

Retroreflective surface optimisation for optical cavities with custom mirror profiles

William J. Hughes^{1,*} and Peter Horak¹

¹*Optoelectronics Research Centre, University of Southampton, SO17 1BJ, UK*

(Dated: January 19, 2026)

Coupling an emitter to a Fabry-Pérot optical cavity can provide a coherent and strong light-matter interface whose performance in a variety of applications depends critically on the emitter-photon coupling strength. Altering the typically spherical profiles of the cavity mirrors can improve this coupling strength, but directly optimising the mirror shape is numerically challenging as the multidimensional parameter space features many local optima. Here, we develop a two-step method to optimise mirror surface profiles while avoiding these issues. First, we optimise the target cavity eigenmode for the chosen application directly, and second, we construct the mirror surfaces to retroreflect this optimised target mode at both ends of the cavity. We apply our procedure to different emitter-cavity coupling scenarios. We show that mirror shaping can increase the cooperativity of coupling to a central emitter by a factor of approximately 3 across a wide range of geometries, and that, for coupling two or more emitters to a single cavity mode, the improvement factors can far exceed an order of magnitude.

I. INTRODUCTION

Fabry Pérot optical cavities can offer greatly-enhanced, and even near-deterministic interfaces between quantum emitters and single photons [1–4] while preserving high optical access [5]. These improvements promise significant performance increases in many applications of quantum technology, such as quantum networking [6–8], modular quantum computation [9, 10], and quantum-enhanced sensing [11]. These cavities have typically been built with spherical mirrors [12, 13], yielding Gaussian cavity modes with well-known properties [14, 15]. However, proposals have suggested that non-spherical mirrors could confer advantageous features to the cavity modes [16–19]. As the field pivots towards microcavities in an effort to reduce footprint, improve robustness and scalability, and to reduce interaction timescales [20–24], alternative methods to traditional superpolishing are required to create the highly curved mirrors [20, 25]. This makes mirror shaping an increasingly compelling prospect, as micromirror fabrication techniques like focused ion beam milling [26] and, to a lesser extent, laser ablation [27, 28] offer control over the surface topography [29–33].

However, the optimisation of a cavity mirror surface is numerically difficult because the relation between the mirror topography and the structure of the resonant modes is not straightforward (outside of perturbation assumptions [34] and specific geometric regimes [35]). This means that any reasonably flexible parameterisation of the surface features a vast parameter space which typically contains many local optima. It is therefore difficult to design the optimum surface for a given task through direct optimisation of the topography without working in a limited parameter space [18] that may hide the possibility for performance improvements.

In this manuscript, we present a two-stage procedure to optimise a cavity for a user-specified goal. The first step is to optimise the cavity mode, instead of the mirror surface, to take advantage of the (typically) simpler optimisation landscape. The second step is to construct the cavity surface to retroreflect this target mode. We show that this approach finds cavities whose performance exceeds any cavity with spherical mirrors, with the potential improvements particularly stark for cavities that are designed to couple to multiple emitters simultaneously.

II. MODE PARAMETERISATION AND OPTIMISATION

Our method will optimise the optical field inside Fabry-Pérot cavities, and we therefore need to describe a general intracavity field. We assume that this intracavity field is monochromatic and not highly divergent from the cavity axis, which is generally a good assumption for Fabry-Pérot cavities that are typically longer than they are wide. Under these assumptions, the intracavity field can be described using the paraxial equation. Formally, the intracavity electric field is a superposition of forward and backward running waves, $\mathbf{E}^{\pm}(\mathbf{r}, t)$ at spatial position $\mathbf{r} = (x, y, z)$ and

* w.j.hughes@soton.ac.uk

time t , which are specified through a scalar field $\epsilon^\pm(\mathbf{r})$ [36] that satisfies

$$\begin{aligned}\mathbf{E}^\pm(\mathbf{r}, t) &= \epsilon^{(u)} \epsilon^\pm(\mathbf{r}) \exp(\mp ikz) \exp(i\omega t), \\ \frac{\partial}{\partial z} \epsilon^\pm(\mathbf{r}) &= \mp \frac{i}{2k} \left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} \right) \epsilon^\pm(\mathbf{r}).\end{aligned}\quad (1)$$

with ω the angular frequency of the field, $k = \omega/c$ the wavevector, $\epsilon^{(u)}$ the fixed, unit-magnitude polarisation of the field, which must be perpendicular to the z -axis, and \pm indicating propagation towards or away from positive z respectively.

The paraxial equation (Eq. (2)) is satisfied by Gaussian beam solutions, which are particularly convenient for Fabry-Pérot cavities because they are transversely localised and satisfy the reflection boundary conditions imposed by spherical mirrors. We will be largely concerned with radially symmetric mirrors, and will therefore use the Laguerre-Gauss modes, which may be written

$$\begin{aligned}\epsilon_{l,m}^{LG,\pm}(\mathbf{r}) &= \sqrt{\frac{2l!}{\pi(l+|m|)!} \frac{1}{w(z)} \left(\frac{\sqrt{2}r}{w(z)} \right)^{|m|}} L_l^{|m|} \left(\frac{2r^2}{w^2(z)} \right) \\ &\exp \left\{ -\frac{r^2}{w^2(z)} - i \left(m\phi \pm \frac{kr^2}{2R_u(z)} \mp (2l + |m| + 1) \psi_G \right) \right\},\end{aligned}\quad (3)$$

where r , ϕ , and z are the conventional coordinates of a cylindrical polar system, with

$$\begin{aligned}w(z) &= w_0 \sqrt{1 + \left(\frac{z}{z_0} \right)^2}, \quad z_0 = \frac{\pi w_0^2}{\lambda}, \\ R_u(z) &= z \left(1 + \left(\frac{z_0}{z} \right)^2 \right), \quad \Phi_G(z) = \arctan \left(\frac{z}{z_0} \right),\end{aligned}\quad (4)$$

where $\lambda = 2\pi/k$ is the wavelength, L_n^α are the associated Laguerre polynomials, and z_0 is the Rayleigh range of the beam. The coordinate origin has been fixed to lie at the centre of the smallest waist of the beam. The Laguerre-Gauss modes are indexed by radial index $l \in \mathbb{N}_0$ and azimuthal index $m \in \mathbb{Z}$.

We will also optimise cavities in one transverse dimension (i.e. suppressing functional dependence in, and derivatives of y). For these cases, the paraxial equation can be solved with Hermite-Gauss modes

$$\epsilon_n^{HG,\pm}(\mathbf{r}) = \sqrt{\sqrt{\frac{2}{\pi}} \frac{1}{w(z) 2^n n!}} H_n \left(\frac{\sqrt{2}x}{w(z)} \right) \exp \left\{ -\frac{x^2}{w^2(z)} - i \left(\pm \frac{kx^2}{2R_u(z)} \mp \left(n + \frac{1}{2} \right) \Phi_G(z) \right) \right\}, \quad (5)$$

where H_n are the physicist's Hermite polynomials and $n \in \mathbb{N}_0$ is the transverse index of the mode. Both Laguerre-Gauss and Hermite Gauss states form complete basis sets (Laguerre-Gauss over a 2D transverse plane, Hermite-Gauss modes over a single transverse dimension) [37]. This means that any field satisfying the paraxial equation can be written as a linear superposition of the basis modes in the appropriate domain,

$$\epsilon^\pm(\mathbf{r}) = \sum_{s=0}^S C_s \epsilon_s^{\Lambda,\pm}(\mathbf{r}), \quad (6)$$

provided the basis size $S+1$ is sufficiently large, where Λ is either LG or HG , the mode indices (l, m) or n are mapped to a single index s , and C_s is the complex coefficient of the basis state $\epsilon_s^{\Lambda,\pm}$. We can thus identify any solution to the paraxial equation with a complex vector \vec{C}

$$\vec{C} = \{C_0, C_1, \dots, C_S\}^T. \quad (7)$$

Our method will optimise the cavity mode for a given task by optimising the coefficients of this vector.

The process by which the coefficients of \vec{C} are optimised depends upon the metric of performance that is being optimised. For the general case where only the performance metric (but not the gradient) can be calculated for a given set of coefficients, one may use local or global optimisation. For a substantial number of interesting cases, the performance metric may be written as the product and quotient of expectation values of Hermitian matrices. In this case, the gradient can be evaluated directly (see App. B), which allows for faster local optimisation.

III. OBTAINING MIRROR TOPOGRAPHY FROM RETROREFLECTION

In order to correspond an intracavity vector field \vec{C} to a mirror surface, we build upon the technique of Karpov et al. [38] for designing surfaces to enhance emitter-cavity coupling strength. In that technique, a target eigenmode, henceforth known as ϵ_T , with strong coupling to the emitter was chosen, and the corresponding mirror *constructed* by matching its surface to the equiphase surface contour of the target mode. If these bespoke mirrors were of infinite extent, they would perfectly retroreflect the chosen mode, which would remain unchanged (up to a phase) upon a round trip of the cavity, making ϵ_T an eigenmode of the cavity. Here, instead of starting with an ansatz for ϵ_T designed to improve coupling to a central emitter, we optimise the ϵ_T to maximise a metric of performance that is chosen according to the application at hand.

To construct the mirror surface, we need to find the equiphase contour of the target electric field. The phase of the electric field is composed of the phase of the paraxial mode ϵ_T , and the propagating wave phase $\exp(\mp ikz)$ (see Eq. (1)). For each coordinate $\mathbf{r} = (\mathbf{r}_\perp, z_{m_\pm})$ on the plane for the positive (negative) z mirror, where \mathbf{r}_\perp is some parameterisation of the transverse coordinate, the incoming paraxial mode ϵ_T^+ has a phase, whose differences across the mirror profile should be offset by the running wave phase incident on the varying surface protrusion $f_\pm(\mathbf{r}_\perp)$ which, by convention, points towards the centre of the cavity. Mathematically,

$$kf_\pm(\mathbf{r}_\perp) \pmod{2\pi} = \mp \arg[\epsilon_T^+(r, z_{m_\pm})]. \quad (8)$$

Strictly speaking, $f_\pm(\mathbf{r}_\perp)$ is only an equiphase contour on the assumption that the phase of ϵ_T changes only negligibly as z changes between the nominal mirror plane and the protruded surface. We do not correct this assumption because the discrepancy generally has little impact, and because the mode mixing method, which we use to calculate the cavity eigenmodes, also makes this assumption [39].

There is, however, a complication with the logic of retroreflective optimisation. The target mode ϵ_T is an eigenmode on the assumption that it can be retroreflected. However, if there is any clipping loss, the chosen mode cannot be perfectly retroreflected because the field leaking beyond the mirror cannot be retroreflected. This means that we must calculate the true eigenmode of the cavity, ϵ_E , instead of assuming that ϵ_T is an eigenmode. Additionally, the performance calculated for ϵ_E might not be strictly the maximum possible value, only a lower bound to it, because it was ϵ_T , not ϵ_E , that was optimised. In practice, most sensible metrics of cavity performance will discourage clipping loss that significantly reduces the number of cavity round trips. This means that the clipping loss of the target mode ϵ_T is likely to be low, so little power in ϵ_T will leak beyond the mirrors, and therefore ϵ_E is likely to be very similar to ϵ_T . Empirical evidence for this is presented in App. C.

Finally, the surface construction method can encounter problems for arbitrary modes with two transverse dimensions, because these can feature phase vortices, which are points where the amplitude goes to zero, but the phase change encountered on a path encircling the vortex does not sum to zero (instead summing to a multiple of 2π). The equiphase surface of a mode with a vortex can require infinitely sharp drops of size $\lambda/2$ in order to maintain phase continuity, often unavoidably through regions of significant mode intensity. While such a mirror surface is a mathematical possibility, it is not a practical solution, and the infinite gradient cliffs would violate the paraxial approximation on which the method rests. To avoid the problem of phase vortices, we restrict our optimisations to situations where the phase only varies with one transverse parameter. For the majority of cases, this parameter will be the cylindrical polar radius r with the assumption of complete azimuthal rotation symmetry, which is guaranteed by using only Laguerre-Gauss modes of $m = 0$. We will also use the Hermite-Gauss modes, whose transverse parameter is the Cartesian x coordinate, for cases that break radial symmetry but maintain translational symmetry along the other transverse direction (y). Though this is a restriction, we note that many cavities are close to satisfying one of these cases, and for those that are not, one may still obtain qualitative or pedagogical insight.

To calculate the non-Gaussian eigenmodes of cavities with non-spherical mirrors, we use mode mixing, as described in [40], which has been used in a variety of theoretical [17, 19, 38, 41, 42] and experimental studies [30, 43, 44]. This requires the evaluation of the surface through Eq. (8) at multiple transverse coordinates. The modulus is taken to give the smoothest surface (smallest vertical change between adjacent points). In our calculations, if the sum of any two adjacent changes exceeded a phase of π , the number of samples was deemed insufficient, and the process repeated with double the sampling density.

IV. EXAMPLE: PHOTON EXTRACTION FROM CAVITIES

A. Identifying the performance metric

We begin with the simple example of a single emitter in the centre of an axially symmetric optical cavity, with our goal to maximise the probability of extracting a single photon from the emitter in a single attempt. This figure

of merit is very important for emitter-cavity systems used as single photon sources [45], but also for systems that produce emitter-photon entanglement to use in quantum repeaters and quantum networks [7], where the operation rate depends strongly on the photon emission probability [8, 9, 46–48].

The interaction between the emitter and the cavity mode is traditionally parameterised by g , the rate of coherent interaction between the emitter excitation and the cavity mode, κ , the amplitude decay rate of the cavity mode field, and γ , the amplitude decay rate of the emitter’s excited state to spontaneous emission. The cavity decay rate can be subdivided into two components $\kappa = \kappa_{\text{int}} + \kappa_T$, where κ_T is the transmissive loss through the output mirror, which can generally be controlled, and κ_{int} is the loss to other sources, which is generally minimised as far as possible. Assuming the transmissive loss κ_T is freely chosen to maximise the photon extraction probability, it can be shown that the single metric determining the maximum photon extraction probability is

$$C_{\text{int}} = \frac{g^2}{2\kappa_{\text{int}}\gamma}, \quad (9)$$

with a maximum extraction probability

$$P_{\text{ext}} = \frac{\sqrt{2C_{\text{int}} + 1} - 1}{\sqrt{2C_{\text{int}} + 1} + 1} \quad (10)$$

where this limit is satisfied in infinite time [49]. In practice, This bound holds well provided the photon is generated over a time-scale longer than the ‘critical time’ $t_{\text{crit}} = \max(1/\kappa, \kappa/g^2)$ [50, 51], which is satisfied in many experimental situations [52]. For a graphical example of how C_{int} affects P_{ext} , see App. A.

An alternative expression for C_{int} emphasises the role of the mode profile through the ‘inverse effective area’, $A_{\text{eff}}^{-1}(\mathbf{r}_e) = \epsilon^{+2}(\mathbf{r}_e) / \int_{P_{z_e}} \epsilon^{+2}(\mathbf{r}) d^2\mathbf{r}$, where P_{z_e} is the transverse plane containing the emitter, which represents the ratio of the intensity of the forward-running cavity mode at the position \mathbf{r}_e of the emitter to the total forward-running power. The expression is

$$C_{\text{int}} = \frac{3\lambda^2 A_{\text{eff}}^{-1}(\mathbf{r}_e)}{\pi \mathcal{L}_{\text{int}}}, \quad (11)$$

where λ is the emitter transition wavelength, and \mathcal{L}_{int} is the round-trip power loss to non-transmissive sources (see [53, 54] for a derivation, where the internal loss \mathcal{L}_{int} , rather than total loss, has been used in order to express C_{int} rather than the cooperativity C , and appropriate quantities have been substituted to match the current notation). Note that the expression assumes, as we shall do throughout, that the emitter is placed perfectly at the antinode of the axial standing wave, which will require a maximum of $\lambda/4$ axial translation of the emitter from its nominal position. It also assumes that the branching ratio of the excited state emission to the cavity mode, which accounts for both the spontaneous emission branching ratio, and the polarisation overlap with the cavity mode, is unity. In the case that the branching ratio is not unity, the right hand side of Eq. (11) should be multiplied by that branching ratio.

For the case of a cavity with one transverse dimension (specifically, a cavity extended in the y -direction so that the y -dependence of the cavity mode is negligible), C_{int} cannot be directly calculated. Instead, we can use the cartesian separability of the mode in this case to assert $A_{\text{eff}}^{-1}(\mathbf{r}_e) = L_{\text{eff}}^{-1}(\mathbf{r}_e) L_{y,\text{eff}}^{-1}$, where the inverse effective length in the x -direction $L_{\text{eff}}^{-1}(\mathbf{r}_e) = \epsilon^{+2}(\mathbf{r}_e) / \int_{X_{z_e}} \epsilon^{+2}(\mathbf{r}) d\mathbf{r}$, where X_{z_e} is the x -axis in the transverse plane containing the emitter. For a cavity that is infinitely extended in y , the inverse effective length in the y direction, $L_{y,\text{eff}}^{-1}$, is zero, but for any real cavity, $L_{y,\text{eff}}^{-1}$ will be finite. We can then optimise the one-transverse dimension cavity using the metric

$$C_{\text{int}}/L_{y,\text{eff}}^{-1} = \frac{3\lambda^2 L_{\text{eff}}^{-1}(\mathbf{r}_e)}{\pi \mathcal{L}_{\text{int}}}. \quad (12)$$

This means that $L_{y,\text{eff}}^{-1}$ never need be calculated to perform optimisation in cases where the y -dependence is negligible, but for a real cavity with non-infinite y extent, the value of $L_{y,\text{eff}}^{-1}$ could be supplied to yield C_{int} .

B. Evaluating the performance metric

To use the retroreflective optimisation method, we need to be able to calculate C_{int} or $C_{\text{int}}/L_{y,\text{eff}}^{-1}$ directly from the cavity mode, and, in order to use more efficient optimisation algorithms, we can calculate their gradients as well. To

do this, we can write the inverse effective area (length) at the emitter $A_{\text{eff}}^{-1}(\mathbf{r}_e)$ ($L_{\text{eff}}^{-1}(\mathbf{r}_e)$) and internal loss \mathcal{L}_{int} as (unnormalised) matrices.

The inverse area and inverse length matrix elements are given

$$A_{\text{eff } s',s}^{-1}(\mathbf{r}_e) = \left(\epsilon_{s'}^{LG,+}(\mathbf{r}_e) \right)^* \epsilon_s^{LG,+}(\mathbf{r}_e), \quad L_{\text{eff } s',s}^{-1}(\mathbf{r}_e) = \left(\epsilon_{s'}^{HG,+}(\mathbf{r}_e) \right)^* \epsilon_s^{HG,+}(\mathbf{r}_e), \quad (13)$$

and the loss matrix elements

$$\begin{aligned} \hat{\mathcal{L}}_{\text{int}} &= \hat{\mathcal{L}}_{\text{int}}^+ + \hat{\mathcal{L}}_{\text{int}}^- \\ \hat{\mathcal{L}}_{\text{int}}^\pm &= \hat{\mathcal{I}} - \left(1 - \frac{1}{2} \mathcal{L}_{\text{abs}} \right) (\hat{\rho}^\pm)^\dagger \hat{\rho}^\pm, \\ \hat{\rho}_{s',s}^\pm &= \int_{\Pi_{m_\pm}} \left(\epsilon_{s'}^{\Lambda,+}(\mathbf{r}) \right)^* \epsilon_s^{\Lambda,+}(\mathbf{r}) d\Pi, \end{aligned} \quad (14)$$

where $\hat{\mathcal{L}}_{\text{int}}^\pm$ is the internal loss matrix of the mirror at positive (negative) z coordinate, $\hat{\mathcal{I}}$ is the identity matrix, \mathcal{L}_{abs} is the round trip absorption loss (including scattering), which is the component of the internal round trip loss that does not include clipping, $\hat{\rho}^\pm$ is the clipping matrix of the aperture defined by the edge of the mirror at positive (negative) z , and Π_{m_\pm} is the finite domain (a surface for simulations with two transverse dimensions, and a length for a single transverse dimension) of the mirror at positive (negative) z . For the spherical apertures considered in the radially symmetric simulations, the elements of $\hat{\rho}^\pm$ can be evaluated using the techniques of Vajente [55]. The details of the matrix derivations are given in App. D.

C. Possible improvements in performance

With the ability to optimise the cavity mode to improve C_{int} , and thus the photon extraction probability, we study the case of an emitter placed in the centre of a cavity of length L with two identical mirrors of diameter D . The scenario represents an ideal geometry to couple strongly to the emitter in the centre of the cavity, and features prominently in experimental design [8, 56]. For the purposes of optimisation, we will take the geometric limits L and D as given since in practice L is often restricted by a minimum emitter-mirror distance to mitigate the impact of electric field noise [57–59], and the diameter D by the maximum extent of machinable surface or requirements for optical access.

We first consider the case that the mirrors are spherical, and thus the cavity modes are Gaussian. Given a cavity length L , there is a minimum viable diameter D_{crit} , below which every possible Gaussian mode sustains clipping losses that exceed the other internal losses of the cavity \mathcal{L}_{abs} . In this case, the cavity performance is low. We can derive an approximate expression for D_{crit}

$$D_{\text{crit}} = 2\sqrt{\frac{\lambda L}{\pi}} \sqrt{-\frac{1}{2} \log(\mathcal{L}_{\text{abs}})}, \quad (15)$$

from the power clipping approximation [28] (see App. E). If the diameter exceeds D_{crit} , then C_{int} may be increased by reducing the mirror radius of curvature, and thus focusing the mode more tightly on the emitter, until the increasingly-divergent mode sustains clipping losses on the mirrors that do become significant when compared to \mathcal{L}_{abs} [60]. The spherical mirror shapes, and consequent Gaussian fundamental mode, mean that only the divergence of the cavity mode can be optimised (which determines the waist of the central focus and the waist on the mirrors); there is no flexibility to further control the transverse profile of the mode as it always has a Gaussian decay.

However, the mirrors could be shaped to improve internal cooperativity in this geometry, a suggestion made in Karpov et al. [38]. There, as here, the surface was constructed to retroreflect a target mode ϵ_T , but in our case ϵ_T is optimised, whereas in their case ϵ_T followed an ansatz. Their ϵ_T is parameterised by a number of occupied basis states N_p , with the N_p lowest-order $m = 0$ Laguerre-Gauss modes given the same amplitude coefficient, and all other basis states zero amplitude. This mode was chosen because it has the maximum intensity at the emitter for a superposition of those basis states. However, this ansatz mode does have increased divergence compared to the fundamental mode of the basis. For the remainder of this section, we will refer to this method of designing the surface as the ‘ansatz method’.

We investigate the potential of retroreflective optimisation to improve the internal cooperativity compared to any cavity with spherical mirrors, and compared to cavities designed through the ansatz method. In Fig. 1, we study an example situation, and perform the optimisation over a range of cavity geometries to explore trends in the optimisation

performance. To perform retroreflective optimisation for a given cavity geometry, we first found the best radius of curvature to use for spherical mirrors, which sets the Laguerre-Gauss basis modes of the optimisation. We then constructed the appropriate matrices to calculate C_{int} (see Sec. IV) and optimised the real parameter vector \vec{C} to find ϵ_{T} with an L-BFGS algorithm [61]. We used the first 171 $m = 0$ Laguerre-Gauss modes to construct the matrices, with the first 136 of these corresponding to free coefficients of \vec{C} (the coefficients of the remaining basis states were set to zero). The number of free coefficients was chosen to be less than the basis size to ensure that the clipping loss of the highest occupied basis state of ϵ_{T} (corresponding to the final free parameter) was correctly calculated by accounting for its overlap with basis states of somewhat larger l -index. The basis size and number of free parameters were chosen to ensure convergence of optimised C_{int} over the range of geometries investigated in Fig. 1. Details of the basis sizes required for convergence and the parameters used for all simulations in this manuscript can be found in App. F.

The first conclusion to draw from Fig. 1 is that the retroreflective optimised surface outperforms any spherical surface, and any surface optimised through the retroreflective ansatz method (Fig. 1a). This is expected, because we optimise the retroreflected mode, rather than simply taking an ansatz that should have favourable properties. The surface that achieves this optimised performance is broadly spherical, but features a very particular pattern of distortions at the level of 10's to 100's of nanometers (Fig. 1b).

We can interpret the performance increase by studying the optimised mode ϵ_{E} (Fig. 1c), which is determined via mode mixing on the surface derived from the target mode ϵ_{T} , but is, in practice, very similar to ϵ_{T} . The intensity of the mode on the emitter is strongly increased (Fig. 1ci), which is the origin of the increased coupling. This tighter focussing is achieved by using a more divergent mode that exploits the entire mirror surface more efficiently than any Gaussian mode, as exemplified by the super-Gaussian decay in the intensity at the edge of the mirror evident in the log intensity plot (Fig. 1cii) lower). The intensity density plots (Fig. 1d) indicate how this optimised mode focusses and diverges strongly within the cavity.

Moving beyond the single example, we can perform the same optimisation procedure for a range of geometries, here varying the length and the mirror diameter. In Fig. 1e), we show the performance achievable with spherical mirrors and with optimised mirrors. The most obvious features are the regions of high and low performance, roughly delineated by the white line depicting D_{crit} , despite its crude derivation that ignores mode-mixing effects. Retroreflective optimisation can subtly move the exact boundary between the regions, but, similar to the case with spherical mirrors, cannot generally find good performance if $D < D_{\text{crit}}$. However, retroreflective optimisation can significantly increase C_{int} in the high performance region ($D > D_{\text{crit}}$). The scale of improvement tends to increase with increasing mirror diameter before saturating at an improvement factor of roughly 3, which is broadly constant across a large range of length scales. This suggests that ideal mirror shaping could increase internal cooperativity by roughly half an order of magnitude in the central emitter scenario provided the mirror is large enough. This increase is considerable, but the optimised mirror shapes are highly complex, and therefore likely not ideal designs for practical manufacture (see Sec. VI for further discussion on this point).

Finally, we note that optimisation also changes the g and κ parameters of the cavity-QED system. Most directly, the more intense focussing of the cavity mode on the emitter increases the emitter-cavity coupling rate g . The optimised cavity κ_{int} typically remains close to the minimum limit set by \mathcal{L}_{abs} , but the increased C_{int} means the optimum transmission of the outcoupling mirror will increase, increasing κ [49]. While the impacts of changes to g and κ on the photon extraction probability are encapsulated by Eqs. (9) and (10), these rates are also important for achieving strong coupling ($g > \kappa, \gamma$), which is important for observing vacuum-Rabi oscillations, as well as affecting the resonant selectivity of the cavity [62], and its sensitivity to birefringence [63]. If the experimental application were to place constraints on the acceptable values of g and κ beyond a desire to maximise P_{ext} , these constraints could be included in the metric optimised to select ϵ_{T} .

D. Non axially-centred emitter

The case where the emitter is in the centre of the cavity is the most natural for achieving strong coupling, but wider experimental factors may impose dissimilar geometric constraints on the mirrors, or place the emitter away from the centre of the cavity. In this case, we expect the optimised cavity to have an asymmetric mode, with the mirrors on the left and right having different profiles. Retroreflective optimisation can also be used in this case

Example data is shown in Fig. 2, which studies a cavity where the emitter is moved towards the right mirror, and additionally where the right mirror is larger than the left. To achieve the strongest focus at the emitter, the cavity with spherical mirror must find the best compromise between two factors. On one hand, the Gaussian mode would ideally focus on the emitter so that the highest intensity of the mode is at the pertinent location. However, in this case, it is the left mirror, which is both smaller and further from the emitter than the right mirror, that limits the mode divergence. By moving the focus towards the left mirror, more strongly focussing modes can fit inside the cavity

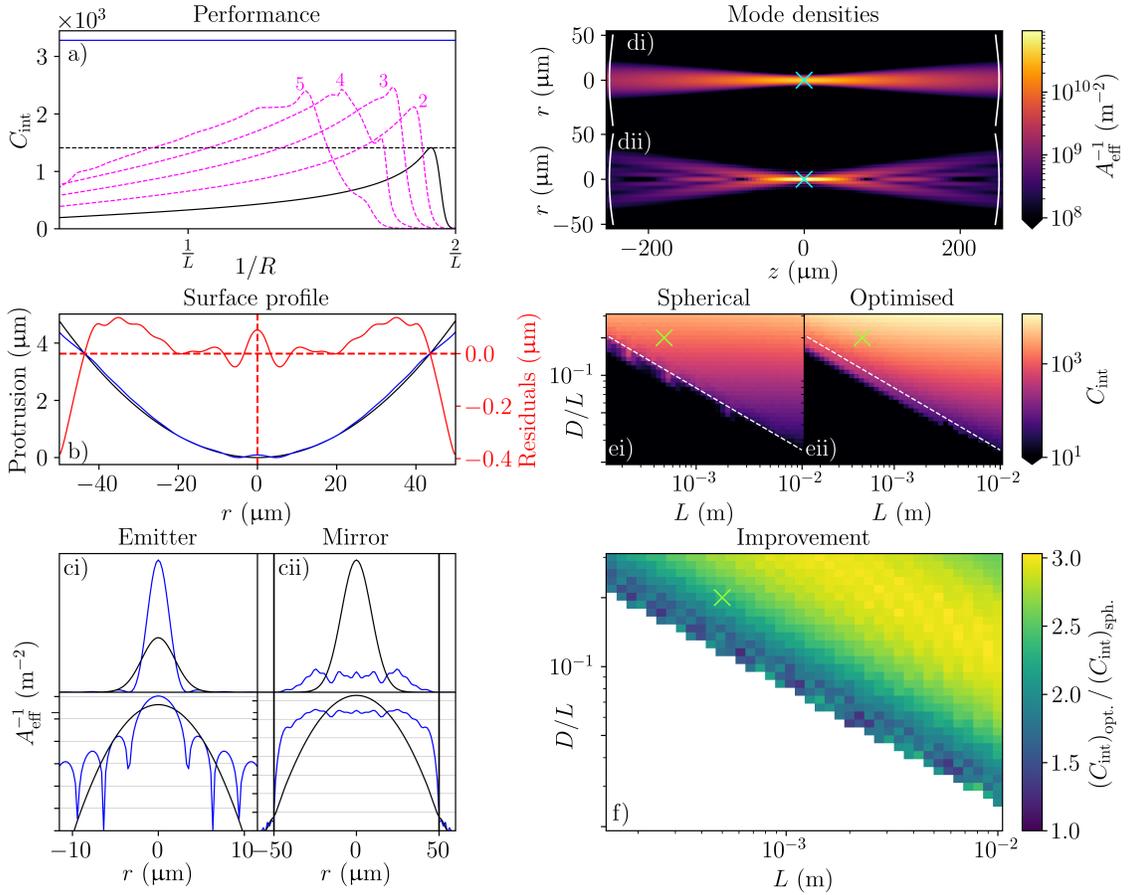


FIG. 1. Example and summary of improvements in C_{int} achievable with retroreflective optimisation for a single emitter in the centre of an optical cavity. The example cavity has a length of $500 \mu\text{m}$, a mirror diameter of $70 \mu\text{m}$, resonant wavelength of 854 nm , and non-clipping internal loss of 20 ppm . a) The C_{int} as a function of nominal mirror radius of curvature R for (black) a spherical mirror and (dashed purple) the mirrors from the ansatz method, where the number of superposed states is labeled around the peak of the corresponding line. Horizontal lines depict the maximum C_{int} achievable with a (dashed black) spherical mirror, and (blue) retroreflective optimised mirror. b) The surface profile of the best spherical mirror (black) and the optimised mirror (blue), with the residuals overlaid (red). c) Comparison of the mode intensities (linear scale in top row, log-10 scale in bottom row with gridlines indicating decades) for the best spherical mirror cavity (black) and the optimised cavity (blue) in the i) central transverse plane containing the emitter and ii) mirror transverse plane, where the vertical lines mark the mirror edges. d) Mode intensities of the i) best spherical mirror cavity, ii) retroreflective-optimised cavity in an xz cross-section of the mode ($y = 0$). The cyan cross indicates the emitter position. e) The maximum C_{int} achievable with spherical mirrors and optimised mirrors. The white dashed line depicts D_{crit} . f) The improvement in C_{int} from choosing the optimised surface over the optimised spherical surface. No data shown where $D < D_{\text{crit}}$, as these cavities do not have high performance. The green crosses in ei), eii), and f) show the configuration exemplified in panels a)-d).

without suffering overwhelming clipping loss. Ultimately, this balance means that the best spherical mirrors produce a Gaussian mode that focuses slightly closer to the centre of the cavity than the emitter (Fig. 2a). The retroreflective optimised mode again fits precisely onto the left mirror, whereas the optimised intensity behaves less dramatically on the right mirror, because clipping here does not limit performance (Fig. 2c and d). The optimised mode is also slightly asymmetric about the basis focal position, something which the Gaussian mode cannot do, although the factor 2.75 improvement in C_{int} is broadly in line with those found in the symmetric case.

E. Transversely-displaced emitter

Having seen that mirror shaping can improve the coupling cooperativity of a single emitter in the centre of a symmetric cavity, and that improvements can also be found when the cavity geometry is not symmetric about the axis, we now discuss whether the cooperativity can be improved if the emitter is displaced from the centre of the

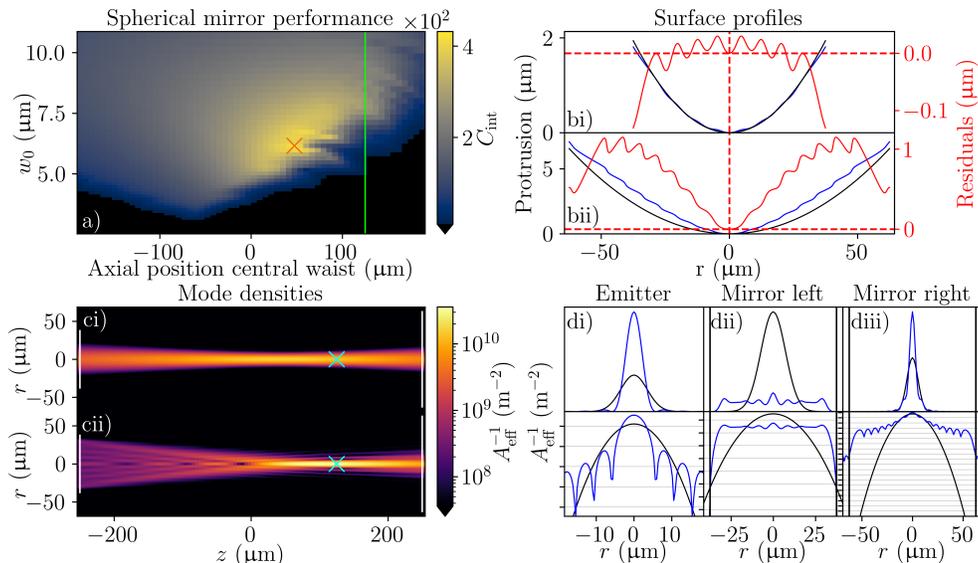


FIG. 2. Example of improvements in C_{int} achievable with retroreflective optimisation for a single emitter displaced from the centre of an asymmetric cavity geometry along the cavity axis. The example cavity has a length of $500 \mu\text{m}$, left and right mirror diameters of $75 \mu\text{m}$ and $125 \mu\text{m}$ respectively, resonant wavelength of 854 nm , non-clipping internal loss of 20 ppm , and an emitter placed $125 \mu\text{m}$ to the right of the cavity centre. a) The C_{int} for a cavity with spherical mirrors, where the curvatures of the two mirrors are set implicitly by the central waist size and position of the corresponding Gaussian beam family. The green vertical line shows the position of the emitter, and the orange cross the spherical mirror configuration with the highest C_{int} . b) The surface profile of the best spherical mirror (black) and the optimised mirror (blue), with the residuals overlaid (red) for the i) left and ii) right mirror. c) Mode intensities of the i) best spherical mirror cavity, ii) retroreflective-optimised cavity in an xz cross-section of the mode ($y = 0$). The cyan cross indicates the emitter position. d) Comparison of the mode intensities (linear scale in top row, log-10 scale in bottom row with gridlines indicating decades) for the best spherical mirror cavity (black) and the optimised cavity (blue) in the i) central transverse plane containing the emitter ii) transverse plane of the left mirror iii) transverse plane of the right mirror, where the vertical lines in ii) and iii) mark the mirror edges.

cavity in a transverse direction. As this breaks the azimuthal symmetry, we consider a cavity with just one transverse dimension, using the Hermite-Gauss modes as our basis states.

We simulate two example cases (Fig. 3), where the top row shows the emitter displaced a small amount from the axis, and the bottom row the emitter displaced close to in line with the edges of the mirrors. In the former case, we see that the best spherical mirrors focus just slightly more centrally in the cavity than the position of the emitter in order to focus most tightly (Fig. 3Ia). The optimised mode again uses the space on the mirror more efficiently to achieve a tighter focus (1.65 times improvement in $C_{\text{int}}/L_{y,\text{eff}}^{-1}$, see Fig. 3Ib), requiring small but precise modifications to the surface profile (Fig. 3Ic).

In the case of large displacement, the physics is somewhat different. Most notably, the spherical mirrors that lead to the highest performance are all found for beam family waists corresponding precisely to transverse degeneracies. The waist values can be predicted as

$$w_{\text{res}}^{q,p} = \sqrt{\frac{\lambda L}{2\pi \tan\left(\frac{\pi q}{2p}\right)}}, \quad (16)$$

where q is the number of free spectral ranges between degenerate modes, and p is the increase in the transverse mode index n between degenerate modes. These degeneracies allow the cavity modes to mix in order to avoid clipping on the mirror edge, as can be seen visually in Fig. 3IIb) (see [41] for further discussion on transverse mode mixing). The optimised mode uses asymmetry precisely to create a tightly focused mode with strong divergence towards the cavity centre, but little divergence away from it Fig. 3IIb) (bottom). This improves $C_{\text{int}}/L_{y,\text{eff}}^{-1}$ by 2.45 times from the best circular mirror case, which itself had precisely selected mirror curvature to satisfy the mode degeneracy condition.

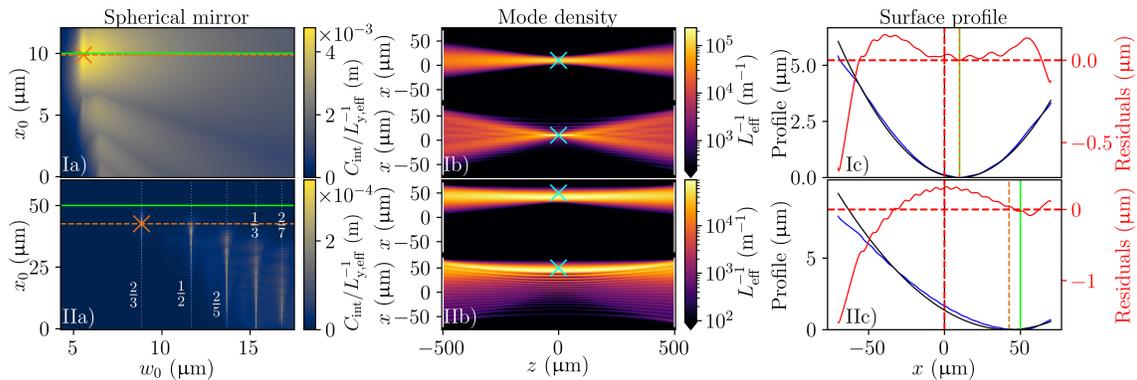


FIG. 3. Example of improvements in $C_{\text{int}}/L_{y,\text{eff}}^{-1}$ achievable with retroreflective optimisation for a single emitter displaced from the centre of a cavity geometry in a transverse direction, where the other direction is assumed large. The example cavity has a length of 1 mm, mirror diameter (full mirror length) of $140\ \mu\text{m}$, resonant wavelength of $854\ \text{nm}$, non-clipping internal loss of 20ppm, and where the emitter is displaced from the centre of the cavity in the transverse direction by row I) $10\ \mu\text{m}$ row II) $50\ \mu\text{m}$. a) The $C_{\text{int}}/L_{y,\text{eff}}^{-1}$ for a cavity with circular mirrors, where the curvature is set implicitly by the central waist size and position of the corresponding Gaussian beam family. The green horizontal line shows the position of the emitter, with the orange cross indicating the best mirror curvature parameters, and orange dotted line the x -coordinate of the centre of the best circular mirror. In IIa), the white dotted lines indicate values of $w_{\text{res}}^{q/p}$, where the corresponding fraction labels q/p , the transverse mode splitting to free spectral range ratio of the degeneracy. b) Mode intensities (per unit length) of the (top) best circular mirror cavity, (bottom) retroreflective-optimised cavity in the xz cross-section of the mode. The cyan cross indicates the emitter position. c) The surface profile of the best spherical mirror (black) and the optimised mirror (blue), with the residuals overlaid (red). The green solid vertical line indicates the transverse position of the emitter, and the orange dotted line the centre of the best circular mirror.

V. EXAMPLE: MULTIPLE EMITTERS IN CAVITIES

The case of the single emitter indicates that mirror surface optimisation can find modes that circumvent traditional limitations of Gaussian beams and consequently improve performance. In that case, the improvement was derived from finding a mode that could fit the shape of the mirror more efficiently than a Gaussian mode, and thus use all the available space on the end mirrors to achieve a tight focus. We now discuss applications involving multiple emitters in a single optical cavity, to ascertain whether improvements could be found for these use cases.

The coupling of multiple emitters to a single cavity mode could unlock two key advantages in quantum information applications. The first is providing direct coupling between these emitters, allowing for quantum gates to be performed [64, 65]. This optical mode interaction could be ‘long range’ compared to traditional interaction vectors, such as motional mode coupling for trapped ions [66] and Rydberg interactions for atoms [67], and has thus been proposed to couple small registers of qubits to increase the qubit capacity of single quantum processors [68]. The second advantage is multiplexing the input/output to an optical cavity over many emitters, and thus increase the rate at which the cavity network node interacts with the external world [10, 69–71]. Owing to the wide variety of schemes utilising direct intracavity interaction [4, 72, 73] and the different approaches to multiplexing, we do not attempt to optimise for all metrics of all possible schemes. Instead, we use the interpretation of the cooperativity as roughly representing the number of coherent interactions between a single photon and an emitter before an incoherent process occurs. This means that high cooperativity across all emitters is likely desirable. Thus, our optimisation aim for the entirety of this section will be to maximise the smallest internal cooperativity across all emitters present in the cavity, henceforth written as $C_{\text{int}}^{\text{min}}$ for brevity.

A. Coupling of two axial emitters

We will start with the simplest case of two identical emitters equally spaced to either side of the centre of the radially symmetric cavity (see Fig. 4a). The enforced axial symmetry of our chosen solutions guarantees that the internal cooperativity is the same on both emitters and therefore we can use the gradient-based optimisation discussed in Sec. II.

We first discuss our qualitative expectations of behaviour with spherical mirrors. If the separation s of the emitters is small, the task of focussing strongly on both emitters is identical to the task of focussing to the centre of the cavity. This is the exact case covered in Sec. IV, where the performance is limited by the diameter D of the mirrors, which

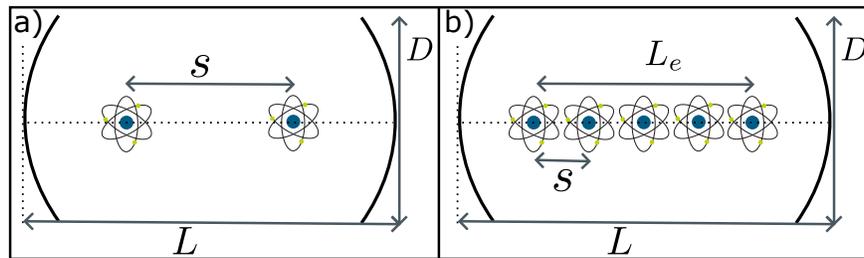


FIG. 4. Schematic diagrams of the cavity geometries discussed in a) Sec. V A and b) Sec. V B. a) Two emitters separated by s are placed symmetrically on the axis of a cavity of length L with mirror diameter D . b) An array of $N_e = 5$ emitters of length L_e , with the separation between emitters again s .

places an upper limit on the divergence of a low loss cavity mode. However, beyond a certain emitter separation s , a cavity mode optimised for central intensity will have diverged significantly at the positions of both emitters, and therefore we would prefer to reduce the mode divergence down from the upper limit. In this regime, the performance is limited by the separation of the emitters, not the diameter of the mirrors. We expect the transition between these two regimes to occur at a separation s_{crit} for which both limitations are saturated simultaneously. In this case, the best Gaussian mode for minimising beam size on both emitters has clipping loss equal to \mathcal{L}_{abs} . This leads to

$$s_{\text{crit}} = \frac{\lambda L^2}{\pi D^2} \{\log(2) - \log(\mathcal{L}_{\text{abs}})\}, \quad (17)$$

derived in App. E. Assuming $s > s_{\text{crit}}$, the best Gaussian mode is limited by the requirement that it does not diverge strongly between the centre of the cavity and the emitters. As derived in App. E, the optimum radius of curvature and internal cooperativity are

$$R = \frac{L}{2} \left(1 + \frac{s^2}{L^2} \right), \quad C_{\text{int}}^{\text{min}} = \frac{6\lambda}{\pi s \mathcal{L}_{\text{abs}}}. \quad (18)$$

We again study the ability of retroreflective optimisation to improve the performance in this scenario by performing example calculations, as displayed in Fig. 5. We again select 136 free optimisation parameters in a larger calculation basis of 171 Laguerre-Gauss basis states to converge over the range of geometries. We choose an example for which $s > s_{\text{crit}}$, and find that retroreflective optimisation does significantly increase performance (Fig. 5a). The optimised mirror shape has an inner and outer radius of curvature, with an abrupt transition between, and interference patterns around that transition (Fig. 5b). The central fit curvature of the left (right) mirror and outer fit curvature of the right (left) mirror have a geometric focus on the left (right) ion. This produces a mode which focusses strongly on both emitters (Fig. 5d). When scanning different geometries Fig. 5e), we see that, where $s < s_{\text{crit}}$, the best spherical and optimised $C_{\text{int}}^{\text{min}}$ are fairly constant. This is because the limitation to performance is the ability of the geometry to produce a strongly focussing mode, which does not depend on the separation of the emitters, and is roughly determined by the ratio D/L , which is constant in the data. However, where $s > s_{\text{crit}}$, the performance of the cavity with spherical mirrors is now limited by the separation of the emitters, and so decreases as s increases. The optimised mode retains high $C_{\text{int}}^{\text{min}}$, because it can focus effectively on both emitters. This can lead to improvements exceeding an order of magnitude in $C_{\text{int}}^{\text{min}}$ over a wide range of geometries (Fig. 5f).

B. Emitters on an axial line

We now discuss the situation of emitters placed on a line in an axial array, again targeting to maximise $C_{\text{int}}^{\text{min}}$. We space N_e emitters equally along a line of total length L_e along the axis of the cavity, with the centre of the array coinciding with the centre of the cavity (see Fig. 4b).

Example data is shown in Fig. 6, where we optimise to couple the cavity mode to an array of $N_e = 10$ emitters. Here, we see three distinct regimes in terms of performance level and improvement, determined by the relation of the length of the array (L_e), and the separation of the emitters ($s = L_e/(N_e - 1)$) to the critical separation s_{crit} (see Eq. (17)). In the short array case ($L_e < s_{\text{crit}}$) the array is so short that the optimum fundamental mode is found by focusing the Gaussian beam as tightly as possible to the centre of the cavity. In this case, the spherical mirror and optimised mirror $C_{\text{int}}^{\text{min}}$ values are, again, roughly independent of the array geometry, with improvement factors around the level found for the case of Sec. IV. In the second ‘intermediate’ regime $s < s_{\text{crit}} < L_e$, the array length necessitates

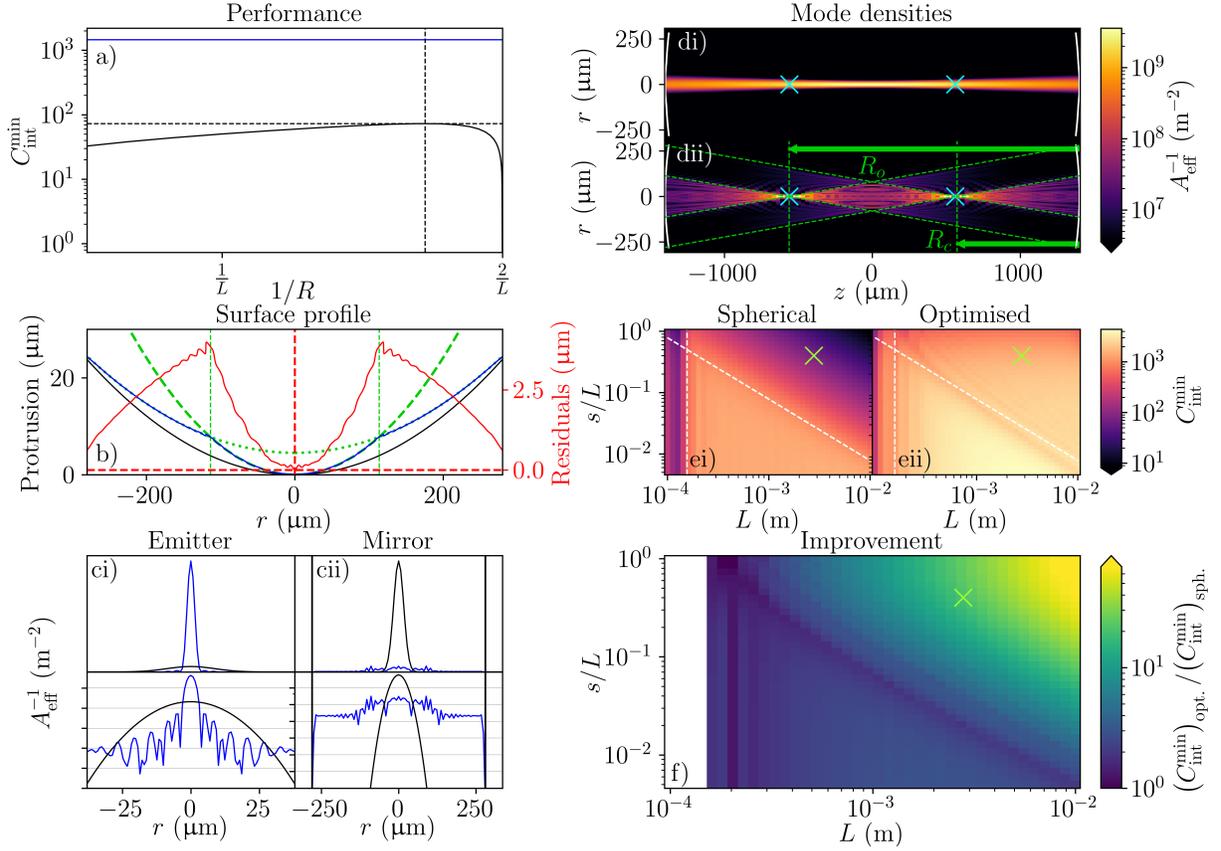


FIG. 5. Example and summary of improvements in $C_{\text{int}}^{\text{min}}$ achievable with retroreflective optimisation for two emitters placed on the axis of the cavity, symmetrically about the central point (see Fig. 4a). The example cavity (panels a-d) has length 2.8mm, mirror diameter 0.56mm, resonant wavelength 854nm, emitter separation 1.12mm, and non-clipping internal loss 20ppm. a) $C_{\text{int}}^{\text{min}}$ across the two emitters for (black line) two identical spherical mirrors as a function of mirror curvature $1/R$ with (dashed black) expected optimum $C_{\text{int}}^{\text{min}}$ and $1/R$ from Eq. (18), and (horizontal blue line) for retroreflective optimised mirrors. b) Surface profile of the (black) best spherical mirror and (blue) optimised mirror, with (red) residuals overlaid. The optimised surface is fit to a profile with an abrupt transition (dashed green vertical line) between an inner (dotted green line) and outer (dashed green line) spherical profile. c) Mode intensities (linear scale in top row, log-10 scale in bottom row with gridlines indicating decades) for the best spherical mirror cavity (black) and the optimised cavity (blue) in i) a transverse plane containing one of the emitters and ii) the mirror transverse plane, with vertical lines marking the edges of the mirror. d) Mode intensities of the i) best spherical mirror cavity, ii) retroreflective-optimised cavity in an xz cross-section of the mode ($y = 0$) with (cyan crosses) the two emitter positions. In ii) The surface fit from b) is indicated through green arrows to show the scale of the central (outer) curvature R_c (R_o) and green dashed lines to show the focal points of the central curvature from the transition radius, and the outer curvature from the mirror radius. e) $C_{\text{int}}^{\text{min}}$ across the two emitters for spherical mirrors and optimised mirrors. The white dashed diagonal line depicts s_{crit} , and the vertical line delineates $D < D_{\text{crit}}$ (left) from $D > D_{\text{crit}}$ (right). f) Improvement in the $C_{\text{int}}^{\text{min}}$ value from retroreflective optimisation. No data shown where $D < D_{\text{crit}}$. The green crosses in ei), eii), and f) show the configuration used for panels a)-d).

that the optimum spherical mirror is less concentric than the short array regime, but because $s > s_{\text{crit}}$, the optimised mode cannot focus on each emitter individually without diverging too strongly to be contained by the mirrors of diameter D . Empirically, we see that the $C_{\text{int}}^{\text{min}}$ of the spherical and optimised mirror cavities drops with increasing array length in this regime. In the final regime, $s < s_{\text{crit}}$, and the optimised mirrors are now able to focus the mode on each emitter individually (Fig. 6IIId). This intensity pattern does look like a standing wave, but it actually arises from interference of transverse modes with a periodicity much bigger than the standing wave. In this regime, the $C_{\text{int}}^{\text{min}}$ of the spherical mirror cavity continues to fall with increasing array length, but the optimised performance plateaus, allowing for larger improvement factors. This metric could not be optimised with a gradient-based optimiser, so we instead used a simplex optimisation strategy, which is likely the reason for the increased numerical noise on the data.

These results suggest two distinct situations in terms of coupling to an array of atoms. If we wish to couple to atoms in specific positions, retroreflective optimisation is able to maintain a high $C_{\text{int}}^{\text{min}}$ even for long arrays where the performance with spherical mirrors is low. This is achieved by finding modes which focus on each emitter individually.

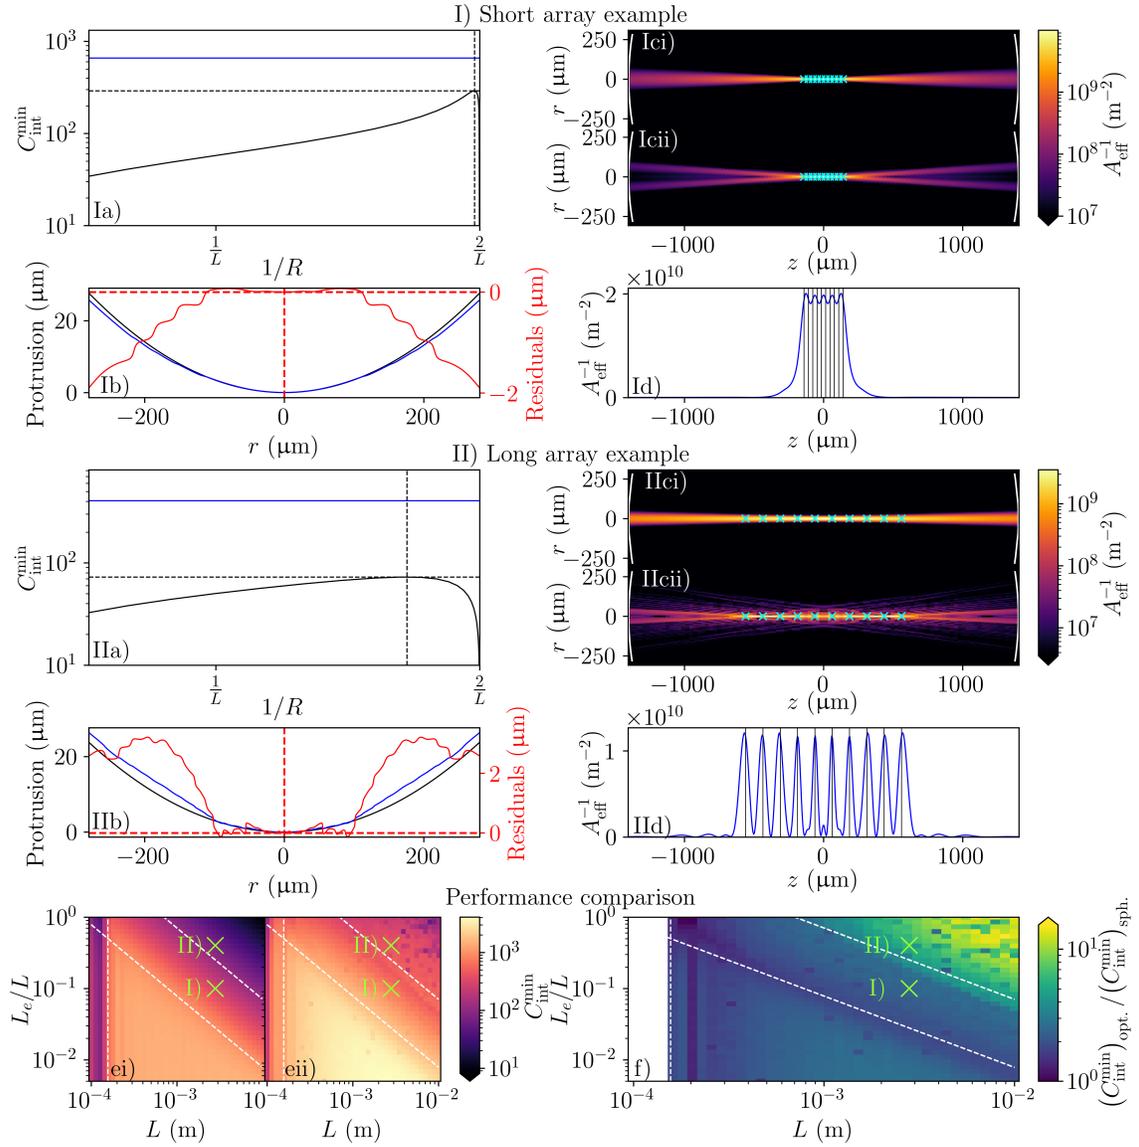


FIG. 6. Example and summary of improving $C_{\text{int}}^{\text{min}}$ across an axial emitter array using retroreflective optimisation. The example cavity has length 2.8 mm, mirror diameter 0.56 mm, resonant wavelength 854 nm, non-clipping internal loss 20ppm, and $N_e = 10$ equispaced emitters. The short (long) array of panels I (II) has length 0.28 mm (1.12 mm). a) Achievable $C_{\text{int}}^{\text{min}}$ for (black line) two identical spherical mirrors as a function of mirror curvature and (horizontal blue line) retroreflective optimised mirrors. The dashed black lines show the expected optimum $C_{\text{int}}^{\text{min}}$ and corresponding radius of curvature from Eq. (18), but where the s in that expression is taken to be L_e , the separation of the two furthest emitters. b) The surface profile of the best spherical mirror (black) and the optimised mirror (blue), with the residuals overlaid (red). c) Mode intensities of the i) best spherical mirror cavity ii) retroreflective-optimised cavity in an xz cross-section of the mode ($y = 0$). The cyan crosses mark the emitter positions. d) The intensity of the optimised mode along the cavity axis. The vertical lines indicate the positions of the emitters. e) The maximum $C_{\text{int}}^{\text{min}}$ across the array achievable with ei) spherical mirrors and eii) optimised mirrors. The lower white dashed diagonal line depicts $L_e = s_{\text{crit}}$, the upper white diagonal line $s = s_{\text{crit}}$, and the white vertical line $D = D_{\text{crit}}$. f) The improvement in $C_{\text{int}}^{\text{min}}$ from choosing the optimised surface. No data shown where $D < D_{\text{crit}}$. The green crosses in ei), eii), and f) show the configuration used for panels a)-d), where the cross label 'I' or 'II' corresponds to the example case.

However, if the atoms are not placed in specific positions, but rather anywhere along the axis, we would effectively like to couple strongly everywhere along the array (i.e. $s \rightarrow 0$ with L_e constant). In this case, we are restricted to operating in the intermediate regime and the potential cooperativity improvements from mirror shaping, though significant, are not orders of magnitude.

VI. PROSPECT FOR MANUFACTURE

The retroreflective optimisation method approximately finds the best possible cavity performance under geometrical constraints for arbitrary mirror shaping. This method is thus well-suited to finding the performance limits imposed by the laws of physics and the corresponding mirror shapes. Perhaps most valuably, it shows how significantly mirror shaping could improve performance in the target application without assuming a limited set of surface parameterisations.

The method does not, however, guarantee that the optimised surfaces are easy to manufacture, and the optimised surfaces (see Figs. 1, 2, 3, 5, and 6) typically harbour high-spatial frequency features. To find the precision of the spatial shaping required to make these mirrors, we take example optimised surfaces, remove frequency components above a cutoff shaping scale with a fourth-order Butterworth filter, and find the performance of cavities with the resulting surfaces. Example results are shown in Fig. 7 which indicate (row ii) very low performance if the cutoff scale is too low and near-optimised performance for high spatial resolution. Across the examples, the transition between these low and high performance regions occurs at a shaping scale of approximately

$$(1/\lambda_s)_{\text{crit}} = \frac{D}{\lambda L}, \quad (19)$$

which is the spatial frequency that would diffract a beam from parallel to the cavity axis to parallel with the line joining the edges of opposite mirrors (i.e. the diffraction allowed by the cavity solid angle). However, the sharp features of the mirror that couples the cavity mode to two emitters (Fig. 7 column c) requires shaping slightly more precise than this.

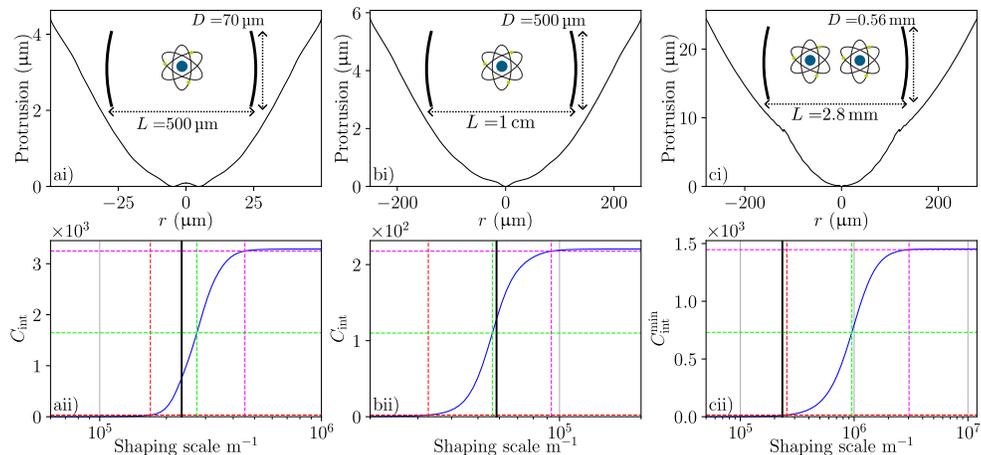


FIG. 7. The scale of surface shaping required to realise different performance levels for retroreflective optimised mirrors. Row i) The surface profiles of the optimised mirrors, with schematic of the cavity geometry and emitter configuration. Row ii) The performance of the cavity when the mirrors are shaped by spatial scale components up to the specified cutoff spatial frequency. The red, green, and purple crosshairs indicate the shaping precision required to reach 1%, 50%, and 99% of the optimised performance respectively. The black vertical line indicates $(1/\lambda_s)_{\text{crit}}$. a) Example for a cavity coupling to a single central emitter used in Fig. 1 ($L = 500 \mu\text{m}$, $D = 70 \mu\text{m}$, $\lambda = 854 \text{ nm}$, and $\mathcal{L}_{\text{abs}} = 20 \text{ ppm}$). b) Example of a different cavity geometry ($L = 1 \text{ cm}$, $D = 500 \mu\text{m}$, $\lambda = 854 \text{ nm}$, and $\mathcal{L}_{\text{abs}} = 20 \text{ ppm}$) again coupling to a single central emitter. c) Example case from Fig. 5 of a cavity coupling to two emitters placed symmetrically about the centre of the cavity on axis ($L = 2.8 \text{ mm}$, $D = 0.56 \text{ mm}$, $\lambda = 854 \text{ nm}$, and $\mathcal{L}_{\text{abs}} = 20 \text{ ppm}$).

These required shaping scales, though still much larger than the smallest beam diameters of focused ion beam milling [29], are very demanding. We would therefore recommend that retroreflective optimisation is used as part of a wider method to find promising mirror shapes for manufacture, rather than designing the surfaces itself. In this approach, the retroreflective optimisation method identifies applications for which mirror shaping holds promise to significantly outperform cavities with spherical mirrors, inspires simpler mirror designs that slightly compromise performance for improved manufacturability, and contextualises the performance of these simpler designs against the ultimate physical limits of the geometry. For a full example of this procedure applied to plano-concave cavities see Hughes and Horak [74].

VII. CONCLUSION

We have developed a method for optimising optical cavities to maximise user-specified metrics. Instead of directly optimising the mirror surface profile, which is vulnerable to finding local optima in a chosen parameterisation, our method directly optimises the cavity mode, before finding the appropriate surface profile to retroreflect that chosen mode. The mode optimisation procedure takes advantage of a typically simple optimisation landscape, even for large parameter spaces, to find solutions that far outperform standard Gaussian modes. The surface construction component typically works well provided the cavity has low clipping loss, which is generally already required of cavities used for quantum technologies.

We applied this method to optimise cavities to couple one or many emitters to the cavity mode. We showed that an optimised surface can significantly improve coupling to a single emitter, either central in the cavity or displaced axially or transversely, by using the available mirror space more efficiently. For cases involving multiple emitters along the axis of a cavity, our method found modes with the ability to focus modes on multiple emitters simultaneously, increasing the performance in certain geometric regimes by more than an order of magnitude.

We anticipate first that our method and results will highlight the general potential of mirror shaping to overcome the limitation imposed by the Gaussian eigenmodes of cavities with spherical mirrors. Secondly, the particularly significant improvements for coupling to multiple emitters, suggest that shaped mirrors could be especially useful for these cases, which find applications in scaling quantum technologies. Finally, though the method tends to produce complex surface profiles that would be challenging to fabricate, these can serve to inspire more manageable designs, and the potential for improvement revealed by retroreflective optimisation could identify applications for which mirror shaping holds promise.

Finally, our method itself relies only upon the paraxial approximation and the assumption of low clipping loss, common operating conditions for Fabry-Perot cavities. This offers considerable flexibility over which metric is optimised, and could allow for the optimisation of cavity geometries in alternative scenarios to those studied here, such as collective coupling to probe ultrastrong coupling [75] or many-body physics [76, 77], using vacuum coupling to control chemical reactions [78–80], or even applications outside cavity quantum electrodynamics.

ACKNOWLEDGMENTS

This work was funded by the UK Engineering and Physical Sciences Research Council Hub in Quantum Computing and Simulation (EP/T001062/1) and Hub for Quantum Computing via Integrated and Interconnected Implementations (EP/Z53318X/1). The authors acknowledge the use of the IRIDIS High Performance Computing Facility, and associated support services at the University of Southampton, in the completion of this work. Data underlying the results presented in this paper are available at Ref. [81]. The authors declare no conflicts of interest.

-
- [1] M. Keller, B. Lange, K. Hayasaka, W. Lange, and H. Walther, Continuous generation of single photons with controlled waveform in an ion-trap cavity system, *Nature* **431**, 1075 (2004).
 - [2] J. McKeever, A. Boca, A. D. Boozer, R. Miller, J. R. Buck, A. Kuzmich, and H. J. Kimble, Deterministic generation of single photons from one atom trapped in a cavity, *Science* **303**, 1992 (2004).
 - [3] J. Bochmann, M. Mücke, C. Guhl, S. Ritter, G. Rempe, and D. L. Moehring, Lossless state detection of single neutral atoms, *Phys. Rev. Lett.* **104**, 203601 (2010).
 - [4] B. Grinkemeyer, E. Guardado-Sanchez, I. Dimitrova, D. Shchepanovich, G. E. Mandopoulou, J. Borregaard, V. Vuletić, and M. D. Lukin, Error-detected quantum operations with neutral atoms mediated by an optical cavity, *Science* **387**, 1301 (2025).
 - [5] L. Hartung, M. Seubert, S. Welte, E. Distante, and G. Rempe, A quantum-network register assembled with optical tweezers in an optical cavity, *Science* **385**, 179 (2024).
 - [6] S. Ritter, C. Nölleke, C. Hahn, A. Reiserer, A. Neuzner, M. Uphoff, M. Mücke, E. Figueroa, J. Bochmann, and G. Rempe, An elementary quantum network of single atoms in optical cavities, *Nature* **484**, 195 (2012).
 - [7] A. Reiserer and G. Rempe, Cavity-based quantum networks with single atoms and optical photons, *Rev. Mod. Phys.* **87**, 1379 (2015).
 - [8] V. Krutyanskiy, M. Canteri, M. Meraner, J. Bate, V. Krcmarsky, J. Schupp, N. Sangouard, and B. P. Lanyon, Telecom-wavelength quantum repeater node based on a trapped-ion processor, *Phys. Rev. Lett.* **130**, 213601 (2023).
 - [9] C. Monroe, R. Raussendorf, A. Ruthven, K. R. Brown, P. Maunz, L.-M. Duan, and J. Kim, Large-scale modular quantum-computer architecture with atomic memory and photonic interconnects, *Phys. Rev. A* **89**, 022317 (2014).
 - [10] Y. Li and J. D. Thompson, High-rate and high-fidelity modular interconnects between neutral atom quantum processors, *PRX Quantum* **5**, 020363 (2024).

- [11] B. Petrak, N. Djeu, and A. Muller, Purcell-enhanced raman scattering from atmospheric gases in a high-finesse microcavity, *Phys. Rev. A* **89**, 023811 (2014).
- [12] C. K. Law and H. J. K. and, Deterministic generation of a bit-stream of single-photon pulses, *J. Mod. Optic.* **44**, 2067 (1997).
- [13] V. Krutyanskiy, M. Meraner, J. Schupp, V. Krčmarsky, H. Hainzer, and B. P. Lanyon, Light-matter entanglement over 50 km of optical fibre, *npj Quantum Inf.* **5**, 72 (2019).
- [14] A. E. Siegman, *Lasers* (University Science Books, 1986).
- [15] A. Yariv, *Quantum Electronics* (Wiley, New York, 1991).
- [16] F. Ferdous, A. A. Demchenko, S. P. Vyatchanin, A. B. Matsko, and L. Maleki, Microcavity morphology optimization, *Phys. Rev. A* **90**, 033826 (2014).
- [17] N. Podoliak, H. Takahashi, M. Keller, and P. Horak, Harnessing the mode mixing in optical fiber-tip cavities, *J. Phys. B: At. Mol. Opt. Phys.* **50**, 085503 (2017).
- [18] D. Karpov and P. Horak, Evolutionary algorithm to design high-cooperativity optical cavities, *New J. Phys.* **24**, 073028 (2022).
- [19] D. V. Karpov, S. Kurdiymov, and P. Horak, Convolutional neural networks for mode on-demand high finesse optical resonator design, *Sci. Rep.* **13**, 15567 (2023).
- [20] M. Steiner, H. M. Meyer, C. Deutsch, J. Reichel, and M. Köhl, Single ion coupled to an optical fiber cavity, *Phys. Rev. Lett.* **110**, 043003 (2013).
- [21] C. Saavedra, D. Pandey, W. Alt, H. Pfeifer, and D. Meschede, Tunable fiber fabry-perot cavities with high passive stability, *Opt. Express* **29**, 974 (2021).
- [22] H. Pfeifer, L. Ratschbacher, J. Gallego, C. Saavedra, A. Faßbender, A. von Haaren, W. Alt, S. Hofferberth, M. Köhl, S. Linden, *et al.*, Achievements and perspectives of optical fiber fabry-perot cavities, *Appl. Phys. B* **128**, 29 (2022).
- [23] M. Teller, V. Messerer, K. Schüppert, Y. Zou, D. A. Fioretto, M. Galli, P. C. Holz, J. Reichel, and T. E. Northup, Integrating a fiber cavity into a wheel trap for strong ion-cavity coupling, *AVS Quantum Sci.* **5**, 012001 (2023).
- [24] R. Zifkin, C. D. Rodríguez Rosenblueth, E. Janitz, Y. Fontana, and L. Childress, Lifetime reduction of single germanium-vacancy centers in diamond via a tunable open microcavity, *PRX Quantum* **5**, 030308 (2024).
- [25] N. Jin, C. A. McLemore, D. Mason, J. P. Hendrie, Y. Luo, M. L. Kelleher, P. Kharel, F. Quinlan, S. A. Diddams, and P. T. Rakich, Micro-fabricated mirrors with finesse exceeding one million, *Optica* **9**, 965 (2022).
- [26] P. R. Dolan, G. M. Hughes, F. Grazioso, B. R. Patton, and J. M. Smith, Femtoliter tunable optical cavity arrays, *Opt. Lett.* **35**, 3556 (2010).
- [27] T. Steinmetz, Y. Colombe, D. Hunger, T. Hänsch, A. Balocchi, R. Warburton, and J. Reichel, Stable fiber-based fabry-pérot cavity, *Appl. Phys. Lett.* **89**, 111110 (2006).
- [28] D. Hunger, T. Steinmetz, Y. Colombe, C. Deutsch, T. W. Hänsch, and J. Reichel, A fiber fabry-perot cavity with high finesse, *New J. Phys.* **12**, 065038 (2010).
- [29] A. A. P. Trichet, P. R. Dolan, D. M. Coles, G. M. Hughes, and J. M. Smith, Topographic control of open-access microcavities at the nanometer scale, *Opt. Express* **23**, 17205 (2015).
- [30] B. T. Walker, B. J. Ash, A. A. P. Trichet, J. M. Smith, and R. A. Nyman, Bespoke mirror fabrication for quantum simulation with light in open-access microcavities, *Opt. Express* **29**, 10800 (2021).
- [31] P. Maier, S. Rupp, N. Lettner, J. H. Denschlag, and A. Kubanek, Fabrication of customized low-loss optical resonators by combination of fib-milling and co2 laser smoothing, *Opt. Express* **33**, 19205 (2025).
- [32] K. Ott, S. Garcia, R. Kohlhaas, K. Schüppert, P. Rosenbusch, R. Long, and J. Reichel, Millimeter-long fiber fabry-perot cavities, *Opt. Express* **24**, 9839 (2016).
- [33] S. Gao, V. Kavungal, S. Oya, D. Okuno, E. Kassa, W. J. Hughes, P. Horak, and H. Takahashi, Profile control of fiber-based micro-mirrors using adaptive laser shooting with in situ imaging, *Opt. Express* **33**, 39009 (2025).
- [34] M. P. van Exter, M. Wubs, E. Hissink, and C. Koks, Fine structure in fabry-perot microcavity spectra, *Phys. Rev. A* **106**, 013501 (2022).
- [35] R. A. Nyman and M. H. Szymańska, Interactions in dye-microcavity photon condensates and the prospects for their observation, *Phys. Rev. A* **89**, 033844 (2014).
- [36] N. Barré, M. Romanelli, M. Lebental, and M. Brunel, Waves and rays in plano-concave laser cavities: I. geometric modes in the paraxial approximation, *Eur. J. Phys.* **38**, 034010 (2017).
- [37] G. Nienhuis and L. Allen, Paraxial wave optics and harmonic oscillators, *Phys. Rev. A* **48**, 656 (1993).
- [38] D. V. Karpov and P. Horak, Cavities with nonspherical mirrors for enhanced interaction between a quantum emitter and cavity photons, *Phys. Rev. A* **105**, 023515 (2022).
- [39] E. S. Hissink, C. Koks, and M. P. van Exter, Nonparaxial corrections for short cavities and fibers, *Phys. Rev. A* **106**, 053505 (2022).
- [40] D. Kleckner, W. T. M. Irvine, S. S. R. Oemrawsingh, and D. Bouwmeester, Diffraction-limited high-finesse optical cavities, *Phys. Rev. A* **81**, 043814 (2010).
- [41] W. J. Hughes, T. H. Doherty, J. A. Blackmore, P. Horak, and J. F. Goodwin, Mode mixing and losses in misaligned microcavities, *Opt. Express* **31**, 32619 (2023).
- [42] W. J. Hughes, T. H. Doherty, J. A. Blackmore, P. Horak, and J. F. Goodwin, Efficient operator method for modeling mode mixing in misaligned optical cavities, *Phys. Rev. A* **109**, 013524 (2024).
- [43] J. Benedikter, T. Hümmer, M. Mader, B. Schlederer, J. Reichel, T. W. Hänsch, and D. Hunger, Transverse-mode coupling and diffraction loss in tunable fabry-pérot microcavities, *New J. Phys.* **17**, 053051 (2015).
- [44] J. Benedikter, T. Moosmayer, M. Mader, T. Hümmer, and D. Hunger, Transverse-mode coupling effects in scanning cavity

- microscopy, *New J. Phys.* **21**, 103029 (2019).
- [45] A. Kuhn and D. Ljunggren†, Cavity-based single-photon sources, *Contemp. Phys.* **51**, 289 (2010).
- [46] N. Sangouard, R. Dubessy, and C. Simon, Quantum repeaters based on single trapped ions, *Phys. Rev. A* **79**, 042340 (2009).
- [47] L. J. Stephenson, D. P. Nadlinger, B. C. Nichol, S. An, P. Drmota, T. G. Ballance, K. Thirumalai, J. F. Goodwin, D. M. Lucas, and C. J. Ballance, High-rate, high-fidelity entanglement of qubits across an elementary quantum network, *Phys. Rev. Lett.* **124**, 110501 (2020).
- [48] V. Krutyanskiy, M. Galli, V. Krcmarsky, S. Baier, D. A. Fioretto, Y. Pu, A. Mazloom, P. Sekatski, M. Canteri, M. Teller, J. Schupp, J. Bate, M. Meraner, N. Sangouard, B. P. Lanyon, and T. E. Northup, Entanglement of trapped-ion qubits separated by 230 meters, *Phys. Rev. Lett.* **130**, 050803 (2023).
- [49] H. Goto, S. Mizukami, Y. Tokunaga, and T. Aoki, Figure of merit for single-photon generation based on cavity quantum electrodynamics, *Phys. Rev. A* **99**, 053843 (2019).
- [50] T. Utsugi, A. Goban, Y. Tokunaga, H. Goto, and T. Aoki, Gaussian-wave-packet model for single-photon generation based on cavity quantum electrodynamics under adiabatic and nonadiabatic conditions, *Phys. Rev. A* **106**, 023712 (2022).
- [51] W. J. Hughes, J. F. Goodwin, and P. Horak, Optimizing finite-time photon extraction from emitter-cavity systems, *J. Opt. Soc. Am. B* **41**, C168 (2024).
- [52] J. Schupp, V. Krcmarsky, V. Krutyanskiy, M. Meraner, T. Northup, and B. Lanyon, Interface between trapped-ion qubits and traveling photons with close-to-optimal efficiency, *PRX Quantum* **2**, 020331 (2021).
- [53] H. Tanji-Suzuki, I. D. Leroux, M. H. Schleier-Smith, M. Cetina, A. T. Grier, J. Simon, and V. Vuletić, Chapter 4 - interaction between atomic ensembles and optical resonators: Classical description, in *Advances in Atomic, Molecular, and Optical Physics*, Advances In Atomic, Molecular, and Optical Physics, Vol. 60, edited by E. Arimondo, P. Berman, and C. Lin (Academic Press, 2011) pp. 201–237.
- [54] K. C. Cox, D. H. Meyer, N. A. Schine, F. K. Fatemi, and P. D. Kunz, Increased atom-cavity coupling and stability using a parabolic ring cavity, *J. Phys. B: At. Mol. Opt. Phys.* **51**, 195002 (2018).
- [55] G. Vajente, Analytical expressions for the clipping of laguerre- and hermite-gauss modes by circular apertures, *J. Opt. Soc. Am. A* **39**, 352 (2022).
- [56] H. Takahashi, E. Kassa, C. Christoforou, and M. Keller, Cavity-induced anticorrelated photon-emission rates of a single ion, *Phys. Rev. A* **96**, 023824 (2017).
- [57] M. Teller, D. A. Fioretto, P. C. Holz, P. Schindler, V. Messerer, K. Schüppert, Y. Zou, R. Blatt, J. Chiaverini, J. Sage, and T. E. Northup, Heating of a trapped ion induced by dielectric materials, *Phys. Rev. Lett.* **126**, 230505 (2021).
- [58] Y.-T. Chen, M. Szurek, B. Hu, J. de Hond, B. Braverman, and V. Vuletic, High finesse bow-tie cavity for strong atom-photon coupling in rydberg arrays, *Opt. Express* **30**, 37426 (2022).
- [59] E. Kassa, S. Gao, S. Teh, D. van Dinter, and H. Takahashi, How to integrate a miniature optical cavity in a linear ion trap: Shielding dielectrics and trap symmetry, *Phys. Rev. Appl.* **23**, 024038 (2025).
- [60] S. Gao, J. A. Blackmore, W. J. Hughes, T. H. Doherty, and J. F. Goodwin, Optimization of scalable ion-cavity interfaces for quantum photonic networks, *Phys. Rev. Appl.* **19**, 014033 (2023).
- [61] C. Zhu, R. H. Byrd, P. Lu, and J. Nocedal, Algorithm 778: L-bfgs-b: Fortran subroutines for large-scale bound-constrained optimization, *ACM Trans. Math. Softw.* **23**, 550–560 (1997).
- [62] T. D. Barrett, D. Stuart, O. Barter, and A. Kuhn, Nonlinear zeeman effects in the cavity-enhanced emission of polarised photons, *New J. Phys.* **20**, 073030 (2018).
- [63] E. Kassa, W. Hughes, S. Gao, and J. F. Goodwin, Effects of cavity birefringence in polarisation-encoded quantum networks, *New J. Phys.* **25**, 013004 (2023).
- [64] L.-M. Duan and H. J. Kimble, Efficient engineering of multiatom entanglement through single-photon detections, *Phys. Rev. Lett.* **90**, 253601 (2003).
- [65] H. Takahashi, P. Nevado, and M. Keller, Mølmer-sørensen entangling gate for cavity qed systems, *J. Phys. B: At. Mol. Opt. Phys.* **50**, 195501 (2017).
- [66] K. Mølmer and A. Sørensen, Multiparticle entanglement of hot trapped ions, *Phys. Rev. Lett.* **82**, 1835 (1999).
- [67] D. Jaksch, J. I. Cirac, P. Zoller, S. L. Rolston, R. Côté, and M. D. Lukin, Fast quantum gates for neutral atoms, *Phys. Rev. Lett.* **85**, 2208 (2000).
- [68] J. Ramette, J. Sinclair, Z. Vendeiro, A. Rudelis, M. Cetina, and V. Vuletić, Any-to-any connected cavity-mediated architecture for quantum computing with trapped ions or rydberg arrays, *PRX Quantum* **3**, 010344 (2022).
- [69] J. Sinclair, J. Ramette, B. Grinkemeyer, D. Bluvstein, M. D. Lukin, and V. Vuletić, Fault-tolerant optical interconnects for neutral-atom arrays, *Phys. Rev. Res.* **7**, 013313 (2025).
- [70] B. Hu, J. Sinclair, E. Bytyqi, M. Chong, A. Rudelis, J. Ramette, Z. Vendeiro, and V. Vuletić, Site-selective cavity readout and classical error correction of a 5-bit atomic register, *Phys. Rev. Lett.* **134**, 120801 (2025).
- [71] S. Kikura, K. Tanji, A. Goban, and S. Sunami, Passive Quantum Interconnects: High-Fidelity Quantum Networking at Higher Rates with Less Overhead, arXiv e-prints , arXiv:2507.01229 (2025).
- [72] L.-M. Duan, B. Wang, and H. J. Kimble, Robust quantum gates on neutral atoms with cavity-assisted photon scattering, *Phys. Rev. A* **72**, 032333 (2005).
- [73] S. Welte, B. Hacker, S. Daiss, S. Ritter, and G. Rempe, Photon-mediated quantum gate between two neutral atoms in an optical cavity, *Phys. Rev. X* **8**, 011018 (2018).
- [74] W. J. Hughes and P. Horak, Advanced mirror shapes for mode enhancement in plano-concave cavities, arXiv e-prints , arXiv:2510.07251 (2025).
- [75] A. Johnson, M. Blaha, A. E. Ulanov, A. Rauschenbeutel, P. Schneeweiss, and J. Volz, Observation of collective superstrong

- coupling of cold atoms to a 30-m long optical resonator, Phys. Rev. Lett. **123**, 243602 (2019).
- [76] M. Lei, R. Fukumori, J. Rochman, B. Zhu, M. Endres, J. Choi, and A. Faraon, Many-body cavity quantum electrodynamics with driven inhomogeneous emitters, Nature (London) **617**, 271 (2023).
- [77] M. Pallmann, K. Köster, Y. Zhang, J. Heupel, T. Eichhorn, C. Popov, K. Mølmer, and D. Hunger, Cavity-mediated collective emission from few emitters in a diamond membrane, Phys. Rev. X **14**, 041055 (2024).
- [78] P. Andrew and W. L. Barnes, Förster energy transfer in an optical microcavity, Science **290**, 785 (2000).
- [79] F. Herrera and F. C. Spano, Cavity-controlled chemistry in molecular ensembles, Phys. Rev. Lett. **116**, 238301 (2016).
- [80] A. Thomas, L. Lethuillier-Karl, K. Nagarajan, R. M. A. Vergauwe, J. George, T. Chervy, A. Shalabney, E. Devaux, C. Genet, J. Moran, and T. W. Ebbesen, Tilting a ground-state reactivity landscape by vibrational strong coupling, Science **363**, 615 (2019).
- [81] W. J. Hughes and P. Horak, Dataset supporting an article "retroreflective surface optimisation for optical cavities with custom mirror profiles" (2026), doi: <https://doi.org/10.5258/SOTON/D3808>.

Appendix A: Role of internal cooperativity in photon extraction

Though the retroreflective optimisation method described in the main text can optimise cavity modes for arbitrary figures of merit, the main text focuses on optimisation of the internal cooperativity due to its critical role in many protocols. Specifically for single photon generation over adiabatic timescales, C_{int} is the single figure of merit that limits the extraction probability

$$P_{\text{ext}} = \frac{\sqrt{2C_{\text{int}} + 1} - 1}{\sqrt{2C_{\text{int}} + 1} + 1}, \quad (\text{A1})$$

originally written as Eq. (10) in the main text [49]. This extraction probability is visualised in Fig. 8 so that the practical impact of optimising the internal cooperativity can be understood. For internal cooperativities below 10, increases in the internal cooperativity result in strong increases in P_{ext} (Fig. 8 blue line). For internal cooperativities above 10, P_{ext} is already a significant fraction of unity, but improvements will dramatically reduce the probability that a photon is not generated (Fig. 8 red line). The benefits to improvements in the internal cooperativity thus depend upon the C_{int} of the system before optimisation. It should, however, be noted that the value for C_{int} quoted in this manuscript (Eq. (11)) assumes unit branching ratio to the final state. A real system with non-unity branching ratio will have C_{int} reduced by a factor of the branching ratio.

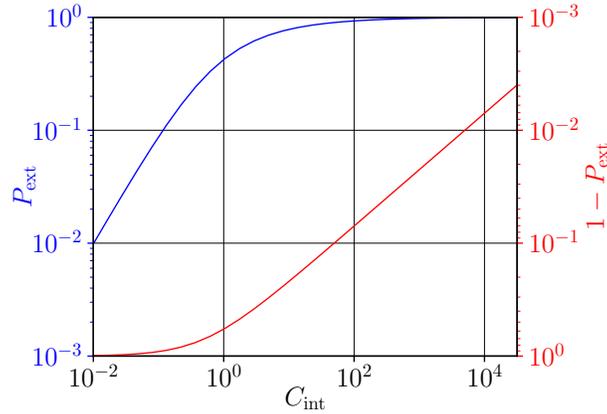


FIG. 8. The adiabatic upper limit to the probability that a photon is extracted (blue), and the lower limit to the probability that a photon is not extracted (red) from an emitter-cavity system with internal cooperativity C_{int} whose output transmission is optimised for photon extraction probability at each datapoint.

Appendix B: Gradient-based retroreflective optimisation

The retroreflective optimisation method is a two-step process that first optimises the target mode ϵ_T by adjusting the coefficients of its corresponding basis vector \vec{C} , and then constructs the surface to retroreflect this mode. The first stage (the optimisation) can proceed through many different optimisation techniques. A convenient and common case is when the metric to optimise m can be written as the product and quotient of expectation values of matrices

with the mode coefficient vector. In this case, it is possible to directly calculate the gradient of m in the parameter space, and thus use optimisation techniques that exploit the direct evaluation of the gradient. In this appendix, we discuss the mathematics of this approach. We begin with a general form of a suitable metric

$$m = p \left[\frac{\prod_{i=1}^{i_M} \vec{C}^T \hat{N}^{(i)} \vec{C}}{\prod_{j=1}^{j_M} \vec{C}^T \hat{D}^{(j)} \vec{C}} \right] \frac{(\vec{C}^T \vec{C})^{j_M}}{(\vec{C}^T \vec{C})^{i_M}}, \quad (\text{B1})$$

where $\hat{N}^{(i)}$ is the i -th Hermitian matrix of i_M whose expectation is on the numerator, and $\hat{D}^{(j)}$ is the j -th Hermitian matrix of j_M whose expectation is on the denominator, and p is a scalar prefactor. The factor on the right hand side is a normalisation so that the metric is independent of the overall scale of the mode coefficients.

For most of the cases in this manuscript, the geometry is symmetric under $z \rightarrow -z$, and we therefore assume that $|\epsilon_T|$ is unchanged upon the transformation $z \rightarrow -z$. In this case, the coefficients of \vec{C} all have the same phase, and thus we choose that all coefficients are real. We will first derive expressions for gradient-based optimisation under this assumption, because it is easy to then adapt those results to cases without this symmetry.

The two derivatives we require are, on the assumption that \vec{C} is a real vector,

$$\vec{\nabla} (\vec{C}^T \vec{C}) = 2\vec{C}, \quad \vec{\nabla} (\vec{C}^T \hat{M} \vec{C}) = 2 \text{Re}(\hat{M} \vec{C}), \quad (\text{B2})$$

where Re signifies the element-wise real part. We can then use the generic formula

$$\vec{\nabla}(m) = m \vec{\nabla}(\ln m), \quad (\text{B3})$$

to yield the gradient

$$\vec{\nabla}(m) = 2m \left(\sum_{i=1}^{i_M} \frac{\text{Re}(\hat{N}^{(i)} \vec{C})}{\langle \hat{N}^{(i)} \rangle} - \sum_{j=1}^{j_M} \frac{\text{Re}(\hat{D}^{(j)} \vec{C})}{\langle \hat{D}^{(j)} \rangle} + (j_M - i_M) \frac{\vec{C}}{|\vec{C}|} \right). \quad (\text{B4})$$

Now let us consider the more general case that \vec{C} has complex coefficients. We can turn this into a real vector \vec{C}^{Re} with twice the number of coefficients by defining

$$\vec{C}^{\text{Re}} = \begin{bmatrix} \text{Re}(\vec{C}) \\ \text{Im}(\vec{C}) \end{bmatrix}, \quad (\text{B5})$$

where Im signifies the imaginary part. We can now take any matrix \hat{A} in the basis of \vec{C} and convert it to a real matrix \hat{A}^{Re} such that all overlap values are preserved. This means that

$$\vec{C}'^T \hat{A} \vec{C} = (\vec{C}'^{\text{Re}})^T \hat{A}^{\text{Re}} \vec{C}^{\text{Re}}, \quad (\text{B6})$$

for all vectors \vec{C} and \vec{C}' . Equating the elementwise coefficients on this condition yields the method to construct \hat{A}^{Re} from \hat{A}

$$\hat{A}^{\text{Re}} = \begin{bmatrix} \hat{A} & i\hat{A} \\ -i\hat{A} & \hat{A} \end{bmatrix}, \quad (\text{B7})$$

This means that we can now use Eq. (B4) with vectors \vec{C}^{Re} and matrices \hat{A}^{Re} in order to optimise complex vectors (i.e. without restriction to axial reflection symmetry). This form is convenient because many numerical gradient-based optimisation routines naturally work with real vectors.

Appendix C: Justification of the use of retroreflective optimisation

As discussed in Sec. III, the conceptual validity of retroreflective optimisation requires that the target mode which we optimise is actually an eigenmode of the cavity when we construct the surface. In truth, the cavity eigenmode

ϵ_E is not quite the same as the target mode ϵ_T due to clipping of ϵ_T on the finite diameter mirror. In this case, retroreflective optimisation is not a strict optimisation, but provided that ϵ_E is similar to ϵ_T , it is likely that the retroreflective procedure will find a result that is close to optimal.

We check the agreement between ϵ_T and ϵ_E empirically in Fig. 9, which uses the scan data from Fig. 1. We see that (Fig. 9a) in the high performance region, where clipping losses are low, the predicted C_{int} and the eigenmode C_{int} are very similar, although in the low performance region (where the clipping losses are much higher), the performance values are quite distinct. We can understand this further by comparing the intensity (Fig. 9b) and round trip losses (Fig. 9c). The intensity comparison shows that ϵ_T and ϵ_E always have very similar intensities at the emitter, with differences at the part-per-thousand level, showing that ϵ_E is very similar to ϵ_T in overall spatial structure. However, the loss comparison indicates that the expected loss (calculated from a power clipping method) does not match the true loss (calculated through a mode mixing method). This loss mismatch is most severe when the clipping loss is not negligible, as it is specifically the clipping loss that is not calculated correctly in the retroreflective method. This is not so much of an issue in practice, because one would rarely be interested in cases where the clipping loss is not low, precisely because there would be high clipping loss. This means that the eigenmode C_{int} matches the expected C_{int} relatively well in the high-performance region; There are differences on the level of 10-20% in the reported value, but the optimisation could find improvements on the level of 200% for this dataset, so it is likely that the retroreflective optimisation procedure does indeed find most of the benefits available from mirror shaping.

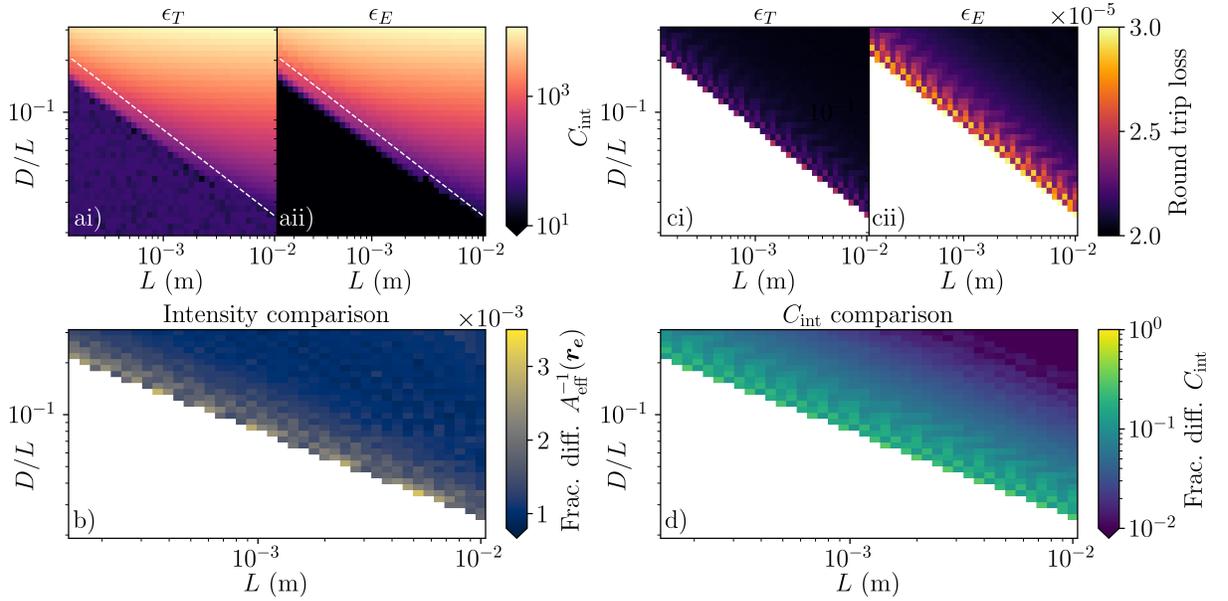


FIG. 9. Comparison of the target mode ϵ_T and cavity eigenmode ϵ_E for the dataset used in Fig. 1. a) The C_{int} of the ϵ_T and ϵ_E . The white dotted line indicates the D_{crit} . b) The fractional intensity difference at the emitter. c) The round trip loss of the target and eigenmodes. d) The fractional difference in C_{int} . In panels b)-d), data is only shown for cases where the mirror diameter exceeds the critical diameter. The nomenclature ‘Frac. diff.’ in colorbar labels is a shorthand for fractional difference. For a generic quantity q taking value q_T for the target mode and q_E for the eigenmode, the meaning of ‘Frac. diff. q ’ is $|q_T - q_E|/q_E$.

Appendix D: Derivation of matrices for the optimising photon extraction

As discussed in Sec. III of the main text, the retroreflective optimisation can proceed more smoothly if quantities within the evaluated metric can be calculated as expectation values of Hermitian matrices. For the case of photon extraction, our chosen metrics are the internal cooperativity C_{int} , and its single-transverse dimension equivalent $C_{\text{int}}/L_{y,\text{eff}}^{-1}$, which have expressions (taken from Eqs. (11) and (12) of the main text, but repeated here)

$$C_{\text{int}} = \frac{3\lambda^2 A_{\text{eff}}^{-1}(\mathbf{r}_e)}{\pi \mathcal{L}_{\text{int}}}, \quad C_{\text{int}}/L_{y,\text{eff}}^{-1} = \frac{3\lambda^2 L_{\text{eff}}^{-1}(\mathbf{r}_e)}{\pi \mathcal{L}_{\text{int}}}, \quad (\text{D1})$$

where λ is the wavelength, \mathcal{L}_{int} is the fractional round trip loss to non-transmissive sources, \mathbf{r}_e is the position of the emitter, and $A_{\text{eff}}^{-1}(\mathbf{r})$ ($L_{\text{eff}}^{-1}(\mathbf{r})$) is the normalised intensity per area (length) of the forward running wave at position \mathbf{r} ,

$$A_{\text{eff}}^{-1}(\mathbf{r}) \text{ or } L_{\text{eff}}^{-1}(\mathbf{r}) = \frac{I(\mathbf{r})}{\int_{\Pi_{z_r}} I(\mathbf{r}') d\Pi_{z_r}}, \quad (\text{D2})$$

where Π_{z_r} is used to unify the notation for the two cases by representing the infinite transverse domain (xy -plane or x -axis for two or one transverse dimensions respectively) that passes through \mathbf{r} , and $I(\mathbf{r})$ is the intensity of the electric field at position \mathbf{r} (per area for two transverse dimensions, and per length for a single transverse dimension). Of the quantities in Eq. (D1), the inverse effective area $A_{\text{eff}}^{-1}(\mathbf{r}_e)$, inverse effective length $L_{\text{eff}}^{-1}(\mathbf{r}_e)$, and internal loss \mathcal{L}_{int} depend upon the cavity mode, and we therefore need to find matrices $\hat{A}_{\text{eff}}^{-1}(\mathbf{r}_e)$ (or $\hat{L}_{\text{eff}}^{-1}(\mathbf{r}_e)$ for single transverse dimension), and $\hat{\mathcal{L}}_{\text{int}}$ such that

$$A_{\text{eff}}^{-1}(\mathbf{r}_e) = \frac{\vec{C}^T \hat{A}_{\text{eff}}^{-1}(\mathbf{r}_e) \vec{C}}{\vec{C}^T \vec{C}}, \quad L_{\text{eff}}^{-1}(\mathbf{r}_e) = \frac{\vec{C}^T \hat{L}_{\text{eff}}^{-1}(\mathbf{r}_e) \vec{C}}{\vec{C}^T \vec{C}}, \quad \mathcal{L}_{\text{int}} = \frac{\vec{C}^T \hat{\mathcal{L}}_{\text{int}} \vec{C}}{\vec{C}^T \vec{C}}, \quad (\text{D3})$$

where the normalisation factor is an acknowledgment that the quantities $A_{\text{eff}}^{-1}(\mathbf{r})$, $L_{\text{eff}}^{-1}(\mathbf{r})$, and \mathcal{L}_{int} are (by design) independent of the total intensity of the mode.

The first example is the inverse effective area matrix (used for the radially-symmetric situations that comprise the majority of the manuscript). Using the expression for electric field intensity

$$I(\mathbf{r}) = \frac{1}{2} c \epsilon_0 E(\mathbf{r})^2, \quad (\text{D4})$$

where $E(\mathbf{r})$ is the temporally constant electric field magnitude at position \mathbf{r} , and Eqs. (1) and (6) from the main text, yields

$$\begin{aligned} \hat{A}_{\text{eff}}^{-1}(\mathbf{r}) &= \frac{\sum_{s'=0}^S \sum_{s=0}^S (C_{s'})^* \left(\epsilon_{s'}^{LG,+}(\mathbf{r}) \right)^* \epsilon_s^{LG,+}(\mathbf{r}) C_s}{\int_{\Pi_{z_r}} \sum_{s'=0}^S \sum_{s=0}^S (C_{s'})^* \left(\epsilon_{s'}^{LG,+}(\mathbf{r}') \right)^* \epsilon_s^{LG,+}(\mathbf{r}') C_s d^2\mathbf{r}'}, \\ &= \frac{\vec{C}^T \hat{A}_{\text{eff}}^{-1}(\mathbf{r}) \vec{C}}{\vec{C}^T \vec{C}}, \\ \left(\hat{A}_{\text{eff}}^{-1}(\mathbf{r}) \right)_{s',s} &= \left(\epsilon_{s'}^{LG,+}(\mathbf{r}) \right)^* \epsilon_s^{LG,+}(\mathbf{r}), \end{aligned} \quad (\text{D5})$$

where the integrals in the denominator of the first line are simplified using the knowledge that $\{\epsilon_s^{LG,+}\}$ is an orthonormal basis in every transverse plane.

An exactly analogous derivation holds for $\hat{L}_{\text{eff}}^{-1}$ on the understanding that, consistent with Eqs. (1) and (3) in the main text and the conventions in this appendix, the dimensions of I and E change upon moving to a single transverse dimension. This means that we can use the more compact expression

$$\left(\hat{A}_{\text{eff}}^{-1}(\mathbf{r}) \right)_{s',s} \text{ or } \left(\hat{L}_{\text{eff}}^{-1}(\mathbf{r}) \right)_{s',s} = \left(\epsilon_{s'}^{\Lambda,+}(\mathbf{r}) \right)^* \epsilon_s^{\Lambda,+}(\mathbf{r}), \quad (\text{D6})$$

on the understanding that $\Lambda = \{LG, HG\}$ is chosen to match the transverse dimensionality of the scenario.

To calculate the internal cooperativity, an evaluation of the internal loss \mathcal{L}_{int} is also required. For this, it is important to clarify the bases of the propagating modes. Conventional mode mixing expresses the mirror at positive (negative) z as a matrix that scatters from the positively (negatively) propagating basis to the negatively (positively) propagating basis. Instead, we want to find our matrices so that they operate purely on the mode coefficients in the positively propagating basis, with the negatively propagating mode completely specified by the assumption of perfect retroreflection (this avoids complications in the Laguerre-Gauss basis where retroreflection swaps coefficients from $m \rightarrow -m$ because retroreflection preserves helicity not angular momentum). With this in mind, it is easiest to start with the mirror at positive z . Imagine replacing the mirror with an aperture of the same size, where light passing through the aperture sustains combined scattering and absorption losses $\mathcal{L}_{\text{abs}}/2$, to match the definition of \mathcal{L}_{abs} from the main text as the **round trip** non-clipping internal loss. The loss upon reflection from the mirror is the same as the loss of power upon this imagined transmission as the only further action of the mirror is to redirect power. The coefficients \vec{C}_{out} of the mode superposition representing this imagined transmission can be written in terms of the

input coefficients (\vec{C}) with

$$\left(\vec{C}_{\text{out}}\right)_{s'} = \sqrt{1 - \frac{1}{2}\mathcal{L}_{\text{abs}}} \sum_{s=0}^S \hat{\rho}_{s',s} \vec{C}_s, \quad \hat{\rho}_{s',s}^+ = \int_{\Pi_{m_+}} \left(\epsilon_{s'}^{\Lambda,+}(\mathbf{r})\right)^* \epsilon_s^{\Lambda,+}(\mathbf{r}) d\Pi, \quad (\text{D7})$$

where $\hat{\rho}^+$ is the clipping matrix of the aperture defined by the edge of the mirror at positive z (which is to what its + superscript refers), and Π_{m_+} is the finite domain of the mirror at positive z . The internal loss matrix of the mirror can then be derived

$$\mathcal{L}_{\text{int}}^+ = \frac{\left(\vec{C}\right)^\dagger \vec{C} - \left(\vec{C}_{\text{out}}^{L/R}\right)^\dagger \vec{C}_{\text{out}}^{L/R}}{\left(\vec{C}\right)^\dagger \vec{C}}, \quad \hat{\mathcal{L}}_{\text{int}}^+ = \hat{\mathcal{I}} - \left(1 - \frac{1}{2}\mathcal{L}_{\text{abs}}\right) \hat{\rho}^{\dagger+} \hat{\rho}^+, \quad (\text{D8})$$

For the mirror at negative z , we can repeat this process with the the negatively propagating mode ϵ^- to see that we would have

$$\left[\hat{\mathcal{L}}_{\text{int}}^-\right]_{\text{neg.basis}} = \hat{\mathcal{I}} - \left(1 - \frac{1}{2}\mathcal{L}_{\text{abs}}\right) \hat{\rho}^{-\dagger} \hat{\rho}^-, \quad \hat{\rho}_{s',s}^- = \int_{\Pi_{m_-}} \left(\epsilon_{s'}^{\Lambda,-}(\mathbf{r})\right)^* \epsilon_s^{\Lambda,-}(\mathbf{r}) d\Pi, \quad (\text{D9})$$

for an expression of $\hat{\mathcal{L}}_{\text{int}}^-$ in the negative propagating basis (whatever those basis states happen to be). To convert this to the positive propagating basis, we note that $\hat{\rho}$ is an integral of the normalised intensity matrix (i.e. inverse effective area or inverse effective length) over the surface of the mirror (see Eq. (D6)). Because reflection does not translate power transversely, within the domain of the mirrors we know that the intensity of the incoming mode and outgoing mode are the same everywhere. Therefore, over the mirror domain we can replace the intensity matrix of $\epsilon^{\Lambda,-}$ with the intensity matrix of $\epsilon^{\Lambda,+}$. Combining these results, we can write the internal loss matrices of both mirrors as

$$\hat{\mathcal{L}}_{\text{int}}^\pm = \hat{\mathcal{I}} - \left(1 - \frac{1}{2}\mathcal{L}_{\text{abs}}\right) \hat{\rho}^{\pm\dagger} \hat{\rho}^\pm, \quad \hat{\rho}_{s',s}^\pm = \int_{\Pi_{m_\pm}} \left(\epsilon_{s'}^{\Lambda,+}(\mathbf{r})\right)^* \epsilon_s^{\Lambda,+}(\mathbf{r}) d\Pi, \quad (\text{D10})$$

which is now a form that takes mode coefficients purely in the positive propagating basis. In the case that the mirror apertures have identical diameter and distance from the origin, and that the mode is restricted to be even symmetric with respect to z -reflection about the origin, the internal loss matrices for the two mirrors are identical. This can help reduce calculation time.

Appendix E: Derivation of critical diameter, separation, and best geometry for coupling to two emitters

The manuscript uses concepts of a critical diameter D_{crit} and critical separation s_{crit} . The critical diameter is the diameter of the mirror for which the clipping losses of the most compact possible Gaussian mode at the mirror equal the non-clipping losses, and is therefore a rough threshold for when a mirror has large enough diameter for the cavity to harbour a mode with negligible clipping loss. The critical separation is used for cases with axially-separated emitters, and it represents the threshold in emitter separation above which the divergence of the mode is limited by the requirement to stay focused on both emitters, rather than being limited by clipping on the mirrors. It turns out that, while deriving D_{crit} and s_{crit} , one can easily find the best geometry for coupling to two axial emitters, and therefore the derivation of all three are combined into a single appendix.

Assume that we have restricted the geometry of a concave-concave cavity to have length L and mirror diameter D . The identical mirrors both have a radius of curvature R , which we assume can be freely set.

A rigorous calculation of the expected clipping loss of the cavity is complicated, and should account for mode-mixing effects [41]. However, in the interests of obtaining analytical expressions that capture the overarching impact of the geometry, we can approximate the clipping loss through the power clipping method [28], which takes the mode of standard cavity theory (derived assuming infinite mirrors), and then calculates the fraction of the incident power lying outside the finite mirror on each reflection. For mirrors of diameter D , this leads to a **round trip** clipping loss

$$\mathcal{L}_{\text{clip}} = 2 \exp\left\{-2 \left(\frac{D}{2w_m}\right)^2\right\}, \quad (\text{E1})$$

where w_m is the waist of the mode at the mirror.

It is well known from standard cavity theory [14] that to minimise the size of the beam waist $w_P = w(z_P)$ at point P on the beam axis with axial coordinate z_P requires a cavity and cavity mode with the following properties

$$z_0 = z_P, \quad R = \frac{L}{2} \left(1 + \left(\frac{2z_P}{L} \right)^2 \right), \quad w_0 = \sqrt{\frac{\lambda z_P}{\pi}}, \quad w_P = \sqrt{\frac{2\lambda z_P}{\pi}}, \quad (\text{E2})$$

which will be useful for the rest of this appendix.

Starting with the derivation of D_{crit} , we are concerned with the minimum possible clipping loss for any Gaussian mode, given the constraints of L and D , which is obtained when w_m is minimised. We aim to find the critical diameter $D = D_{\text{crit}}$ such that the clipping loss $\mathcal{L}_{\text{clip}}$ and non-clipping loss \mathcal{L}_{abs} balance. This is because, if $\mathcal{L}_{\text{clip}}$ dominates, the total cavity losses are high, and the performance low for any spherical mirror, whereas if $\mathcal{L}_{\text{clip}}$ is negligible, it is possible to achieve loss at the minimum level of \mathcal{L}_{abs} , and thus achieve high cavity performance. The critical diameter achieving this balance is therefore

$$D_{\text{crit}} = \sqrt{\frac{\lambda L}{\pi}} \sqrt{2(\log(2) - \log(\mathcal{L}_{\text{abs}}))}, \quad (\text{E3})$$

obtained through rearrangement of Eq. (E1) with $w_m = \sqrt{\lambda L/\pi}$ as the minimum waist at mirror from Eq. (E2).

We now derive the critical separation s_{crit} , defined such that the mode that has the minimum waist size at both emitters separated by the critical separation s_{crit} , must have clipping loss on the mirror \mathcal{L}_{abs} . We know (Eq. (E2)) that the smallest beam waists in the emitter planes (separated by s), are found with central waist $w_0 = \sqrt{\lambda s/(2\pi)}$. This leads to a waist at the mirror of

$$w_m = \sqrt{\frac{\lambda(L^2 + s^2)}{2\pi s}} \approx \sqrt{\frac{\lambda L^2}{2\pi s}}, \quad (\text{E4})$$

where the approximation that $L^2 \gg s^2$ has been made in the right hand expression for algebraic convenience. This approximation is often reasonable, as the emitter separation could be far smaller than the cavity length, but not bigger. However, we would expect to see some discrepancies when $s \sim L$. Taking this assumption, substituting w_m into Eq. (E1) where the clipping loss should equal the non-clipping internal loss ($\mathcal{L}_{\text{clip}} = \mathcal{L}_{\text{abs}}$), and rearranging, yields

$$s_{\text{crit}} = \frac{\lambda L^2}{\pi D^2} \{ \log(2) - \log(\mathcal{L}_{\text{abs}}) \}. \quad (\text{E5})$$

Finally, we discuss the large mirror diameter D limit to C_{int} when coupling to two emitters separated by s . As mirror diameter is large, the clipping losses will be negligible, so our only concern is minimising the waist on both emitters. From Eq. (E2), we know the minimum waist at the emitters is $w_e = \sqrt{\lambda s/\pi}$.

The inverse effective area at the emitter for this mode is

$$A_{\text{eff}}^{-1} = \frac{2}{\pi} \frac{1}{w_e^2}. \quad (\text{E6})$$

Additionally, as the clipping losses are expected to be negligible, we expect $\mathcal{L}_{\text{int}} = \mathcal{L}_{\text{abs}}$. Applying Eq. (11) of the main text results in an internal cooperativity of

$$C_{\text{int}} = \frac{6\lambda}{\pi s \mathcal{L}_{\text{abs}}}. \quad (\text{E7})$$

Appendix F: Numerical parameters for simulations

The data presented in the main text compares retroreflective optimised cavities against the cavities with the best possible spherical/circular mirrors for the same geometric constraints. This procedure has three steps. The first is to scan the possible spherical/circular mirrors in order to find both the best performance with spherical mirrors, and the basis functions to use for retroreflective optimisation. These calculations were performed using mode mixing using N_{sph} basis states and numerical integration with s_{sph} samples in the Riemann sum (linearly spaced in either r or x coordinate). The second step is the first component of the retroreflective optimisation i.e. the determination of ϵ_T . This optimises the chosen metric on the assumption of perfect retroreflection, where the chosen metric is evaluated in a basis of size N , of which we choose to optimise the coefficients of the first N_{opt} basis states, and leave the remainder

zero. The third step is to take the optimised ϵ_T , and use the surface reconstruction method (Sec. III main text), where the surface is initially sampled (again, linearly in r or x) using s small elements, and mode mixing is performed in a basis of size N . A summary of the parameters used for the simulations can be found in Table I. From this, we highlight that the optimisation for Fig. 3, where the emitter was displaced towards the edge of the cavity, required a very large number of samples for its mode mixing integrals to converge, but such sampling was not necessary for cases with the emitter closer to the cavity centre.

Example	$N_{\text{sph.}}$	$s_{\text{sph.}}$	N	$N_{\text{opt.}}$	s
Fig. 1	171	15000	171	136	15000
Fig. 2	171	15000	171	136	15000
Fig. 3	50	300	250	200	500000
Fig. 5 examples	171	15000	171	136	15000
Fig. 5 scan	171	5000	171	136	5000
Fig. 6 examples	171	15000	171	136	15000
Fig. 6 scan	111	5000	111	81	5000

TABLE I. Parameters used for the retroreflective optimisations conducted in the main text. The columns indicate the parameter, and the rows discuss the figure, or part of the figure, where the relevant optimisation is used.

The second part of this procedure (i.e. the first part of the retroreflective optimisation method) optimises ϵ_T by changing the coefficients of $N_{\text{opt.}}$ states in a basis of size N . The number of coefficients ultimately corresponds to the number of degrees of freedom in the target mode, and consequently the resultant mirror surfaces. If $N_{\text{opt.}}$ is too low, there will be not enough freedom to describe the optimum ϵ_T , and thus the method will not report the optimum performance. Data indicating how the retroreflective performance (i.e. the cavity performance assuming perfect retroreflection) varies with $N_{\text{opt.}}$ and N are shown in Fig. 10. Here we see that, for most examples, ≤ 50 basis states and optimised coefficients are sufficient to find the optimised performance to a high degree of accuracy.

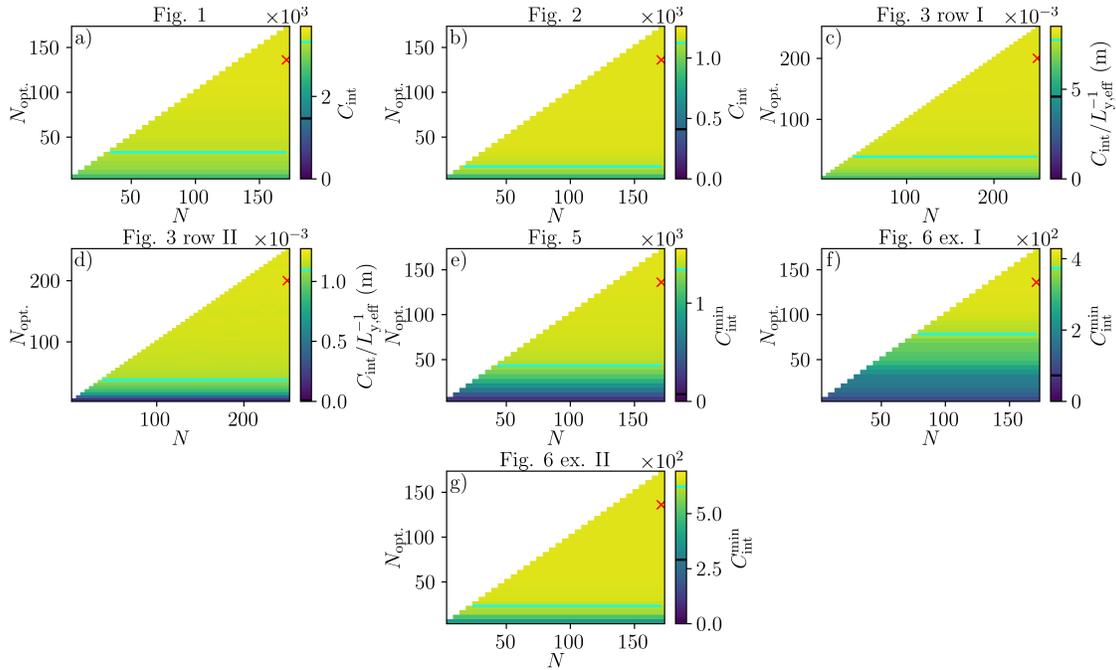


FIG. 10. The dependence of the cavity performance assuming perfect retroreflection on the basis size N and number of variable coefficients $N_{\text{opt.}}$, where $N_{\text{opt.}} \leq N$. Note that, because this studies the convergence of the first part of retroreflective optimisation only, no mode mixing calculation is conducted. Each panel shows the convergence data for the titled figure. The black band on the colour bar indicates the retroreflected expected performance for the best spherical mirror. The cyan band in the colour bar and as a contour on the performance map indicates the boundary above which the reported performance has increased by 90% of the total increase from the spherical case to the optimised value reported for the largest basis sizes in that panel's data. The red cross indicates the parameters chosen for the case study examples in each figure. Note that the black band in panel d) is very close to zero, because the best spherical mirror cavity produced a highly mixed mode, where the nominal fundamental mode was strongly clipped.