CNRS
Centre National de la Recherche Scientifique

INFN
Istituto Nazionale di Fisica Nucleare

Longitudinal noise subtraction: the *alpha*-, *beta*- and *gamma*-technique

VIR-050A-08

Bas Swinkels, Enrico Campagna, Gabriele Vajente, Lisa Barsotti and Matt Evans

Issue: 1

Date: 9th June 2008

VIRGO * A joint CNRS-INFN Project
Via E. Amaldi, I-56021 S. Stefano a Macerata - Cascina (Pisa)
Secretariat: Telephone (39) 050 752 521 * FAX (39) 050 752 550 * Email W3@virgo.infn.it

Contents

1 Introduction

The Virgo interferometer – from the point of view of the longitudinal control – is controlled in 4 degrees of freedom: the difference between the two short arms (MICH), the length of the power-recycling cavity (PRCL), the difference in length of the long arms (DARM) and the sum of the two long arm lengths (CARM). These consist of various combinations of the movements of the 6 main mirrors in the interferometer. This might seem like an arbitrary choice, but they form a convenient decomposition that matches the various optical signals coming out of the interferometer. Moreover, one of them (DARM) is directly proportional to the gravitational-wave signal. [1] Unfortunately, this 'diagonalization' is not perfect and considerable cross-coupling between the degrees of freedom is present. As long as these couplings are small, they form no major problem for keeping the interferometer locked at its operational point. They do, however, form a problem when trying to achieve the best possible sensitivity for gravitational waves. Without any compensation, the cross-coupling of the control noise from the MICH loop (and to a lesser extent from the PRCL and CARM loop) to the important DARM loop dominates the sensitivity at low frequencies (up to 200 Hz).

The obvious goal is to minimize the coupling between these loops as much as possible. This can be achieved by a noise subtraction technique, which consists of adding a certain amount of the actuation signal of the MICH loop to the actuation signal of the DARM loop. Originally, this was simply a constant factor, which was called *alpha* for no particular reason. At the moment, this has evolved to a frequency dependent filter and the trick as a whole is called the *alpha-technique*. Similar schemes are used to prevent control noise of PRCL and CARM from entering DARM, which are called the *beta-technique* and the *gamma-technique*. This note will mainly discuss the alpha-technique, but the whole discussion is directly applicable to the beta- and gamma-technique.

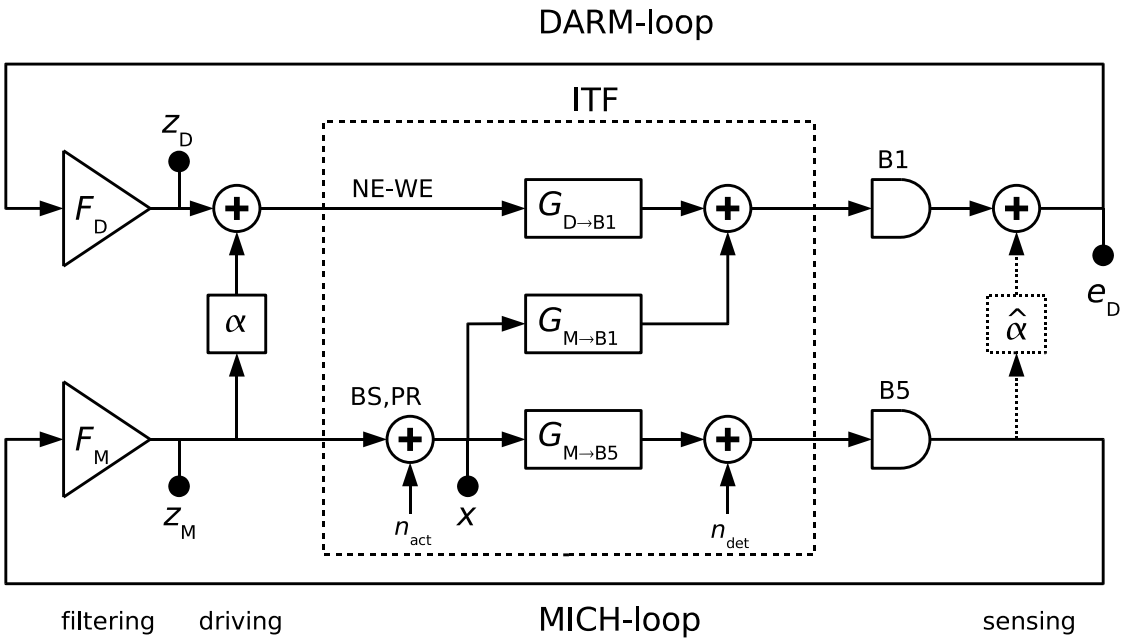

Figure 1: Simplified scheme of the two control loops involved in the alpha-technique. The part in the dotted box represents the physical signals inside the interferometer, the remaining part represents digital signals inside Global Control. See text for the definition of the various signals.

2 Theory

The theory behind the alpha-technique can be understood using the simplified scheme of the DARM and MICH loops of Fig. 1. The signal to control the DARM loop is received by the B1 detector (`Pr_B1_ACp`). This signal is amplified by the DARM loop filter F_{D} and is sent to the interferometer by actuating the NE and WE mirrors in opposite directions. This signal couples back to the detector with a transfer-function $G_{\mathrm{D}\to\mathrm{B1}}$. Similarly, in the MICH loop the control signal is detected by the B5 detector (`Pr_B5_ACq`),[1] amplified by the loop filter F_{M}, sent to the interferometer via the BS and PR mirrors and finally couples back to the detector with a transfer-function $G_{\mathrm{M}\to\mathrm{B5}}$. The cross-coupling from the MICH input to the B1 detector is modeled by a transfer-function $G_{\mathrm{M}\to\mathrm{B1}}$. The noise in the MICH loop is assumed to enter at two points: the actuation noise n_{act} (e.g. *coil-driver* noise and seismic noise entering via the mirror suspension) and the detection noise n_{det} (e.g. shot noise of the detector). Note that all the mentioned elements have frequency dependent transfer-functions. The following analysis is thus assumed to take place in the frequency domain, without any explicit notation.

As a start, the case is analyzed without any noise-subtraction (so $\alpha = 0$). The signal at the point labeled x in Fig. 1 is then

$$x = F_{\mathrm{M}}(G_{\mathrm{M}\to\mathrm{B5}}x + n_{\mathrm{det}}) + n_{\mathrm{act}}. \tag{1}$$

Solving for x gives

$$x = (F_{\mathrm{M}}n_{\mathrm{det}} + n_{\mathrm{act}})/(1 - F_{\mathrm{M}}G_{\mathrm{M}\to\mathrm{B5}}) = G_{\mathrm{Mcl}}(F_{\mathrm{M}}n_{\mathrm{det}} + n_{\mathrm{act}}), \tag{2}$$

where $G_{\mathrm{Mcl}} = 1/(1 - F_{\mathrm{M}}G_{\mathrm{M}\to\mathrm{B5}})$ is the closed-loop transfer-function of the MICH loop. Then consider how this noise arrives at the B1 detector

$$B1 = G_{\mathrm{D}\to\mathrm{B1}}F_{\mathrm{D}}B1 + G_{\mathrm{M}\to\mathrm{B1}}x \tag{3}$$

Solving for $B1$ and inserting Eq. 2 yields

$$B1 = G_{\mathrm{M}\to\mathrm{B1}}x/(1 - G_{\mathrm{D}\to\mathrm{B1}}F_{\mathrm{D}}) = G_{\mathrm{Dcl}}G_{\mathrm{M}\to\mathrm{B1}}G_{\mathrm{Mcl}}(F_{\mathrm{M}}n_{\mathrm{det}} + n_{\mathrm{act}}) \tag{4}$$

[1] Recently, we started using `Pr_B2_8MHz_ACp` for controlling MICH. This does not change the following discussion, other than replacing B5 with B2.

where $G_{\mathrm{Dcl}} = 1/(1 - F_{\mathrm{D}} G_{\mathrm{D}\to\mathrm{B1}})$ is the closed-loop transfer-function of the DARM loop.

2.1 Noise suppression by alpha

The noise cross-coupling from the MICH to the DARM loop can be reduced by adding the MICH actuation signal z_{M} multiplied by a transfer-function α to the one of DARM. Tracing all the noise contributions to the B1 detector in a way similar to above now gives

$$B1 = G_{\mathrm{Dcl}} G_{\mathrm{Mcl}} \Big[\left(\alpha G_{\mathrm{D}\to\mathrm{B1}} + G_{\mathrm{M}\to\mathrm{B1}}\right) F_{\mathrm{M}} n_{\mathrm{det}} + \left(\alpha F_{\mathrm{M}} G_{\mathrm{D}\to\mathrm{B1}} G_{\mathrm{M}\to\mathrm{B5}} + G_{\mathrm{M}\to\mathrm{B1}}\right) n_{\mathrm{act}} \Big] \tag{5}$$

The coupling of the detection noise n_{det} can be canceled completely by setting the term $(\alpha G_{\mathrm{D}\to\mathrm{B1}} + G_{\mathrm{M}\to\mathrm{B1}})$ to zero. The ideal value of α is thus

$$\alpha = -\frac{G_{\mathrm{M}\to\mathrm{B1}}}{G_{\mathrm{D}\to\mathrm{B1}}}. \tag{6}$$

Inserting the nominal value back into equation Eq. 5 shows the effect on the coupling of the actuation noise n_{act} to B1:

$$B1 = G_{\mathrm{Dcl}} G_{\mathrm{M}\to\mathrm{B1}} n_{\mathrm{act}}. \tag{7}$$

Compared to Eq. 4, the actuation noise is thus not suppressed, but has actually increases by a factor of $1/G_{\mathrm{Mcl}}$. The different result for the two noise sources can be understood using Fig. 1: there are two branches from z_{M} to $B1$, one with transfer function $G_{\mathrm{M}\to\mathrm{B1}}$ and one with transfer function $\alpha G_{\mathrm{D}\to\mathrm{B1}}$. Setting alpha to its nominal value, any noise present at z_M cancels perfectly at $B1$. This is not the case for the actuation noise n_{act}, which enters on only one of the two branches.

2.2 Alternatives

Instead of canceling the detection noise, it is also possible to cancel the actuation noise n_{act} by setting alpha to

$$\alpha = -\frac{G_{\mathrm{M}\to\mathrm{B1}}}{F_{\mathrm{M}} G_{\mathrm{D}\to\mathrm{B1}} G_{\mathrm{M}\to\mathrm{B5}}}. \tag{8}$$

Substituting in Eq. 5 shows that the remaining noise at the B1 detector is now

$$B1 = -\frac{G_{\mathrm{Dcl}} G_{\mathrm{M}\to\mathrm{B1}}}{G_{\mathrm{M}\to\mathrm{B5}}} n_{\mathrm{det}}. \tag{9}$$

Since, in practice, the detection noise dominates the effect of the actuation noise (especially above 10 Hz), setting alpha to this alternative value is not an option.

Instead of compensating at the level of the *driving*, it is also possible to obtain similar results by compensating at the level of the *sensing*. This alternative technique $\widehat{\alpha}$ is indicated by the dotted lines in Fig. 1. With a reasoning similar to above, the detection noise n_{det} can be totally canceled at the level of the B1 detector by setting $\widehat{\alpha}$ to

$$\widehat{\alpha} = -\frac{F_{\mathrm{M}} G_{\mathrm{M}\to\mathrm{B1}}}{F_{\mathrm{D}} G_{\mathrm{D}\to\mathrm{B1}}}. \tag{10}$$

The disadvantage of this solution is mainly the complexity of the resulting formula compared to Eq. 6. In this case, $\widehat{\alpha}$ would have to be changed every time the MICH or DARM filter changes. A variant of this idea is obtained by not trying to minimize the cross-coupling at the B1 detector, but to minimize the coupling at the error point of the DARM loop, which is indicated by e_{D} in Fig. 1. In this case, the detection noise is canceled by setting $\widehat{\alpha}$ to

$$\widehat{\alpha} = -F_{\mathrm{M}} G_{\mathrm{M}\to\mathrm{B1}}. \tag{11}$$

The large disadvantage of this solution is that a signal other than the B1 detector would be the new low-noise channel, which is against the common practice.

2.3 Frequency dependence of alpha

According to Eq. 6, the ideal value of alpha is determined by the ratio of two transfer-functions. Since both are frequency dependent, alpha should also be made frequency dependent for best cancellation. The transfer-function $G_{\mathrm{D}\to\mathrm{B1}}$ could be split in a part representing the mechanics of the end-mirrors (which converts an applied voltage to a mirror displacement in meters) and an optical part (which converts meters into Watts on the B1 photo-diode). Similarly, the transfer-function $G_{\mathrm{M}\to\mathrm{B1}}$ could be split in a part representing the mechanics of the beam-splitter (and to a lesser extend the PR-mirror) and an optical part. From theory, it is known that both optical transfer-functions are flat in frequency up to the *cavity pole* at 500 Hz. The ratio between these two should scale with the finesse of the *arm-cavities*: the larger the finesse, the smaller the cross-coupling from MICH to DARM. The main part of the frequency dependence of alpha comes thus from the difference in mechanics of the end-mirrors compared to the beam-splitter (differences in weight, eddy-current effect and reallocation). This is confirmed by the experience that alpha changes a lot every time a change is made in the driving of the mirrors (i.e. a change in the re-allocation of the force to the *marionetta*), but is relatively stable otherwise.

2.4 How to measure alpha?

To calculate the optimal value of alpha, the two transfer-functions of Eq. 6 should be known. $G_{\mathrm{D}\to\mathrm{B1}}$ can be determined directly by injecting noise into the DARM loop and measuring the transfer-function between `Gc_DARM_z` and `Pr_B1_ACp` signals. The transfer-function $G_{\mathrm{M}\to\mathrm{B1}}$ from the MICH to the DARM loop, however, cannot be measured directly, since any noise added to the DARM loop will be suppressed by the feedback. This has to be compensated in the calculation.

First assume that currently no noise cancellation is applied ($\alpha = 0$). The relation between the detector signal $B1$ and the MICH correction signal z_{M} is then given by

$$B1 = G_{\mathrm{M}\to\mathrm{B1}} z_{\mathrm{M}} + G_{\mathrm{D}\to\mathrm{B1}} F_{\mathrm{D}} B1 \tag{12}$$

Solving this for the *measured* transfer-function $M_{\mathrm{M}\to\mathrm{B1}}$ between `Gc_MICH_z` and `Pr_B1_ACp` signals gives

$$M_{\mathrm{M}\to\mathrm{B1}} = \frac{B1}{z_{\mathrm{M}}} = \frac{G_{\mathrm{M}\to\mathrm{B1}}}{1 - G_{\mathrm{D}\to\mathrm{B1}} F_{\mathrm{D}}} = G_{\mathrm{Dcl}} G_{\mathrm{M}\to\mathrm{B1}}. \tag{13}$$

Combining this with Eq. 6 yields the measured transfer-function α_{meas}

$$\alpha_{\mathrm{meas}} = -\frac{M_{\mathrm{M}\to\mathrm{B1}}}{G_{\mathrm{D}\to\mathrm{B1}} G_{\mathrm{Dcl}}}. \tag{14}$$

The difference with Eq. 6 is the correction for the DARM closed-loop transfer-function G_{Dcl}. Note that the simplified model of Fig. 1 assumes that there is no cross-coupling from the DARM loop back to the MICH loop, which might complicate the measurement of alpha. No such cross-coupling was observed in recent measurements of the total 'optical matrix' [2], so this is probably a valid assumption.

2.4.1 Iterative method

Now assume that there is already a preliminary correction α_{old} present. In this case, the signal at the B1 detector is

$$B1 = G_{\mathrm{M}\to\mathrm{B1}} z_{\mathrm{M}} + G_{\mathrm{D}\to\mathrm{B1}} (F_{\mathrm{D}} B1 + \alpha_{\mathrm{old}} z_{\mathrm{M}}). \tag{15}$$

Similar to above, the *measured* transfer-function between z_{M} and $B1$ is obtained as

$$M_{\mathrm{M}\to\mathrm{B1}} = \frac{B1}{z_{\mathrm{M}}} = \frac{G_{\mathrm{M}\to\mathrm{B1}} + G_{\mathrm{D}\to\mathrm{B1}} \alpha_{\mathrm{old}}}{1 - G_{\mathrm{D}\to\mathrm{B1}} F_{\mathrm{D}}} \tag{16}$$

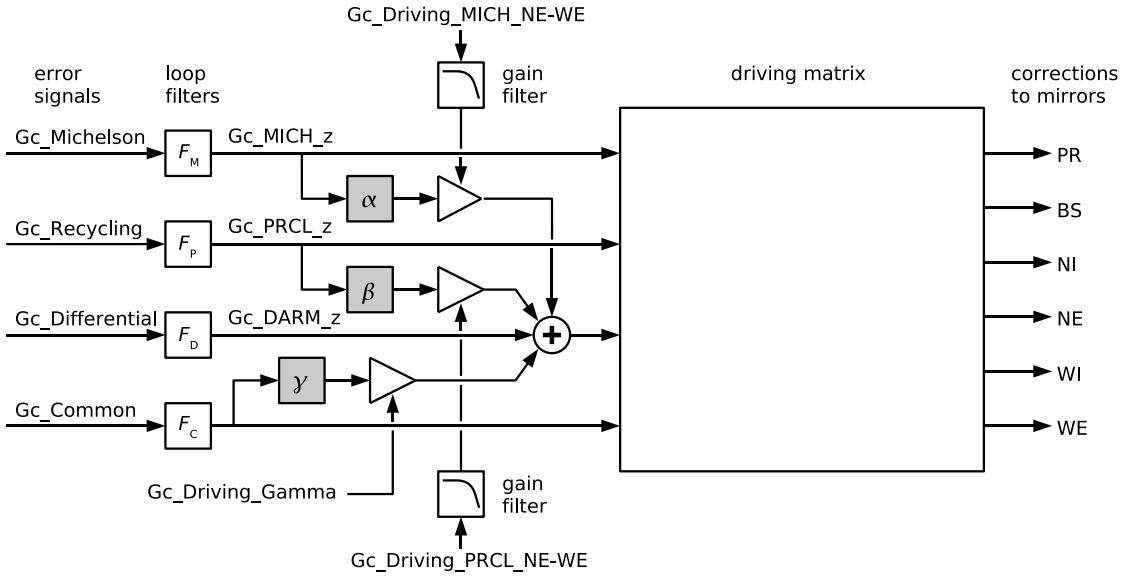

Figure 2: Partial scheme of Global Control, the alpha-, beta- and gamma-filter are indicated in gray. Note that applying the corrections just before the driving matrix effectively creates a matrix with off-axis, frequency-dependent elements.

and finally the measured transfer-function as

$$\alpha_{\mathrm{meas}} = -\frac{M_{\mathrm{M}\to\mathrm{B1}}(1 - G_{\mathrm{D}\to\mathrm{B1}}F_{\mathrm{D}}) - G_{\mathrm{D}\to\mathrm{B1}}\alpha_{\mathrm{old}}}{G_{\mathrm{D}\to\mathrm{B1}}} = \alpha_{\mathrm{old}} - \frac{M_{\mathrm{M}\to\mathrm{B1}}}{G_{\mathrm{Dcl}}G_{\mathrm{D}\to\mathrm{B1}}}. \tag{17}$$

This formula allows a measurement of a new alpha while an old version of the filter is already running. Note that the second term is directly proportional to the remaining cross-coupling and equals zero in case of a perfectly tuned alpha. If the error in calculating this term scales with the magnitude of the term itself, this formula can be used in an *iterative* scheme: first implement a preliminary filter and then repeat making a measurement and updating the filter until the measurement converges. This is of great practical importance: The term $G_{\mathrm{D}\to\mathrm{B1}}G_{\mathrm{Dcl}}$ is hard to measure at low frequencies due to the large gain of the DARM loop. Instead of measuring it, it can be calculated using an approximate model. The iterative scheme ensures that the error due to inaccuracies of the model vanishes as the alpha-filter converges.

To evaluate the performance of a filter, a suppression S_{meas} is defined as the ratio between the remaining cross-coupling and the cross-coupling in the case of alpha set to zero:

$$S_{\mathrm{meas}} = \left|\frac{M_{\mathrm{M}\to\mathrm{B1}}}{M_{\mathrm{M}\to\mathrm{B1}}|_{\alpha=0}}\right| = \left|\frac{M_{\mathrm{M}\to\mathrm{B1}}}{G_{\mathrm{Dcl}}G_{\mathrm{M}\to\mathrm{B1}}}\right| = \left|\frac{M_{\mathrm{M}\to\mathrm{B1}}}{-G_{\mathrm{Dcl}}G_{\mathrm{D}\to\mathrm{B1}}\alpha}\right| \approx \left|\frac{M_{\mathrm{M}\to\mathrm{B1}}}{G_{\mathrm{Dcl}}G_{\mathrm{D}\to\mathrm{B1}}\alpha_{\mathrm{meas}}}\right|. \tag{18}$$

3 Technical implementation

The alpha-filter is implemented as a digital filter with 8 *second-order-stages* inside *Global Control*. The input of the filter is the MICH correction (`Gc_MICH_z`). The output of the filter is added to the DARM correction (`Gc_DARM_z`) just before the *driving matrix*. The filter coefficients can be changed while the interferometer is locked, by iterating the *on-fly* parameter `Gc_Driving_Alpha_SOS` from 1 to 8, while setting the filter coefficients with `Gc_Driving_Alpha_Zero_Freq`, `_Zero_Q`, `_Pole_Freq` and `_Pole_Q`. The overall gain of the filter can be set by `Gc_Driving_MICH_NE-WE`. This value is low-pass-filtered by a *gain-filter*, which allows changing the gain without causing glitches. See Fig. 2 for a scheme of Global Control indicating the position of the alpha-filter.

4 Measurement procedure

To measure the transfer-functions used for calculating alpha, it is required to add noise to the MICH and DARM degrees-of-freedom. This is achieved by performing a *longitudinal noise-budget* [3], which already includes these noise injections. After this measurement, the transfer-function of alpha is calculated with Eq. 17. Since the DARM measurement used to show no coherence at low frequencies (below 30 Hz), the terms $G_{\mathrm{D}\to\mathrm{B1}}$ and G_{Dcl} were based on a model at low frequencies and on the measurements at high frequencies. Recently, a change was made in Global Control that allows the injection of better *shaped* noise. This allowed for measurements with more coherence at low frequencies, so the model of the DARM loop is no longer required. After large changes in the configuration of the interferometer (e.g. a different reallocation of the force sent to the marionette), the measurement usually converges after implementing 2 or 3 new filters.

4.1 Fitting a transfer-function

Once the frequency dependence of the ideal alpha-filter is calculated, it has to be implemented inside Global Control using a transfer-function of finite order. Since Global Control is currently running at the limit of its performance, only 8 *second-order-stages* have been allocated for the alpha-filter. Of these, 2 stages are used to match a notch filter (for the beam-splitter violin-modes at 167 Hz), which is inside the driving part of Global Control and thus a known function. This leaves 6 stages, which means that the remaining function must be fitted with a transfer-function with not more than 12 poles and 12 zeros. In the past, this fitting was performed using the `invfreqs` function of *Matlab*. Since this function gives poor results when fitting a high-order function at once, the fitting was performed by splitting the total frequency range in several pieces, each time fitting a low order function to the residual of the previous fit. With some effort, this resulted in a fitting error of slightly less than one percent over the relevant frequency range. Recently, this complex procedure was substituted with the superior *Vector Fitting* algorithm [4]. This algorithm is able to very accurately fit a high-order transfer-function over the entire frequency range at once. The obtained relative fitting error is usually better than 1 part in 1000, which is basically to within the measurement noise.

The accuracy of the fit directly influences the performance of the alpha-filter. The obtainable suppression S_{misfit} is limited by the relative fitting error

$$S_{\mathrm{misfit}} = \left| \frac{\alpha_{\mathrm{fit}} - \alpha_{\mathrm{meas}}}{\alpha_{\mathrm{meas}}} \right| . \tag{19}$$

The calculation and fitting of alpha is implemented in a few Matlab scripts. The whole procedure of launching a noise budget, calculating, fitting and uploading the new filter takes about half an hour, which is mainly limited by waiting for the data.

5 Results

Figure 3 shows an example of a measurement of the alpha-filter and a fitted 10th-order transfer-function. The resulting suppression for the old filter, the predicted suppression based on the fitting error and the suppression of the new filter are shown in Fig. 4. The real effect of the alpha-technique can be observed best by comparing figures 5 and 6. In the former, the alpha-filter is switched off. The noise from the MICH loop can be seen limiting spectrum of the B1 detector up to a frequency of 300 Hz. In the latter, the alpha-technique is operational. In this case, the MICH noise is only limiting below 40 Hz. Note that both plots are obtained with data from October 2007. As discussed later in section 10, the noise from the MICH loop is now no longer limiting the sensitivity.

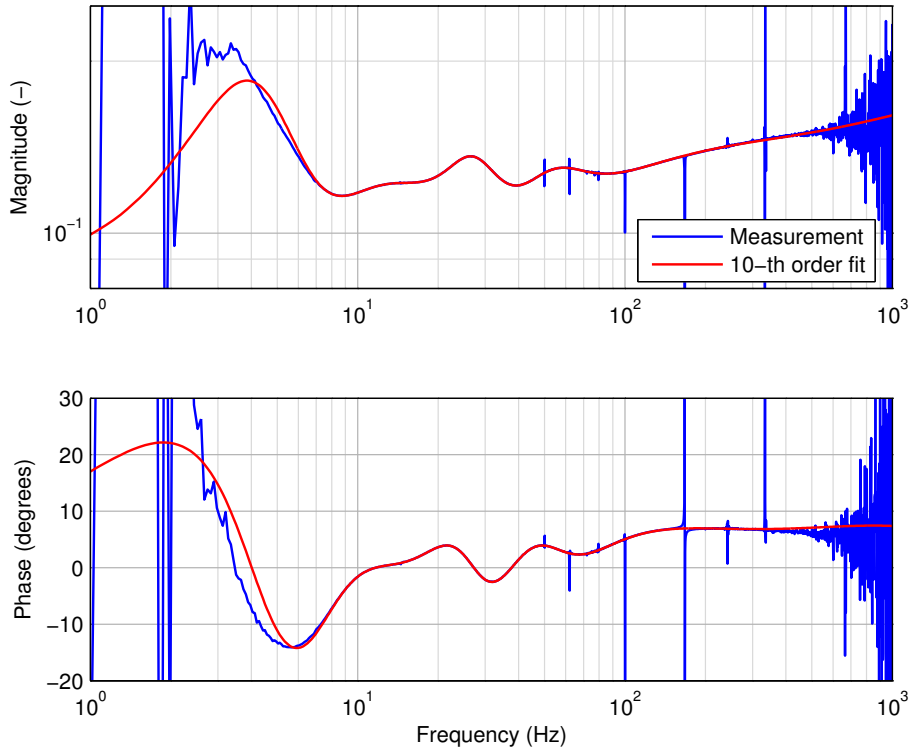

Figure 3: Measurement of the ideal transfer-function of alpha and a fit with a 10th-order filter. The fit is performed over a limited frequency range and only includes points where the measurement has enough coherence.

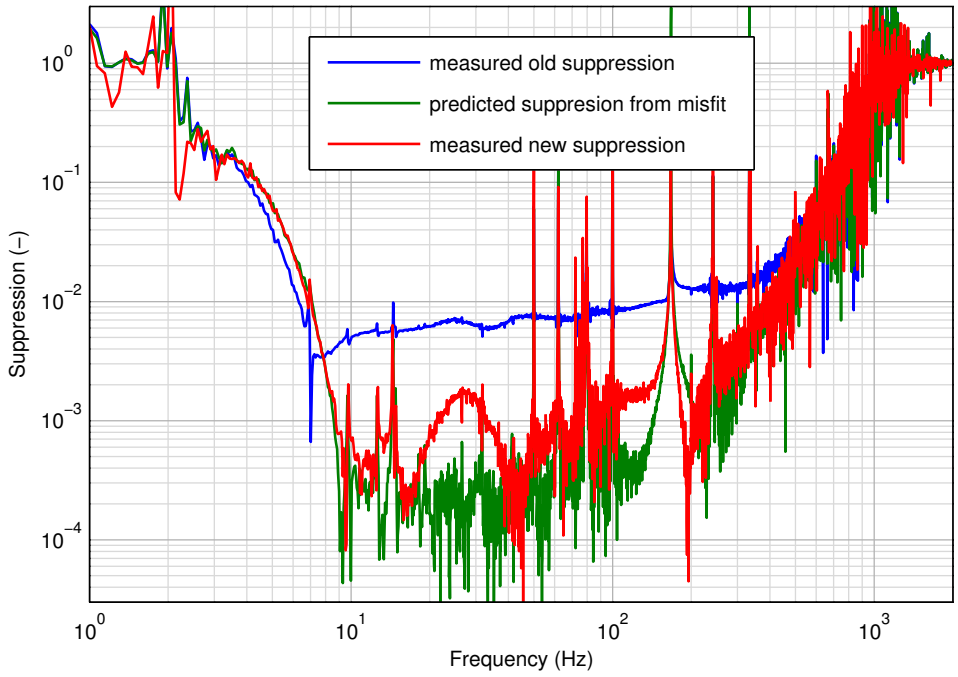

Figure 4: Measured suppression of the alpha-filter before and after retuning, calculated with Eq. 18. Also shown is the suppression predicted from the fitting error, which is calculated with Eq. 19. Note that some of the structure in the new suppression is clearly predicted by the fitting error.

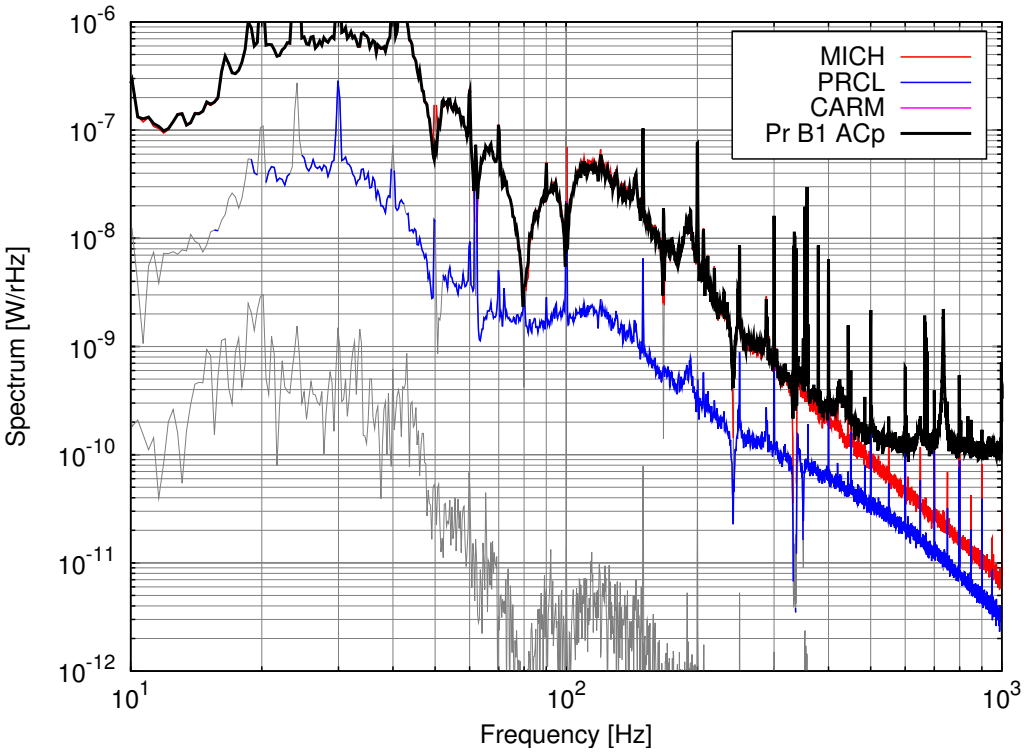

Figure 5: Longitudinal noise budget with both alpha and beta corrections switched off. The cross-coupling from the MICH loop is dominating the dark fringe up to 300 Hz.

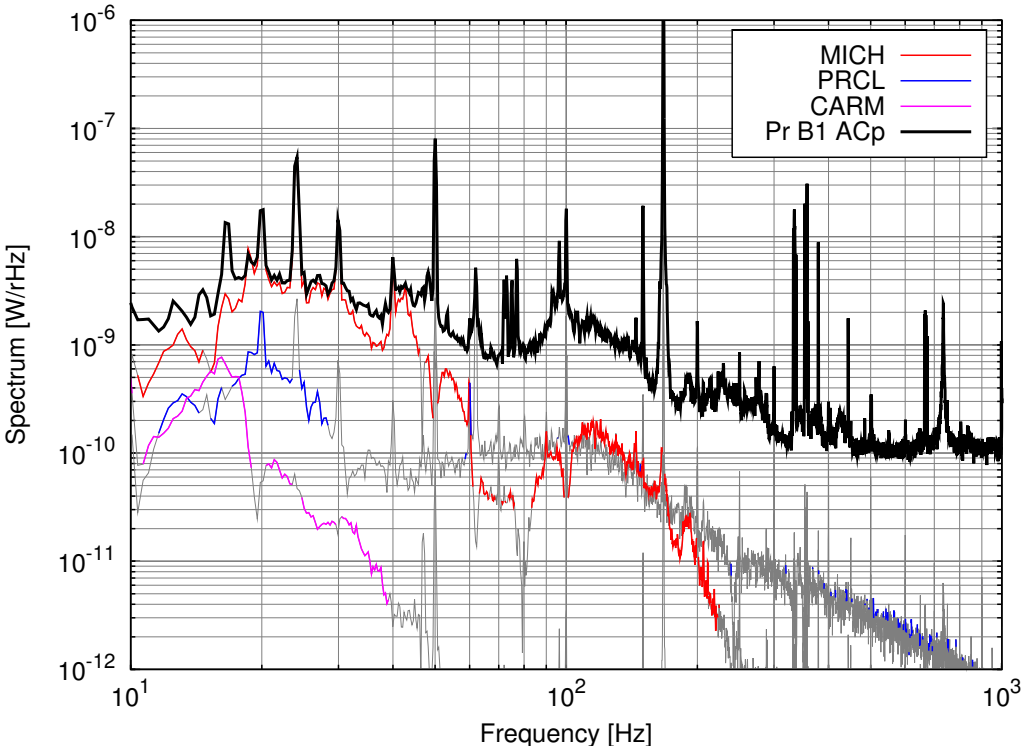

Figure 6: Longitudinal noise budget with alpha and beta corrections in operation.

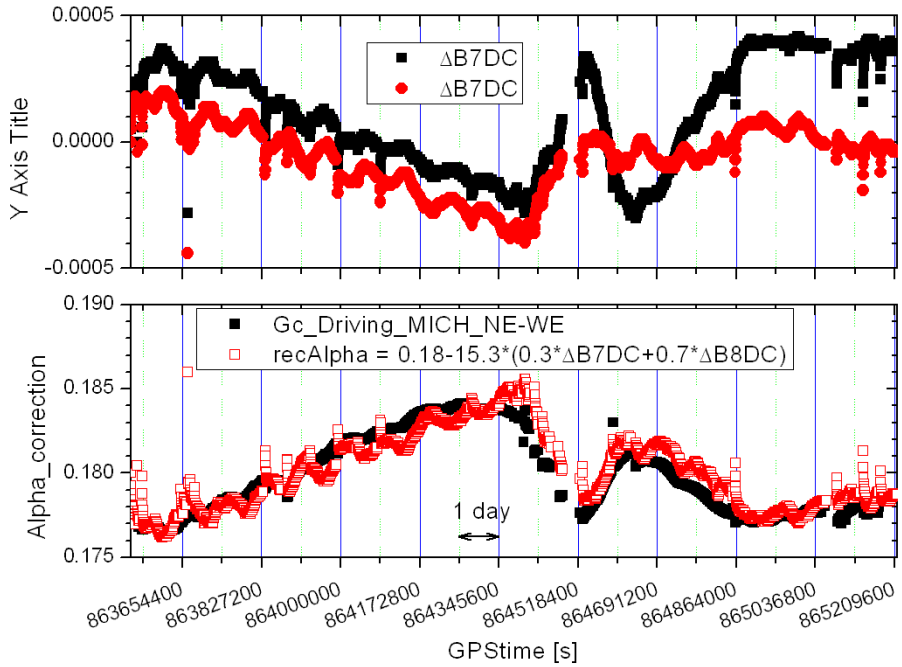

Figure 7: Correlation between the gain of the alpha-filter and the etalon effect during the first weeks of VSR1. Top graph: transmitted power of the two arms, which should scale with the finesse of the corresponding cavities. Bottom graph: Gain of alpha (driven by the servo) compared to a linear combination of the powers (related to the average finesse). Plot taken from [5].

5.1 Repeatability of the measurement

Repeated measurements of the alpha transfer-function showed that it is very stable over time, except for a change in gain of a few percent. This was later explained by a change of the *optical gain* due to a spurious *etalon* in the input mirrors, which changes the finesse of the arm-cavities as a function of the temperature of these mirrors [5]. Such a drift would limit the long term suppression to a factor of 30 with a fixed alpha-filter. Compared to the typical suppression of better than 300 measured with a recently tuned filter, more than an order of magnitude of suppression would be lost.

6 Servoing the gain

The observed drift of the alpha-filter can be compensated by a servo that slowly adjusts the gain of the filter in Global Control. An error signal for this servo can be obtained using a 24 Hz calibration line, which is injected in the MICH loop by Global Control. (The same line is also used for calculating the unity-gain-frequency.) Demodulating this line in the signal of the B1 detector gives a direct measurement of the remaining cross-coupling between MICH and DARM at 24 Hz. This servo is implemented in *AlpLock*, which executes every 20 seconds. The servo is simply integrating the error-signal (`Alp_Lock_SIG_ALPHA`), with the servo-gain adjusted to achieve a time-constant of a few minutes. First tests with the servo showed that the demodulation phase must be accurately tuned to obtain the right error signal. Due to technical reasons, this phase is only set approximately. The fine-tuning (which is necessary every time a new filter is uploaded) is done by adding an offset to the error signal (`ALPHA_OFFSET` in AlpLock).

The effect of the etalon effect can be seen clearly in data taken during VSR1, where the servo adjusts the gain of alpha to track the change in optical gain, see Fig. 7.

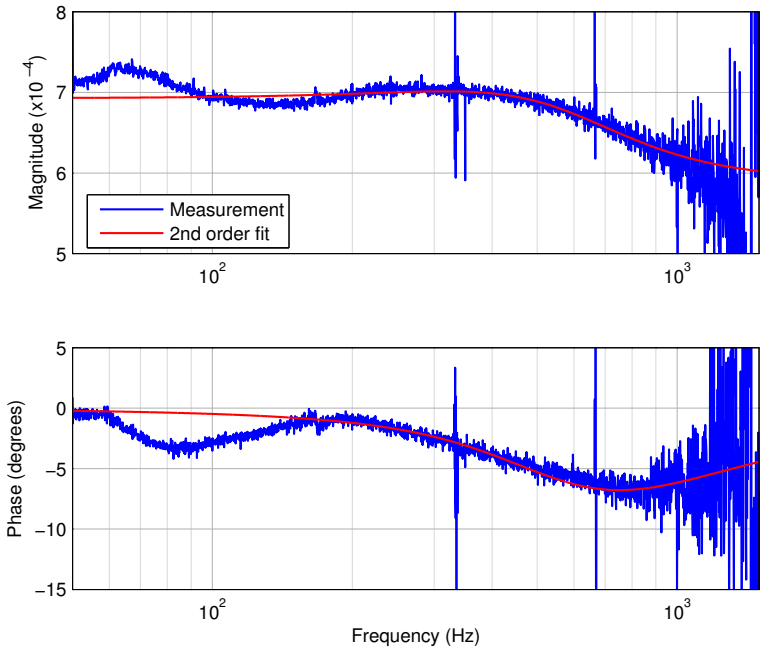

Figure 8: Measurement of the ideal transfer-function of beta (shown in blue) and a fit with a 2nd-order filter (shown in red).

7 Beta

Apart from a cross-coupling from the MICH to DARM, there is also significant cross-coupling from PRCL to DARM. As with MICH, this can be solved by adding a proper amount of PRCL correction signal to the DARM correction. This scheme is known as the *beta-technique*. The theoretical details of this technique will not be discussed, since they are essentially the same as for the alpha-technique except for an appropriate renaming of the signal names.

Beta was first implemented as a constant. Although the transfer-function of the ideal beta was reasonably flat, it was shown that a small phase mismatch limited the suppression to only slightly better than 10. Only recently, beta was changed to a frequency dependent filter. See Fig. 8 for a measurement and a fit of the filter. The noise of the PRCL loop is now far from limiting the sensitivity.

Unlike the gain of the alpha-filter, which typically varies by less than 10 percent over a few weeks, the gain of beta can actually change sign. Similar to alpha, a servo is therefore implemented to change the gain of the filter. The value of beta has shown a clear correlation to finesse-asymmetry of the arm-cavities, which changes due to the etalon effect [5]. This can be seen clearly in data taken during VSR1, where the servo adjusts the gain of beta to track the finesse-asymmetry, see Fig. 9.

7.1 Technical details

Beta is implemented as a digital filter inside Global Control, which currently has three second-order-stages. The filter can be changed using the same parameters as for uploading alpha, but with `Gc_Driving_Alpha_SOS` varying from 11 to 13. The input of the filter is the correction signal for PRCL (`Gc_PRCL_z`) and the output of the filter is added to the DARM correction (`Gc_DARM_z`). The gain of the filter can be set using `Gc_Driving_PRCL_NEWE`, which is low-pass-filtered to prevent glitches when changing the gain. See Fig. 2 for a scheme of Global Control indicating the position of the beta-filter. A slow servo to control the amplitude of beta is implemented

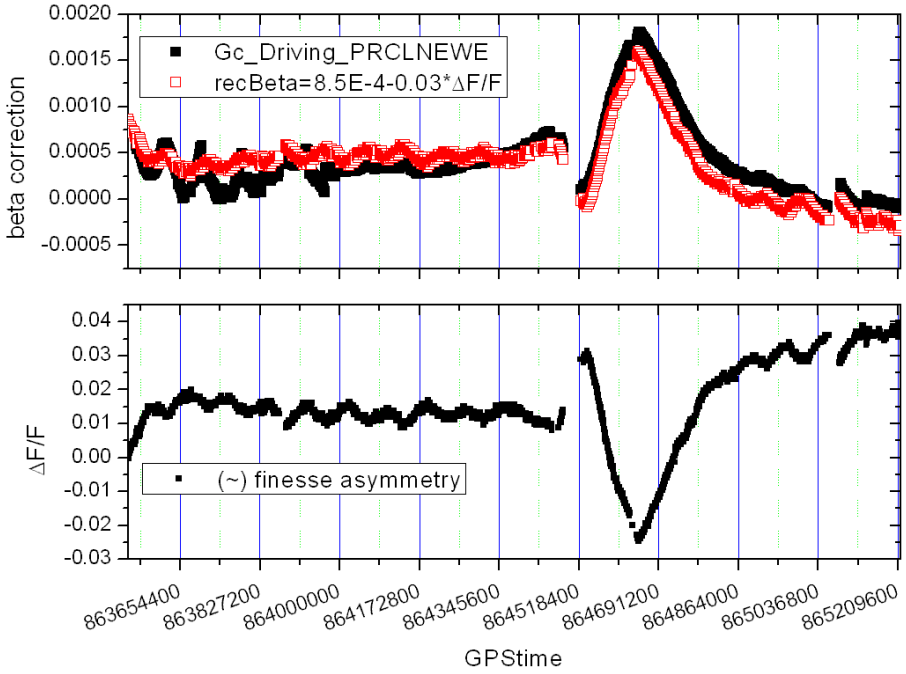

Figure 9: Gain of the beta-filter being servoed to compensate for the etalon effect during the first weeks of VSR1. Top graph: Gain of beta compared to a scaled version of the finesse asymmetry. Bottom graph: Finesse asymmetry calculated from the power transmitted by the arm-cavities. Plot taken from [5].

in AlpLock. The error signal (`Alp_Lock_SIG_BETA`) is obtained by demodulating the 62 Hz PRCL line in the signal of the B1 detector. The typical suppression of the cross-coupling of the PRCL noise by the beta-filter is a factor of 20 around 100 Hz.

8 Gamma

Finally, also the noise of the CARM degree of freedom can couple to the DARM loop. Note that the common mode signal is split into a high-frequency part (using the SSFS) and a low-frequency one. Here, only the low-frequency part is considered, which actuates both end-mirrors with the same sign to keep the laser locked onto the reference cavity. In this scheme, any mis-balancing of the end-mirrors (wrong gain, different frequency dependency, ...) will directly couple a part of the CARM signal to DARM. In the past, this was solved by implementing the so-called *West-End Filter* [2], which was located after the driving matrix on the path to the west-end mirror. The problem with this scheme was that any change in the West-End filter required retuning of the alpha- and beta-filter. For this reason, it was decided to replace the West-End filter by one located just before the driving matrix, very similar to alpha and beta. Unsurprisingly, this filter is now called the *gamma*-filter. The advantage of the old West-End filter was that the response of both end-mirrors is equalized directly, so apart from reducing the coupling from CARM to DARM, it also reduces the reverse coupling from DARM to CARM. This advantage is lost with the gamma-technique, but this is not considered a problem since the CARM loop is not very critical. As with beta, the theory describing the alpha-technique is directly applicable to the gamma-technique after a proper renaming of signal names.

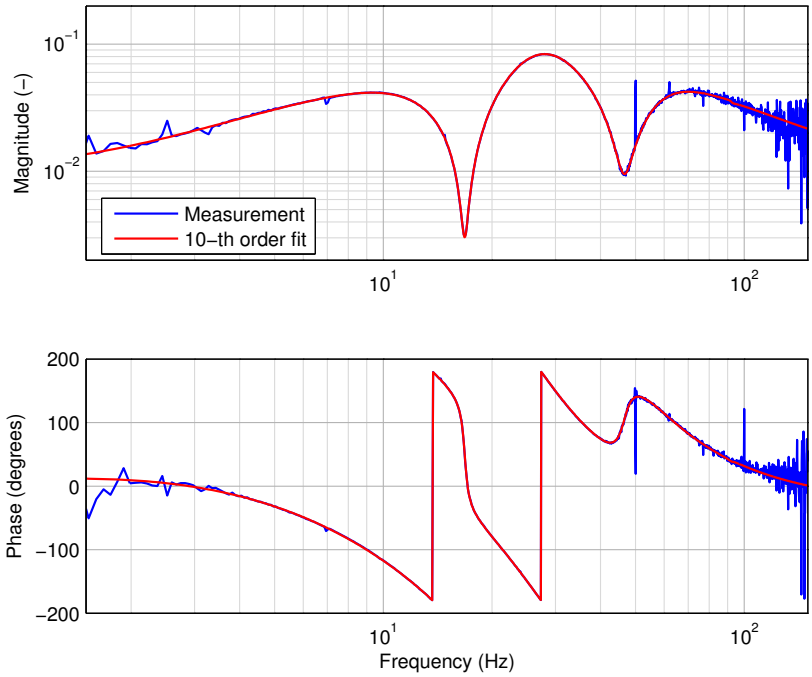

Figure 10: Typical measurement of the transfer-function of gamma (shown in blue) and a fit with a 10th-order filter (shown in red). The distinct notches might be caused by filters located in the DSP of the end-mirrors, which are present to suppress some resonances of the marionetta.

8.1 Technical details

The gamma filter is currently implemented as a 10th-order digital filter in Global Control. The filter can be changed using the same parameters as for uploading alpha, but with `Gc_Driving_Alpha_SOS` varying from 21 to 25. The input of the filter is the correction signal for CARM and the output of the filter is added to the DARM correction (`Gc_DARM_z`). The overall gain of the filter is set using `Gc_Driving_Gamma`. See Fig. 2 for a scheme of Global Control indicating the position of the gamma-filter. The transfer-function of gamma is entirely determined by the mis-balance in driving the end-mirrors, so it includes differences in mechanics, electronics and possibly filters implemented in the two DSPs. Since all of these should be stable over time and do not depend on optical properties like the etalon effect, the gain is kept fixed. See Fig. 10 for a measurement and a fit of the filter. The typical suppression of the cross-coupling of the CARM noise by the gamma-filter is a factor of 100 around 10 Hz.

9 Interaction with the *h*-reconstruction

The ultimate output of a detector for gravitational waves is a strain h as a function of time. In Virgo, this is obtained by transforming the signal of the B1 detector to a displacement with a calibrated transfer-function. The mirrors of the interferometer are, however, not isolated masses, but are being actuated by the longitudinal control system using coils and magnets. Instead of modeling the complete control system, the actuation signals of the different mirrors are transformed by calibrated transfer-functions to mirror displacements, which are then subtracted from the gravitational signal. See e.g. [6] for more details.

It has been argued that since the alpha- and beta-technique add a signal to the actuation of the end-mirrors, the h-reconstruction has to 'work harder' to subtract the effect of the actuation. At least at low frequencies, this is not true. Since these techniques greatly reduce the noise at the B1 detector, the correction signal of the

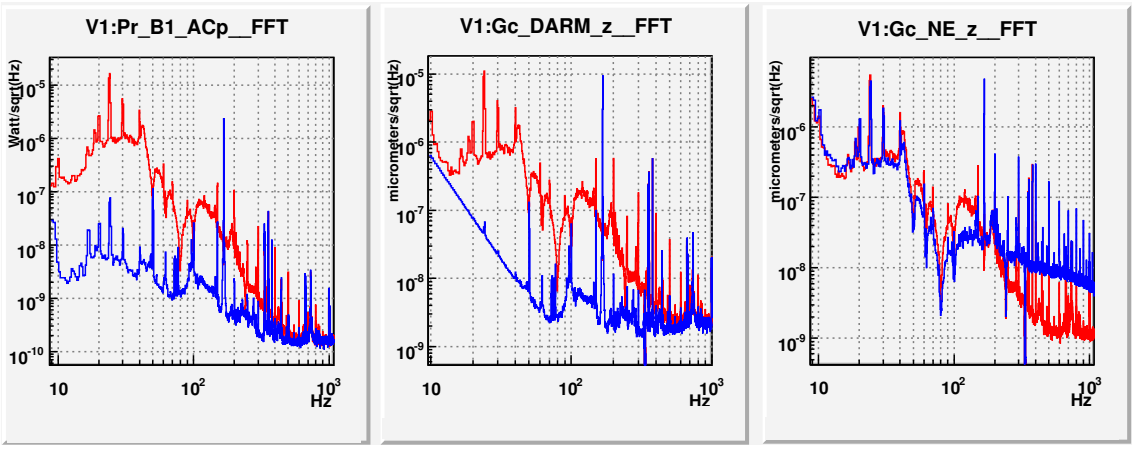

Figure 11: Spectra of the B1 detector, the DARM correction and the signal sent to one of the end-mirrors, with alpha and beta in operation (in blue) and with both of them switched off (in red). Note that the straight line at low frequencies in the DARM spectrum is an artifact of the Fourier window, the spectrum should look similar to the spectrum of B1.

DARM loop (`Gc_DARM_z`), will be equally lower. To this signal, the alpha- and beta-corrections are added, so in the end the correction sent to the end- mirrors is almost equal with or without these techniques. This effect can be clearly seen at low frequencies in Fig. 11. At higher frequencies, the alpha-technique causes a higher level of actuation, but this effect might be neglectable compared to the shot-noise of the B1 detector. Only around 100 Hz (the unity-gain-frequency of the DARM loop) there are some differences.

In principle, the h-reconstruction could achieve the same subtraction of noise from the MICH and PRCL loops offline, as alpha and beta achieve online. In practice, the transfer-functions used in the reconstruction are of low order, so this subtraction is not very effective below 40 Hz. It might be interesting to somehow incorporate the accurate knowledge obtained by the tuning of the alpha- and beta-filters into the reconstruction of gravitational wave channel.

10 Conclusion

In this note, we showed that the alpha-, beta- and gamma-technique can lead to a large gain in sensitivity at low frequencies. With the alpha-technique, a suppression of the MICH noise with a factor around 500 is now routinely achieved. Similarly, the beta-technique now suppresses PRCL noise with more than a factor 10 and the gamma-techniques suppresses CARM noise with a factor of 100, but these are less critical. Because of these results and the recent changes in sensing the photo-diodes and fine-tuning of the driving matrix (see [2]), the noise coming from the auxiliary longitudinal loops is now no longer limiting the sensitivity, see Fig. 12 for a recent noise projection. This has been essential for discovering other problems limiting the sensitivity at low frequency, such as noise caused by stray light, the coil-drivers and the alignment system.

References

[1] L. Barsotti, *"The control of the Virgo interferometer for gravitational wave detection"*, Ph.D. Thesis, Universitá di Pisa (2006).

[2] G. Vajente, E. Campagna and B. Swinkels, *"Locking characterization"*, Virgo note VIR-005A-08 (2008).

[3] G. Vajente, *"Measurement of control noise budgets"*, Virgo note VIR-003A-08 (2008).

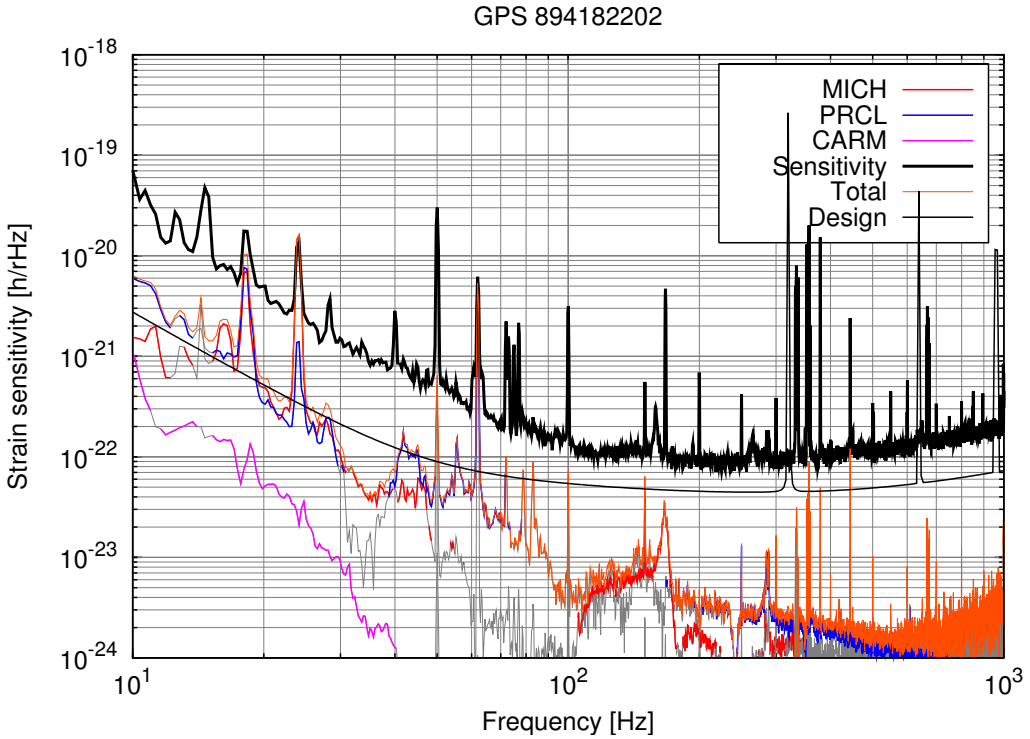

Figure 12: Recent noise budget (May 2008) with alpha, beta and gamma well-tuned. Except for the calibration lines at 24 and 62 Hz and a single line at 18 Hz (probably seismic noise from the laser lab), the noise of the auxiliary loops is not limiting the sensitivity and is only slightly above the Virgo design.

[4] B. Gustavsen and A. Semlyen, *"Rational approximation of frequency domain responses by vector fitting"*, IEEE Trans. Power Delivery, vol. 14, no. 3, pp. 1052-1061 (1999).

[5] M. Punturo, *"Etalon, alpha and beta corrections"*, Virgo logbook entry 16962 (2007).

[6] F. Beauville, *"Prélude à l'analyse des données du détecteur Virgo"*, Ph.D. Thesis, chapter 4, Laboratoire d'Annecy-le-Vieux de Physique des particules (2005).