Met Office

SWS-SHIMS Calibration and Uncertainties

Interim report, to Christmas, 2017

OBR Technical Note 90
December 2017

A. K. Vance, J. Norwood-Brown, D. O'Sullivan

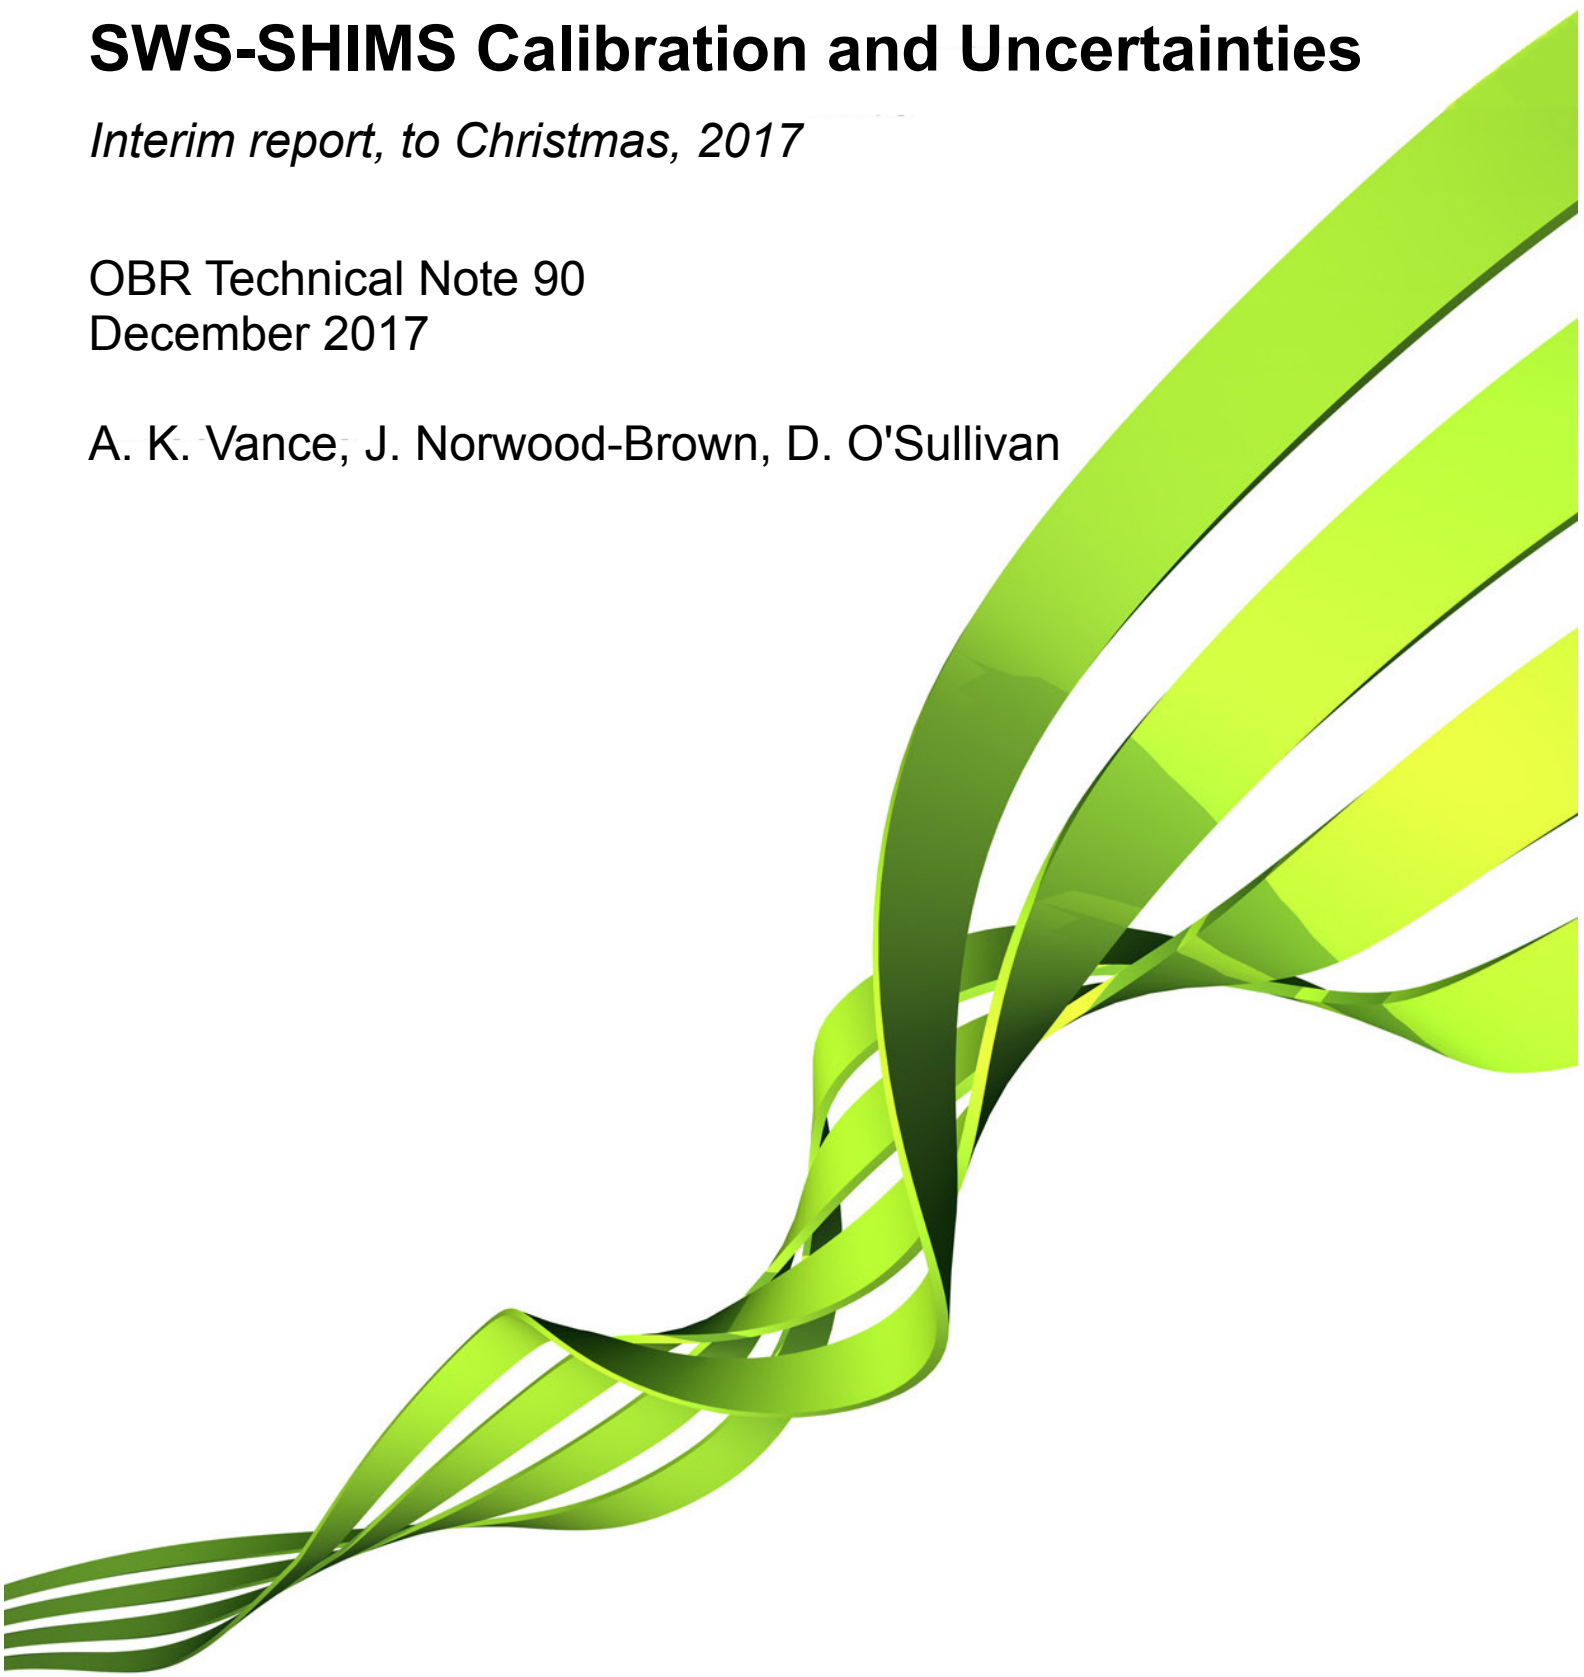

If printing double sided you will need this blank page. If printing single sided, please delete this page.

© Crown copyright 2017

Contents

© Crown copyright 2017 Met Office

Abstract.

Various investigations have been carried out on the SWS-SHIMS instrument with a view to improving its characterisation. Although some serious failings are reported, it is also believed that with greater attention to detail in its operation, substantially better data may be obtained than have often been, in the past. Where the aspects of the system were common to both SWS and SHIMS, darkroom investigations have generally been carried out with the SWS fore optics, for practical reasons.

Illumination of the SWS head by the direct solar beam appears to give rise to bias. The IDL processing code can produce good data but has inadequate filtering of dark spectra, and is unstable. Temporal linearity of the Zeiss modules is not good, but calibration at all required integration times will prevent errors from this source. The radiometric linearity of the NIR modules is questionable at low light levels. The long term stability of local calibration standards appears to be very good, but uncertainty in the repeatability of laboratory calibrations appears to be around 7%. Although strong evidence has been found of a serious stray light problem in SWS, in side-by-side comparison with a Cimel sun photometer, very good agreement was seen; the reason for these contradictory results is unknown. It has been found that even very small displacements of a SHIMS dome from the transfer standard can produce large errors but, properly calibrated, SHIMS appears to agree with the aircraft's broadband radiometers to about 10%. Likely variations in the temperature of the cooler will have no noticeable effect on data. Only very small errors are expected to be introduced through bending of SWS' fibres as the head rotates, but there is concern that significant errors in SHIMS data could result from mishandling of fibres during other aircraft work.

Acknowledgements

Thanks to Joelle Buxmann, Clare Ryder, Stuart Rogers and Ian Rule for their assistance with the work described herein.

© Crown copyright 2017 Met Office

1. Background.

Much work has been done by various people over many years and, whilst there are no grounds to question their work, it is plain that something has been missed, as robust, absolute calibration of SWS-SHIMS appears never to have been achieved. SWS-SHIMS data have been published at various times, and again, whilst these publications may appear to be sound, they are reliant upon various assumptions about the function of SWS-SHIMS which, to date, we have been neither proved nor disproved; that this has remained the case for so long is suggestive that these assumptions may not be entirely valid. The storage and documentation of historical data has frequently been inconsistent which has hampered investigations significantly. New processing code, written in Python, was believed to be at an advanced stage of development but this appears to have actually been the result of poor internal communications and, in effect, it does not exist. Archiving and characterisation work has almost invariably been made a secondary activity, behind scientific flying, which has, without doubt, contributed to the lack of progress in the past.

Particularly notable amongst previous work is Debbie O'Sullivan's 2012 report *Performance and Calibration of the SWS and SHIMS Instruments*. Although there is some overlap between O'Sullivan 2012 and this report, the former focuses more specifically on the performance of the Zeiss spectrometer modules upon which SWS-SHIMS is based, whilst here the concern is more with the fore optics and with operation of the system, as a whole, in the field. O'Sullivan 2012 is, therefore, recommended to the reader, both as a source of background information, and as a complimentary work.

Much of the work described herein was carried out using SWS but a significant amount of this (e.g. linearity) is equally relevant to SHIMS; the decision to use SWS for such items was made on the grounds that its fore optics simplified darkroom work.

2. Initial SWS case studies, B723 & B886, and IDL processing.

Flight B723 (CAVIAR, Bristol Channel, July 2012) was selected for having a high level box pattern likely to be clear overhead. A simulation using a standard atmosphere in HTFRTC, however, showed very large differences (figure 2.1); large differences also were evident between the UV/visible (visible) and near infrared (NIR) modules. Some confidence in the use of HTFRTC had been obtained by running simulations in the mid IR, and by superimposing a curve produced by modifying a 5800 K Planck function by λ^{-4}; these suggested that the simulation was unlikely to be wildly erroneous. It was also noted that significantly different spectra were obtained on opposing cross-sun legs of the box pattern.

© Crown copyright 2017 Met Office

There was a belief that a new, Python-based processing suite was nearing completion, but this appears to have not been the case, and to have been the result of poor communication. The most recent complete suite appears to be an IDL version written mostly in 2012. Confidence in the use of the IDL code was obtained by running it on the raw data from B723, labcal094 and labcal095, to obtain gains and calibrated data which where a very good match for the pre-existing calibrated data.

B886 (figures 2.2-2.4) was a specific, SWS test flight carried out over the Bristol Channel in January 2015, and consisting solely of one high level box pattern with a near-coincident CALIOP overpass. Figure 2.2 shows the ground tracks of both the aircraft (grey) and the satellite (red), with the periods for which there are SWS data highlighted in blue. Figure 2.3 shows the attenuated backscatter from CALIOP. The aircraft was able to penetrate the tropopause, and the CALIOP data indicate this to be at around 8 km, with no obvious aerosol above this, although it is possible that small but significant amounts of aerosol could still be present; the sky was reported as clear by the mission scientist. This flight had been preceded by a laboratory calibration (labcal102) with the integrating sphere, this having been calibrated in 2014. B886 was processed using the recent laboratory data, and compared to an HTFRTC simulation (using a standard, mid latitude, winter atmosphere) and, again, if was found that SWS disagreed with HTFRTC by a large amount at almost all wavelengths. Figure 2.4 shows the HTFRTC simulation in red (degraded to SWS' resolution), with two versions of the SWS data; light colours show the results of the standard IDL code whilst the dark blue and green show carefully selected spectra calibrated 'by hand'. Note that this version of HTFRTC included a bug relating to the treatment of near-zero radiances, which is what gives rise to the odd band-like depressions in the curve. The upper edge of the simulation is valid, however, and provides a useful baseline for comparison.

From the foregoing it would appear that although the existing IDL code works, in the sense that it can convert raw SWS data into something approximating to atmospheric spectra but what it produces can be non-physical (B886 NIR). These discrepancies could arise as a result of various things, including the laboratory calibration, instrumental factors during the flight, the processing software or combinations of these. Single spectra were selected from B723 and B886 and the entire calibration process was carried out from 'first principles' with new code written in Gnu Octave, without reference to the existing IDL code. The philosophy of this was that the chances of two people making errors in code written in two different languages, without reference to each other, and getting the same answers would be negligible. In figure 2.4, showing spectra calibrated using both the IDL and Octave codes, the results for the UV/visible module are very similar, but there is a substantial change in the NIR spectra. The Octave code, however, produces spectra closer to the model and, more significantly, which agree reasonably well with the UV/visible module in the overlap region. Clearly this exercise has not produced a resolution to the calibration problems but the good agreement

© Crown copyright 2017 Met Office

between the two calibrations in the visible gives confidence in the IDL code, while the poor agreement in the NIR is due mainly to the inability of the IDL code adequately to reject bad spectra, although it is possible that non linearity in the detector may also contribute.

In both flights it is apparent that on legs presenting the starboard side of the aircraft to the sun, the zenith radiance drops more slowly, as wavelength increases, than on port legs, as illustrated in figure 2.5 (B886). The fine lines show HTFRTC simulations for the three cross-sun legs of the B886 box pattern, coloured according to the side presented to the Sun: legs 1 and 3 present the port side of the aircraft to the Sun, whilst leg 5 presents the starboard side. The box pattern was misaligned, being rotated by approximately -5° (heading minus solar azimuth) such that the two cross-sun legs (3 and 5) are not true reciprocals (relative to the solar azimuth) and so one might expect some difference in the observed radiances, but as the heading on leg 1 differs from that on leg 3 by some 20°, it might be expected that this would result in a more significant difference. The thin lines show HTFRTC simulations for the three legs (assuming a clean atmosphere, with only Rayleigh scattering), and it is clear that there is minimal difference between these. The thicker lines show the 'calibrated' radiances from SWS, and it is immediately apparent that there is minimal difference between the two port runs (SWS head shadowed by fuselage) but that these differ substantially from the starboard (SWS head in direct sunlight) run showing that precise heading is of less significance the gross orientation of the aircraft.

Plots of aircraft attitude and heading show that, during B886 the aircraft generally flew with a 0.5° list to starboard, rolling ±1°, that the heading changed relative to the solar azimuth by up to -6° and that the pitch dropped by about 2° during the course of the box pattern. Simulations with HTFRTC suggest that this sort of variation in viewing angle will have had no significant effect on the measured radiances in clear sky conditions.

B723 and B886 show evidence of poor calibration and port/starboard bias in SWS, but give confidence in IDL processing code despite inadequate filtering of bad data.

3. Historical gains

Figure 3.1 and 3.2 show a number of past sensitivities (1/gain) for UV/visible and NIR modules. Large changes are seen in both but, from `~swsshims/record_of_data_location.ods` it would appear that from labcal001 to labcal049, the UV/vis module was unchanged and the NIR module was only changed twice: from labcal012 to labcal021 a different module was used. In terms of these plots, labcal001-labcal049 covers curves 8-16. A new NIR module was inserted on 19th June 2013, between the last two gain curves plotted, here. It can be seen that there is

© Crown copyright 2017 Met Office

significant variation during this period but curve 8 in the NIR (the only curve plotted from the 'other' NIR module) is not an outlier, and although there is a large change across the change of modules in 2013, it is not the largest change seen. Currently, the source of the gain curves dated earlier than labcal001 (curves 1-7, here) is not known. Although it is a concern that the gains, which should remain reasonably constant with time, do not, this is not necessarily a problem, providing that they are accurate.

Much greater variability than expected, is seen in past gain curves.

4. Detector Linearity

SWS is required to work over a very large dynamic range; in order to accommodate this it is common to integrate for various different lengths of time, according to the scene radiance - 500 ms in the case of B886, for example. The calibration is, however, carried out at only one light level and the range of integration times possible without saturating the detectors is restricted to less than 50 ms. As a result it is critical that the detectors' linearity is well understood and, although papers have been published using SWS data which rely on a linear response, to date, no assessment of linearity has been found. Linearity can be looked at in two ways: 'temporal', considering the relationship between the recorded counts and the integration time, whilst viewing a source of constant radiance, and 'radiometric', considering the relationship between the recorded counts and the source radiance, for any given integration time.

4.1 Temporal Linearity

The relationship between counts and integration time, for any given level of illumination, needs to be understood, stable, and (ideally) linear. Firstly, considering only dark measurements, these being the least dependent on external factors. If (1) spectra were inspected and any spectrum where an incompletely closed shutter was suspected, was rejected, (2) all spectra of the same integration time, from that calibration, are averaged, (3) the electronic offset (calculated by extrapolating back a linear fit to counts vs integration time) is removed, (4) each measurement is normalised by division by its integration time, and (5) the normalised counts for the shortest integration time is subtracted from each pixel, then, if the detectors are perfectly linear, and in the absence of noise, the result for each integration time should be identically zero. In practice there will be noise, but one would hope that there would be no discernible trend. Figure 4.1 shows the results for four sets of dark measurements , illustrating that this is not the case. The amount by which the dark counts differ from the those obtained at the shortest integration times are plotted as a percentage of the offset-corrected counts at the shortest integration times.

© Crown copyright 2017 Met Office

Figure 4.2 shows the results of repeating this exercise with measurements of the sphere; the difference of the dark-corrected counts per second from the value at the shortest integration times are shown as a percentage, as above. In this case the situation is complicated by some pixels saturating at the longer integration times which means that it is not appropriate to average the entire spectra. Two lines for each calibration are plotted: the thin lines are calculated using pixels selected where the recorded counts are low enough that measurements should still be valid at 500 ms. The thick lines come from pixels which saturate at integrations greater than about 250 ms. No data are included where any of the pixels are within 10 % of the detectors' maximum value.

Although the non linearity is not desirable, and the reason for it is not known, it should not cause any problems as long as the instrument is only used in the field at integration times which may be calibrated - extrapolation to other integration times is inadvisable.

4.2 Radiometric Linearity

Assessment of radiometric linearity is less straightforward. The integrating sphere is equipped with four bulbs which may be switched individually. Although it is not normal to operate the sphere without all four bulbs lit, it has been done in the past, and two calibrations of the sphere have been carried out recording the output of the sphere with different numbers of bulbs lit; one has 1, 2, 3, and 4 bulbs, and one has just 2 and 4. Unfortunately the dates of these calibrations are not known. Figure 4.3 shows a comparison of these calibrations, from which the output of the sphere would appear to be linear with bulb number to about 5 % over most of its range. Note that older calibrations were carried out at a lower resolution than more recent ones and so the former do not capture some of the radiometric structure apparent in the latter.

Given reasonable confidence in the linearity of the integrating sphere output, it was possible to run the sphere with 1, 2, 3 and 4 bulbs, and compare the radiance recorded by SWS under otherwise normal conditions; figure 4.4 shows the results. The spectra presented are as dark-corrected counts per second per lit bulb, expressed as a percentage so that figure 4.4 indicates the drop below the maximum value observed. Data are omitted when the detectors were saturated. Clearly, the ideal situation would be indicated by a single line through 0 %. For each of the four radiances, for each detector, a number of spectra are plotted; these correspond to the different integration times used. Again, these being normalised, a single line through zero would indicate perfection; the spread of values may be taken as an indication of the imperfect temporal linearity, discussed above.

In both detectors we see evidence of lower radiances being under read, to a small extent in the visible module but to a much greater extent in the NIR. This is a puzzle, as it

© Crown copyright 2017 Met Office

would suggest that, as the scene radiance falls, the agreement between the visible and NIR modules should worsen, but this has not yet been noted. Likewise, if SWS is susceptible to stray light, that effect might, under certain circumstances, be masked by this apparent drop in sensitivity, but as the visible module does not appear to be affected, differences would be expected in the overlap region.

That some of the NIR spectra exhibit a significantly different shape to others is also a concern.

Temporal linearity is not good, but issue can be avoided with appropriate calibration. Radiometric linearity is better in the visible module than in the NIR,and at higher radiances. For low radiance scenes uncertainties of 3 % and 15 % might be assumed.

5. Accuracy of Calibration of Standards.

Irrespective of anything done within the Met Office, a fundamental limit on the accuracy of SWS-SHIMS measurements is imposed by the accuracy of the calibration of the standards used to calibrate it. These are shown in figure 5.1, with the data provided by Labsphere for the integrating sphere plotted in black. The situation with the bulb used to calibrate SHIMS is less straightforward, however, and comes in various tables, dealing with different aspects, including two tables, both labelled as 'NIST uncertainties', which are different; both are plotted. The combined uncertainty, plotted in pink in figure 5.1, has been arrived at by the addition of tables T1, T2, T3 and the 1 % transfer uncertainty, in quadrature. As this figure is dominated by the NIST uncertainties, over which there is some doubt (due to the two sets) the figure could be lower, but it might be prudent to assume that, regardless of any improvements to the instruments or procedures, accuracy of better than ~0.8 % (SWS) or ~1.75 % (SHIMS) is not possible.

Accuracy of local calibration standards used for SWS and SHIMS are around 0.8 % and 1.75 %, respectively, over most of their ranges.

6. Stability of Calibrations.

Calibration certificates for the local standards (bulb and sphere) show these to be stable to a few percent over long periods. The transfer standard should be equally stable (or almost) over similar periods as it is based upon a similar bulb, fed from a highly stable power supply, although we don't have independent verification of this.

There have been questions raised over whether SWS-SHIMS has a different sensitivity every time it is set it up, so to simply look at a long series of standard measurements

© Crown copyright 2017 Met Office

(local or transfer) is likely to provide more information about the instrument than about the standards themselves, although the effects of both will, clearly, be involved. If, however, the instruments remain stable whilst they remain set up in the darkroom, on each individual occasion, and their performance does not change when they are moved from local to transfer standards, then the differences that arise from the set-up for that particular calibration session (e.g. from dust on optical elements or differences in alignment) should be the same for measurements of both standards. If this is the case (the standards and the instruments are stable for the duration of one calibration session), and the set-up used is adequately repeatable, then it might be supposed that the ratio of the counts (minus relevant dark current) obtained from the local standard, to the counts-minus-dark for the transfer standard, should not change with time.

Figure 6.1 shows pairs of standard measurements, four pairs for upper SHIMS (upper SHIMS) and SWS, and six pairs for lower (lower SHIMS), where there were no changes made (each pair done in same session). Spectra were inspected, suspect data were eliminated, and then spectra were averaged according to integration time. The relevant dark currents were subtracted, and then the measurement of the local standard was divided by the measurement of the transfer standard (figure 6.1 shows only data for 30 ms integrations). The precise shape of the resulting curves is not considered important providing that they are the same every time; if they are, it might be taken as confirmation that the standards are stable, as they are currently used. There are clearly insufficient data here to draw robust conclusions, but if these cases are representative they would seem to suggest that the repeatability of our *laboratory* calibrations is about generally ±5-8 %.

Stability of laboratory calibrations may be taken as ~7 %.

7. Courtyard Cimel Comparison.

The opportunity occurred on the morning of 9th June to compare SWS with a recently calibrated Cimel CE318 sun photometer at Exeter. The Cimel and SWS were set up in the courtyard, zenith viewing, along with OBR's Leosphere aerosol lidar. Conditions were not ideal as there was broken low cloud throughout the morning but there were sufficient gaps in the cloud to allow good measurements. Coincident, range corrected lidar signal and relative depolarisation ratio show a well defined boundary layer capped by broken cloud. From the lidar (figure 7.1), and by eye, there was no evidence of of cloud or aerosol aloft, although this cannot be ruled out. The presence of small amounts of scatterers should not, however, affect the comparison providing that their distribution is adequately uniform and slow-changing.

© Crown copyright 2017 Met Office

Figure 7.2 shows six comparisons made during this exercise. The Cimel produces two radiances at each wavelength, shown in red and orange in each case; these are measured using different gain settings, and the different wavelengths are measured sequentially, not simultaneously, with the result that the data in each of the individual axes are spaced over about forty seconds. In ideally clear conditions, these two radiances will be identical but if there are significant changes within the field of view (FoV) of the instrument over this period, the two values will differ, perhaps substantially (e.g. bottom right axes). The dark blue and green curves are an average of the SWS visible and NIR data, respectively, recorded in the forty seconds following the Cimel start time (noted in each plot); the number of spectra in each average is also noted in each plot. The light blue and light green lines indicate the highest and lowest radiance measured at each wavelength in that forty seconds, and so are are indicative of the maximum variations observed by SWS. The figure labelled as 'cloud based detection' refers to the percentage of lidar measurements from ten seconds before the Cimel start time to fifty seconds after it, where a cloud base height could be retrieved.

In general (surprisingly) the agreement is very good both between the SWS modules and between SWS and the Cimel, particularly earlier in the comparison. In the top left axes, for example, 310 SWS spectra contribute to each average, with very little spread (giving good confidence of an unchanging, clear sky), which agree extremely well with the Cimel. There is only a small difference between two Cimel radiances. A possibly interesting feature, however, is that later in the day SWS does appear to report higher radiances than the Cimel; note that at 113822 (middle, right), there, again, appears to be very little change in the fields of view, but there is a definite difference between the Cimel and SWS. Although one could not state without doubt that there is a definite trend, here, one might suspect this to be possible.

When SWS was first installed on G-LUXE, its pointing head had an entrance window on the surface of the head but this was replaced with the current arrangement (short, tubes with windows at their inner ends) as it was noticed that the external window would flare badly whenever exposed to the direct solar beam. By locating the first transmissive element at the bottom of short tubes illumination by the direct solar beam is prevented until the solar zenith angle is less than 26° for the UV/visible module and less than 20° for the near infrared (assuming the instrument is pointed to the zenith). As the inside surfaces of these tubes are Lambertian (to a reasonable approximation), it will still be the case that 'direct' solar radiation may reach the windows after being scattered on the tube wall (figure 7.3), but it has been assumed, to date, that the amount of radiation that might enter the system as a result of being scattered on the tube wall and then scattered again on or in the window would be negligible. It was suspected that this might not be the case, however, as the difference in radiance between the diffuse field (of interest here) and the direct beam is many orders of magnitude, with the result that only a very small amount of radiation might be needed to produce a noticeable error. Between 1006

© Crown copyright 2017 Met Office

and 1138, the solar zenith angle reduces from around 37° to 29° with the result that the flux entering the mouths of the tubes increases by about 10 %, lowest directly illuminated point in the tubes descends 10 mm, and the maximum plane angle subtended by the window at the lowest illuminated point increases from 30° to 37° for the NIR and 52° to 76° for the visible. This possible effect remains to be investigated and quantified, but crude back-of-the-envelope estimates suggest that the difference seen at 113822 could be equivalent to less than 0.001 % of the measured radiation having come from the direct beam.

Zenith sky data from a side-by-side comparison of a Cimel sun photometer and SWS show very good agreement, although there may be some evidence of the direct solar beam being scattered into the instrument.

8. Stray Light in SWS.

Further evidence that SWS might be susceptible to stray light was obtained at Praia airport, Cape Verde, during the ICE-D experiment. The aircraft performed three pirouettes at the end of flight B934, one of them while SWS was pointed to the zenith sky. Figure 8.1 shows that an average of roughly 400-500 nm, being the peak of the spectrum, which has been adjusted to account for the setting of the Sun, by assuming that the drop in radiance was linear over the period of each individual pirouette, and linearly increased the radiances through the pirouettes. There does appear to be some evidence of stray light. The effects of roll and pitch, due to Praia's heavily cambered taxiway, can be seen in the radiances, but the effects do not seem to be large. The roll does not appear to correlate with features in the radiance trace, and, as the maxima of pitch line up with both the maximum and minimum radiances, this is unlikely to be a major contributor to the changes in radiance. The almucantar pirouettes (not shown) have, of course, the maximum radiance occur with the aircraft is pointing directly into, and away from, the Sun, so, if the telescopes were significantly off zenith, given that the only way they could be canted is fore or aft, the peak radiance might be expected to occur with the aircraft, again, pointing into or out of Sun, but it actually occurs about 30° aft of the starboard beam, when the head is in direct sunlight, with the minimum 180° from that, when it is in shadow.

Accurate calculation of the amount of 'direct' radiation which may enter the instrument by this route is not trivial but the problem has been assessed experimentally. On 28th January 2016, SWS was set up outside at Exeter with obscurers fitted over the designed/supposed FoV (figure 8.2). The obscurers were constructed from cardboard tubes, crimped at one end and blackened internally, to ensure that they were not, themselves, able to scatter significant radiation into the FoV. A plastic collar was produced that fitted around the SWS head, and the obscurers were mounted from this

© Crown copyright 2017 Met Office

with bull wire such that the head could be rotated in both axes while the obscurers maintained their position relative to the entrance tubes. That the obscurers, so fitted, were blocking the designed FoV was confirmed in the darkroom, using a torch. SWS was then set to various angles relative to the direct solar beam while the radiance was being recorded. Due to a problem with the NIR module at the time, data were collected for both telescopes using the visible module. As the interest here is in relative signal rather than absolute, this was deemed sufficient.

Two plots are presented. Figure 8.3 is a time series showing averaged raw counts on the abscissa while time progresses down the ordinates. The ordinates are labelled with the corresponding notes from the SWS instrument log book. Normal dark measurements are shown in red. At various times four thicknesses of black adhesive tape was stuck over the entrance tubes, as a 'sanity check'; these should match the dark currents, in terms of counts, and do.

Some clarification of some of the event labels may be helpful:

- *dead down-sun* - head pointing exactly away from sun so that both the entrance tubes and the obscurer itself are shadowed by the head.
- *some fiddling with obscurer* - making adjustments, not relevant to this report.
- *bodge tape* - black adhesive tape covering entrance tubes.
- *~90° to sun - face and mouths shadowed* - axis of instrument at 90° to direct beam; neither the face of the head through which the entrance tubes pass nor the mouth of the obscurer are in the direct beam.
- *tubes off, wire still there* - obscurer removed, but bull wire (weathered galvanising) frame remains in place.
- *tube walls in sun, windows shadowed* - the walls of the entrance tubes are illuminated by the direct beam, but this dos not impinge on the windows at the bottom of the tubes.
- *as above, but no tubes* - obscurer removed.
- *some sun on windows* - the direct beam now strikes the edge of the window.
- *dead into sun, none on windows* - head directed directly at the Sun, but the entrance tubes are shadowed by the obscurer.
- *as above but shadowing window blank with notebook* - a notebook was used to cast shadow on the otherwise directly lit, white, window blank, which might well scatter towards the entrance tubes.
- *~90° - face of head shadowed, some sun on obs mouth* - due to the Sun's precession, the head had to be re-orientated. As the mouth of the obscurer was not precisely parallel with the face of the head, in this reversed orientation, it was not possible to achieve shadow on both the obscurer mouth and the face of the head.

© Crown copyright 2017 Met Office

- *dead into sun, obscurer tubes in place* - no direct sun on entrance tubes.
- *dead into sun, obscurer tubes in off* - full solar beam on entrance tubes and windows. The detectors were saturated and the data are therefore beyond the range of this plot.
- *approx Z with tubes* - head is pointed to approximate zenith with the obscurer in place.
- *approx Z without tubes* - as above, but without obscurer. This is similar to the situation that existed during comparison with the Cimel, mentioned above.

Figure 8.4 shows, in six axes, raw spectra (dark-corrected) recorded during this exercise. Spectra may be labelled 'visible' or 'NIR' although all are clearly in the range covered by the visible module; the labels, in this case, refer to the input telescope through which radiation is being received, rather than the recording module or spectral range. It should be borne in mind that there is variability in the measurements, in addition to that which was being investigated, due to there being some broken cloud present.

It is immediately obvious that no matter where the head is pointed, something resembling sunlight can be seen at levels significantly above the dark level, despite the obscurers. Top right suggests that at 90° to the sun, 10-15 % of the signal might be stray light; down-sun, things are less significant, but the stray light level is still several times the dark current. The visible and NIR show similar behaviour, but not identical; this would not be expected as the telescopes are different and the NIR fibre is optimised for that region of the spectrum, rather than the visible region which is actually being recorded, and presented here. The closer the head is pointed to the Sun, the worse the situation becomes, and once the direct beam is incident upon the windows the situation becomes substantially worse, again.

The legends are all in chronological order, although some of the axes overlap in time, and labels link to figure 8.3. The middle, right plot shows four spectra with the head pointing dead into the solar beam, which it might be hoped would show something very much like a dark current, as the designed FoV of the telescope should be entirely occupied by the inside of the obscurer, itself in shadow; very clearly, however, it does not. The three spectra labelled 'dead into sun' increase because the sky was clearing during that time, but the radiance can be seen to drop for *'book shading window blank'*, when the otherwise-illuminated window blank is put into shadow.

Some additional data were recorded in the darkroom. SWS was set up so as to view the bulb normally used to calibrate SHIMS, and at various angles away from it, up to 90 degrees (straight up). Ideally, the signal might be expected to drop to zero when the bulb was out of the designed FoV, but this is not the case. In figure 8.5, the two curves

© Crown copyright 2017 Met Office

are two repetitions of the measurements (red, then blue). The counts shown are a mean of the brightest (approximately) eighty pixels of the visible module with an integration time of 500 ms. There is still detectable signal, well off axis but what happens beyond 60° is not yet understood. It is suspected that SWS may be picking some reflected light from the darkroom ceiling. These measurements were made without the obscurer in place and no attempt was made to prevent the bulb from illuminating the darkroom itself, and could usefully be repeated with greater rigour.

The possibility was considered, that the Cimel itself might have a heretofore unacknowledged susceptibility to radiation outside its nominal FoV, and cursory attempts have been made to assess this. Initially this was done in a similar manner to that described for SWS, using the SHIMS calibrated bulb. This was however ineffective as the Cimel appears to be much less sensitive than SWS; the recorded signal when the Cimel was directed at the bulb was so low that now signal would be expected off axis, even if there was a significant susceptibility. A second attempt was therefore made, using the integrating sphere.

The integrating sphere has a large exit port which may be considered, in this context, to be an isotropic, hemispherical source. It follows that a Cimel, designed to have a narrow FoV, with a sharp cut-off, should record the same radiance coming from the sphere, irrespective of its distance from it, providing that its FoV is entirely occupied by sphere's exit port. If the Cimel's FoV (at the sphere) is larger than the exit port, then a lower radiance will be recorded, and so if a Cimel is used to measure the sphere, from different distances, in the event of it having a strong sensitivity to light outside its designed FoV, a drop in the recorded radiance will be seen closer to the sphere than would be expected. One might expect to see a small increase in radiance when the Cimel is very close to the sphere, due to radiation being reflected back into the sphere by the Cimel. The observations of this exercise are presented in figure 8.6. Due to the constraints of the darkroom, only a small range of distances were actually measured but it is clear that only very close to the sphere is any change seen, and, then, only at longer wavelengths where the Cimel might be expected to have a greater impact on the sphere output. At the shorter wavelengths, nearer to the peak of solar output, no change is seen. Although there are very clear limitations and omissions in these measurements, and further work is required, it would seem to be clear that SWS has a significant susceptibility to stray light, but that the Cimel does not. These data clearly contradict the results of the courtyard Cimel comparison, and how it might be that SWS could appear to agree very well with such a well respected instrument if it were as susceptible to stray light as it appears to be, is a question which remains to be answered.

This work would suggest, however, that stray light could easily cause errors of 50 %, but it would seem reasonable to assume that the errors caused by stray light might vary greatly, depending on the atmospheric and instrumental conditions. It is clear that the

© Crown copyright 2017 Met Office

magnitude of the problem is dependent upon the angle between the solar beam and the principal axis of SWS, but it might also be expected to depend upon aerosol loading. In conditions of heavy aerosol, the strength of the balance of direct to diffuse radiance will be shifted towards the latter, and so the stray light should be less significant, but dust on the SWS entrance tubes head will increase the amount of light scattered, and dust will exacerbate flare on the entrance windows.

SWS measurements made outside and in the darkroom show very strong evidence of a severe stray light problem in SWS, although no evidence of similar issues has been found with the Cimel.

9. Seating of SHIMS Domes Within Transfer Standard.

Concern has been expressed in the past over the care with which the ground crew carry out transfer calibrations of upper SHIMS, but possibly of greater concern is how lower SHIMS transfer calibration measurements are made. The orientation of the transfer standard relative to the domes is fixed by two protruding bolt heads which limit the possibilities to either the correct orientation, or 180° from that; any other orientation will result in an axial displacement of some 3.5 mm - this being the thickness of the bolt heads. From this, correct location of the standard on upper SHIMS is fairly straightforward, but a fundamental and hitherto unreported issue, in the past, has stemmed from the method by which the transfer standard has been held against the lower SHIMS dome, and the fact that the aircraft, having pneumatic tyres and suspended wheel struts, can move relative to the ground or hangar floor during a measurement. The long standing procedure has been to support the transfer standard from below, with a laboratory jack. Although this should be adequate if the aircraft does not move, it may well do so due to the addition or removal of personnel and/or equipment, and, in one extreme case, to the addition of fuel, which lead to the destruction of the lower SHIMS dome. It was noted, following the post-ICE-D transfer measurements that there was a gap of some 5 mm between the top of the transfer standard and the bottom of the flange around the dome. This was assumed to be inconsequential, but subsequent measurements suggest that this is not the case.

Figure 9.1 shows spectra recorded with a 30 ms integration time, with a SHIMS dome seated both correctly and incorrectly. It is relatively unlikely that any measurements would have been made with any orientation other than the correct one, or 180° from it, but measurements were made at ±90° as well. That the greatest signal is seen with the dome seated correctly is, of course, no surprise, but that the signal might fall by ~25 % if the dome is displaced by only 3.5 mm is a matter of some concern, as it is entirely possible that this sort of displacement might have affected lower SHIMS transfer calibrations on the aircraft. Transfer measurements have been made for many years but

© Crown copyright 2017 Met Office

it is thought that the ICE-D is the first time that they have actually been applied to data and, although the effect described here will have rendered the lower SHIMS transfer calibration worthless, the SWS and upper SHIMS ones should be valid. The intention was to replace the laboratory jack used to support the transfer standard with an air jack which should both hold the lamp securely against the lower SHIMS dome, and allow the aircraft to move, but initial experiences have not been positive, so the laboratory jack remains in use and a solution remains to be devised.

Figure 9.2 shows the spectra from these same tests, with each of the six arrangements of transfer standard and SHIMS dome in a separate pair of axes. In this case, however, spectra are presented for integration times of 10 - 500 ms, ordered chromatically, with 10 ms being violent and 500 ms being red. The ordinate, here is now dark-corrected counts divided by integration time, and so, if the instrument was entirely stable and linear, all the spectra in each axes would over overlie each other; clearly this is far from the case, and so calls into question the linearity of the instrument. This is also illustrated in figure 9.3 which shows mean dark-corrected counts/ms of the brightest ~120 pixels in each detector as a function of integration time, but expressed as percentage difference from the 10 ms figure. Again, these curves should overlie each other, as identical straight lines; that they do not may be taken as further indication of the importance of calibrating the instrument at any and every integration time used for measurement.

Although it is possible that large errors could have occurred in the past, if sufficient care is taken, these should be entirely avoidable in the future.

Failure to have SHIMS correctly seated in the transfer standard can cause large errors. Unseating of lower SHIMS due to equipment or personnel movement in the aircraft could well cause errors exceeding 25 %.

10. SHIMS on ICE-D, and Transfer Calibrations.

Although the importance of transfer calibrations has been mentioned in the past, and transfer measurements have been made, routinely, for many years, it is believed that the ICE-D experiment is the first occasion on which these have actually been applied to scientific data. The transfer curves used for the ICE-D data are plotted in figure 10.1. Three sets of curves are shown, with blue being from upper SHIMS, red from SWS and green from lower SHIMS. The curves themselves are the ratio of the laboratory measurement to the measurement taken on the aircraft, and so should be the factor by which the flight data require to be multiplied in order to negate any changes relating to the instrument being fitted to the aircraft. In each set, curves are presented for integration times of 5, 10, 20, 30, 40, and 50 ms, with the following exceptions: SWS also has curves for 75 ms but lacks a 5 ms NIR curve; lower SHIMS also lacks a 5 ms

© Crown copyright 2017 Met Office

NIR curve, but has curves for 500 ms. It is immediately obvious that the sensitivity of lower SHIMS appears have dropped substantially between being measured in the laboratory and on the aircraft. Although damage to a fibre was previously suspected, it is now thought that this apparent drop is, at least in part, the result of lower SHIMS not being fully seated in the transfer standard when the on-aircraft measurements were being made, due to movement of equipment and personnel.

Figure 10.2-10.5 (Clare Ryder, Reading University) show comparisons between SHIMS data, recorded on 12th August 2015 (flight B923) at an altitude of 6 km, the broadband radiometers (BBR) and the Edwards-Slingo radiative transfer model. The model used a McClatchey tropical standard atmosphere (without aerosol) and so, although reasonable simulation of downwelling radiation might be expected, the upwelling simulations will be less reliable as the standard atmosphere will not match the actual one on the day. In the preparation of the data, the same attitude corrections were applied to the SHIMS data as to the BBR data, and the same cosine response was assumed, and corrected for. Although the assumption of a common attitude should be sound, it is unlikely that the cosine response assumed will be correct for SHIMS, and, although the impact of this is not known, it should be borne in mind as an additional source of uncertainty. The comparison between the BBRs and SHIMS is complicated by the fact that, not only are the BBRs not spectrally resolved, but the wavelength ranges of the clear- and red-doomed BBRs do not match those of the visible and NIR SHIMS modules. The SHIMS data presented here are, therefore, integrated across the relevant spectral window to match the BBRs ranges. The BBRs ranges are 300-3000 nm of the clear dome and 700-3000 nm for the red dome and so 'red' and 'visible', here, correspond to the red BBR, and the clear-minus-red. AS SHIMS does not measure beyond 1710 nm, a small increment (calculated using the model) has been added between 1710 and 3000 nm; this is however, the lowest radiance part of the window of interest, and the numbers involved are small.

Figure 10.2 shows a comparison of the downwelling visible (300-700 nm). The model is shown in black diamonds, BBR data are shown in red with the pink band indicating ±10 %, and upper SHIMS visible is shown in blue, becoming cyan when the attitude correction is applied. Although there are clearly anomalies at various times, and there is an obvious difference in how SHIMS and the BBRs vary with time, the agreement is generally within 10 %.

With the downwelling red (700-3000 nm), however, the situation is not quite so rosy (figure 10.3). The model and BBR data are shown, as before, but the SHIMS data (blue), here, are the NIR + the relevant portion of visible + the model increment covering 1710-3000 nm; cyan shows these after the attitude correction. In addition the upper SHIMS NIR only, and upper SHIMS NIR + part of visible are shown in green. Again,

© Crown copyright 2017 Met Office

there are some obvious anomalies, the cause of which is not, currently, known, but apart from these the agreement with the BBR and model is, again, surprisingly reasonable. It is interesting to note that the shape of the SHIMS time series matches that of the model slightly better than the BBR, although the BBR is closer in terms of irradiance.

Comparisons of upwelling radiation are more difficult, as previously noted. Figures 10.4 and 10.5, respectively, show the corresponding upwelling visible and red comparisons (colours as for downwelling). Due to the limitations of the atmospheric profile used in the forward model, little can be said about these data. Although features appear in both the BBR and SHIMS data at the same times and appear to match, approximately, in proportion and sense, SHIMS does appear to be reporting almost twice the visible irradiance of the BBRs, and almost four times the red. It should be noted, however, that there are suspicions over the validity of the lower SHIMS transfer correction in this case, because the displacement of the aircraft, as mentioned above. That the transfer ratio for lower SHIMS should be more than twice that for upper SHIMS, despite their physical and optical similarities must also raise questions about whether the discrepancy between the BBRs and lower SHIMS is actually more the result of a poor transfer correction than of a poorly performing instrument, in itself. It is, however, entirely possible that the performance of lower SHIMS, itself, is substantially better than would seem to be the case from these data.

Figure 10.6 shows further, more basic, comparisons between SHIMS and the BBRs. Here, the BBR data are simply the clear dome minus red dome, plotted with the appropriate part of the SHIMS visible spectra. B928 compares similarly to B923 and, indeed, the rest of ICE-D (not presented), with agreement to around 10-20 % for upper SHIMS, but with lower SHIMS reporting roughly twice the irradiance indicated by the BBRs. B776 is also presented, although not an ICE-D flight. Here, by contrast, lower SHIMS appears to agree to 10-15 % through most of the flight while upper SHIMS shows the larger differences. That a similar pattern of agreement exists on all ICE-D flights and but that others (using different calibration sessions) show quite different results may suggest that the major factor in determining the accuracy of the data is the quality of the calibration, rather than any inherent weakness in the instrument. The reason why the sense of the difference between the BBRs and upper SHIMS is not constant, however, is not known.

Through much of ICE-D, upper SHIMS appears to agree with the BBRs to around 10 %, but lower SHIMS reports close to twice the BBRs. This may, in part, be explained by a poor lower SHIMS transfer calibration. B776 show better agreement in upwelling than downwelling, suggesting calibration rather than instrument problems.

© Crown copyright 2017 Met Office

11. The Effect of Transport.

The occasion of the NERC 50th anniversary celebrations at Cranfield afforded the opportunity to investigate whether the very act of moving SWS-SHIMS between the laboratory and the aircraft has any noticeable effect on the instrument. SWS-SHIMS was packed up, driven to Cranfield, unloaded, unpacked & set up as a static display item, packed, driven back, and once more set up in the darkroom. At either end of this exercise, measurements were made of the transfer standard. Figure 11.1 shows, for each of the six modules, the difference between the pre- and post-50th transfer measurements as a percentage of their mean. Although the use of the transfer standard should remove any transport effects, it is is important to have some confidence that the instrument does not change significantly as a result of vibration. It is immediately obvious that there is a large change in the upper SHIMS visible module, but this is suspected to be the result of extraneous material having been trapped when the upper SHIMS fibre was being connected to the visible module; the need for more attention to be paid to this possibility has been noted. Amongst the other five modules, the differences are less significant, although not entirely unconcerning. Work remains to quantify the likely differences resulting simply from the process of connecting and disconnecting the optical fibres.

No definitive conclusions can be drawn from this, but it would not seem unreasonable to expect a few percent change due to the transporting of the instrument. This should, however, be accounted for when the transfer calibration is applied.

Changes may occur in transit, but should be small, and should disappear upon calibration.

12. Changes in Module Temperature.

Measurements were made in the darkroom of both dark and scene views of the integrating sphere while the temperature within the cooler, which contains the six spectrometer modules, was varied between 283 K and 293 K. This was carried out on two occasions: firstly as the temperature was reduced by the usual refrigeration system, from 293 K to 287 K, and on the second, as the temperature rose from 283 K to 290 K, when the refrigeration was turned off. Representative examples of the data recorded are shown in figure 12.1, as time series of counts, averaged across wavelength; data are only used from pixels which remain unsaturated throughout, and so the plots of data from the visible module do not contain data from pixels 55-149; all NIR data were usable. There are sixteen plots, presented in four groups: visible modules, top eight, NIR modules, bottom eight, and data recorded at 30 ms, left, 50 ms, right. In each of the four groups, data from the first measurements (rising temperature) are shown in the left

© Crown copyright 2017 Met Office

two plots, with data from the second measurements (falling temperature) on the right two. The upper two plots are of scene measurements, while the lower two show dark measurements. In all plots a percentage is quoted - this is the difference between the maximum and minimum values presented in that plot, as a percentage of the mean scene counts. All abscissae are labelled simply with spectrum number. Although data are plotted in chronological order, that chronology is discontinuous due to the the omission of spectra recorded at different integration times, giving rise to the 'steps', particularly evident in the visible, scene plots. It is important to bear in mind that as dark measurements and scene measurements cannot be made simultaneously, although scene and dark data are interleaved, data points in the scene plot may not line up precisely with the corresponding data in the dark plot. It is clear that there is a variation of output with temperature, but it is also clear that this does not exceed 1 %, even for a temperature change of 7 K. It follows, therefore, that variations of ±1 K (much larger than usually seen in the cooler) should be inconsequential.

The effects of temperature in the cooler varying by ±1 K are negligible.

13. The effect of Bending Optical Fibres

Optical fibres are used to connect all fore optics to all modules. In the case of SWS, the ability of the fibres to bend is used to cope with the changing orientation of the head, relative to the rest of the instrument and, hence, the bending changes almost constantly, in use. In the case of both SHIMS, once the instrument is fitted, the optical fibres should not again be bent, but this cannot be ruled out, as other work on the aircraft may result in this. It is well known that the transmission of optical fibres changes when they are bent or otherwise deformed but the extent to which the bending experienced by SWS-SHIMS fibres may affect instrument performance has not previously been assessed. The fibres currently in use are incoherent bundles, but single-core fibres have also been purchased in the past, although we believe that they have never been used. Figure 13.1 shows time series of the means of the brightest ~100 pixels from each detector (blue for visible, red for NIR) as a loop in the fibre is released, whilst viewing the transfer standard, with the single core being shown in the upper axes, and the multi core in the lower. That the single core should show greater changes when it's bent, than does the bundle, comes as no great surprise, but that the signal in one module should rise, is unexpected, and requires further investigation. Again, that the sense of the change in both modules should be the opposite in the two types is also puzzling. Although this effect clearly requires more rigorous investigation, this cursory check suggests that gross errors are unlikely to arise from this source, providing that the single core-fibres are not used, and that the fibres are not damaged or bent excessively.

© Crown copyright 2017 Met Office

This latter point is a concern in the case of SHIMS, in two areas. Although the SHIMS fibres are believed to be secured for most of the distance from the domes to the rack, there are loose lengths at both ends. There is free fibre which must be connected to the dome before it is inserted into its mount; the excess fibre lies between the pressure hull and the skin of the aircraft, but how the fibres are stowed, and what bends they are subject to, are not known. The second point of concern is that the inboard ends of the fibres pass through the trunking along the edge of the cabin sole, with other instrument wiring. How much of the fibres are stowed there, and how they are stowed is, again, unknown. It seems likely, however, that the fibres will be disturbed when adjacent racks are being fitted or removed, and previous experience with ARIES' optical fibres suggests that there is a high likelihood of fibres being treated with less respect than their due; this is a cause of concern both because of the risk of excessively tight bends, and because significant interference with a fibre could invalidate the calibration. Further investigation is clearly warranted.

Changes in SWS signal of ~1 % might be expected due to fibres bending as telescope rotates. SHIMS signal will not change in flight but variations may occur if fibres are disturbed by other work. These should be small, but this cannot be guaranteed, and changes may not be removed by transfer calibration.

14. SHIMS on B886

Due to dark measurements being missing from the B886 dataset, it is not possible to calibrate the SHIMS irradiances fully, but qualitative comparisons may still be made. Figure 14.1 shows an overlay of upper SHIMS and BBR data, and curious variations are seen in all instruments. The plot covers the short, straight and level 'approach run' (approximately cross-sun, Sun to port), and the four legs of the box pattern (up-sun, port, down-sun, starboard) - although not explicitly indicated, the turns are clearly visible. It might be hoped that measurements of the hemispherical irradiance, in clear sky conditions, should be the similar on all four legs of a box pattern, as this will be dominated by the unobstructed, direct, solar beam, rendering small changes in pitch and roll negligible, and that there should be no change along individual legs, but clearly this is not the case. Figure 14.1 shows a number of interesting features. Note that SHIMS data are not calibrated and so whilst relative changes may be noteworthy, the absolute values will be incorrect.

As SHIMS and both BBRs are mounted together at the same, fixed orientation, it is unexpected that the differences between the three instruments should differ from one leg to another, or along individual legs, as the three domes cannot move relative to each other. Oscillations during runs are seen, on occasions, in all instruments. Figures 14.2 and 14.3 show irradiance and attitude time series for the flight, showing some correlation

© Crown copyright 2017 Met Office

between aircraft orientation and reported irradiance. Although some correlation may be seen between irradiance and roll and heading, as the BBRs (and SHIMS) behave differently at different times, it is unclear to how great an extent the connection is causal. On the second leg of the box (~800 min), for example, SHIMS shows substantial oscillations, while the BBRs' oscillations are much smaller, although all are in phase with each other and approximately in phase with aircraft roll, but on the following (down-sun) run, the red dome shows no oscillations while the clear dome BBR and SHIMS, do, although these show some discrepancy in phase.

It is well known that the leading edge of the domes can accumulate dirt, and can be abraded, but as the domes should have been cleaned before the flight, the former should not be an issue in this case. Damage or dirt on the leading edge of the dome, however, might be thought of as only likely to be a problem in clear sky conditions, when the aircraft is flying directly into the Sun but, here, it can be seen that the largest oscillations in SHIMS data occur on one cross- and the down-sun run, when the direct beam should have been clear of any affected glass. Inspection of drawings shows that the direct beam was not obstructed by any part of the aircraft during the box (figures 14.4 and 14.5). This is clearly an issue which requires further investigation, for the BBRs as well as for SHIMS.

It is normal practice to apply attitude corrections to BBR and SHIMS data. This has not been carried out in this case, but will not affect the relative comparisons being made. It is worthy of note that the correction adjusts for the orientation of the dome to the direct solar beam, it does not take any account of the tilting of the plane of the dome relative to the horizon.

In B886, unexplained variations in signal, of several percent, from upper SHIMS and clear- and red domed BBRs are seen which seem to show no definite correlation to each other, to Sun position or aircraft attitude.

15. Table of Uncertainties

The figures presented here are intended only as a rough guide and summary - they are approximations and liable to change upon further investigation. Numerical values are not readily specified for some sources of uncertainty.

Since many sources of uncertainty in SWS-SHIMS data are highly situation-dependent, producing a single figure to describe the instrument's performance under generalised conditions is a rather futile exercise but it would seem reasonable to assume that, for SHIMS, a reasonable estimate for whole-system uncertainty might be around 10 %. The problems with stray light in SWS make such an assessment (for a general case) even

© Crown copyright 2017 Met Office

more problematic. It is conceivable that in some, highly specific cases, the errors might be very small, but as the possible errors will vary with the orientation of Sun, aircraft and telescopes, and with the aerosol loading at the time of measurement, these cases will be few.

Source	SWS	SHIMS	Comments
temporal linearity	potentially large		entirely avoidable
radiometric linearity	up to ~3 % (visible), ~14 % (NIR)		worst at low (ir)radiance
local standard lights	0.8 %	1.75 %	
laboratory calibrations	7 %		repeatability
stray light	may be very large	n/a	strongly scene dependent
seating of transfer std.	0	potentially large	entirely avoidable
transport	~2 %		removed by transfer calibration
cooler temp. drift	0		
fibre bending	1 %	n/k	Insufficient information about SHIMS fibre routing.
§14 variations	0	~5 %	
post processing	0-large		poor filtering of spectra.

16. Conclusions and Further Work

It appears that the IDL software, whilst imperfect and in need of work, particularly in terms of it's rejection of bad spectra, and selection of integration times, can produce good results.

The repeatability of lab calibrations appears to have been poor in the past, although more recent calibrations appear to be more repeatable.

The temporal linearity appears to be poor, which could have a substantial impact on the calibration and field measurements, but this problem can be circumvented by restricting field measurements to integration times at which the instrument has been calibrated.

© Crown copyright 2017 Met Office

Radiometric linearity is better in visible module than NIR,and at higher radiances. For low radiance scenes errors of 3 % and 15 %, respectively, might be assumed. Linearity appears to be better at higher, non saturating, radiances. Further analysis is required to determine the extent to which this effect could be corrected for in post processing.

There appears to be a systematic difference in airborne SWS data depending on which side of the aircraft is presented to the Sun on cross-sun runs; which is believed to be cause by stray light.

Local calibration standards appear to be reasonably stable over long periods but the calibrations carried out in the laboratory may carry an uncertainty of ~7 %

Zenith sky data from a side-by-side comparison of a Cimel sun photometer and SWS show very good agreement, but measurements made outside, and in the darkroom, show very strong evidence of a severe stray light problem in SWS. No evidence of similar issues has been found with the Cimel further work is required, as these results are contradictory.

Failure to have a SHIMS dome correctly seated in the transfer standard can give rise to substantial errors. Due to manner in which transfer measurements are made on the aircraft, it is almost certain that this has happened on occasions in the past. It is critical that all engaged in the transfer calibration process are made aware of its sensitivity.

SHIMS has been observed to agree with the BBRs to ~10%. That the good agreement with upper SHIMS, and poor agreement with lower SHIMS appears reasonably constant through out ICE-D (where the lower SHIMS transfer calibration is known to have been bad) is illustrative of the criticality of this process, but also gives some confidence in the stability of the instrument.

Changes may occur as a result of the instrument being packed, and transported between the laboratory and the aircraft, but these will be small, and dealt with by the transfer calibration.

The effect of small variations in the temperature of the cooler containing the spectrometer modules is negligible. It would seem reasonable that in the event of a major redesign of the rack, significant weight and space could be saved but reducing the refrigeration equipment.

The bending to SWS optical fibres essential to the operation of the instrument may well cause variable errors of up to1 %. Neither the magnitude nor stability of SHIMS fibre related errors are known. Although the transmission of the fibres should remain constant once the instrument is fitted, there is concern that other work on the aircraft may result in

© Crown copyright 2017 Met Office

bending of the fibres, with a consequential change in transmission which is not captured in the transfer calibration. Although it is expected that this effect might be small, there are no data to confirm this, and further investigation would be prudent.

B886 permitted downwelling measurements to be made in near-ideal, clear sky conditions, but variations were seen in upper SHIMS, clear- and red dome BBR data. Variations in one instrument occasionally correlate with variations in the aircraft attitude, or the output of another instrument but there is no consistent pattern. Further investigation is required to resolve this.

It is proposed that the use of SWS-SHIMS should be limited to 10 ms integrations, only, and that the control software be altered to enforce this. This will remove the possibility that a required integration time be missed at one of the calibration stages, as has happened at various points in the past. In conjunction with this, it is proposed that the control software be altered to include the periodic recording of dark currents, automatically, but with the option given to the operator to suppress this behaviour for a short period, should this be necessary. As individual 10 ms spectra will often have very low signal to noise, it is further proposed that the real-time display be altered to display averages of the data from the preceding 1 second, updating at 1 Hz.

It is proposed that, pending a redesign of the SWS fore optics, to mitigate the stray light issue identified, the SWS head should be withdrawn from the standard aircraft installation of SWS-SHIMS, and its installation be made subject to a specific request, as useful calibration of SWS data, recorded with the current fore optics is not usually possible.

Although the work reported herein represents a considerable step forward in the understanding of the SWS-SHIMS instrument, many of the investigations, as noted above, are rudimentary and it is recommended that consideration be given to repeating them prior to any major work on the instrument, or decisions relating to its future.

© Crown copyright 2017 Met Office

17. Figures

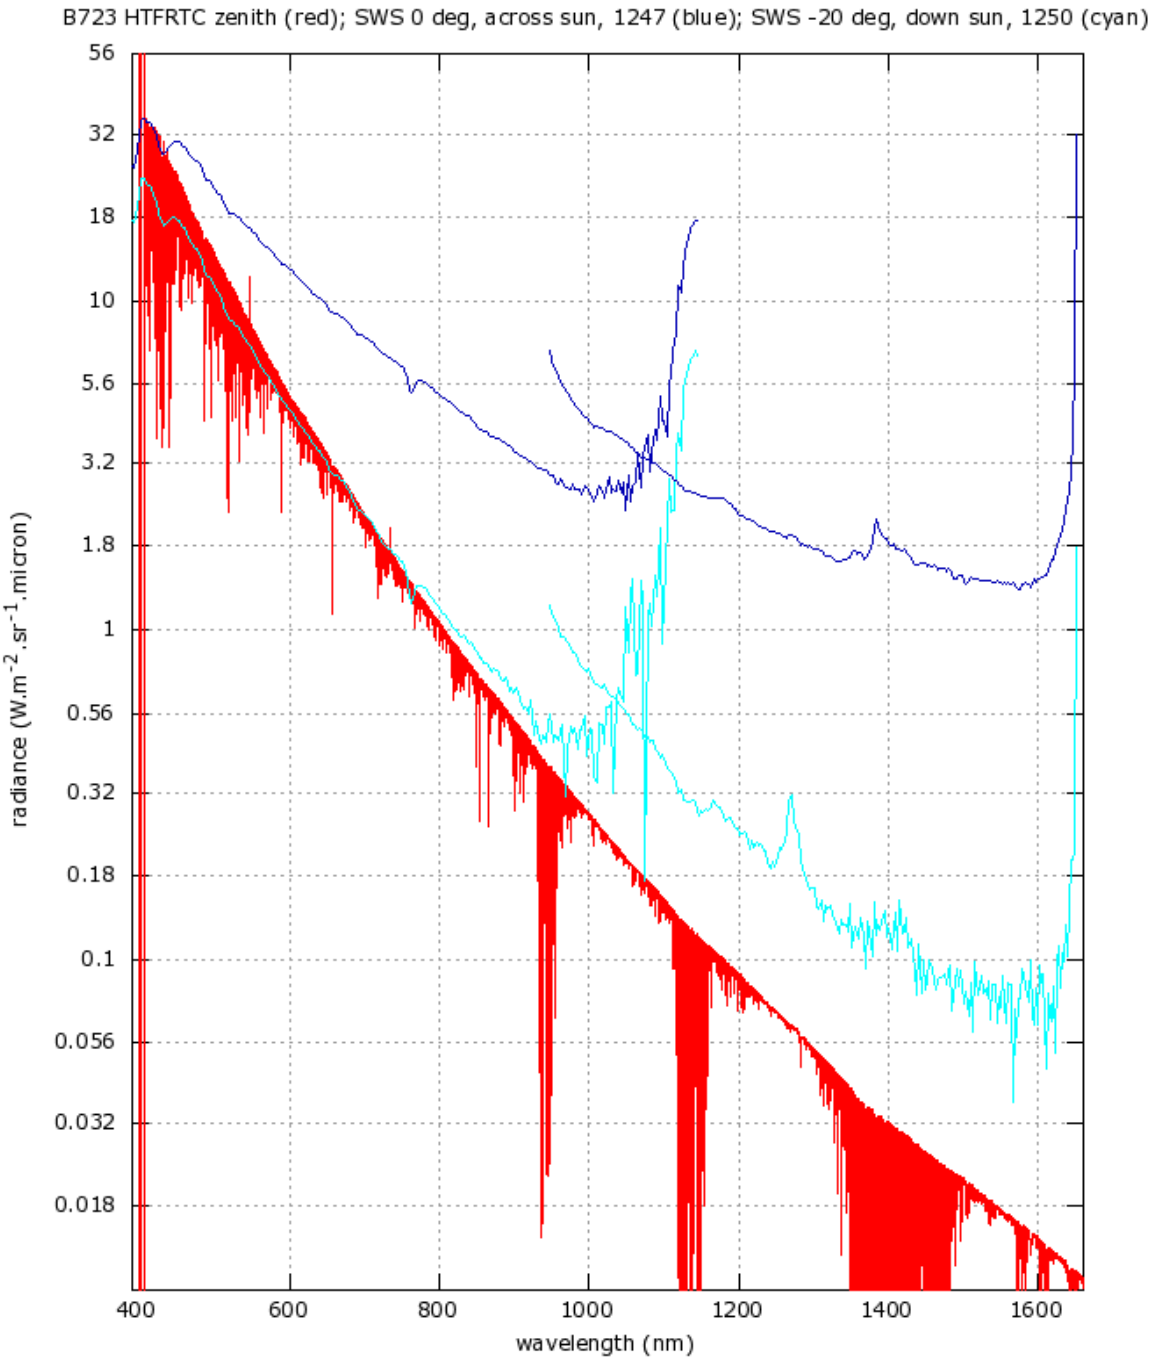

Figure 2.1. SWS spectra from off-zenith down Sun, and zenith cross Sun runs, with high resolution HTFRTC simulation of latter

© Crown copyright 2017 Met Office

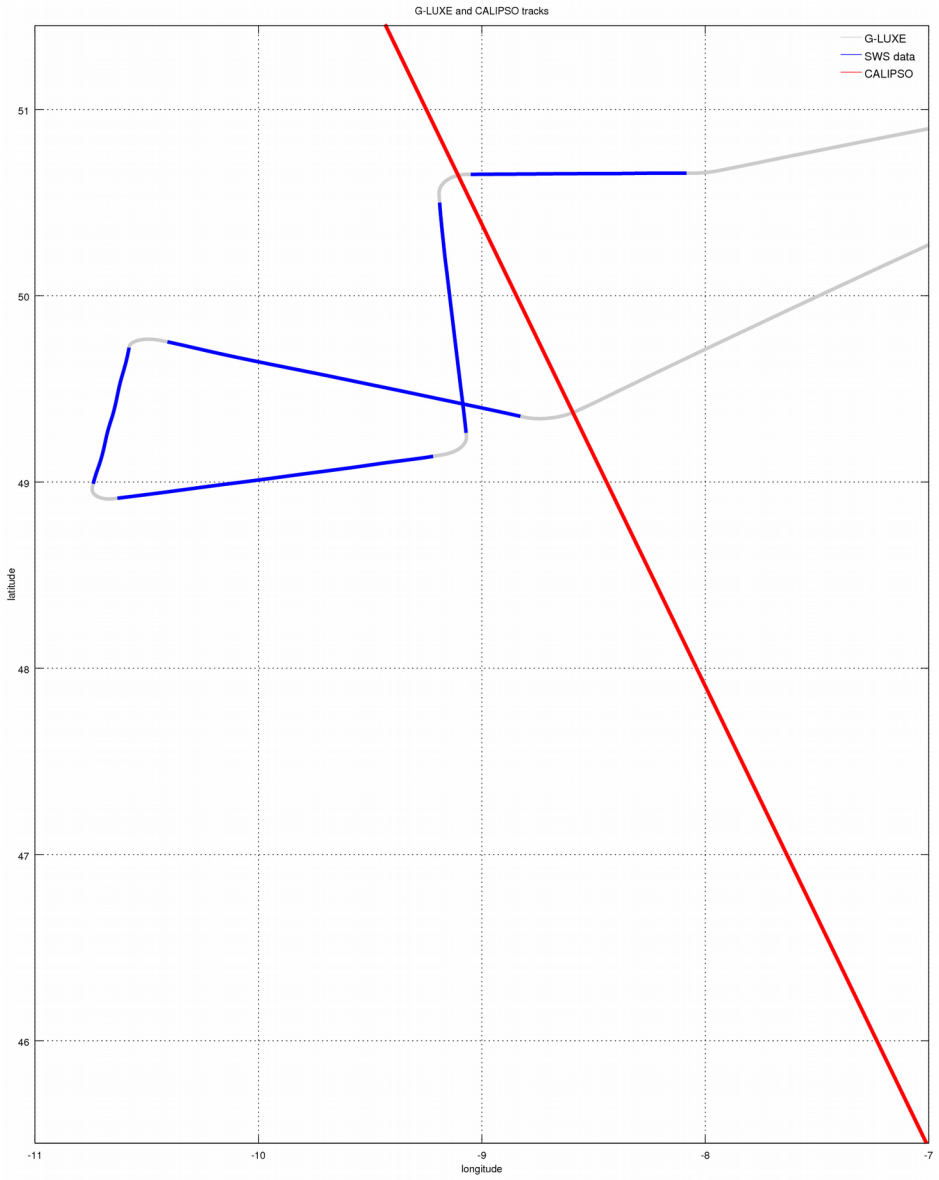

Figure 2.2. B886 track plot with location of SWS measurements in dark blue and CALIOP ground track in red.

© Crown copyright 2017 Met Office

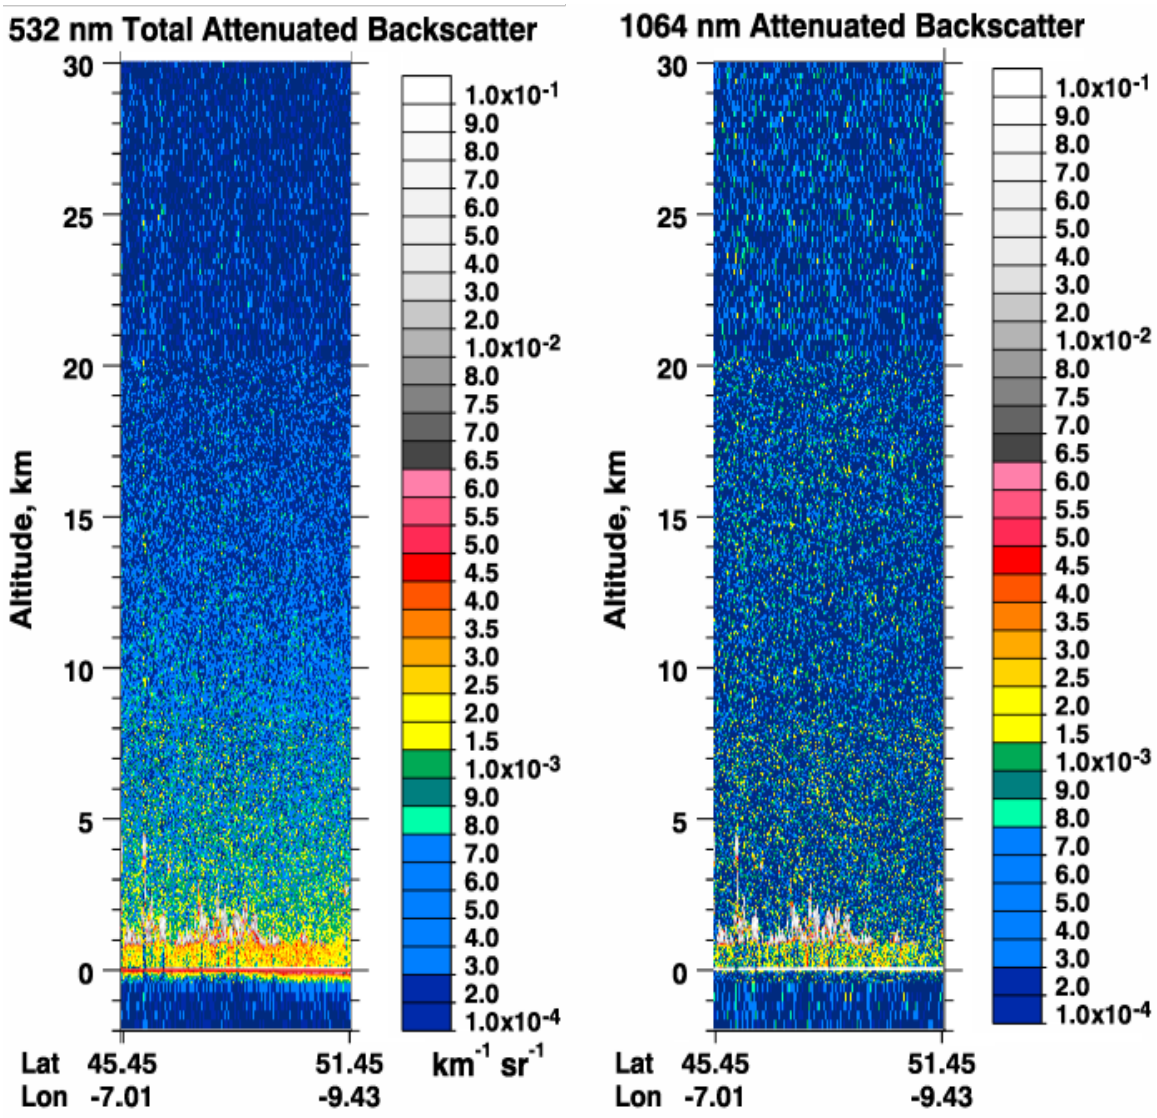

Figure 2.3. CALIOP attenuated backscatter at 532 adn 1064 nm. Horizontal span of data presented here matches extent of ground track shown in fig. 2.2.

© Crown copyright 2017 Met Office

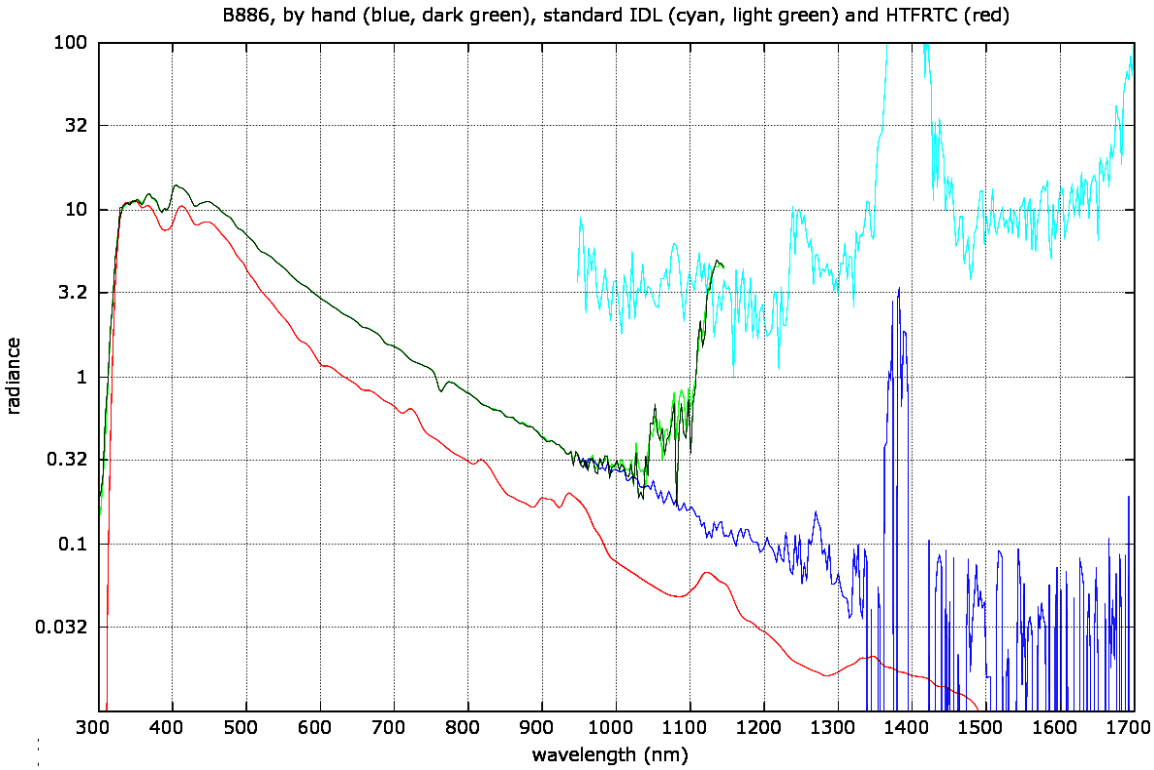

Figure 2.4. B886 calibrated using the standard code (cyan, light green) and manually (blue, dark green), with an HTFRTC simulation in red.

© Crown copyright 2017 Met Office

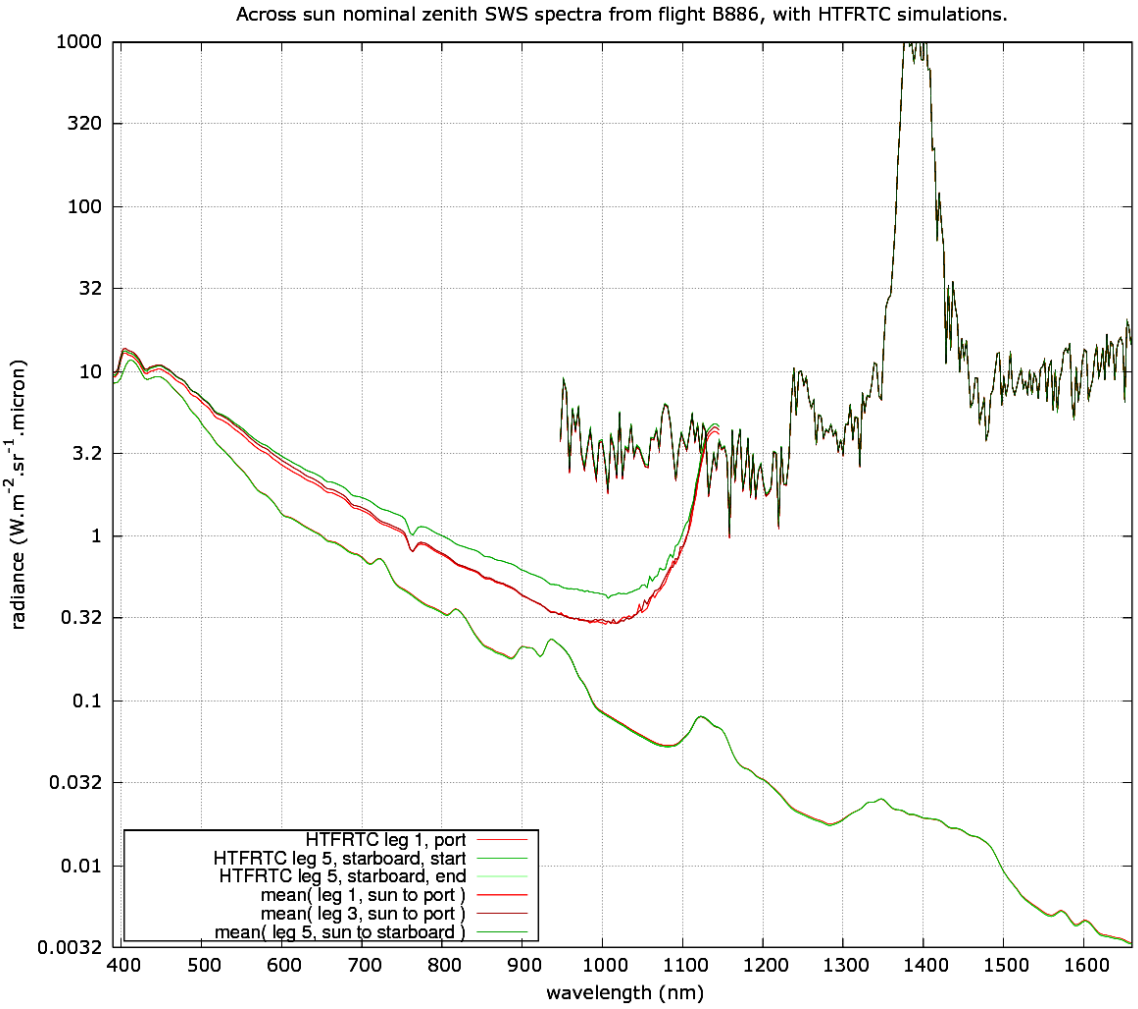

Figure 2.5. Port/starboard bias in SWS cross_Sun spectra from B886. Data are coloured according to the side of the aircraft presented to the Sun.

© Crown copyright 2017 Met Office

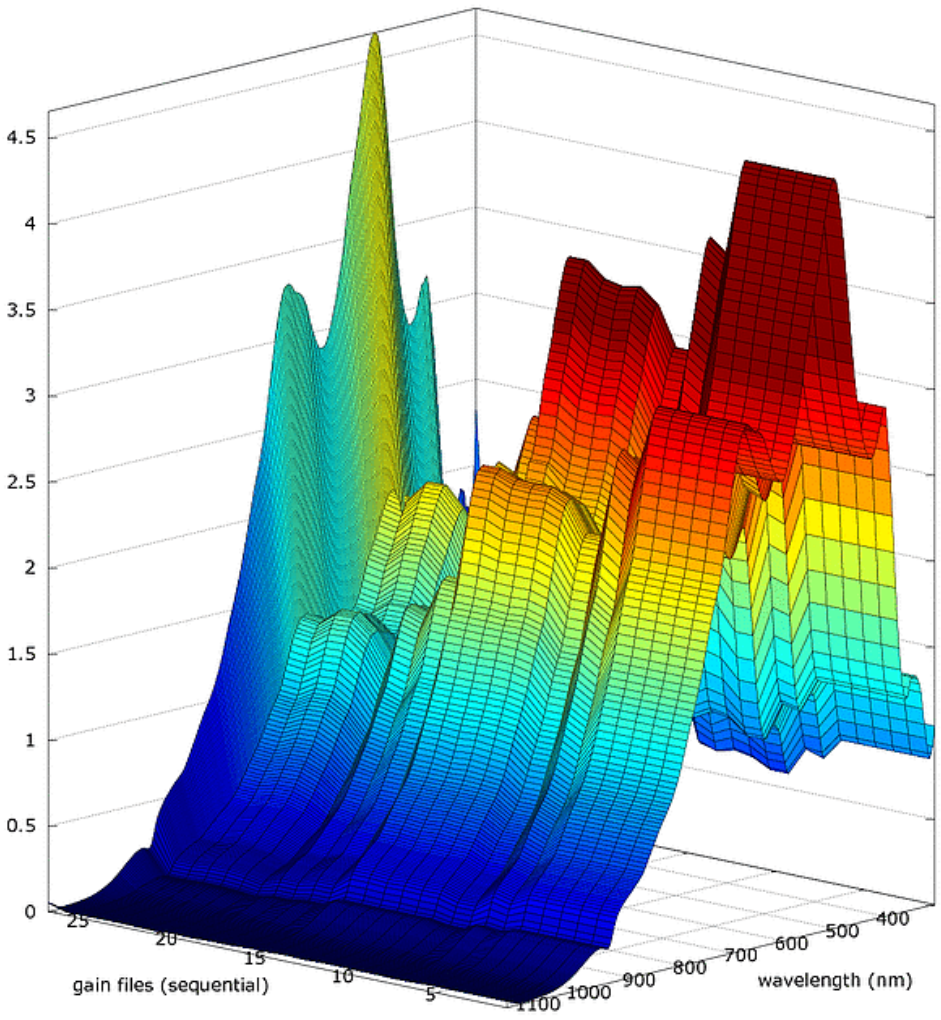

Figure 3.1. Sensitivities from SWS visible module.

© Crown copyright 2017 Met Office

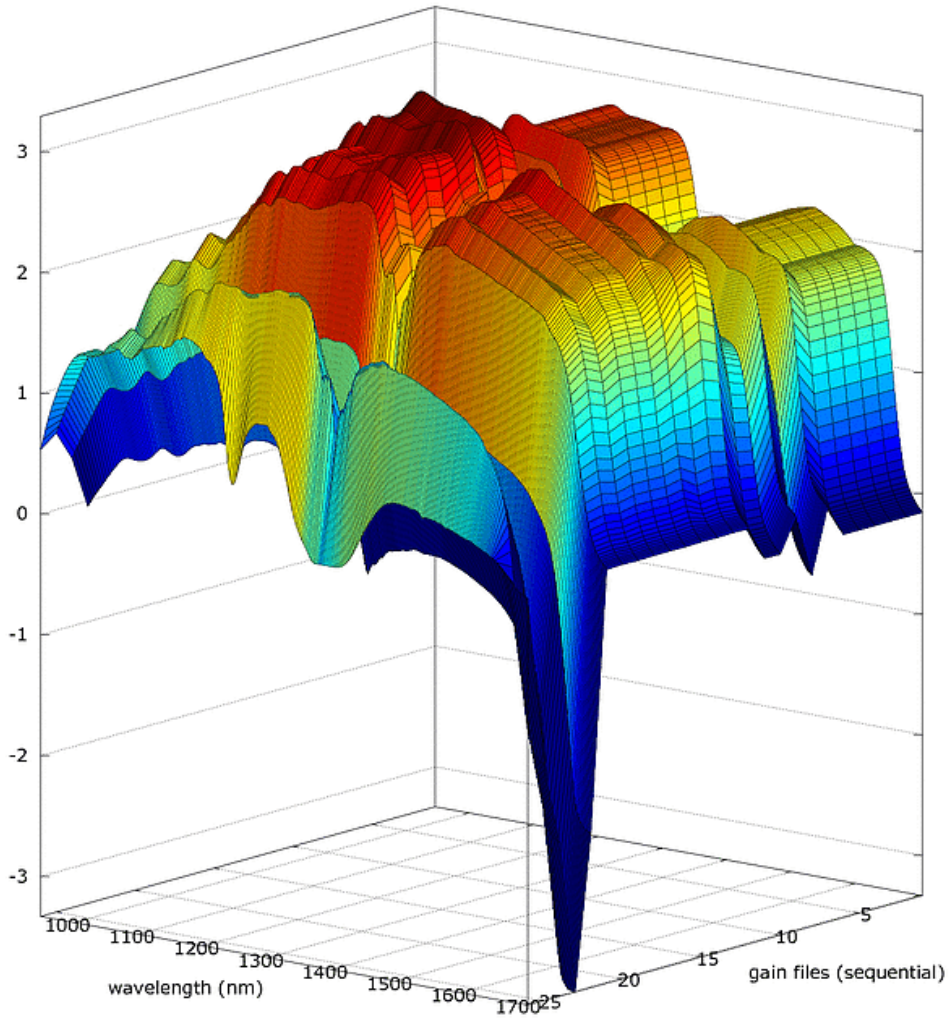

Figure 3.2. Sensitivities from SWS NIR module.

© Crown copyright 2017 Met Office

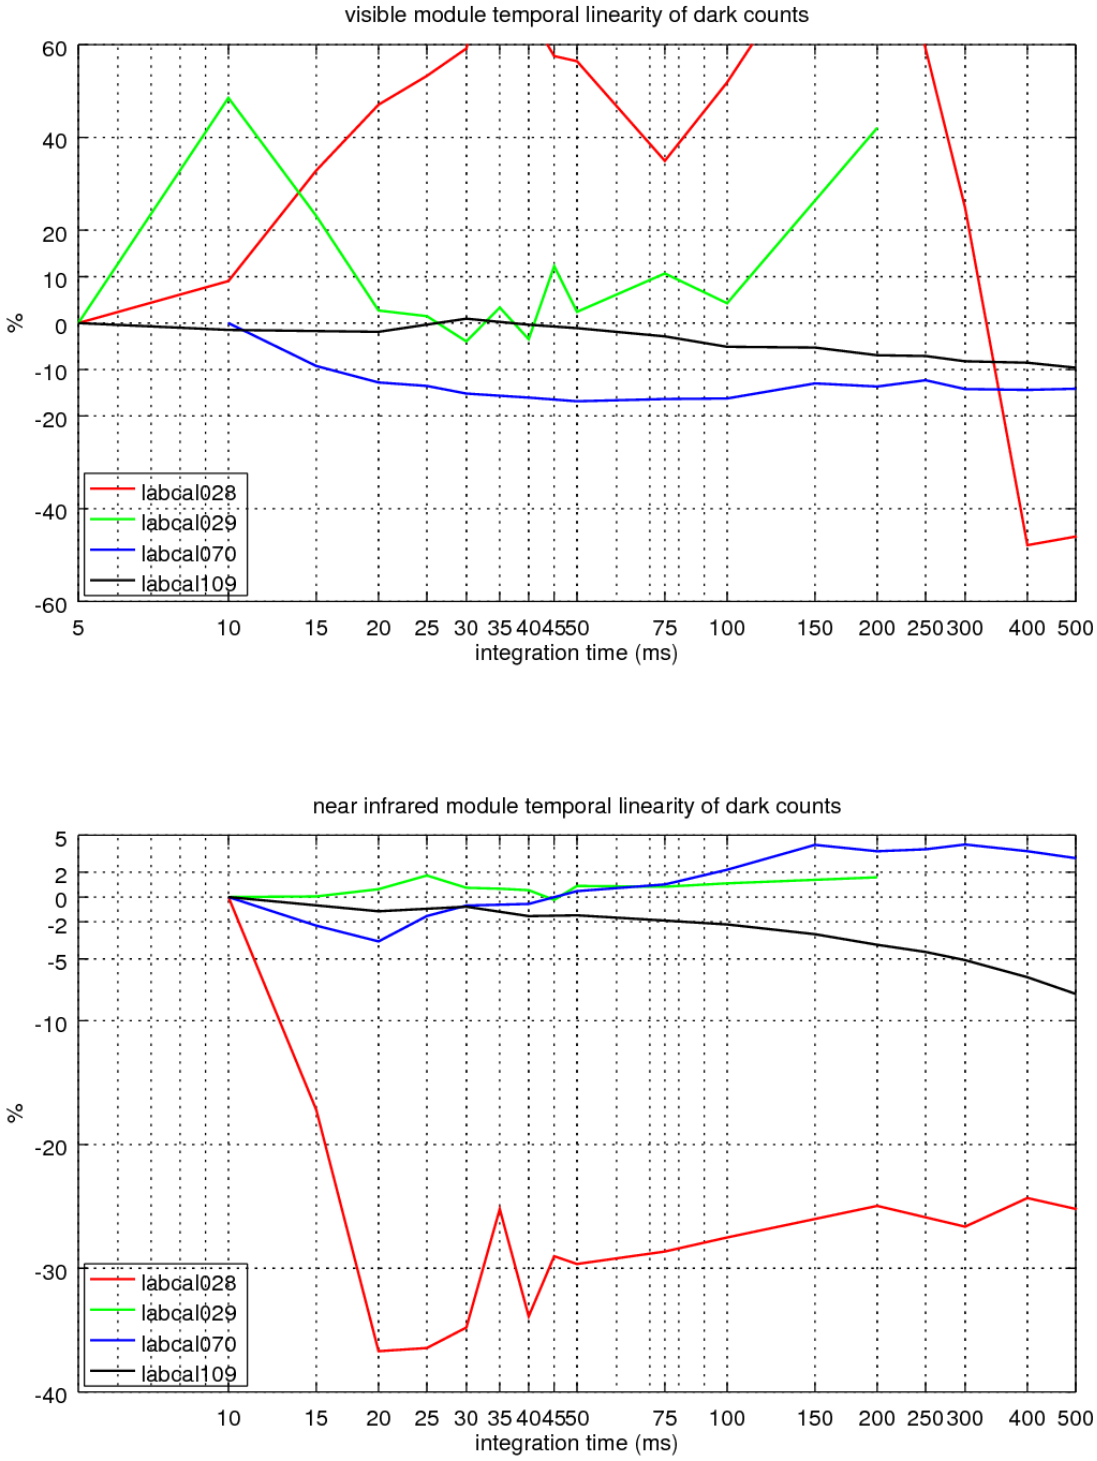

Figure 4.1. Temporal linearity of dark currents for SWS modules. The difference of counts/second at each integration time from the shortest is plotted as a percentage of the value at the shortest time.

© Crown copyright 2017 Met Office

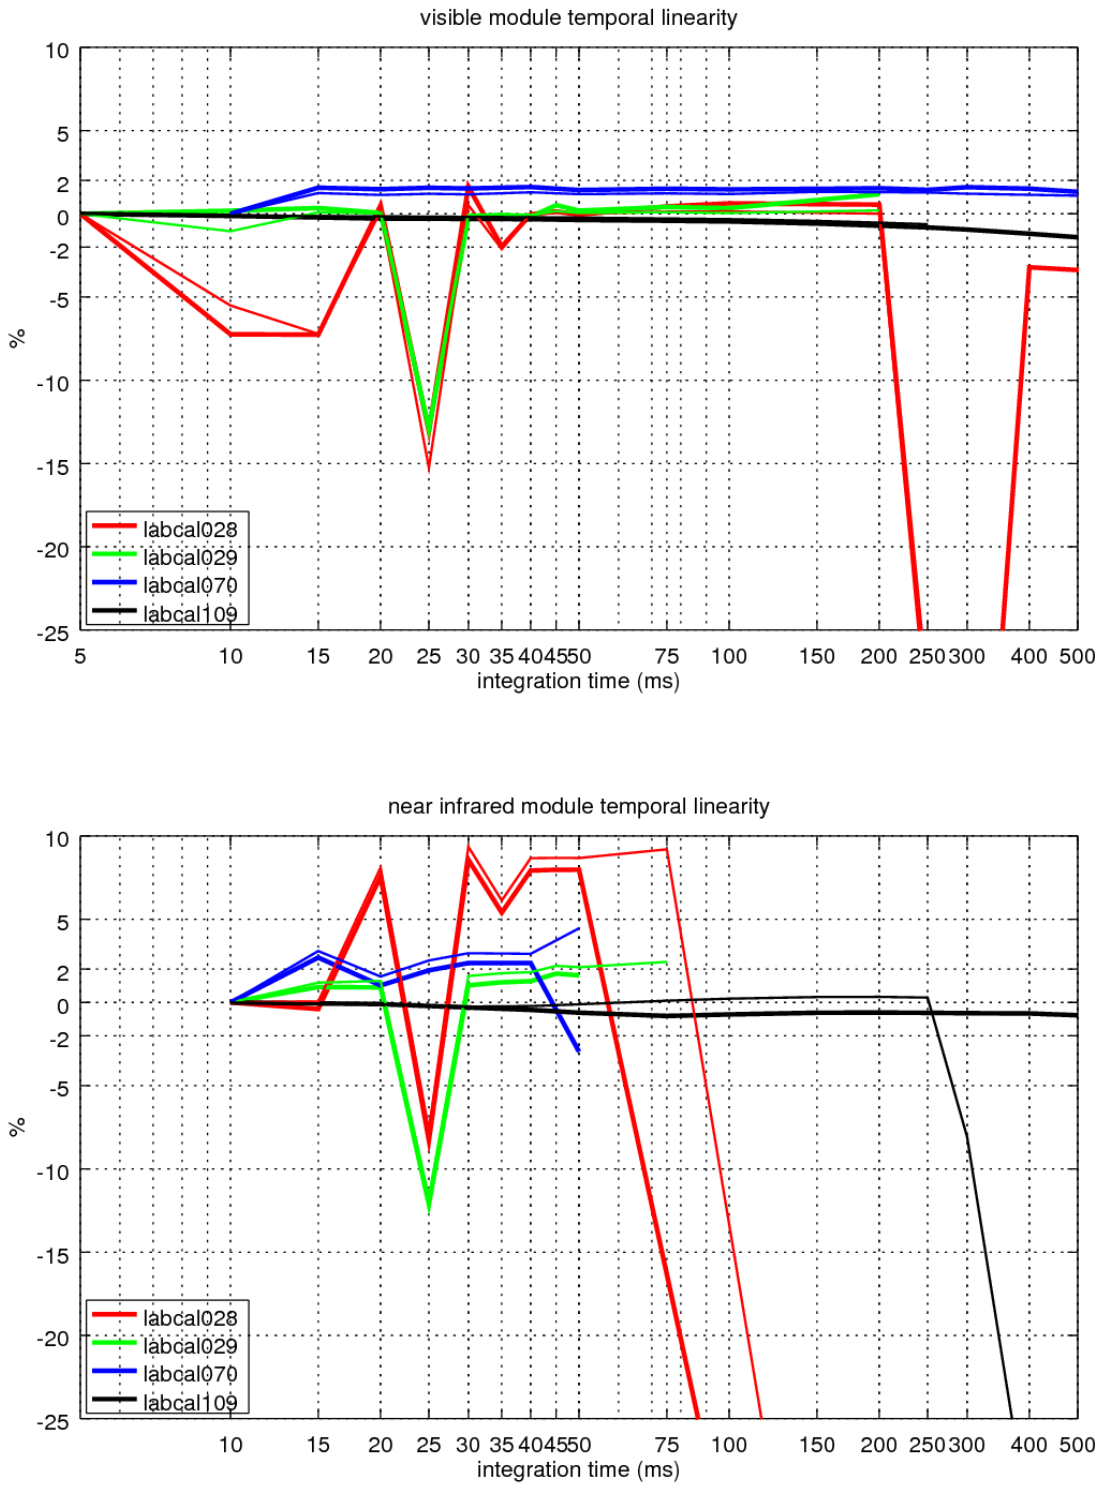

Figure 4.2. Temporal linearity of dark-corrected measurements of integrating sphere. The difference of counts/second at each integration time from the shortest is plotted as a percentage of the value at the shortest time.

© Crown copyright 2017 Met Office

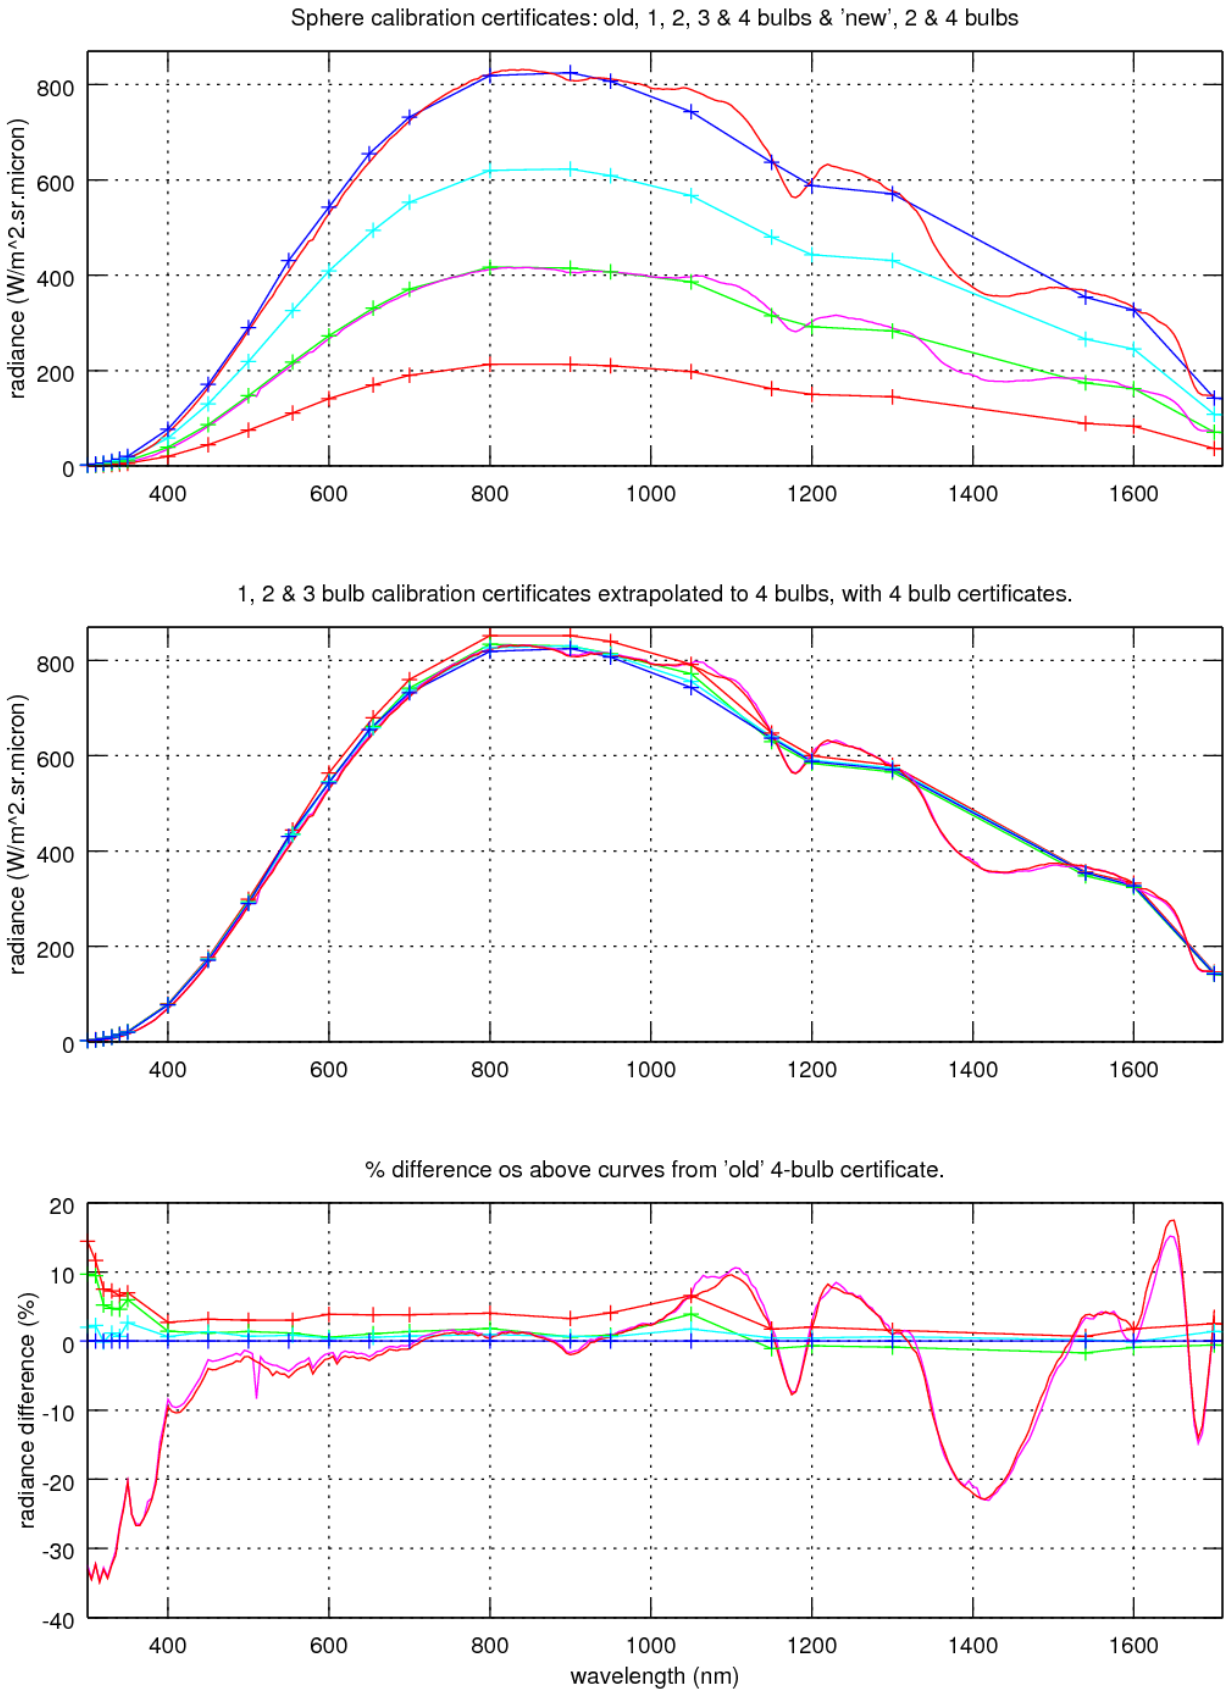

Figure 4.3 Linearity of sphere output. Top, calibration certificates produced with 1, 2, 3 and 4 bulbs illuminated. Middle, certificares for 1,2 and 3 bulbs extrapolated to 4 bulbs. Bottom, difference of other certificates from the oldest 4-bulb certificate.

© Crown copyright 2017 Met Office

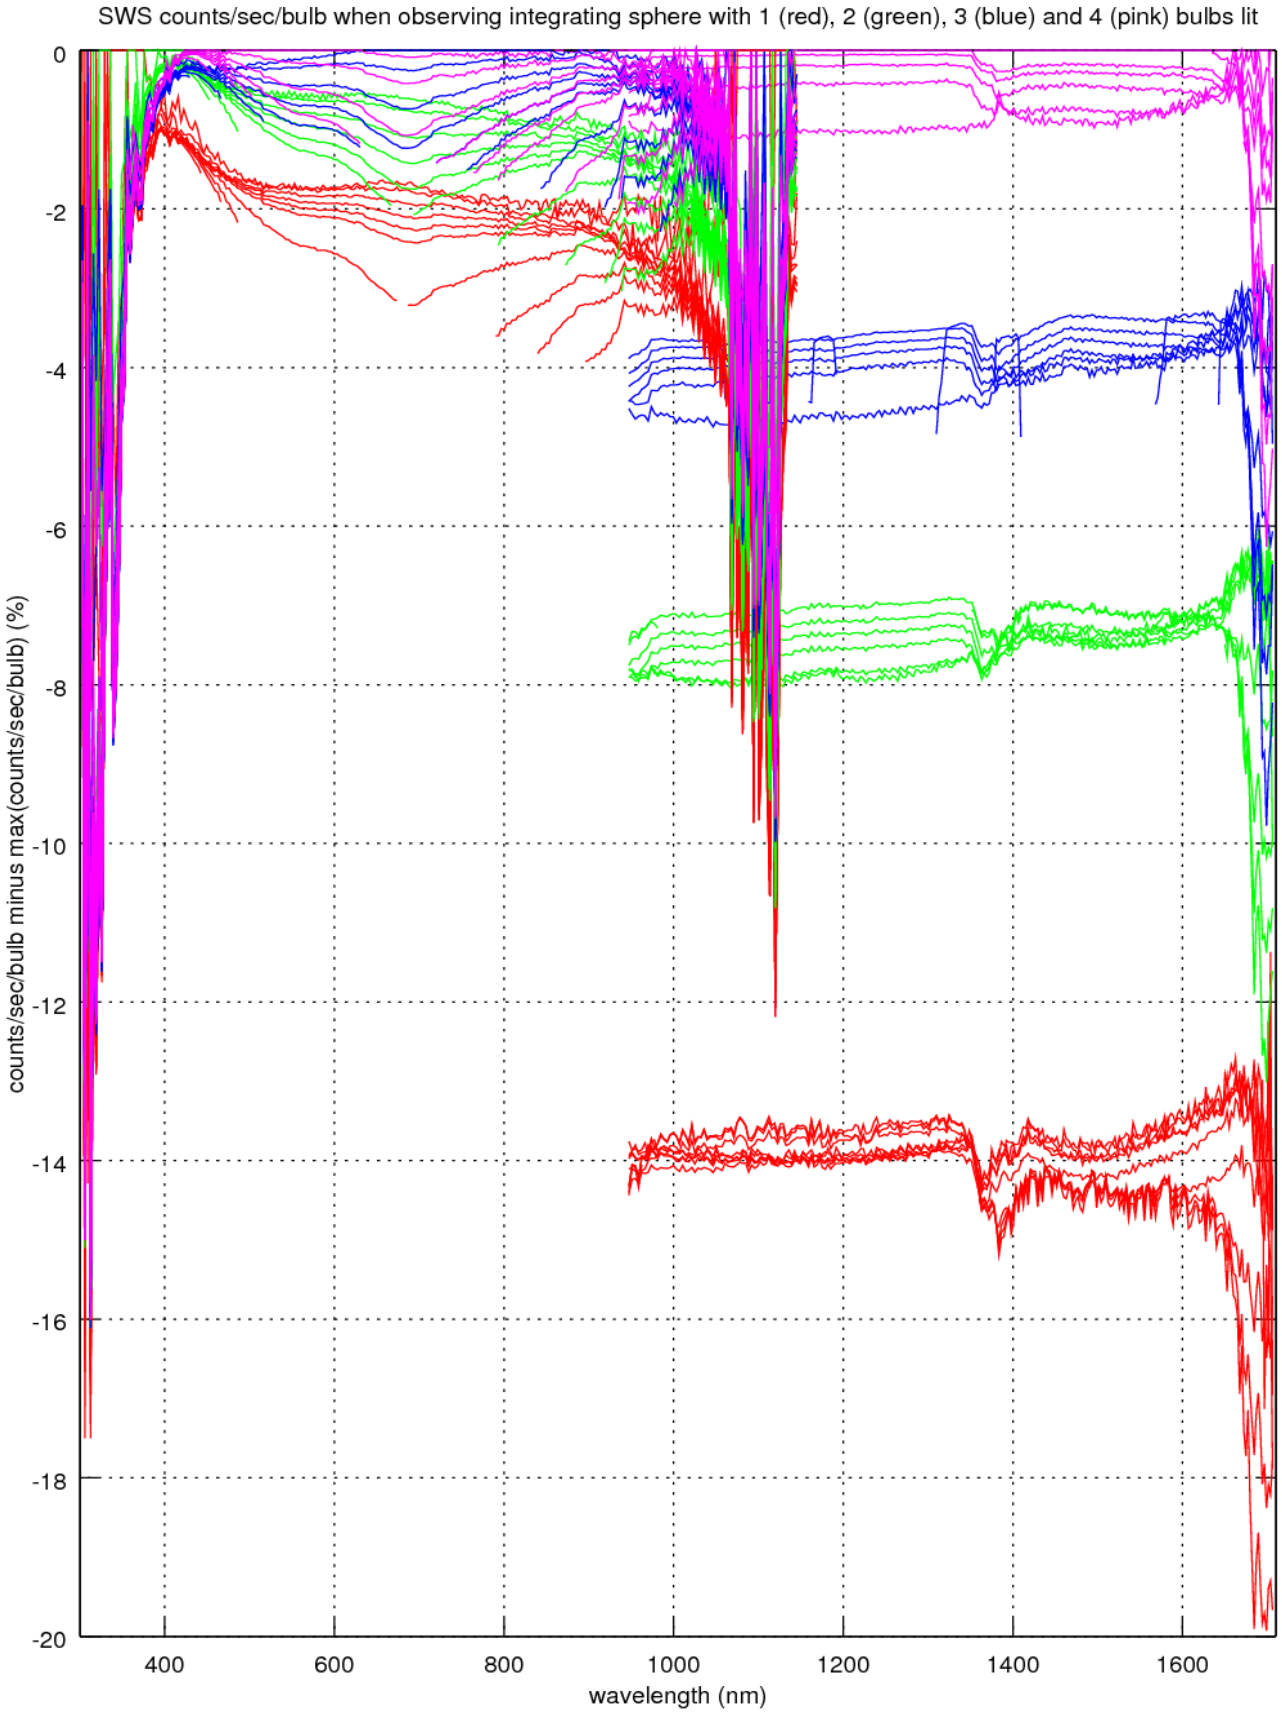

Figure 4.4. Comparison of SWS measurements of sphere output with different numbers of bulbs.

© Crown copyright 2017 Met Office

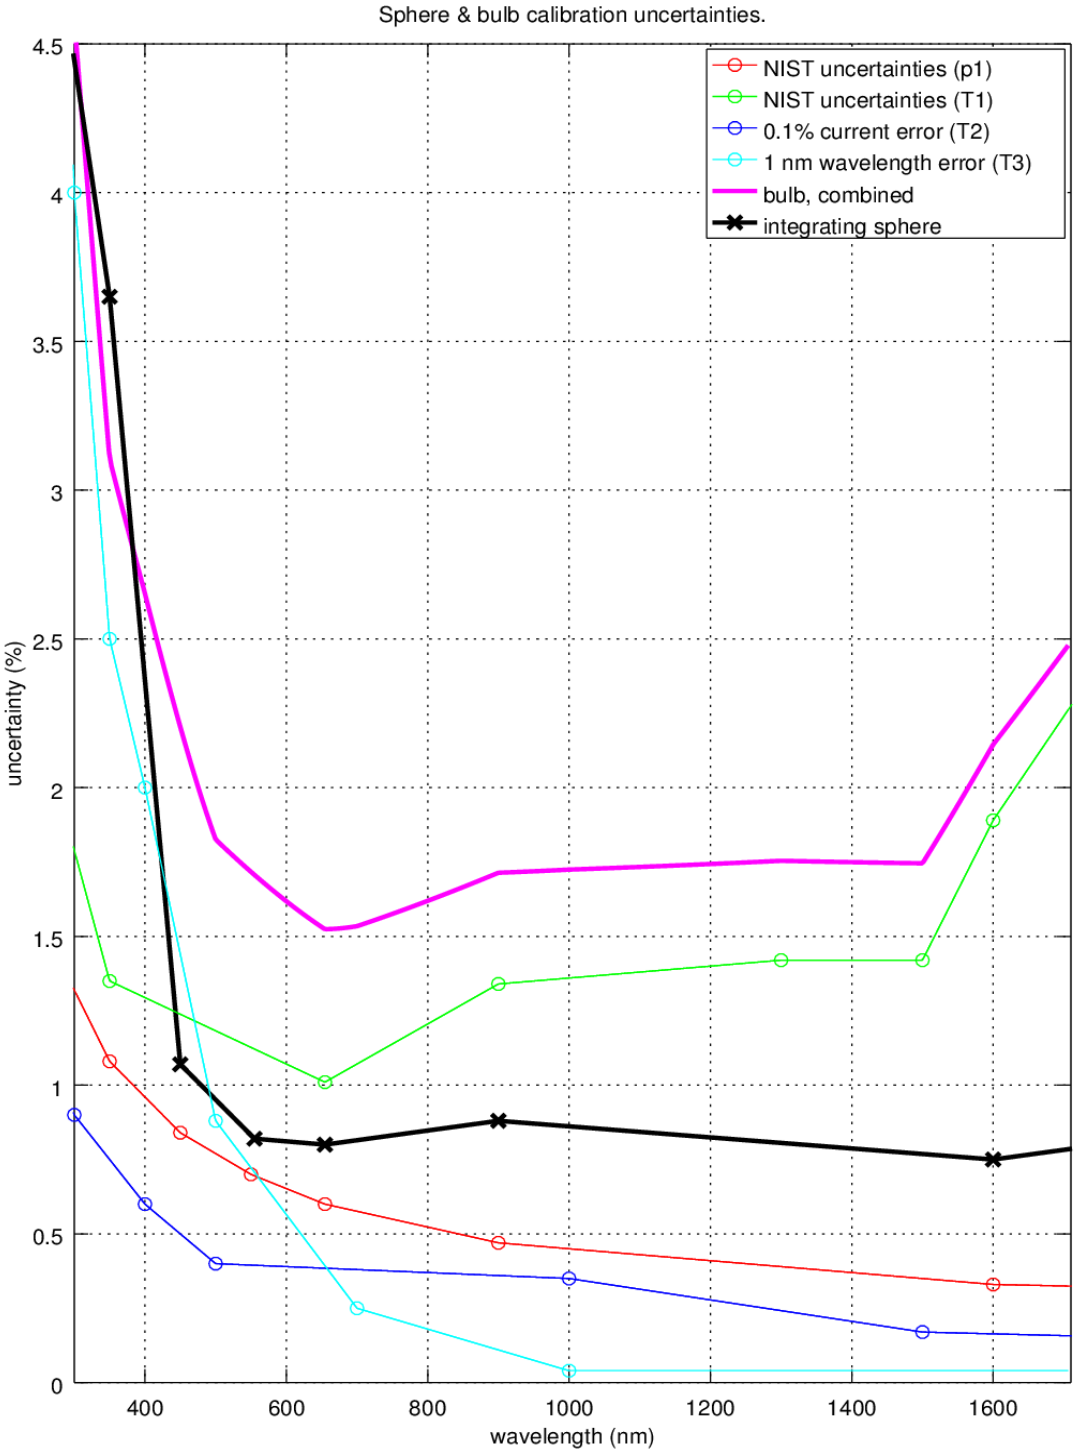

Figure 5.1. Uncertainties from calibration of local standards.

© Crown copyright 2017 Met Office

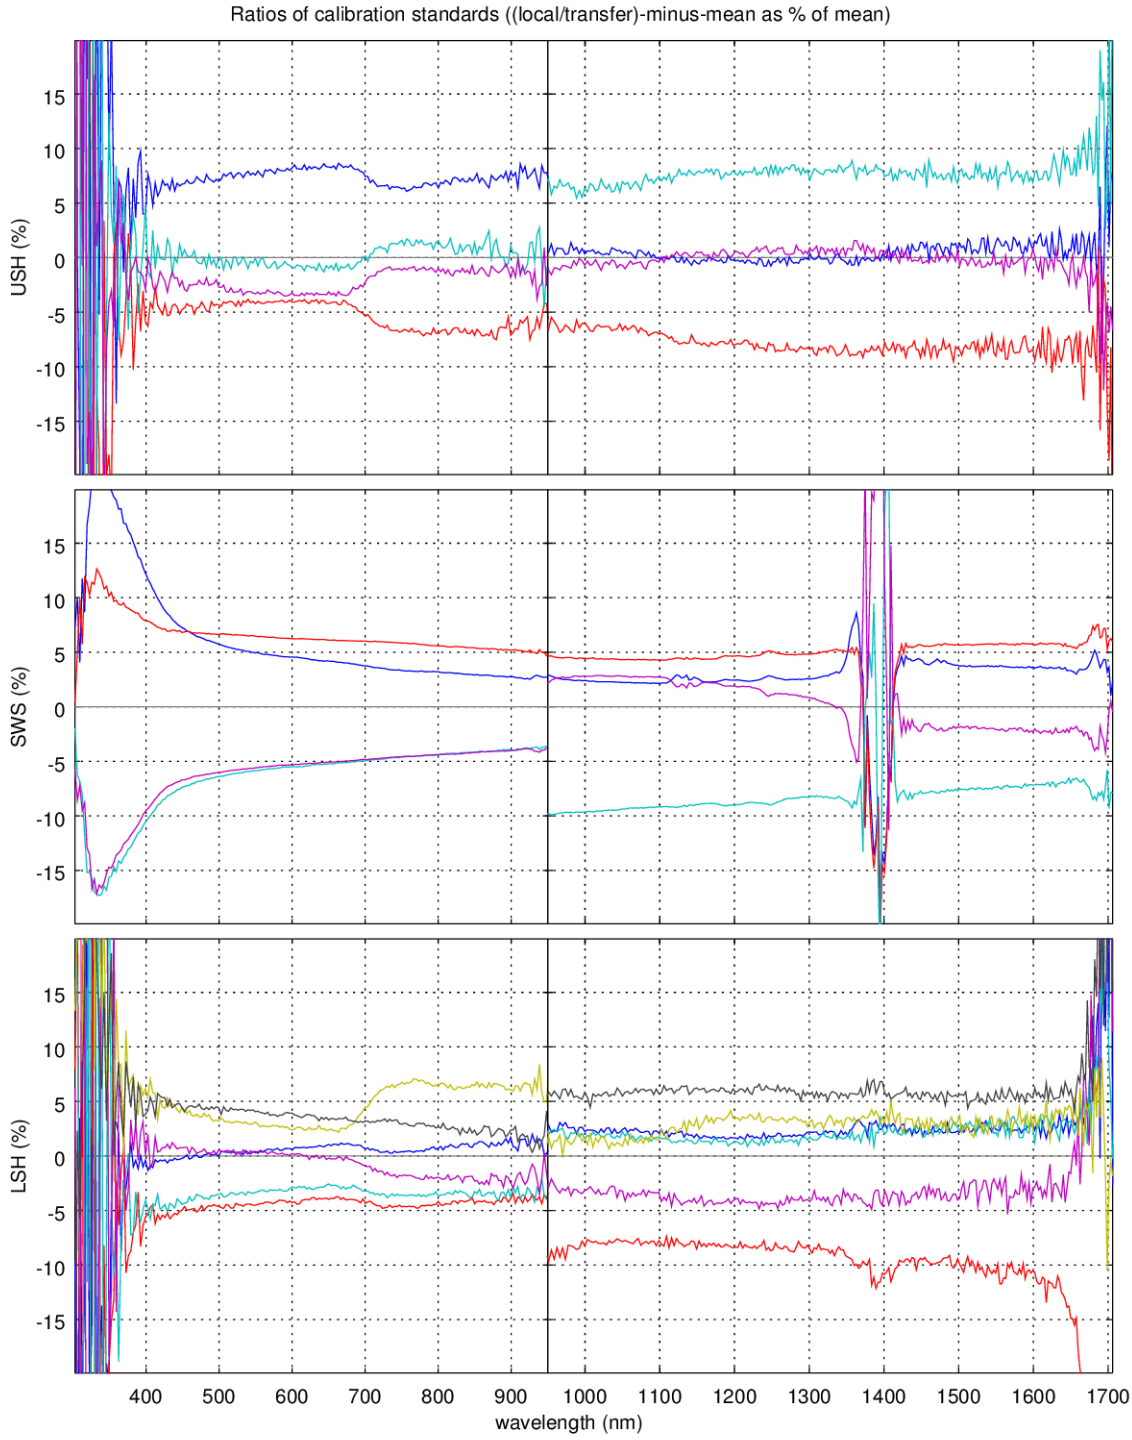

Figure 6.1. Comparison of successive measurements of lab- and transfer standards.

© Crown copyright 2017 Met Office

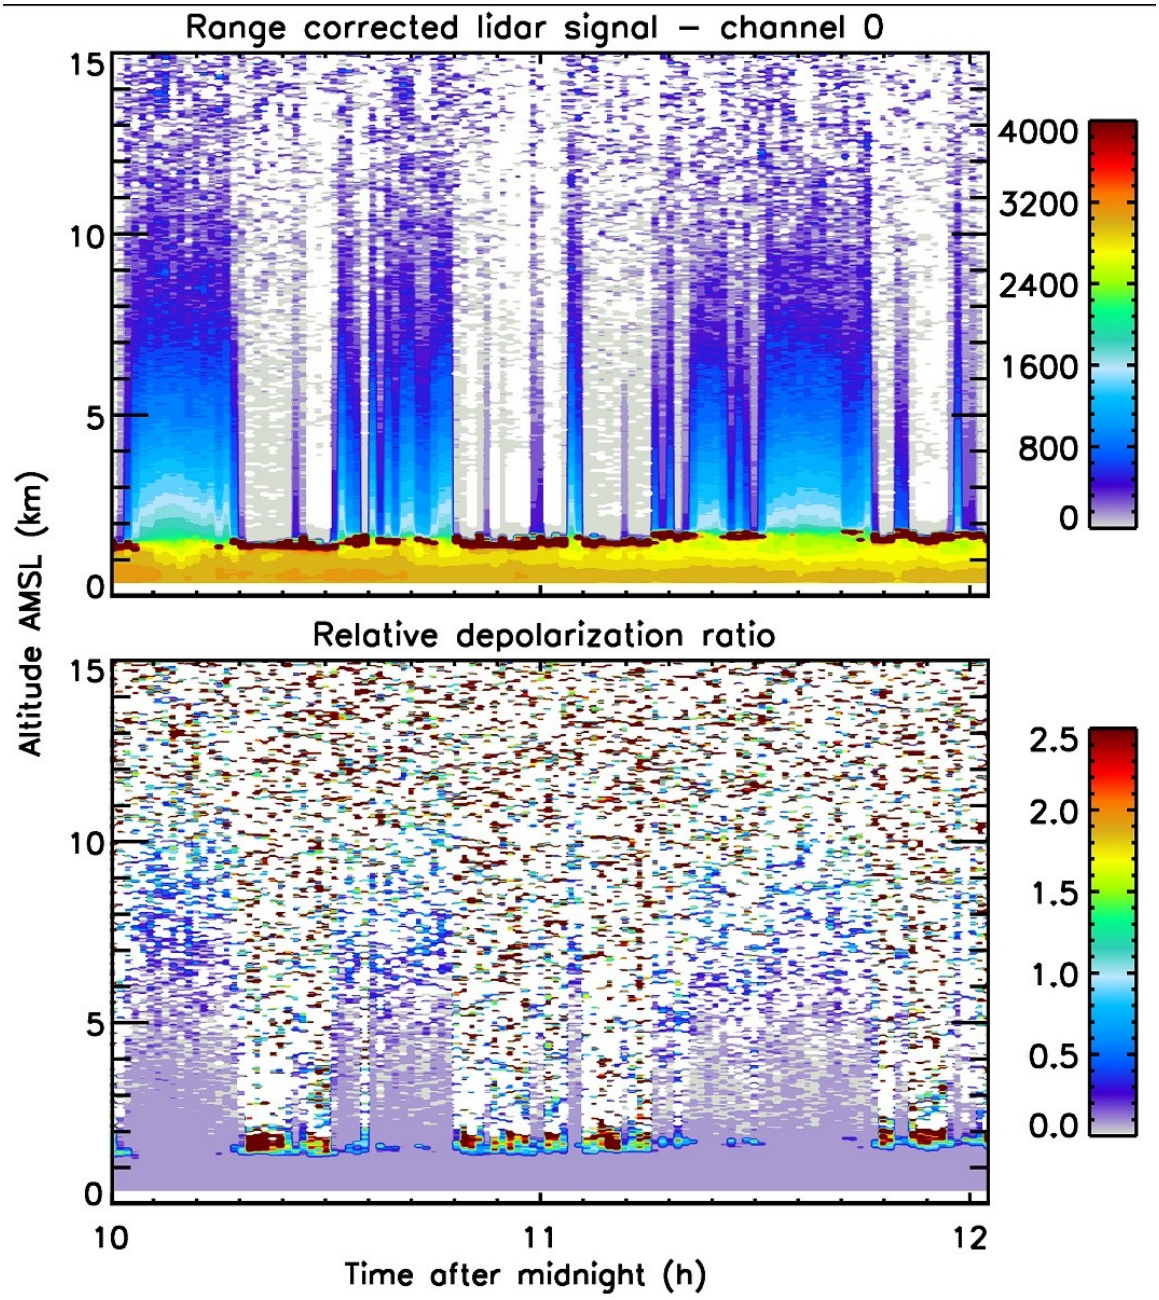

Figure 7.1. Lidar data for Cimel-SWS comparison

© Crown copyright 2017 Met Office

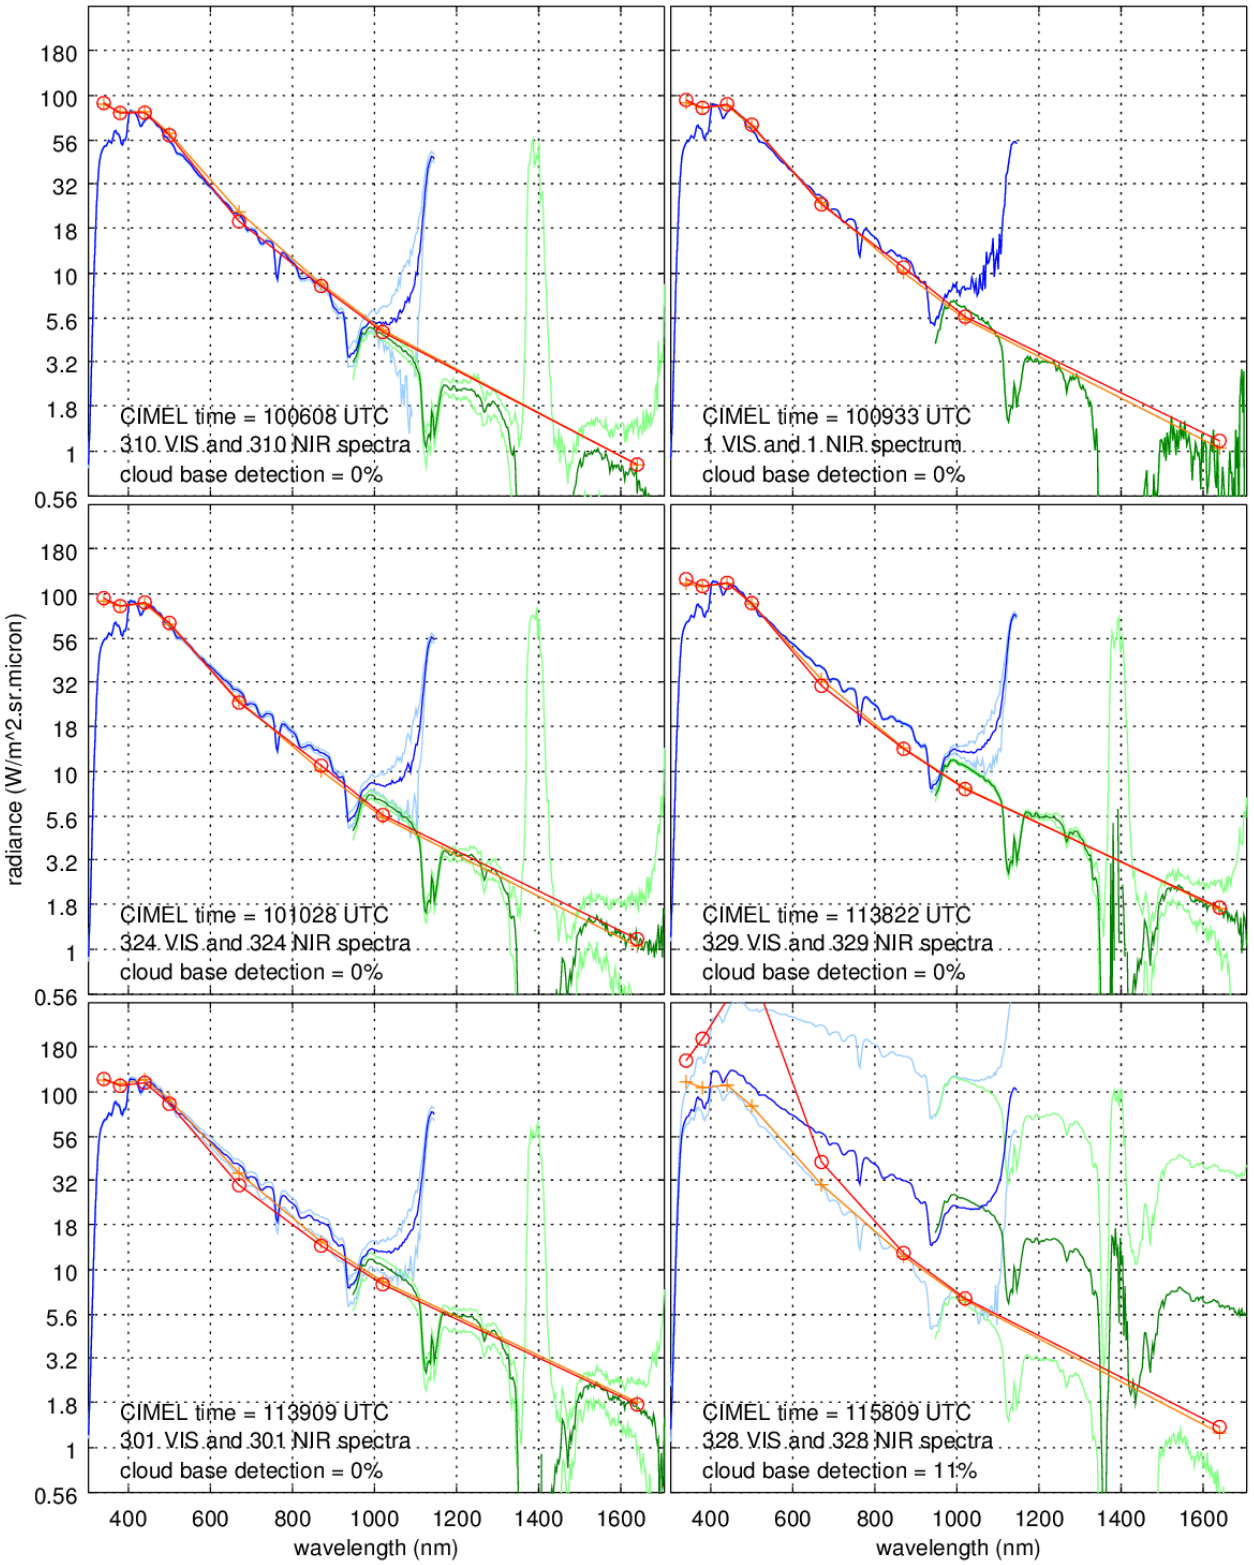

Figure 7.2. Six Cimel-SWS comparisons.

© Crown copyright 2017 Met Office

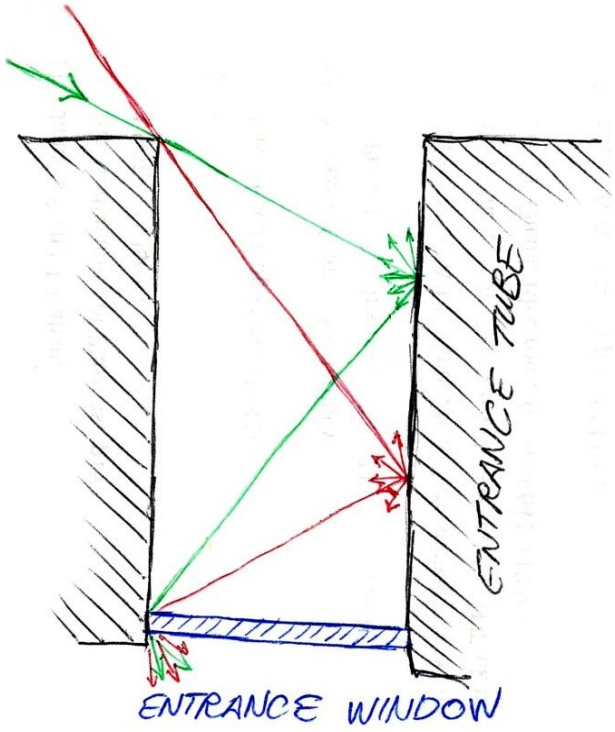

Figure 7.3. Cartoon illustrating stray light acquisition.

© Crown copyright 2017 Met Office

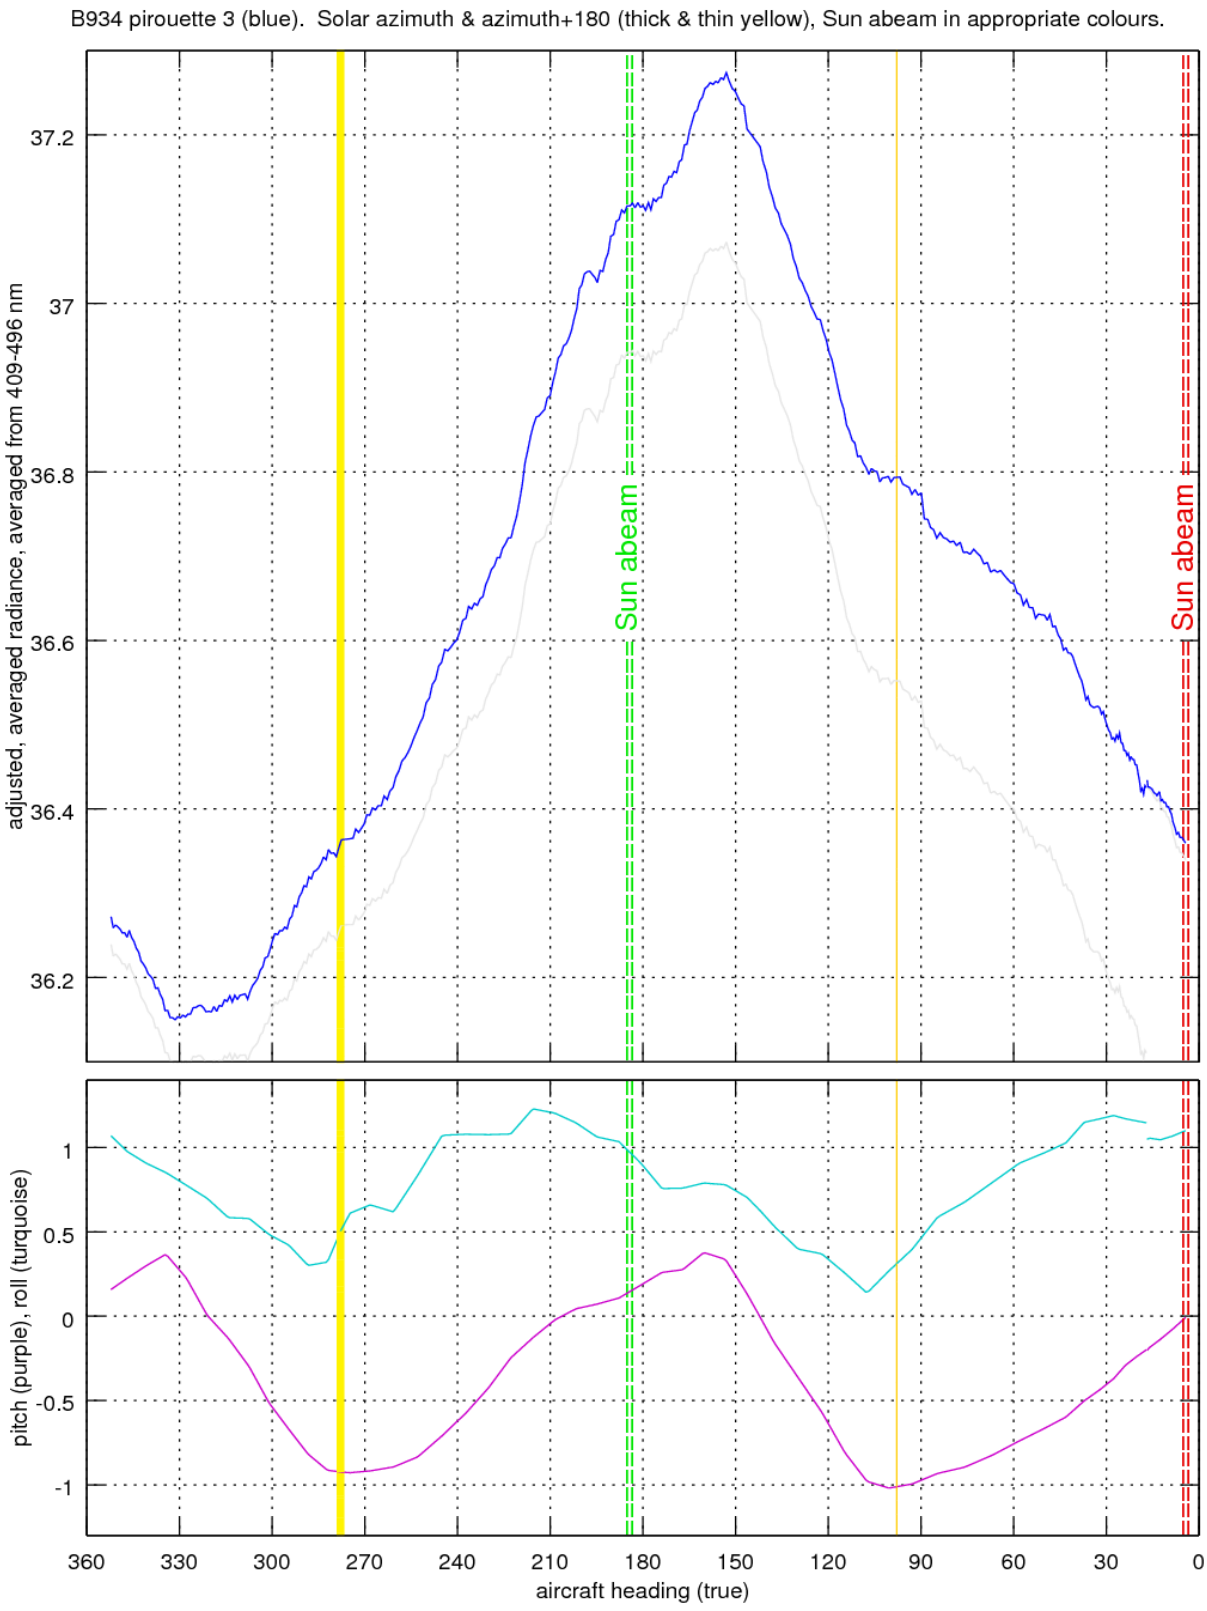

Figure 8.1. Averaged SWS zenith sky radiance during a pirouette during ICE-D. Upper panel: radiance adjusted for increasing zenith angle (blue) and unadjusted (grey). Lower panel: pitch and roll. The aircraft pointing directly into the Sun is shown in thick yellow, out of Sun in thin yellow, and Sun abeam in appropriate colours.

© Crown copyright 2017 Met Office

Figure 8.2. SWS set up outside with obscurer fitted.

© Crown copyright 2017 Met Office

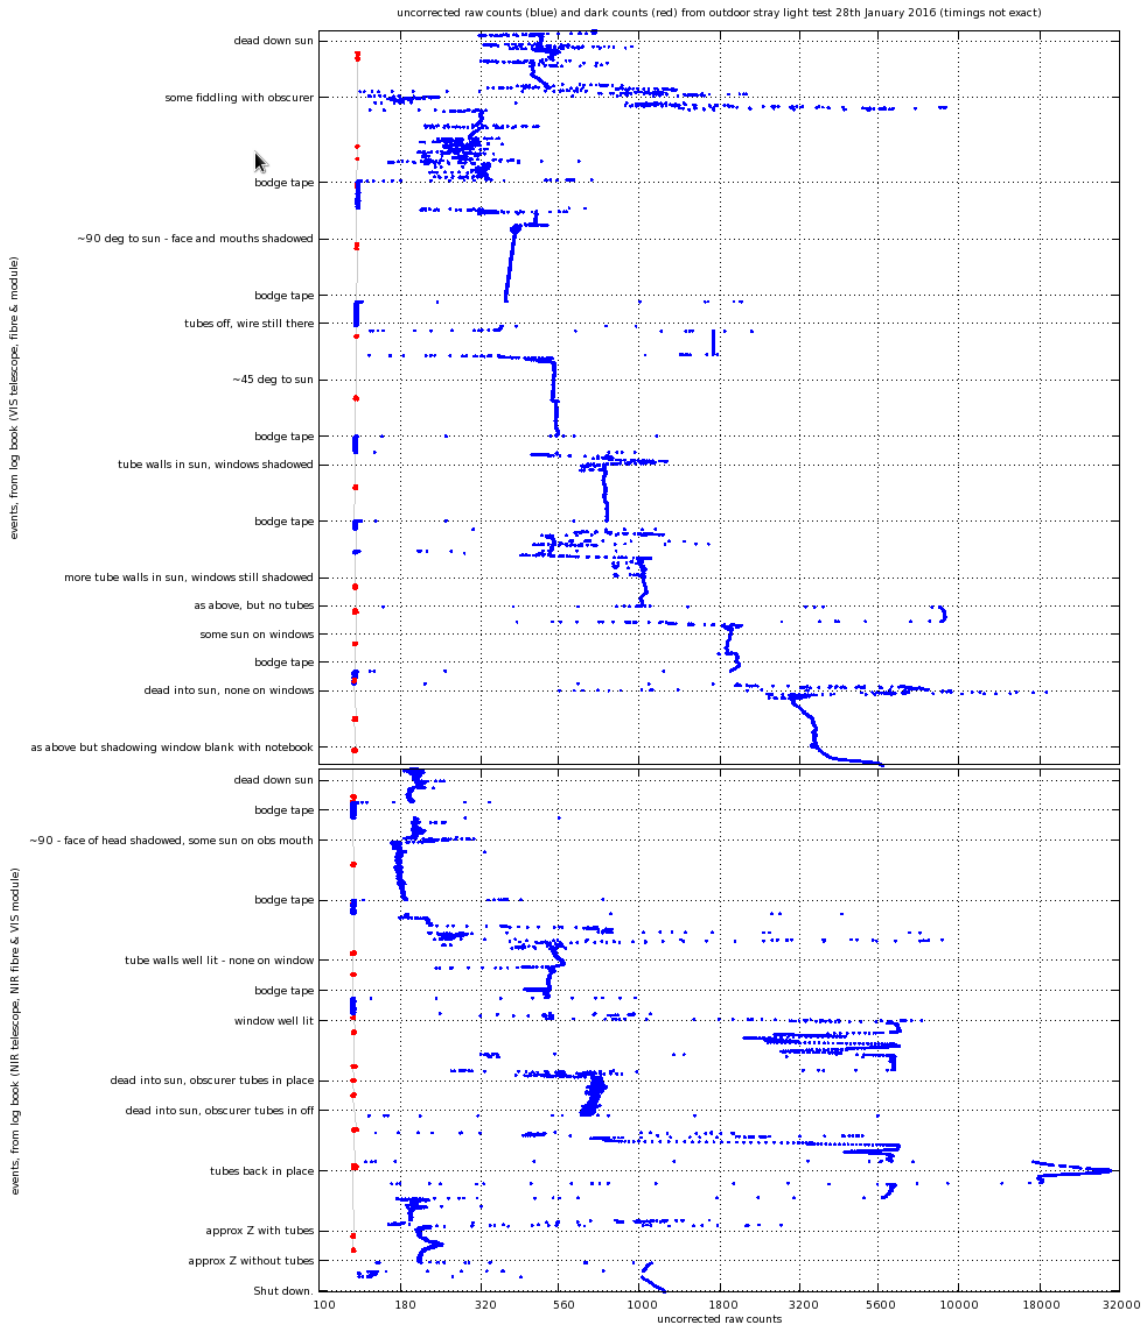

Figure 8.3. Time series of obscurer measurements, outdoors.

© Crown copyright 2017 Met Office

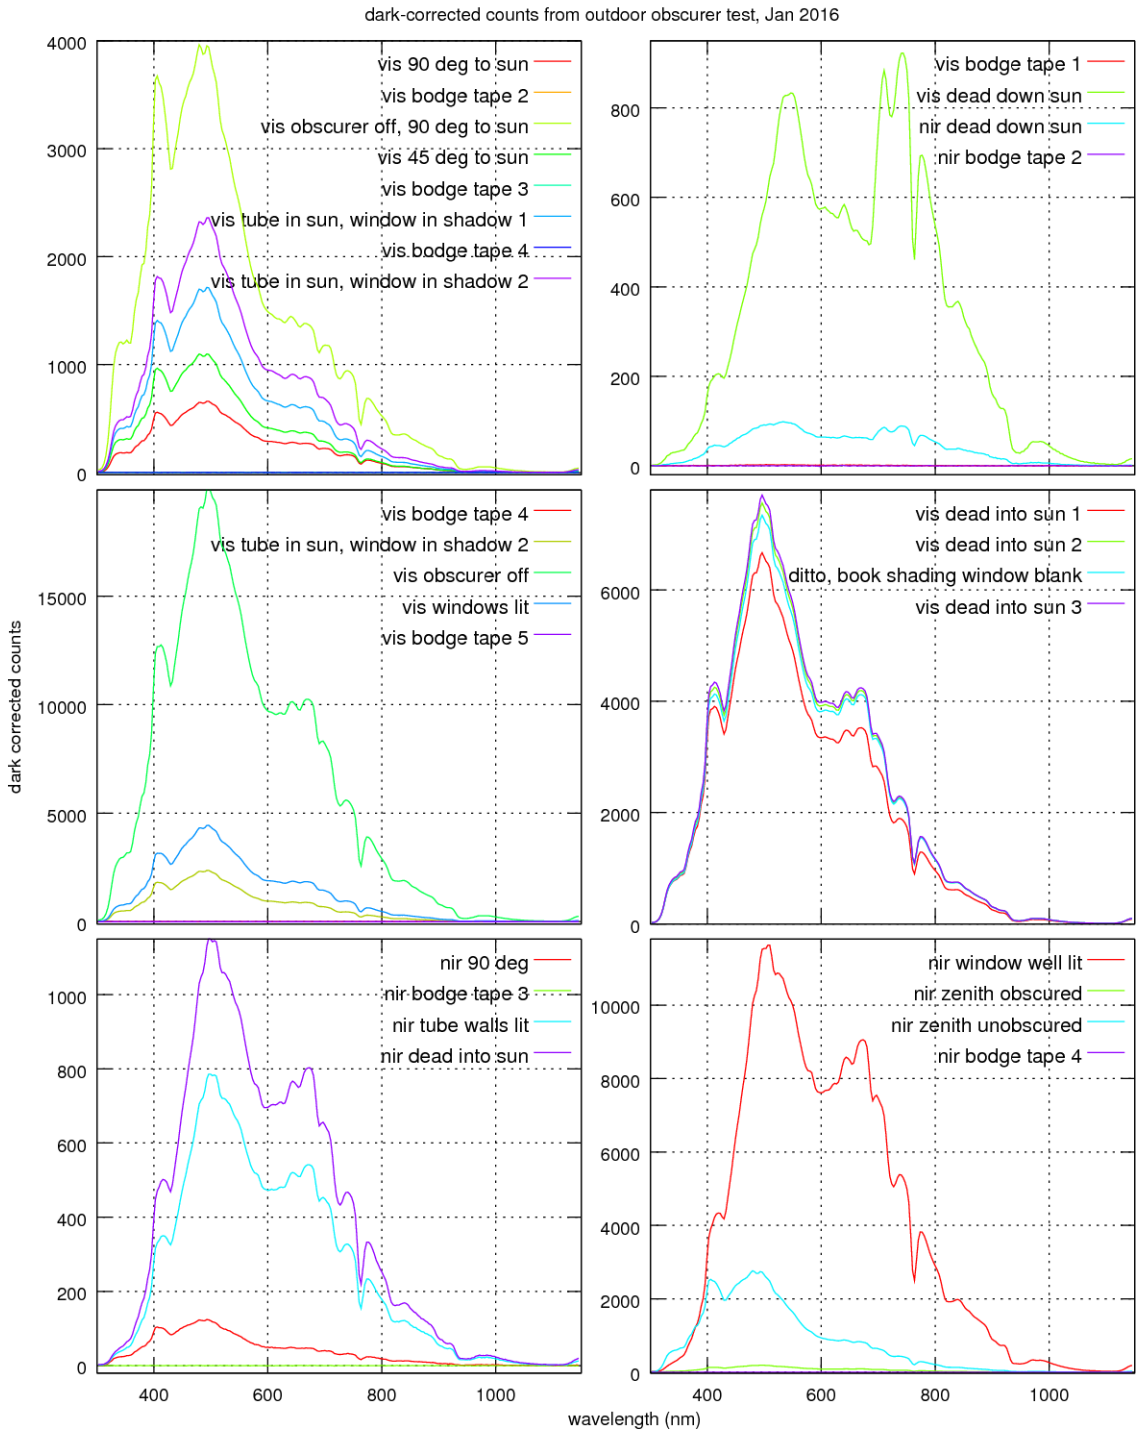

Figure 8.4. Raw spectra recorded during outdoor stray light test.

© Crown copyright 2017 Met Office

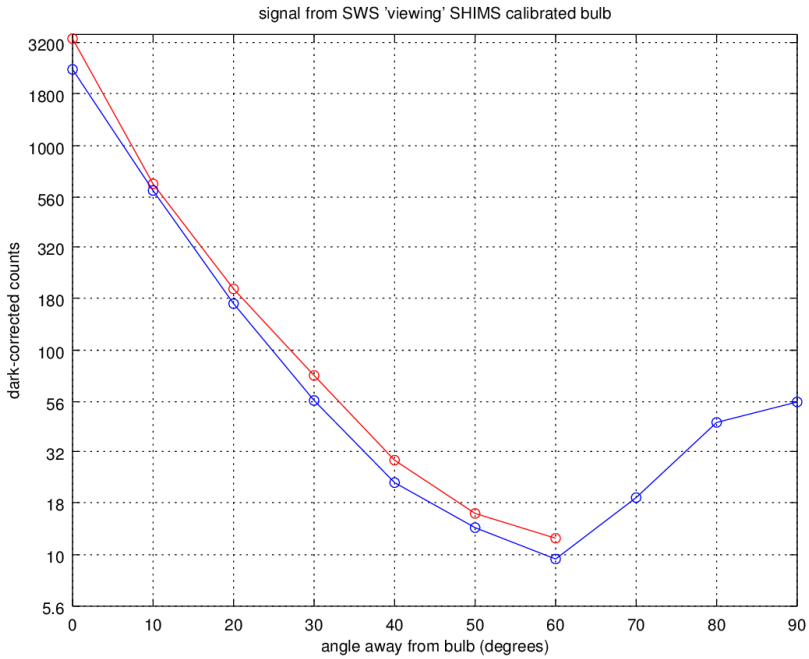

Figure 8.5. SWS measurements of SHIMS' calibrated bulb.

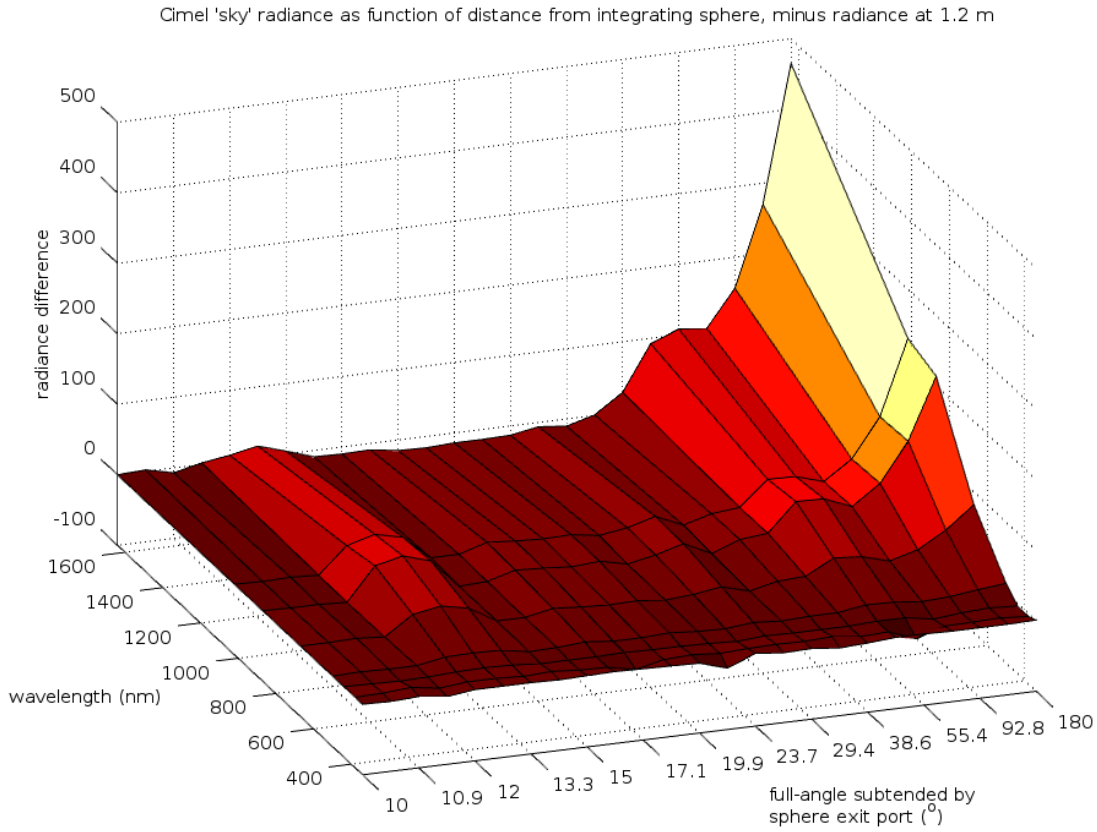

Figure 8.6. Change in radiance recorded using a Cimel to view the SWS integrating sphere at different distances.

© Crown copyright 2017 Met Office

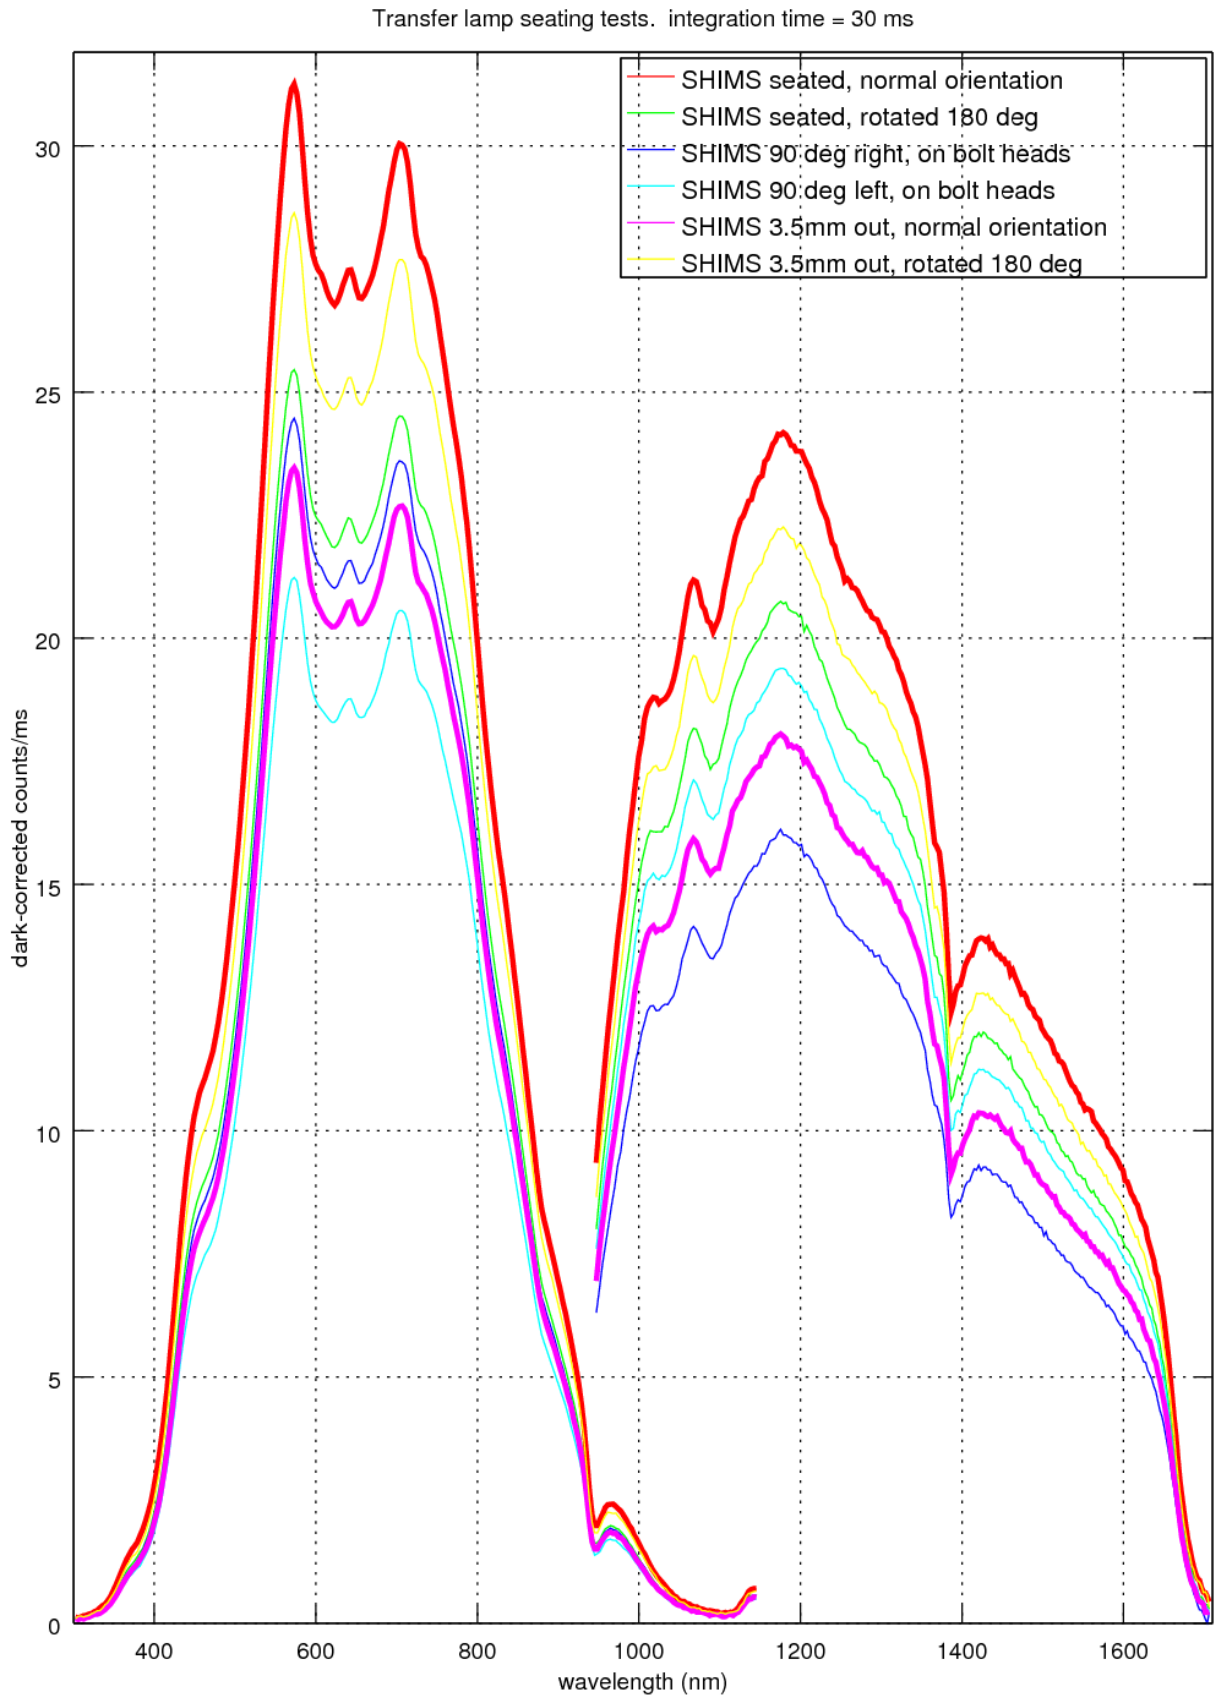

Figure 9.1. Transfer standard seating - 30 ms integration.

© Crown copyright 2017 Met Office

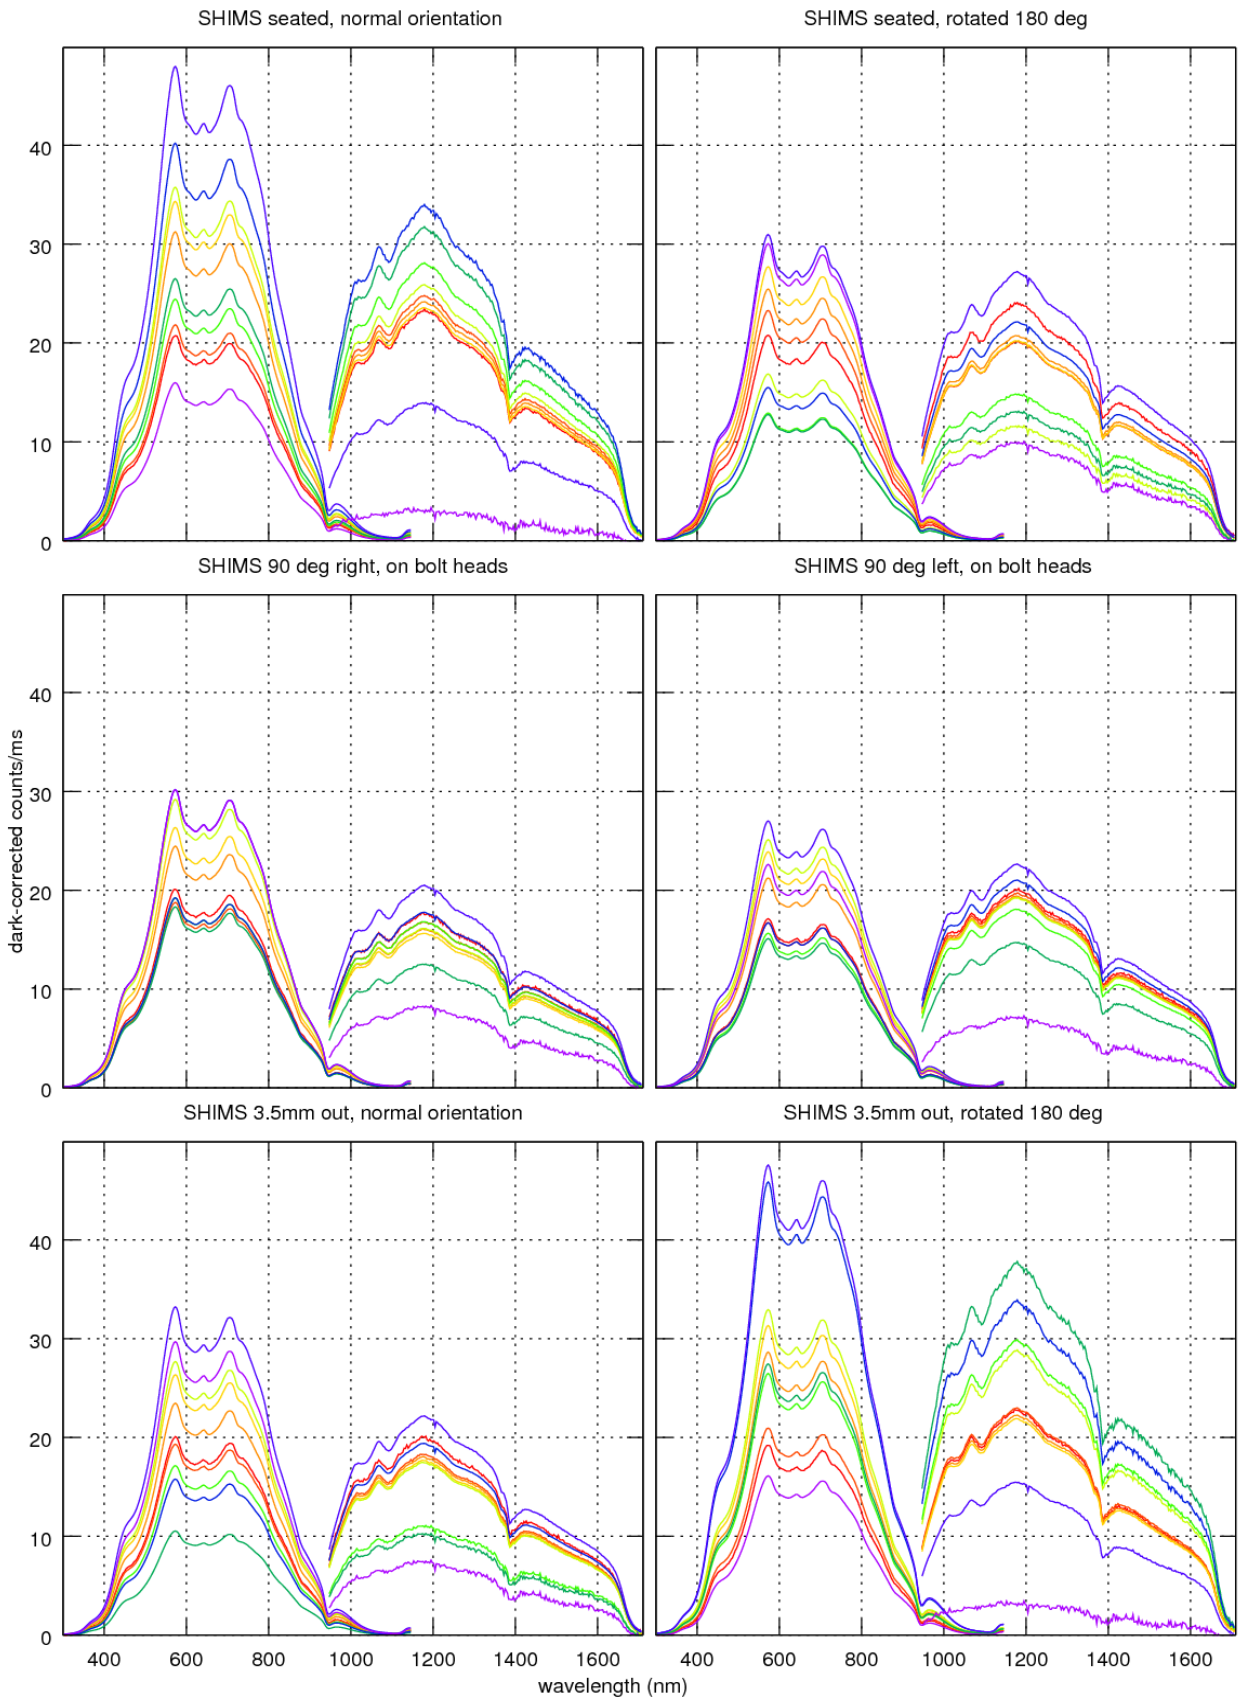

Figure 9.2. Transfer standard spectra recorded with various seating arrangements and integration times.

© Crown copyright 2017 Met Office

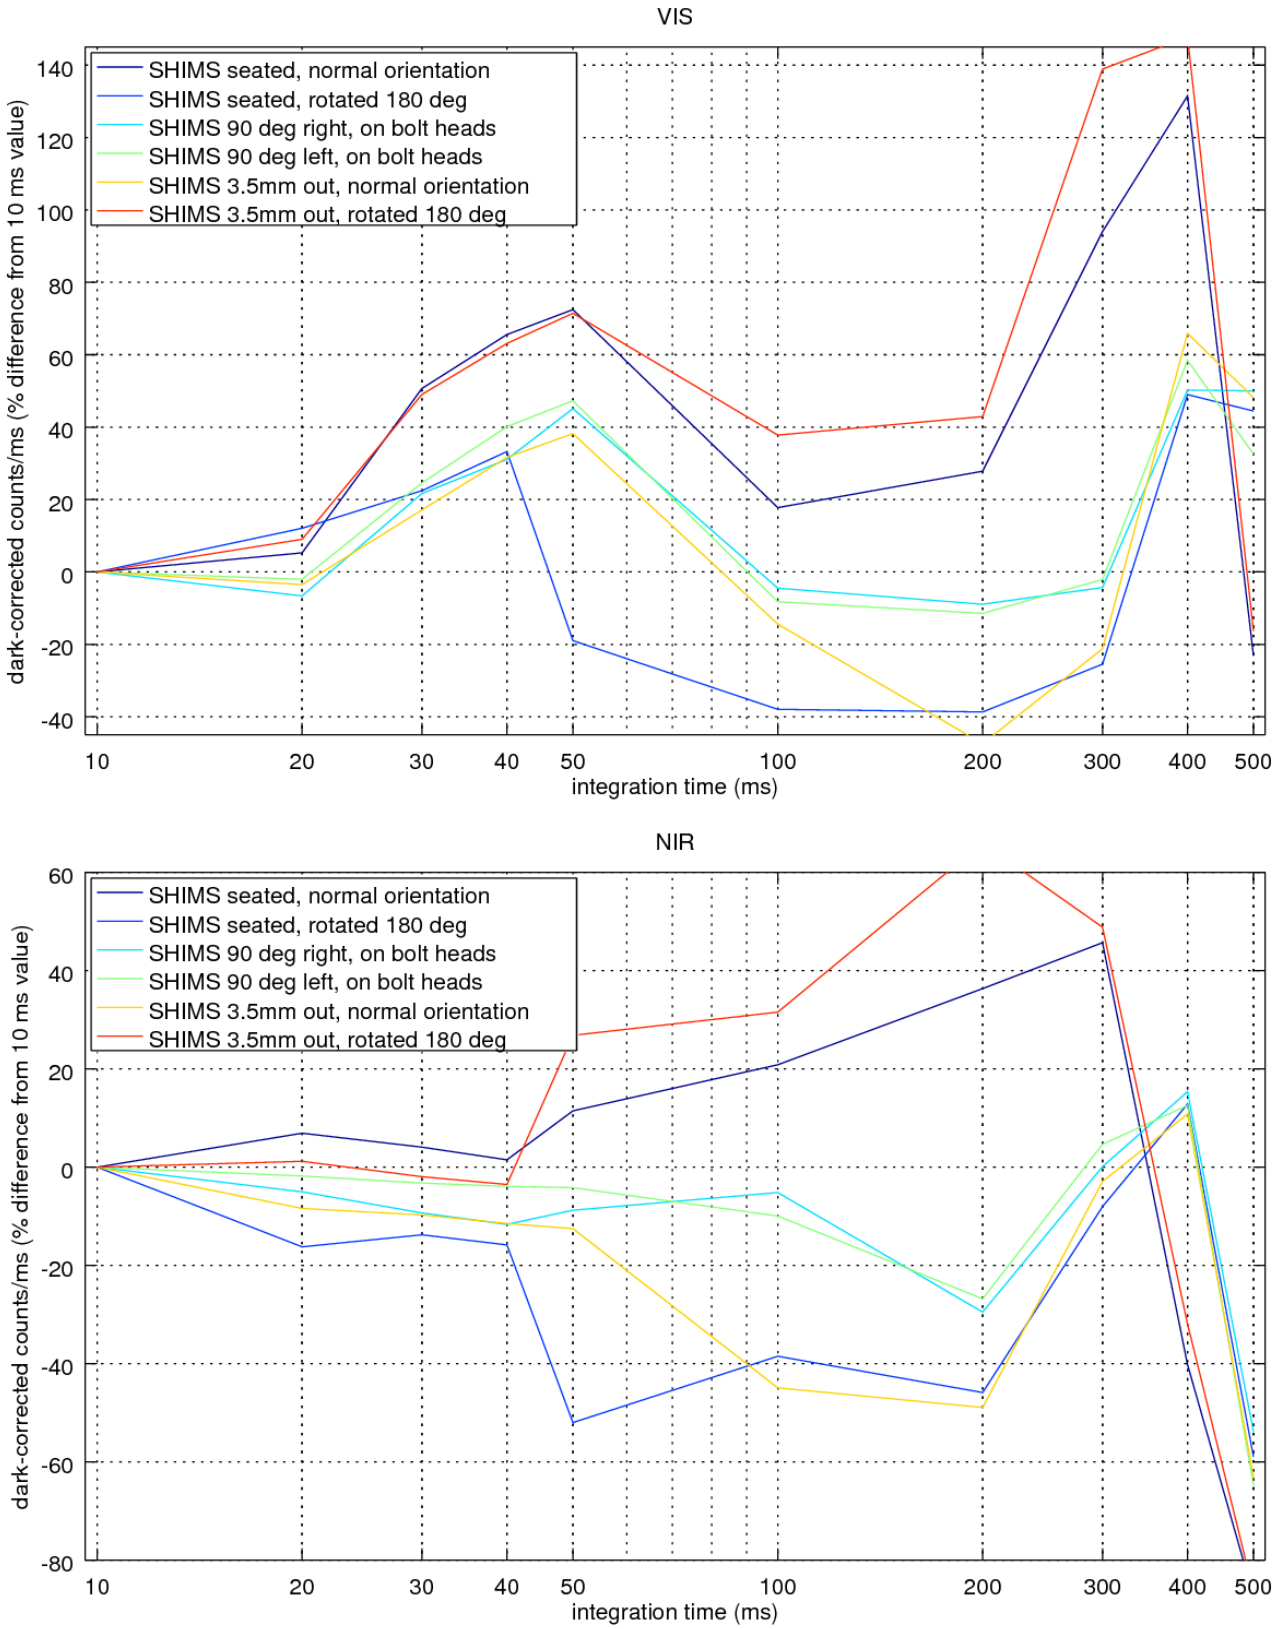

Figure 9.3. Transfer standard averaged, dark corrected counts/ms with various seating arrangements, plotted as a function of integration time.

© Crown copyright 2017 Met Office

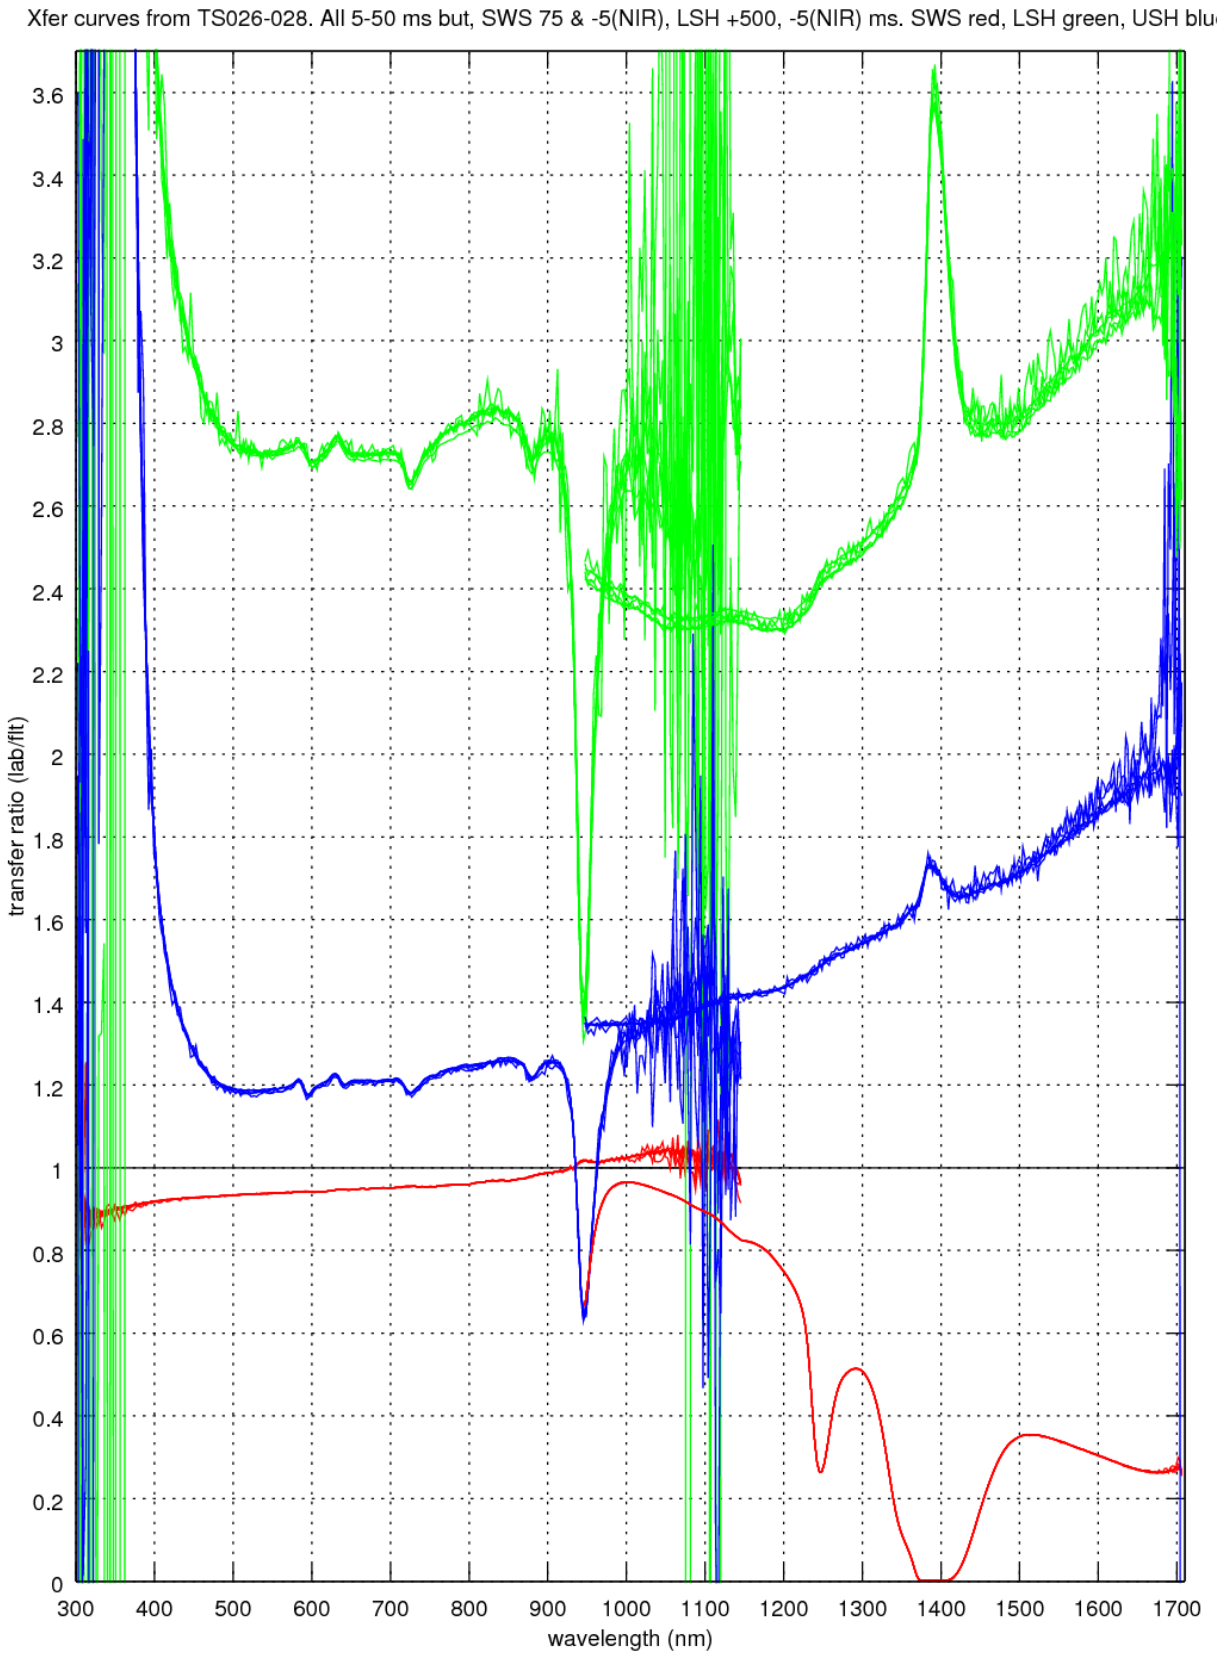

Figure 10.1. ICE-D transfer curves.

© Crown copyright 2017 Met Office

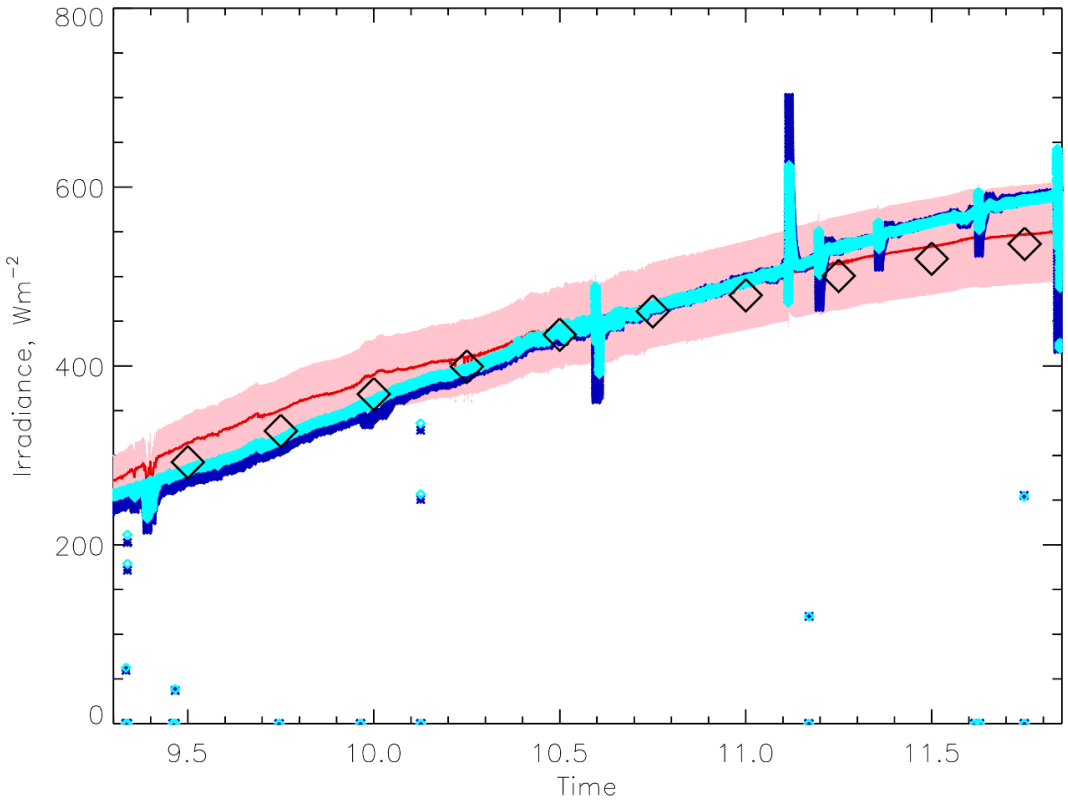

Figure 10.2. B923 downwelling visible; SHIMS, BBRs and model (Ryder).

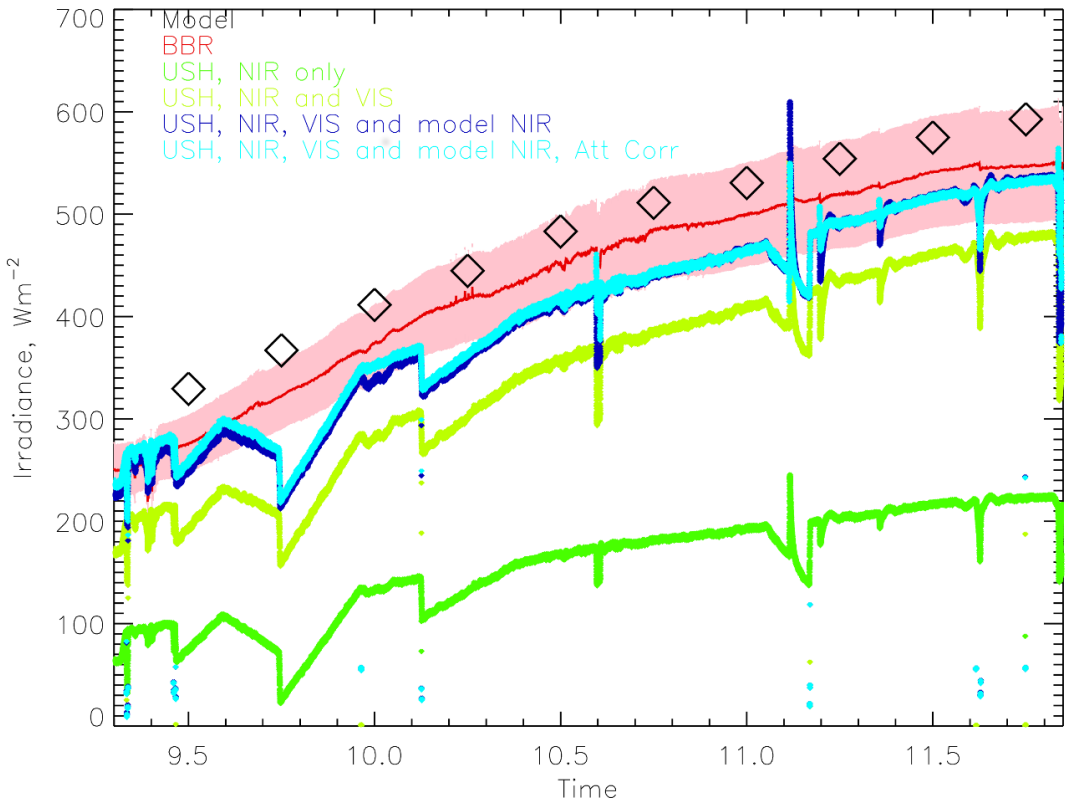

Figure 10.3. B923 downwelling red; SHIMS, BBRs and model (Ryder).

© Crown copyright 2017 Met Office

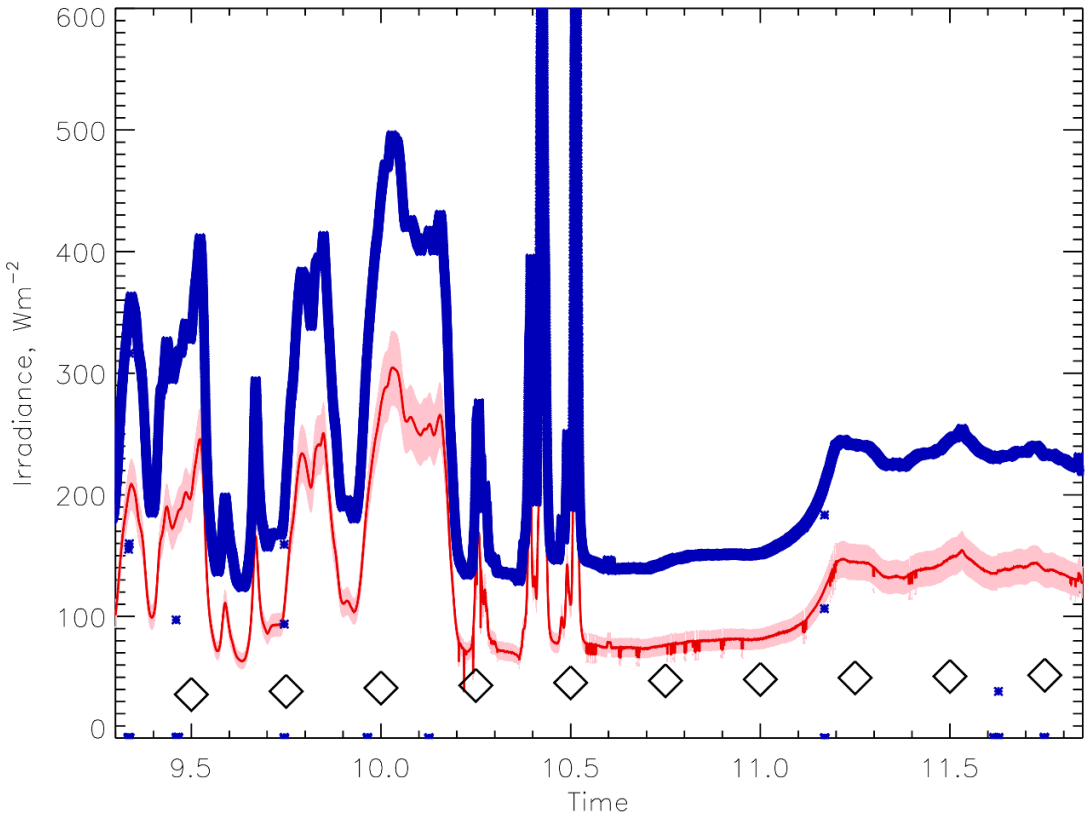

Figure 10.4. B923 upwelling visible; SHIMS, BBRs and model (Ryder).

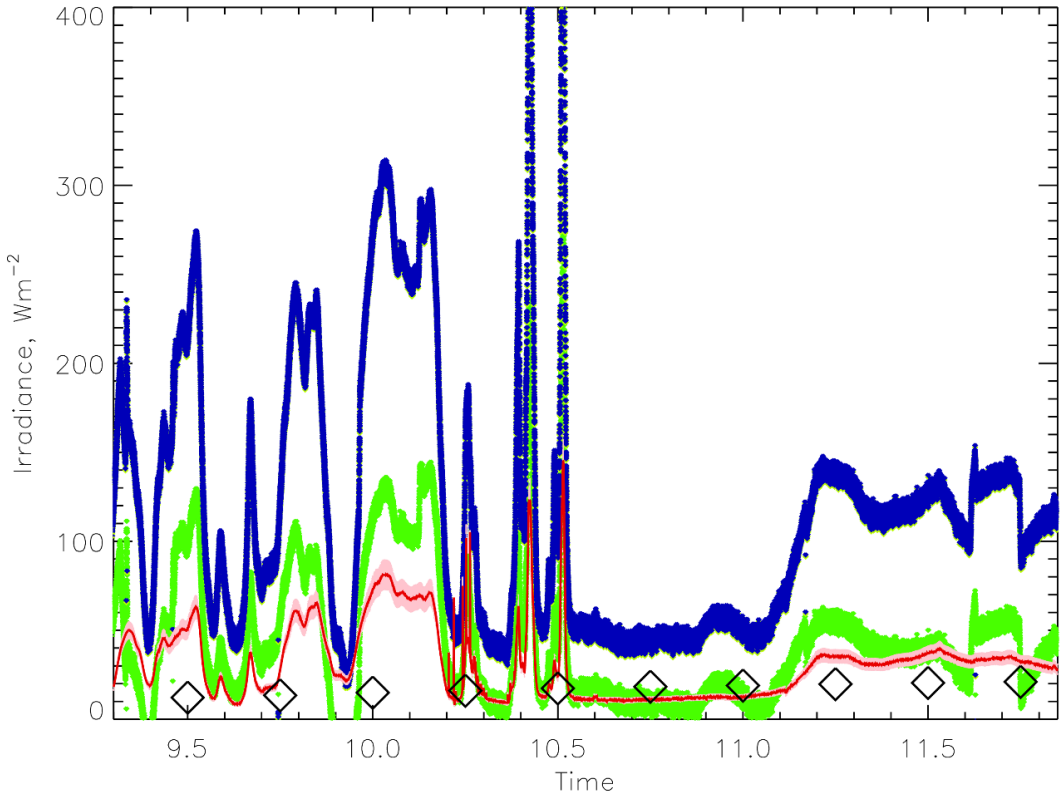

Figure 10.5. B923 upwelling red; SHIMS, BBRs and model (Ryder).

© Crown copyright 2017 Met Office

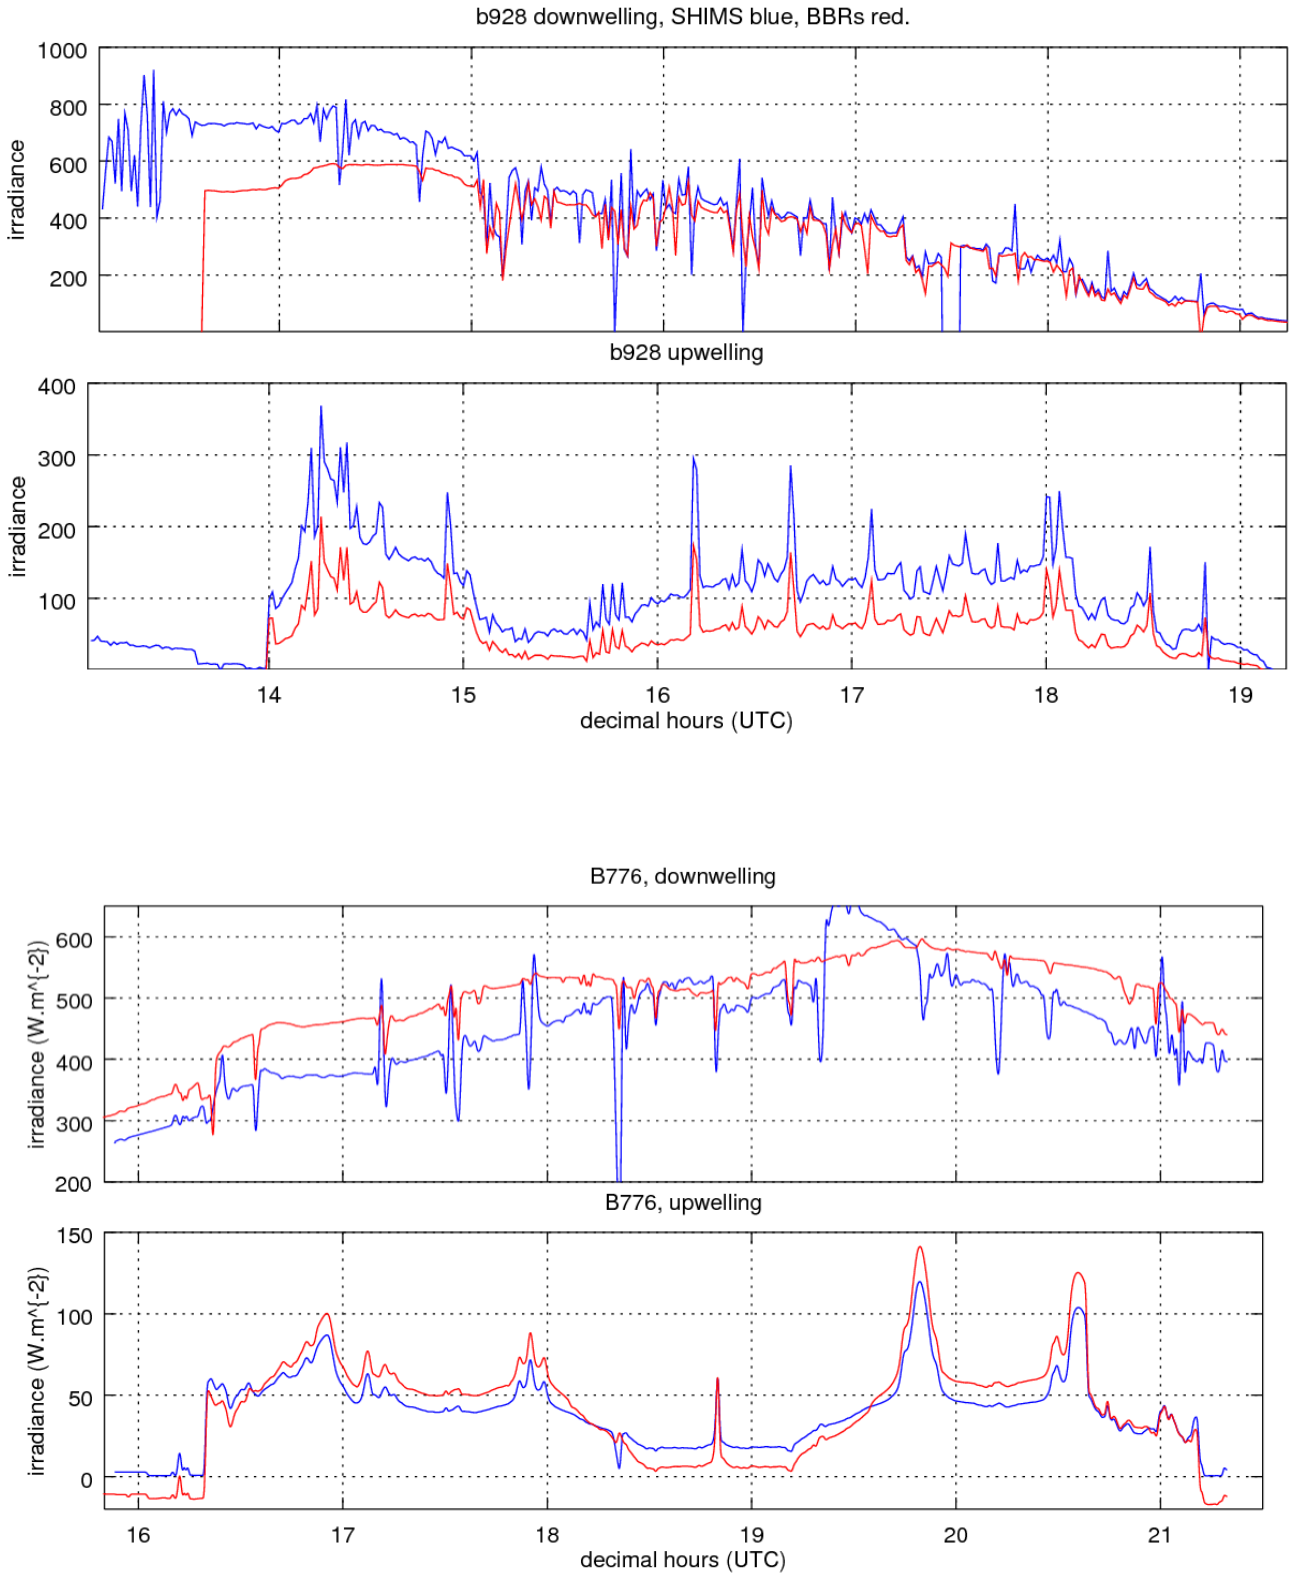

Figure 10.6. SHIMS-BBR comparisons for flights B928 (top) and B776 (bottom).

© Crown copyright 2017 Met Office

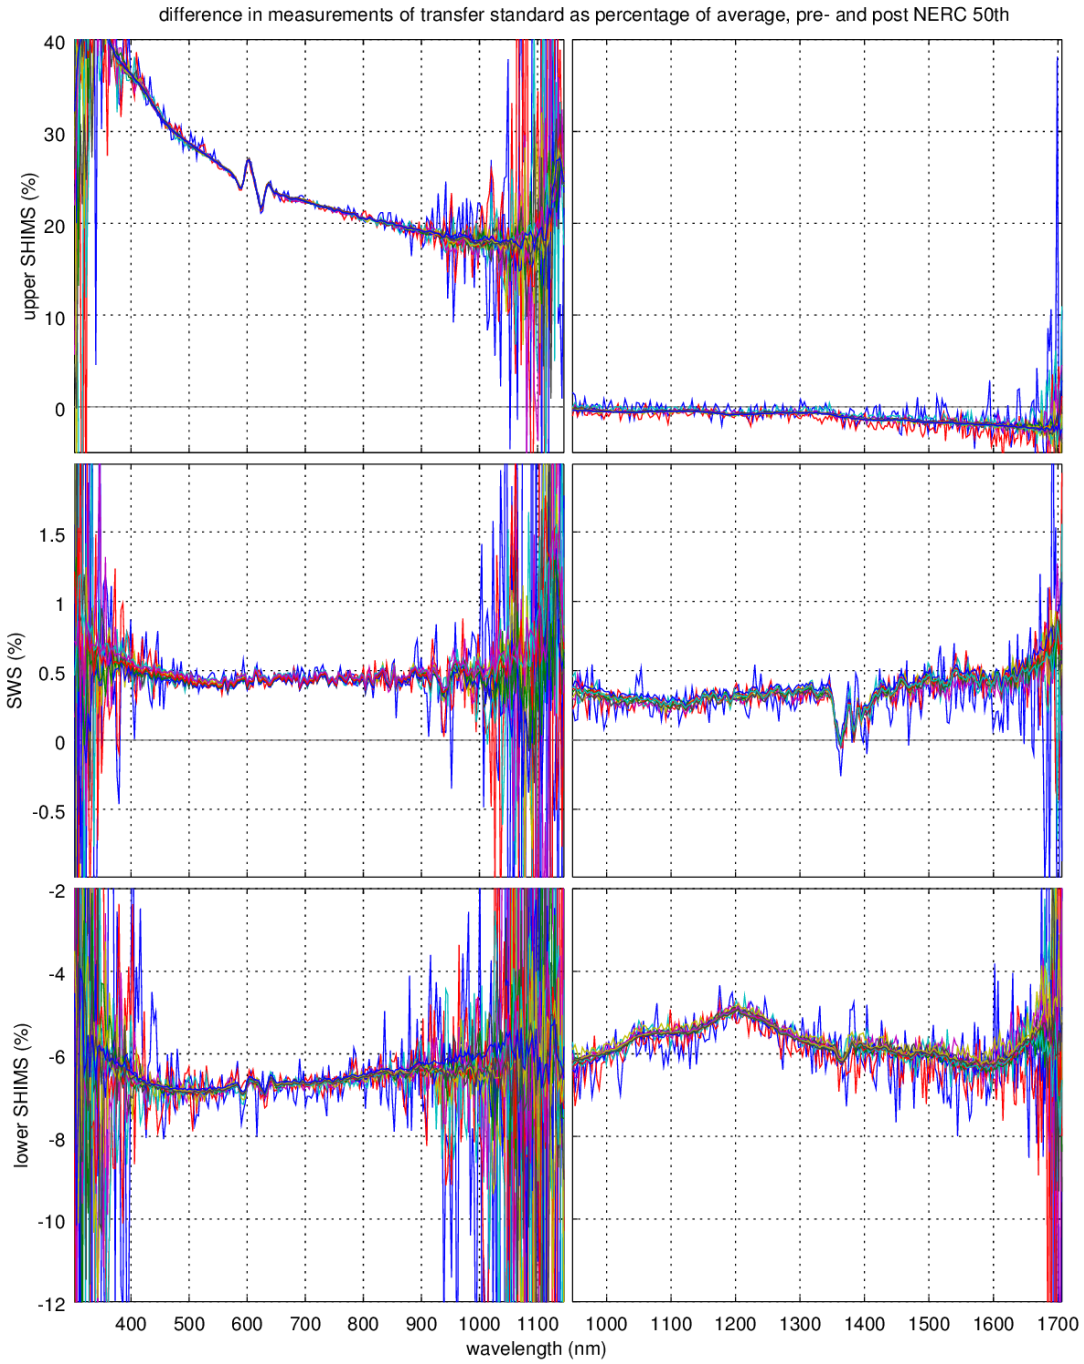

Figure 11.1. Comparison of SWS-SHIMS transfer measurements before and after NERC 50th.

© Crown copyright 2017 Met Office

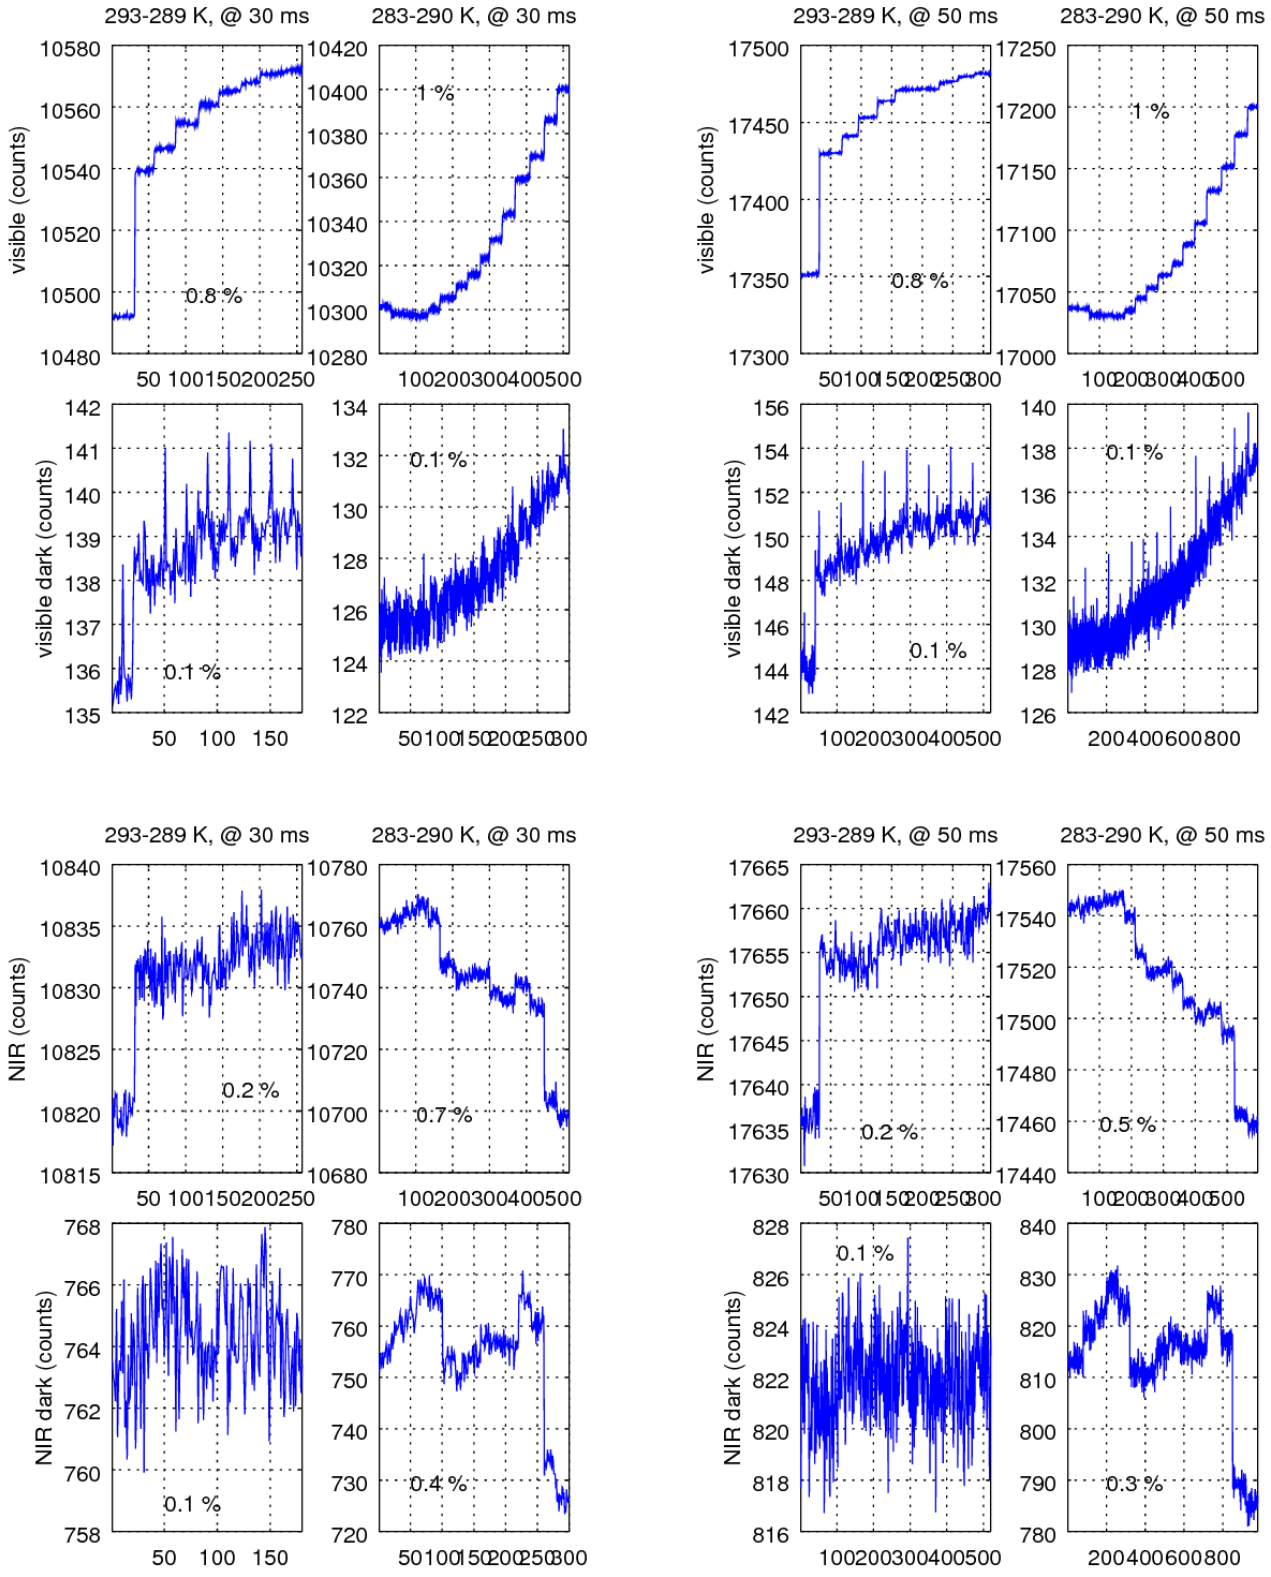

Figure 12.1. The effect of cool box temperature variations on SWS output - details in text.

© Crown copyright 2017 Met Office

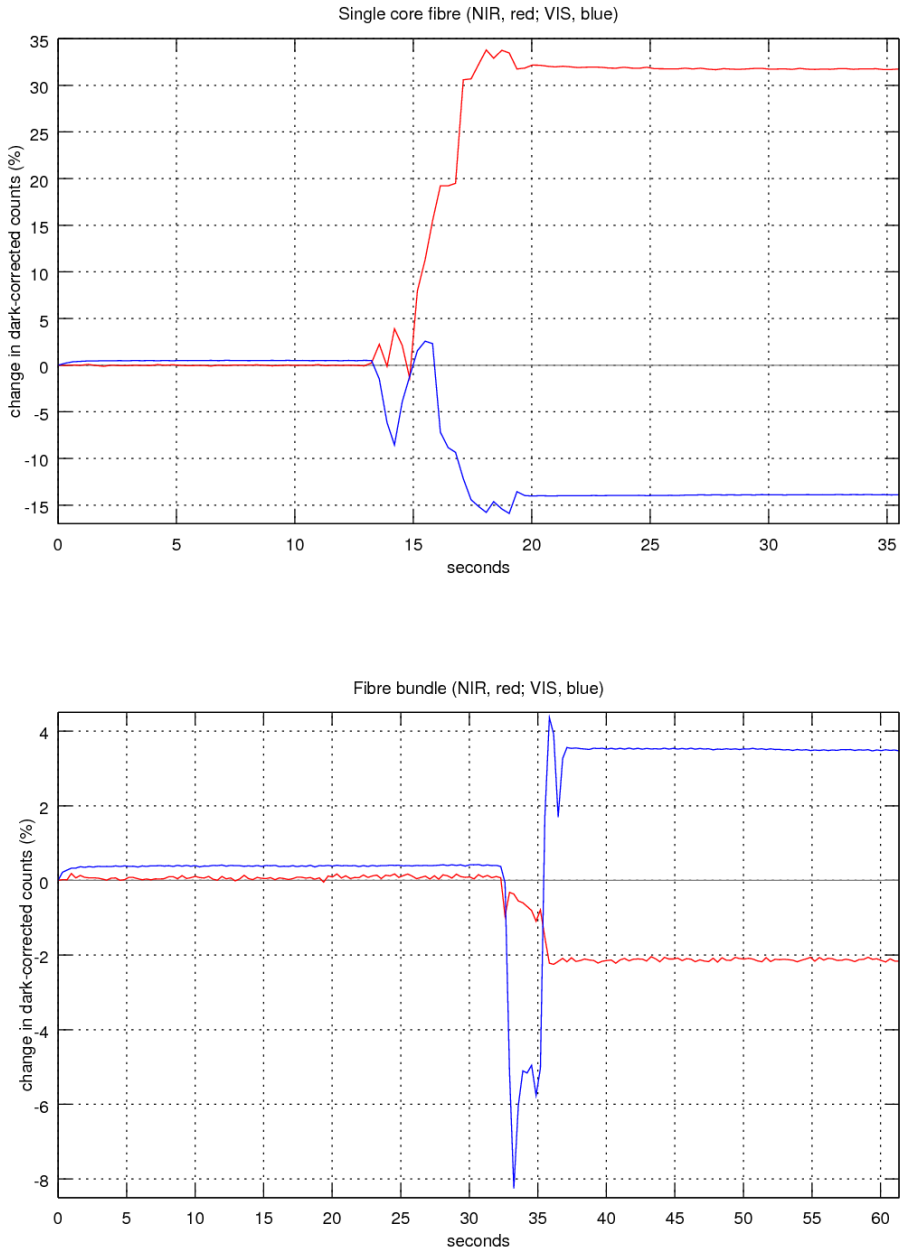

Figure 13.1. Changes in recorded signal resulting from the release of one coil of optical fibre.

© Crown copyright 2017 Met Office

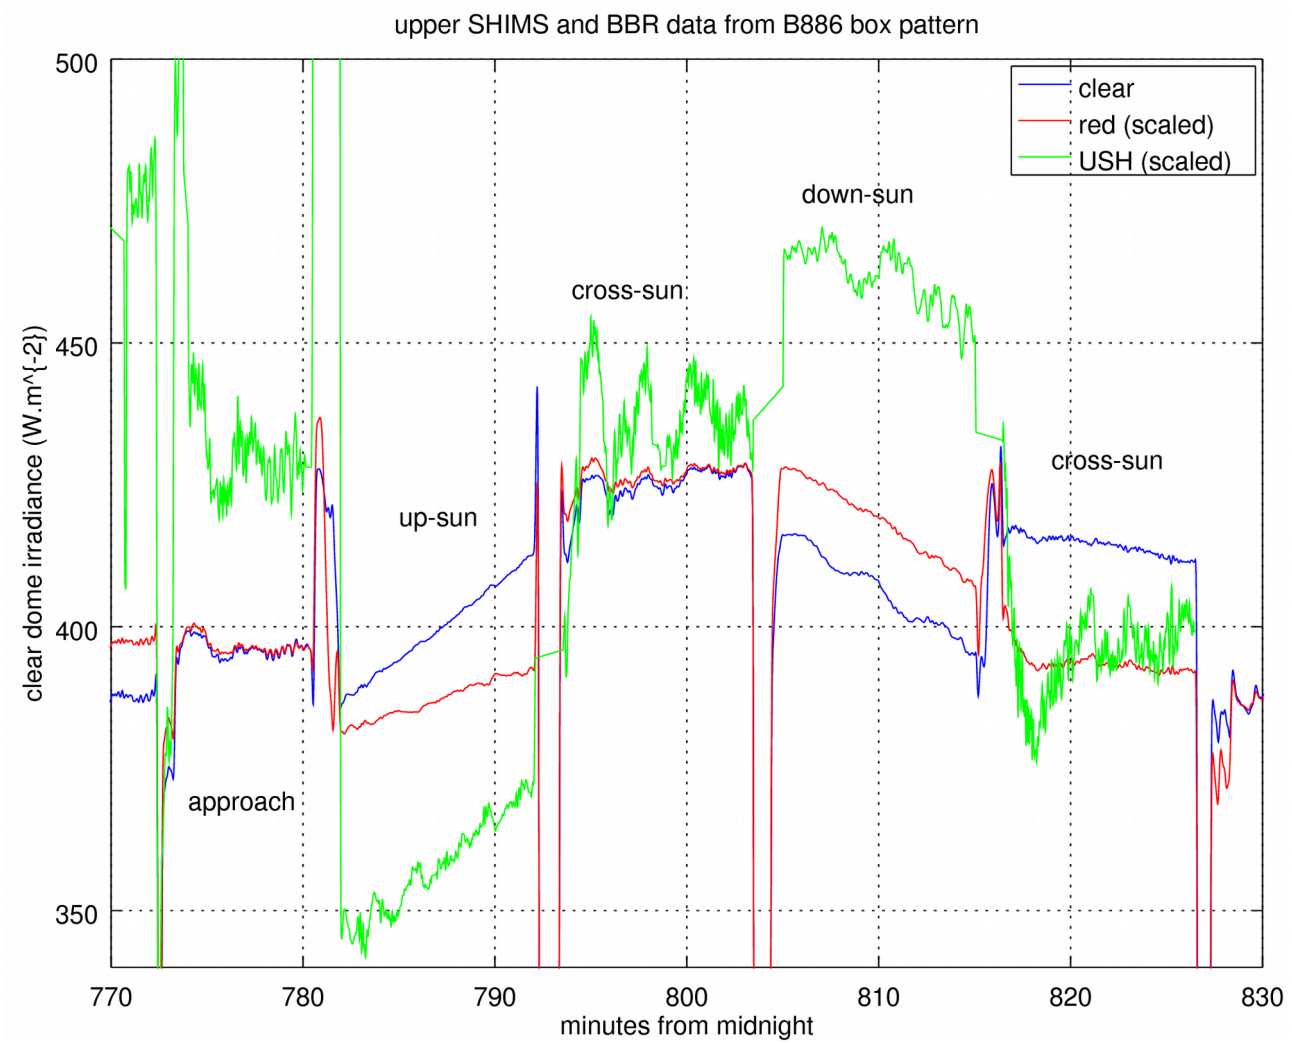

Figure 14.1. Overlay of upper SHIMS and BBR data, with red dome BBR in red, clear dome BBR in blue, and scaled upper SHIMS data in green.

© Crown copyright 2017 Met Office

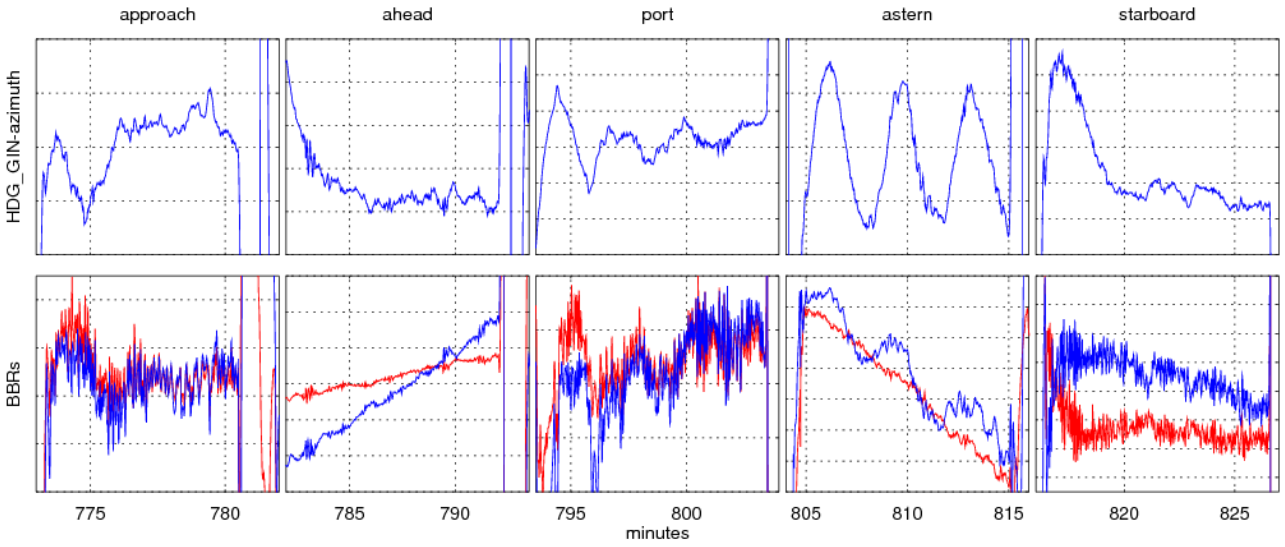

Figure 14.2. Time series of aircraft heading minus solar azimuth and BBR irradiances (all data scaled independently, so as to illustrate corellation between aircraft orientation and BBR irradiance fluctuations

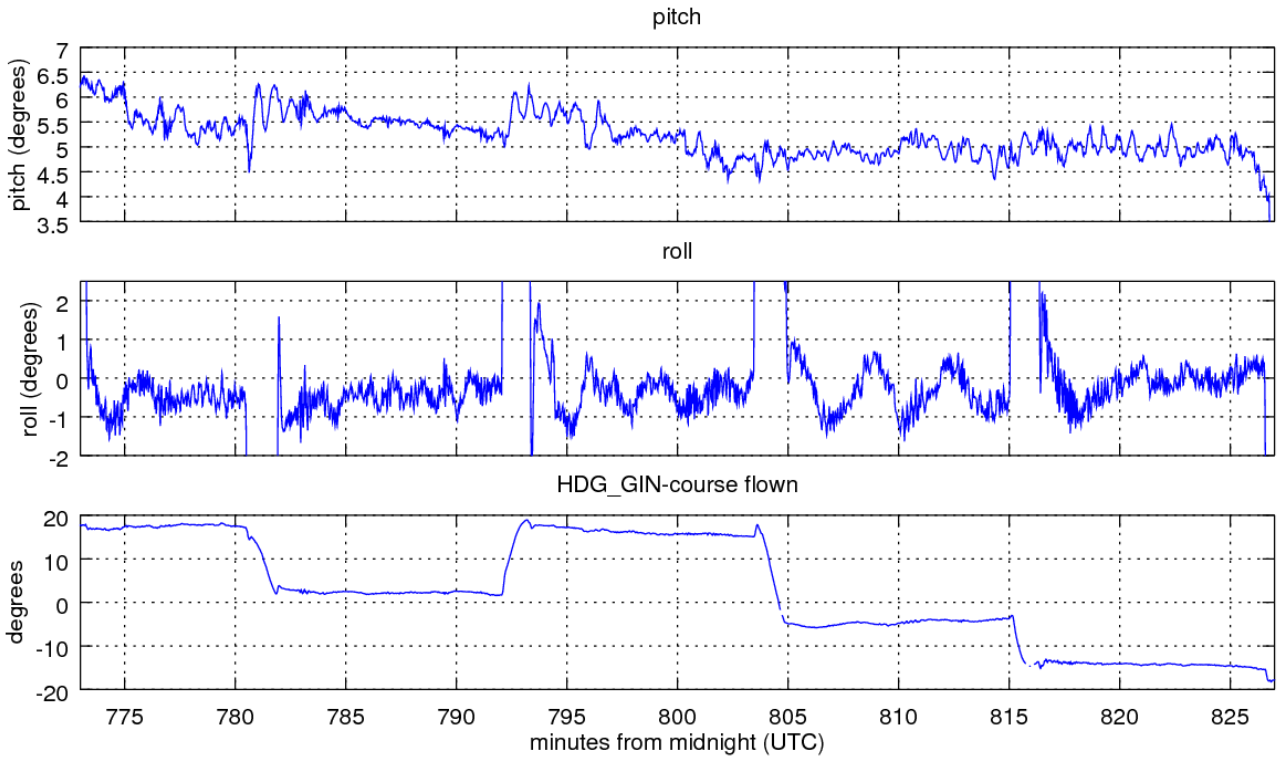

Figure 14.3. Time series of aircraft pitch, roll and yaw during B886 box pattern and approach.

© Crown copyright 2017 Met Office

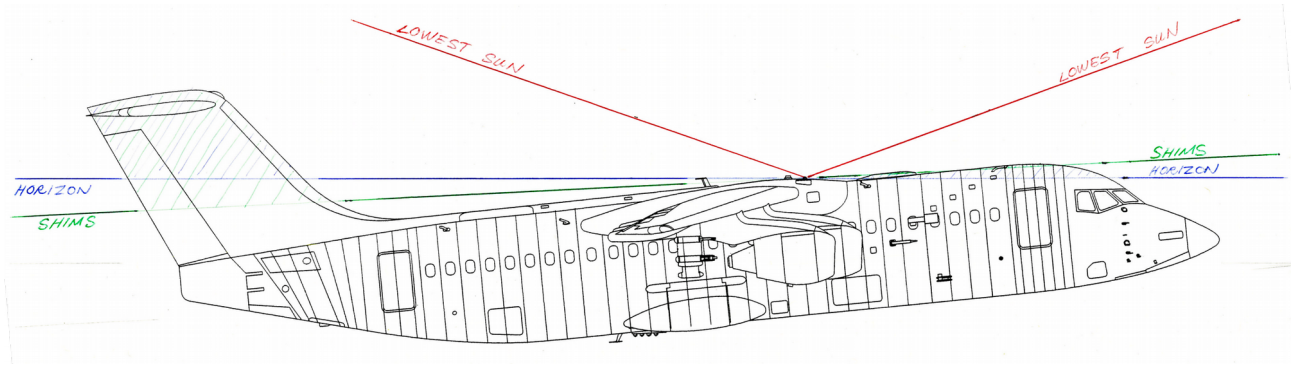

Figure 14.4. Starboard elevation of aircraft with approximate, true horizon added in blue, the plain of SHIMS and BBRs in green, and the lowest direct solar beam in red.

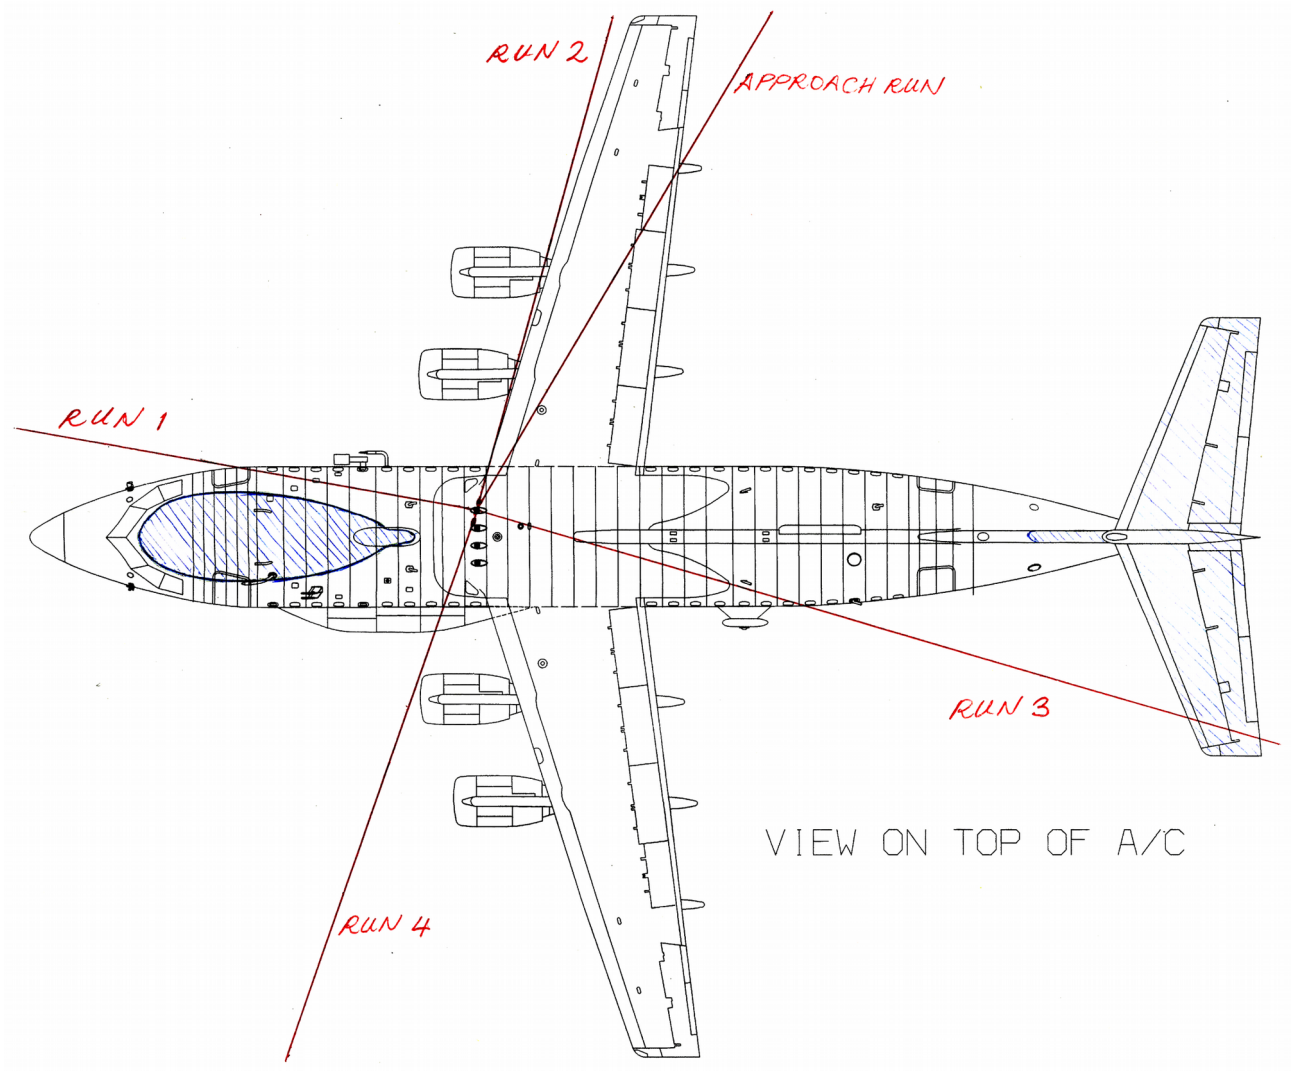

Figure 14.5. Plan of aircraft with approximate, areas above true horizon hatched in blue and the direction of the Sun on each run, viewed from SHIMS, added in red.

© Crown copyright 2017 Met Office

Appendix: End-of-section summaries..

2. B723 and B886 show evidence of poor calibration and port/starboard bias in SWS, but give confidence in IDL processing code despite inadequate filtering of bad data.

3. Much greater variability than expected, is seen in past gain curves.

4. Temporal linearity is not good, but issue can be avoided with appropriate calibration. Radiometric linearity is better in visible module than NIR,and at higher radiances. For low radiance scenes uncertainties of 3 % and 15 % might be assumed.

5. Accuracy of local calibration standards used for SWS and SHIMS are around 0.8 % and 1.75 %, respectively, over most of their ranges.

6. Stability of laboratory calibrations may be taken as ~7 %.

7. Zenith sky data from a side-by-side comparison of a Cimel sun photometer and SWS show very good agreement, although there may be some evidence of the direct solar beam being scattered into the instrument.

8. SWS measurements made outside and in the darkroom show very strong evidence of a severe stray light problem in SWS, although no evidence of similar issues has been found with the Cimel.

9. Failure to have SHIMS correctly seated in the transfer standard can cause large errors. Unseating of lower SHIMS due to equipment or personnel movement in the aircraft could well cause errors exceeding 25 %.

10. Through much of ICE-D, upper SHIMS appears to agree with the BBRs to around 10 %, but lower SHIMS reports close to twice the BBRs. This may, in part be explained by a poor lower SHIMS transfer calibration. B776 show better agreement in upwelling than downwelling, suggesting calibration rather than instrument problems.

11. Changes may occur in transit, but should be small, and be calibrated out.

12. The effects of temperature in the cooler varying by ±1 K are negligible.

13. Changes in SWS signal of ~1 % might be expected due to fibres bending as telescope rotates. SHIMS signal will not change in flight but variations may occur if fibres are disturbed by other work. These should be small, but this cannot be guaranteed, and changes may not be removed by transfer calibration.

14. In B886, unexplained variations in signal, of several percent, from upper SHIMS and clear- and red domed BBRs are seen which seem to show no definite correlation to each other, to Sun position or aircraft attitude.

© Crown copyright 2017 Met Office