An Overview of Quality Control at Herstmonceux

P. Gibbs, D. Benham, R. Sherwood, P. Standen, D. Walters, R. Wood
(NERC Space Geodesy Facility, Herstmonceux, UK)
&
G.M. Appleby
(NERC Space Geodesy Facility, Monks Wood, UK)

Abstract

SLR stations aim to produce high quality data with long-term stability. These goals can only be achieved by constant attention to detail and close monitoring of system performance. We describe here the collection of techniques employed at Herstmonceux to fulfil these aims: regular monitoring of hardware consistency; daily inspection of prediction trends; real-time control of data gathering; post-observation data screening; and continuous checks on meteorological data and epoch stability.

Introduction

The Herstmonceux SLR system comprises a mode-locked Neodymium-Yag laser of pulse length (FWHM) about 100ps, a SPAD detector and four Stanford SR620 interval timers. For all observing (calibration and satellites) we adhere to a strict single photon policy. Over many years we have investigated the ways in which various critical features of our system change in response to variations in both external influences (such as air temperature) and instrumental settings (such as dye strength). We have then used this knowledge to establish a set of working constraints within which we can be confident that our data are consistently of the highest quality. In this paper we discuss the complete range of monitoring processes and operational procedures that we use routinely to assess overall system performance and maintain long-term stability. A more detailed description of our analysis of our compensated SPAD appears elsewhere in these proceedings (Gibbs & Wood 2001). A full analysis of the inter-comparison of our four SR620 timers has been given by Appleby *et al* (1999).

Laser Stability

Overview

Variations in the behaviour of the laser have been found to affect measured range and RMS, beam pointing and the characteristics of the semi-train of pulses. In this section we outline the steps we take to stabilise laser performance and thus minimize any variations.

Laser repetition rate

We now use a repetition rate of 10Hz for all satellites and for all calibration targets. We find that tuning the laser for just this one rate makes it more stable and easier to maintain. But for high satellites this repetition rate is too fast for our sequential system. So for GLONASS, ETALON and GPS ranging we collect data at a rate of 5Hz whilst continuing to fire the laser at the standard rate of 10Hz and ignoring every other pulse. This allows us to use exactly the same hardware and software configurations as for low satellites.

Dye strength

Experience has shown that variations in dye strength affect all of the following: single-shot precision, measured range, laser-beam pointing, achieved laser repetition rate and semi-train quality. To minimize the potential errors from this source, we use a computer-readable photo-diode to monitor the strength of the dye (against an arbitrary scale) whenever the laser is in use, and have established fairly tight working limits for the dye strength. Care is always taken to calibrate the system immediately before and immediately after modifying the dye strength.
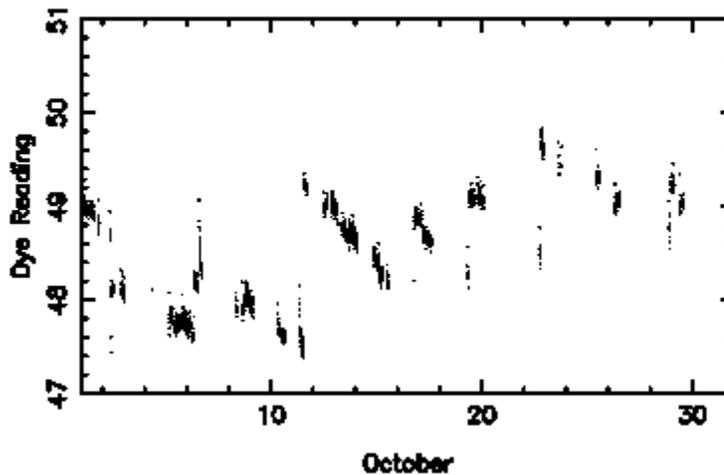

Figure 1. A plot of dye strength (arbitrary units) showing the gradual weakening of the dye with time and the upward jumps in strength when more dye is added. We find, interestingly, that the rate of decay is little affected by how much the laser is used.

Laser temperature

The laser has its own cooling system that keeps the temperature stable when the laser is in use. However human error and hardware failure have been known to occur. If the laser gets too hot it not only affects the range measurements (in a similar fashion to dye-strength variations) but also takes many hours to recover properly, during which time system drifts can take place. We therefore monitor the temperature of the laser and, if the temperature strays outside normal

working limits, warn the observer to take remedial action *and* set a software flag to prevent further use until the system returns to normal.
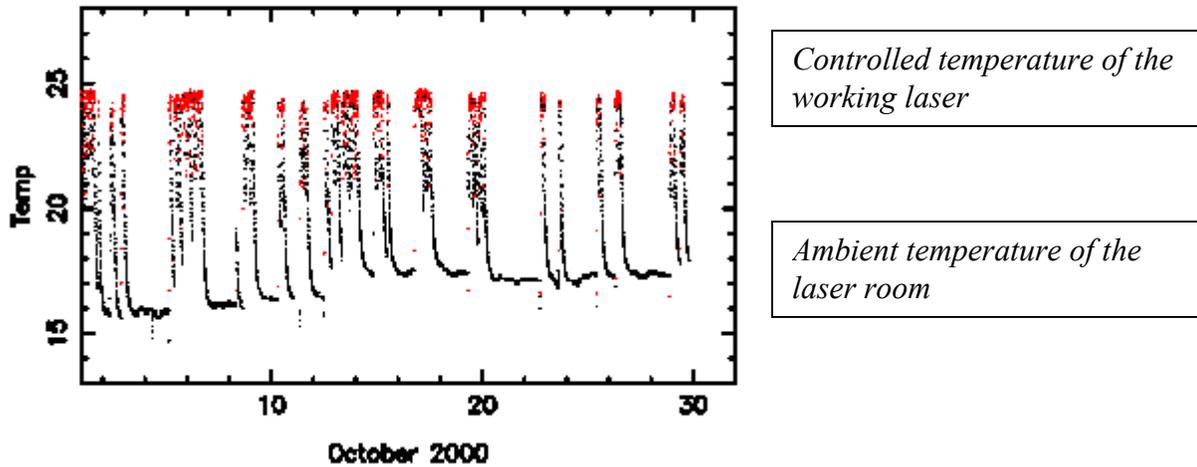

Controlled temperature of the working laser

Ambient temperature of the laser room

Figure 2. *A plot showing a typical series of laser temperatures. The red points indicate the laser is firing. The plot clearly shows that an upper limit of 25°C is reached when the system is behaving normally; an 'alarm' limit of 27°C has been set.*

Semi-train separation

Each time the laser is fired we emit a semi-train of 5 or 6 pulses separated from each other by about 9 nanoseconds. Measurement of the exact separation of the pulses has proved to be difficult and was originally estimated by differencing the calibration values derived using each individual pulse in the semi-train. However, we found that the separation value so derived was not consistent (at the few picosecond level) and depended on the particular hardware configuration. After much investigation we discovered that the cause lay in the behaviour of the SR620 interval counters themselves (see Appleby *et al* 1999). This method of determining pulse separation was thus ruled out.

In parallel with this investigation we attempted a direct physical measurement of the effective laser cavity length, including all the optical components. However, uncertainties in the optical properties of some of them limited the accuracy - although we did obtain a value very close to our subsequently adopted value.

We now use the orbit-fitting program **SOLVE** (Appleby & Sinclair 1992) to analyse the range measurements from each pulse in the semi-train, on a pass-by-pass basis. Having fitted an orbit to all the pulses, using an a-priori value for the inter-pulse separation, we are able to use the data for each individual pulse to determine a pass-averaged 'range bias' for each in order to improve the a-priori value. From a series of such values using observations of Lageos 1&2, Stella, Starlette and ERS-2 we are able to determine the value to a precision of a few picoseconds.

Once we have determined a value for the semi-train separation by this method, we are able to confirm it using our short-arc analysis system (Hausleitner *et al* 1998) that is daily used for quality assessment of the EUROLAS stations. In this approach observations from each pulse are

treated independently and the short arc analysis is used to determine any residual differences in range between the pulses. If our value for the inter-pulse separation is accurate, the derived differences will be close to zero. However if there is any error in the value it will be clearly seen in the results since each consecutive pulse increases the apparent range bias; pulse 2 has bias equal to the error, pulse 3 has twice the error, and so on.

We now monitor these measured inter-pulse separations on a pass-by-pass basis and, at reduction time, flag the observer if there is any deviation from the current value. After any work is done on the laser we monitor very closely to look for changes and are able to see changes at a level of about 10ps. A 10ps error in the semi-train separation value translates to a one-way range error of 1mm for the mean result based on the current relative contributions of each pulse to the overall data yield. Figure 3 shows an example of a separation change after laser work.
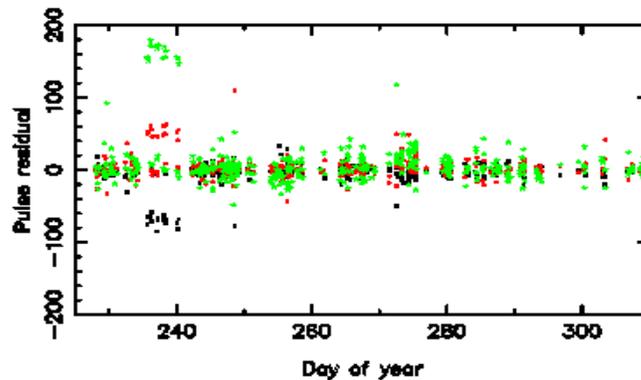

Figure 3. Residual plots after fitting an orbit for the first three pulses of the semi-train.

The data for days 235-240 show the results of a large change in the value of the pulse gap after a major laser service with both the pre and post service values applied. The plot also shows a small change in value just after day 270 when we did some minor adjustments to the laser. This shows up quite clearly an error of just 10ps in the semi-train pulse separation value.

SR620 Counters

Counter comparisons

Since 1994 we have used a suite of four SR620 interval counters and have made regular delay-dependent comparison checks of their performance against each other and also against an HP timer and a PPET. Figure 4 gives the results of the regular tests we have made during the year 2000 for the three counters named SRa, SRb and SRd, counter SRa being the counter exclusively used for standard operations. Clearly shown in the plots is the same repeat pattern for each device relative to Sra, determined to a level of precision of about 30ps.
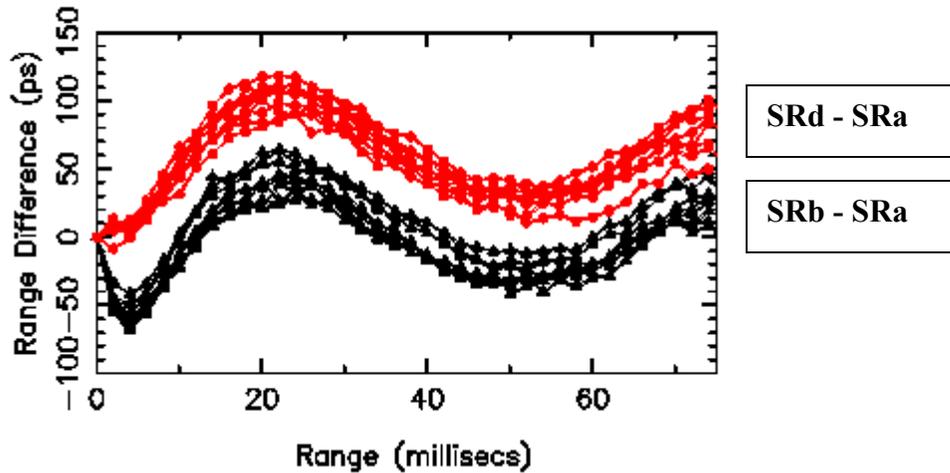
Figure 4. *Comparisons of SR620s B and D against A. Regular comparisons of C have not been possible because of problems with that device.*

Temperature control

Although we have not conducted a series of definitive tests to evaluate the effect, we are aware that the SR620s behaviour changes according to temperature. The few tests we have conducted after power cuts, when the laser room temperature drops, would indicate that there is a 1mm/°C change in our calibration value. More work needs to be done to quantify the changes, but we will continue to keep within the strict limits we have set, however small the effect may turn out to be.
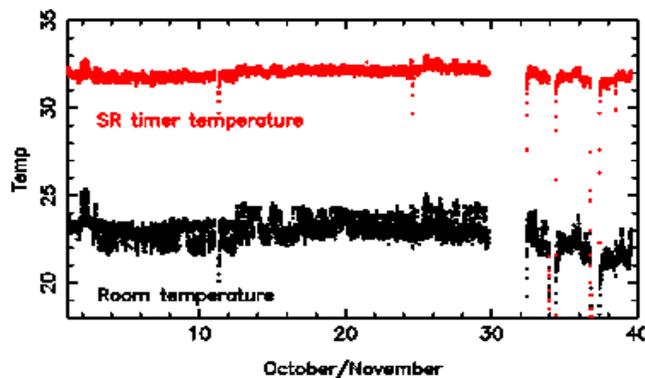

Figure 5. *Temperature plots of the air being drawn in and expelled from a SR620 timer. The clear drops in temperature are due to a series of power cuts.*

Optimising Observing

Daytime pointing

To correct the pointing errors of the telescope we have traditionally observed a series of stars at night and formed a pointing model for the telescope. However, as the telescope temperature changes during the daytime its pointing behaviour changes from that determined during the night. This sometimes makes daytime tracking for Lageos and the higher satellites very difficult. In an attempt to quantify this effect, we have carried out some experiments to observe a set of bright stars during the daytime. By introducing a red filter in front of our TV camera we are able to see stars to about 3^{rd} magnitude during the day and thus derive pointing model parameters appropriate to a variety of ambient temperatures. From this we hope we can model the telescope behaviour for a whole range of temperatures and derive a universal pointing model.

Daytime Coude alignment

Whenever any maintenance work is carried out on the laser, the alignment between the emitted beam and the optical axis of the emitter telescope is invariably lost, and adjustment of the Coude chain mirrors is required. This operation is best done at night when the beam is visible from light back-scattered by the atmosphere. However, the alignment again appears to be temperature dependent: as the telescope system heats up or cools down the beam alignment changes slightly, compounding the difficulty of daytime ranging.

To overcome this problem we need to be able both to see the beam in the daytime and to steer the beam easily. To see the beam we have used a CCD camera in our detector box in conjunction with frame-grabber software. By gating the frame grabber with the laser start signal we are able to take a single frame of short duration at the time the laser is "visible" as back-scattered light in the sky. To steer the beam into the alignment we have replaced the final mirror in our Coude chain with a steerable mirror controlled from the user interface. Thus during a pass it is possible for the observer periodically to view the alignment of the beam using the daytime TV, and apply any necessary corrections to the final mirror to re-centre it. This can be done very quickly (~30seconds) and is particularly effective for Lageos and the higher satellites.
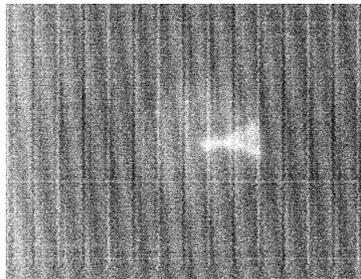

Figure 6. *The beam, centred in the iris, as seen by the frame grabber and CCD camera during daytime operation.*

Using the "best" predictions

For some time we have inter-compared the various sets of IRVs which are now available daily through ILRS prediction centres. Our recommendations for particular satellites are as follows:

Satellite	*IRV provider*
Champ, ERS-2	GFZ
GFO, Topex	Honeywell
GPS and Glonass	CODE/NERC
All other satellites	Honeywell *or* NERC daily *or* NERC weekly + time bias

Although this priority list is the default used by our prediction system, our observer may override to select any other IRV set and time bias value if they look better at the time.

Near real-time EUROLAS display

Using the EUROLAS near-real-time status display (Gurtner 1999) enables us to get information, virtually instantaneously, from other participating stations. For the more difficult satellites, and at times when predictions are poor because of high solar activity, successful observations from one station can be used for *immediate* time bias correction elsewhere. We strongly recommend it.

Real-time observing

Our policy is to observe as close to a single photon level of return as possible for all satellites (Gibbs & Wood 2001). Extensive tests have led us to believe that for a single-photon detector, even if it is compensated for high-level returns and at the expense of some loss of single-shot precision, the greatest consistency can be achieved only at this level. Poisson statistics indicate that for return rates less than about 20%, each return will, on average, be a single photon.
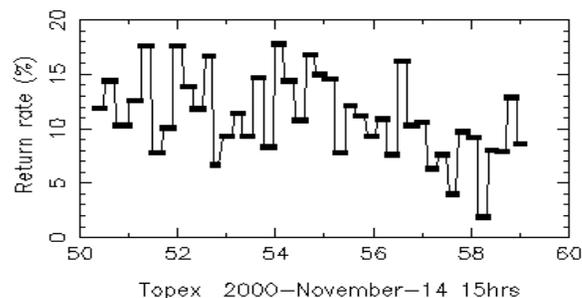

Figure 7. *Plot of the measured, real-time return rate during a typical pass of Topex. The jumps in rate are caused by changes in the value of neutral density inserted automatically into the return path to keep the return rate close to 12%.*

To keep to this ideal we must, in real time, detect and count the real returns amongst the noise events, count the number of pre-return noise events and from the results calculate the return rate.

Using this estimate of return rate the software automatically moves a neutral density filter wheel within the detection system so that we maintain, on average, a return rate of about 12%, well within the single photon regime. Because the return rate calculation is less reliable for steep tracks, we use the best predictions and apply additional time bias corrections to flatten the track of residuals, and close our working range-gate to eliminate as much noise as possible.
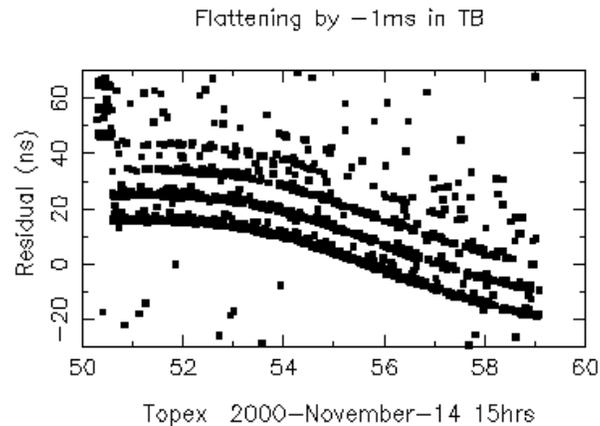
Flattening by −1ms in TB

Figure 8. *An example of track flattening in real time.*

Data processing

Overview

We have automated our reduction procedure as much as possible but we still require the user intervene at critical points in the reduction, and when the automatic system detects a possible source of error. We do **not** propose to go over to complete automation.

We have been gathering a large variety of statistics from the reductions over several years. When combined with a knowledge of the expected data distributions from our system for each satellite, have been able to define working limits within which the vast majority of observations should lie. These limits have been coded into the reduction system to provide a first-line data quality check, the actual range of acceptable values of the parameters of course being station-dependent. Data outside these limits are not sent to the data centres until they have been examined more closely; and only then is a decision made on whether to release them.

Current checks

Normal points: we reject any normal points with fewer than five returns. We do make exceptions, notably for GFZ-1 and CHAMP.

Orbital parameters: we check some of the parameters from our orbit fitting and flag anomalous values to provoke user intervention.

Semi-train pulse separation: a health and consistency check for the laser, particularly important if any work has been done on the laser.

Semi-train pulse size: we discard data from any pulses within the semi-train contributing fewer than 15 points to the whole pass.

Single photon policy: although we try to maintain single photon detection in real-time, inevitably some multi-photon returns are recorded; such data are rejected at the reduction stage.

Peak-LEHM (Leading Edge Half Maximum):   using subroutine DISTRIB (Sinclair 1996) we examine the distribution of range residuals in each satellite pass to determine the positions of the peak and of the LEHM. Each satellite has its own characteristic distribution of "Peak-LEHM" values and we check that the observed value lies between assigned upper and lower limits. Typical histograms of the Peak-LEHM values for Lageos1, Ajisai, ERS-2 and Starlette are shown in Figure 9.
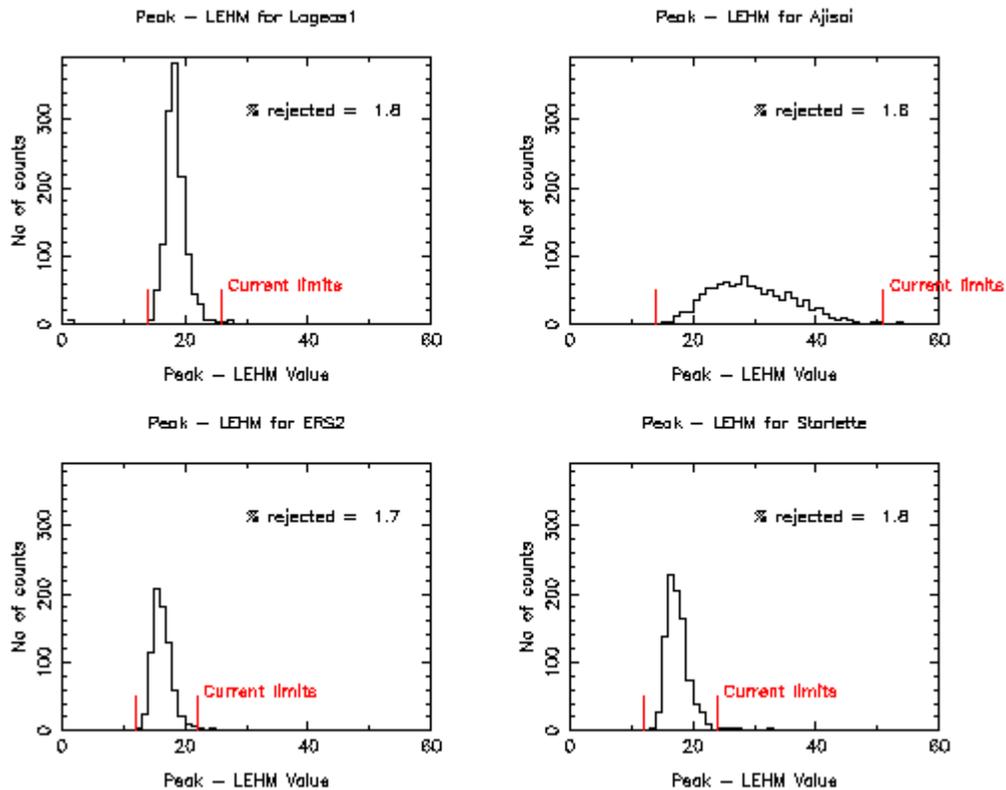

Figure 9 *Histograms of Peak-LEHM values for all passes of Lageos-1, Ajisai, Starlette and ERS-2. Based on these plots we have set upper and lower limits for Peak-LEHM. This is a check on the shape of the residual distribution. Not only do we insist that the whole pass falls within these limits, we also insist that each individual pulse from the semi-train produces residuals that fall within the limits.*

RMS: similarly each satellite has its own characteristic distribution of RMS values and we again adopt upper and lower limits for acceptability. Typical histograms of the RMS distribution are shown in Figure 10.
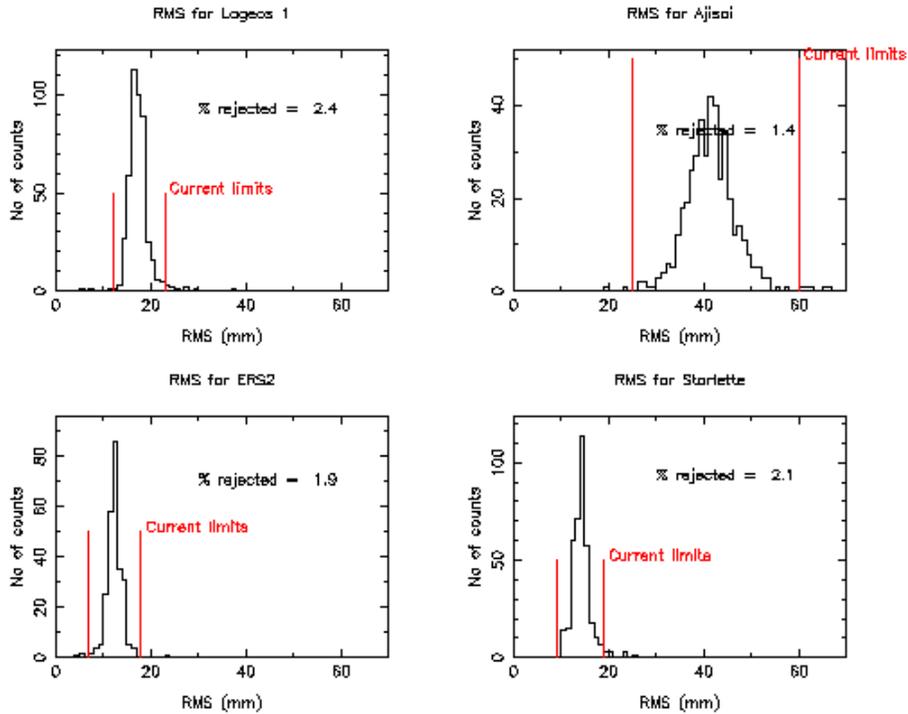

Fi*gure 10* *Histograms of RMS values for all passes of Lageos 1, Ajisai, Starlette and ERS-2. The RMS for individual satellites varies a lot for a single photon system due to the signature of the satellites, the effects of which are now well understood. Testing the actual value from each pass provides an overall health check on the pass.*

Orbital analysis

Long arc analysis

Even with all the above checks, we still generate the occasional bad data point. The six-day-arc analysis (Hausleitner, Appleby & Sinclair 1998) that we carry out automatically every day enables us to check for errors in our Lageos data, and is sensitive at a level of about 2 cm.

Short-arc analysis

We have used this method to analyse our multi-timer system and our semi-train inter-pulse gap. It could also be used to analyse multi-detector and multi-colour systems at some later stage. We are currently working to extend both the long-arc and short-arc analyses to include other groups of stations and other satellites.
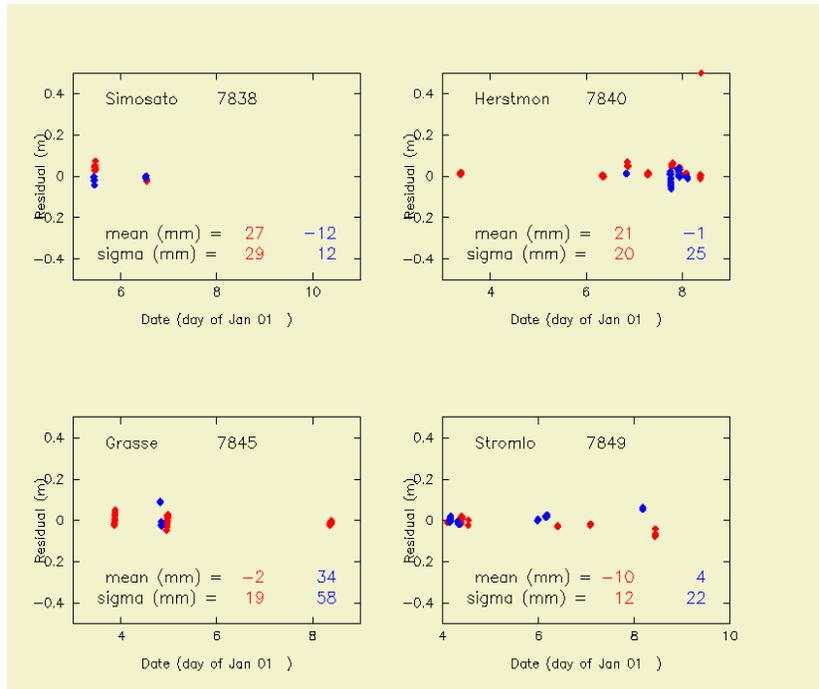
Figure 11. *Range residuals for four stations from a six-day global orbital fit to Lageos-1 (red) and Lageos-2 (blue). The single outlier from Herstmonceux is apparent (top right).*
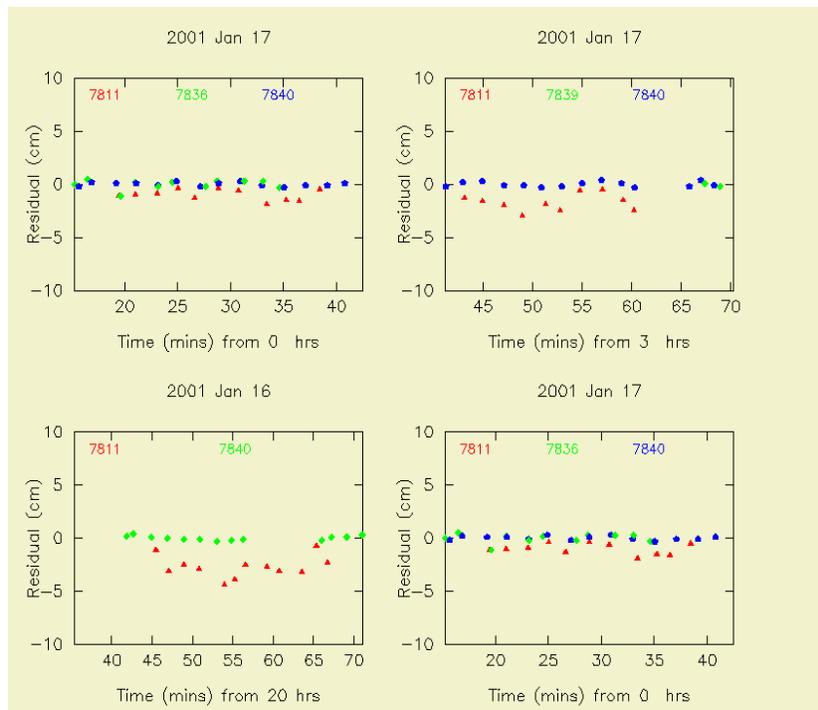

Figure 12. *Range residuals from four short-arc orbital solutions using single passes of Lageos-1 obtained simultaneously from two or more EUROLAS stations.*

Calibrations

Ground targets

We currently have five local ground targets. Three targets are to the East at a distance of ~600m: one flat board and two retro-reflectors. Two targets are to the West at ~120m: both are retro-reflectors. Two of the Eastern targets and one of the Western targets have been accurately surveyed on two occasions with ties to the telescope and various other reference marks, including the IGS marker. We now regularly make calibration measurements to all five targets although we always use the same target for the reduction of our published data. We have also experimented with an "internal" target mounted on the telescope.

The calibration targets are also used to investigate various other aspects of system performance: for example by carrying out long series of calibrations whilst varying system parameters of interest such as dye strength, C-SPAD temperature, return rate *etc*.

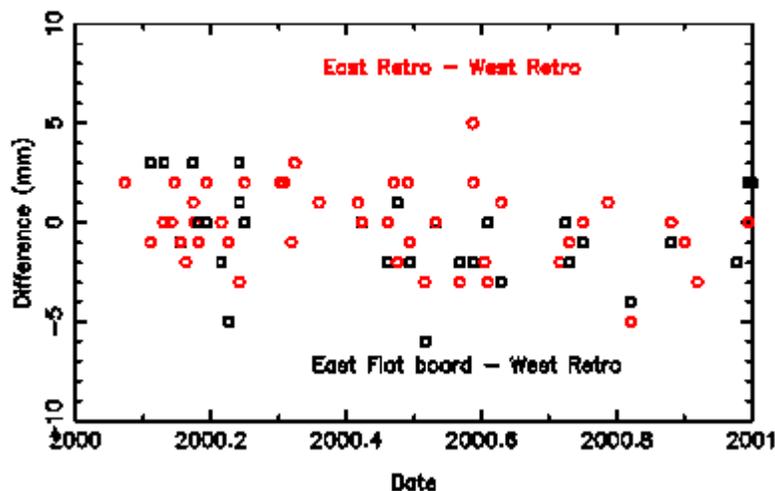

Figure 13 *Comparison of calibration ranges for the 3 surveyed targets.*

Calibration variations

There are many factors that can cause changes in the system calibration value. At Herstmonceux one of the main contributors is the temperature effect on the C-SPAD detector (Gibbs & Wood 2001). To minimise this effect we try to take a calibration measurement before and after each pass, and always at least once an hour when passes are bunched.

Figure 14 shows measured variation in system calibration value as a function of time and detector temperature when using the C-SPAD uncompensated channel. The compensated channel behaves in exactly the same way except that the effects apply in the opposite sense (Gibbs & Wood 2001).
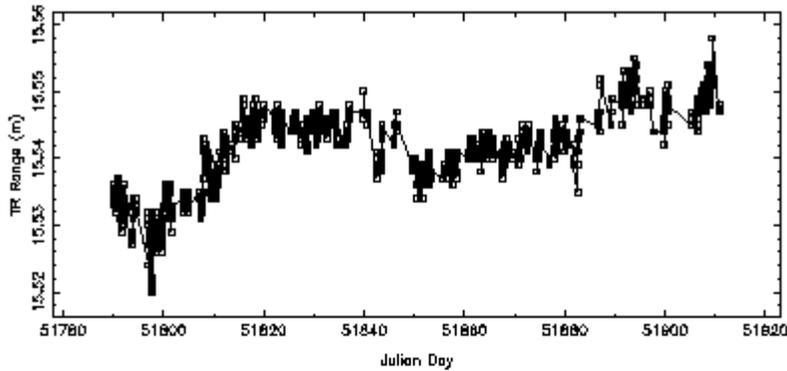
Figure 14a Calibration ranges taken over 3 months showing a 3-4 cm variation and a diurnal effect.
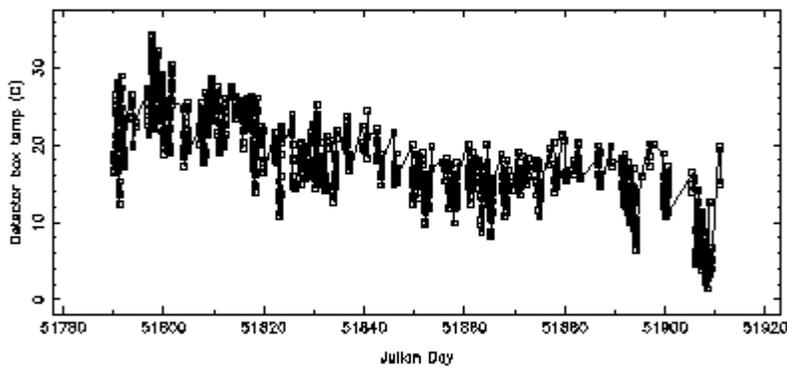

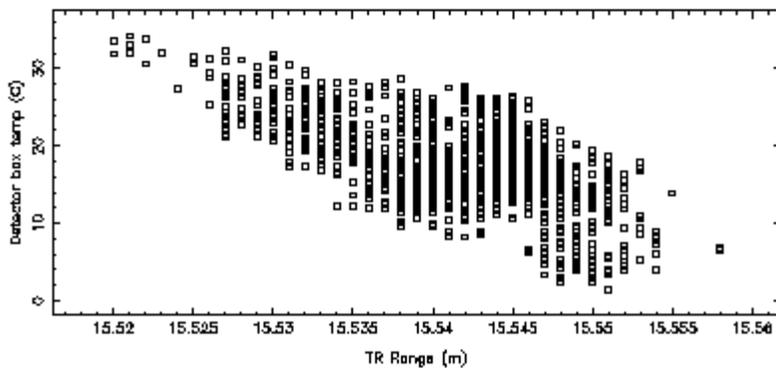*Figure 14b* The temperature of the C-SPAD casing within the detector box as a function of time.

Figure 14c Plot showing the effect of temperature on the measured range to the calibration target.

Meteorological data

Data are gathered every few minutes from each of two separate temperature, pressure and humidity devices and significant differences between them are flagged in real time. We also run parallel software that generates an alarm condition if the meteorological data gathering ceases for any reason. Once a week our readings for that week are checked against those of the local UK Meteorological Office; and on a daily basis the pressure devices are checked against a mercury barometer.

System Clock

The primary system frequency at 10MHz is derived from a dedicated GPS receiver and disciplined by a high quality quartz oscillator. Epoch is read from an in-house counter board driven by the 10MHz signal and synchronised to UTC using a GPS 1-pps. The integrity of the counter board is checked against GPS before and after each pass, or once per hour in non-observing times. An alarm is generated automatically if epoch read from the counter board differs by 0.5 microseconds or more from the GPS.

Conclusions

We believe that **self-monitoring:**
- is better than relying on feedback from analysis or data centres;
- helps you to understand the "personality" of your system; and
- pinpoints problems as they happen.

Automated recording of time-series of relevant data quantities pays huge dividends for both short-term and long-term quality control.

The better the prediction, the better the observation and thus the better the final product.

It is **quality** not **quantity** that really counts – we always have the words of Andrew Sinclair ringing in our ears "**No** data is better than **poor** data".

References

G.M. Appleby & A.T. Sinclair, *Formation of on-site Normal Points*, Proceedings of the 8[th] International Workshop on Laser Ranging Instrumentation, Annapolis 1992.

G.M. Appleby, P. Gibbs, R.A. Sherwood, & R. Wood, *Achieving and maintaining sub-centimetre accuracy for the Herstmonceux single-photon SLR facility*, Proceedings of Laser Radar Ranging and Atmospheric Lidar Techniques II, Florence 1999.

P. Gibbs & R. Wood, *The C-SPAD as a single-photon detector,* these proceedings, Matera 2000.

W. Gurtner, *Near-realtime status exchange,* SLRMAIL 0372, 1999.

W. Hausleitner, G.M. Appleby & A.T. Sinclair, *Rapid quality checks within the EUROLAS cluster,* Proceedings of 11[th] International Workshop on Laser Ranging, Deggendorf, 1998.

A.T. Sinclair, *Survey of SLR data pre-processing by single-photon stations,* SLRMAIL 0008 and 0014, 1996. (ftp://cddisa.gsfc.nasa.gov/pub/reports/slrmail/).