Classical and Quantum Gravity

PAPER

OPEN ACCESS

RECEIVED
14 October 2025

REVISED
28 November 2025

ACCEPTED FOR PUBLICATION
7 January 2026

PUBLISHED
29 January 2026

Parametrization selection and online adaptation of subtraction filters for gravitational wave detectors*

Mathyn van Dael[1,3,**], Marjolein Daanen[1], Koen Tiels[1], Gert Witvoet[1,2], Bas Swinkels[3], Diego Bersanetti[5], Julia Casanueva[4], Manuel Pinto[4], Maddalena Mantovani[4], Piernicola Spinicelli[4], Camilla de Rossi[4], Mattia Boldrini[6], Paolo Ruggi[4] and Tom Oomen[1,7]

[1] Department of Mechanical Engineering, Control Systems Technology, Eindhoven University of Technology, Eindhoven, The Netherlands
[2] TNO, Optomechatronics Department, Delft, The Netherlands
[3] Nikhef, Amsterdam, The Netherlands
[4] European Gravitational Observatory (EGO), I-56021 Cascina, Pisa, Italy
[5] INFN, Sezione di Genova, I-16146 Genova, Italy
[6] INFN, Sezione di Roma, Piazzale Aldo Moro 2, I-00185 Roma, Italy
[7] Delft Center for Systems and Control, Delft University of Technology, Delft, The Netherlands
**Author to whom any correspondence should be addressed.

E-mail: m.r.v.dael@tue.nl

Keywords: gravitational waves, Advances Virgo Plus, online subtraction, adaptive filtering, parametrization selection

Abstract

Noise from auxiliary subsystems, amplified by their own control system, can couple to the output signal of gravitational wave detectors, limiting the maximum attainable sensitivity. Subtraction filters can be used to mitigate this coupling of noise by adding a secondary disturbance path with the purpose of canceling the noise in the output of the detector. The aim of this paper is to develop a systematic approach for the design and online adaptive estimation of subtraction filters. The proposed method adaptively updates the subtraction filter without the need for external perturbations to the system, providing a robust approach towards handling the time-varying couplings in the system as well as reducing the need for detector downtime. The method is validated on a representative simulation of the Advanced Virgo+ gravitational wave detector, illustrating that the method is capable of suppressing the coupling of noise from an auxiliary subsystem while the coupling varies over time.

1. Introduction

Gravitational wave (GW) detectors such as Advanced Virgo+ (AdV+) (Acernese *et al* 2023) and Advanced LIGO (Aasi *et al* 2015) are large optical interferometers consisting of a vast number of mechanics and optics, which are subject to a wide variety of disturbances that can limit the sensitivity of the detector (Bersanetti *et al* 2021, Capote *et al* 2025). Beyond the common approach of improving the hardware, data-driven approaches are used that use a measurement of some noise to subtract it in the output of the detector (Vajente 2008, Coughlin *et al* 2016, Tapia *et al* 2023). A prime example is noise amplified by the control loops of auxiliary degrees of freedom (DoFs) coupling to the main DoF containing the GW signal. Here, a subtraction filter is required that operates between the two feedback loops to cancel the coupling of the auxiliary DoF to the main DoF during operation of the system (Swinkels *et al* 2008, Meadors *et al* 2014). Because the dynamics of these loops are time-varying (van Dael *et al* 2024), the subtraction filter requires regular updating to maintain optimal performance.

The baseline approach is to identify the subtraction filters by injecting a perturbation signal and fitting a filter on the measured frequency response of the subtraction filter (Swinkels *et al* 2008, Meadors

* This work has been funded by the Netherlands Organisation for Scientific Research (NWO) under Grant Number 680.92.18.02.

et al 2014, Tapia *et al* 2023). While this method typically achieves a good performance, it requires updating the filter faster than the time-variation of the system dynamics, which is practically infeasible when the dynamics fluctuate in a matter of minutes. Additionally, this method requires injecting external perturbations, which leads to downtime of the detector. There are several approaches that do not use external perturbations such as in Giaime *et al* (2003), Tiwari *et al* (2015), Coughlin *et al* (2016) for linear dynamics and Vajente *et al* (2020) for non-linear dynamics, but these approaches are not designed to work in-loop during operation of the system. In Driggers *et al* (2012), an adaptive approach is presented to estimate the subtraction filter using the least mean square (LMS) algorithm. This approach shows promising results, but the algorithm does not consider the closed-loop aspects of two coupled DoFs, which may potentially lead to suboptimal performance and instability of the loops.

Although significant advances have been made in estimating subtraction filters online, there is no approach that can be used to subtract noise from an auxiliary loop to the primary loop containing the GW signal during operation of the system, without the need for dedicated external perturbations. The aim of this paper is therefore to develop a framework for adaptive online estimation of the subtraction filter, using pre-existing signals in the detector, with the objective of minimizing the time-varying coupling from an auxiliary DoF to the sensitivity of the detector.

The contribution of this paper is threefold.

(1) Exploration of design considerations for a static subtraction filter, providing a practical guide for selection and design.
(2) Development of an adaptive control approach to estimate the subtraction filter online using only pre-existing detector signals.
(3) Validation of the approach on a representative simulation of the AdV+ detector.

The problem of designing the subtraction filter admits several perspectives as a class of control problems. First, as a decoupling problem, where typically static input-output transformations are designed to approximately diagonalize the multiple-input multiple-output (MIMO) dynamics, either locally in a frequency range with similar dynamics (MacFarlane 1970, Owens 1978), or in a broader frequency band when the dynamics differ (Vaes *et al* 2004, Stoev *et al* 2017). This idea is extended to some of the auxiliary loops in AdV+, which exhibit time-varying dynamics, for which the static decoupling matrix was iteratively updated (van Dael *et al* 2024). Second, as a MIMO feedback design problem that treats the subtraction path as the off-diagonal controller element K_{21}, for which $\mathcal{H}_2/\mathcal{H}_\infty$ formulations (Doyle *et al* 1989, Skogestad and Postlethwaite 2005) could be used, or data-driven iterative feedback tuning (Hjalmarsson *et al* 1998) to accommodate the time-varying dynamics. Third, as a disturbance feedforward problem, where a measured exogenous disturbance is used to identify a feedforward filter. Several approaches have been proposed for this class of problems, such as in Sternad (1990) where an LQG-like approach is used to estimate the feedforward filter, or variations of the LMS algorithm (White and Tomizuka 1997, Driggers *et al* 2012, Hakvoort and Beijen 2023). And finally as a synchronization problem between subsystems (Mishra *et al* 2008, Navarrete *et al* 2015, Evers *et al* 2019), where an off-diagonal control term is used to minimize the relative error between subsystems.

The approach presented in this paper is based on the disturbance feedforward and synchronization class of approaches. The synchronization perspective is used as a basis for the parametrization of the subtraction filter, and iterative learning control (ILC) (Bristow *et al* 2006) is a disturbance feedforward based approach used to recursively estimate the subtraction filter. ILC updates a feedforward signal from batches of data and basis-functions ILC (BILC) (Phan and Frueh 1996, van de Wijdeven and Bosgra 2010) is adopted, which represents the filter as an FIR expansion whose coefficients are learned. BILC does not require dedicated perturbations, accommodates time variation through iterative updates, and provides a transparent weighting-based tuning suited to GW-specific requirements. The objective is to apply BILC to estimate the subtraction filter in GW detectors, providing a framework for its implementation on which extensions can be built. Establishing benchmarks with respect to other methods is left to future work.

The outline of the paper is as follows. Section 2 formalizes the control problem. Section 3 discusses the design considerations for each parametrization of the subtraction filter. Section 4 presents the adaptive control approach to estimate the subtraction filter online. Section 5 validates the approach on a representative simulation of two DoFs in the AdV+ detector and section 6 concludes the paper and discusses future work.

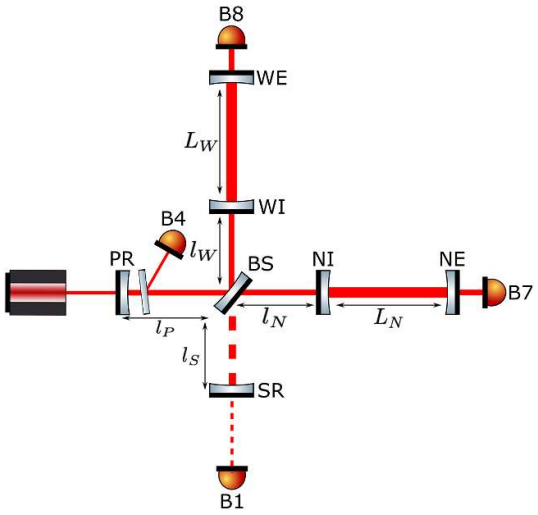

Figure 1. Part of the optical configuration of Advanced Virgo Plus for the O4 science run. The B1 photodiode measures the interference pattern between the two arms ($l_W + L_W$ and $l_N + L_N$), which change in length with opposite signs when a gravitational wave passes.

2. Problem formulation

This section first introduces the working principle of interferometric GW detectors and the challenge posed by noise coupling between control loops. It then motivates the need for a new approach to estimate the subtraction filter by outlining the limitations of current approaches, and finally formally states the problem addressed in this work.

2.1. System description

Figure 1 shows a simplified scheme of the optical configuration of AdV+. The detector measures a relative length change between the two 3 km long arms L_W and L_N when a GW passes, which induces a change in the interference pattern on the B1 photodiode. This is the so-called Differential Arm (DARM) DoF, i.e.

$$L_{\text{DARM}} = L_N - L_W, \tag{1}$$

which is the most important DoF as it contains the GW signal. Additionally, there are four other so-called longitudinal DoFs, which are defined as

$$\begin{aligned} L_{\text{CARM}} &= \frac{L_N + L_W}{2}, \\ L_{\text{MICH}} &= l_N - l_W, \\ L_{\text{PRCL}} &= l_P + \frac{l_N + l_W}{2}, \\ L_{\text{SRCL}} &= l_S + \frac{l_N + l_W}{2}, \end{aligned} \tag{2}$$

which serve as auxiliary DoFs. The Michelson (MICH), power recycling cavity length (PRCL), and signal recycling cavity length (SRCL) DoFs are actively controlled using a feedback loop that uses PDH signals (Black 2001) to determine a control signal that is distributed among the mirrors to control the cavity lengths, while the common arm (CARM) loop uses a more complex control scheme (van Dael *et al* 2025), the discussion of which is out of scope for this paper.

Beyond the basic design requirement that these feedback controllers should achieve a certain RMS error on the cavity length, the noise coupling from the auxiliary DoFs to DARM should also be minimized. In figure 2, the spectrum of DARM and the noise contributions from the auxiliary DoFs to DARM are shown, while using the pre-existing subtraction filters. The detection band of AdV+ is between 10 Hz and 10 kHz and though the auxiliary longitudinal DoFs are not limiting the spectrum of DARM except for a few frequency bins, they are still close to limiting, which might become a problem in the future when the sensitivity of the detector improves. Furthermore, since the coupling varies over time, some DoFs may become arbitrarily close to limiting the sensitivity.

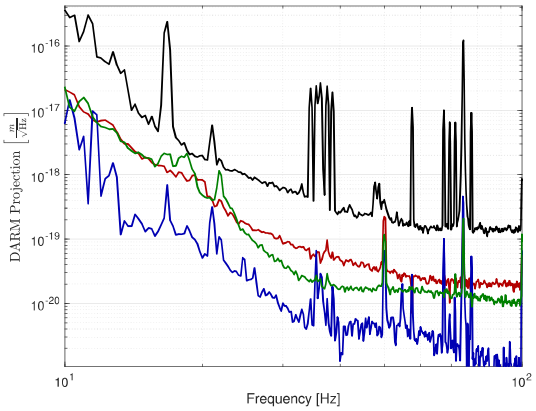

Figure 2. Noise projection of the auxiliary longitudinal degrees of freedom to DARM as present for AdV+ in O4 science run, with (▬) the spectrum of DARM and the contributions to DARM of (▬) MICH, (▬) PRCL and (▬) SRCL, while using a subtraction filter for MICH. Both MICH and SRCL are close to limiting the sensitivity of DARM in the 10 Hz to 30 Hz region.

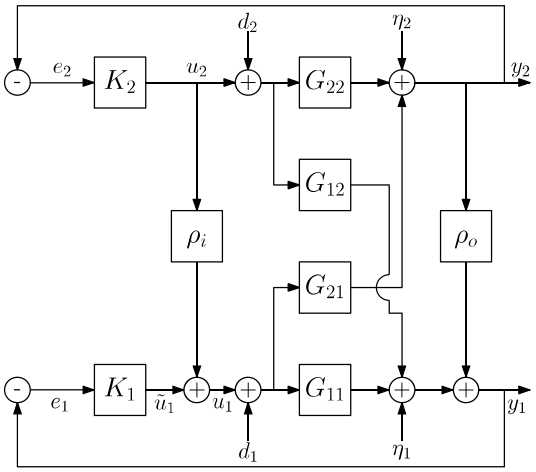

Figure 3. Block diagram for a general two degree of freedom multiple-input multiple-output system. The subtraction filter can be parametrized to be between the control signals u as ρ_{i} or between the output signals y as ρ_{o}, to suppress either the input disturbance d_2 or the output disturbance η_2, depending on the choice of parametrization.

2.2. Subtraction filter design

The block diagram in figure 3 shows a simplified version of the feedback loops for the DARM DoF as the lower loop (loop 1) and one of the auxiliary DoFs as the upper loop (loop 2). The plant G_{ij} represents the opto-mechanical response from DoF j to DoF i and both loops are controlled by a feedback controller K_i. Both loops are furthermore subject to input disturbances d_i and output disturbances η_i and the feedback controllers are designed to reject the disturbances in their respective error signals e_i.

The output signal y_1 of the DARM loop in figure 3 additionally contains the GW signal. To minimize the noise propagation from d_2 or η_2 to y_1, the subtraction filter ρ_{i} or ρ_{o} is used depending on which noise source is dominant. In AdV+, the output disturbance η_2 (sensor noise) is typically the limiting noise source for the auxiliary DoFs in the 10 to 50 Hz region and the subtraction filter ρ_{i} is used to minimize the coupling of η_2 to y_1 by parametrizing ρ_{i} as

$$\rho_{\mathrm{i}} = -\frac{G_{12}}{G_{11}}. \tag{3}$$

Note that the method of obtaining this parametrization will be discussed in section 3. To obtain the optimal filter in practice, standard frequency response function (FRF) identification is used to obtain an FRF of G_{12} and G_{11} and a parametric model is then fitted onto the FRF of (3) to obtain the subtraction filter in the frequency band where the auxiliary DoF is limiting. This method works well in practice, reducing the coupled noise by up to 2 orders of magnitude for some DoFs (Swinkels *et al* 2008).

2.3. Problem statement

However, the dynamics of the longitudinal DoFs exhibit time-varying behavior (van Dael *et al* 2024), stemming from thermal effects as well as slow drifts in the alignment of the mirrors. Consequently, to maintain the optimal performance over time, the filter in (3) should also be time-varying. Modeling this effect is practically impossible due to the complexity of the system, so in the current procedure, the filter is regularly updated by re-identifying it using the FRF method. While the attained performance over time is still good, the efficacy of the method is limited by the time-varying behavior, which if relatively fast cannot be accounted for using the current method. Continuously updating the subtraction filter using the current method is not possible either, since the injection of noise obscures the GW signal.

The problem addressed in this paper is two-fold. The first is the problem of designing the subtraction filter, reviewing its requirements, selection of parametrization and design considerations, since the method is already widely adopted in the GW community. The second problem is to develop a method for the online estimation of the subtraction filter that can recursively adapt to the time-varying behavior of the system without requiring the injection of external perturbation signals.

3. Parametrizing subtraction filter

This section first presents the possible parametrizations of the subtraction filter, after which several design considerations are discussed as well as the possibility of simultaneously attenuating both input and output disturbances. The section concludes with a practical guideline on how to select and design the subtraction filter. Throughout this section, a uni-directional parametrization is assumed, i.e. using only a subtraction filter from loop 2 to loop 1. The implications of using bi-directional parametrizations where also a subtraction filter from loop 2 to loop 1 is added are discussed at the different design considerations but not further explored.

To simplify the derivations of the transfer functions and to discuss the design considerations, the MIMO representation of the two-dimensional system is used, which is shown in figure 4. Here, $G \in \mathbb{C}^{2\times 2}$ and $K \in \mathbb{C}^{2\times 2}$ are the two-dimensional optomechanical plant and controller, respectively, and $T_u \in \mathbb{R}^{2\times 2}$ and $T_y \in \mathbb{R}^{2\times 2}$ the input and output transformations given by

$$T_u = \begin{bmatrix} 1 & \rho_\mathrm{i} \\ 0 & 1 \end{bmatrix}, \quad T_y = \begin{bmatrix} 1 & \rho_\mathrm{o} \\ 0 & 1 \end{bmatrix}, \tag{4}$$

with ρ_i and ρ_o as derived in (6), respectively. The MIMO transfer matrices based on figure 4 from d and η to y are

$$y = \underbrace{\left(I + T_y G T_u K\right)^{-1} T_y G}_{SG} d + \underbrace{\left(I + T_y G T_u K\right)^{-1} T_y}_{S} \eta, \tag{5}$$

with $I \in \mathbb{R}^{2\times 2}$ the identity matrix.

3.1. Possible subtraction filter parametrizations

Throughout this section it is assumed that the objective is to minimize the contributions of the disturbances d_2 and η_2 to the error signal y_1, since this is the measured variable containing the GW signal[8]. The contribution of d_2 and η_2 to y_1 are obtained by taking the $1,2$ elements of SG and S as defined in (5), which are given by

$$y_1 = \frac{G_{12} + G_{22}\rho_\mathrm{o} - G_{11}G_{22}K_2\rho_\mathrm{i} + G_{12}G_{21}K_2\rho_\mathrm{i}}{X} d_2 + \frac{\rho_\mathrm{o} - G_{12}K_2 - G_{11}K_2\rho_\mathrm{i}}{X} \eta_2, \tag{6}$$

with

$$X = 1 + G_{11}K_1 + G_{22}K_2 + G_{11}G_{22}K_1K_2 - G_{12}G_{21}K_1K_2 + G_{21}K_2\rho_\mathrm{i} + G_{21}K_1\rho_\mathrm{o}. \tag{7}$$

Equation (6) shows either ρ_i or ρ_o can be used to minimize the coupling of d_2 and η_2 to y_1. To find the optimal filters, the terms in (6) are set to zero and the resulting filter parametrizations are shown in table 1, where using only ρ_i or ρ_o is considered to suppress either d_2 or η_2.

[8] In practice $\tilde{u}_1$ is also used to extract the GW signal. Since any noise subtracted from y_1 will also not be present in $\tilde{u}_1$, only y_1 is considered here.

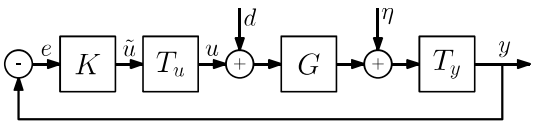

Figure 4. Block diagram for a general Multiple-Input Multiple-Output system. The subtraction filter ρ_i and ρ_o are part of the input and output transformation matrices T_u and T_y, respectively.

Table 1. Optimal subtraction filters ρ_i and ρ_o for input and output disturbances d_2 and η_2.

	d_2	η_2
ρ_i	$\dfrac{G_{12}}{G_{11}G_{22}K_2 - G_{12}G_{21}K_2}$	$-\dfrac{G_{12}}{G_{11}}$
ρ_o	$-\dfrac{G_{12}}{G_{22}}$	$G_{12}K_2$

Beyond these four uni-directional parametrizations, there are many more parametrizations by selecting different combinations of input and output signals for the subtraction filters. They however do not offer any additional benefits, which is why the paper is restricted to these four parametrizations. Nevertheless, the derivation of the design considerations discussed next readily extend to any other possible parametrization. Bi-directional parametrizations offer added benefits in terms of noise suppression (Evers *et al* 2019) at the expense of a more complex design procedure, but this is out of scope for this paper.

3.2. Design considerations

Throughout the remainder of this section, several design considerations will be discussed that should be taken into account when selecting the subtraction filter parametrization.

Amplification. While the subtraction filter minimizes the propagation of one disturbance from loop 2 to loop 1, it can simultaneously amplify another. Furthermore, when $G_{21} \neq 0$, the subtraction filter potentially amplifies disturbances from loop 1 to loop 2, which may be undesired.

To analyze these effects, the amplification matrix $\mathcal{A}_{ij}(T_u, T_y)$ is used, which is defined as follows.

Definition 1. Let $M(T_u, T_y)$ denote a transfer matrix that depends on designable transfer functions T_u and T_y, and let $M(I,I)$ be the same transfer function matrix when both T_u and T_y are identity. The elements $\mathcal{A}_{ij}$ of the amplification matrix $\mathcal{A}$ are then defined as

$$\mathcal{A}_{ij}\left(T_u, T_y\right) = \frac{M_{ij}\left(T_u, T_y\right)}{M_{ij}\left(I,I\right)}, \tag{8}$$

where M_{ij} are the elements of M.

Remark 1. If $M_{ij}(I,I) = 0$, the amplification matrix is undefined and $\mathcal{A}_{ij}(T_u, T_y)$ should therefore be assessed separately.

Using the parametrization $\rho_\mathrm{i} = -\frac{G_{12}}{G_{11}}$ to minimize the coupling of η_2 to y_1, the amplification matrix of η in y is then according to definition 1 given by

$$\begin{aligned}\mathcal{A}^{\eta}\left(T_u, I\right) &= \frac{M(T_u, I)}{M(I,I)} \\ &= \begin{bmatrix} \frac{1}{1+G_{11}K_1}Q & 0 \\ \frac{G_{11}}{G_{11}+K_2G_{11}G_{22}-K_2G_{12}G_{21}}Q & \frac{G_{11}}{G_{11}+K_2G_{11}G_{22}-K_2G_{12}G_{21}}Q \end{bmatrix},\end{aligned} \tag{9}$$

with

$$Q = \frac{1 + G_{11}K_1 + G_{22}K_2 + G_{11}G_{22}K_1K_2 - G_{12}G_{21}K_1K_2}{1 + G_{11}K_1}, \tag{10}$$

and where based on (5)

$$M(T_u, I) = \left(I + GT_uK\right)^{-1}. \tag{11}$$

Given the assumption that $G_{21} = 0$, (9) reduces to

$$\mathcal{A}^{\eta}(T_u, I) = \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix}, \tag{12}$$

illustrating that for uni-directional coupling, the parametrization for ρ_{i} perfectly cancels η_2 in y_1 without amplifying other disturbance paths. However, when $G_{21} \neq 0$, the other disturbance paths may be amplified depending on the system dynamics. The significance of this amplification should therefore be assessed per case.

To illustrate the impact of amplifying the other disturbance in loop 1, consider again the case where $\rho_{\mathrm{i}} = -\frac{G_{12}}{G_{11}}$ is used to cancel η_2 and first consider the case where $G_{21} = 0$. The resulting amplification matrix of d in y is then given by

$$\mathcal{A}^{d}(T_u, I) = \frac{M(T_u, I)}{M(I, I)} = \begin{bmatrix} 1 & 1 + G_{22}K_2 \\ 1 & 1 \end{bmatrix}, \tag{13}$$

showing that the coupling of d_2 to y_1 has worsened as a result of the subtraction filter. The amplification term $1 + G_{22}K_2$ is furthermore large below the bandwidth of loop 2, since there $G_{22}K_2 \ll 1$. This means d_2 is significantly amplified below the bandwidth of loop 2, with the quantity depending on the design of K_2 and the dynamics of G_{22}. A solution could be to have a bandpass in the subtraction filter around the frequencies where η_2 is limiting, which would work if this frequency region is around or above the bandwidth of the second loop.

To illustrate the impact of the subtraction filter choice on the second loop, now assume that $G_{21} \neq 0$. The amplification matrix of d in y is then given by

$$\mathcal{A}^{d} = \frac{M(T_u, I)}{M(I, I)} = \begin{bmatrix} \frac{G_{11}}{G_{11} + K_2(G_{11}G_{22} - G_{21}G_{12})}Q & Q \\ \frac{G_{11}}{G_{11} + K_2(G_{11}G_{22} - G_{21}G_{12})}Q & \frac{G_{11}}{G_{11} + K_2(G_{11}G_{22} - G_{21}G_{12})}Q \end{bmatrix}. \tag{14}$$

While it is difficult to assess the exact amplification due to the complexity of the amplification matrix, it is clear that designing ρ_{i} to minimize the coupling of η_2 to y_1 also has an impact on the coupling of d_1 and η_1 to y_2 and even of d_1 to y_1.

When a parametrization is chosen, the amplification of the other disturbance paths is one of the main priorities in assessing the best parametrization. The impact depends on the requirements of each loop, the system dynamics and the size of the different disturbances, and this should therefore be assessed for each case separately.

MIMO stability. When the system has uni-directional coupling ($G_{21} = 0$), the stability of the system remains stable when the selected subtraction filter parametrization is also stable, because there is no loop between the two individual loops that can become unstable. However, when the system has bi-directional coupling ($G_{21} \neq 0$), the stability of the system is affected by the choice of subtraction filter. The MIMO poles of a system are given by

$$\begin{aligned} X = \det\left(I + T_y G T_u K\right) &= (1 + G_{11}K_1)(1 + G_{22}K_2) \\ &\quad - G_{12}G_{21}K_1K_2 + G_{21}K_2\rho_{\mathrm{i}} + G_{21}K_1\rho_{\mathrm{o}}, \end{aligned} \tag{15}$$

with T_u and T_y parametrized as in (4). If, e.g. as parametrization $\rho_{\mathrm{o}} = G_{12}K_2$ is used, (15) becomes

$$\begin{aligned} \det\left(I + T_y G T_u K\right) &= (1 + G_{11}K_1)(1 + G_{22}K_2) - G_{12}G_{21}K_1K_2 + G_{21}K_1G_{12}K_2 \\ &= (1 + G_{11}K_1)(1 + G_{22}K_2). \end{aligned} \tag{16}$$

The last equality shows that, for this particular parametrization, stability only depends on the stability of the individual loops, which are designed to be stable. For other parametrizations, (15) may not necessarily reduce to this form, requiring stability to be assessed using MIMO stability criteria, see e.g. Skogestad and Postlethwaite (2005). But even in the case where (15) does reduce to this form, it is important to note that it only holds when a perfect fit of the optimal parametrization is obtained. Given the time-varying dynamics of a system, this assumption is not guaranteed in practice and stability issues may therefore still ensue.

Other design considerations. Beyond the amplification and stability considerations, there are other aspects that should be considered. Since the current method relies on FRF identification, an obvious precondition is that the subtraction filter can be identified in the frequency range where the disturbance

is limiting. Some parametrizations might therefore be more beneficial since they only contain terms that are easy to identify. Furthermore, the order of the subtraction filter might also impact the choice of parametrization, since higher order parametrizations with complex dynamics may be more difficult to fit. However, it could also be that some terms share common dynamics and thus cancel, which would simplify the required fit and potentially yield better subtraction levels. Choosing the best parametrization is therefore completely dependent on the system dynamics and disturbances present.

3.3. Simultaneous cancellation of disturbances

As shown next, unfortunately, there exists no parametrization for which both the input and output disturbance are canceled simultaneously. Consider the case where the goal is to have both disturbances d_2 and η_2 perfectly canceled in y_1 and suppose T_u and T_y are parametrized such that $S_{12} = 0$, meaning η_2 is perfectly canceled in y_1. To then also perfectly cancel d_2 in y_1 requires the $1,2$ element of SG to be zero, i.e.

$$\begin{bmatrix} S_{11} & 0 \\ S_{21} & S_{22} \end{bmatrix} \begin{bmatrix} G_{11} & G_{12} \\ G_{21} & G_{22} \end{bmatrix} = \begin{bmatrix} S_{11}G_{11} & S_{11}G_{12} \\ S_{21}G_{11} + S_{22}G_{21} & S_{21}G_{12} + S_{22}G_{22} \end{bmatrix}. \tag{17}$$

The only way to achieve this is if $S_{11} = 0$, which would induce rank deficiency in the matrix on the right of the equality sign in (17). Practically, this would mean that the first DoF is not controlled, which for obvious reasons is not desirable. Note that no assumption on the structure of T_u and T_y is made, so even a parametrization in T_u and T_y that uses all of its elements would not work.

Perfect cancellation of both disturbances is considered here, while typically the requirement for GW detectors is that the disturbances should just be sufficiently low as to not limit the sensitivity. Therefore, for this situation, it is likely possible to find a parametrization that sufficiently suppresses both disturbances. Additionally, only uni-directional parametrizations are considered here, while bi-directional parametrizations provide more flexibility and frameworks for the designs of such parametrizations already exist, see e.g. Doyle *et al* (1989), Skogestad and Postlethwaite (2005), Evers *et al* (2019). These frameworks could be explored to find bi-directional parametrizations that sufficiently suppress both disturbances.

3.4. Design guidelines

Based on the design considerations presented in this section, a brief guideline on the choice of parametrization is given in a step-by-step process.

(1) Determine which disturbance (d_2 or η_2) is more critical to suppress. If both need to be suppressed in the same frequency band, extensions of the current design methods should be explored.
(2) If the system is bi-directionally coupled, choose a parametrization that does not affect the MIMO stability of the system.
(3) The choice of parametrization then depends in order of importance on the following considerations:
    (a) Use the amplification matrix to determine which of the remaining parametrizations may lead to an undesired level of amplification of the other disturbance.
    (b) Choose the parametrization that requires the least number of parameters to fit, since this is likely to yield a more accurate estimate. The parametrization $\rho_\mathrm{o} = G_{12}K_2$ likely requires the least number of parameters to estimate option since K_2 is already known, but the parametrization $\rho_\mathrm{i} = -\frac{G_{12}}{G_{11}}$ may for example lead to cancellation of dynamics and could therefore be a better option.

4. Adaptive estimation of subtraction filter

This section will derive a recursive approach for an adaptive estimation of the subtraction filter online using pre-existing perturbations in the system.

Throughout this section, the convolution matrix of a transfer function is used. Consider a transfer function G with finite impulse response (FIR) of length N and input and output signals u and y, respectively. The finite time response of G is given by

$$y[k] = \sum_{l=1-N}^{k} g[l]\, u[k-l], \tag{18}$$

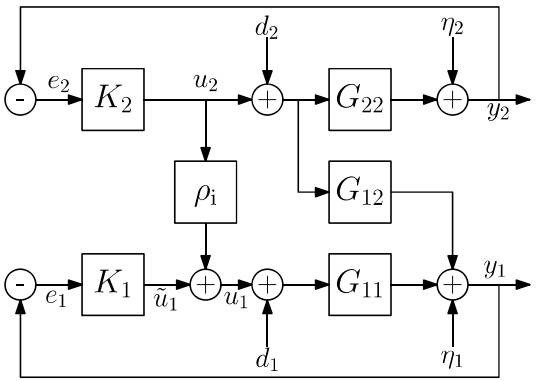

Figure 5. Simplified block diagram of the control problem where uni-directional coupling is assumed, i.e. $G_{21} = 0$. The objective of the ILC controller is to find the optimal subtraction filter ρ.

with $0 \leqslant k \leqslant N$ and the finite-time convolution matrix is given by

$$J = \begin{bmatrix} g[0] & 0 & \dots & 0 \\ g[1] & g[0] & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ g[N-1] & g[N-2] & \dots & g[0] \end{bmatrix}, \tag{19}$$

where J is the convolution matrix representation of G.

4.1. Design assumptions and control problem

In this section, the assumption is made that the system is uni-directionally coupled, i.e. $G_{21} = 0$. This assumption typically holds true for GW detectors since the DARM DoF for which the subtraction method is almost exclusively used has actuators and sensors which have orders of magnitude lower noise floors than those of the auxiliary DoFs. In practice, the coupling from DARM to the auxiliary DoFs is therefore always negligible, i.e. $G_{21} \approx 0$. An advantage of this assumption is that it simplifies an important aspect, since for this case u_2 will be uncorrelated with any disturbance from loop 1, i.e. d_1 and η_1. While an analysis of the implications for the method when this assumption does not hold is beyond the scope of this paper, this assumption is thus not expected to limit the application of this method to GW detectors.

Furthermore, since in AdV+ the most common parametrization is to use $\rho_i = -G_{12}/G_{11}$ to suppress η_2 in y_1, the same use-case is considered here for the derivation of the method. If $G_{21} = 0$ is substituted in (6), the following coupling is obtained

$$y_1 = S_1 S_2 \left(G_{12} - G_{11} G_{22} K_2 \rho_i\right) d_2 - S_1 S_2 K_2 \left(G_{12} + G_{11} \rho_i\right) \eta_2, \tag{20}$$

with

$$S_i = \frac{1}{1 + G_{ii} K_i}, \quad i = 1, 2. \tag{21}$$

The block diagram of figure 3 can then be simplified to the block diagram in figure 5. Note that in this block diagram the only available signals are u_2, $y_1 = -e_1$ and u_2, so no direct measurements of the disturbances are available.

To find the optimal subtraction filter ρ_i, a batch of data is taken containing the input signal of loop 2 of finite length N, i.e, $\mathbf{u_2^j} = [u_2^j(0), u_2^j(1), \dots, u_2^j(N-1)]^\top$, and output signal $\mathbf{y_1^j} = [y_1^j(0), y_1^j(1), \dots, y_1^j(N-1)]^\top$, where j is the batch index. The objective is to find ρ^{j+1} using $\mathbf{u_2^j}$, $\mathbf{y_1^j}$ and ρ^j, where ρ^{j+1} approaches the optimal filter $\rho^{\text{opt}} = -G_{12}/G_{11}$ as $j \to \infty$.

4.2. Subtraction filter parametrization

To estimate the subtraction filter online, the filter is parametrized using a set of basis functions and corresponding coefficients, i.e.

$$f^j = \Psi^j \theta^j, \tag{22}$$

with f^j the output of ρ^j, Ψ^j a set of basis functions constructed using the input signal $\mathbf{u_2^j}$ and $\theta^j \in \mathbb{R}^{n_\theta}$ the parameter vector containing the coefficients of the basis functions. The objective of the algorithm is then to find the optimal coefficients θ^{j+1} at each iteration j such that the output signal $\mathbf{y}_1^{\mathbf{j+1}}$ is minimized.

There are many possible choices for the basis functions Ψ^j, such as polynomial basis functions (van de Wijdeven and Bosgra 2010) and rational basis functions (Bolder and Oomen 2015, Blanken *et al* 2017, Poot *et al* 2023, Ickenroth *et al* 2025) with different basis function structures depending on the application and model structure. The choice of basis functions used here is to use polynomial basis functions, i.e.

$$f^j(k) = \sum_{l=0}^{n_\theta - 1} \theta_l^j q^{-l} u_2^j(k) = \sum_{l=0}^{n_\theta - 1} \theta_l^j u_2^j(k-l), \tag{23}$$

with q the shift parameter in time, which represents an FIR structure with θ^j the FIR coefficients. The choice for this parametrization is motivated by the fact that the filter is inherently stable and straightforward to implement.

4.3. ILC with basis functions

This section presents the ILC algorithm used to obtain the model parameter estimates θ^{j+1}, which are obtained by minimizing a cost function $\mathcal{J}(\theta^{j+1})$ containing a batch of data, i.e.

$$\theta^{j+1} = \arg\min_{\theta^{j+1}} \mathcal{J}\left(\theta^{j+1}\right). \tag{24}$$

The cost function $\mathcal{J}(\theta^{j+1})$ is based on the norm-optimal ILC approach (Lee *et al* 2000, Gunnarsson and Norrlöf 2001, van de Wijdeven and Bosgra 2010) and is defined as the weighted contributions of the output signal y_1^j and the subtraction filter output f^j, i.e.

$$\mathcal{J}\left(\theta^{j+1}\right) = \|y_1^{j+1}\|_{W_y}^2 + \|f^{j+1}\|_{W_f}^2 + \|f^{j+1} - f^j\|_{W_{\Delta f}}^2, \tag{25}$$

with $\|x\|_W^2 = x^\top W x$ and user-defined weighting matrices $W_y \succ 0$, W_f, $W_{\Delta f} \succeq 0$. This cost function is linear in the parameter vector θ^{j+1} and therefore has the following closed-form solution to update the parameter vector θ^{j+1}

$$\theta^{j+1} = Q\theta^j + L\mathbf{y}_\mathbf{1}^\mathbf{j}, \tag{26}$$

with

$$\begin{aligned} Q &= \left(\Psi^\top \left(J^\top W_y J + W_f + W_{\Delta f}\right)\Psi\right)^{-1} \Psi^\top \left(J^\top W_y J + W_f\right)\Psi, \\ L &= \left(\Psi^\top \left(J^\top W_y J + W_f + W_{\Delta f}\right)\Psi\right)^{-1} \Psi^\top J^\top W_y, \end{aligned} \tag{27}$$

where $L \in \mathbb{R}^{n_\theta \times N}$ and $Q \in \mathbb{R}^{N \times N}$ are the learning and robustness matrices, respectively and J is the finite-time convolution matrix of the process sensitivity PS given by

$$\mathrm{PS} = \frac{G_{11}}{1 + G_{11}K_1}. \tag{28}$$

After obtaining a batch of data, the new set of basis functions Ψ^j is constructed using this data and the parameter vector θ^{j+1} is then updated using (26), which can all be done offline.

4.4. Enforcing filter structure

In certain situations it is desired to force a particular structure in the filter. One example is when the subtraction filter should only work in a certain frequency band, which can be achieved by enforcing a bandpass structure outside this frequency range. Another example is when it is known that only parts of the dynamics change, e.g. the gain of the filter. For this situation a static filter may be implemented and only the gain is then updated using the adaptive algorithm.

For both situations, consider the following structure

$$\rho = \rho_{\text{static}} \cdot \rho_{\text{ILC}}, \tag{29}$$

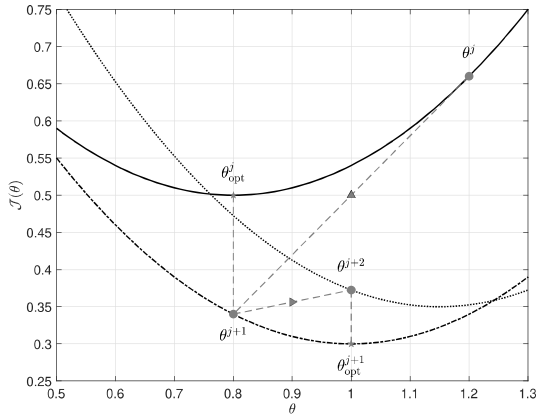

Figure 6. Illustration of fluctuations of ILC error after convergence using three example cost functions $\mathcal{J}(\theta)$ at trial j (▬), trial $j+1$ (▪ ▪) and trial $j+2$(⋯⋯). Starting at trial j, the parameter update for θ^{j+1} is computed using the input signal $\mathbf{u}_2^{\mathbf{j}}$ and the output signal $\mathbf{y}_1^{\mathbf{j}}$, which leads to a new parameter θ^{j+1}, which is at the optimal value θ^j_{opt} for trial j. The cost in this instance is lower for the cost function at trial $j+1$, but the same procedure for trial $j+1$ to trial $j+2$ leads to a higher cost even though the optimal parameter $\theta^{j+1}_{\text{opt}}$ is computed. The estimated parameter vector θ^j is therefore expected to fluctuate in a cloud of optima after convergence.

with ρ_{static} the static filter and ρ_{ILC} the adaptive filter. The same adaptive algorithm can be used for this situation by prefiltering the input signal $\mathbf{u}_2^{\mathbf{j}}$ with the static filter ρ_{static}, i.e.

$$\tilde{\mathbf{u}}_2^{\mathbf{j}} = \rho_{\text{static}}\, \mathbf{u}_2^{\mathbf{j}}, \tag{30}$$

where now $\tilde{\mathbf{u}}_2^{\mathbf{j}}$ is used to form the basis functions Ψ^j in (22) and the same update rule in (26) is used to update the adaptive filter ρ_{ILC}.

4.5. Convergence of the ILC algorithm

Under the condition that the input signal $\mathbf{u}_2^{\mathbf{j}}$ does not change, i.e. $\mathbf{u}_2^{\mathbf{j}} = \mathbf{u}_2^{\mathbf{j+1}}$ for all j, the algorithm has monotonic convergence if the following condition is satisfied

$$\bar{\sigma}\left(Q - LJ\Psi\right) < 1, \tag{31}$$

with

$$\bar{\sigma}\left(Q - LJ\Psi\right) = \bar{\sigma}\left(\left(\Psi^\top \left(J^\top W_y J + W_f + W_{\Delta f}\right)\Psi\right)^{-1} \Psi^\top W_{\Delta f}\Psi\right), \tag{32}$$

where $\bar{\sigma}(\cdot)$ denotes the maximum singular value of the matrix. Following this condition, if

$$\Psi^\top \left(J^\top W_y J + W_f + W_{\Delta f}\right)\Psi \succ 0, \tag{33}$$

then the Hessian of the cost function in (25) is positive definite with respect to θ^{j+1} and the cost function is therefore convex. This is the standard convergence condition for ILC with basis functions (van de Wijdeven and Bosgra 2010), but the condition $\mathbf{u}_2^{\mathbf{j}} = \mathbf{u}_2^{\mathbf{j+1}}$ does not hold for this system since the input signal is a stochastic signal.

As $\mathbf{u}_2^{\mathbf{j}}$ varies over time, the cost function $\mathcal{J}(\theta^{j+1})$ also changes over iterations j since it depends on $\mathbf{u}_2^{\mathbf{j}}$ through Ψ. Consequently, the parameter vector θ^j computed based on $\mathbf{u}_2^{\mathbf{j}}$ and $\mathbf{y}_1^{\mathbf{j}}$ is not necessarily optimal for the next iteration $j+1$ since $\mathbf{u}_2^{\mathbf{j+1}} \neq \mathbf{u}_2^{\mathbf{j}}$. This is graphically illustrated in figure 6, where the cost function $\mathcal{J}(\theta)$ is shown for different iterations j of the ILC algorithm. The cost function is convex but is different for each iteration j since $\mathbf{u}_2^{\mathbf{j}}$ is a realization of a random process. However, under this assumption that $\mathbf{u}_2^{\mathbf{j}}$ is a realization of a filtered white noise process, the optimal parameter vector θ^{opt} is expected to form a cloud of optima in the parameter space, and the algorithm is expected to converge to this cloud of optima. So while the algorithm will not have monotonic convergence, it is still expected to converge to this cloud of optima for the first few trials j, after which the parameter vector θ^j will fluctuate in this cloud of optima, as illustrated in figure 6. Though the mathematical proof is out of scope for this paper, this convergence behavior has been observed in simulations as will be shown in section 5, and is therefore expected to hold in practice.

4.6. Selecting weighting matrices

The weighting filters W_y, W_f, and $W_{\Delta f}$ in (25) are used to tune the ILC algorithm. The weighting filters are typically selected as diagonal matrices determining the relative importance of the different elements in the cost function, but a frequency-dependent weighting can also be applied to emphasize specific frequency regions.

The weighting matrix W_y emphasizes the relative importance of minimizing the output signal y_1^j in the cost function. In the case of a GW detector, subtraction filters are typically designed to minimize the impact of $\mathbf{u}_2^j$ in y_1^j in a limited frequency band, e.g. in the region where GWs are detected. The weighting matrix W_y can therefore be selected as a bandpass filter in this frequency region, which will then emphasize the minimization of the output signal in this frequency region.

The weighting matrix W_f is used to penalize the subtraction filter output f^j in the cost function and to increase robustness for model uncertainty. Particularly for the case of GW detectors where subtraction is only needed in a limited frequency band, it is also desired to have the subtraction filter behave as a bandpass filter such that it does not significantly worsen y_1^j outside the desired frequency band. To penalize large values of f^j outside the desired frequency band, the weighting matrix W_f can be selected as the inverse of a bandpass filter, i.e. large amplitudes outside the desired frequency band such that ρ is small in this region.

The weighting matrix $W_{\Delta f}$ is typically chosen as a static matrix and is used to penalize large changes in the subtraction filter between iterations, which influences the convergence speed and the sensitivity to trial varying disturbances.

4.7. ILC on output

The algorithm could also be applied using the subtraction filter on the output ρ_o. The only difference in the ILC algorithm is that J is replaced with the convolution matrix of the sensitivity function S_1 as defined in (21) instead of the process sensitivity in (28). The update rule and design of the weighting matrices all remain the same in this case.

5. Simulation results

This section presents a simulation of the ILC algorithm for the subtraction of SRCL noise coupling to DARM, using representative dynamics for this coupling of the AdV+ detector. The motivation for choosing SRCL as the auxiliary DoF is that this DoF is known to vary significantly over time, more than most other auxiliary DoFs. The simulated models used for the coupling term are therefore based on measurements of the coupling of SRCL to DARM in AdV+, providing a relevant use-case to assess the potential of the proposed approach. The proposed approach is nevertheless designed to be generally applicable to any auxiliary DoF coupling to DARM.

The same design assumptions as noted in section 4.1 are used, i.e. the system is uni-directionally coupled and the subtraction filter ρ_i is used to suppress η_2 in y_1. Additionally, because of the uni-directional coupling, the estimated subtraction filter does not influence stability, see section 3.2.2, so this aspect is not further considered in this section. The goal of the simulation is to show that the ILC algorithm can be used to estimate the subtraction filter ρ_i such that the coupling of η_2 to y_1 is minimized. In particular, it will be shown that the ILC algorithm can track the changes in the optimal subtraction filter by iteratively updating the filter based on a set of data. An important assumption that is made here is that the subtraction filter is constant within a set of trial data, i.e. the changes in the optimal subtraction filter are slow compared to the duration of a trial. This assumption is expected to hold true in practice since the dynamics of GW detectors typically vary on timescales of minutes to hours, while the duration of a trial is envisioned to be in the order of seconds.

5.1. Simulation setup

A representative simulation of the AdV+ dynamics and in particular the SRCL to DARM coupling is made using the block diagram in figure 5. The transfer functions G_{11} and G_{12} are estimated based on noise injections in AdV+ and the models are shown in figure 7. For the SRCL to DARM coupling G_{12}, three models are derived which represent a similar situation to AdV+ where the anti-resonance location shifts over time. The model for G_{12} is given by

$$G_{12}(s,\omega_{ar}) = 2\frac{\left(s^2+0.2\omega_{ar}s+\omega_{ar}^2\right)}{\left(s^2+31.42s+987\right)\left(s^2+10.05s+157.9\right)}, \tag{34}$$

with $\omega_{ar} \in 2\pi\,[20, 30, 40]$. For simulation purposes, the gains of the modelled transfer functions and disturbances are not calibrated, but this does not affect the purpose of the simulation. The disturbances

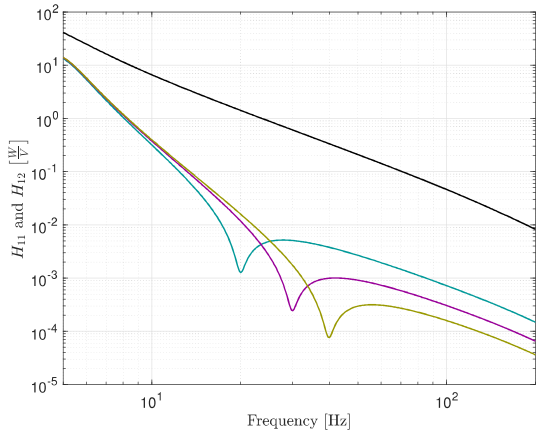

Figure 7. Modeled transfer function for G_{11} representing the DARM transfer function (▬) and models for G_{12} with 20 Hz anti-resonance (▬), 30 Hz anti-resonance (▬) and 40 Hz anti-resonance (▬).

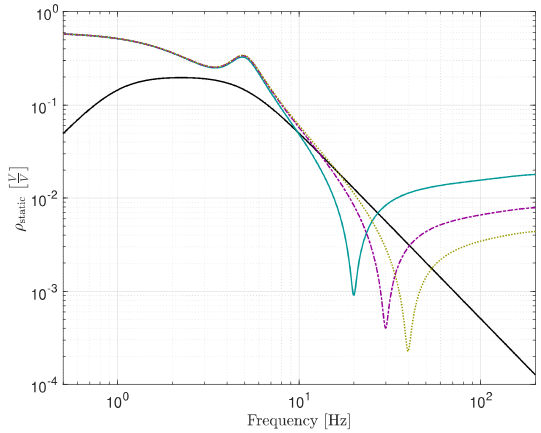

Figure 8. Pre-filter ρ_{static} (▬) and optimal subtraction filters $-G_{12}/G_{11}$ for the 20 Hz anti-resonance model (▬), 30Hz anti-resonance model (▬) and 40 Hz anti-resonance model (▬).

d_1 and η_1 are furthermore also included in the simulation. For d_1, the model as used in van Dael *et al* (2025) is used, where for numerical purposes only a single suspension stage is considered, and the shot noise values for η_2 are chosen such that η_2 only limits y_1 between roughly 10 and 70 Hz. The goal is to estimate ρ_i such that the coupling of η_2 to y_1 is minimized in the 10 to 70 Hz frequency region.

5.2. Pre-filter and weighting functions

The goal of the pre-filter and the weighting functions are to ensure that the subtraction filter ρ_i has a bandpass structure around the 10 to 70 Hz frequency region.

The pre-filter ρ_{static} and the optimal $\rho = -G_{12}/G_{11}$ are shown in figure 8 and the motivation for the design is twofold. First, the pre-filter is designed such that the gain rolls-off below 1 Hz and also has a negative slope for high frequencies to minimize the amplification of d_2. Second, ideally as few parameters n_θ as possible are used to estimate the subtraction filter, since this simplifies the optimization and minimizes the risk of overfitting or excessive filter behavior. Since the optimal subtraction filters all have a -2 slope above 10 Hz and only the anti-resonance frequency changes, the pre-filter is designed such that it has the same slope between 10 and 70 Hz.

In figure 9, the weighting filters W_e and W_f used in the ILC algorithm are shown. The weighting filter W_e is designed as a bandpass filter between 20 and 70 Hz to emphasize a good fit of the subtraction filter in this region. The weighting filter W_f is chosen as an inverse bandpass filter with a high gain outside the 8 to 70 Hz frequency region to penalize high gains of the subtraction filter in this region.

Finally, to set the relative importance of the three terms in the cost function, the gains of the matrices are chosen as $W_e = 10^3$, $W_f = 10^{-4}$ and $W_{\Delta f} = 10^{-1}$. These parameters are emperically determined, where the gain of W_e is left as a default value, the gain of W_f is chosen such that the bandpass structure of the filter is sufficiently enforced by iterating on the simulation and observing whether ρ

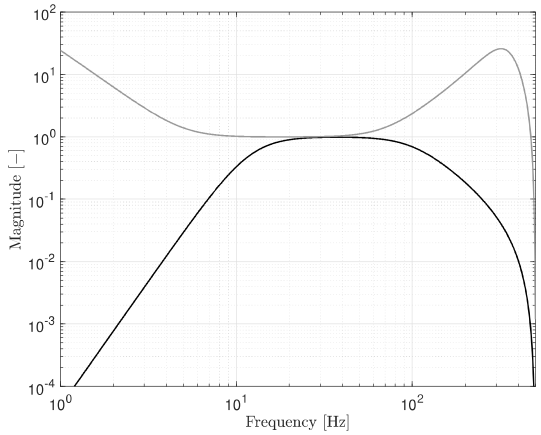

Figure 9. Weighting filter W_e (—) is designed as a bandpass filter to emphasize the frequency region where the subtraction filter should achieve high attenuation. The weighting filter W_f (—) penalizes a high gain of the subtraction filter by having an inverse bandpass structure, i.e. high values outside the frequency region where the subtraction filter should be effective. This ensures the subtraction filter does not amplify other noises outside the frequency region where it should be effective.

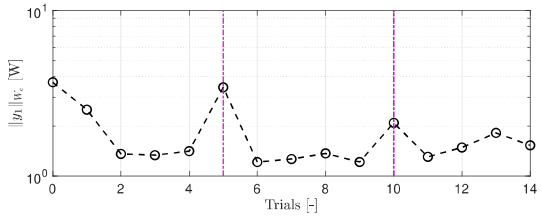

Figure 10. The filtered 2-norm of y_1 (—) for 15 trials of the ILC algorithm. At trials $j = 5$ and $j = 10$ (—), the model of G_{12} is changed, which yields a spike in the error since the feedforward signal was derived for the previous model of G_{12}. The algorithm converges within roughly two trials to the new optimal filter structure.

has the desired bandpass structure. The gain of $W_{\Delta f}$ is chosen such that the convergence speed is not too slow but also not within a single step such that a better performance is achieved. The term $W_{\Delta f}$ furthermore can also be used to improve the robustness of the algorithm to glitches in the data, which occasionally occur in GW detectors (Dooney *et al* 2025). These glitches may directly impact the estimation of the subtraction filter, leading to potentially incorrect parameter estimates. By increasing the gain of $W_{\Delta f}$, a glitch in one of the trials has less impact on the estimation of the subtraction filter since large changes in the parameter estimates between trials are penalized.

Note that for the pre-filter and weighting functions, small differences in the pole and zero frequencies are present. These changes stem from an empirical tuning process, where these frequencies are slightly adjusted to improve the performance of the algorithm.

5.3. Performance analysis

For the simulation, 15 trials are performed with θ_0 a vector of zeroes for the first trial. At trial $j = 5$ and trial $j = 10$, the model of G_{12} is changed to respectively the 30Hz and 40Hz anti-resonance models to simulate changes in the plant model during the iterations. Each trial uses 4.5 s of data sampled at 1 kHz and $n_\theta = 10$ FIR parameters are estimated for the subtraction filter ρ_i.

In figure 10, the filtered 2-norm of y_1 is shown for all trials, where the filtered 2-norm is essentially the RMS of the output signal y_1 weighted by the weighting filter W_e, to measure if η_2 is suppressed in y_1. Note here that this is different from the true cost function $\mathcal{J}(\theta^{j+1})$ in (25), since there the contributions of f_j and the rate of change of f_j are also included.

The filtered 2-norm of y_1 converges within roughly two iterations of the algorithm. At $j = 5$ and $j = 10$ where the model of G_{12} is changed, the filtered 2-norm of y_1 spikes due to the model of G_{12} changing, but quickly converges again. The transfer functions are in practice not expected to change instantaneously, so these large spikes are not representative for the expected behavior when this method is applied in practice. It is nevertheless beneficial to have the algorithm converge quickly to the new optimal filter in case the timescales of the changes of the plant are faster than the length of a dataset. Note furthermore that the filtered 2-norm of y_1 fluctuates around the optimal value, which is likely for

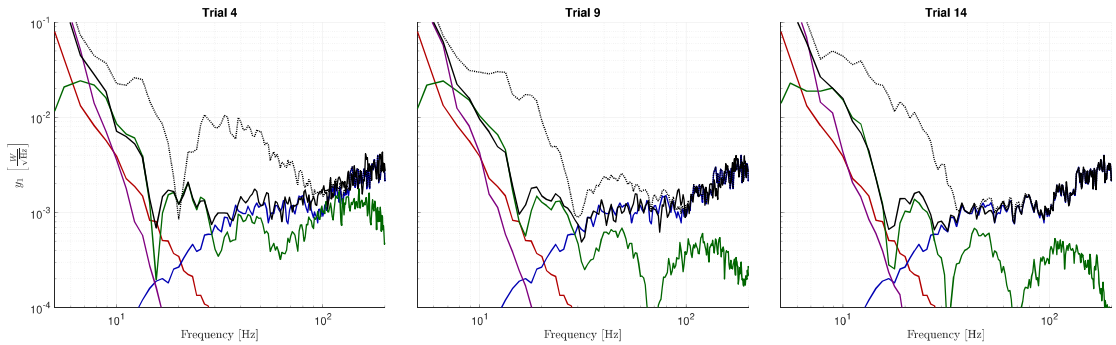

Figure 11. Noise budget of y_1 at trial $j=4$, $j=9$ and $j=14$, with y_1 using the subtraction filter (▬), y_1 without the subtraction filter (⋯), the contribution of u_2 in y_1 (▬), the contribution of η_1 in y_1 (▬), the contribution of d_1 in y_1 (▬). The contribution of u_2 is reduced by up to a factor 8 in the 10 to 70 Hz frequency region for each trial.

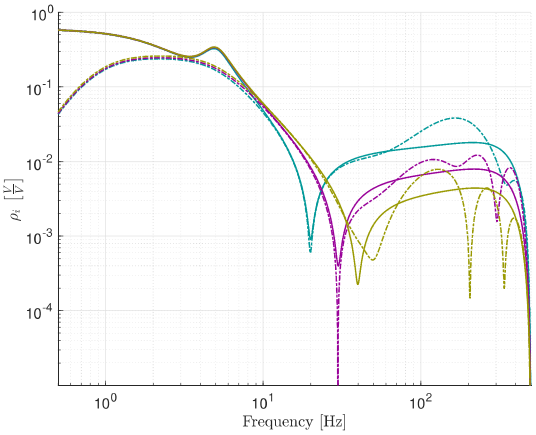

Figure 12. The optimal subtraction filter $-\frac{G_{12}}{G_{12}}$ (▬) and the estimated subtraction filter ρ_i (▬) at trial $j=4$, The optimal subtraction filter $-\frac{G_{12}}{G_{12}}$ (▬) and the estimated subtraction filter ρ_i (▬) at trial $j=9$ and the optimal subtraction filter $-\frac{G_{12}}{G_{12}}$ (▬) and the estimated subtraction filter ρ_i (▬) at trial $j=14$. The subtraction filter fits the optimal filter reasonably well inside the 10 to 70 Hz frequency region and the gain decreases outside this frequency region.

the reasons explained in section 4.5, next to the fact that the realization of the noise in u_2 changes, thus also marginally changing the RMS error in y_1.

In figure 11, the noise budget is shown for trial $j=4$, $j=9$ and $j=14$, i.e. the last trials for each model G_{12}. The estimated subtraction filter suppresses the contribution of η_2 in y_1 by up to a factor 8 in the 10 to 70 Hz frequency region after the algorithm has converged compared to the situation where no subtraction filter is used. Ideally the subtraction filter suppresses the contribution of η_2 such that it is not limiting anymore in y_1, which is the case below 20 Hz where d_2 and η_2 have roughly the same contribution in y_1 and above 30 Hz where η_1 is limiting. Due to the penalty on the high gain above 70 Hz, the contribution of η_2 is furthermore not limiting y_1 above 70 Hz.

In figure 12, the estimated subtraction filter ρ_i is shown for trials $j=4$, $j=9$ and $j=14$. The estimated subtraction filter fits the optimal subtraction filter $-\frac{G_{12}}{G_{11}}$ reasonably well inside the 10 to 70 Hz frequency region. For trial $j=4$, even the fit of the anti-resonance is quite accurate, but this worsens for trial $j=9$ and $j=14$, which is due to the poor signal-to-noise ratio at this frequency combined with the fact that for trial $j=14$, the anti-resonance is closer to the 70 Hz cut-off of W_e and W_f, i.e. less emphasis is placed on a good fit in this region combined with a penalty on a high gain around 100 Hz. The subtraction filter furthermore rolls off above 200 Hz as desired due to the weighting filter W_f and also rolls-off below 1 Hz due to the high-pass behavior in the pre-filter ρ_{static}. Note that in practice the amplification around 200 Hz will not pose any problems, as the feedback loop itself has a large roll-off in this frequency range. It is nevertheless a good illustration that this behavior can be enforced at any frequency in case it does pose practical problems. Around 150 Hz, a peak in the gain can be observed which also yields a poor fit and thus amplifies the noise in this frequency region. This is inherent to the causal feedforward approach, which results in the *waterbed effect* akin to feedback control (Beijen *et al* 2015), where the subtraction is improved at one frequency range at the expense of another frequency range.

The ILC algorithm thus successfully manages to suppress the contribution of η_2 in y_1 in the desired frequency range without worsening it outside this frequency range such that it limits y_1. The algorithm does however seem limited to suppressing η_2 enough that it is at the same level as the disturbances d_1 and η_1, but it is not able to get it below the level of these disturbances by ideally an order of magnitude. This is likely due to the fact that 4.5 s of data is used per trial, which is not enough to obtain sufficient accuracy of the subtraction filter. Since the algorithm works with time domain signals and requires inverting matrices in (27), large data stretches become computationally expensive, which limits the size of the dataset and therefore likely also the performance of the algorithm.

6. Conclusions and future work

A framework for the design and estimation of the subtraction filter is presented in this work. First, a guideline is presented to determine the optimal uni-directional parametrization of the subtraction filter, which is based on several criteria. Second, an adaptive algorithm is presented for estimating the selected subtraction filter without the use of external perturbations. The algorithm shows promising results, achieving up to a factor 8 suppression of noise between 10 and 80 Hz from an auxiliary DoF to the output signal of the detector in simulation. The approach may therefore be able to replace the existing approach, which would reduce the commissioning time and the downtime of the detector.

For future work, several directions for future research are being explored. First, the parametrization of the subtraction filter is currently restricted to the uni-directional case, whereby only a filter from the auxiliary DoF to the main DoF is used. Using bi-directional subtraction filters have been shown to yield better performance and could therefore be further explored in this context. Additionally, the adaptive algorithm for the estimation of the subtraction filter could be further improved by exploring more flexible parametrizations as well as optimizing the algorithm to be able to use longer datasets, resulting in a more accurate estimate of the optimal subtraction filter.

Data availability statement

All data that support the findings of this study are included within the article (and any supplementary files).

Acknowledgments

This work has been funded by the Netherlands Organisation for Scientific Research (NWO) under Grant Number 680.92.18.02. The authors gratefully acknowledge the Italian Istituto Nazionale di Fisica Nucleare (INFN), the French Centre National de la Recherche Scientifique (CNRS) and the Netherlands Organization for Scientific Research, for the construction and operation of the Virgo detector and the creation and support of the EGO consortium. The authors also gratefully acknowledge research support from these agencies as well as by the Spanish Agencia Estatal de Investigación, the Consellera d'Innovació, Universitats, Ciència i Societat Digital de la Generalitat Valenciana and the CERCA Programme Generalitat de Catalunya, Spain, the National Science Centre of Poland and the Foundation for Polish Science (FNP), the European Commission, the Hungarian Scientific Research Fund (OTKA), the French Lyon Institute of Origins (LIO), the Belgian Fonds de la Recherche Scientifique (FRS-FNRS), Actions de Recherche Concertées (ARC) and Fonds Wetenschappelijk Onderzoek—Vlaanderen (FWO), Belgium.

Author contributions

Mathyn van Dael 0000-0002-6061-8131
Conceptualization (equal), Data curation (lead), Formal analysis (lead), Investigation (supporting), Methodology (equal), Software (supporting), Supervision (lead), Validation (equal), Writing – original draft (lead)

Marjolein Daanen
Conceptualization (equal), Data curation (supporting), Formal analysis (lead), Investigation (lead), Methodology (equal), Software (lead), Validation (lead)

Koen Tiels 0000-0001-9279-110X
Conceptualization (supporting), Formal analysis (supporting), Investigation (supporting), Methodology (supporting), Software (supporting), Supervision (supporting), Writing – review & editing (equal)

Gert Witvoet 0000-0003-3560-5942
Conceptualization (supporting), Investigation (supporting), Supervision (supporting), Writing – review & editing (equal)

Bas Swinkels 0000-0002-3066-3601
Conceptualization (equal), Data curation (supporting), Formal analysis (supporting), Supervision (supporting), Validation (supporting), Writing – review & editing (equal)

Diego Bersanetti 0000-0002-7377-415X
Conceptualization (supporting), Data curation (supporting), Investigation (supporting), Resources (equal), Writing – review & editing (equal)

Julia Casanueva 0000-0002-2948-5238
Conceptualization (supporting), Data curation (supporting), Resources (equal)

Manuel Pinto
Data curation (supporting), Resources (equal)

Maddalena Mantovani 0000-0002-4424-5726
Data curation (supporting), Resources (equal)

Piernicola Spinicelli 0000-0001-8078-6047
Data curation (supporting), Resources (equal)

Camilla de Rossi 0000-0002-5825-472X
Data curation (supporting), Resources (equal)

Mattia Boldrini 0000-0001-9861-821X
Data curation (supporting), Resources (equal)

Paolo Ruggi 0009-0000-9362-053X
Data curation (supporting), Resources (equal)

Tom Oomen 0000-0001-7721-4566
Conceptualization (supporting), Formal analysis (supporting), Methodology (supporting), Supervision (supporting), Writing – review & editing (equal)

References

Aasi J, Abbott B P, Abbott R, Abbott T, Abernathy M R and Ackleyv K (The LIGO Scientific Collaboration) 2015 Advanced LIGO *Class. Quantum Grav.* **32** 74001

Acernese F *et al* 2023 Advanced Virgo Plus: Future Perspectives *J. Phys.: Conf. Ser.* **2429** 12040

Beijen M A, Heertjes M F, Butler H and Steinbuch M 2015 Performance trade-offs in disturbance feedforward compensation of active hard-mounted vibration isolators *2015 American Control Conf. (ACC)* pp 2149–54

Bersanetti D, Patricelli B, Piccinni O J, Piergiovanni F, Salemi F and Sequino V 2021 Advanced virgo: status of the detector, latest results and future prospects *Universe* **7** 322

Black E D 2001 An introduction to Pound–Drever–Hall laser frequency stabilization *Am. J. Phys.* **69** 79–87

Blanken L, Boeren F, Bruijnen D and Oomen T 2017 Batch-to-batch rational feedforward control: from iterative learning to identification approaches, with application to a wafer stage *IEEE/ASME Trans. Mechatronics* **22** 826–37

Bolder J and Oomen T 2015 Rational basis functions in iterative learning control—with experimental verification on a motion system *IEEE Trans. Control Syst. Technol.* **23** 722–9

Bristow D A, Tharayil M and Alleyne A G 2006 A survey of iterative learning control *IEEE Control Syst. Mag.* **26** 96–114

Capote E and Jia W *et al* 2025 Advanced LIGO detector performance in the fourth observing run *Phys. Rev. D* **111** 62002

Coughlin M, Mukund N, Harms J, Driggers J, Adhikari R and Mitra S 2016 Towards a first design of a Newtonian-noise cancellation system for Advanced LIGO *Class. Quantum Grav.* **33** 244001

Dooney T, Narola H, Bromuri S, Curier R L, Van Den Broeck C, Caudill S and Tan D S 2025 Time-domain reconstruction of signals and glitches in gravitational wave data with deep learning *Phys. Rev. D* **112** 44022

Doyle J, Glover K, Khargonekar P and Francis B 1989 State space solution to standard H2 and H∞ control problem *IEEE Trans. Autom. Control* **34** 831–47

Driggers J C, Evans M, Pepper K and Adhikari R 2012 Active noise cancellation in a suspended interferometer *Rev. Sci. Instrum.* **83** 24501

Evers E, van de Wal M and Oomen T 2019 Beyond decentralized wafer/reticle stage control design: A double-Youla approach for enhancing synchronized motion *Control Eng. Pract.* **83** 21–32

Giaime J A, Daw E J, Weitz M, Adhikari R, Fritschel P, Abbott R, Bork R and Heefner J 2003 Feedforward reduction of the microseism disturbance in a long-base-line interferometric gravitational-wave detector *Rev. Sci. Instrum.* **74** 218–24
Gunnarsson S and Norrlöf M 2001 On the design of ILC algorithms using optimization *Automatica* **37** 2011–6
Hakvoort W B J and Beijen M A 2023 Filtered-error RLS for self-tuning disturbance feedforward control with application to a multi-axis vibration isolator *Mechatronics* **89** 102934
Hjalmarsson H, Gevers M, Gunnarsson S and Lequin O 1998 Iterative feedback tuning: theory and applications *IEEE Control Syst. Mag.* **18** 26–41
Ickenroth T, van Haren M, Kon J, van Meer M, van Hulst J and Oomen T 2025 Automatic basis function selection in iterative learning control: a sparsity-promoting approach applied to an industrial printer *2025 American Control Conf. (ACC)* pp 2931–6
Lee J H, Lee K S and Kim W C 2000 Model-based iterative learning control with a quadratic criterion for time-varying linear systems *Automatica* **36** 641–57
MacFarlane A G J 1970 Commutative controller: a new technique for the design of multivariable control systems *Electronic lett.* **6** 121–3
Meadors G D, Kawabe K and Riles K 2014 Increasing LIGO sensitivity by feedforward subtraction of auxiliary length control noise *Class. Quantum Grav.* **31** 105014
Mishra S, Yeh W and Tomizuka M 2008 Iterative learning control design for synchronization of wafer and reticle stages *2008 American Control Conf.* pp 3908–13
Navarrete M O, Heertjes M F and Schmidt R H M 2015 Common zeros in synchronization of high-precision stage systems *2015 IEEE Int. Conf. on Mechatronics (ICM)* pp 602–7
Owens D H 1978 *Feedback and Multivariable Systems* (Peter Peregrinus Ltd.)
Phan M Q and Frueh J A 1996 Learning control for trajectory tracking using basis functions *Proc. 35th IEEE Conf. on Decision and Control* vol 3 pp 2490–2
Poot M, Portegies J, Kostić D and Oomen T 2023 Rational basis functions in iterative learning control for multivariable systems *2023 62nd IEEE Conf. on Decision and Control (CDC)* pp 4644–9
Skogestad S and Postlethwaite I 2005 *Multivariable Feedback Control: Analysis and Design* (Wiley, West Sussex)
Sternad M 1990 The use of disturbance measurement feedforward in LQG self-tuners *IFAC Proc.* **23** 31–36
Stoev J, Ertveldt J, Oomen T and Schoukens J 2017 Tensor methods for MIMO decoupling and control design using frequency response functions *Mechatronics* **45** 71–81
Swinkels B, Campagna E, Vajente G, Barsotti L and Evans M 2008 Longitudinal noise subtraction *Technical Report* VIR-050A-08 (Virgo Collaboration)
Tapia E N, Valentini M, Bersanetti D, Was M, Maggiore R, van Dael M, Perry J W, Swinkels B L, Mantovani M and Freise A 2023 A MIMO approach for longitudinal sensing and control noise projections of Advanced Virgo gravitational wave detector *Class. Quantum Grav.* **40** 185008
Tiwari V *et al* 2015 Regression of environmental noise in LIGO data *Class. Quantum Grav.* **32** 165014
Vaes D, Swevers J and Sas P 2004 Optimal decoupling for MIMO-controller design with robust performance *Proc. 2004 American Control Conf.* vol 5 pp 4601–6
Vajente G 2008 Analysis of sensitivity and noise sources for the Virgo gravitational wave interferometer *PhD Thesis* (Scuola Normale Superiore di Pisa)
Vajente G, Huang Y, Isi M, Driggers J C, Kissel J S, Szczepanficzyk M J and Vitale S 2020 Machine-learning nonstationary noise out of gravitational-wave detectors *Phys. Rev. D* **101** 42003
van Dael M *et al* 2024 Online decoupling of the time-varying longitudinal feedback loops for improved performance in Advanced Virgo Plus* *Class. Quantum Grav.* **41** 215008
van Dael M *et al* 2025 Control of the laser frequency in the Virgo interferometer: Dynamic noise budgeting for controller optimization *Astropart. Phys.* **164** 103028
van de Wijdeven J and Bosgra O H 2010 Using basis functions in iterative learning control: analysis and design theory *Int. J. Control* **83** 661–75
White M T and Tomizuka M 1997 Increased disturbance rejection in magnetic disk drives by acceleration feedforward control and parameter adaptation *Control Eng. Pract.* **5** 741–51