Enabling quantum optomechanics for phononic crystals and superfluid helium films

Wei, X.

Chapter 3

Photon Filtering System

In order to perform quantum optomechanical experiments, the quantum excitation of the phononic crystal membrane must be controlled and measured [15, 56]. Such mechanical quantum excitations, namely phonons, are very difficult to measure directly due to their low energy. Therefore, the phonon information must be transferred to the optical field, which provides a much more sensitive readout channel via single-photon detection. This is the same underlying technique used for optical cooling and heating, but here it is operated at the single-photon level. The phonon-to-photon transduction relies on the same cavity-enhanced sideband-scattering mechanism that is used for optical cooling and heating, but here we detect the interaction event by event, at the level of individual scattered photons [57].

The most challenging part of this experiment is to separate the single photons at the cavity resonance frequency from the large number of pump photons required to achieve an efficient interaction. This separation is achieved using a filter-cavity system consisting of four coupled Fabry-Perot cavities. The four cavities are frequency-locked sequentially, with each cavity locked to its neighbor. The first filter cavity is initially locked to the cryogenic science cavity. The output of the filter-cavity system is then directed to a high-efficiency superconducting nanowire single-photon detectors (see Chapter 2).

3.1 Filter cavity design

Separating the few photons that are up or down converted via optomechanical interactions with a membrane from the strong pump field requires high-performance spectral filtering because the relevant frequency separation is only of the order of 1 MHz. In this near-degenerate regime, even a small residual pump leakage can dominate the detected signal. To sufficiently filter out the pump laser, a suppression of approximately 140 dB at a 1 MHz frequency offset is required.

Fabry–Pérot filter cavities meet these requirements by acting as tunable narrowband resonators. By setting the cavity linewidth and tuning the resonance frequency via the cavity length, the transmission peak can be aligned to the scattered photons while the pump field is strongly rejected at a 1 MHz detuning. To quantify the achievable suppression, we derive the transmission for one cavity and then generalize the result to two and n cavities.

Initial theoretical calculations for the filter cavity design have been performed by Bas ten Haaf. In this chapter, the transmission properties are re-derived using the more general transfer-matrix model of coupled cavities.

3.1.1 Single-cavity model including mirror-coating loss

A Fabry–Pérot cavity provides frequency selectivity because only optical fields that satisfy the resonance condition build up coherently inside the resonator. In the frequency domain, the transmission exhibits a set of narrow resonances separated by the free spectral range

$$\mathrm{FSR} = \frac{c}{2L}, \tag{3.1}$$

where L is the cavity length. Fields detuned from resonance are therefore suppressed, which forms the basis for spectral filtering.

The width of each resonance sets the frequency resolution of the filter. We denote by κ the full width at half maximum of the intensity transmission peak (in units of Hz), which is related to the finesse $\mathcal{F}$ by

$$\kappa = \frac{\mathrm{FSR}}{\mathcal{F}}. \tag{3.2}$$

Since the required suppression is specified at a detuning of order 1 MHz, the achievable filtering performance is directly constrained by the cavity linewidth and by any loss that limits the finesse.

We now derive the transmitted field for a single Fabry–Pérot cavity while including mirror-coating loss only. The cavity consists of two mirrors with amplitude reflection and transmission coefficients (r_1, t_1) and (r_2, t_2). Coating-related loss is incorporated phenomenologically at the field level by introducing amplitude attenuation factors $(1-$

ℓ_1) and $(1-\ell_2)$, applied once per interaction with mirror 1 and mirror 2, respectively. With this convention, summing all multiple round trips yields

$$\begin{aligned} E_{\text{trans}} &= E_0\, t_1 t_2\, (1-\ell_1)(1-\ell_2)\, e^{-i\psi/2} \Big[1 + r_1 r_2 (1-\ell_1)(1-\ell_2) e^{i\psi} \\ &\qquad\qquad + \big(r_1 r_2 (1-\ell_1)(1-\ell_2) e^{i\psi} \big)^2 + \cdots \Big] \\ &= E_0\, t_1 t_2\, (1-\ell_1)(1-\ell_2)\, e^{-i\psi/2} \frac{1}{1 - r_1 r_2 (1-\ell_1)(1-\ell_2) e^{i\psi}}. \end{aligned} \tag{3.3}$$

Here E_0 and E_{trans} are the incident and transmitted field amplitudes, and

$$\psi = 2kL = \frac{2\omega L}{c}, \tag{3.4}$$

is the round-trip phase for optical angular frequency ω and wave number $k = \omega/c$. The global phase factor $e^{-i\psi/2}$ does not affect the transmitted intensity and is retained only for completeness. The parameters $\ell_{1,2}$ represent effective losses associated with the mirror coatings, such as absorption and scattering, modeled here as a multiplicative field reduction per mirror interaction.

For the present application, the relevant detuning is only $\sim 1\,\text{MHz}$. Achieving $\sim 140\,\text{dB}$ pump suppression at such a small offset requires an ultranarrow cavity linewidth, which, for a single Fabry–Pérot cavity of fixed length, translates into an extremely high finesse. In a lossless model this requirement can be met only if the mirror transmission is pushed to an extraordinarily small level, at the $\sim 10^{-3}\,\text{ppm}$ scale for the $L = 40\,\text{cm}$ cavity considered here. Figure 3.1(a) illustrates this idealized scenario, where the resonance is sufficiently sharp to provide the desired off-resonant suppression.

However, a single-cavity solution becomes impractical once realistic mirror-coating loss is taken into account. Even a typical coating loss of order 1 ppm per mirror interaction severely limits the intracavity buildup and simultaneously reduces the resonant throughput. As shown in Figure 3.1(b), the transmission peak is strongly diminished in this regime, implying that most photons at the desired frequency would be lost. This is incompatible with our goal, which requires not only high suppression of the pump but also appreciable transmission of the scattered photons.

These considerations lead to a central conclusion. A single Fabry–Pérot cavity cannot simultaneously provide the required $\sim$140 dB suppression at 1 MHz detuning and maintain acceptable on-resonance transmission under realistic coating-loss levels. We

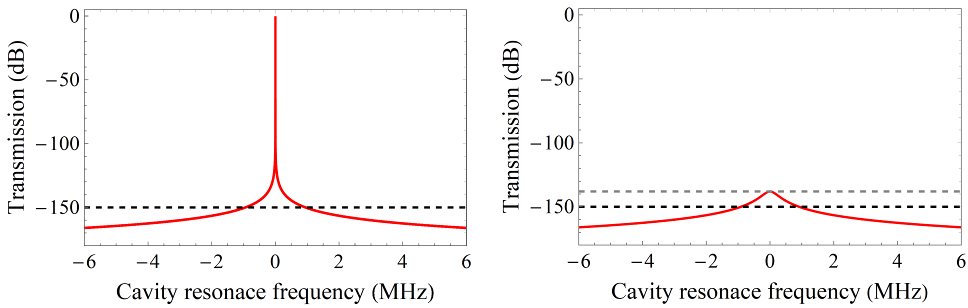

(a) Ideal mirrors without coating loss **(b)** Single-mirror coating loss of 1 ppm

Figure 3.1: Simulated transmission response of a single Fabry–Pérot cavity ($L = 40\,\mathrm{cm}$) with two identical mirrors. The lossless case illustrates that reaching strong suppression at a detuning of 1 MHz requires ultralow mirror transmission at the $\sim 10^{-3}$ ppm scale. Including a realistic coating loss strongly reduces the resonant transmission peak, indicating severe insertion loss even when high off-resonant suppression is obtained.

therefore turn to multi-cavity filter configurations, which allow the required extinction to be achieved while distributing the filtering burden across several resonators, thereby relaxing the mirror-transmission requirement for any individual cavity and mitigating the insertion-loss penalty.

3.1.2 Transfer-matrix model of coupled cavities

We now analyze the two-cavity Fabry–Pérot configuration. A convenient description, which generalizes directly to larger cavity numbers, is given by the transfer-matrix formalism. This formalism expresses the linear relation between forward- and backward-propagating field amplitudes across each optical element, and the response of the full structure follows from simple matrix multiplication. Additional effects, such as mirror loss or unequal spacings, are introduced by modifying the corresponding elementary matrices [58].

At a given reference plane z, we group the complex amplitudes of the two counter-propagating components as $\mathbf{E}(z) = (E_+(z), E_-(z))^\mathsf{T}$, corresponding to propagation along $\pm z$. We assume a monochromatic field at optical angular frequency $\omega = 2\pi\nu$. A polarization-fixed optical element is described by a 2×2 transfer matrix $\mathbf{M}$ relating the field vectors at two planes $z_1 < z_2$ along the optical axis, $\mathbf{E}(z_1) = \mathbf{M}\,\mathbf{E}(z_2)$. For a sequence of elements, the total transfer matrix is the ordered product of the individual matrices, taken in the same spatial order.

Propagation through free space, giving phases $\pm\phi$ to (E_+, E_-), is represented by

$$\mathbf{M}_{\mathrm{f}s}(\phi) = \begin{pmatrix} e^{i\phi} & 0 \\ 0 & e^{-i\phi} \end{pmatrix}. \tag{3.5}$$

For the length L and refractive index n (typically $n \simeq 1$), the single-pass phase is defined as $\phi(\nu) = (2\pi\nu/c)\, nL$, with c the speed of light in vacuum.

A mirror is specified by an amplitude reflectivity r and transmissivity t. For a lossless mirror, $r^2 + t^2 = 1$, so $t = \sqrt{1 - r^2}$. The corresponding interface matrix is

$$\mathbf{M}_{\mathrm{m}}(r) = \frac{1}{t} \begin{pmatrix} 1 & r \\ r & 1 \end{pmatrix}. \tag{3.6}$$

To include coating loss, we assign each mirror an intensity loss fraction l per interaction, accounting for absorption and scattering loss. In the reduced two-port description this is captured by relaxing the lossless relation to $r^2 + t^2 + l = 1$. For a given amplitude reflectivity r we therefore take $t = \sqrt{1 - r^2 - l}$ and keep the same interface-matrix form,

$$\mathbf{M}_{\mathrm{ml}}(r, l) = \frac{1}{t} \begin{pmatrix} 1 & r \\ r & 1 \end{pmatrix}, \qquad t = \sqrt{1 - r^2 - l}. \tag{3.7}$$

Consider four mirrors labeled $1, 2, 3, 4$ from left to right. Mirrors 1 and 4 are the outer mirrors with amplitude reflectivity r_{out}, and mirrors 2 and 3 are the inner mirrors with amplitude reflectivity r_{in}. The three gaps of lengths L_{12}, L_{23}, and L_{34} have single-pass phases ϕ_{12}, ϕ_{23}, and ϕ_{34}. For the symmetric configuration used in coupled filter cavities, we set $\phi_{12} = \phi_{34} \equiv \psi$ and $\phi_{23} \equiv \phi$. Here ψ is the intracavity phase for the two cavities (between mirrors 1–2 and 3–4), while ϕ is the inter-cavity spacing phase (between mirrors 2–3).

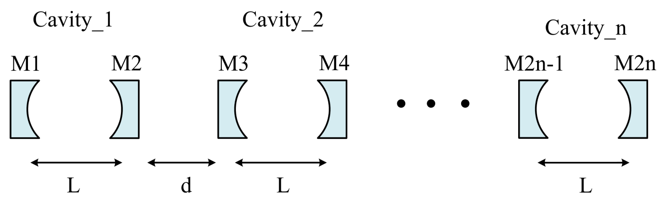

Figure 3.2: Schematic of a multi-cavity filter system formed by $2n$ mirrors. The gaps alternate between the cavity length L and the inter-cavity spacing d.

The total transfer matrix from the right of mirror 4 to the left of mirror 1 is then

$$\mathbf{M}_4(\phi,\psi) = \mathbf{M}_\mathrm{m}(r_\mathrm{out})\,\mathbf{M}_\mathrm{f}(\phi)\,\mathbf{M}_\mathrm{m}(r_\mathrm{in})\,\mathbf{M}_\mathrm{f}(\psi)\,\mathbf{M}_\mathrm{m}(r_\mathrm{in})\,\mathbf{M}_\mathrm{f}(\phi)\,\mathbf{M}_\mathrm{m}(r_\mathrm{out}). \tag{3.8}$$

We denote the matrix elements of $\mathbf{M}_4(\phi,\psi)$ as $\mathbf{M}_4(\phi,\psi) = \begin{pmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{pmatrix}$. For a unit-amplitude wave incident from the left and no incident laser from the right, the transmitted amplitude is $t_\mathrm{tot} = 1/M_{22}$, and the power transmission is $T(\nu) = |t_\mathrm{tot}|^2 = 1/|M_{22}|^2$.

We now generalize to $2n$ mirrors forming n coupled cavities. The mirrors are labeled $1, 2, \ldots, 2n$ from left to right, with mirrors 1 and $2n$ as the end mirrors and mirrors 2 to $2n-1$ as intermediate mirrors. Figure 3.2 defines the geometry and the two relevant gap lengths. The gaps alternate as

$$L_{(2k-1)(2k)} = L \quad (k = 1, \ldots, n), \qquad L_{(2k)(2k+1)} = d \quad (k = 1, \ldots, n-1).$$

Here L is the length of each cavity and d is the spacing between neighboring cavities.

For the two spacings L and d, we define the corresponding single-pass phases $\phi(\Delta\nu)$ and $\psi(\Delta\nu)$. For a small laser detuning $\Delta\nu$ around resonance, the phases vary linearly,

$$\phi(\Delta\nu) = \phi_0 + \frac{2\pi L}{c}\,\Delta\nu, \qquad \psi(\Delta\nu) = \psi_0 + \frac{2\pi d}{c}\,\Delta\nu,$$

where ϕ_0 and ψ_0 are the phases at $\Delta\nu = 0$.

The total transfer matrix for $2n$ mirrors is

$$\mathbf{M}_{2n}(\Delta\nu) = \mathbf{M}_\mathrm{m}(r_\mathrm{out}) \prod_{k=1}^{2n-1} \Big[\mathbf{M}_\mathrm{f}\big(\Phi_k(\Delta\nu)\big)\,\mathbf{M}_\mathrm{m}(r_{k+1})\Big], \tag{3.9}$$

where the product is evaluated from left to right as k increases, and $\Phi_k(\Delta\nu) = \phi(\Delta\nu)$ for $k = 1, 3, \ldots, 2n-1$ while $\Phi_k(\Delta\nu) = \psi(\Delta\nu)$ for $k = 2, 4, \ldots, 2n-2$. Here $r_1 = r_{2n} \equiv r_\mathrm{out}$ and $r_i \equiv r_\mathrm{in}$ for $i = 2, \ldots, 2n-1$. The corresponding power transmission is

$$T_{2n}(\Delta\nu) = \frac{1}{\big|[\mathbf{M}_{2n}(\Delta\nu)]_{22}\big|^2}. \tag{3.10}$$

Figure 3.3 shows the simulated power transmission of a series-coupled Fabry–Pérot filter with two, three, and four cavities, corresponding to 4, 6, and 8 mirrors in total. In the model, each mirror has a power transmissivity of 250 ppm and an additional

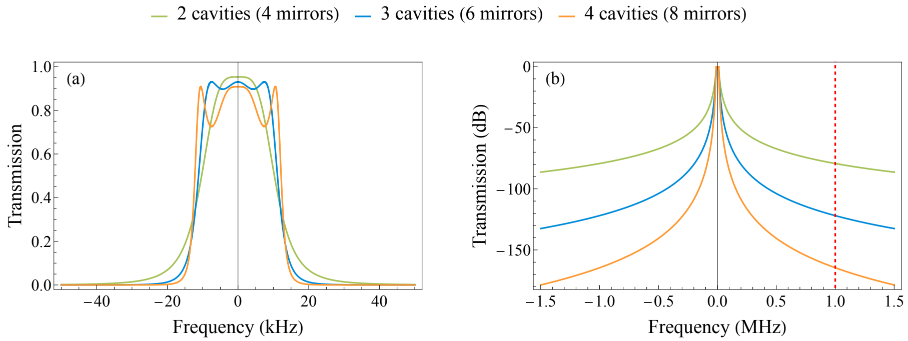

Figure 3.3: Simulated power transmission of a series-coupled Fabry–Pérot filter with two, three, and four cavities, corresponding to 4, 6, and 8 mirrors in total. Each mirror is modeled with a power transmissivity of 250 ppm and an intensity loss of 1 ppm. (a) Transmission lineshapes near resonance. The peak transmission decreases as more mirrors are added. (b) Transmission over a wider detuning range. The red dashed line marks a detuning of 1 MHz, where the suppression improves as more cavities are added and exceeds 160 dB for the four-cavity design.

3

mirror intensity loss of 1 ppm, which represents coating absorption and scattering loss. For an ideal lossless and impedance-matched cavity, the on-resonance transmission can approach unity, whereas any intracavity loss reduces the achievable peak transmission and introduces insertion loss.

Figure 3.3(a) shows the transmission response near resonance. The peak transmission remains close to 100% for all three configurations and decreases systematically as the number of mirrors increases. This reduction is expected because each additional mirror introduces extra absorption and scattering, which adds up to a larger total insertion loss for the system. This behavior sets the maximum throughput of the filter. In the four-cavity design, the peak transmission is about 91% in the idealized loss model, which is acceptable for our purposes.

Figure 3.3(b) shows the transmission over a wider detuning range, and the red dashed line marks a detuning of 1 MHz from the filter cavity system resonance. In the experiment, a single-frequency laser is injected into the membrane-in-the-middle science cavity. The output field is dominated by the pump at the laser frequency, while the interaction with the mechanical mode generates only a small fraction of photons shifted by the mechanical frequency near 1 MHz. These two frequency components propagate into the filter system. The filter cavities are tuned to be resonant with the scattered photons, so that the desired sideband is transmitted at the filter resonance while the pump is strongly attenuated. We therefore use the transmission at 1 MHz detuning to quantify the pump suppression.

As the number of cavities increases, the transmitted power at this offset drops, indicating stronger suppression for pump laser. For the four-cavity design, the transmission at a detuning of 1 MHz is below 10^{-16} relative to the peak transmission. Here the transmission is plotted in decibels as $10 \log_{10} T$. For the four-cavity design, the value at 1 MHz detuning lies more than 160 dB below the resonant peak, meeting our design requirement.

The assumed per-mirror loss of 1 ppm represents an idealized value. In the experiment, coating absorption and scattering can be higher, which would reduce the peak transmission and should be treated as an important tolerance in the experimental design.

3.1.3 Practical optical design with circulator

The transfer-matrix treatment in the previous section models the multi-cavity filter as a fully coupled optical system. In that idealized limit, fields reflected from downstream mirrors can propagate back toward upstream cavities and interfere with earlier round trips. The resulting spectrum represents a useful upper benchmark because it captures the strongest possible multi-pass interference, which maximizes out-of-band suppression for a given set of mirror parameters.

In the experiment, however, the filter must be operated under active frequency stabilization. Each cavity is held on resonance relative to a common reference set by the cryogenic science cavity. Under these conditions, uncontrolled optical feedback between cavities is more than a modeling detail. It can couple nominally independent control loops through parasitic back-reflections and fluctuations in mode matching, which modifies the reflection signals used for locking and reduces operational robustness. The experimental design therefore suppresses backward-propagating fields between neighboring cavities while preserving high forward transmission.

Figure 3.4 shows the coupled-cavity configuration discussed in the previous section, in which back-reflections are allowed to re-enter upstream cavities. In the apparatus we implement the feedback-blocked configuration in Figure 3.5. Between neighboring cavities we insert a polarizing beam splitter together with two quarter-wave plates, forming a polarization circulator that redirects retro-reflected light into a separate port rather than reinjecting it into the preceding cavity. This routing also provides convenient access to the reflection from each cavity for frequency stabilization. In addition, two lenses are placed between cavities to realize spatial mode matching,

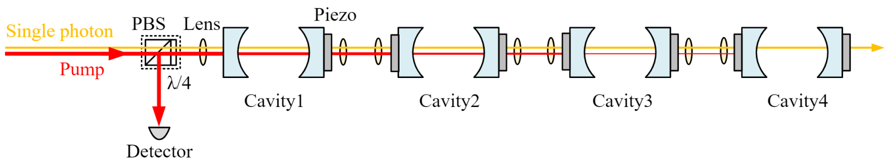

Figure 3.4: Simplified schematic of the four-cavity filter system operated without polarization circulation optics. Two incident beams are indicated at the input. Each cavity mirror is mounted on a piezoelectric actuator to provide fine tuning of the cavity resonance for frequency stabilization. Two lenses are placed between neighboring cavities to achieve spatial mode matching.

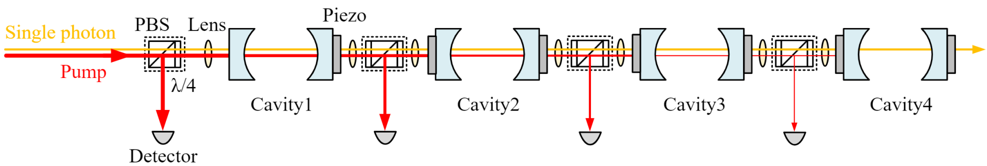

Figure 3.5: Simplified schematic of the four-cavity filter with polarization circulators inserted between cavities. The same two incident beams, piezo-actuated mirrors, and inter-cavity mode-matching lenses as in Figure 3.4 are shown. In addition, each inter-cavity section includes a PBS together with two quarter-wave plates, forming a polarization circulator that redirects back reflection light into a detector.

which is required to maintain high coupling into the fundamental cavity mode and to preserve usable on-resonance throughput.

The feedback-blocked configuration motivates a simplified transmission model. Each cavity is treated independently using the same transfer-matrix formalism, yielding a 2×2 matrix $\mathbf{M}_i(\Delta\nu)$ and a power transmission

$$T_i(\Delta\nu) = \frac{1}{\left|[\mathbf{M}_i(\Delta\nu)]_{22}\right|^2}. \tag{3.11}$$

The key difference from the coupled case is that no field reflected from cavity $i+1$ is allowed to re-enter cavity i. The total transmission therefore factorizes as a product of single-cavity transmissions,

$$T_{\mathrm{tot}}(\Delta\nu) = \eta \prod_{i=1}^{n} T_i(\Delta\nu), \tag{3.12}$$

where η is the total transmission from one cavity to the next, including circulator insertion loss and all other possible losses.

Figure 3.6 compares the simulated transmission of a two-cavity filter in the back-reflection-coupled case (green) and the back-reflection-blocked case (blue). Figure 3.6(a) plots the transmission within ±50 kHz of resonance. Figure 3.6(b) plots the transmission over ±1.5 MHz, with a vertical red dashed line at 1 MHz and horizontal dashed lines indicating the transmission values at this detuning for the two cases. Blocking back reflections narrows the resonance, but it degrades the suppression at a 1 MHz detuning. Allowing back reflections broadens the linewidth, while the suppression ratio improves.

In the back-reflection-blocked case the total transmission reduces to a product, $T_{\text{tot}}(\Delta\nu) \propto T_1(\Delta\nu)T_2(\Delta\nu)$, and for two identical cavities this becomes $T_{\text{tot}} \propto [T(\Delta\nu)]^2$. Multiplying the single-cavity transmissions makes the overall resonance narrower, even when the linewidth of each cavity is unchanged.

Allowing back reflections introduces feedback paths that couple the two cavities into a single multi-mirror interferometer. The transmission is reshaped by interference rather than by a simple squaring of the single-cavity response. The resulting lineshape is broader, with a flatter top near resonance. This flatter peak can be useful because it provides a wider frequency range with nearly constant transmission and reduces sensitivity to small frequency drifts.

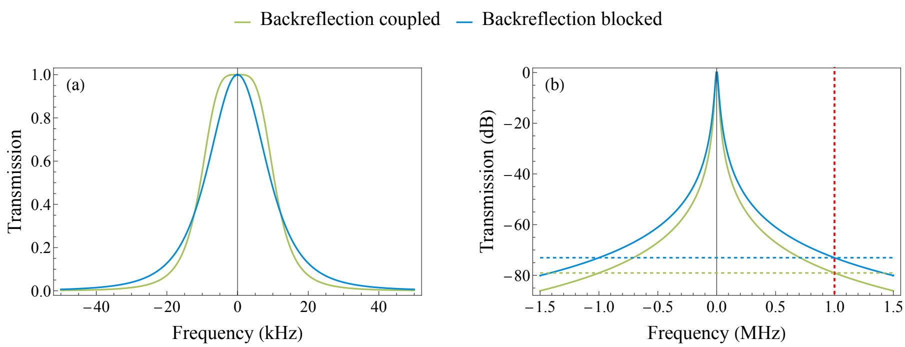

Figure 3.6: Simulated power transmission of a two-cavity filter, comparing the back-reflection-coupled case (green) with the back-reflection-blocked case (blue). Each mirror is modeled with an intensity transmissivity of 250 ppm. (a) Transmission within ±50 kHz of resonance. (b) Transmission over ±1.5 MHz. The vertical red dashed line marks the detuning of 1 MHz, and the horizontal dashed lines indicate the transmission values at this detuning for the two configurations.

To obtain strong pump rejection at a fixed detuning of 1 MHz while retaining a usable resonant passband, we implement a four-cavity cascade. Figure 3.7 compares the simulated power transmission of the four-cavity filter for the back-reflection-coupled

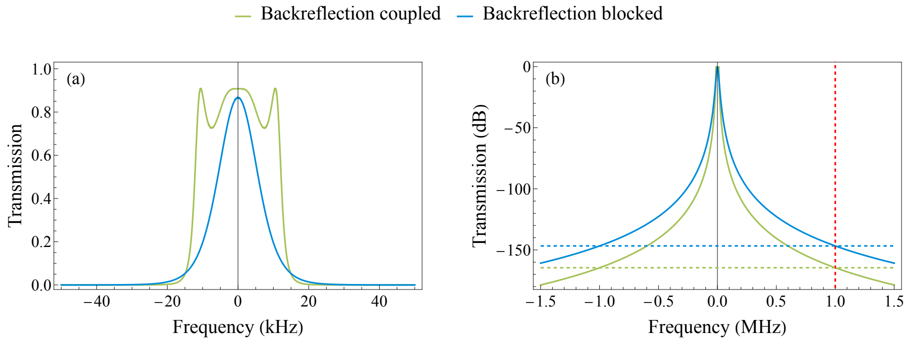

Figure 3.7: Simulated power transmission of a four-cavity Fabry–Pérot filter system, comparing the back-reflection-coupled configuration (green) with the back-reflection-blocked configuration (blue). Each mirror has an intensity transmissivity of 250, ppm and an intensity loss of 3, ppm per mirror. The insertion loss $1-\eta$ of each optical circulator is 0.015. (a) Transmission within ± 50 kHz of resonance. The coupled configuration reaches a peak on-resonance transmission of $\sim 91\%$, compared with $\sim 87\%$ for the blocked configuration. (b) Transmission over ± 1.5 MHz. The vertical red dashed line marks a detuning of 1 MHz. The suppression exceeds 160 dB in the coupled case and 140 dB in the blocked case.

and back-reflection-blocked configurations.

In the back-reflection-blocked configuration, reverse-propagating fields are prevented from re-entering previous cavities. For mirrors with intensity transmissivity 250 ppm and intensity loss 3 ppm per mirror, together with an insertion loss of 0.015 for each optical circulator, the peak on-resonance transmission is $\sim 87\%$. Here the PBS–$\lambda/4$ sets used in the four-cavity system have a measured total intensity transmission of $\sim 94\%$, corresponding to an intensity transmission of $\sim 98.5\%$ per circulator (i.e., an insertion loss of ~ 0.015). Allowing back-reflection increases the peak transmission to $\sim 91\%$. At a detuning of 1 MHz, the coupled configuration provides > 160 dB suppression, compared with > 140 dB for the blocked configuration.

Taken together, Figure 3.7 sets the expected operating condition of the four-cavity filter. With the mirror and circulator parameters, the cascade filter cavities provides a resonant transmission of 87% while achieving pump suppression beyond 10^{14} at a detuning of 1 MHz. The back-reflection-coupled configuration yields both a higher peak transmission and stronger suppression at the detuning of interest. In the experiment, we begin with the back-reflection-blocked implementation using optical circulator to simplify the initial alignment and locking.

3.2 Filter cavity alignment

Reproducing the response in Figure 3.7 requires stable alignment and efficient mode matching across the entire four-cavity system. We therefore start from the back-reflection-blocked implementation using optical circulators, so that each cavity can be aligned and mode matched independently before the full cascade is assembled.

The alignment was implemented as a reproducible procedure compatible with high vacuum operation and aimed at maximizing coupling to the fundamental transverse cavity mode. The cavity structure and mirror mounts were custom fabricated and assembled. To ensure vacuum compatibility, all components were cleaned prior to mirror installation using isopropanol and ultrasonic bath cleaning. Ultrasonic cleaning was typically performed for 1–2 h per batch. Small threaded parts such as screws were cleaned for longer because residues are more readily trapped in threads and recesses.

After mechanical assembly and before installing the mirrors, the cavity parts were placed in the vacuum chamber, evacuated, and gently baked to promote desorption and to verify vacuum integrity. Under these conditions the vacuum chamber reached pressures on the order of 10^{-5} mbar, which was sufficient to validate vacuum compatibility and to reduce the risk of time-dependent contamination that can increase scatter and degrade long-term stability.

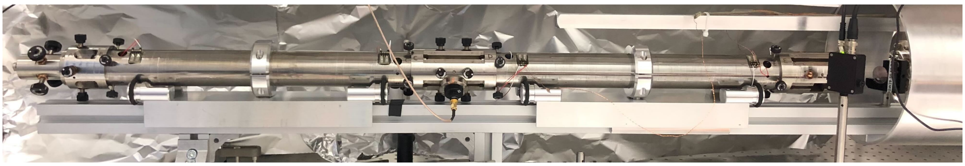

Figure 3.8: Photograph of the filter-cavity mount, showing the first two cavities and the inter-cavity spacer.

Optical alignment was carried out in two stages. A visible red beam was first used for coarse resonance acquisition by scanning the cavity length with the mirror piezos until transverse-mode resonances were observable in transmission. Any stable transmitted mode was taken as the acquisition criterion and a suitable starting point for fine alignment. The system was then switched to a narrow-linewidth 1064 nm laser for fine optimization, where mirror angles and mode-matching optics were adjusted to maximize coupling to the TEM_{00} mode.

The cavity is formed by two concave mirrors with radius of curvature $R = 250$ mm and length $L = 40$ cm. Spatial mode matching to the cavity eigenmode was achieved using two $f = 150$ mm lenses. Both lenses were mounted in adjustable holders that allowed transverse (x, y) alignment and longitudinal (z) positioning along the optical axis.

For coarse alignment, a visible red beam was injected into the cavity. Because the mirrors are designed for high reflectivity at 1064 nm with an intensity transmissivity of $T \simeq 250$ ppm. The transmission increases at visible wavelengths, providing a stronger transmitted signal for resonance acquisition during the initial alignment stage. The cavity length was scanned with the mirror-mounted piezoelectric actuators and the transmitted field was monitored on a camera placed after the output mirror. Coarse alignment was considered successful once repeatable resonant features and recognizable transverse-mode patterns were observed during cavity length scans, indicating that the injected beam was close to the cavity axis.

Fine alignment was then performed with a narrow-linewidth 1064 nm laser by scanning the laser frequency across cavity resonances while optimizing transverse centering of the cavity mirrors. The relevant degrees of freedom were the (x, y) positions of the optics, together with the axial spacing of the mode-matching lenses, which were adjusted to maximize the TEM_{00} coupling in transmission. Although mirror tilt was not adjustable in this design, the mechanical assembly fixed the cavity axis sufficiently well that the mismatch was dominated by transverse offsets, which can be corrected by centering the optics to suppress coupling to higher-order modes.

The transmitted power was recorded on a photodetector and used as the primary optimization signal. The coupling to the TEM_{00} mode was quantified from the mode-resolved transmission spectrum by comparing the fundamental-mode contribution to the summed contribution of higher-order modes within the same scan window,

$$\eta_{00} = \frac{P_{00}}{P_{00} + \sum_{m,n\neq 0} P_{mn}}. \tag{3.13}$$

Here P_{00} denotes the photodetector response associated with the TEM_{00} resonance and $\sum_{m,n\neq 0} P_{mn}$ is the summed response from all resolved higher-order transverse modes. After optimization, η_{00} exceeded 0.99. The reflected signal exhibited an on-resonance dip near the detector background, consistent with near-optimal mode matching and input coupling.

For multi-cavity assemblies, alignment was performed sequentially. The first cavity was optimized at 1064 nm and stabilized, and its transmitted beam was used to align the second cavity, followed by the remaining stages. This staged strategy reduces the number of simultaneously varying degrees of freedom and preserves the mode purity established upstream, which is essential for achieving high extinction at small

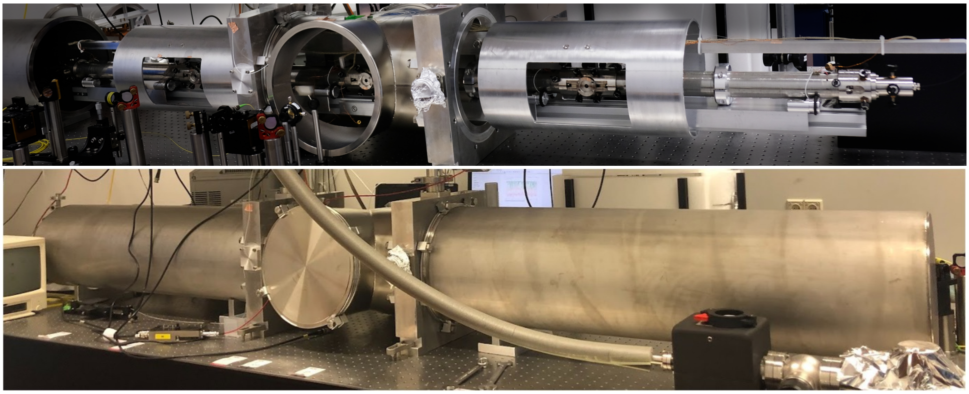

Figure 3.9: Photographs of the filter-cavity system. Top: the fully assembled four-cavity mount. Bottom: the system installed with the vacuum chamber sealed.

detunings. A photograph of the fully aligned four-cavity system is shown in Figure 3.9. The filter cavity system has been a major development project together with the Fine Mechanics Department and the Electronics Department. The system is completely made out of Invar and is a true piece of art, designed and machined by Harmen van der Meer. As a reference, the overall layout of the four-cavity filter system is shown in Figure 3.10.

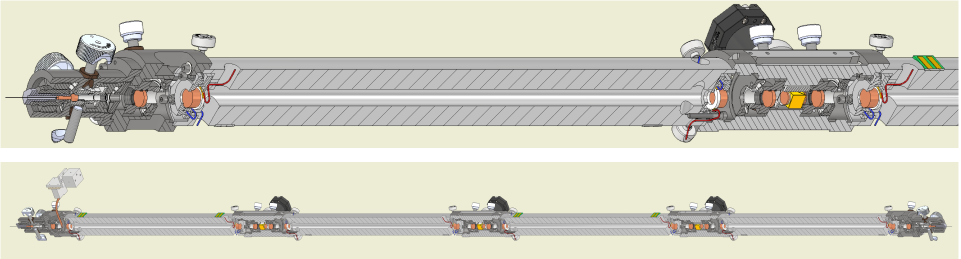

Figure 3.10: Schematic of the four-cavity filter system. Top: configuration of a single cavity. Bottom: configuration of the full four-cavity system, in which cavities 1 and 4 are defined by fiber couplers and the central section is implemented entirely in free space. The mechanical components were designed and machined by Harmen van der Meer.

With the spatial mode matching established, the next subsection describes the dither-locking scheme used to stabilize the cavities on resonance.

3.3 Dither locking

Several locking schemes can be used to stabilize a Fabry–Pérot cavity, including Pound–Drever–Hall (PDH) locking and dither locking [59, 60]. PDH generates an error signal by phase modulating the incident light and demodulating the detected response at the modulation frequency, which relies on the reflected optical sidebands. Dither locking modulates the cavity length at a small amplitude and demodulates the resulting modulation of the transmitted or reflected power to generate the error signal.

For the present multi-cavity filter, the main requirement is to maximize the transmitted optical power. PDH locking of each additional cavity requires phase-modulation sidebands on the incident field. With the first cavity locked, its transmitted light provides the input to the next cavity. Implementing PDH on each cavity would therefore require an electro-optic modulator (EOM) between cavities to apply phase modulation to the transmitted light, introducing sidebands and increasing the insertion loss of the filtering system. To avoid additional EOMs between cavities and the associated loss, we implement dither locking for stabilization of the cascaded cavity system.

The dither lock is implemented by applying a small sinusoidal length modulation to the cavity using the piezoelectric actuator on the mirror mount. The demodulation yields an error signal that changes sign across resonance and crosses zero at the transmission maximum, without introducing additional optical frequencies. The cavity length modulation can be written as,

$$L(t) = L_0 + \Delta L \cos(\omega_m t), \tag{3.14}$$

which modulates the cavity resonance and hence the detected signal $P(t)$ (transmission or reflection). For sufficiently small ΔL, we expand $P(L)$ to first order about the operating point L_0,

$$P(t) = P\big(L(t)\big) \simeq P(L_0) + \left.\frac{\mathrm{d}P}{\mathrm{d}L}\right|_{L_0} \Delta L \cos(\omega_m t), \tag{3.15}$$

neglecting higher-order terms in ΔL.

The error signal is obtained by mixing $P(t)$ with a reference at ω_m, followed by low-pass filtering,

$$\epsilon(t) \equiv \mathrm{LPF}\{P(t)\cos(\omega_m t + \phi)\}\,, \tag{3.16}$$

where ϕ is the reference phase and $\mathrm{LPF}\{\cdot\}$ denotes low-pass filter. Substituting

Eq. (3.15) into Eq. (3.16) and using $\cos a \cos b = \frac{1}{2}[\cos(a-b) + \cos(a+b)]$ gives

$$P(t)\cos(\omega_m t+\phi) \simeq P(L_0)\cos(\omega_m t+\phi) + \frac{1}{2}\left.\frac{\mathrm{d}P}{\mathrm{d}L}\right|_{L_0} \Delta L[\cos\phi + \cos(2\omega_m t+\phi)]\,. \quad (3.17)$$

The low-pass filter removes the terms oscillating at ω_m and $2\omega_m$, leaving the near-DC component

$$\epsilon \simeq \frac{1}{2}\left.\frac{\mathrm{d}P}{\mathrm{d}L}\right|_{L_0} \Delta L \cos\phi. \quad (3.18)$$

With the reference phase adjusted to $\phi \simeq 0$, the error signal is proportional to the local slope $\mathrm{d}P/\mathrm{d}L$ and therefore changes sign across resonance.

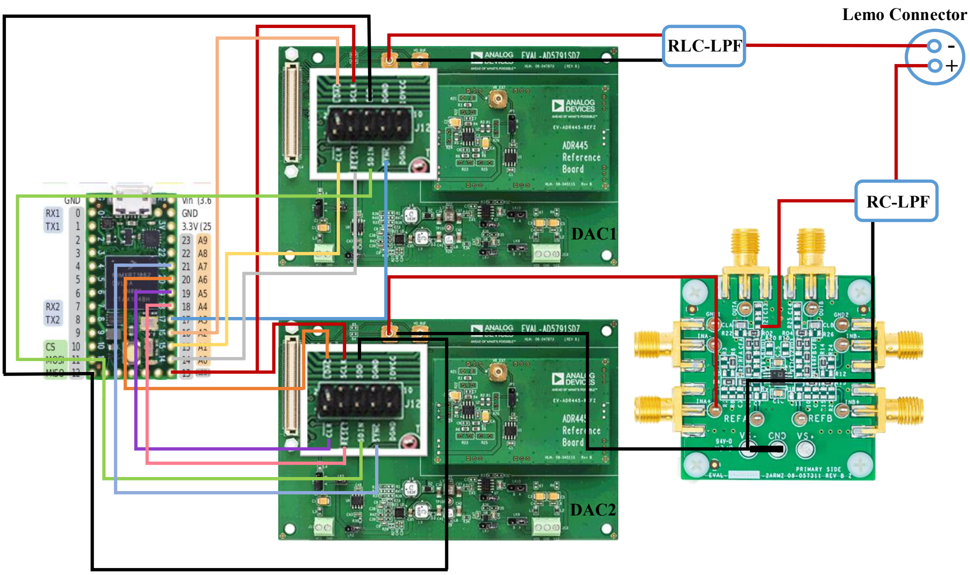

Figure 3.11: First-generation lockbox for cavity dither locking. A Teensy® 4.0 microcontroller generates the feedback signal, which is applied to the piezo via two evaluation boards (EVAL-AD5791ARDZ). DAC2 provides a DC bias, and DAC1 provides the real-time correction.

With a 40 cm length and 250 ppm mirrors, the Fabry–Pérot cavity has a finesse of $\mathcal{F} \approx 1.3 \times 10^4$ and a correspondingly narrow resonance linewidth of about 30 kHz (FWHM). This narrow linewidth means the lock must keep the cavity near resonance while adding negligible frequency noise.

The cavity length is controlled with a PZT that provides 3.3 μm of displacement over a 230 V drive range, corresponding to $\mathrm{d}L/\mathrm{d}V \approx 14.3$ nm/V. This conversion factor allows the cavity linewidth to be mapped onto a voltage range, which then sets the required PID output-noise level.

At $\lambda = 1064\,\mathrm{nm}$, one free spectral range corresponds to a length change of $\lambda/2 = 532\,\mathrm{nm}$, and the free spectral range is $375\,\mathrm{MHz}$. A linewidth of $30\,\mathrm{kHz}$ (FWHM) corresponds to an effective length change of about $43\,\mathrm{pm}$, which maps to a voltage range of about $3\,\mathrm{mV}$. For reference, achieving $\sim 1\%$ of this linewidth corresponds to an output-noise level below $\sim 30\,\mu\mathrm{V}$ over the feedback bandwidth. A custom low-noise PID servo was built to drive the cavity piezo. Commercial controllers were not well matched to this operating point, mainly because their output noise was too high.

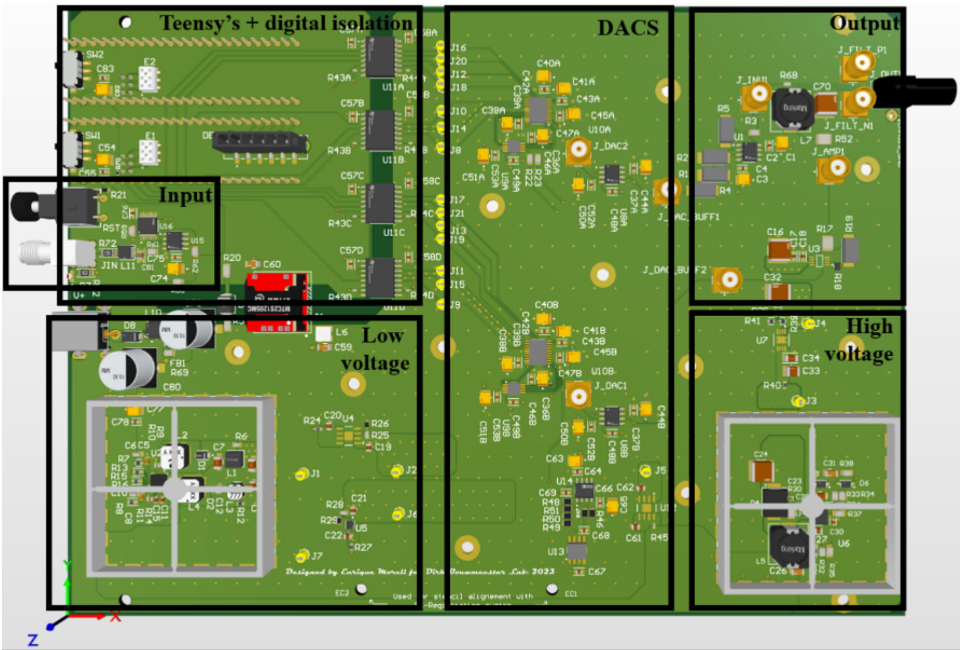

Figure 3.12: PCB layout of the second-generation lockbox, which integrates the functionality of the first-generation prototype into a single board. This board has been designed and constructed by my PhD colleague Enrique Morell.

Figure 3.11 shows the first-generation lockbox. The real-time controller is implemented on a Teensy® 4.0 microcontroller. Two high-resolution voltage-output DAC evaluation boards, EVAL-AD5791SDZ, generate the analog drive signals labeled DAC1 and DAC2 in Figure 3.11. In our setup the cavity is mounted in vacuum, and experimentally we find that the relevant disturbance spectrum is slow enough that a closed-loop bandwidth of a few hundred hertz is sufficient. We therefore use

DAC2 only to set the static operating point, while all dynamic corrections required for frequency stabilization are applied through DAC1.

The piezoelectric actuator is driven differentially, so that the effective drive voltage can be written as $V_{\mathrm{PZT}}(t) = V_{+}(t) - V_{-}(t)$. The positive terminal is driven by DAC2 and the negative terminal by DAC1. During normal operation, DAC2 is held fixed and provides only a DC bias that sets the working point, while DAC1 carries the full real-time correction computed from the demodulated error signal. With V_{+} held constant, small variations satisfy $\delta V_{\mathrm{PZT}} \simeq -\delta V_{-}$. The DC bias from DAC2 is amplified by a low-noise high-voltage amplifier (EVAL-ADA4522-2ARMZ), providing a voltage range on the order of 50 V. The on-board gain-setting resistors were replaced to set the amplifier gain.

To suppress electronic noise delivered to the piezo, both DAC outputs are filtered before entering the PZT inputs. In our first version this is implemented as a low-pass filter with a cutoff near 150 Hz. This filtering reduces broadband noise.

The demodulated error signal is processed by a digital PID controller running on the microcontroller, and the resulting correction is applied through DAC1. PID control is a standard approach for stabilizing slowly drifting systems while maintaining strong suppression of low-frequency disturbances. In our implementation, the dither modulation is applied at 1 kHz for synchronous demodulation, while the digital control loop updates at ~ 10 kHz.

In practical operation, lock acquisition proceeds by setting the DC bias on DAC2 to bring the cavity close to resonance and then enabling the dither demodulation and PID feedback on DAC1. Once locked, DAC1 continuously compensates the drift within the chosen bandwidth, while DAC2 remains fixed. This separation simplifies tuning and noise budgeting because only the feedback channel contributes appreciably within the lock bandwidth.

This first-generation design provided stable operation in the experiment. Based on this working prototype, we consolidated the same architecture into a dedicated PCB, whose layout is shown in Figure 3.12. Building on this PCB, the final assembled lockbox is shown in Figure 3.13.

3.4 Frequency stabilization of filter cavities

With the control scheme in place, we now show how the system behaves in practice. We begin with a single cavity, where the controller captures the reflection minimum and then turns on the PID feedback to stabilize the cavity at resonance, and then

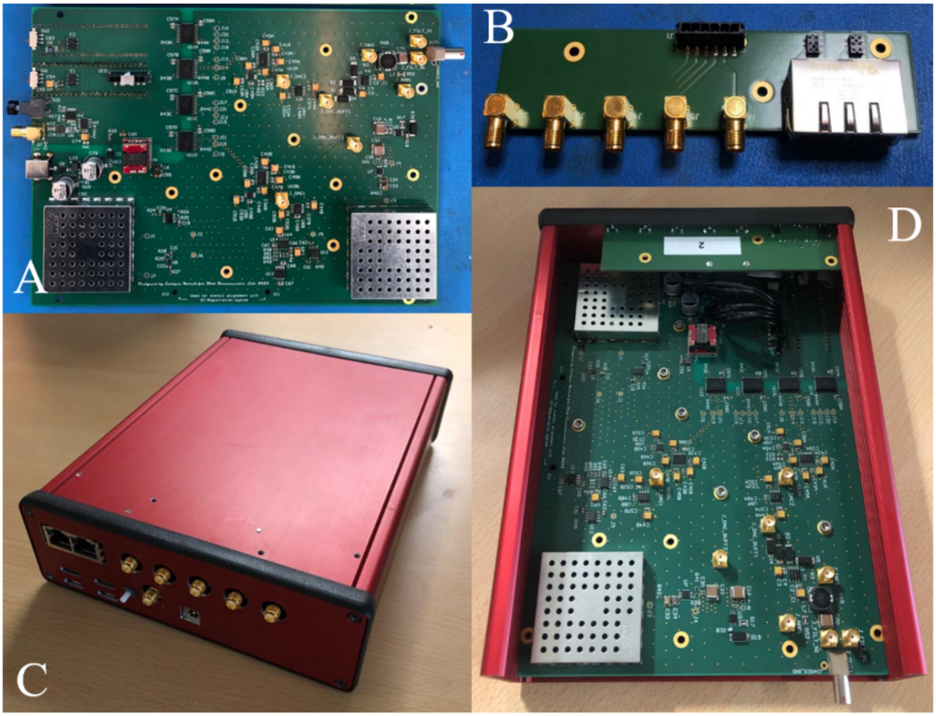

Figure 3.13: Photographs of the assembled lockbox based on the V2 PCB design by Enrique Morell.

apply the same strategy to the four-cavity filter by locking the cavities one by one. The locking protocol used for the measurements in this section was further developed by Enrique Morell and Hidde Kanger.

Figure 3.14 shows a representative single-cavity locking trace. The red curve in the top figure is the slow scan voltage generated by DAC2, which sweeps the cavity length to search for a resonance in cavity reflection signal. The black curve is the fine scan from DAC1, which is enabled only after the controller identifies a reflection dip during the DAC2 sweep. Starting from $t = 0$, the controller ramps DAC2 until the reflection signal exhibits a clear dip. It then holds DAC2 at the corresponding voltage and switches to the fine scan. During this fine scan, a dip-recognition triggers the feedback, where the PID output is routed to the piezo. In the example shown, the resonance is reached at about 1 s, and then the DAC2 value is captured and fixed near 10 V. The subsequent suppression of the fine scan indicates that the feedback has taken over and the cavity remains near the resonant operating point.

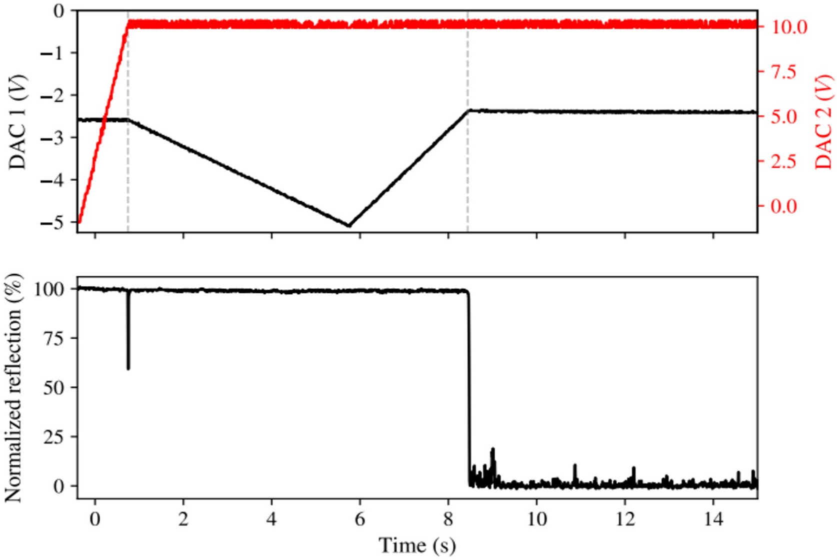

Figure 3.14: Single-cavity locking time series. The red trace (DAC2) performs a slow scan to locate the reflection dip. The black trace (DAC1) provides a short-range fine scan and triggers the PID once the dip is re-identified.

We next apply the same locking strategy to the filter cavity system. The cavities are locked one after another, and the next cavity is addressed only after the previous one has stabilized under feedback. Figure 3.15 displays the signals recorded by the four detectors. Reflection signals from all four cavities are used for locking, and an additional detector at the fiber coupled output after the fourth cavity measures the total transmitted optical power. In Figure 3.15, R_1, R_2, and R_3 approach zero after their respective cavities are locked on resonance.

To test short-term stability without active correction, we stop updating the PID outputs at $t = 0$ and keep the actuator drive voltages fixed. Even without active correction, the system maintains high throughput for several seconds. The measured total transmission exceeds 40%, close to the designed value, with the remaining discrepancy attributed primarily to optical alignment imperfections and residual environmental noise.

In the trace shown, the overall transmission remains above 90% of its locked value for roughly 4 s. This observation indicates that the differential length drift of the four cavities is sufficiently small on this timescale to preserve the aligned resonance condition, providing a practical operating window for experimental sequences that temporarily disable feedback.

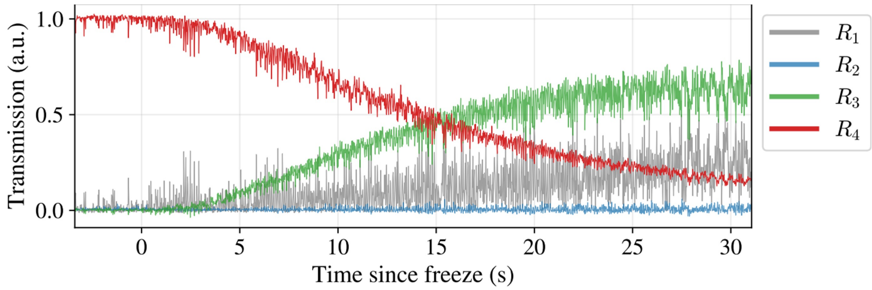

Figure 3.15: Four-cavity sequential locking and drift test. The red curve is the transmission signal through the final cavity and the other three curves are the reflection signals from the first three cavities. At $t = 0$ the feedback outputs are frozen, and the total transmission remains $> 90\%$ for $\sim 3\,$s.

3.5 Conclusions

We have presented the design and implementation of a four-cavity Fabry–Pérot filter, including the optical model and a fully custom locking system consisting of dedicated control electronics, firmware, and locking logic. A sequential reflection-based locking strategy was developed to reliably bring the cavities into resonance and establish stable operation of the complete filter system.

The measured time traces demonstrate sequential locking of the cavity chain and an optical transmission of approximately 40%, currently limited primarily by optical alignment and locking precision. In addition, a lock-freeze procedure was implemented, allowing the feedback loops to be disabled while retaining approximately 90% of the locked-state transmission for up to 4 s. The filter developed in this chapter provides the high-transmission optical filtering required for the photon-counting projection measurements needed for performing and confirming phonon quantum projection measurements. For this we first need to reach close to the quantum ground state of a mechanical system coupled to an optical cavity, as addressed in the next chapter of this thesis.