Met Office

Observations Based Research

Capabilities and operation of the FAAM EZlidar

OBR Technical Note No. 79

25 March 2010

Franco Marenco

Contents

1 System description

The lidar system, an ALS450 manifactured by Leosphere, is an elastic backscattering lidar with daytime capability, suitable for aerosol and thin cloud observations. This system is light, compact and simple to operate, is designed to be operated on the ground on a 24-hour basis, is both weatherproof and eye-safe, and is meant to keep alignment for a very long time (Note: eye-safety is at the moment to the European norm; the system can in principle be made compliant to the U.S. norm as well). Very little hardware maintenance is in principle required, as it is limited to flashlamp replacement, and paying attention that the coolant does not go below freezing point during both operations and storage.

2 Aircraft fitting and operation

In the present configuration, the lidar is mounted on the aircraft in a nadir-viewing geometry, with a backward tilt of 4.2°. As the aircraft normally flies with a forward pitch, this ensures a near-vertical orientation. The nadir viewing geometry is much more efficient than the traditional zenith-viewing geometry, since the backscatter coefficient will typically increase with range and partially counter the effect of distance and extinction; this reduces the dynamic range of the signal hitting the receiver.

As the full overlap between the emitted beam and the receiver field-of-view is achieved ~300 m below the aircraft, one should aim at flying at least that much above the target. If quantitative information is seeked, it has also to be mentioned that far more accurate postprocessing algorithms

Table 1: Nominal specifications of the ALS450 lidar system.

Emitted wavelength	354.7 nm
Receiver bandwith	0.36 nm
Pulse Energy	16 mJ
Pulse repetition Frequency	20 Hz
Pulse duration	4 ns
Pulse-to-pulse stability	5–7%
Beam diameter	2.5 mm (laser), 30 mm (beam exp.)
Beam divergence	$\leq$ 1 mrad (laser), 0.2 mrad (beam exp.)
Vertical resolution	1.5 m
Overlap range	150 m (95%), 300 m (100%)
Maximum range	user defined, typ. $\leq$ 15 km
Integration time	user defined, typ. 5 – 30 s
N.D. filter optical density	0.7 $\parallel$, 0.3 $\perp$
Channel 0	analog, $\parallel$
Channel 1	analog, $\perp$
Channel 2	photon count, $\parallel$
Channel 3	photon count, $\perp$
Flashlamp lifetime	$30 \cdot 10^6$ shots (415 hours operation)
Coolant	water, or ethalene glycol 20% solution
Coolant freezing point	0°C or -8.9°C, respectively
Eye-safety compliance	EN60825-1

Figure 1: The ALS450 lidar head.

can be applied if a sufficient layer of clear air (purely molecular atmosphere) is available above the target; this implies flying 600–1000 m above the target boundary (e.g., if studying the surface layer, one will want to fly 600–1000 m above the top of the PBL aerosol layer, and not simply 600–1000 m above the surface layer).

The main choice the lidar operator has when running the lidar is the integration time. A shorter integration time gives in principle a better resolution, but this is not always true as the signal-to-noise ratio is then increased, particularly at the far range and when observing optically thin layers (the signal can be later degraded in resolution, which improves the signal-to-noise ratio, but it is a waste of resources not to aim at the correct resolution from the beginning). A very short integration time will also have the following disadvantages: (a) dramatically slow down the real-time viewer software; (b) increase the amount of data that have to backed up (the lidar data take up a lot of disk space!); and (c) require additional effort and computing power for the subsequent data analysis. It is therefore preferable to match the integration time to the scale of the atmospheric feature under study and to the correct signal-to-noise requirements.

The horizontal resolution is derived by combining the speed of the aircraft (typically 100–150 m/s) with the integration time set for acquisition by the user. The vertical resolution is 1.5 m when flying at a constant level; when profiling it is degraded, and can be determined by combining the integration time with the vertical ascent or descent rate (typically 5–10 m/s). Due to these large speeds, a reasonable integration time would lie between 2 and 5 s, the larger value being preferred when spatial scale is not critical. Note that this is already shorter than typical integration times suggested by the lidar manifacturer (5–30 s).

When turning, the aircraft roll can vary very quickly, affecting the viewing geometry quite rapidly, and this has been seen to introduce distortion and anomalous features into the profiles. One should therefore not excessively rely on the lidar during turns.

3 On-board viewer

The acquisition program supplied by the manifacturer includes signal real-time processing and display; however this has not been found satisfactory for two reasons. The first reason is that the software was designed for a ground-based zenith viewing geometry and is unable to take the aircraft position and attitude into account. The second reason is that only a custom designed software can give us the necessary flexibility for quick adaptations and upgrades, and full control on the data processing algorithms. We have therefore developed a bespoke real-time viewer.

The viewer program runs on the lidar computer, simultaneously with the Leosphere acquisition program (which remains in charge of running the instrument and recording the data). RemoteDesktop permits the lidar user to sit elsewhere on the aircraft and operate from there. The viewer is a suite of IDL routines that will automatically read any new lidar data file appearing on disk, retrieve aircraft information from the HORACE host through on-board networking, and then perform some basic processing and displaying. No derived data are recorded, as processing at this stage is considered preliminary and subject to re-processing later on the ground; however the on-board operator can take screenshots, useful for highlighting the main features identified during flight. Example screenshots can be seen in Fig. 2 and 3.

Since the viewer is to be used for real-time operation, it should ideally be robust; however it is not easy to foresee all different possible cases that can arise and therefore a continuous maintenance and upgrade will be carried out whenever necessary. So far the on-board viewer has experienced

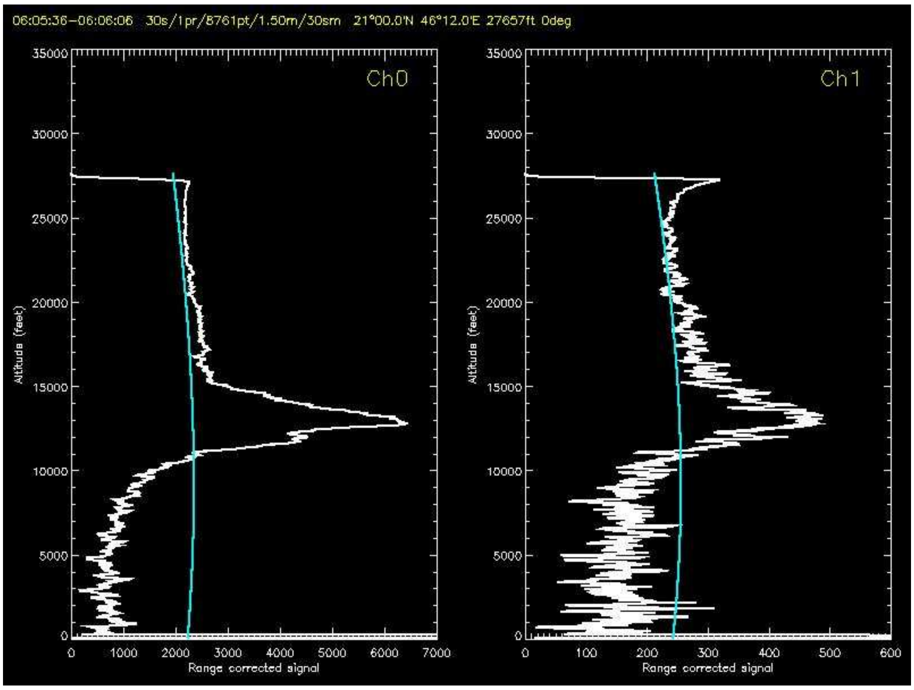

Figure 2: Example Lidardisplay on-board viewer screenshot: single profile showing the non-depolarized (Channel 0) and depolarized (Channel1) range corrected signal.

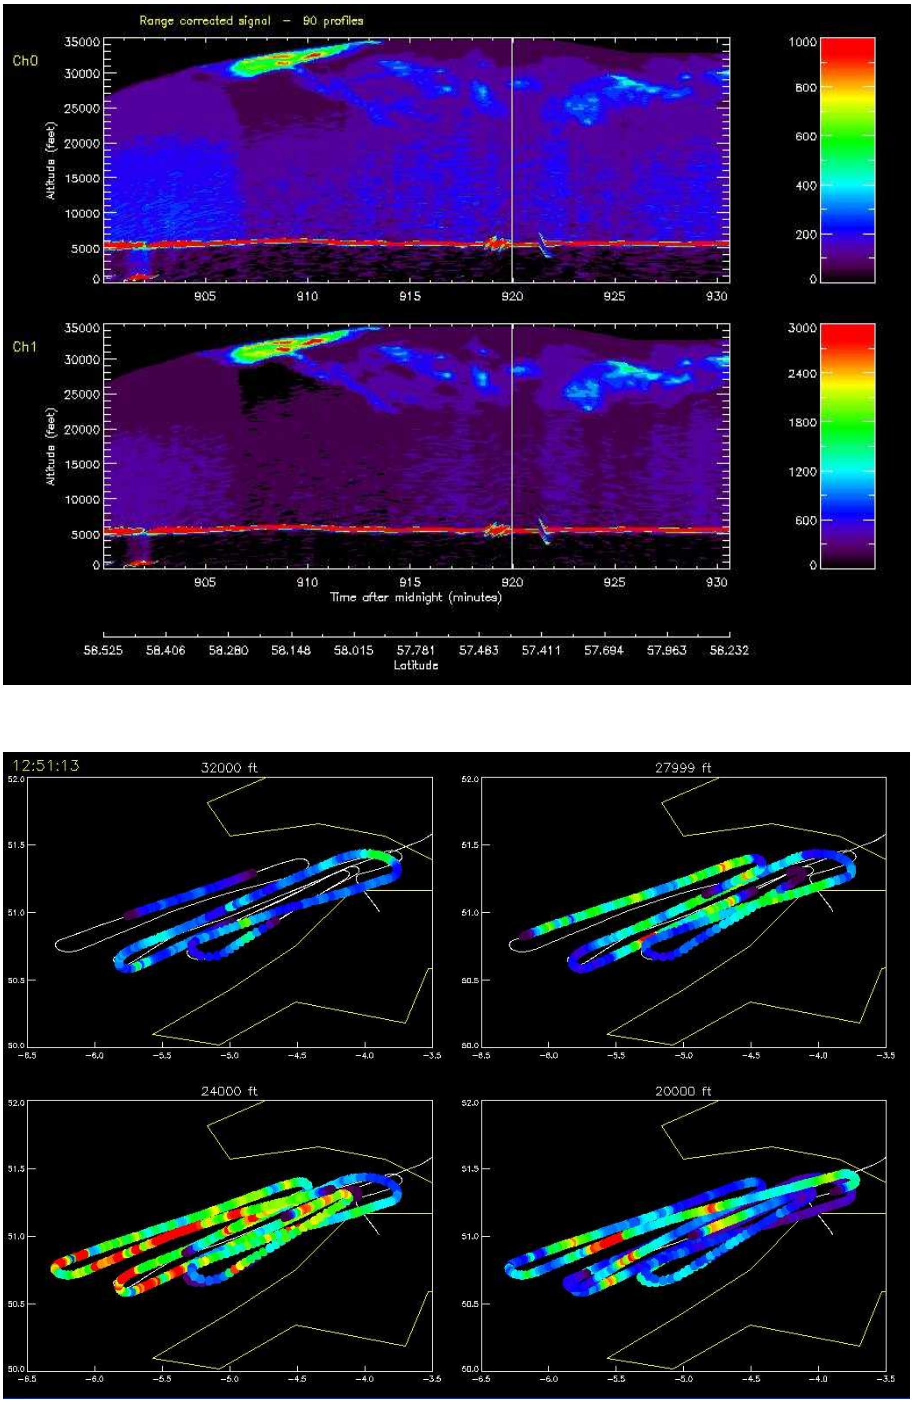

Figure 3: Example Lidardisplay on-board viewer screenshots: (top) Channel 0 and Channel 1 range corrected signal represented as a contour plot representing 20 minutes of data (note the cirrus and its associated fall streaks); and (bottom) geographic maps of range corrected signal for a selection of altitudes (32, 28, 24 and 20 Kft), summarizing a whole flight.

only one dramatic in-flight crash due to HORACE and network down time: the whole IDL session had to be killed and restarted, which currently requires the user to get to the back aircraft seat and access the lidar computer directly. Rectification, introducing a timeout, has been programmed into the viewer, but remains to be tested.

The viewer runs in a continuous loop, and checks the disk for new lidar datafiles to display. The user can interact with it by passing commands into a file, in order to change settings, scales and options as needed. As the first version of the viewer was very slow, a large effort has been put into optimizing it: dramatic speed increases are possible when using the correct IDL programming techniques (up to a factor of 20 for certain portions of the code).

4 Qualitative lidar data analysis

Basic lidar analysis will at first concentrate on the signal itself and the signal-to-noise ratio. Statistical (random) noise will in general be reduced by integrating (i.e. worsening the temporal resolution, and with it the horizontal resolution) and smoothing (i.e. worsening the vertical resolution). The amounts to which these processes will be applied is to be chosen, based on both the signal noise amplitude and the expected scale for the physical phoenomenon under study.

Geometrical computations based on the viewing geometry will also apply (aircraft height and attitude) in order to reference the lidar signal as a function of time, altitude above sea level, and the geographical coordinates (georeferencing).

The lidar signal processed in this way can already be used to visually infer very useful information on atmospheric targets, as illustrated here in a few examples. Fig. 4(a) shows an example of a nearly molecular lidar profile; the light blue line depicts the ideal lidar return from a purely Rayleigh scattering atmosphere. Note the near-range (top part) of the profile, where the lidar signal rises from zero to reach the blue line; this is the area below overlap, i.e. the transmitted beam is not yet fully in the receiver field-of view, and as such the signal intensity is not linear. This same feature is zoomed in Fig. 4(b), and the two green horizontal lines indicate the aircraft height and the nominal full overlap range. The same lidar profile is plotted again in Fig. 4(c), this time with the vertical scale extended down to the surface; this highlights two new features: a gradual increase in lidar signal, with respect to the purely molecular case, indicating the presence of boundary layer aerosols extending from the surface up to $\sim$ 900 m, and the surface peak itself, which is very sharp and strong (peak range corrected signal intensity in this case is at 33000).

Another example of the behaviour of the lidar signal with aerosols is displayed in Fig. 4(d), where a dust layer from the MEVEX campaign, centered at a 4000 m altitude, is shown. Note that nowhere in the profile is there a layer where the signal seems to follow the molecular profile, i.e. we must infer that aerosols are present everywhere between the ground and the aircraft. The molecular profile shown in the plot being arbitrarily normalized, the absolute position of the latter with respect to the first is irrelevant; a layer where the two go parallel would rather have to be seeked. Note also that

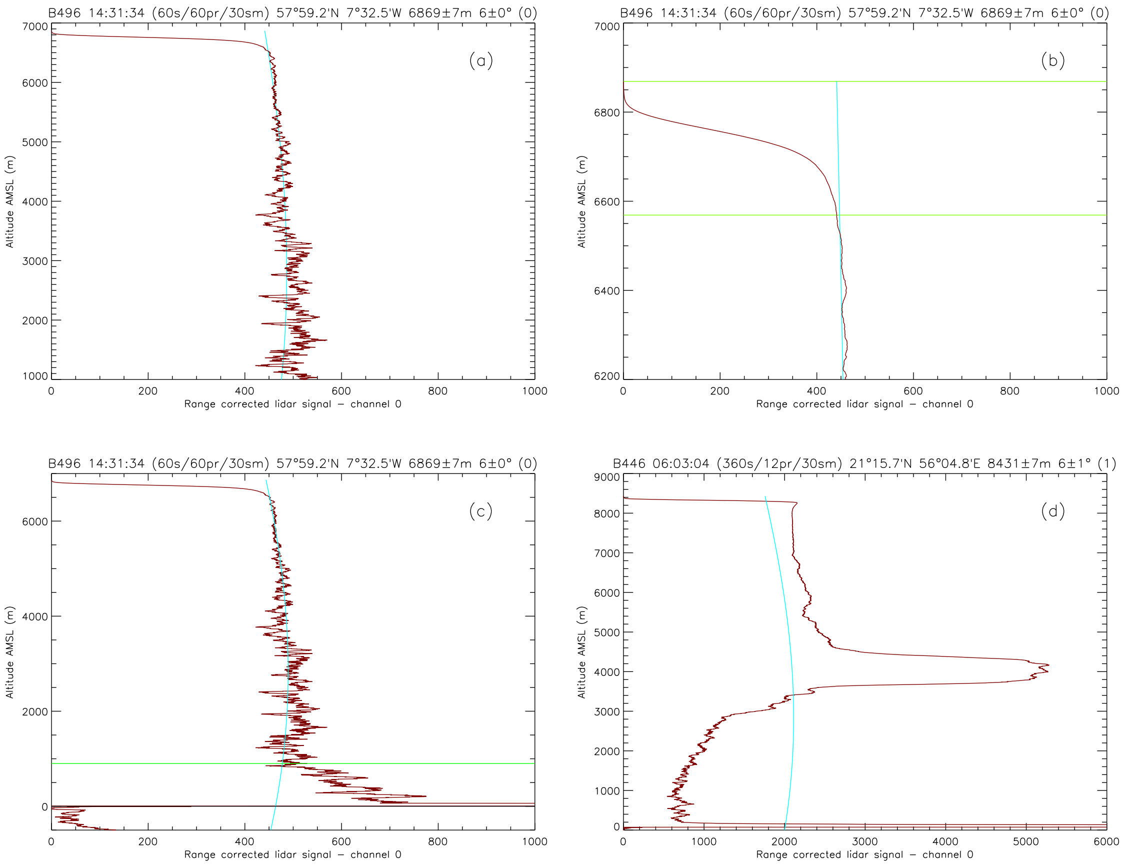

Figure 4: Lidar signal examples in absence of clouds: (a) a nearly molecular scattering profile; (b) a detail of the returns from ranges nearer than full overlap; (c) increased signal due to aerosol near ground, and surface spike; and (d) a dust aerosol layer. Light blue line: molecular (Rayleigh scattering) profile.

the signal returns from the atmosphere beyond the main aerosol layer is smaller (by a factor ~ 3) with respect to the signal returns above that layer: this is the effect of extinction.

But let us now have a look at how cloud signals will appear. Fig. 5(a) shows the example of a thin cirrus: note the sharp cloud boundaries. Note also all the little peaks in the part of the atmosphere above and below this cloud: the profile does not appear very smooth, and all those “little peaks” indicate actually the formation of other (subvisual) cloud layers, the first one identified even before overlap.

Apart from the surface, the targets examined so far are optically thin, i.e. their extinction is sufficiently small as to allow the signals from the layers beyond to be identified. Most clouds, however, will be optically thick, i.e. a coherent return will only be detectable from the first cloud boundary. Fig. 5(b) shows the signal obtained above a stratocumulus; the signal resembles very much the return from the surface: a very strong peak, with no signal beyond; and this despite the fact that this particular cloud extended down to ~ 900 m (green horizontal line). Obviously, in this case the only useful information in the signal is cloud top height. But the effect of extinction is displayed even

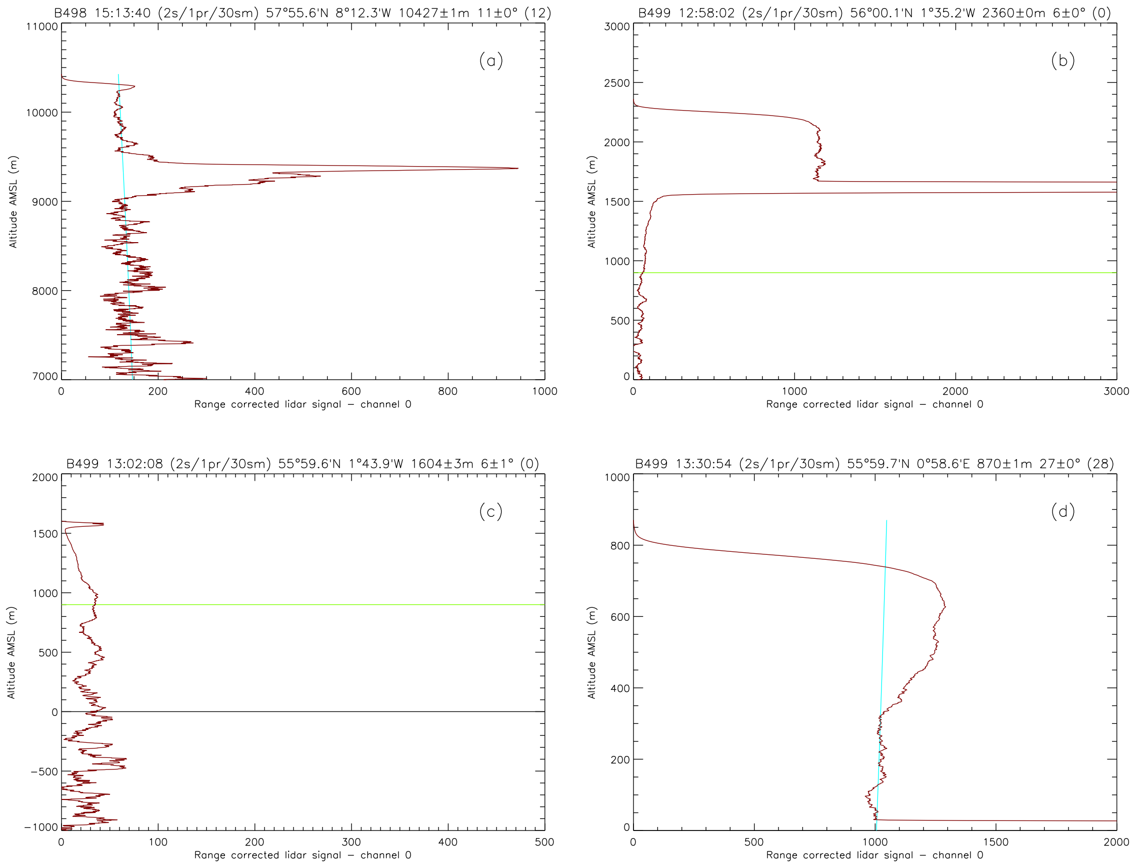

Figure 5: Cloud lidar signal examples: (a) flying above cirrus; (b) flying above stratocumulus; (c) flying in stratocumulus; and (d) flying at cloud base, within precipitation.

more dramatically in Fig. 5(c), taken while flying inside that same cloud: the whole signal is reduced to near zero, as the laser beam is estinguished before even reaching the overlap with the receiver field-of-view.

Another interesting case is shown in Fig. 5(d), obtained when flying at stratocumulus cloud base: an increased signal (with respect to purely molecular returns) is seen, and has been interpreted as precipitation thanks to simultaneous information from the cloud-physics probes. The altitude where the precipitation signal disappears ($\sim$ 320 m) is supposed to indicate where evaporation occurs. Note the importance of interpreting these data together with the knowledge that we are flying at cloud base, and together with the information that there is precipitation coming from the cloud physics probes. From the lidar dataset alone it would be quite difficult to identify this increased signal as being ascribed to precipitation.

To conclude this section on qualitative lidar analysis, it should be highlighted that from signal alone it can be hard to identify what the different possible targets actually are (aerosols, thin clouds, thick clouds, precipitation, surface), and that continuous logging of information on targets from the operator is very precious (this includes visual observations and information provided by other on-board probes).

5 Aerosol optical properties

What can be directly measured with a lidar is:

$$P \cdot R^2 = K\,\beta\,e^{-2\int_0^R \alpha dR},$$

where P is the lidar signal (expressed as light intensity, photon counts, detector current, or any other units), R is the range (distance between target and lidar), β is the atmospheric volume backscattering coefficient, α is the atmospheric volume extinction coefficient, and K is an instrumental constant (*lidar constant*). The above equation is known as the lidar equation, and it links the *range corrected signal* $P \cdot R^2$ with the quantity $\beta \cdot e^{-2\int_0^R \alpha dR}$, known as *attenuated backscatter*.

The lidar constant can drift and is usually not known with great accuracy; this is why in the previous section we have been happy with profiles of the range corrected signal, which can be computed directly. As shown above, these plots allow retrieval of a lot of information on the geometrical properties of the target; they do not however directly translate into the atmospheric optical properties at the lidar wavelength, β and α, unless a relationship between the two is known. In case of a two-component atmosphere (molecular Rayleigh scattering + aerosols or thin cloud), the knowledge of the molecular contribution is another essential input to the data analysis scheme; this however is quite straightforward by using either a standard atmosphere or a measured density profile (take-off/landing profile or dropsonde). Until now my approach has been to use a standard atmosphere, as this does not affect the lidar retrieval as much as the other choices outlined below.

a) Classical scheme

The solution of the lidar equation for a two-component atmosphere and an uncalibrated lidar has been widely discussed in the literature, and the two principal papers on it are those by Fernald (1984) and Klett (1985). In addition to the knowledge of the molecular profile as discussed above, the solution scheme they propose is based on the following assumptions:

1. a reference height must be set, where the backscatter ratio (ratio of total backscatter to molecular backscatter) is known;

2. the aerosol (or cloud) extinction-to-backscatter ratio (or *lidar ratio*) is known.

The first assumption is needed because the lidar constant is unknown, i.e. every signal profile will be “calibrated” at processing time against the atmosphere at a given height. Actually, I implement a slight modification of this algorithm, and instead of a single height I use a height interval as this helps reduce the influence of the signal statistical noise. In general, this interval will be chosen in a portion of the atmosphere dominated by Rayleigh scattering, as identified by qualitative data

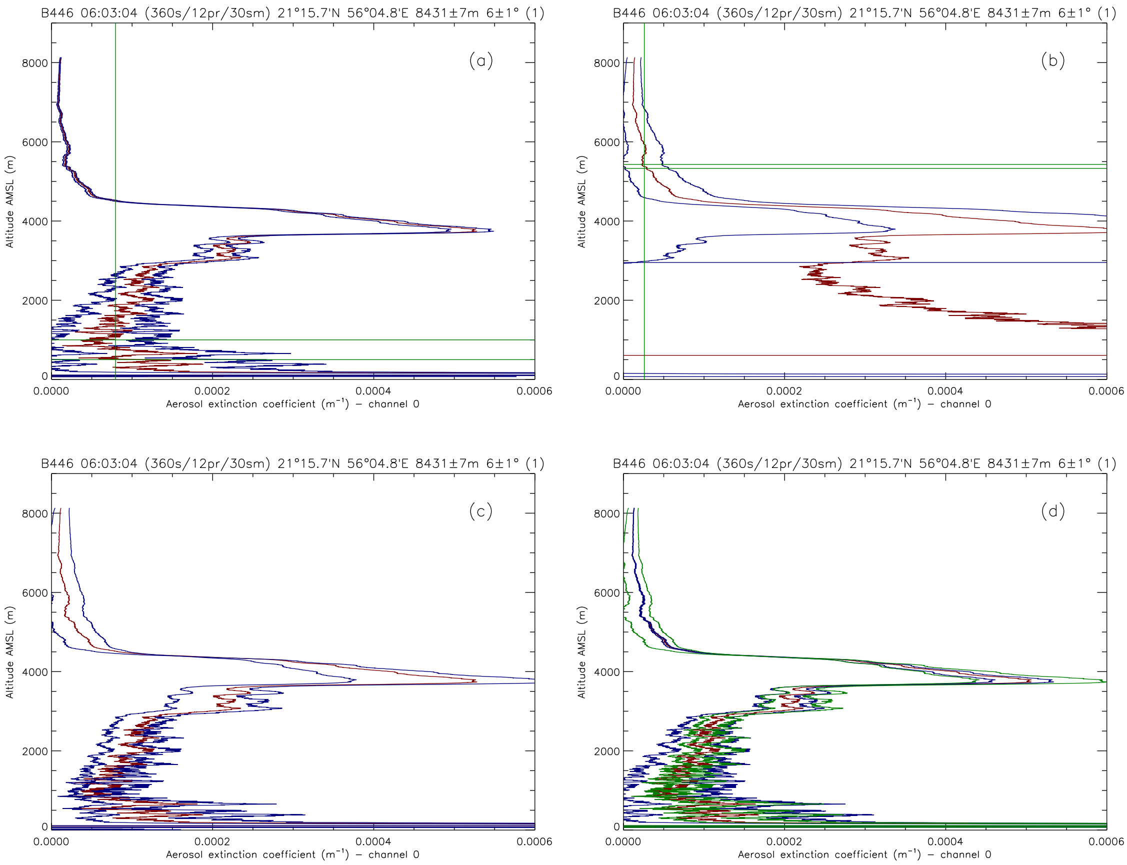

Figure 6: Aerosol extinction coefficient sensitivity test upon assumptions, within the classical inversion scheme: (a) reference backscatter ratio, far range; (b) reference backscatter ratio, near range; (c) lidar ratio. Panel (d) depicts the result of the iterative solution with different choices of the reference backscatter ratio at both near and far range.

analysis; this then allows us to set the reference backscatter ratio to a value near 1 or slightly larger (e.g. 1.05). If no clear air portion is available in the profiles, ancillary information is needed, for instance from the nephelometer and/or from particle counters.

It has to be stressed, however, that the effect of this "calibration" is not a simple multiplication onto the final product (extinction or backscatter profile). Rather, the lidar equation converges to a solution which is independent of the reference backscatter ratio when moving inward from the reference point, and it diverges outward. In other words, the derived backscatter or extinction profile is stable if the reference height interval is chosen beyond the aerosol/cloud layer (below in a nadir-viewing geometry), whereas the solution is unstable when the reference point is chosen near the lidar. This is clearly shown in Fig. 6(a) and (b), where the effect of a 100% error on the assumed aerosol contribution at the reference point is shown. The green horizontal lines denote the reference height interval, and the green vertical line the *a priori* reference aerosol extinction. In both cases a lidar ratio of 30 sr was assumed, whereas the reference backscatter ratio has been derived from the nephelometer profile: 1.36 at 500–1000 m in panel (a) and 1.19 at 5330–5430 m in panel (b).

Table 2: Lidar ratio database recommended by Leosphere.

Aerosol or cloud target	LR (sr)
background	15
maritime	20
maritime polluted — Japan	67
maritime polluted — US	50
arctic clean	20
arctic polluted	35
continental clean	40
continental clean up to 3 km	33
continental rural up to 4.5 km — US	36
continental urban up to 4.5 km — US	36
continental polluted dry	60
continental polluted wet	50
continental very wet	50
continental very dry	100
continental polluted up to 3 km	77
polluted strongly absorbing	90
dust in PBL	45
dust above 3 km	30
forest fire	65
rain forest	35
water cloud clean high LWC	50
water cloud clean low LWC	80
water cloud polluted	80
water cloud clean high IWC	15
cirrus cloud at 9–10 km	9.1
cirrus cloud at 8 km	7.2
cirrus arctic cloud	12

Therefore, although the nadir viewing geometry is penalized into choosing a reference point beyond the target because there is often no aerosol-free area near the surface, and thus ancillary data will be needed, we conclude that the first of the Fernald-Klett assumptions is not very critical, provided that the reference point is chosen at the far range (below target).

Let us now examine Fig. 6(c), where the consequence of picking an *a priori* lidar ratio is displayed: the dark red line is based on a lidar ratio of 30 sr, whereas the blue lines are based on lidar ratios 15 and 45 sr. This is a rather controversial issue, and lots of research exists on how to pick a lidar ratio, but the fact is that this choice is in general rather arbitrary, and even aerosols of a similar nature could show very different lidar ratios (see e.g. Table 2). It can be seen from Fig. 6(c) that although a direct proportionality between the lidar ratio and the resulting extinction coefficient is not verified, the whole of the profile is affected.

b) Iterative scheme

As the choice of the lidar ratio is rather arbitrary, there is lots of literature on different methods of constraining the solution of the elastic backscatter lidar equation. Some of the studies rely on

Raman lidar, which can separate extinction from backscatter, but may be limited in range due to much smaller signal returns; some others will rely on ancillary measurements, such as i.e. the co-located total-column optical depth derived from a sun-photometer at the same wavelength.[1] Finally, I'll mention a particular case where the elastic backscatter lidar signal itself is sufficient for inferring a lidar ratio (assumed to be constant with height). This occurs when it is possible to set *two* reference heights with their accompanying backscatter ratio, one above the aerosol layer, and the other one below: this is sufficient to constrain the lidar equation. The method is briefly outlined in Di Girolamo *et al.* (1994) for a similar configuration, and details of the equations can be found in Marenco *et al.* (1997); those equations can be extended to our case although in the latter paper a different configuration was studied.

Fig. 6(d) shows the result, when the reference values derived from the nephelometer are set below the aerosol layer as in panel (a) and above the aerosol layer as in panel (b). This result is shown as a dark red curve; the resulting lidar ratio is 27.3 sr, and it is this time *measured*, although still vertically averaged. Sensitivity tests for completely wrong guesses of the reference aerosol backscatter ratios are also shown (100% error). The blue lines show the profiles derived by perturbing the far end backscatter ratio (the resulting lidar ratio gets perturbed by $\pm$ 1 sr, in this case), whereas the green lines show what happens when perturbing the near end backscatter ratio (resulting lidar ratio $\pm$ 10 sr in this case).

c) Aerosol retrieval strategy

Whereas visual inspection of every single profile ensures data quality, when analysing large datasets a certain degree of automation must be seeked.

In this respect, the classical scheme with a far range reference point is certainly much easier to implement, as the solution is not too much affected by the choice of the reference backscatter ratio, and moreover additional on-board instruments (such as the nephelometer or a particle counter) can help into this choice. However, the profiles still have to be checked to make sure that surface or cloud returns do not occur at the reference, as this might lead to unphysical results. The critical point in this approach remains the choice of the lidar ratio.

The iterative scheme requires additional care and human intervention, in order to set a second reference point at the near range; and moreover it would not in principle apply to all profiles, as two "relatively clean" layers below and above the target must be assumed; these would not be available e.g. when flying inside the target layer. It can however be applied to selected profiles sampled from sufficiently far above the target layer, ideally with a useful clear air layer at the near range; this way the lidar ratio can be inferred, and then generalised for the application of the classical scheme to the rest of the data.

[1] On some campaigns we'll try to get additional information from a sun-photometer on the ground; in SAMBBA, for instance, we plan on having a Microtops operated at the base airport.

6 Comparison with the nephelometer

After all this work for setting the reference point and the lidar ratio, it might be interesting to compare with the nephelometer. The comparison is displayed in Fig. 7 (dust aerosols from MEVEX), and shows a large discrepancy between the two instruments (factor of 2). More work on the instruments and on the data inversion is needed to asses the reason for this.

It might be thought that this difference is ascribed to the fact that absorption properties of dust in the UV may be different than in the visible; this seems however to be ruled out if one looks at the spectral behaviour of the dust extinction coefficient displayed in Fig. 8. According to it, dust extinction should be 7% larger at 355 nm than at 550 nm, and not 50% smaller.

As in Fig. 7 the lidar extinction coefficient is derived from Channel 0 only (non-depolarized), an improvement might arise if Channel 0 and Channel 1 are analyzed together, taking thus into account total backscatter rather than non-depolarized backscatter only; this is however not believed to alter the result by more than 10–20%.

7 Work plan

The work on the lidar data analysis is to be continued in preparation for the SAMBBA campaign; moreover as the lidar is back on the aircraft new data wait to be processed. It is therefore necessary to concentrate on the following; as it is a large workload, the priority will be set on the first three points.

1. Define the optimal method for exploitation of the lidar signal, after the meeting held at Leo-

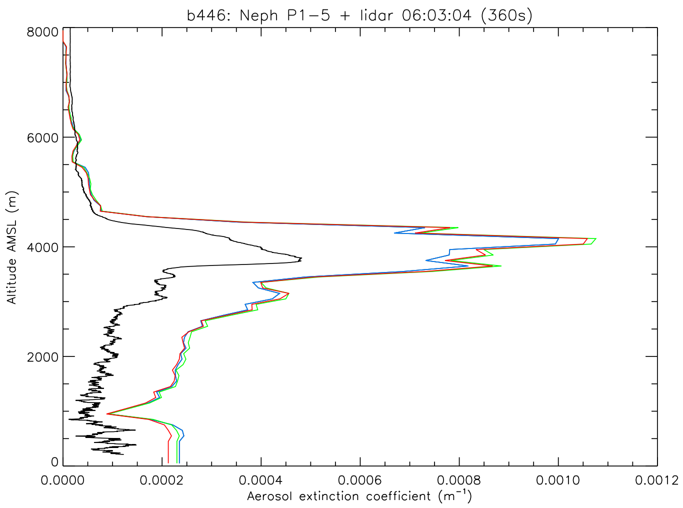

Figure 7: Lidar extinction coefficient of dust at 355 nm (black line), compared to the nephelometer scattering coefficients at 440 nm (blue line), 550 nm (green line) and 670 nm (red line). These measurements were taken during MEVEX. The nephelometer profiling of the aerosol layer occurred ~45 min. after the lidar measurement.

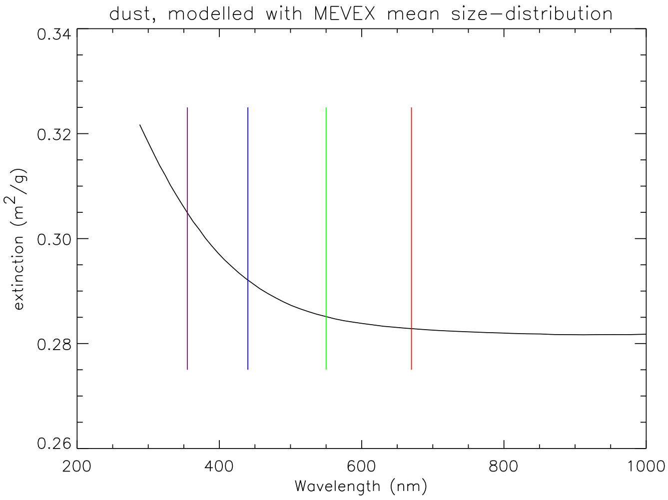

Figure 8: Wavelength dependency of the modelled dust extinction coefficient, as derived from the MEVEX size-distribution. The vertical coloured lines indicate the following wavelengths: 355 nm (purple), 440 nm (blue), 550 nm (green) and 670 nm (blue).

sphere in March.

2. Define the data analysis software tools, with a view towards efficiency and automation.

3. Use of the depolarization information, and its calibration.

4. Data analysis of the MEVEX aerosol data, as a test for the strategy outlined above regarding lidar ratios and automation of the target retrieval scheme.

5. Data analysis for selected recent cirrus / contrails flights.

8 Future developments

The lidar capability could be extended by adding a N_2 Raman channel; this would enable the separate determination of backscattering and extinction without the need for a lidar ratio, and thus would simplify the implementation of an automated data analysis scheme. Raman options are already available with Leosphere, either by replacing the depolarized channel with a Raman channel, or by enlarging the system to allow three channels (in the latter case, the lidar fitting within the aircraft might have to be modified). Note that as Raman signal returns are rather weak, the system performance with this option will have to be assessed before taking a decision.

While awaiting the installation of the large DIAL system onto the aircraft (which has presently been delayed), one could also temporarily place the present system horizontally within the aircraft cabin, and develop a mirror device enabling it to use the existing top and bottom windows, thus changing it into a dual pointing system. This would enable the collection of useful data even when flying at low altitude, as well as enable studies focused on the upper troposphere and lower stratosphere. Another alternative would be the purchase of a second system, which would give an almost complete picture of the whole column at any time (excluding the pre-overlap layers $\pm$300 m below and above the aircraft).

9 Flying recommendations

For a better use of the lidar, I outline a few recommendations that could be applied to the flying activity:

1. Regularly monitor the lidar temperatures when flying. If the lidar is fitted but not in use, everytime the aircraft flies it should still be heated and its temperature monitored. When not flying, it must be made sure that the lidar and the aircraft are stored above freezing point (winter detachments must thus provide a heated hangar). When the above cannot be ensured, the lidar must be drained.

2. When requesting the lidar for a flight, indicate: (a) what type of target you are interested in (e.g. cirrus, stratocumulus, dust, marine aerosol, etc.); and (b) what product you expect (e.g. PBL height, cloud top, aerosol extinction, cloud depolarization, etc.). By knowing, we'll be able to plan the best use of it.

3. When requesting the lidar for a flight, indicate the horizontal, vertical and temporal scale of the features to be studied, as best known from your scientific background.

4. Fly at least 300 m above your target, as there is not full overlap between emitter and receiver at a nearer range. For an optically thin target (aerosols, cirrus) consider flying 600–1000 m above the target at some point during the flight, as this will help deriving an extinction-to-backscatter ratio and thus reduce the uncertainty on the optical properties.

5. When sampling an area with large aerosol load (e.g. dust campaigns like MEVEX), if possible take a few profiles of the atmosphere from top (600–1000 m above plume) to bottom (near ground) in the area of interest: this will yield nephelometer and particle counter profiles that will be useful in the lidar inversion and/or for instrument intercomparison. Mind that the takeoff and landing areas might have different atmospheric properties and thus could be inappropriate. When the high flight level is reached, turn around and do a high altitude constant level run over the same airmass, while taking lidar profiles.

6. For aerosol campaigns, auxiliary information from a ground-based calibrated sun-photometer (e.g. Microtops in SAMBBA) can be used together with the lidar. In that case, take the time for a high altitude constant level run the nearest possible to the area sampled by the photometer. Ideally, this should also be the area where the aerosols are going to be sampled during most of the flight (rather than a fixed location), but this in practice might be difficult.

7. Save our resources and time by **not** requesting a lidar for low-level flying, as this would provide no useful data. If for instance you are sampling the surface layer, you could switch the lidar off during most of the flight, and have it switched on for collecting data during a high altitude overpass of the area of interest both before and after the in-situ sampling. Flying within fogs and optically dense clouds would also produce no useful lidar data.

8. If you are the lidar operator, try to continuosly log information that will later help identify what the targets seen in the lidar signal are: cloud, aerosol, precipitation, etc. Ancillary information from visual observation and from other probes (relative humidity, precipitation from cloud physics probes, etc.) can be very useful and should be logged as well. Take a screenshot if during flight you see some feature of particular interest, and archive it together with the data and possibly with your comment attached.

References

Di Girolamo, P., M. Cacciani, A. D. Sarra, G. Fiocco, and D. Fuà, “Lidar observations of the Pinatubo aerosol layer at Thule, Greenland,” *Geophys. Res. Lett.* **21**, 1295–1298 (1994).

Fernald, G. F., “Analysis of atmospheric lidar observations: some comments,” *Appl. Opt.* **23**, 652–653 (1984).

Klett, J. D., “Lidar inversion with variable backscatter/extinction ratios,” *Appl. Opt.* **24**, 1638–1643 (1985).

Marenco, F., V. Santacesaria, A. F. Bais, D. Balis, A. di Sarra, A. Papayannis, and C. Zerefos, “Optical properties of tropospheric aerosols determined by lidar and spectrophotometric measurements (Photochemical Activity and Solar Ultraviolet Radiation campaign),” *Appl. Opt.* **36**, 6875–6886 (1997).

Met Office
FitzRoy Road, Exeter
Devon, EX1 3PB
UK

Tel: 0870 900 0100
Fax: 0870 900 5050
enquiries@metoffice.gov.uk
www.metoffice.gov.uk