INSTITUTE OF PHYSICS PUBLISHING CLASSICAL AND QUANTUM GRAVITY

Class. Quantum Grav. **21** (2004) S1711–S1722 PII: S0264-9381(04)79351-X

Calibration of the dual-recycled GEO 600 detector for the S3 science run

M Hewitson[1], S Babak[4], R Balasubramanian[4], K Danzmann[1], H Grote[1], G Heinzel[1], J Hough[2], H Lück[1], M A Papa[3], J R Smith[1], K A Strain[2], H Ward[2], B Willke[1] and G Woan[2]

[1] Max-Planck-Institut für Gravitationsphysik (Albert-Einstein-Institut) und Universität Hannover, Außenstelle Hannover, Callinstr. 38, 30167 Hannover, Germany
[2] Physics and Astronomy, University of Glasgow, Glasgow G12 8QQ, UK
[3] Max-Planck-Institut für Gravitationsphysik, Albert-Einstein-Institut, Am Mühlenberg 1, 14476 Golm, Germany
[4] Department of Physics and Astronomy, Cardiff University, PO Box 913, Cardiff, CF2 3YB, UK

E-mail: martin.hewitson@aei.mpg.de

Received 19 April 2004
Published 28 September 2004
Online at stacks.iop.org/CQG/21/S1711
doi:10.1088/0264-9381/21/20/014

Abstract

The GEO 600 interferometric gravitational detector took part in an extended coincident science run of the LIGO Scientific Collaboration (S3) that started in November 2003. GEO had recently been upgraded to be the first large-scale fully suspended dual-recycled interferometer in the world and was in the early stages of commissioning in this configuration. In order to prepare the GEO 600 data for the possible extraction of science results and for exchange between analysis groups of the gravitational wave community, the data need to be accurately calibrated. An online, time-domain calibration scheme that was initially developed to calibrate the power-recycled GEO 600 configuration, has been extended to cover the significantly more complicated case of calibrating the dual-recycled interferometer, where the optical response of the instrument is much more difficult to measure and calibrate out online. This report presents an overview of this calibration scheme as it was applied to calibrating the GEO S3 science run data. In addition, results of the calibration process are presented together with some discussion of the accuracy achieved.

PACS numbers: 04.80.Nn, 95.55.Ym, 95.75.Kk

(Some figures in this article are in colour only in the electronic version)

 © 2004 IOP Publishing Ltd Printed in the UK

1. Introduction

The GEO 600 interferometric gravitational wave detector is located near Hannover, Germany and is currently in the final stages of commissioning. The optical layout of GEO 600 is based on the Michelson interferometer with two additional mirrors that are used to enhance the sensitivity of the detector. A power-recycling mirror forms an optical cavity in which the circulating power can build up thus improving the shot-noise performance of the detector. Another mirror, the signal-recycling mirror, is placed at the output port of the detector, leading to a resonant enhancement of potential gravitational wave signals in a particular (chosen) frequency band.

Prior to installing the signal-recycling mirror, GEO was operated for some time as a power-recycled interferometer (see [1] for details of the power-recycled operation of GEO). During this time, GEO took part in an extended science run called S1, details of which can be found in [2]. The most recent work at GEO was concerned with the installation of the signal-recycling mirror followed by the development and debugging of a new locking scheme suitable for operating the detector in a dual-recycled configuration. Details of signal recycling can be found in [3–6]. The locking scheme for the dual-recycled configuration of GEO, together with other steps that have led to the operation of the detector, is described in [7].

During November 2003, GEO participated for one week in a long-term coincident science run with other gravitational wave detectors around the world; this science run is termed S3. After this first week of participation, GEO was taken offline for further commissioning. After an improvement of almost a factor 10 in sensitivity was achieved, GEO rejoined the science run and took science quality data for a further two weeks (see [8]). The two parts of S3 participated in by GEO are termed S3 I and S3 II within the GEO collaboration.

The calibration of the dual-recycled GEO 600 for these two parts of the S3 science run is an essential part of preparing the data for distribution to analysis groups who hope to detect gravitational wave signals or at least set upper limits on the strengths of signals from possible known sources. Calibrating the dual-recycled GEO is significantly more complicated than calibrating the power-recycled configuration, however, the methods developed for calibrating the power-recycled detector can be carried forwards and extended to the case of the dual-recycled detector. The principles of the calibration scheme developed for the power-recycled detector are discussed in [9].

2. Principles of calibrating the dual-recycled GEO 600

Extending the calibration scheme used to calibrate the power-recycled GEO 600 configuration involves the development of a parametrized model of the more complicated response of the detector to gravitational wave signals. This model of the instrument response is used, together with an optimization algorithm, to estimate, online, the response of the detector to gravitational wave signals. In addition, a few of the techniques that were employed in the power-recycling calibration scheme were replaced with more sophisticated methods. A brief overview of the power-recycled calibration scheme and its extension to the calibration of the dual-recycled interferometer is given below; its application to the S3 science run follows in the subsequent sections.

2.1. Calibration of the power-recycled GEO

Calibration of the power-recycled GEO configuration was performed using a time-domain signal-processing pipeline that is described in [9] and shown implemented for the S1 science

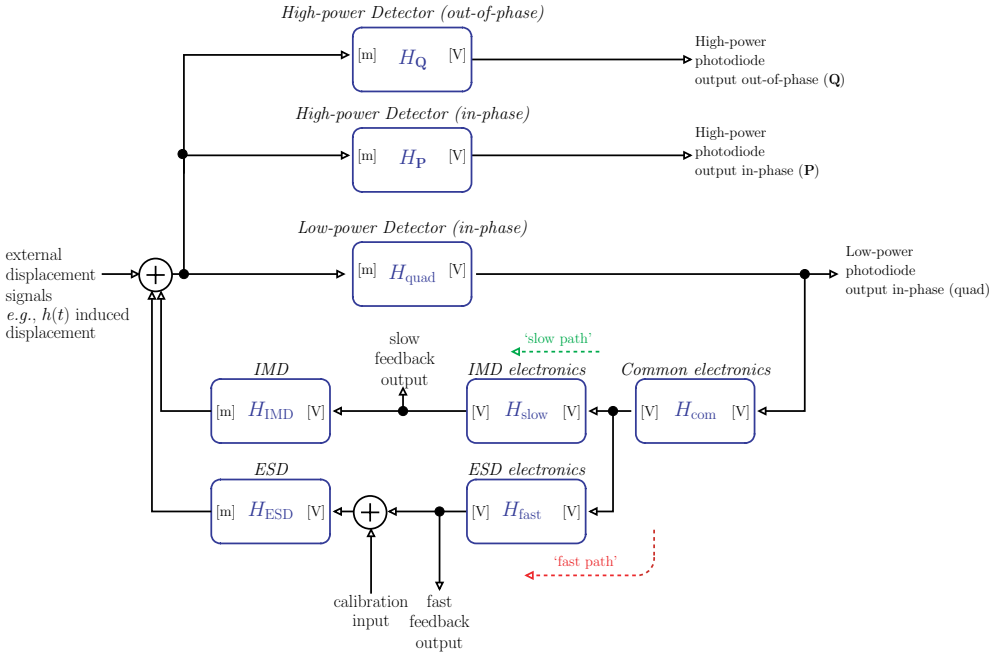

Figure 1. A schematic of the length control servo and differential displacement sensing set-up of GEO for S3. The recording points of some relevant signals are also indicated.

run in [2]. The basic idea is to correct for any transfer functions that convert differential arm-length changes of the instrument to the detector output signal. In the case of GEO, there are two main effects: the optical response of the detector and the Michelson length-control servo. The optical response function for the power-recycled GEO is independent of frequency (at frequencies in the detection band) but varies with time; the Michelson length-control servo is a frequency-dependent gain that comprises the (assumed fixed) responses of the electronic servo filters and the actuator response function, together with the optical response. By continually inducing a known displacement at a frequency well above the unity gain point of the Michelson length-control servo, the optical gain of the instrument can be estimated online, at regular intervals. Calibrating the output of the detector by correcting for this optical gain gives a signal that at high frequencies (outside the Michelson loop bandwidth) represents any differential arm-length changes of the instrument. By feeding the output of the detector through model time-domain filters of the split-feedback electronics and actuators (see figure 1), two more signals can be generated that represent the feedback signals of the length-control servo. When these calculated feedback signals are added to the signal got from calibrating out the optical gain of the instrument, we correct for the effect of the Michelson length-control servo. The resulting signal now represents differential arm-length changes at all frequencies where the detector model is valid (from 50 Hz to 6 kHz) and can be converted to a strain signal by dividing by the length of one arm of GEO, 1200 m.

2.2. Extending the calibration scheme to the dual-recycled detector for S3

Extending the calibration scheme to the dual-recycled detector involves a detailed model of the (now) more complex optical response of the detector. For the power-recycled GEO, the transfer function from differential arm-length changes (and hence gravitational wave strain) to the detector output signal was flat in frequency. In the dual-recycled configuration, this is no

longer the case and the optical transfer function must be modelled as a frequency-dependent gain function (see section 2.3.1). In addition, the demodulation of the output photodiode signal is now also frequency dependent such that, for any one particular demodulation phase, gravitational wave signals will be spread between the two demodulation quadratures. By choosing this demodulation phase, gravitational wave signals of a particular frequency can be maximized in one quadrature and minimized in the other; for all other frequencies, the ratio between the two quadratures is different. In general, the quadrature that maximizes the signal content at a chosen frequency is termed the **P** or 'in-phase' quadrature, while the other quadrature is termed the **Q** or 'out-of-phase' quadrature. A detailed discussion of the principles involved in calibrating the dual-recycled GEO detector using time-domain methods is given in [10].

In S3 I, the demodulation phase was set such that the **P** quadrature had maximum signal content at dc. For simplicity, the S3 implementation of the dual-recycled GEO calibration scheme concentrated on calibrating only the **P** quadrature. In S3 II, the demodulation phase was changed to give the maximum signal content in **P** around the peak sensitivity of the instrument (1000 Hz).

Other improvements in filtering and up-sampling methods common to both the power- and dual-recycled calibration schemes are discussed in section 3.1 where details of the software implementation are given.

2.3. Detector details and set-up for S3

The calibration scheme relies heavily on a sufficiently detailed and accurate model of the detector subsystems that directly affect the response of the detector to gravitational waves, namely, the main optical subsystem (the dual-recycled interferometer) and the Michelson length-control servo. These subsystems are coupled in the control scheme of GEO. A model of the detector, including the optical system, the signal readout scheme and the length control servo, is shown in figure 1.

2.3.1. The optical response of the detector. The detector has an optical response that can be modelled as three separate elements. The in-loop optical response (labelled H_{quad} in figure 1) represents the transfer function from differential arm-length changes to detector output signal as seen by a low-power quadrant photodiode demodulated in only one quadrature. This photodiode output signal provides the error-signal that is filtered and fed back to the length-control actuators to keep the detector at its optimal operating point; it is also used in the automatic alignment system of GEO. The other two optical elements (labelled $H_{\mathbf{P}}$ and $H_{\mathbf{Q}}$ in figure 1) represent the equivalent transfer function as seen by a high-power photodiode that is demodulated in both quadratures, **P** and **Q**. As was already mentioned above, for the S3 science data, calibration focused only on the **P** quadrature for simplicity. In principle, in order to properly correct for the effect of the length-control servo, we would need to determine the optical response of the instrument as seen by the low-power quadrant photodiode. In these initial experiments, however, this response was assumed to be the same as the **P** response with the exception of an overall gain difference which was measured prior to the science run and assumed fixed thereafter. Some discussion about the effect of this assumption is given in 4.1.

The optical response of the instrument can be modelled using a simple pole/zero model. This response can be written as

$$H_{\mathbf{P}}(s) = G_{\mathbf{P}} \frac{(s - z_{\mathbf{P}})}{(s - p_{\mathbf{P}})(s - p_{\mathbf{P}}^{*})}, \tag{1}$$

where '$*$' denotes complex conjugation, $z_{\mathbf{P}}$ is the zero in the **P** response, $p_{\mathbf{P}}$ is the complex

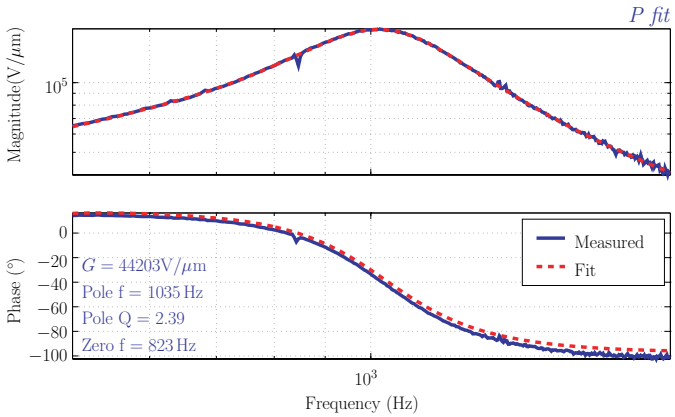

Figure 2. A transfer function of the optical response of the detector as seen by the high-power diode **P**. A pole/zero fit to the magnitude is also shown together with the recovered parameters. This transfer function measurement was made prior to S3 II using a standard spectrum analyser; the calibration lines were removed during the measurement. The fit was made using a simple nonlinear least-squares algorithm in MATLAB. Such a fit provides a base from which to check the optical parameters that are recovered using the online method of measuring the optical transfer function (see section 3).

pole in the **P** response, and $G_{\mathbf{P}}$ is an overall gain; this response can be parametrized more conveniently using a frequency/Q representation of the pole and zero.

It is possible to measure the optical response of the instrument to differential arm-length changes by injecting white noise into the electrostatic actuators and measuring the transfer function between this injected noise and the detector output signal measured at **P**. Correcting this measurement for the response of the electrostatic actuator, which is modelled as the response of a simple pendulum to a force applied to the suspended mass, gives the measured optical response shown in figure 2. This example measurement was made prior to S3 II and shows a fit to the data using the model described by equation (1). The fit is made to the magnitude using a nonlinear least-squares algorithm in MATLAB. This figure shows that the model of the optical response is a good representation of the real optical response. Since the injection of broad-band white noise effectively blinds the detector to gravitational waves, the online measurement of the optical transfer function used in the calibration scheme is done at only a few spot frequencies using injected calibration lines (see section 2.3.4).

2.3.2. The Michelson length-control servo. The Michelson length-control servo uses a split feedback system where low-frequency (up to around 20 Hz) feedback signals are applied, after appropriate filtering, to actuators at the intermediate stage of the triple pendulums from which the main test mass mirrors are suspended; high-frequency feedback signals (up to few hundred Hz) are applied directly to the test masses using electrostatic actuators mounted on reaction pendulum chains suspended behind the main pendulum chains. These two feedback paths are shown in figure 1 as three electronic elements (H_{com}, H_{slow} and H_{fast}), and two actuators, H_{IMD} and H_{ESD}. Here, 'slow' refers to the low-frequency feedback path and 'fast' refers to the high-frequency feedback path.

For the two parts of GEO's S3 participation (S3 I and S3 II), separate electronic models were developed to reflect the different states of the electronics at the two times. These models were designed by making pole/zero fits to measured transfer functions of the electronics and were typically accurate to better than 10% over their frequency band of influence.

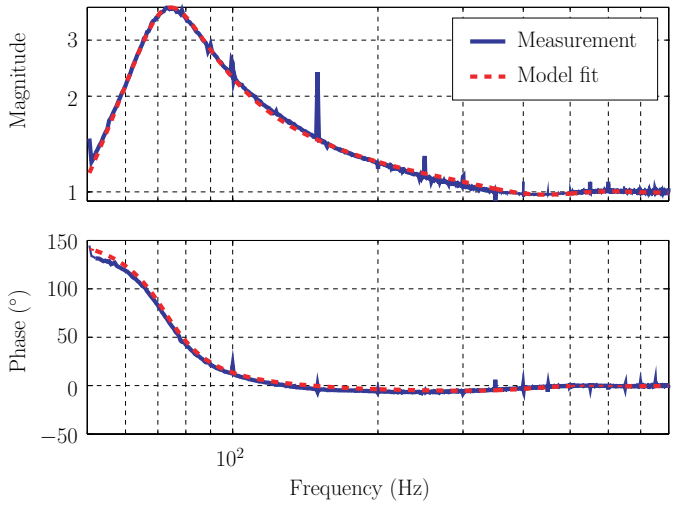

Figure 3. A measurement of the closed-loop transfer function of the Michelson servo used during S3 II. A model fit is also shown. The unity gain point of the servo is around 50 Hz.

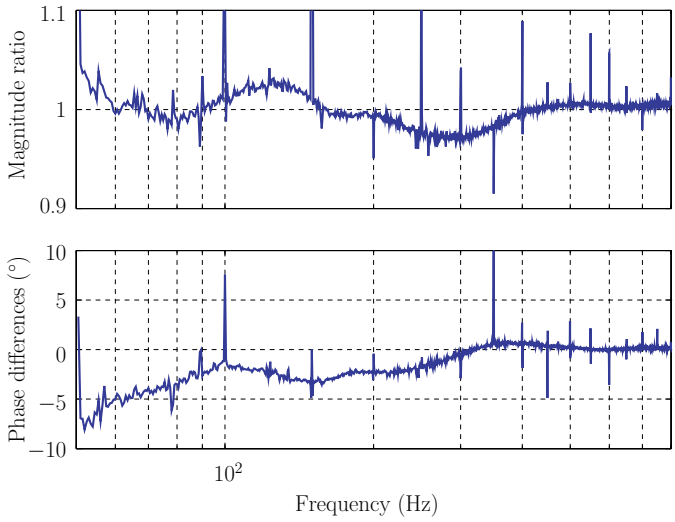

Figure 4. The magnitude and phase differences between the measured and modelled closed-loop transfer function of the Michelson servo for S3 II.

2.3.3. Michelson closed-loop transfer function. The closed-loop transfer function of the Michelson length-control servo is easily measured while the detector is operating in normal conditions. Since this transfer function includes all aspects of the detector model, it provides a useful test. In addition, there are two gains (one in each of the common electronics and the fast path electronics) that are not included in the measurement of the electronic transfer functions; these gains can be determined by fitting a model of the closed-loop transfer function to the measurement. Such a fit for the S3 II configuration is shown in figure 3; the magnitude ratio and phase difference between the model and the measurement are shown in figure 4.

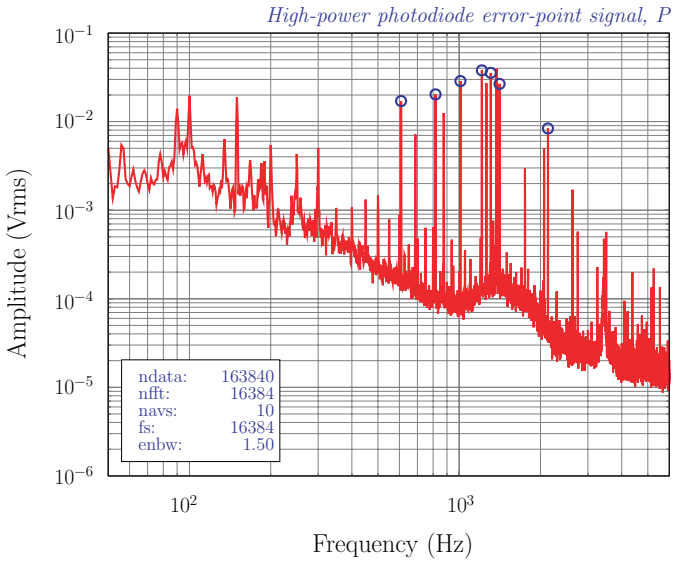

Figure 5. A snap-shot spectrum of the high-power error-point signal during S3 I. The injected calibration lines are highlighted with (blue) circles.

2.3.4. Injected calibration lines. In order to determine the four parameters of optical transfer function of the detector, we require at least four measurements of transfer function. In other words, we need to inject at least four calibration lines. In order to map out the frequency response, we need the flexibility to inject lines at arbitrary frequencies in the detection band with a range of amplitudes. To do this, a signal injector was developed that allows, through software control, the injection of multiple lines, each with a user definable phase, frequency and amplitude. The signal generator, which is locked to a GPS frequency standard to ensure the frequency stability of the injected lines, was used to continually inject calibration lines for the entire S3 run. The output of the generator was recorded at all times to allow the transfer function from the injection point to the high-power error-point to be continually evaluated.

For S3 I and S3 II, different sets of calibration lines were injected to reflect the change in the detector response due to the different tuning of the signal-recycling mirror and the improvement in sensitivity of the detector. Figures 5 and 6 show snap-shot spectra of the high-power error-point signal for S3 I and S3 II respectively, clearly showing the injected calibration lines.

3. Calibration procedure applied to the S3 science run

The calibration procedure detailed in [10] and summarized above was implemented as a full software pipeline suitable for calibrating the S3 science run data. The details of this pipeline are covered in the following section.

3.1. The calibration pipeline applied to the S3 science run detector configuration

The calibration pipeline that is described in [10] was implemented as a stand-alone application that takes in GEO raw output data and calculates a calibrated strain signal together with some diagnostic and intermediate data products. A simplified schematic of this calibration pipeline

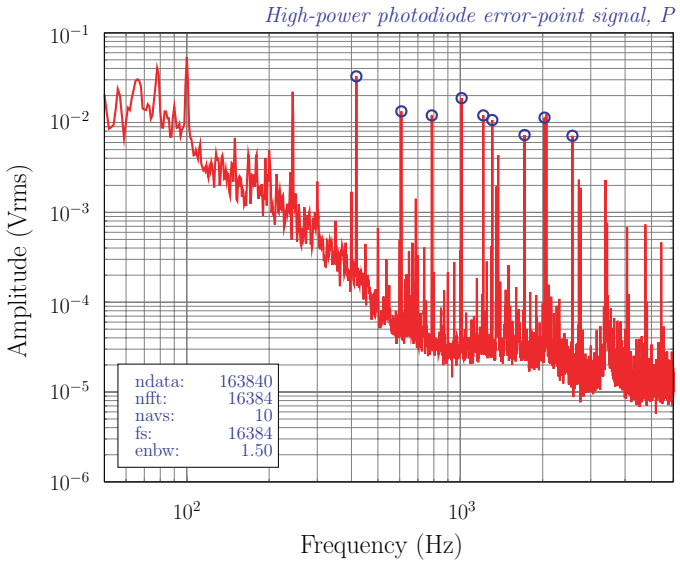

Figure 6. A snap-shot spectrum of the high-power error-point signal during S3 II. The injected calibration lines are highlighted with (blue) circles.

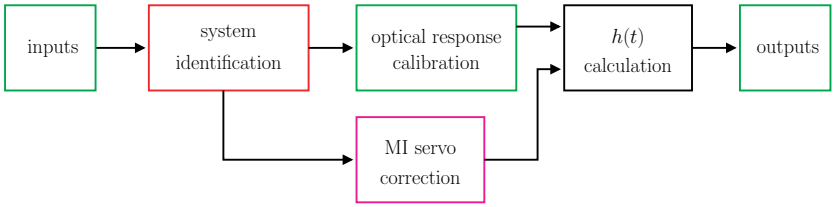

Figure 7. A schematic of the tasks involved in the calibration of the dual-recycled GEO.

is shown in figure 7. The details of each of these steps are discussed below in relation to the S3 calibration.

3.1.1. System identification. The system identification routine identifies the parameters of the **P** optical transfer function. The optical transfer function was estimated once per second by fitting a detailed frequency domain model of the detector to the transfer function measurements made from the injected calibration lines measured in both the injected signal and the **P** error-point signal. This fitting was done using a least-squares optimization routine which minimizes the differences between the complex measurements of the transfer function and the complex detector model by varying the parameters of the model optical response. When a set tolerance for successive changes in the quality of the fit is reached, the routine returns its best guess parameter set (the pole frequency, the pole Q, the zero frequency and the overall gain factor). In addition, the routine is restricted to a certain number of iterations to ensure that it returned a parameter set in a small fraction of a second. (This is needed if we are to ensure the calibration pipeline runs online in real time. So far, this condition has never arisen in normal operating conditions.)

The optimization routine also returns a measure of success (χ^2) by forming the sum of the squares of the differences between the final, best model guess and the measured transfer

function. Each value in the sum is normalized by the estimated variance of the measured transfer function points; this variance is in turn based on the signal-to-noise ratio of the particular calibration line in the detector error-point spectrum. No normalization is done for the number of parameters being fitted. This χ^2 measure of success gives an indication of how well the model of the detector describes the real instrument response for any particular time.

3.1.2. Optical response calibration. Calibrating out the optical response was done using time-domain IIR filters. These inverse optical filters were designed online using the bilinear transform to convert the parametrized frequency-domain model returned by the system identification routine to filter coefficients. The filters were designed to operate at $16 \times 16\,384$ Hz so that the Nyquist rate of the filters was well above the highest frequency where an accurate representation of the instrument response is required. (The bilinear transform produces filters that typically deviate from the desired response around the Nyquist frequency.) The filters were allowed to slowly vary in time by low-pass filtering the estimated optical parameters with a two-pole filter of corner frequency 0.01 Hz before calculating the filter coefficients; this results in changes of the filter coefficients of much less than 1% between successive updates (for long-term variations, see section 4.1).

In order to use these filters, the detector output signal (**P**) had to be up-sampled to the correct sample rate; this was done using a band-limited interpolation method. After filtering, the (displacement) signal was down-sampled back to 16 384 Hz.

3.1.3. Correcting for the response of the Michelson length-control servo. The effect of the Michelson length-control servo was corrected for using a set of time-invariant IIR filters based on detailed models of the servo electronics and using the method described in section 2.1. Application of this method leads to two time-series that represent differential arm-length changes in different frequency bands. A different set of filters was developed for S3 I and S3 II to reflect the changes of the electronics between the two times.

3.1.4. Calculating the strain signal. Calculation of the strain signal is done by summing the three (displacement) time-series. One displacement signal is the result of calibrating out the effect of the optical response from the de-whitened error-point signal. The other two displacement signals are the simulated feedback signals got from filtering the de-whitened error-point signal through the models of the fast and slow feedback chains. The resulting time-series represents the differential arm-length changes of the detector at all frequencies where the model of the detector is accurate (50–6000 Hz). This time-series is converted to a strain signal by dividing by 1200—the arm-length of GEO 600.

4. The results of calibrating the S3 science data

The calibration pipeline described above was applied to the two data sets of S3. Slightly different versions of the pipeline were implemented to reflect the differences in the detector status for the two parts (S3 I and S3 II) of S3. The following sections show some of the results of calibrating the S3 data. A brief discussion is given of the accuracy achieved in the calibration process.

4.1. Recovered optical parameters

Figures 8 and 9 show the time evolution of the optical parameters estimated by the system identification routine from the S3 I and S3 II data sets. Included in the plots is a lock status

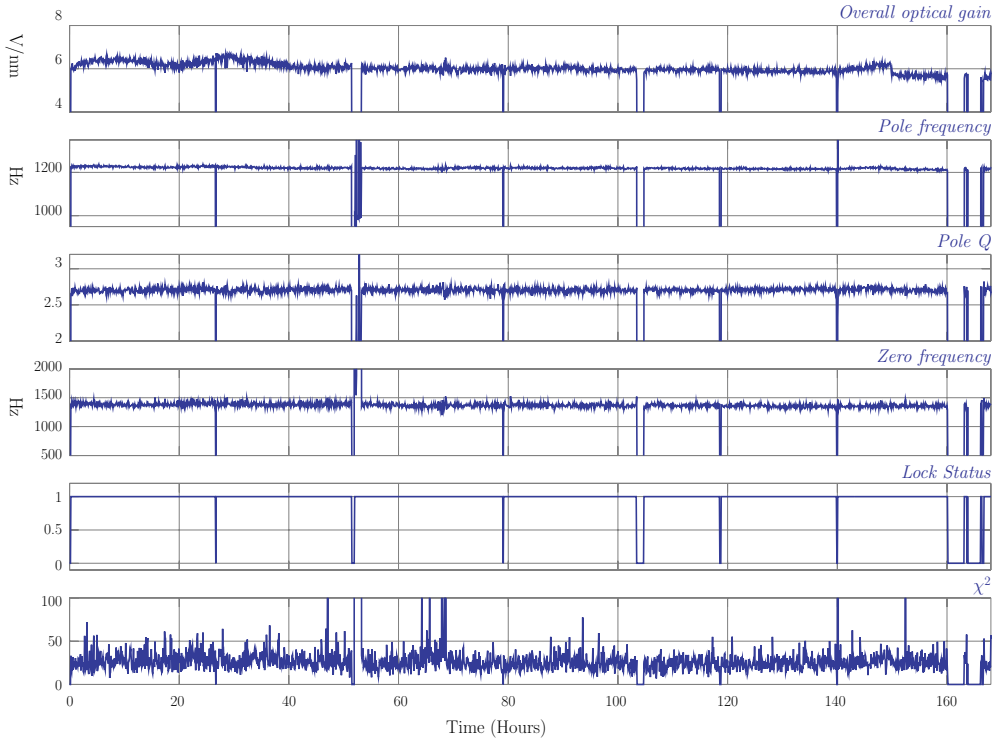

Figure 8. The recovered optical parameters from the S3 I science run. The data are plotted every 30 s for the duration of the run. The bottom two traces show the lock status of the detector (1 is locked) and the recovered χ^2 values respectively.

indicator: when this has a value of 1, the detector was locked at its nominal operating point. The parameter values are shown plotted every 30 s for the duration of the two data sets. Some clear differences are noticeable between the two recovered parameter sets. Firstly, there is a significant difference between the recovered overall optical gain during S3 I compared to S3 II. This difference arises from changes in the optical readout system between the two data recording periods. We also see a change in the other optical parameters due to the change in the tuning frequency of the signal-recycling response. In S3 I, the signal-recycling cavity was tuned to have its peak amplification around 1200 Hz; in S3 II, this was changed to around 1 kHz.

A significant increase in the nominal χ^2 values returned by the system identification routine can also be observed. This is believed to come from a less accurate detector model of the feedback electronics. In addition, the choice of demodulation phase that maximizes **P** around 1 kHz in S3 II means that the assumption that the in-loop and out-of-loop photodiode responses are equivalent is no longer valid; this could easily lead to the observed high χ^2 values though more investigation is needed to be sure of the cause. Experiments after S3 show that when the in-loop diode is correctly modelled, χ^2 values consistent with S3 I can be recovered. Other possible influences on the χ^2 values include the presence of varying lines and spectral features in the neighbourhood of the calibration lines. More investigation into these effects is required to more fully understand the information contained in the χ^2 values.

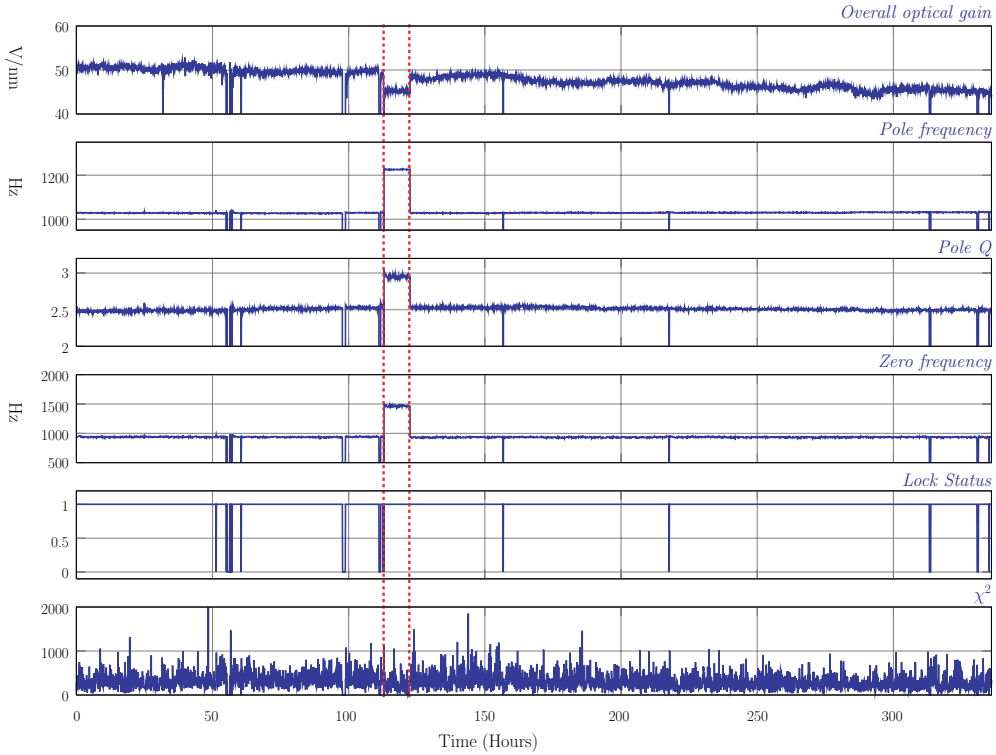

Figure 9. The recovered optical parameters of the S3 II science run plotted every 30 s for the duration of the run. The detector lock status and the recovered χ^2 values from the system identification routine are shown on the bottom two traces. The red, vertical, dotted lines indicate the period of time when the detector was detuned to a different operating point: 1200 Hz instead of 1000 Hz.

In the S3 II parameter set, a change in the recovered parameter values can be seen between hours 110 and 125 (marked by the vertical dotted lines). This is due to the 'accidental' tuning of the signal-recycling mirror to the wrong frequency (the S3 I value). However, we see that the calibration scheme was able to work over this period and continued to estimate the correct parameters and calibrate the detector output.

4.2. Calibration accuracy

The overall calibration accuracy for the S3 data is determined by a complicated combination of many individual errors and considerations. To get a feeling for the scale of the calibration accuracy we can confine ourselves to looking at the more dominant errors in the calibration process. The systematic errors in the process are dominated by the inaccuracies in the model of the closed-loop transfer function of the Michelson length-control servo. These are up to 3% in magnitude at frequencies from 50 Hz to 800 Hz; the corresponding phase error is of the order 5° towards lower frequencies, and essentially zero above 500 Hz. The uncertainties in the process are dominated by the uncertainties in the calibration of the electrostatic actuator used to inject the calibration lines. The magnitude of the ESD calibration is believed to be accurate to within 5%; no error in phase due to possible deviation from the $1/f^2$ response is expected in the frequency band from 50 Hz to 4 kHz. Additional noise from the up-sampling

of the overall optical gain estimates introduces a further random error at the level of 3% in magnitude and 3° in phase. Other sources of error are identified but are at a much lower level than those discussed above; they are therefore omitted from this simple consideration of the errors. From these considerations, we conclude that the calibration is accurate to within 10% in magnitude and 3° in phase at frequencies from around 500 Hz to 6 kHz and to within 15% in magnitude and 8° in phase at frequencies down to 50 Hz.

4.3. Summary

A calibration pipeline was implemented and run over the two S3 data sets. The process was able to recover and calibrate out, in real time, an accurate estimate of the optical response of the detector. The result was a data set containing $h(t)$ data (together with a data quality indicator with 1 Hz resolution) for the entire time that the detector was locked and tuned to its nominal operating point.

The calibration pipeline can be refined and extended to include the calibration of the other output quadrature of the detector (**Q**); the combination of the two calibrated output streams can lead to an optimal calibration in the future. Further improvements can be made to some of the signal processing techniques that are used so that the errors in the process can be minimized. In addition, significant effort needs to be invested in understanding the information contained in the χ^2 values and how the values depend on different states of the interferometer as well as on the injected calibration lines.

References

[1] Hewitson M *et al* 2003 A report on the status of the GEO 600 gravitational wave detector *Class. Quantum Grav.* **20** S581–91
[2] Hewitson M *et al* 2003 Calibration of GEO600 for the S1 science run *Class. Quantum Grav.* **20** S885–93
[3] Strain K A and Meers B J 1991 *Phys. Rev. Lett.* **66** 1391–4
[4] Heinzel G *et al* 1998 *Phys. Rev. Lett.* **81** 5493–6
[5] Heinzel G *et al* 2002 Dual recycling for GEO 600 *Class. Quantum Grav.* **19** 1547–53
[6] Grote H, Freise A, Malec M, Heinzel G, Willke B, Lück H, Strain K A, Hough J and Danzmann K 2004 *Class. Quantum Grav.* **21** S473–80
[7] Willke B *et al* 2004 Status of GEO 600 *Class. Quantum Grav.* **21** S417–23
[8] Smith J R *et al* 2004 *Class. Quantum Grav.* **21** S1737
[9] Hewitson M *et al* 2003 Calibration of the power-recycled gravitational wave detector, GEO600 *Rev. Sci. Instrum.* **74** 4184
[10] Hewitson M *et al* 2004 Principles of calibrating the dual-recycled GEO 600 *Rev. Sci. Instrum.* submitted