevation) are well known from interferometry, it is possible in most cases to use this information to determine the true range unambiguously.

Other special requirements of the new design included very detailed real-time meteor echo discrimination algorithms, as well as various forms of on-line "post-detection" software, including determination of middle atmosphere winds (80–100 km altitude), mapping of radiant sources during shower conditions, and determination of ambipolar diffusion coefficients and temperatures. Determination of diffusion coefficients (Tsutsumi et al., 1994) and temperatures (Hocking et al., 1997; Hocking, 1999a) are in fact very recent developments. It was also intended that the radar could measure meteor entrance speeds (e.g. Cervera et al., 1997), and substantial developments have been implemented here to permit real-time determinations. The system is also sufficiently flexible that new developments can be incorporated into the system where appropriate.

A persistent and underlying philosophy has been to ensure that, as much as possible, the data acquisition and analysis are fully integrated onto a single platform. Off-line analysis is reduced because much of the data reduction is performed in real time. Results of the on-line analysis can be obtained with relative ease. At the same time, the raw data associated with any events of possible significance are stored to file in real time, and so are accessible to the user should further detailed analysis be required.

In the following sections, the implementation of these objectives will be described. Following that, examples of typical outputs will be discussed and displayed.

3. Hardware

The system hardware comprises antennas, cables, a transmitter and a receiving/digitizing unit. The antenna layout is illustrated in Fig. 1, where the five receiving antennas are arranged in the form of an asymmetric cross, with arms of lengths of either 2 wavelengths or 2.5 wavelengths, as shown. The location of the transmitter antenna is not critical, although it should not be too close to any of the receivers. Each receiving antenna is connected to a separate receiver with cables of equal phase-length — typically 70 m or so. Systems generally operate at a fixed frequency which is

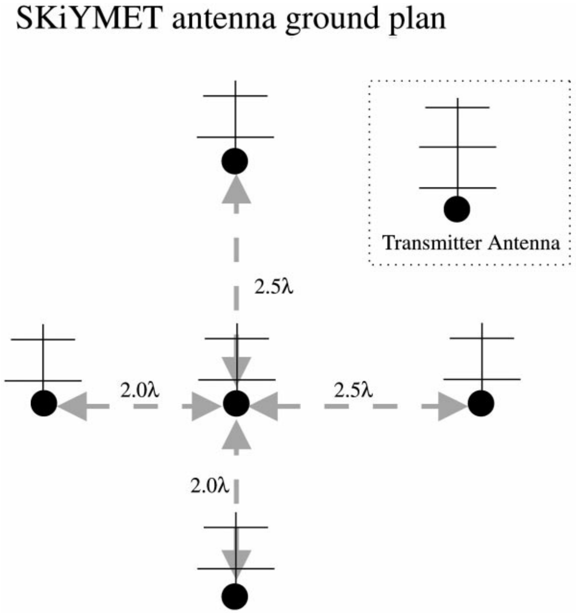

Fig. 1. Plan view of the antenna arrangement for the radar system. The location of the transmitter antenna is not critical and can be placed in any convenient location. The receiving antennas all need to be in a horizontal plane. The symbol λ represents the radar wavelength.

selected by the user, but is normally in the range from 20 to 50 MHz. The transmitter antenna is connected to the transmitter with low loss cable. It is also possible, in cases of very long receiver-antenna cables, to place pre-amplifiers at the antennas to enhance the signal for transmission back to the receivers.

The transmitter is a solid state unit comprising (typically) six 1 kW modules, although larger numbers of modules can be used. It is supplied by Tomco Electronics Pty. Ltd. of Australia. It is physically compact, being approximately a cube in shape, and having a length of about 60 cm per side. The peak power transmitted is generally 6 kW, and a selection of pulses are available. They may have a variety of shapes, including Gaussian, square, and square pulses with shaped tapering on the ends. Pulse lengths can also be selected, and can vary from a few hundred metres up to several kilometres. Generally an effective pulse length of about 2 km is the most common choice. The transmitter also includes protection against excessive standing-wave ratios and over-heating.

The next hardware unit to consider is called the Radar Data Acquisition System (RDAS). This unit includes five identical receivers, each connected to a separate antenna, which are then interfaced to a digitization system. It also contains the "Frequency Synthesizer Unit" (FSU), which provides various reference frequencies for the transmitter and receivers. The whole system is in turn driven by a personal computer running under a FreeBSD UNIX operating system. Many parameters can be selected under computer control, including filter bandwidths: up to four different filters may be chosen. Typical base-band filter half-widths are 25, 50, 125 and 250 kHz. For a 2 km pulse, the 25 kHz filter approximately optimizes the signal-to-noise ratio. Other selectable parameters include pulse length, pulse shape, pulse repetition frequency, receiver gain (up to 122 dB), sampling interval, and number of coherent integrations. The system also contains a large self-diagnostic capability, enabling thorough checks on the system functionality.

The use of a UNIX operating system is an important advantage. It enables fast and efficient multi-tasking, and also makes it easy to link the system to modems and the Internet. The multi-tasking capability makes it possible to stream the data into memory on a continuous basis, while at the same time analyzing the incoming data for the occurrence of meteor echoes. This minimizes the system "blind-time", so that the system is able to detect several thousand useful meteor echoes per day. The data are then written to hard disk. It is possible to store all data to disk, for detailed studies at a later time, or to record only the raw data associated with individual detections. It is most common to use the latter option. Typically, the data stored are those for 1 s before the meteor occurrence, and 3 s afterwards, at the range of detection. These data are then further analyzed to determine meteor range, angular location, height, strength of backscatter, lifetime, radial drift velocity (and its error), and meteor entrance speeds (where possible). This processing once again takes advantage of the multi-tasking capabilities of the UNIX environment, so that the processing can proceed independently while the radar continues to acquire new data. A wide variety of determinations are then possible, such as upper atmosphere wind speeds, temperatures in the meteor region, ambipolar diffusion coefficients, pressure in the meteor region, and locations of shower radiants. A flow chart describing the interactions between the various units, the data acquisition, and the software algorithms, is shown in Fig. 2.

In the following sections, we will describe the various software algorithms in some detail. We will begin with the initial detection software, and then progress to the higher level analyses.

4. Detection algorithms

One of the most important tasks performed by the SKiYMET software is meteor detection and discrimination. In the following sections, these algorithms will be outlined.

Meteor detection is performed with two successive processes. The first algorithm in the meteor identification sequence is termed the "detector", and this performs a "first-pass" examination of the in-phase and quadrature time series, identifies potential meteors for further analysis, and stores the data associated with them to data files, which we shall call "Preliminary EVent Files" (PEVs). These files typically include 1 s of raw data prior to the meteor

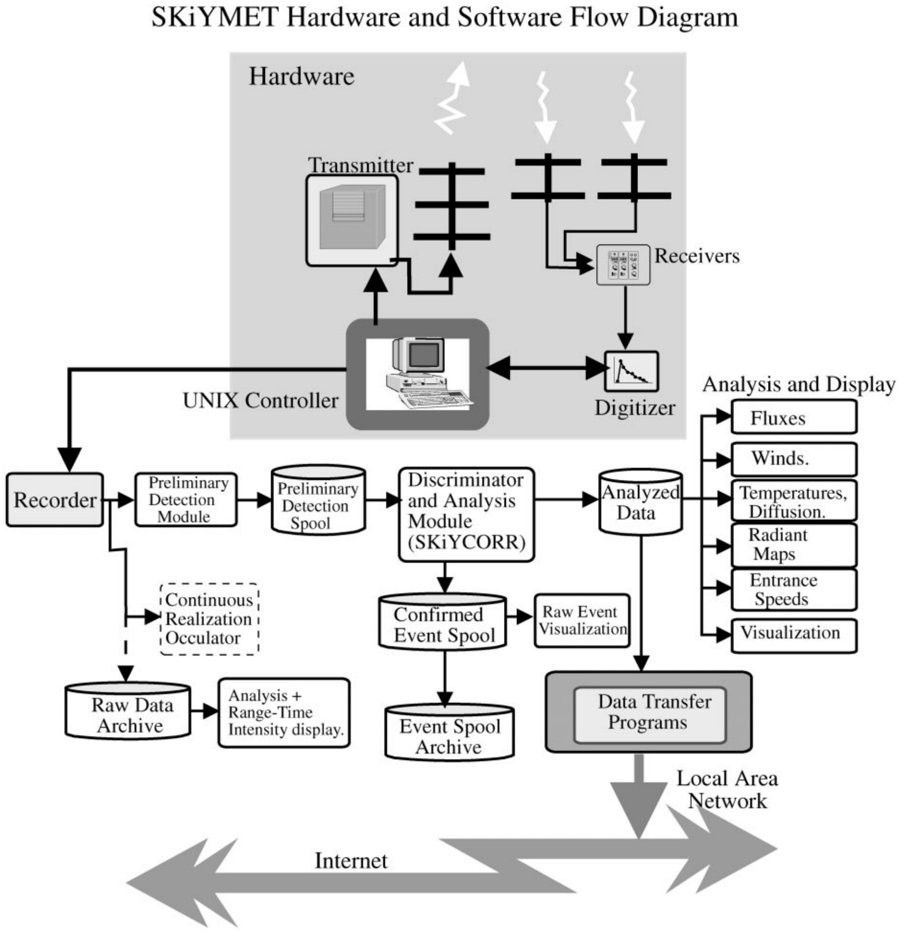

Fig. 2. Flow chart showing the interaction between and within the various hardware and software components of the system.

peak, and 3 s after it, at the range at which the meteor was detected. As a rule, this results in a reduction of storage space by a factor of about 100 (or more) relative to that which would have resulted if all data were stored.

The second-stage analysis (termed "skiycorr") re-examines these same meteors, but in much greater detail, and confirms them as meteors where possible. This strategy has been adopted so that the first-stage software (detector) has simpler tests to perform and so requires only minimal tasks to undertake during the initial detection. In cases of extremely high data rates, the detector is therefore not encumbered by more stringent tests. The second-stage (confirmation) program then analyzes these "PEV" files, but the time constraints for it are not so severe. Analysis is performed in the background, with the first-stage detector always having the highest priority. The confirmation program analyzes the PEVs whenever there is spare time on the machine. As a rule, with modern computers this second-stage analysis actually completes within seconds of the data acquisition, but in principle it is possible that it could lag the initial detection by any length of time. Thus, if there are bursts of meteor activity, the first-stage detector can function most efficiently for detection, and the second-stage confirmation can proceed at a later time when the meteor activity has diminished.

4.1. First-stage detection

One of the most difficult processes in the meteor analysis sequence is simply the task of finding meteor echoes and discriminating them from other short-duration signals. The features which distinguish meteors (or at least underdense meteors; e.g. see Jones, 1975) include their relatively short duration, their rapid onset, and their quasi-exponential decay time. Once detected, it is relatively simple to perform cross- and auto-correlation functions on them, but the intial detection can be quite difficult, especially if the meteor signal is weak. Therefore, various signal-to-noise improvements are necessary for this initial stage of detection. The "detector" performs the following tests.

Step 1. First, the in-phase and quadrature components are averaged in time bins, thereby implementing a (temporary) coherent integration on each receiver. This is in addition to any coherent integration which has already been performed at the digitization stage. Typically, we average over a sequence of 2–4 points in length (where the actual value is user-specified). The number should be selected in such a way that the integration time is substantially less than the period of the fastest oscillations expected due to the radial drift of the meteor trails. For a PRF of 2144 Hz (which is very common), with an initial coherent integration of 4, an extra (temporary) coherent integration of 2–4 points is a very reasonable number.

Step 2. Next, the amplitude of the signal on each receiver is determined from the in-phase and quadrature averages determined above, and then the program incoherently averages the amplitudes across all five receivers. This optimizes our signal for initial detection purposes. We use the incoherent average across receivers, rather than the coherent one, because a coherent average may well sum to zero, or a small value, since the phases on all five receivers are different. Then, the magnitude of each such sum is compared to the mean value of the amplitude in the previous "n" points, where the value of n is user-determined and is typically of the order of 10 or 20. If the current value significantly exceeds the standard deviation of the previous n points, then the point is considered as a possible meteor. If this is not so, the program proceeds to the next point. The meaning of "significant" is user-defined — the user specifies a number which represents the multiple of standard deviations by which the peak must exceed the standard deviation of the previous points. If the point does satisfy this criterion, then a flag is set to indicate a possible meteor. The next point is then examined; it is necessary that this must also exceed the noise floor by the above-specified multiple. If this is not so, the "meteor-present" flag is turned off. If the next point also satisfies this test, the program proceeds to further tests. Individual large spikes usually relate to lightning, or impulsive interference, and if these occur, they can be removed and replaced by the average of the value on either side. This latter process can be selected at the user's discretion.

Step 3. It is now necessary to proceed to further tests, in order to determine if the signal is truly a meteor. It is therefore determined whether the large value really does rise abruptly out of the noise. If the rise is slow, the data are rejected as aircraft, or E-region echoes, or some other type of contaminant. It is also necessary to ensure that the echo is reasonably short lived (typically less than 2 or 3 s). For example, aircraft generally take a few seconds (typically 10 or more) to pass through the beam so the signal is tested to see if it has returned to "normal" after about 3 s. If this is so, the signal is more likely to be a real meteor, so the test progresses to checking that the mean amplitude level has returned to a value comparable to the noise level within typically 3 s of the peak. This procedure also helps eliminate other slowly fading events such as ionospheric echoes. However, this test alone is insufficient to ensure that the signal is indeed a real meteor. It is also necessary to check that the meteor suddenly "turns on" and that the mean level for the cluster of three successive points at a location 0.3 s prior to peak does not have a mean value which exceeds the previously assigned trigger level. This effectively constraints the rise time to 0.3 s or less.

Experience has also shown that lightning can be a nuisance, both in the form of lightning-induce RF interference and radar scatter from lightning plasma trails. If these appear in the time series as individual amplitude spikes, they can be removed using the spike-rejection algorithms previously discussed. However, the duration of the lightning signals often exceeds the interpulse period, and their effect can be found on several successive pulses. When this occurs, the signals tend to show substantial amplitude variation throughout their lifetime, rather than the moderately smooth decay associated with meteors. This is true for both RF interference and scatter from lighting plasma-tubes. Other forms of signal also exist which alternate up and down in amplitude on time scales of a few tenths of a second, and these also need to be eliminated. All of these types of echoes can be eliminated with the following test.

Step 4. The data are searched for occurrences where the amplitude fluctuates up and down as a function of time in the 900 ms following the peak, using the following algorithm. First, the place where the amplitude has fallen to 0.3 times the peak value at the initial meteor onset is located. Next, the meteor echo amplitude beyond this point is examined to determine whether the amplitude rises to above 0.7 times the initial amplitude. If this occurs, it is assumed that the echo has non-meteor characteristics. As a consequence, the signal is rejected. This algorithm has proved very effective at removing lightning as well as other non-meteor echoes such as sporadic E-region reflections (which tend to show substantial amplitude variability). It does have the disadvantage that occasionally deeply fading oscillatory over-dense meteor echoes may be rejected, but this likelihood can be accommodated by adjusting the depth of the fading permitted (i.e. the values of 0.3 and 0.7 as previously discussed). This procedure will also remove signals associated with overdense meteor trails, but since these are only a small percentage of the total number of meteors, the loss is not too significant. The rejection of a great many false echoes more than compensates for the few overdense meteors which may be rejected by this algorithm.

Step 5. Despite the care taken in these tests, experience has shown that non-meteor events can occasionally slip through, and therefore there is one more important test which can be applied. This test also rejects very weak meteors which cannot be used because their signal-to-noise ratio is too small to allow their locations to be determined. The actual test involves forming the cross-correlation function between all pairs of antennas and examining the variance of the rate-of-change of phase of these correlation functions. When the standard deviation of the rate-of-change of phase at zero lag is found across all receiver-pair combinations, a very useful decision-making parameter is obtained. If the value is small, it is relatively certain that we have a good meteor, while if the value is large, it is likely that the "signal" is in fact very noisy and should be discarded. The division between a "large" and "small" standard deviation must be determined by the user, but visual examination of many meteors has permitted us to determine that a value of about 6 m/s (where the rate of change of phase has been converted to a radial drift velocity) is reasonable. However, in the detector algorithm, the parameter can be user-defined. Indeed in some cases, values of up to 12 m/s are accepted, so that even some very weak meteors will not be rejected. More rigorous tests are applied in the following, second-stage detector, program, which will be discussed shortly. It should be noted that the cross-correlation function, rather than the autocorrelation function, is used, since the autocorrelation function has a troublesome noise-spike at zero lag which can complicate calculations. The cross-correlation function is less contaminated by such spikes.

The above tests have been found to be very effective in removing non-meteor echoes. As an example, during a recent meteorological storm, 4000 echoes due to lightning were recorded and were initially accepted as meteors; yet after running through these final two identification processes, it was found that there were only 10 left. About six of these were real meteors, and the other four were lightning spikes superimposed on simultaneous short-lived aircraft occurrences. This is a relatively rare event, and four such "false detections" in the course of a day's worth of meteors represents a very small error rate. However, it is emphasized that these detection criteria can be controlled, to some extent, by the user. Some parameters are adjustable, and some tests can be turned on or off. Often the criteria are relaxed so that a higher percentage of "doubtful" signals are accepted. This is done so that these echoes can be visualized with the various graphics packages available; it is often a good idea to be able to investigate "doubtful meteors" from time to time, in order to keep a visual check on the rejection criteria. It must be remembered that all these detections will still be subjected to a much more rigorous series of tests in the second-stage confirmation program, so false detections will be removed there. Indeed, the criteria used in that program even allowed removal of the four curious lightning-cases "left over" in the discussion above.

Step 6. Finally, if an echo is detected and passes all of these tests, the data associated with this echo are saved to a PEV file. In particular, a 4-s data stream containing all of the in-phase and quadrature components for each receiver is saved, where the data stream starts 1 s prior to the peak of the meteor, and finishes 3 s after the peak, at the range of detection. Data are not stored at the resolution used for the tests discussed earlier, but rather at the (much higher) temporal resolution of the original data acquisition specifications.

4.2. Second-stage analysis: meteor confirmation

4.2.1. Data storage and general strategy

The second stage of identification and confirmation can now begin. The program name used for these studies is "skiycorr". It is possible to run this analysis in parallel with the first-stage detector, with this new program identifying new 4-s records and analyzing them as soon as the detector has produced them. However, it is also possible to save all 4-s files and analyze them at a later time. The purpose of this program is two-fold. First, it is used to further subject the original meteor-detections to more rigorous testing, to be sure that the "possible meteors" really are valid meteors. Files which pass these tests are written to new files called "Confirmed EVent" files (CEVs). Then, secondly, it is used to determine the location of the meteors in the sky (zenith, azimuth and range), and then perform other calculations such as determinations of decay times, radial velocities, and entrance speeds. The results of these analyses (including meteor trail drift velocities) are then written to an ASCII text file for further study. Once again, these analysis procedures will be described in a sequential manner.

Before applying these tests, it is important to register the intrinsic phase delays between the antenna/receiver combinations. These delays must be known in order for the algorithm to determine meteor positions. They arise from a variety of factors, but are mainly due to receiver delays, since all cables are generally cut to equal phase lengths, and all receiving antennas are identical. As a rule these phase delays are measured at the time of installation, and also at regular subsequent intervals. They are stored to file, from which the skiycorr program can read them as required. We have found that the receivers are very stable with respect to time, and drift by only a few degrees over the course of a year. Nevertheless, regular phase checks are highly advised.

Skiycorr repeats some of the steps in the detector, but with more aggressive rejection criteria. Some new tests have also been added and these will be discussed in due course. It also has the option to include its own "spike-rejection" criteria, which can be used either to supplement that used in the detector, or to replace it entirely.

The confirmation program, skiycorr, then begins to search for meteors within each 4-s record produced by the detector. The 4-s PEV files also contain some header information which tells the skiycorr program where (within the 4-s record) to look for the meteor. Following this, skiycorr performs a double-check of the position of the meteor peak. This check begins with a search of the incoherently averaged data points, seeking out pairs of points which stand out above the noise, just as in the detector, but with more stringent acceptance criteria.

Once a possible meteor has been accurately located, a series of more thorough tests then ensues. Some of these repeat algorithms used in the first stage of detection, but with more stringent parameters. However, the majority of these tests are new. They are listed below. It should be noted that some of the tests may not seem especially "stringent". This is true, and it is a deliberate strategy, with the intent being to apply many tests of a modest nature throughout the program, rather than one or two extreme tests. When combined, our tests then make a powerful discriminator of meteor echoes.

4.2.2. Second-stage tests

The tests performed in "skyicorr" are now described.

Step 1. Three separate data intervals of width 0.25 s are isolated at the following locations. The first interval is found at the very start of the record (usually 1.0 s prior to the meteor peak, and therefore covering a range from 1.0 to 0.75 s prior to the peak). The second interval is chosen just prior (0.35–0.1 s) to the peak. The third is found starting 0.7 s after the peak. The root-mean-square (RMS) values of the incoherently averaged data within these three bins are then found and compared. The largest RMS (root-mean-square) value of these is selected and is compared to the value of the peak of the meteor (where the peak value has been found after applying both the coherent and incoherent integrations described earlier). It is required that this peak value exceed the largest RMS value within the other three bins by at least a factor of two. (It is worth noting that when the incoherent amplitude averages are formed, the mean value is offset from zero by an amount which can substantially exceed the standard deviation of the amplitudes. This offset is removed from all the data before any tests are performed, and it is especially important to be aware of this for the above tests.) A factor of two may not seem large, but is sufficient to eliminate cases where the "peak" detected by the detector turns out to be little more than a noise burst.

Step 2. The next step involves cross-correlating the data on the different channels prior to the meteor peak. In this case, rather than requiring high levels of correlation, low levels are sought, so it can be determined that the meteor truly has occurred suddenly. This therefore takes advantage of one of the most unique features of meteors — the fact that meteor signals rise very rapidly out of the noise. This test successfully removes cases of moderate E-region reflections which are persistently present but which suddenly increase in strength, for example, because the data prior to the sudden increase are still correlated. It is also a very good way to remove aircraft contributions which somehow passed through the earlier tests. Specifically, the variance of the rate of change of phase of the cross-correlation functions between all pairs of antennas is examined near zero lag. The test is very similar to the cross-correlation procedure described earlier for the detector, but in this case small values are rejected because this indicates highly correlated data just prior to the meteor. In reality, the rates of change of phase are converted to radial drift velocities (the units of which are m/s) and those cases in which the standard deviation for the mean is less than 2.0 m/s are rejected.

Step 3. Following step 2, the same type of tests are applied yet again, but on this occasion the interval of time used includes that in which the meteor occurred. In this case, data which are well correlated are sought. A data bin of typically 1.0–2.0 s is used, depending on the meteor duration, and cross-correlation functions are formed between all pairs of receivers. It has been found that data are acceptable if the standard deviation of the mean radial drift velocity is less than 5.5 m/s. This has been determined largely by the visual inspection of many hundreds of meteors, and this test works well for radar frequencies above 30 MHz. Larger values are acceptable for lower frequencies.

If the data passes all of the above tests, an acceptable meteor has been identified; if not, the data are rejected and ignored and the program passes on to examine the next 4-s data set.

5. Meteor parameters

The next step is to determine parameters which describe the meteor and its plasma trail. These are found in the skycorr program, after the meteor has been confirmed. Clearly, the peak amplitude and the lifetime are easily found — the latter is determined by finding the time for the cross-correlation function to fall to 0.5 times its value at zero lag.

One of the most important pieces of information is, of course, the location of the meteor in the sky. This is found by comparing phase differences of the meteor signal at zero time lag in the cross-correlation functions, after compensation for any intrinsic phase differences between the receivers. These phase differences may then be inverted to determine a direction of arrival of the reflected radio wave by standard interferometric techniques (e.g. Roettger and Ierkic, 1985; Larsen and Roettger, 1991; Larsen et al., 1992). The novel interferometric antenna arrangement shown in Fig. 1 removes (in principle) any angular ambiguities. However, if the meteor signal-to-noise ratio is poor (leading to errors in the phase determinations), it can happen on occasion that the program cannot decide between two possible meteor positions in the sky. If this occurs, both options are saved, and the meteor is assigned an "ambiguity level" which specifies the number of acceptable angular positions. All this information is saved to the final output file so that the subsequent user can decide whether to employ this information or not. In general, only a small percentage of meteors show such angular ambiguities.

At this juncture, there is another calculation which needs to be made, and this is the determination of the height of the meteor. If the range is known, this is relatively trivial, since the angle from zenith may be used to determine the height. The curvature of the Earth is allowed for in this calculation. However, there are times when this calculation is more complicated; this occurs if the pulse repetition frequency of the system exceeds 1500 Hz. Such high PRFs are often used to enhance the signal-to-noise ratio, and also to permit determination of quantities which require higher temporal resolution (such as calculation of meteor entrance velocities).

If a high PRF is used, it is possible that the meteor will be range-aliased. For example, a PRF of 2144 Hz (a commonly used value) produces an aliasing range of $\sim$70 km, so a meteor at 150 km range appears at a range of 10 km. Likewise meteors at 80 km range, or 220 km range also appear at 10 km. Skyicorr therefore determines all possible ranges from which the meteor might have come, and then uses this information to determine possible heights of the meteor. Ranges which do not produce a height in the region 70–110 km (where these limits can be user-specified) are rejected. In the majority of cases (except for the lowest elevation angles), this procedure produces only one possible range, thus allowing unambiguous location of the meteor. In cases where there are range ambiguities, the "ambiguity level" is increased to reflect this fact, and the possible range-angle combinations are printed to file. Thus, these are still available to subsequent users should they be needed.

The final parameter which is determined is the speed at which the meteor entered the atmosphere. Determination of this parameter is not always possible, but when it can be determined, it represents very important information. The details about this determination will be given later.

Finally, the data are written to two files. One is a binary file, and the other is a text file. This dual option offers some extra security against file deletion, and also permits users to use either as input for subsequent analyses. The following data are written to file. First, the date and time are given (year, month, day, hour, minute, second, and millisecond). Then a unique identifier is given which allows the user to identify the 4-s record which holds the raw data (CEV file) for this meteor; this can be useful if the user wishes to go back and examine the raw data associated with the meteor in more detail. Following this, the range and height of the meteor are listed. Then the mean radial drift velocity and its associated error for the mean are written. Following this are the angle from zenith, and the azimuth angle anti-clockwise from due East. The next parameter is the ambiguity level — if this is 1, the data are unambiguous, whilst if it is greater than 1 then there exists the possibility that the meteor has been wrongly located. In the case of ambiguous echoes, other possible locations will be suggested in the lines immediately preceding or following the current line. Following this, a value representing the phase errors between antenna pairs is stored. This is found by determining the phase differences which should exist between all antenna pairs for the specified meteor location, and then finding the difference between the actual value and this "theoretical" value. The largest such difference is then recorded. This can sometimes be used to resolve between several possible ambiguous meteor locations. The next parameter is a 2-digit number which specifies which antenna pair has this maximum phase error. Other important parameters which are written to file include the meteor amplitude (digital units), the meteor decay time, and the meteor entrance speed into the atmosphere (if available).

6. Scientific measurements

We now turn to consideration of the scientific parameters deduced in the analyses. Some of these are determined in skiycorr, but the system also includes other algorithms which accept the output of skiycorr and determine higher level information, which relates to the essential science produced by the radar. The SKiYMET radar affords measurements of a variety of parameters, including both "traditional" ones, and some very new capabilities. Of the "traditional" ones, meteor fluxes, upper atmosphere winds and ambipolar diffusion coefficients are the most predominant, while newer parameters include absolute measurements of atmospheric temperature and pressure, as well as astronomical quantities like meteor entrance speeds and radiant mapping. In the following subsections, we will describe the various techniques used to determine and demonstrate their application. It is not possible to describe each technique in detail, but we will present a brief overview, and refer the reader to more detailed articles where appropriate.

6.1. Meteor fluxes

Perhaps the simplest parameter to measure with a radar of this type is meteor fluxes. In its simplest form, this consists simply of "counting" meteor occurrences. However, in order to be useful for astronomy purposes, it is important to make certain adjustments. Firstly, it is necessary to compensate for angular biases in detection rates. Such biases occur due to radar polar diagram effects, range effects, and atmospheric effects. The strengths of the meteor echoes should be converted to more useful parameters like particle mass. In order to do this, it is necessary to utilize the meteor amplitude, and use the known transmitted power, receiver gains and efficiencies, range effects and atmospheric effects. Atmospheric biases include angle of entry effects. For example, meteors tend to be detected more commonly at angles beyond 25° from zenith, and rather more rarely from overhead. The meteor echo strength varies proportionally to the meteor range to the power of -3. Examples of the types of calculations involved in these conversions can be found in Brown and Jones (1995) and Brown et al. (1998a, b), and references therein.

The SKiYMET software does not do these conversions directly, since the determinations are best done on a site-dependent basis, but it does provide all the information needed to do such conversions. Information which is supplied includes echo amplitude, range, zenith angle, and azimuth. Fig. 3 shows plots of typical meteor fluxes (before correction) and typical angular and height distributions. Notice that in this case there is a preponderance of meteors in two directions: this arises because the antennas which were used were planar Yagi antennas. If crossed antennas are used, fed with a 90° phase difference, this asymmetry disappears, and options exist to use such antennas with the SKiYMET system.

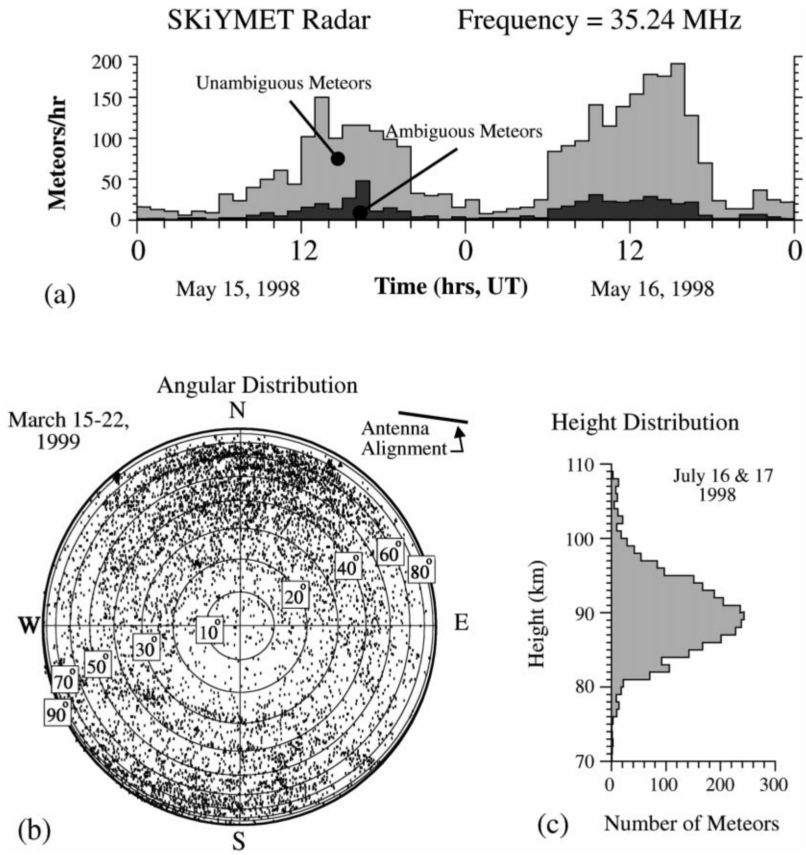

Fig. 3. Displays of typical flux data produced by the radar. (a) Typical meteor count rates as a function of time of day over a period of two days. A pronounced diurnal cycle is apparent, although the day-to-night variation is not always this severe. (b) Typical azimuthal and zenithal distribution of several thousand meteors, taken with a system in New Mexico. Note that the excess of meteors to the north and south are due to the antenna alignment. If crossed antennas are used, the distribution becomes more rotationally symmetric. (c) Typical height distribution of the meteors recorded with the radar.

6.2. Upper middle-atmosphere winds

Upper atmosphere winds have always been a main-stay of radio-meteor research. Measurements of this parameter are accomplished by measuring the radial velocity of every meteor detected, and then combining these measurements in an all-sky manner to determine upper level winds. In the case of SKiYMET, radial velocities are determined within skycorr by using both auto- and cross-correlation functions associated with meteor detections, and using the rate of change of phase near zero lag to determine the radial velocity. All possible cross correlations between all signals from all five receiving antennas are used, as well as the auto-correlation functions on each receiver, and then the average is taken. The standard deviation for the mean is also found. This latter parameter is used as an estimate for the error, and is also stored with the radial velocity for subsequent analysis. If, however, the standard deviation for the mean is excessively large (typically greater than 5.5 m/s — see previously) then the "meteor" is rejected entirely.

The on-line all-sky least-squares fitting routine currently assumes a uniform wind $\mathbf{u} = (u, v, w)$ and then minimizes the quantity

$$\sum_i [\{\mathbf{u} \cdot \mathbf{r}_i^u\} - v_{ri}]^2, \tag{1}$$

where i refers to the meteor number in a specified height and time window. Typically, such a window would cover a height region of 3–4 km, and a duration of about 1.5 h. Such windows are stepped at time steps of 1 h, and height steps of 3 km. The vector $\mathbf{r}_i^u$ is a unit vector pointing from the radar to the ith meteor trail. The value v_{ri} is the measured radial velocity, and $\mathbf{u} \cdot \mathbf{r}_i^u$ is a dot-product.

In the simplest case, it is assumed that $w = 0$; in more complex cases, it is possible to also assume that there are gradients as a function of position in the sky. The SKiYMET radar produces an on-line solution of Eq. (1) for the special

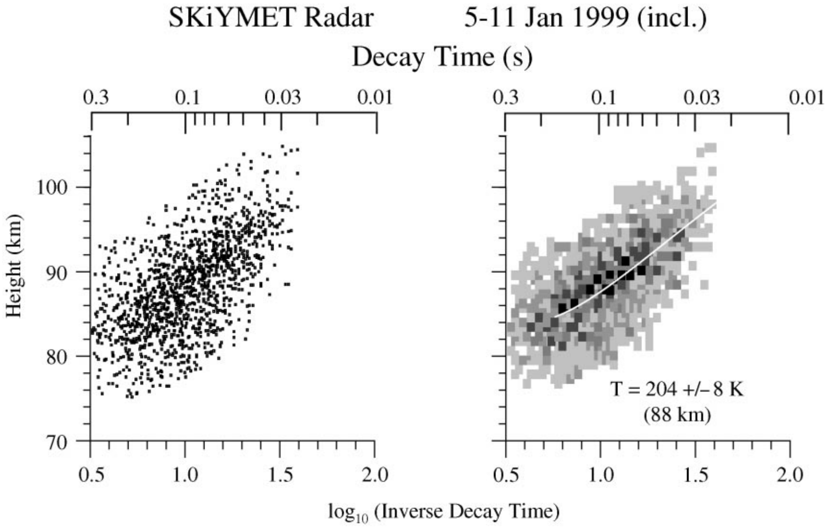

Fig. 4. Typical graphs of decay times vs. height produced by a SKiYMET radar. The graph on the left shows a scatter plot of individual meteors, while that on the right shows a density plot for the same data. The density shades are not given in absolute units since only the relative densities are important for the curve fitting. A best-fit fourth-order polynomial is shown as the white line in the right-hand side graph.

case that $w = 0$, but provides all the relevant information for the user to apply more sophisticated least-squares fitting algorithms should they so choose.

There is one more recommended process which should be applied in meteor winds determinations, which is also applied with the on-line SKiYMET determinations. This is a double-iteration of Eq. (1). In this process, Eq. (1) is first applied to all data in the specified time- and space-bin, to produce (u, v, w). Then, for each meteor, the radial velocity which should have been observed, assuming that the true mean wind was indeed (u, v, w), is determined. This is then compared to the actual measured radial velocity. If this difference exceeds some user-specified value (typically 30 or 40 m/s) then this meteor velocity is rejected as an outlier. This rejection value does not need to be specified too carefully — its purpose is simply to remove truly errant radial speeds. Then, Eq. (1) is repeated, but only using meteors which pass the above test. This procedure tends to "clean up" the data. It does not necessarily mean that rejected meteors are erroneous measurements, but rather that they are not typical of the "average" conditions in the sky at the time. They could, for example, come from a region and time where there were large, short-lived perturbations in the wind fields. Such events are often deserving of further study, but such investigations must be left to the scientists using the instruments. For first-order estimates of the mean wind, they are best ignored.

It is also useful to store the value of $\{[\sum_i [\mathbf{v}_{\mathrm{rm}(i)} - \mathbf{v}_{\mathrm{r}i}]^2]/N\}^{1/2}$, where N is the total number of points, $v_{\mathrm{rm}(i)}$ is the model radial velocity determined for the ith meteor assuming that the meteor's angular location has been correctly located and that the mean wind is correctly specified by (u, v, w), and $v_{\mathrm{r}i}$ is the measured radial velocity for this meteor. This parameter is called the "residual", and gives a measure of fluctuations of the wind speeds about the mean. It serves as a crude indicator of gravity wave and turbulence strengths.

Because the wind speed is a somewhat common parameter, and has been determined with many previous radars, we will not present graphical examples here. Examples can be found in Hocking and Thayaparan (1997), among others.

6.3. *Ambipolar diffusion coefficients, temperatures and pressures*

Another fundamental parameter that is produced by meteor signal analyses is the decay time of the amplitude. The time for the amplitude to fall to one-half of the peak amplitude is given by

$$\tau_{1/2} = (\lambda^2 \ln 2)/(16\pi^2 D_\mathrm{a}), \tag{2}$$

where D_a is called the ambipolar diffusion coefficient and λ is the radar wavelength. Studies of D_a have been undertaken by several authors (e.g. Tsutsumi et al., 1994; Nakamura et al., 1997). It has been shown (e.g. Jones and Jones, 1990; Hocking, 1999a) that the ambipolar diffusion coefficient is related to the atmospheric temperature (T) and pressure (P) at the height of the meteor trail by

$$D_\mathrm{a} = K_0 T^2/P, \tag{3}$$

where K_0 is a constant. A graph of the logarithm of the inverse-decay–time as a function of height is a monotonically increasing function with substantial scatter. An example is shown in Fig. 4. By fitting a polynomial to plots like

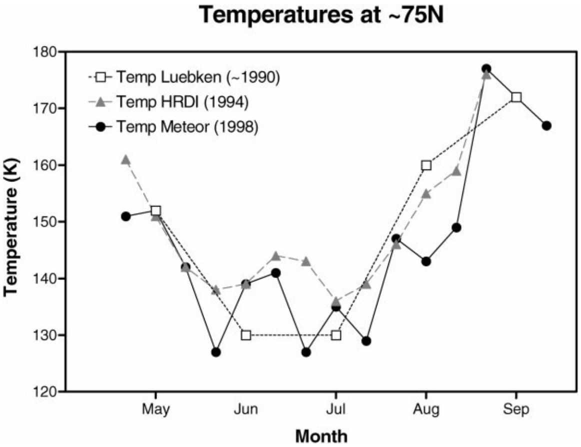

Fig. 5. Temperatures measured by a SKiYMET-like radar at Resolute Bay in northern Canada, and compared to other methods, during the summer time.

this as a function of height, Hocking et al. (1997) were able to produce reasonable estimates of T^2/P as a function of height and month, which compared moderately well with the COSPAR International Reference Atmosphere (CIRA). However, because of errors in the CIRA pressures, it was not possible to use these estimates to produce accurate absolute temperatures. However, by using a different approach, which involves a scale-height analysis of D_a, Hocking (1999a) was able to deduce absolute temperatures at the height of peak meteor count rates, and this process required no knowledge of the atmospheric pressure. This procedure appears to be robust, and Fig. 5 shows an example of temperatures measured at Resolute Bay, in northern Canada, compared to other reference measurements. After determination of absolute temperatures, and determination of T^2/P from Eq. (2), it is also possible to derive absolute pressures at this same height. Determination of pressures is a very novel capability which is rare in almost any ground-based instrument.

SKiYMET radars are capable of all these calculations, and have already produced a large number of measurements of these quantities at a variety of sites. Detailed studies using these results will be forthcoming in future publications. It should be especially noted that the meteor radar operates throughout both day and night, and this gives the system the ability to study tidal temperature oscillations. For example, by superposing decay times from common hours throughout a full month, a composite day can be produced, with temperatures determined for each hour of the day. Harmonic fitting can then be applied, to give diurnal and semidiurnal temperature tides. This opens new possibilities for studies of atmospheric tides which have previously not been available, since many of the instruments which have produced mesopause-level temperatures in the past have been optical in nature, and could not operate outside of cloud-free, moon-free, night-time conditions.

6.4. Radiant location

In astronomical studies of meteor activity, a common requirement is determination of the sources of meteors. Meteors can be either sporadic, or shower, types. In the former case, they tend to come from a broad range of locations in the sky, while in the latter case they come from a single radiant. Although it is not possible to determine the radiant of any individual meteor using the SKiYMET radar, it is nonetheless possible to determine locations of radiants when many meteors have a common origin. This can be done in the following way.

When a meteor enters the atmosphere, it is seen by the radar if it produces a plasma trail which is aligned perpendicular to a vector originating from the radar and pointing to the meteor trail. Thus, there is a “great circle” in the sky, centered around the radar, along which meteors from a particular radiant can be detected. Conversely, if a meteor trail is detected at a particular azimuth and zenith angle, then it is known that the meteor radiant must have been at some point on another great circle aligned perpendicularly to the vector from the radar to the trail. The exact location of the source on this great circle is unknown. However, the great circle can easily be plotted in celestial equatorial co-ordinates. If we now turn to another meteor, it also has its own great circle of “possible radiants”, and this can be plotted. It will generally differ from that of the first meteor. Such great circles can be plotted in celestial equatorial co-ordinates for every meteor detected. This requires compensation for the time of day of the detection, as well as compensation for latitude and longitude, but it is not a difficult procedure. If there were a significant single source of meteors in the sky (as during a meteor shower), then many of these great circles will cross at a common right-ascension and declination, and this point will indicate the source of the shower.

This procedure has been automated with the SKiYMET radar, and Fig. 6 shows two examples of this procedure. The first shows a typical situation when there is no dominant source in the sky. It can be seen that there is a broad region which is slightly stronger than its surrounds, but the region is diffuse. Such diffuse “sources” are common for the case where the sporadic background is the main origin of most meteors. The second shows the case when there is a strong meteor source in the sky, and it correctly locates the position of the shower radiant to within 3°. More detailed analysis allows even better localization. Other versions of similar algorithms have been presented by Jones and Morton (1982) and Jones and Brown (1993, 1994).

It should be emphasized that even in the case of sporadic meteors, the SKiYMET radar gives useful information about the distribution of meteors. Studies using SKiYMET-like radars are already underway to further understand both meteor sources and this diffuse background (e.g. see Brown et al., 1998a,b). Accurate location of radiant sources, and studies of the sporadic background, are therefore both areas

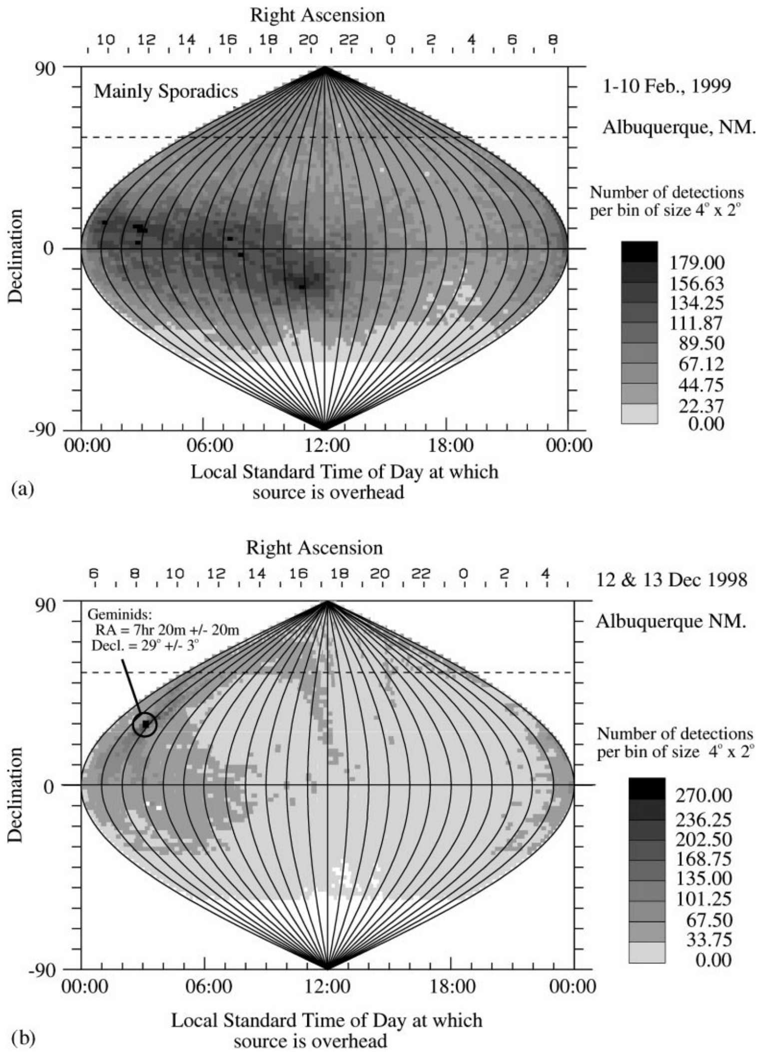

Fig. 6. (a) Results of "radiant mapping" determinations for data from early February, 1999. This is a period when there are no major sources in the sky, so that this is a typical case for a sporadic background of meteors. Note that the abscissa uses two scales — one is the right ascension, the other is the time of day at which radiant (if any) is overhead. To use this axis, apply the scales to the equator, then follow the curved lines of right ascension to (or from) the appropriate declination. (b) Results of "radiant mapping" determinations for data from December 12 to 13, 1998. This is a period when the Geminids stream was the major source in the sky. Notice that the strongest density is at the point indicated by the circle in the top left-hand corner. The correct value for the location of the Geminids, according to the Norton's Star Atlas (Norton, 1973), is RA = 7 h 28 min, Decl. = +32°. This agrees with the value determined by the analysis to within the error quoted on the figure. It should be emphasized that the errors shown arise because of our choice of grid size — higher accuracy is indeed possible (to about 1.5° or better) if smaller grids are used. We have chosen to use the coarser resolution because this graph emulates a real-time display on the system. Use of a smaller grid size produces software which is too slow for on-line implementation, but can certainly be applied off-line.

where the SKiYMET radar is making, and will continue to make, significant contributions.

6.5. Meteor entrance speeds

Meteors enter the atmosphere with speeds in the range of typically 10–100 km/s. There are therefore at least two speeds associated with a meteor trail. One speed is the drift of the trail as it is blown around by the atmospheric wind, but there is also a second. This latter speed is the speed at which the head of the meteor moves through the atmosphere, and it is very similar to the speed of entry of the meteor into the atmosphere. Whilst it is common to measure the drift speed due to the wind, it is also possible, although more difficult, to measure the speed of the head of the echo. This determination can be undertaken as follows.

If a meteor trail being formed in the atmosphere is illuminated by radio waves impinging from a direction perpendicular to the direction of alignment of the trail, then a Fresnel diffraction pattern is produced at the radar by the backscattered radio waves. This pattern sweeps past the radar (as the front of the trail proceeds to move forward) with the speed of the incident meteoroid which formed the trail. This pattern has a well-known shape, which is similar to that formed when light is diffracted around a sharp edge (e.g. Hecht and Zajac, 1974, p. 386). By measuring the complex amplitude as a function of time, it is possible to measure the speed of the pattern past the radar, and thence that of the meteor itself. It should be emphasized that this diffraction pattern first appears before the meteor signal peaks in amplitude, and persists after the peak. In the past, the amplitude oscillation occurring just after the peak has been used to measure the meteor speeds (e.g. McKinley, 1961), and more recently Cervera et al. (1997) have used the phase oscillations prior to the peak. The SKiYMET radar is unique in that it employs both these parts of the signal time series, and uses them to determine the entrance speed. The determination is fully automated. We will not describe the method in full detail here; it is discussed in much more detail by Hocking (1999b). However, we will demonstrate some results produced by the SKiYMET radar.

Fig. 7 shows a sequence of typical distributions of meteor speeds measured with a SKiYMET radar near Adelaide, Australia. The maximum speeds are around 40 km/s, but this is a limitation of the sampling strategy used, rather then the technique itself. In this case, we used a PRF of 2144 Hz, and a 4-point coherent integration. Higher speeds can be measured if less coherent integrations are used. The basic shape is broadly similar from one time interval to the next, but there are also some subtle differences. One difference of note is that there is a secondary peak at speeds of 7–8 km/s, and this peak shows a temporal variation. This is the speed which dust and debris would have in orbit around the Earth, and so it is possible that these points could be due to space debris falling into the atmosphere, or dust from a diffuse cloud orbiting the Earth. This particular group of meteoroids deserve much further study.

On occasions, there is an extra feature which shows in these graphs, and that is the presence of a narrow "spike" of velocities superimposed on the general shape. This arises when meteor showers are present and dominant. The speed associated with such spikes is that of the entrance speeds of meteors from the shower, and there is usually only a very small spread in values. Indeed, it appears (from multi-frequency comparisons of the same meteors) that the method can determine the speeds to accuracies of less than 0.5 km/s. By combining information about entrance speeds, and radiant determinations, it is possible to determine the orbits of meteor streams. This is very important knowledge for meteor astronomers, since it allows them to associate particular meteor streams with particular comets. An outstanding example of this appears in the paper by Arlt et al. (1999), where a study of the June Bootids meteor shower was able to prove for the first time that this meteor stream was associated with a comet called Pons-Winecke. Such studies are relatively simple with the SKiYMET radar.

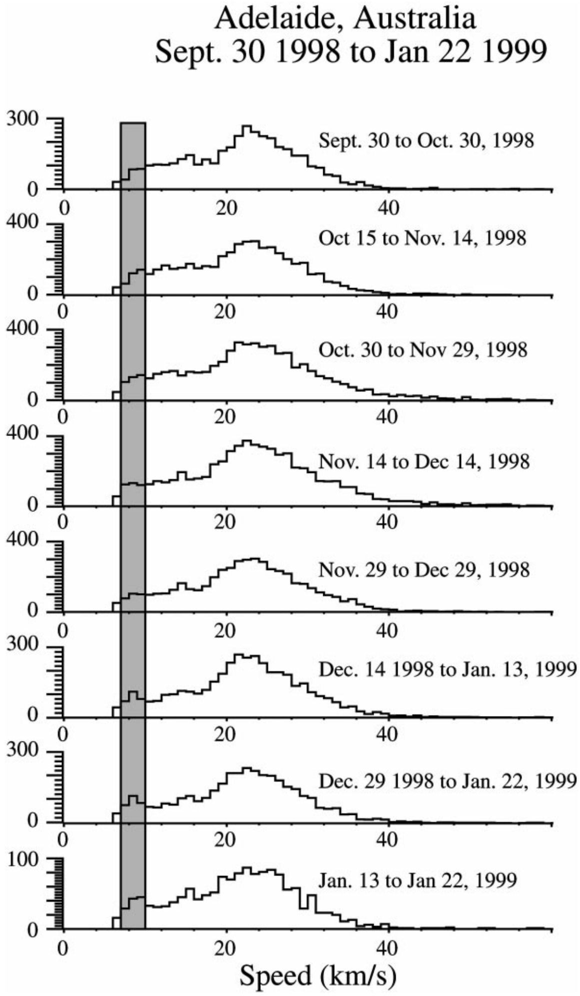

Fig. 7. Typical distributions of the entrance speeds of meteors as they enter the atmosphere, for a PRF of 2144 Hz, and a 4-point coherent integration. Speeds of around 7–8 km/s are highlighted. These data were produced at a site called Delamere, near Adelaide, Australia, but similar graphs have also been produced at other sites.

7. Conclusions

We have demonstrated how, by combining new technology and both old and new concepts in meteor physics, it has been possible to develop a new generation of meteor radar. The principles used in meteor selection, and examples of applications of the data, have been demonstrated. Examples have included both atmospheric and astronomical application of this instrument.

Acknowledgements

This project involved help and support from many people. We are especially grateful to Michael Smith, Shane Dillon, Adrian Murphy, Daniel O'Connor, Jacqui Nelson, Anna Hocking, Suzanne Hocking, David Hocking and Jeffrey Hocking. Useful discussions with Peter Brown and Jim Jones are also acknowledged. Support from Alan Manson and Chris Meek during an installation at Saskatoon is also acknowledged, as is the help from the staff at Starfire Optical Base in New Mexico during an installation of a SKiYMET radar at that site. Interactions with Kerry Ellis and Alan Webster are also recognized.

References

Arlt, R., Rendtel, J., Brown, P., Velkov, V., Hocking, W.K., Jones, J., 1999. The 1998 outburst and history of the June Bootid meteor shower. Monthly Notices of the Royal Astronomical Society 308 (3), 887.

Avery, S.K., Riddle, A.C., Balsley, B.B., 1983. The Poker Flat, Alaska, MST radar as a meteor radar. Radio Science 18, 1021–1027.

Avery, S.K., Vincent, R.A., Phillips, A., Manson, A.H., Fraser, G.J., 1989. High latitude tidal behaviour in the mesosphere and lower thermosphere. Journal of Atmospheric and Terrestrial Physics 51, 595–608.

Briggs, B.H., 1980. Radar observations of atmospheric winds and turbulence: a comparison of techniques. Journal of Atmospheric and Terrestrial Physics 42, 823–833.

Briggs, B.H., 1984. The analysis of spaced sensor records by correlation techniques. In: Vincent, R.A. (Ed.), Handbook for MAP, Vol. 13, Ground Based Techniques. Univ. of Illinois, Urbana, pp. 166–186.

Brown, P., 1999. The Leonid meteor shower: historical visual observations. Icarus 138, 287–307.

Brown, P., Hocking, W.K., Jones, J., Rendtel, J., 1998a. Observations of the Geminids and Quadrantids using a stratosphere–troposphere radar. Monthly Notices of the Royal Astronomical Society 295, 847–859.

Brown, P., Jones, J., 1995. A determination of the strengths of the sporadic radio-meteor sources. Earth, Moon and Planets 68, 223–245.

Brown, P., Simek, M., Jones, J., 1997. Radar observations of the Leonids: 1964–1995. Astronomy and Astrophysics 322, 687–695.

Brown, P., Simek, M., Jones, J., Arlt, R., Hocking, W.K., Beech, M., 1998b. Observations of the 1996 Leonid meteor shower by radar, visual and video techniques. Monthly Notices of the Royal Astronomical Society 300 (1), 244–250.

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

Cervera, M.A., Elford, W.G., Steel, D.I., 1997. A new method for the measurement of meteor speeds: the pre-t_0 phase technique. Radio Science 32, 805–816.

Hecht, E., Zajac, A., 1974. Optics. Addison-Wesley, London, p. 565.

Hocking, W.K., 1997a. System design, signal processing procedures and preliminary results for the Canadian (London, Ontario) VHF Atmospheric Radar. Radio Science 32, 687–706.

Hocking, W.K., 1997b. Recent advances in radar instrumentation and techniques for studies of the mesosphere, stratosphere and troposphere. Radio Science 32, 2241–2270.

Hocking, W.K., 1999a. Temperatures using radar–meteor decay times. Geophysical Research Letters 26, 3297–3300.

Hocking, W.K., 1999b. Real-time meteor entrance speed determinations made with interferometric meteor radars. Radio Science, in press.

Hocking, W.K., May, P., Roettger, J., 1989. Interpretation, reliability and accuracies of parameters deduced by the spaced antenna method in middle atmosphere applications. Pure and Applied Geophysics 130, 571–604.

Hocking, W.K., Thayaparan, T., 1997. Simultaneous and co-located observation of winds and tides by MF and Meteor radars over London, Canada, (43°N, 81°W) during 1994–1996. Radio Science 32, 833–865.

Hocking, W.K., Thayaparan, T., Jones, J., 1997. Meteor decay times and their use in determining a diagnostic mesospheric temperature-pressure parameter: methodology and one year of data. Geophysical Research Letters 24, 2977–2980.

Jones, J., 1975. On the decay of underdense radio meteor echoes. Monthly Notices Royal of the Astronomical Society 173, 637–647.

Jones, J., Brown, P., 1993. Sporadic meteor radiant distributions: orbital survey results. Monthly Notices of the Royal Astronomical Society 265, 524–532.

Jones, J., Brown, P., 1994. The radiant distribution of sporadic meteors. Planetary and Space Science 42 (2), 123–126.

Jones, J., Morton, J.D., 1982. A method for imaging radio meteor radiant distributions. Monthly Notices of the Royal Astronomical Society 200, 281.

Jones, J., Webster, A.R., 1992. Forward-scatter radiant mapping. Asteroids, Comets, Meteors. Lunar and Planetary Institute, Houston, pp. 273–276.

Jones, J., Webster, A.R., Hocking, W.K., 1998. An improved interferometer design for use with meteor radars. Radio Science 33, 55–65.

Jones, W., Jones, J., 1990. Ionic diffusion in meteor trails. Journal of Atmospheric Terrestrial Physics 52, 185–191.

Larsen, M.F., Palmer, R.D., Fukao, S., Woodman, R.F., Yamamoto, M., Tsuda, T., Kato, S., 1992. An analysis technique for deriving vector winds and in-beam incidence angles from radar interferometric measurements. Journal of Atmospheric and Oceanic Technology 9, 3–14.

Larsen, M.F., Roettger, J., 1991. VHF radar measurements of in-beam incidence angles and associated vertical-beam radial velocity corrections. Journal of Atmospheric and Oceanic Technology 8, 477–490.

McKinley, 1961. Meteor Science and Engineering. McGraw-Hill, Toronto.

Nakamura, T., Tsuda, T., Tsutsumi, M., 1997. Development of external interferometer for meteor wind observation attached to the MU radar. Radio Science 32, 1203–1214.

Norton, A.P., 1973. Norton's Star Atlas, 16th Edition. Gall and Inglis, Edinburgh, p. 153.

Roper, R.G. (Ed.), 1987. Handbook for MAP, Vol. 25. Scostep Secretariat, University of Illinois, Urbana, IL, USA, 422 pp.

Roettger, J., 1981. Investigations of lower and middle atmosphere dynamics with spaced antenna drifts radars. Journal of Atmospheric Terrestrial Physics 43, 277–292.

Roettger, J., Ierkic, H.M., 1985. Postset beam steering and interferometer applications of VHF radars to study winds, waves and turbulence in the lower and middle atmosphere. Radio Science 20, 1461–1480.

Tsuda, T., Nakamura, T., Kato, S., 1987. Mean winds observed by the Kyoto Meteor radar in 1983–1985. Journal of Atmospheric and Terrestrial Physics 49, 461–466.

Tsuda, T. et al., 1995. A preliminary report on observations of equatorial atmospheric dynamics in Indonesia with radars and radiosondes. Journal of the Meteorological Society of Japan 73, 393–406.

Tsutsumi, M., Tsuda, T., Nakamura, T., 1994. Temperature fluctuations near the mesopause inferred from meteor observations with the middle and upper atmosphere radar. Radio Science 29, 599–610.

Woodman, R.F., 1985. Spectral moment estimation in MST radars. Radio Science 20, 1185–1195.

Zrnic, D.S., 1979. Estimation of spectral moments for weather echoes. IEEE Transactions on Geoscience Electronics GE-17, 113–128.