

# Optimization of Scalable Ion-Cavity Interfaces for Quantum Photonic Networks

Shaobo Gao<sup>1</sup>, Jacob A. Blackmore<sup>1,\*</sup>, William J. Hughes<sup>1</sup>, Thomas H. Doherty<sup>1</sup>, and Joseph F. Goodwin<sup>1,†</sup>  
*Department of Physics, University of Oxford, Clarendon Laboratory, Parks Road, Oxford OX1 3PU, United Kingdom*

 (Received 10 December 2021; revised 25 August 2022; accepted 21 November 2022; published 11 January 2023)

In the design optimization of ion-cavity interfaces for quantum networking applications, difficulties occur due to the many competing figures of merit and highly interdependent design constraints, many of which present “soft limits,” which are amenable to improvement at the cost of engineering time. In this work, we present a systematic approach to this problem that offers a means to identify efficient and robust operating regimes and to elucidate the trade-offs involved in the design process, allowing engineering efforts to be focused on the most sensitive and critical parameters. We show that in many relevant cases it is possible to approximately separate the geometric aspects of the cooperativity from those associated with the atomic system and the mirror surfaces themselves, greatly simplifying the optimization procedure. Although our approach to optimization can be applied to most operating regimes, here we consider cavities suitable for typical ion-trapping experiments and with substantial transverse misalignment of the mirrors. We find that cavities with mirror misalignments of many micrometers can still offer very high photon extraction efficiencies, offering an appealing route to the scalable production of ion-cavity interfaces for large-scale quantum networks.

DOI: [10.1103/PhysRevApplied.19.014033](https://doi.org/10.1103/PhysRevApplied.19.014033)

## I. INTRODUCTION

There are many situations in which the interaction of atoms and light may be usefully augmented by placing the atomic system within an optical cavity, ranging from highly efficient optical-lattice generation [1,2] to the creation of spin squeezed states for improved metrology [3–5], as well as a wealth of quantum information processing applications [6–10]. Among these, one application where cavities seem poised to play a critical role is in the efficient extraction of single photons from trapped atoms or ions. Such systems will likely be central to matter-light interface design in quantum networks [11–15] and, at a more basic level, can also be used as extremely high quality sources of indistinguishable single photons [11,16–23].

Despite a decade of significant advances [24–27] in the use of cavities for network applications, the plethora of varied designs seen in today’s experiments [28] suggests an interesting question: Given a few simple experimental constraints, is it possible to identify an optimal atom-cavity system for quantum networking; and what defines “optimal,” given the many competing figures of merit? Fidelity, attempt rate, and success probability are clearly important considerations but for practical systems, simplicity

of construction and stability are also paramount. Even if a figure of merit can be identified, the problem is complicated by the many different design parameters and a variety of often experiment-specific parameter constraints and manufacturing tolerances. It is to our benefit that few of these constraints are absolute but such “soft limits” make simple many-parameter optimization a poor substitute for a real understanding of where the easiest routes to improvement lie.

Identifying the dominant drivers of performance in atom-cavity systems is crucial if such devices are to be produced at scale for the realization of quantum networks of useful size. The requirement that a large number of such devices work simultaneously requires each to be extremely reliable, a phrase rarely heard in relation to today’s cavity-networking experiments. Achieving this means, on the one hand, identifying parameter regimes where robust performance may be realized and, on the other, minimizing the experimental and engineering complexity of the systems used.

In this paper, we set out a systematic approach to addressing the first of these challenges. We focus particularly on the use of such systems in the rapid generation of entanglement between remote pairs of atomic ions via a measurement-based entanglement-swapping scheme, a key resource for distributed quantum computation [29–31]. While many of the ideas and approaches presented may be readily applied to other use cases and types of quantum

\*Corresponding author. [jacob.blackmore@physics.ox.ac.uk](mailto:jacob.blackmore@physics.ox.ac.uk)

†[joseph.goodwin@physics.ox.ac.uk](mailto:joseph.goodwin@physics.ox.ac.uk)

emitter (e.g., neutral atoms), care should be taken to ensure that the assumptions made still apply. Most notably, many trapped-atom applications require very low line-width cavities to ensure sufficient resonant selection of the desired atomic transition, meaning that it may not be possible to optimize the cavity geometry independently of the mirror coating, as we do in this work.

We start in Sec. II by providing a detailed definition of the problem, describing the photon-production scheme, the cavity geometry, and the “real-world” limits and tolerances of the system. In Sec. III, we then consider performance metrics for ion-cavity systems in the context of quantum networking. Despite the complexity of the general case, we show that with appropriate choice of atomic levels and driving scheme, a much simpler performance optimization is sufficient, namely, maximizing the probability of extracting a single photon from the cavity. In Sec. IV, we show that the optimization of extraction probability is well approximated as a separable problem, with the atomic, geometric, and mirror-interface parameters considered in turn. We then show how this can be used to identify optimal geometries and mirror coatings given limits on mirror misalignment or scattering loss. In Sec. V, we use this separability to study the impact of realistic constraints on the optimization and how the performance of the cavity is impacted by errors in manufacture or in estimation of the magnitude of losses and misalignment. We conclude in Sec. VI with a summary of the operating regimes where performance is achievable with the lowest engineering demands. Most notably, we find that (with some caveats) substantial transverse misalignments of the cavity mirrors are indeed tolerable, underlining the potential for simple passively aligned cavity assemblies suitable for scalable quantum networks.

## II. DESCRIPTION OF THE PROBLEM

### A. Photon-production scheme

There are several possible ways to extract a photon from an atomic system into an optical cavity mode [20,32]. We choose to assume the use of a vacuum-mediated stimulated Raman adiabatic passage (vSTIRAP) scheme [11,33–35], a simplified illustration of the level structure of which is given in Fig. 1. In this scheme, the ion is initialized in level  $|0\rangle$  and transferred to state  $|1\rangle$  via emission of a single photon near the  $|e\rangle \rightarrow |1\rangle$  transition frequency. States  $|0\rangle$  and  $|1\rangle$  are generally stable or metastable and not directly connected by an electric dipole (E1) transition; instead, each one is connected by E1 transitions to an excited state  $|e\rangle$ , creating a  $\Lambda$ -type system. The transition  $|0\rangle \rightarrow |e\rangle$  is driven by an external laser source with Rabi frequency  $\Omega$ ; we label this arm of the transition the “drive.” The other arm is the “cavity” transition from  $|e\rangle \rightarrow |1\rangle$ , driven by the cavity vacuum field with a strength quantified via the ion-cavity coupling  $g$ . For efficient transfer, both arms of

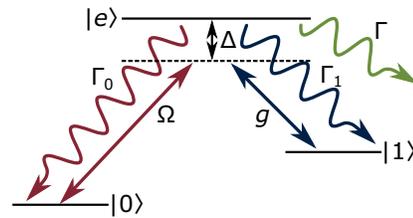


FIG. 1. A simplified energy-level structure for the atomic part of our system. A pair of (meta)stable energy levels,  $|0\rangle$  and  $|1\rangle$ , are connected by a pair of dipole-allowed optical-frequency transitions via a common excited state  $|e\rangle$ . The  $|0\rangle \rightarrow |e\rangle$  transition is driven by an external laser field, with Rabi frequency  $\Omega$ . The transition  $|e\rangle \rightarrow |1\rangle$  is driven by the vacuum mode of the cavity, with ion-cavity coupling strength  $g$ . The laser and cavity fields may be tuned away from resonance with the bare atomic level by an equal amount  $\Delta$  to reduce the population of  $|e\rangle$ . From  $|e\rangle$ , we assume that decays are possible to  $|0\rangle$  with rate  $\Gamma_0$ , to  $|1\rangle$  with rate  $\Gamma_1$ , and to possible other levels (not shown) to give a total upper-state decay rate of  $\Gamma$ . Transfer of population to  $|1\rangle$  is associated with emission of a single photon of frequency  $\omega_{e1} - \Delta$ , which must occur into the cavity mode to be useful for networking.

the transition must have equal detuning  $\Delta$  such that the Raman condition is met [35]. This family of schemes has many advantages over simple two-level  $\pi$ -pulse schemes: control of the photon wave-packet shape [16,36]; compatibility with almost any dipole-allowed  $\Lambda$  system in the atom, enabling a wide range of network photon wavelengths; compatibility with bichromatic driving schemes [37]; and maximization of the photon extraction probability, provided that one is unconcerned with the length of the photon (see the discussion of this point in Sec. III B and Appendix B).

Measurement-based remote-entanglement schemes must first produce a bipartite entanglement between an atomic qubit and a photonic qubit. To encode and populate these qubits, we require that  $|1\rangle$  represents a manifold of several distinct sublevels; typically, this is also the case for  $|0\rangle$  and  $|e\rangle$ . Many protocols are possible [38] but each involves driving two discrete  $\Lambda$  systems. These may or may not share an initial state and may be driven simultaneously with a bichromatic drive [31,37,39] (for polarization-encoded schemes) or individually in sequence [40] (for time-bin encoded schemes).

The presence of the more complex electronic structure necessary to produce spin-photon entanglement introduces extra conditions that must be met to avoid fidelity loss due to off-resonant transfers and/or spectator decay outside the qubit subspace [41] (see Sec. III). However, provided that these are satisfied, the performance of each of the pairs of  $\Lambda$  systems is well described by the model given in Fig. 1. For the purposes of this paper, we optimize this simple

system only, with the results applicable—with appropriate caveats—to a diverse range of remote-entanglement schemes.

### B. The cavity

Having defined the photon-production scheme within the electronic structure of the emitter, we now describe the properties of the Fabry-Perot cavity used to provide the coupling,  $g$ . We consider the symmetric cavity geometry shown in Fig. 2, consisting of two concave spherical mirrors with equal radii of curvature  $R$  and equal diameter  $D$ . In the ideal system, the mirrors are coaxial and separated by  $L$  and the ion sits on the cavity axis at  $L/2$ . The volume of the mirrors, modeled as a spherical cap, is given by

$$V = \frac{\pi h}{6} \left( \frac{3D^2}{4} + h^2 \right), \quad (1)$$

where the saggita,  $h$ , is defined as  $R - \sqrt{R^2 - D^2/4}$ . Within this framework, the fundamental lower limit to  $L$  is  $2h$ , corresponding to two mirrors that touch when not misaligned.

We consider coupling to the fundamental transverse electric  $TE_{00}$  mode of the cavity and neglect higher-order components, which can be ensured by tuning the fundamental mode to Raman resonance and choosing a suitable  $L/R$  ratio to break the transverse-mode degeneracy. Our mode then forms a  $TE_{00}$  Gaussian mode with waist  $w_0$  at the location of the ion, with a far-field divergence half-angle  $\theta$ , where

$$\theta = \frac{\lambda}{\pi w_0}, \quad (2a)$$

$$w_0 = \sqrt{\frac{\lambda}{2\pi} \left[ L(2R - L) \right]^{1/4}}, \quad (2b)$$

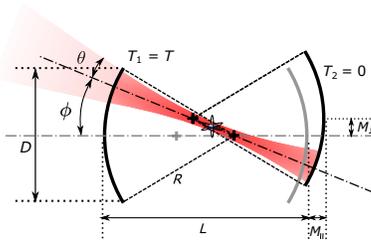


FIG. 2. The geometry describing a misaligned optical cavity. The cavity consists of two concave spherical mirrors with equal radii of curvature ( $R$ ) and supports a mode with divergence half-angle  $\theta$ . In the optimum configuration (shown in gray), the overall length of the cavity is  $L$  and the two mirrors are perfectly aligned to the cavity axis. In the misaligned case, one mirror is displaced along the cavity axis by  $M_{\parallel}$  and transversely by  $M_{\perp}$ ; we assume for simplicity that  $M_{\parallel} = M_{\perp} = M$ . This misalignment tilts the overall mode by an angle  $\phi$  from the cavity axis and slightly increases  $\theta$ .

and  $\lambda$  is the wavelength of the mode. At the center of the cavity, an emitter located at an antinode of the cavity mode couples to the cavity with a rate  $g_0$  inversely proportional to the square root of the mode volume  $V = \pi w_0^2 L/4$ :

$$g_0 = \sqrt{\frac{3\lambda^2 c \gamma_1}{\pi^2 L w_0^2}} = \sqrt{\frac{3\lambda^2 c \gamma_1}{4\pi V}}, \quad (3)$$

where  $c$  is the speed of light in vacuum and  $\gamma_1 = \Gamma_1/2$  is the partial (half-) line width of the atomic decay channel coupled to the cavity. Note that if the vSTIRAP transfer is between sublevels of the initial, final, and excited states, this represents the partial line width of the specific transition, including angular-momentum considerations. Furthermore, we assume that the cavity and magnetic field are arranged such that the cavity vacuum mode couples optimally to the transition in question; where this is not the case [42],  $g_0$  is suppressed accordingly.

Residual translational misalignment of the cavity mirrors can occur in the transverse and longitudinal directions, with magnitudes  $M_{\perp}$  and  $M_{\parallel}$ , respectively. For cavity configurations far from concentricity,  $M_{\parallel}$  has negligible effect, while  $M_{\perp}$  primarily introduces a tilt  $\phi$  to the cavity mode axis, measured relative to the unperturbed value. However, for near-concentric cavities,  $M_{\parallel}$  significantly changes the sensitivity of the system to transverse misalignment and, furthermore, leads to substantial changes in  $\theta$ :

$$\phi = \arctan \left( \frac{M_{\perp}}{[2R - L] - M_{\parallel}} \right) \approx \frac{M_{\perp}}{[2R - L] - M_{\parallel}}, \quad (4a)$$

$$L \rightarrow L' = 2R - \left( \frac{2R - L - M_{\parallel}}{\cos \phi} \right) \approx L + M_{\parallel}, \quad (4b)$$

$$\theta \rightarrow \theta' = \left[ \frac{2\lambda}{\pi \sqrt{L'(2R - L')}} \right]^{1/2}. \quad (4c)$$

These distinct effects mean that the two forms of misalignment should, in general, be considered independently; to simplify our optimization, we choose to define a single misalignment parameter  $M = M_{\parallel} = M_{\perp}$ . This combination of the two parameters represents a worst-case scenario, where increasing axial misalignment moves the cavity closer to concentricity, increasing the sensitivity to transverse misalignment. This is itself compounded by radial misalignments increasing  $\phi$ . The action of both is to increase the clipping losses, as shown in Fig. 3. Note that for the purposes of this work, we assume that the ion position is appropriately adjusted to remain aligned to the mode center when this shifts due to mirror misalignment.

To ensure that photons are extracted from the cavity into a single free-space mode, one mirror is coated for maximum reflection ( $T_2 \approx 0$ ), while the other “outcoupler” mirror has a finite transmission  $T_1 = T$ . The value of  $T$

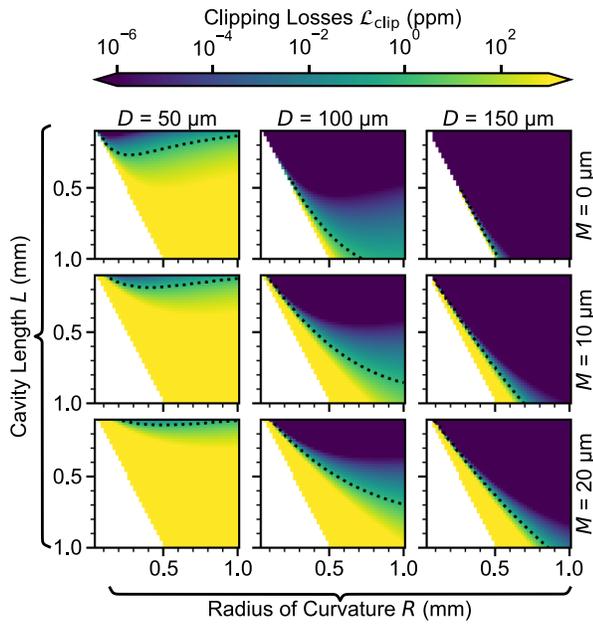


FIG. 3. The dependence of round-trip clipping losses, due to cavity misalignment, on the cavity geometry. Each panel shows the value of  $\mathcal{L}_{\text{clip}}$  as a function of the radius of curvature of the mirrors and the cavity length for a given combination of mirror diameter ( $D$ ; columns) and misalignment ( $M$ ; rows). The white region indicates where the cavity does not support stable modes or where the cavity cannot be assembled due to geometric constraints. Misalignment of the mirrors causes the mode to tilt at an angle  $\phi$ , with the effect being more pronounced for cavities approaching the concentric condition ( $R = L/2$ ). For small  $\phi$ , the mode remains well within the spherical region and losses are minimal, but as the mode approaches the edge of the mirror, these rapidly increase. This leads to a clipping-loss “threshold,” which is particularly sharp for large and well-aligned mirrors. In all panels, we show the 1 ppm clipping contour as a black dotted line.

is chosen to optimize the performance of the system (see Ref. IV C) and is typically several hundred parts per million (ppm). Considering these transmissive losses along with the *intrinsic* round-trip losses  $\mathcal{L}_{\text{in}}$  due to absorption, scattering, or clipping of the cavity mode, the total decay rate of the cavity field amplitude  $\kappa$  and the finesse  $\mathcal{F}$  are thus given by

$$\kappa = \frac{c}{4L} (T + \mathcal{L}_{\text{in}}), \quad (5a)$$

$$\mathcal{F} = \frac{2\pi}{T + \mathcal{L}_{\text{in}}}. \quad (5b)$$

We consider two distinct contributions to  $\mathcal{L}_{\text{in}}$ , assuming that other losses, such as absorption or scattering in the region between the cavity mirrors, are negligible for short cavities under ultra-high vacuum conditions:

(1)  $\mathcal{L}_{\text{clip}}$  describes losses due to clipping of the edge of the mode by the finite diameter of the cavity mirror [43].

(2)  $\mathcal{L}_{\text{scat}}$  describes scattering losses due to rms mirror surface roughness ( $\sigma$ ), typically approximated [44] by

$$\mathcal{L}_{\text{scat}} \approx 2 \left[ 1 - e^{-(4\pi\sigma/\lambda)^2} \right]. \quad (6)$$

As we show in Sec. IV, the assumption that these two loss mechanisms are independent underpins our approach to optimization.

### C. Practical considerations and assumptions

Although the methods described in this paper are reasonably general and may easily be adapted to most emitter-cavity systems, we make a number of assumptions regarding the geometry of the cavity and mirrors motivated by our own experiments, which we outline here.

The cavity is assumed to be arranged around the center of an ion trap and as such is limited in its minimum length and maximum mirror diameter, due to the need to minimize exposure of the ion to the mirror dielectric, which can disrupt the trapping potential [28]. The ion is assumed to be located at the center of the cavity mode and perfectly localized.

The mirrors are assumed to exhibit significant relative transverse misalignment—although the treatment holds for more precisely aligned mirrors, the range of  $M$  we consider is more representative of the offsets that might occur in simple passively aligned systems. Meanwhile, angular misalignment of the mirrors is considered to be insignificant and the length of the cavity is assumed to be stable and set at an appropriate value for the necessary cavity-Raman resonance, achieved in our system by active closed-loop control of this degree of freedom.

The mirrors themselves are assumed for simplicity to be perfectly spherical, with uniform scattering loss across the surface. This is a good approximation for some fabrication methods (mechanical or focused ion-beam milling) but traditional laser-ablation methods typically yield Gaussian mirror profiles and even modern dot-milling ablation techniques lead to significant deviations from a spherical profile away from the mirror center [45]. Empirically, however, “clipping”-type losses still grow rapidly with increasing  $\phi$  or  $\theta$  and to a suitable approximation [see Eq. (15) and associated discussion], Gaussian and other nonspherical mirrors can be well modeled by a hard-edged spherical profile [46].

For the purposes of the analysis presented in this paper, we consider the machining of mirrors to be volume limited, i.e., there is a cost associated with removing a certain volume of the mirror substrate, be it in one or more of milling time, equipment limitations, surface-quality degradation, etc. Bounded by this limit, we assume that during

the manufacture one can freely vary  $R$  and  $D$  to maximize performance.

### III. PERFORMANCE METRICS

The overriding aim of a cavity-mediated quantum network is to produce entanglement between pairs of remote nodes at a high fidelity and rate. Currently, the state-of-the-art fidelity is for ion-ion remote entanglement at 96.0(1)%, which has been achieved at a rate of  $100 \text{ s}^{-1}$ , collecting photons with a high-numerical-aperture lens [47]. Systems for practical future quantum networks will require substantially higher rates, and similar or higher fidelities, to ensure that entanglement between nodes can be efficiently distilled to fidelities approaching those of local gates [48]. However, how these two key metrics should be measured against one another is a nontrivial question and is highly architecture- and application-dependent. Furthermore, many factors contribute to the performance, including purely atomic processes such as state preparation and mapping, and purely photonic considerations such as imperfections in the photonic Bell-state analyzer. Given these concerns, it is reasonable to question whether a universal approach to network cavity-design optimization exists.

In this section, we consider the impact on remote-entanglement fidelity and rate due to properties of the ion-cavity system itself and demonstrate that for suitable choices of atomic system and driving scheme, it is possible to identify a single metric of performance, independent of the finer details of the network architecture.

#### A. Fidelity

We aim to achieve high-fidelity entanglement between pairs of ions located in remote nodes via a two-photon entanglement-swapping protocol [12,31]. This requires the production of single photons from each node, with one photonic degree of freedom (e.g., polarization or time-bin) entangled with high fidelity to the final electronic state of the atom. In order for the path information of the two photons to be erased and the entanglement swapped from ion-photon to ion-ion, each photon must also be highly indistinguishable. We now consider three mechanisms that can reduce the fidelity of ion-photon entanglement or the indistinguishability of the photons produced.

First, we note that the fidelity of the ion-photon state may be reduced due to off-resonant coupling during the vSTIRAP transfer. As discussed in Sec. II A, transfer to an (entangled) electronic qubit superposition state requires two  $\Lambda$  systems to be driven and unless these two processes are made sufficiently distinct via frequency and/or polarization selection, the ion and photon will only be partially entangled. Additionally, off-resonant decays to other sublevels beyond the electronic qubit subspace introduce problematic “loss” errors (for more details, see

Refs. [41,49]). Avoiding these effects generally demands that the qubit levels are well resolved, i.e., the final qubit splitting is much larger than  $\kappa$ . This is an important consideration but such errors can typically be made negligible via suitable choice of atomic system and bias field.

To produce indistinguishable photons, each node must emit well-defined and identical wave packets. Interestingly, provided that the photon detectors are time resolving and with sufficiently low jitter, it is not necessary for the photons to be at identical frequencies [50] but only for any difference to be stable and well characterized. However, any mismatch in the wave-packet envelopes between nodes provides which-way information during the entanglement-swapping step, reducing the final fidelity. The use of vSTIRAP schemes permits the shape of the wave-packet envelope to be tuned via suitable modulation of the drive pulse, allowing such effects to be suppressed. A more significant problem occurs in the event of spontaneous scattering back to the initial state during the Raman transfer, the stochastic nature of which leaves the photonic wave packet in a mixed state. We give an overview of this problem in Appendix A, with detailed discussions also found in Refs. [23,41]. While vSTIRAP pulse shaping can help reduce such spontaneous scattering, it is also crucial to bias the system toward transfer into the target state. One route to achieving this is to use a very high ion-cavity coupling  $g$  but this places unnecessary demands on the precision of mirror alignment. A better approach is to simply use a  $\Lambda$  system with a low branching ratio ( $\alpha$ ) on the “drive” arm, which keeps fidelity losses due to spontaneous scattering  $< 1\%$  across the parameter ranges considered here [41], and often far lower, permitting us to conclude our consideration of this issue.

Finally, we note that birefringence in the cavity-mirror substrates or dielectric coatings can cause the polarization of the intracavity field to rotate upon each reflection, introducing a time dependence to the output-photon polarization. In polarization-encoded schemes, although the fidelity of the ion-photon entanglement is only modestly reduced, the phase of this entangled state becomes time dependent, which can significantly impact the achievable fidelity of the final ion-ion entangled state. We have recently studied the impact of such effects [51], concluding that without excessively complex mitigation strategies, maintaining a fidelity loss  $< 1\%$  requires birefringence  $\Omega_B \ll \kappa/10$ . This places considerable demands on the quality of mirrors used and a more suitable approach may be to use time-bin encoded photonic qubits, a technique that also prevents fiber birefringence from limiting fidelities over longer link distances. For the purposes of this work, we assume that the information is not encoded in the polarization degree of freedom and so the impact of any birefringence is negligible.

In summary, although these mechanisms should be considered carefully in the design of a photonic network,

each is best mitigated through appropriate choice of the atomic emitter, driving scheme, and photonic encoding. With these choices made, the impacts on fidelity due to the cavity itself are expected to be lower than those due to the photonic network, entanglement-swapping process, and local operations on the atom and we can consider only the maximization of the photon-production rate in our optimization process.

### B. Rate

For a two-photon entanglement scheme [12,31], the rate at which we produce remote ion-ion Bell pairs is given by

$$R_{\text{Bell}} = \frac{1}{2} \times \frac{(P_{\text{ext}}\epsilon_{\text{net}}\epsilon_{\text{det}})^2}{\tau_{\text{att}}}, \quad (7)$$

where  $P_{\text{ext}}$  is the probability of single-photon extraction from the cavity at each node,  $\epsilon_{\text{net}}$  and  $\epsilon_{\text{det}}$  are the photon network transmission and the detector efficiency,  $\tau_{\text{att}}$  is the attempt cycle period, and the factor of 1/2 accounts for anticoincidence events and is intrinsic to the protocol [12,31]. Note that due to the two-photon herald, the success rate is quadratically sensitive to the probability of single-photon detection from each node. The network and detector efficiency terms depend on physical components of the system and lie beyond the scope of this paper but we note for completeness that the wavelength of the network photon dictates the choice of materials and is thus a critical determinant of the performance. Optimizing an ion-cavity system for a high entanglement rate thus depends on maximizing the attempt rate  $1/\tau_{\text{att}}$  and, most crucially, the photon extraction probability  $P_{\text{ext}}$ .

We first consider contributions to the attempt rate, considering contributions to the attempt cycle period:

$$\tau_{\text{att}} \sim \tau_{\text{prep}} + \tau_{\text{lat}} + \tau_{\gamma} + \frac{\Delta x_{\text{net}}}{c}. \quad (8)$$

Here,  $\tau_{\text{prep}}$  is the time needed to optically pump and transfer population into the initial state,  $\tau_{\text{lat}}$  is the sum of all optical and electronic latencies in the attempt loop, and  $\Delta x_{\text{net}}/c$  is the combined quantum and classical signal-propagation time for two nodes separated by optical length  $\Delta x_{\text{net}}$ . The sum of these times limits the attempt rate to approximately 1 MHz and so by requiring the characteristic length of the photon  $\tau_{\gamma} \ll 1 \mu\text{s}$ , we can ensure that the duration of photon production is insignificant. Most fundamentally, the rate at which a photon can be extracted from the cavity is limited by  $\kappa$  but rapid production of a photon also demands a reasonable  $g$  and a sufficiently strong Raman drive. However, driving the system too rapidly breaks the adiabaticity required for efficient population transfer, so some care must be taken in the choice of transfer pulse. In Appendix B, we consider the effects

of nonadiabatic driving and show that for all the operating regimes outlined in this paper, high-probability Raman transfer can be achieved while ensuring that the photon length makes an insignificant contribution to the attempt cycle period.

With appropriate care taken in the choice of atomic system, entanglement protocol, and Raman transfer pulse, we thus identify a single dominant figure of merit for our ion-cavity system, namely, the probability of photon extraction per attempt cycle,  $P_{\text{ext}}$ .

To calculate  $P_{\text{ext}}$ , we first consider the probability that a photon is generated in the cavity,  $P_{\text{gen}}$ . Accounting for losses due to spontaneous emission, we can place an upper bound on  $P_{\text{gen}}$ , which saturates in the case of perfect adiabatic transfer:

$$P_{\text{gen}} \leq \frac{2C}{2C+1}, \quad (9a)$$

$$C = \frac{g^2}{2\gamma\kappa}. \quad (9b)$$

Here, we define the cooperativity  $C$ , which is a function of  $g$ , and  $2\gamma = \Gamma$ , the total spontaneous decay rate of the excited atomic state. The overall extraction probability is then the product of  $P_{\text{gen}}$  and the probability that the photon exits the outcoupler mirror before being absorbed or scattered out of the cavity mode:

$$P_{\text{ext}} = P_{\text{gen}} \left( \frac{T}{T + \mathcal{L}_{\text{in}}} \right) \leq \frac{2C}{2C+1} \left( \frac{T}{T + \mathcal{L}_{\text{in}}} \right). \quad (10)$$

Following Ref. [52],  $P_{\text{ext}}$  can be shown to be bounded by

$$P_{\text{ext}} \leq 1 - \frac{2}{1 + \sqrt{1 + 2C_{\text{in}}}}, \quad (11)$$

in which the *intrinsic* cooperativity  $C_{\text{in}}$  is defined as

$$C_{\text{in}} = \frac{g_0^2}{2\gamma\kappa_{\text{in}}}, \quad (12)$$

where  $g_0$  is the value of  $g$  at the central cavity antinode. The intrinsic loss rate within the cavity is defined as

$$\kappa_{\text{in}} = \frac{\mathcal{L}_{\text{in}}c}{4L}. \quad (13)$$

The bound in Eq. (11) is only saturated when the overall cavity decay,  $\kappa$ , is appropriately optimized by considering the optical properties of the cavity, as is discussed further in Sec. IV C. As maximizing finesse is not generally important, we assume that it is always possible to achieve this optimum  $\kappa$ , although this does require an accurate estimate of the intrinsic loss prior to coating the mirrors.

Note that in this treatment, we assume that any photon populating the free-space mode leaking from the cavity

outcoupler can be coupled into the photonic network with constant high efficiency. This is a reasonable assumption for a typical  $TE_{00}$  cavity mode imaged appropriately into a single mode fiber but for systems with large tilt  $\phi$  and divergence  $\theta$ , this assumption may break down for certain cavity-mirror substrate geometries.

#### IV. CAVITY-DESIGN OPTIMIZATION

In the following three subsections, we discuss how the geometric and optical properties of the cavity can be optimized to maximize the photon extraction probability, for cavity geometries similar to those in Fig. 2. In Sec. IV A, we discuss the parameters considered in the optimization of the performance and our strategy for optimization. In Sec. IV B, we show how consideration of clipping losses defines the optimal cavity geometry and then use this result, along with the known mirror scattering loss, to optimize outcoupler transmission (and hence cavity finesse) in Sec. IV C. In Sec. IV D, we combine these results to understand how  $P_{\text{ext}}$  is optimized through the choice of  $L$ ,  $R$ , and  $T$  and offer some “rules of thumb” for identifying optimal operating regimes in the presence of certain practical constraints.

##### A. Optimization scheme

In this subsection, we briefly outline our approach to design optimization. In Sec. III, we show, for certain common choices of driving and entanglement schemes, that the optimization of the ion-cavity system is reduced to a maximization of the photon extraction probability. As shown in Sec. III B,  $P_{\text{ext}}$  is less than  $P_{\text{gen}}$  due to intrinsic losses in the cavity and quantifying and minimizing these losses is central to our optimization scheme.

The ability to split the sources of loss into independent constituents [see Eq. (6)] enables the use of a two-part optimization strategy, which we describe in the following three sections. The first stage is to optimize the geometry by reducing the optimization scheme to a maximization of  $g$  within well-defined bounds, where  $\mathcal{L}_{\text{clip}} \ll \mathcal{L}_{\text{scat}}$ . The second stage is to then optimize  $T$  given the fixed value of  $\mathcal{L}_{\text{scat}}$ .

The parameters considered in the optimization can be broadly split into two distinct categories. The first, which we call “controlled parameters,” are those over which the system designer can tune to optimize performance. We consider the controlled parameters to be the cavity length  $L$ , the mirror radii of curvature  $R$  and diameters  $D$ , and the outcoupler transmission  $T$ . We assume the ability to set these parameter freely up to some limits defined by the experimental system, e.g.,  $L \geq L_{\text{min}}$  dictated by ion-trap–electrode distances, or the relationship between  $R$  and  $D$  defined by the total cavity-mirror volume  $V$ .

Another important “controlled” parameter is the atomic branching ratio ( $\alpha = \Gamma_1/\Gamma$ , the ratio of the cavity arm

transition decay rate to the total decay rate of the intermediate state  $|e\rangle$ )—this value is set by the choice of atomic species and level scheme and although not continuously tunable, a range of values are possible.

The second group of parameters are the “parasitics,” all of which would ideally be zero. Here, we assume that the system designer has limited control or knowledge of the actual value but is able to establish an upper bound or tolerance on each. For our analysis of the ion-cavity system, the parasitics that we consider are the root-mean-squared roughness of the mirror surfaces  $\sigma$  and the misalignment  $M$ . There are naturally many parameters that are beyond the direct control of the experimentalists, such as noise on the magnetic field or laser intensity; however, these have much wider impact than on the ion-cavity system alone and lie beyond the scope of this paper.

The advantages of our two-stage optimization scheme over naive many-parameter optimization schemes are readily apparent when we take into account the dimensionality of our parameter space. Given a certain set of parasitics and constraints on the range of each controlled parameter, it is not unreasonable to perform a many-parameter optimization of the photon extraction rate to identify the ideal cavity design. Provided that a suitable algorithm is used, this family of optimization schemes can provide both the optimal values for the controlled parameters and some estimate of the local shape of the cost function. While this provides a measure of how stable the system performance is against small perturbations, the broader landscape of the cost function remains opaque, preventing the identification of robust near-optimal operating regions. Alternatively, the full cost landscape may be exhaustively calculated but this comes at considerable computational overhead and, ultimately, the designer will likely end up making lower-dimensional cuts through the parameter space in a bid to understand the structure. In contrast with such global methods, the scheme we present in the following sections is both extremely simple and immediately yields the pertinent cuts through the higher-dimensional parameter space. This additional information allows an experimentalist to select where to most efficiently expend effort to improve the parasitics with their system.

##### B. $\mathcal{L}_{\text{clip}}$ and optimization of cavity geometry

We begin by optimizing the geometric properties of the cavity ( $L$  and  $R$ ) to maximize the ion-cavity coupling rate while simultaneously minimizing the loss due to clipping of the mode by mirrors of finite diameter  $D$ . Independent optimizations of these two metrics drive the cavity geometry in broadly opposing directions, so we must consider how to optimally trade performance in one sector against that in the other.

Our first step is to rewrite Eq. (12) using Eqs. (13), (2a), and (3), to express this as a product of a purely atomic term, the cavity-arm branching ratio  $\alpha$ , a purely geometric term quantified by the divergence half-angle  $\theta$ , and the reciprocal of the intrinsic round-trip cavity loss  $\mathcal{L}_{\text{in}}$ :

$$C_{\text{in}} = 6\alpha\theta^2 \frac{1}{\mathcal{L}_{\text{in}}}. \quad (14)$$

As discussed in Sec. II, the intrinsic cavity loss  $\mathcal{L}_{\text{in}}$  is the sum of both clipping loss  $\mathcal{L}_{\text{clip}}$  and mirror surface-scattering loss  $\mathcal{L}_{\text{scat}}$ . In principle, to maximize  $P_{\text{ext}}$ , one needs to optimize the cavity geometry for both divergence  $\theta$  and clipping loss  $\mathcal{L}_{\text{clip}}$  simultaneously, which implicitly demands an involved procedure, accounting for  $\mathcal{L}_{\text{scat}}$  and introducing mirror transmission as an additional dimension of optimization. However, we now show that the geometric element of the optimization is well approximated by the maximization of a single parameter,  $\theta$ , within a boundary imposed by  $D$  and  $M$ , which ensures that  $\mathcal{L}_{\text{clip}}$  remains small. The remaining optimization of  $\kappa$  via outcoupler transmission then has an analytic solution.

To calculate the clipping loss, we assume that the small tilt due to misalignment causes a lateral shift,  $\Delta x = \phi L'/2$ , of the cavity mode at the mirror. As we are considering the perturbative regime of small clipping loss, and because we assume the cavity is operated far from transverse-mode degeneracy, mode-mixing effects can be ignored and the clipping is modeled simply as loss of amplitude [53]. By extending the work of Hunger *et al.* [54] from the axial case,  $\mathcal{L}_{\text{clip}}$  is approximated by considering the overlap of the shifted unity-normalized TE<sub>00</sub> Gaussian mode  $G'$  on the cavity-mirror surface  $S$ :

$$G' = G_0 e^{-\frac{(x-\Delta x)^2 + y^2}{w^2}} = G_0 \exp \left[ -\frac{(x - \phi \frac{L'}{2})^2 + y^2}{(\frac{\lambda}{\pi\theta})^2 + (\frac{L'}{2}\theta)^2} \right],$$

$$\mathcal{L}_{\text{clip}}(L, R, M, D) = 1 - \int_{S(D)} |G'|^2 dA, \quad (15)$$

where  $G_0$  is a normalization factor,  $w = w(\theta, \phi, L')$  is the cavity mode size at the cavity-mirror surface, and  $\phi = \phi(M, L, R)$  and  $\theta = \theta(M, L, R)$  are the tilt angle and the divergence half-angle.

Figure 3 shows the dependence of  $\mathcal{L}_{\text{clip}}$  on the cavity-geometry parameters  $L$  and  $R$ , for a range of mirror diameters  $D$  and misalignments  $M$ . Naturally, clipping losses are higher for small-diameter mirrors and large misalignments but for many cavity geometries the tilted mode remains well confined within the extent of the mirror and losses are negligible. However, as the cavity approaches the concentric configuration ( $R = L/2$ ), the mode tilt and divergence rapidly increase, along with the resultant losses. This leads to a threshold effect, which is particularly pronounced in the case of large-diameter mirrors and small misalignment.

Noting this threshold, we can set a maximum tolerable loss  $\mathcal{L}_{\text{clip}}^{\text{max}}$ , defining a critical  $R$ - $L$  contour. On one side of the contour, clipping loss will remain negligible for all  $\{R, L\}$ , but once crossed, losses will rapidly increase to an unacceptable level. Qualitatively, this critical geometry describes how concentric one can safely make the cavity before mode tilt and divergence make the optical mode hit the edge of the mirror. Here, we choose  $\mathcal{L}_{\text{clip}}^{\text{max}} = 1$  ppm, ensuring that clipping losses will be insignificant for even the highest-finesse cavities considered. We note that while we choose an arbitrary threshold for our analysis, the exact value makes little difference to the overall properties of the cavity. This is because, for all but the smallest mirrors, the onset of clipping loss is so rapid as to make our “critical contour” nearly independent of  $\mathcal{L}_{\text{clip}}^{\text{max}}$ .

Within this boundary  $\mathcal{L}_{\text{clip}} \ll \mathcal{L}_{\text{scat}}$  and so  $\mathcal{L}_{\text{in}} \approx \mathcal{L}_{\text{scat}}$ , which is independent of the cavity geometry. The optimization of clipping loss is thus reduced to a numerically calculated bound on the single-parameter optimization of  $\theta$ . Under this approximation, the intrinsic cooperativity becomes

$$C_{\text{in}} \approx \frac{6\alpha\theta^2}{\mathcal{L}_{\text{scat}}}, \quad (16)$$

where the dependence of  $\theta$  on the cavity geometry is given by Eq. (4c).

Figure 4 shows the dependence of  $\theta$  on the cavity geometry in the presence of varying amounts of misalignment. The clipping-loss boundaries for a range of mirror diameters are given as red contours, while the dashed black contours give the ratio of mode tilt to divergence angle,  $\phi/\theta$ , indicating which of these two effects dominates the clipping loss for a given configuration.

The optimal cavity geometry is therefore found by optimization of  $L$  and  $R$  along the boundary associated with the available mirror diameter. While  $\theta$  diverges at the concentric limit  $L \rightarrow 2R$ , this very divergence along with increased tendency for high  $\phi$  means that the clipping limit restricts how closely we can approach this condition. We note for completeness that  $\theta$  also diverges for  $L \rightarrow 0$  but there is always a practical limit on minimum cavity length, as discussed in Sec. I.

### C. $\mathcal{L}_{\text{scat}}$ and optimization of mirror transmission

In this section, we consider the impact of the optical properties of the cavity mirror itself; specifically, the scattering losses due to surface roughness  $\mathcal{L}_{\text{scat}}$  and the transmission coefficient of the outcoupler,  $T$ . The geometric optimization in Sec. IV B gives the maximum possible  $\theta$  and thus, through Eq. (16), also maximizes  $C_{\text{in}}$ . Maximizing  $P_{\text{ext}}$  now depends upon choosing a suitable value of  $T$ . We desire that photons generated within the cavity are transmitted through the outcoupler rather than scattered,

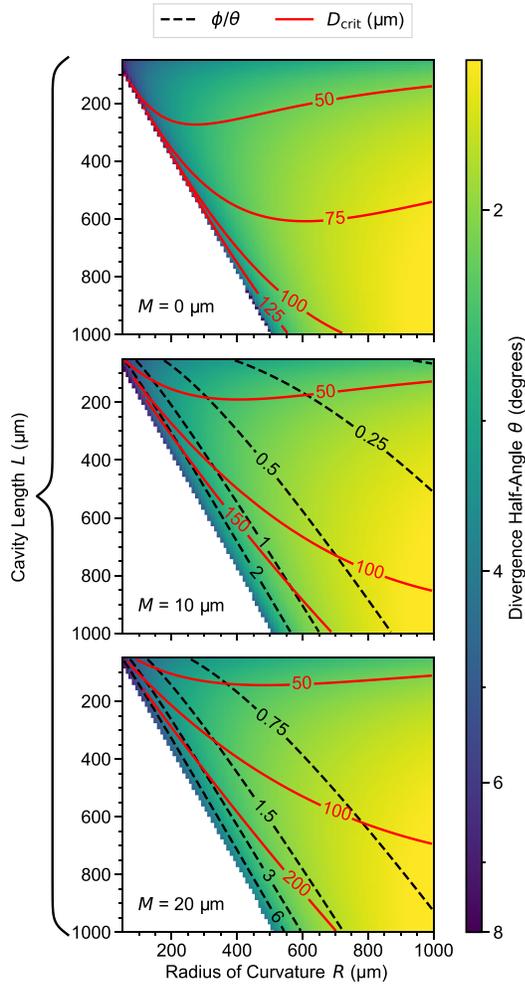


FIG. 4. Maximization of the geometric component of the cooperativity with finite-diameter mirrors. In each panel, the color map shows the divergence angle  $\theta$  as a function of the mirror radii of curvature and the nominal cavity length.  $\theta$ , and thus the geometric component of the cooperativity, increases strongly toward the concentric limit. The white region indicates where the cavity does not support stable modes. For three values of misalignment  $M$ , we show the minimum mirror diameter  $D$  for which  $\mathcal{L}_{\text{clip}}$  is 1 ppm (red contours) and the ratio of tilt angle to divergence  $\phi/\theta$  (black dashed contours). Note that the red contours are discontinuous where the cavities formed are geometrically limited. For a given diameter and misalignment, this allows us to identify the critical cavity geometry where clipping losses become significant and to attribute the source of clipping loss to either mode tilt ( $\phi \gg \theta$ ) or divergence ( $\theta \gg \phi$ ).

which is achieved in the limit  $T \gg \mathcal{L}_{\text{in}}$ . However, applying this limit means that  $C \ll C_{\text{in}}$ , which will reduce  $P_{\text{gen}}$ . Achieving optimal performance requires us to balance these two considerations. With  $\mathcal{L}_{\text{clip}} \approx 0$  after geometric optimization, the transmission optimization now depends only on the scattering loss. Assuming that there is transmission only through the outcoupler mirror, the optimal value for this parameter can be determined analytically

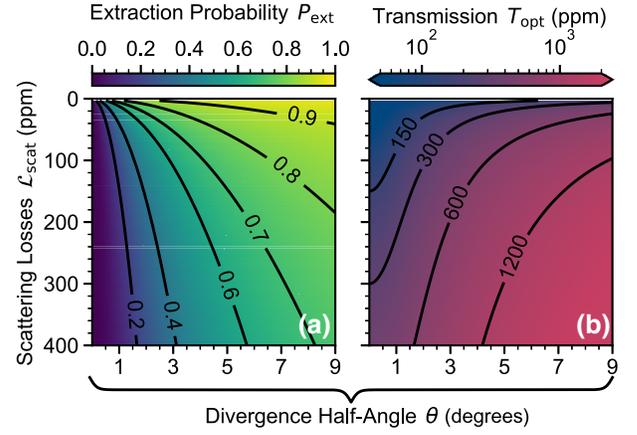


FIG. 5. Performance optimization versus mode divergence and round-trip scattering loss. (a) The optimized photon extraction probability,  $P_{\text{ext}}$ . The divergence angle is set by the geometry, as shown in Fig. 4, while the scattering losses  $\mathcal{L}_{\text{scat}}$  are fixed by the mirror manufacturing process. We assume an atomic transition branching ratio of  $\alpha = 1/20$ . (b) The corresponding optimal output coupler mirror transmission,  $T_{\text{opt}}$ , needed to achieve the value of  $P_{\text{ext}}$  in (a). The contours in each panel follow the color maps.

(following Ref. [52]):

$$T_{\text{opt}} = \mathcal{L}_{\text{in}} \sqrt{1 + 2C_{\text{in}}} \approx \mathcal{L}_{\text{scat}} \sqrt{1 + 2C_{\text{in}}}. \quad (17)$$

We apply this approximation and Eq. (11) to optimize  $P_{\text{ext}}$  for a range of values of both  $\theta$  and  $\mathcal{L}_{\text{scat}}$ , the results of which are shown in Fig. 5. As can be seen from the figure, for larger  $\theta$  it is possible to achieve higher extraction efficiencies for a given  $\mathcal{L}_{\text{scat}}$  by increasing mirror transmission. However, it is important to note that when scattering losses are sufficiently low, the high cooperativity saturates  $P_{\text{ext}}$ , making performance in this regime only weakly dependent on the cavity geometry.

#### D. Combined optimization of $P_{\text{ext}}$

In this section, we combine the results of Secs. IV B and IV C to establish an optimal approach to cavity design in the presence of known engineering tolerances.

Examples of typical geometric optimization landscapes are given in Figs. 6(a)–6(c) for three different values of scattering loss and a constant misalignment of  $M = 5 \mu\text{m}$ . In each panel, we show the optimum value of  $P_{\text{ext}}$  for a given  $R$  and  $L$  as a color map, with accompanying dashed black contours, while the dashed red contours in (a) show contours of fixed mirror volume,  $V$ , for which clipping losses, as defined in Sec. IV B, are fixed to 1 ppm. These plots resemble the divergence-angle color map of Fig. 4(b) but now show how these translate to overall  $P_{\text{ext}}$  performance when combined with a range of  $\mathcal{L}_{\text{scat}}$ . Note that while the qualitative behavior is similar in each case, in the presence of higher scattering losses, the efficiency is both

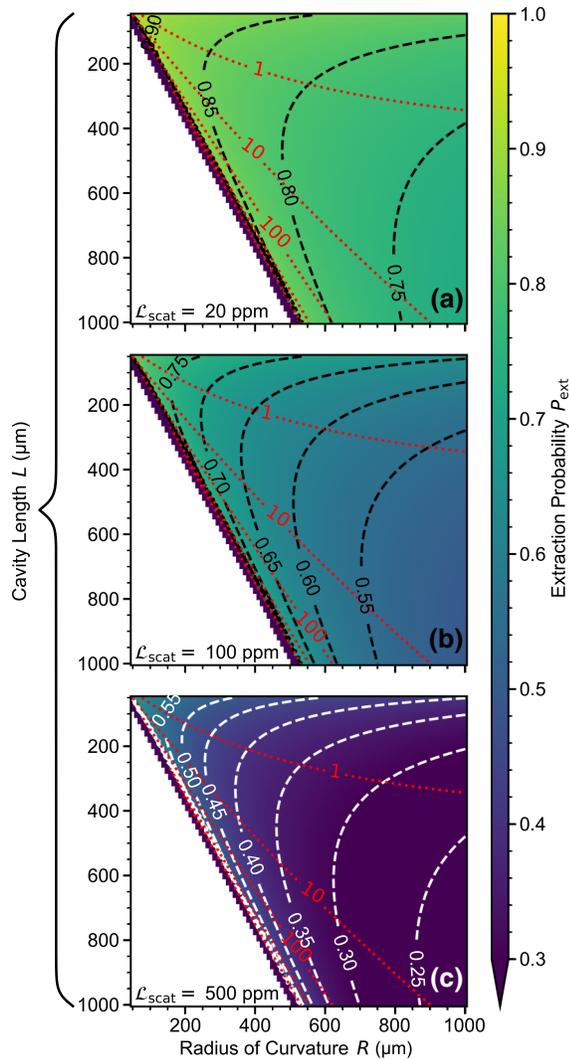


FIG. 6. The parameter space searched for the optimal ion-cavity system. The color mapping shows the performance as a function of the cavity length  $L$  and the mirror radius of curvature  $R$ , for fixed misalignment  $M = 5 \mu\text{m}$  and atomic branching ratio  $\alpha = 1/20$ . (a)–(c) The cases with round-trip scattering loss  $\mathcal{L}_{\text{scat}}$  equal to (a) 20 ppm, (b) 100 ppm, and (c) 500 ppm. The dashed black (white) contours show the optimized photon extraction probability,  $P_{\text{ext}}$ . The dotted red contours in (a) indicate the critical value of the mirror volume,  $V$ , in picoliters ( $1 \text{ pl} = 10^{-15} \text{ m}^3$ ), while the white region indicates where the cavity mode is unstable. Note that the stability region and the isovolume contours are consistent across all three configurations.

lower and more strongly dependent on the cavity geometry. We note that this optimum value of  $P_{\text{ext}}$  depends heavily on the value of the atomic branching ratio  $\alpha$  through the cooperativity. As shown in Eq. (16),  $C_{\text{in}} \propto \alpha$  and  $\alpha$  is independent of the cavity geometry, so this is a simple rescaling of the maximum achievable photon extraction probability. Hereafter, for the purposes of numerical evaluation, we fix

the value of the branching ratio to  $\alpha = 1/20$ , which appropriately matches schemes in alkali-like atoms and singly ionized alkaline-earth-like atoms where  $|0\rangle$  is a stable  $S$  state,  $|e\rangle$  is an short-lived  $P$  state, and  $|1\rangle$  is a metastable  $D$  state.

We can see from Fig. 6 that  $P_{\text{ext}}$  is maximized at different geometric positions depending on the clipping-loss boundary, set by mirror volume  $V$ . By following the isovolume contours, in the direction of decreasing  $L$  the extraction probability increases. The maximum  $P_{\text{ext}}$  thus lies at the top left corner of the region bounded by the clipping-loss boundary and the minimum cavity length. Therefore, when there is no further constraint on  $R$ , the optimal cavity length is *always* the minimum length compatible with the experiment in question.

In some situations, the assumption of an unconstrained optimization of  $R$  will not hold. This might occur if using preexisting mirrors with fixed values of both  $R$  and  $D$  or it might be a result of the challenges of making mirrors with low radii of curvature or with large diameters. In a typical case where  $R \simeq L_{\text{min}}$  and both  $R$  and  $D$  are fixed, the optimal length is, instead, the longest permissible without exceeding the clipping-loss threshold [55]. However, as we later show in Sec. VB, the resulting increase in  $P_{\text{ext}}$  is fairly limited and for robust operation it may still be preferable to use  $L = L_{\text{min}} \approx R$  in this scenario.

## V. OPTIMIZATION SUBJECT TO CONSTRAINTS

In Sec. IV, we have determined a strategy for optimizing the photon extraction probability in the presence of limitations, such as scattering loss and lateral misalignment of the cavity mirrors. In this section, we discuss the expected performance of systems optimized following this strategy. In Sec. VA, we consider such an optimization, subject to limitations on the cavity length and the volume of the mirror. In Sec. VB, we consider the impacts of inaccurate estimation of the manufacturing tolerances and deviation from the optimal design parameters to quantify the robustness of the optimized system.

### A. Performance versus parameter constraints

Previously, we have placed no explicit limitations on the geometry of the cavity system, other than requiring that the mirrors be physical and the resonator be stable. In practice, however, this is not achievable. In any realistic experimental scenario, there will be additional constraints applied beyond simple geometrical limits. In this subsection, we explore how the optimal cavity is impacted by experimentally motivated limits on mirror volume and minimum cavity length.

In the absence of any constraints, the optimal geometry will tend toward  $(L, R) \rightarrow 0$  but, as discussed in Sec. I, such a geometry is neither practical nor physical. Due to the presence of other experimental features within the

apparatus, e.g., ion-trap electrodes or focused laser beams, there is usually a minimum length that one can realistically achieve before impacting the overall functionality of the experiment. In Fig. 6, we show the parameter space over which one would optimize a cavity with no external constraints on  $L$  or  $R$ . Additionally, as dotted red contours, we show how the parameter space is further constrained when we consider the additional external constraints of fixed mirror volume and  $\mathcal{L}_{\text{clip}} \leq 1$  ppm. Combining this limitation with a minimum length constraint provides a unique maximum. To avoid biasing our results, we allow the optimizer to freely vary  $L$ . We find that, within the region that we explore, the optimal mirrors have a saggita that is significantly less than the minimum length of the cavity. Furthermore, the optimal length is always the minimum imposed, with the radius of curvature tending toward concentricity. Ultimately, the optimal cavities are often far from the concentric limit, due to the fixed mirror volume and the necessity of avoiding excessive clipping loss. As the available values of  $R$  and  $D$  are interrelated via the fixed mirror volume  $V$ , movement towards lower clipping losses, through an increase in  $D$ , requires an increase in  $R$ . This increase in  $R$  both reduces the divergence angle and moves the cavity away from concentric operation, assuming a fixed cavity length.

In Figs. 6(a), 6(b), and 6(c), we show the same geometries but for progressively higher scattering losses (20 ppm, 100 ppm, and 500 ppm, respectively). As is to be expected, there is no change in the general properties of the cavity,

with regard to the clipping-loss boundaries and the stability criterion. Crucially, for a given volume, the peak extraction probability occurs for the same geometry no matter the value of the scattering losses. However, as scattering losses increase, there is a significant decrease in the value of this maximal extraction probability.

Having identified the four parameters that define the optimization landscape,  $V$ ,  $L_{\text{min}}$ ,  $M$ , and  $\mathcal{L}_{\text{scat}}$ , we now explore the full four-dimensional parameter space. However, following from our preceding analysis, we have established that  $\mathcal{L}_{\text{scat}}$  simply reduces the maximal value of  $P_{\text{ext}}$  and the optimal value of transmission. As such, we are able to restrict ourselves to presenting results for a single value of  $\mathcal{L}_{\text{scat}}$  without changing the optimal cavity geometry presented. We have also already established, and shown in Fig. 4, that increasing misalignment corresponds to introducing nonzero  $\phi$  and modification of the position of clipping-loss contours. Naturally, this does lead to changes in the presented optimal geometry as well as the maximal value of  $P_{\text{ext}}$ . Numerically, however, we find that is a slow variation across a small range of parameters. As such, we can consider variation across a limited range of discrete values of  $M$ . Across the two remaining parameters,  $L_{\text{min}}$  and  $V$ , we expect a more complex trade-off for optimal  $P_{\text{ext}}$ : as such, we present these as continuous axes.

In Fig. 7, we show the parameters related to the optimum cavity for a given pair of  $L_{\text{min}}$ , the minimum length, and  $V$ , the milling volume. To determine the optimum, we use a Nelder-Mead algorithm with a random pair of

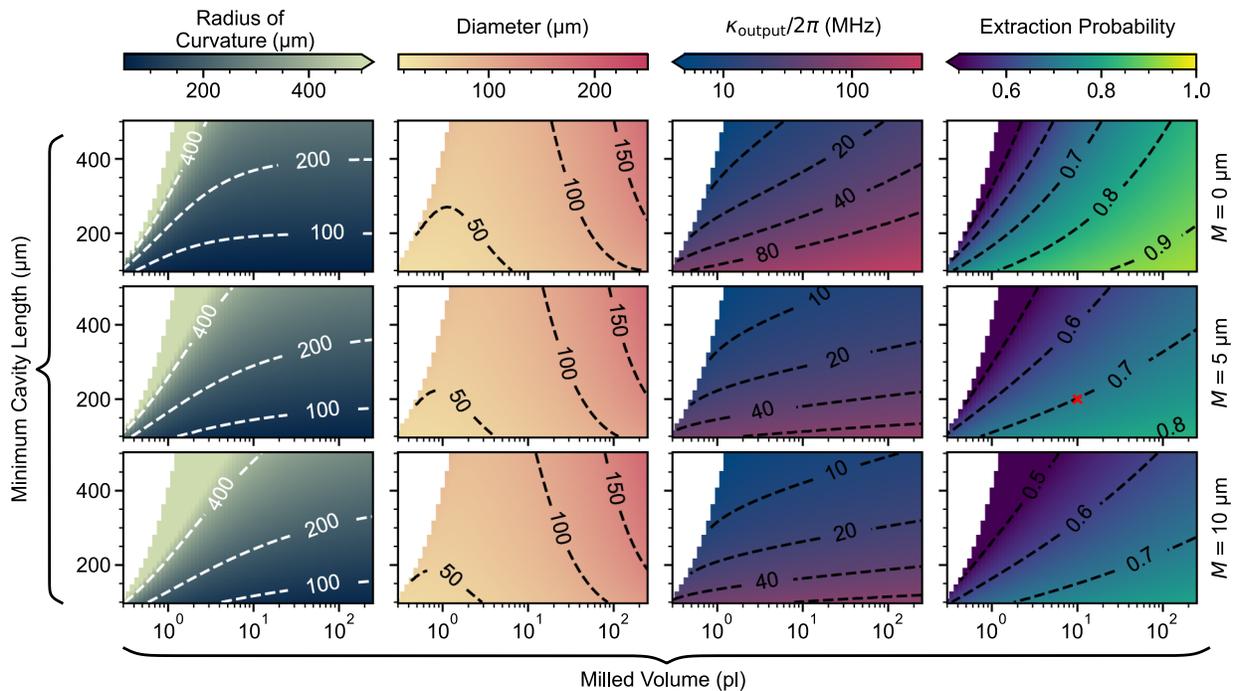


FIG. 7. The optimum cavity performance for a fixed milling volume and minimum lengths. For all configurations shown, the value of the scattering loss is fixed to 100 ppm. Regions where the cavity optimizer fails to find a suitable optimum, subject to these constraints, are shown in white. The red cross (“x”) shows the geometry considered in Fig. 8.

initial parameters  $L$  and  $R$ , chosen within the stability region. To prevent false local optima from being presented, we repeat this process with multiple initial conditions until a pair of optima are returned with less than 0.1% variation in the objective function. To aid convergence, we limit the maximum values of  $L$  and  $R$  to 3 mm.

Considering first the panels showing  $R$  and  $D$ , we see that for all but the lowest milling volumes and longest lengths, it is advantageous to maintain  $R \ll 500$   $\mu\text{m}$ . This increases the mode divergence and gives a smaller mode waist to increase coupling—however, as the cavity length is kept short, the mode size at the mirror is small enough to prevent clipping losses. As the maximum volume increases, it is more optimal to increase  $D$  than it is to decrease  $R$ , this being because as the cavity is designed to operate close to concentricity, changes in  $R$  then lead to the rapid onset of clipping losses.

A key consideration for using cavities in quantum networking is the rate at which a photon can be effectively extracted from the cavity. This is governed chiefly by  $\kappa_{\text{output}}$ , which is related to the transmission  $T$  by

$$\kappa_{\text{output}} = \frac{Tc}{4L}. \quad (18)$$

In Appendix B, we show that for, reasonable vSTIRAP pulse parameters, the adiabatic limit can be achieved with pulses of duration approximately  $10/\kappa$ . The lowest transmissions that we find are for cavities with long minimum lengths, small mirror volumes, and high misalignment; these yield values of  $\kappa_{\text{output}}/2\pi \approx 10$  MHz or extraction times of order 1  $\mu\text{s}$ . As discussed in Sec. III B, for the majority of experiments, attaining an attempt time on the order of 1  $\mu\text{s}$  is challenging due to necessary state preparation or cooling steps: as such, we do not consider these low  $\kappa$  solutions to be inherently limiting.

## B. Robustness of performance

While the optimum value for both  $M$  and  $\mathcal{L}_{\text{scat}}$  is zero, we optimize our system on the basis of an assumed achievable value for these parameters. These tolerances are only estimated before manufacture and in the case of  $M$  the true value will not be known until the cavity is operational. It is therefore crucial to consider the impact of incorrectly estimating these tolerances and to ensure that small errors of this sort will not substantially degrade the performance. In this section, we first consider a system in which the cavity length, outcoupler transmission, and mirror geometry are optimized, as shown previously, to maximize  $P_{\text{ext}}$  for fixed values of  $M$ ,  $\mathcal{L}_{\text{scat}}$ ,  $V$ , and  $L_{\text{min}}$ . We then consider how the performance of this optimized system is impacted if the estimates of  $M$  and  $\mathcal{L}_{\text{scat}}$  used during the optimization are incorrect and how manufacturing tolerances on the mirror assembly, through  $L$ ,  $T$ ,  $R$ , and  $D$ , impact performance.

In Fig. 8(a), we first consider how variations in length and transmission affect the cavity performance. Unsurprisingly, if the cavity is made longer than the design length, which in this case is also the minimum allowed, then the cavity enters a region of poor-performance operation, with clipping losses rapidly reaching greater than the 1 ppm threshold that we have previously imposed. In the converse case, in which the cavity is made shorter than designed (putting aside that our limitation is due to performance of other aspects of the system), we observe a more gradual drop off in the performance of the cavity. However, even in the more sensitive direction (where  $L > L_{\text{opt}}$ ), the performance remains at 99% of the optimal value for a 3% change in length. Considering now the effect of transmission, at the optimal length, the 1% degradation contour spans nearly a 400-ppm change in transmission, requiring only a 30% tolerance, far in excess of typical tolerances on high-quality dielectric coatings—which are typically specified to a few parts per million.

The other important cavity-manufacture parameters are those related to the mirror geometry:  $R$  and  $D$ . In Fig. 8(b), we show the parameter space spanned by these two parameters, noting that we have lifted our constraint that the milled volume is fixed, instead indicating this limitation with a contour of constant  $V$ . The dependence of  $P_{\text{ext}}$  on diameter is minimal, demonstrating the rapid onset of clipping losses, as shown first in Fig. 3. Additionally, we see a general increase in  $P_{\text{ext}}$  toward smaller values of  $R$ , again bounded by clipping losses. However, because of our constraint on milled volume, the optimum is found along a contour that is near perpendicular to that for clipping loss. This implies that the trade-off for reduced sensitivity to clipping is a loss of  $P_{\text{ext}}$ . As such, when subject to the constraint on maximum volume, the opportunity cost of conservative mirror design, i.e., larger diameters, is significant in the available extraction probability. For instance, a 5  $\mu\text{m}$  increase in diameter from optimal causes a 5% drop in extraction probability due to the large increase in  $R$  needed to maintain the constraint on design volume.

Finally, we consider the possibility that the estimates made of the values of  $M$  and  $\mathcal{L}_{\text{scat}}$  are incorrect. This is not an unrealistic scenario, as both parameters need to be estimated *a priori*. As is to be expected, when both parameters are initially underestimates, that is, the true values of  $M$  and  $\mathcal{L}_{\text{scat}}$  are greater than used for the optimization, the cavity performance is degraded. With increased  $\mathcal{L}_{\text{scat}}$ , the cavity remains stable but with reduced extraction probability; however, for  $M$ , the cavity rapidly loses performance as it approaches a region of large clipping losses. In the converse case, where the two parameters are overestimates, the cavity continues to function at, or beyond, its design level. With lower  $\mathcal{L}_{\text{scat}}$ , the extraction probability increases, there are no performance gains. However, for overestimating  $M$ , as clipping losses have already been made negligible, further reductions provide no gains. From

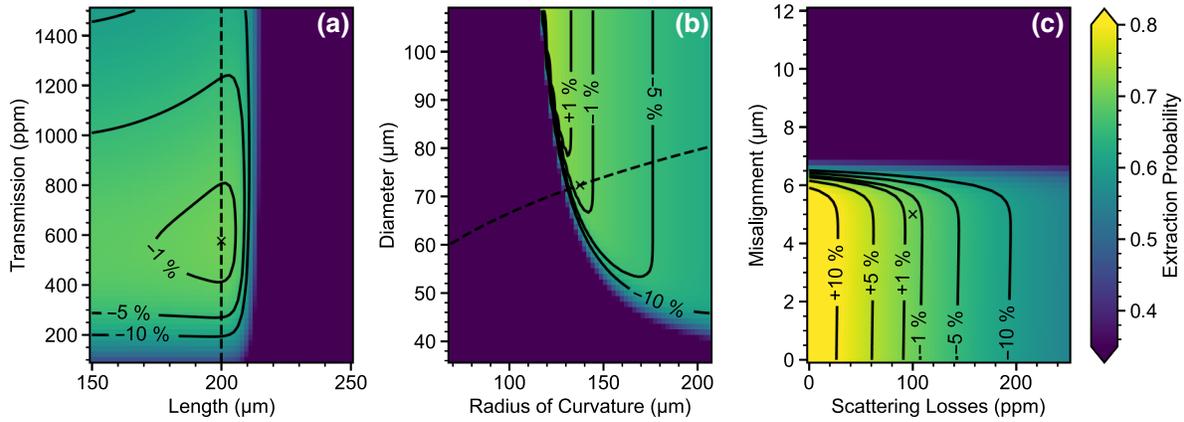


FIG. 8. The performance of an optimized ion-cavity system subject to manufacturing defects. In all three panels, the extraction probability is shown by the color map, the cross (“ $\times$ ”) represents the optimum set of parameters, and the solid black contours show the  $\pm 1\%$ ,  $5\%$ , and  $10\%$  levels in extraction probability. The cavity mirrors have a fixed design volume of  $10\text{ pl}$ ,  $\mathcal{L}_{\text{scat}} = 100\text{ ppm}$ , and  $M = 5\text{ }\mu\text{m}$ . (a) The cavity extraction probability as a function of the cavity length and transmission, keeping  $R$  and  $D$  constant. The dashed vertical line indicates the fixed value of the minimum length used to optimize the mirror geometry. We see that the maxima are not strongly constrained by  $T$ ; however, for an approximately  $5\%$  increase in  $L$ , the extraction probability drops away; the converse is not true with a  $1\%$  error on  $P_{\text{ext}}$  for an approximately  $10\%$  reduction in  $L$ . (b) Changes in  $P_{\text{ext}}$  due to mirror fabrication, with the volume constraint lifted. The dashed black line is the isovolume contour for the design volume. Increasing  $R$  degrades performance rapidly, while changes that decrease  $R$  lead to unstable operation. Increasing  $D$  has only a small impact on performance, while decreasing  $D$  also leads the cavity toward high clipping losses and poor extraction probabilities. (c) The change in extraction probability due to poor estimation of  $M$  and  $\mathcal{L}_{\text{scat}}$ . Underestimates of  $M$  result in unstable cavity configurations, while underestimates of  $\mathcal{L}_{\text{scat}}$  simply degrade performance. Overestimates of  $\mathcal{L}_{\text{scat}}$  improve performance; however, there is no such change for overestimates of  $M$ .

this analysis, we conclude that making accurate estimates of  $M$  is critical when building static monolithic cavities and so to ensure stable performance, the value used for optimization should be the  $2\text{-}\sigma$  upper bound or greater. However, state-of-the-art performance is achievable within the expected experimental tolerances.

## VI. SUMMARY AND OUTLOOK

In this paper, we consider how the design of “real-world” ion-cavity systems can be tailored to optimize their performance for quantum networking applications; in particular, the generation of single photons from ions for measurement-based remote-entanglement schemes. We begin by defining several figures of merit for the performance of such schemes. We show that through suitable choice of the ion species, the photon-generation scheme, and the driving pulse, it is possible to ensure that the impact of the cavity design on the entanglement fidelity and the attempt rate can be made negligible. This allows us to identify a single figure of merit for our optimization, namely, the probability of photon extraction from the cavity per attempt cycle,  $P_{\text{ext}}$ .

To optimize  $P_{\text{ext}}$ , we show that in this context, the cooperativity can be well approximated as a product of three independent components, depending, respectively, on the cavity geometry, the electronic structure of the ion, and scattering from the mirror surface. The separability of the problem allows us to perform the optimization process in

a serial manner, first numerically optimizing the geometry (for a given set of limits) before identifying the optimal mirror parameters analytically. The significant reduction in the dimensionality of this optimization problem versus the more general case helps provide intuitive demonstrations of the trade-offs inherent in such ion-cavity system designs and highlights regions likely to yield robust performance.

We explore these trade-offs subject to realistic constraints, coming from minimum cavity length and the maximum available milling volume. By considering radii of curvature and diameters of mirrors to be freely exchangeable for a fixed volume, we find that the optimum microcavity is always the shortest permitted. The cavity-mirror radius of curvature then tends toward the concentric limit ( $R \rightarrow L/2$ ) with the diameter increased to prevent clipping at the  $1\text{ ppm}$  level.

Provided that the maximum transverse misalignment of the cavity mirrors is well characterized and accounted for, we show that very high extraction probabilities can be achieved. The relaxation of alignment requirements makes possible a range of much simpler techniques for the fabrication of microcavities, a step on the route to scalable quantum networks.

By studying the robustness of these optimal configurations with respect to errors in the machining and uncontrolled parameters, we find that should our goal be to maximize extraction probability, efforts should always be expended on reducing scattering losses but residual misalignment can be accounted for at the design level,

provided that it can be predicted. Additionally, we find that there is little reward for conservative mirror design with respect to clipping losses, as performance rapidly degrades as the radius of curvature moves from the optimum. More stable operation can be achieved by increasing the diameter of the mirror but at a significant opportunity cost.

Our conclusions regarding “optimal” cavity design do not, broadly speaking, contradict contemporary trends but the significant range of geometries and operating regimes seen in groups around the world indicates that the field is some way from consensus on the topic. We intend our work to help other researchers to approach this multifaceted problem systematically, adapting our approach to the parameter regime defined by their experimental constraints.

The simulation code used in this paper is available under an open-source license [56] and we encourage others to use and modify it toward the design of their own systems.

### ACKNOWLEDGMENTS

We would like to thank Peter Horak for stimulating discussions. This work was funded by the United Kingdom Engineering and Physical Sciences Research Council Hub in Quantum Computing and Simulation (Grant No. EP/T001062/1) and the European Union Quantum Technology Flagship Project AQTION (Grant No. 820495). The data and analysis presented is available online [57] or from the authors at reasonable request.

### APPENDIX A: TEMPORAL MIXING

For many quantum information applications, it is important that the photons extracted from the cavity are indistinguishable from each other. Where this is the case, care must be taken in the design of the vSTIRAP-type schemes that we consider, where spontaneous emission to the initial state during the photon-production process can significantly enhance distinguishability.

The mechanism for this process is described in detail in Refs. [23,41] but can be summarized as follows. During the application of the Raman driving pulse, there is a finite chance that the ion will spontaneously emit a photon into free space and decay to the initial state,  $|0\rangle$ , interrupting the population transfer to the target state  $|1\rangle$ . Once in this initial state, the vSTIRAP transfer begins again, now driven by the remaining portion of the driving pulse (which may result in cavity emission with or without one or more further spontaneous emission events). Note that if spontaneous emission does not occur directly to  $|0\rangle$  or to another state that itself rapidly decays to  $|0\rangle$ , then the transfer process terminates; this has an impact on the average production rate but not on indistinguishability and is generally far more tolerable.

In each photon-production attempt resulting in cavity emission, the emitted photon has a waveform determined by the time of the last spontaneous emission event (relative to the start of the driving pulse). On average, the emission of the cavity is a mixture of these photons, weighted by the probability of observing the corresponding final spontaneous emission time. Photons produced after spontaneous emission exhibit distorted and retarded photon wave packets, with the effects being more severe the later the final spontaneous emission event occurs. This means that two photons from the same system may have highly distinguishable temporal wave packets.

Combating this source of error requires that the last spontaneous emission times do not vary or that spontaneous emission followed by cavity emission happens very rarely. There are three main ways to achieve this:

(1) Increasing the cavity cooperativity: increasing  $C$  will increase the ratio of cavity emission to spontaneous emission, increasing the likelihood that the cavity photon is emitted without a spontaneous emission event during the production process and thus decreasing distinguishability. Increasing the cooperativity is, however, challenging in the presence of misalignment and finite mirror size (as discussed in the main body of the text from the perspective of increasing total extraction) and therefore increasing the cooperativity to levels where this error is insignificant may not be technically feasible. Even where possible, this adds additional complication to the cavity optimization process, involving trading off indistinguishability and production rate, a balancing act that depends greatly on the specific system and application.

(2) Fast excitation: if the excitation pulse is fast compared to the spontaneous lifetime of the excited-state and cavity decay time (e.g., by utilizing a picosecond pulsed laser), spontaneous emission and subsequent cavity emission can only occur during this very short excitation window and all photons emitted from the cavity effectively have the same starting time. Note that in the case of fast excitation, production of the photon no longer proceeds via a vSTIRAP process and the corresponding  $P_{\text{gen}}$  is lower and obeys a distinct expression [41]; if such a scheme is used, the optimization steps in this paper would need to be significantly modified.

(3) Using a system with a favorable branching ratio: if the levels for the scheme are chosen such that the branching ratio to the initial state is much lower than that to the final state, i.e.,  $\Gamma_0 \ll \Gamma_1$ , the temporal distinguishability introduced by this process can be rendered negligible at even modest cooperativities. In this case, there is no need to drive the system very quickly to avoid the degradation of indistinguishability and an adiabatic drive can be used without problems.

In our own experiments, we plan to use a  $\Lambda$  system that satisfies the third of these criteria and thus we need not consider issues of indistinguishability further during the optimization of our design. For the purposes of the work presented in this paper, we assume that this criterion is met; we believe that efficient and scalable systems will be built on intrinsically robust production schemes and that this matter should be addressed independently of the optimization of extraction efficiency.

## APPENDIX B: ADIABATIC DRIVING

The formulas used for the photon extraction probability in this paper are derived in the adiabatic limit, assuming an infinitely slow Raman drive pulse. This is algebraically straightforward and avoids the need to consider the details of the driving pulse while optimizing the cavity parameters (once the cavity parameters are chosen, the drive can then be optimized according to experimental requirements). However, practical schemes must operate in finite time and, indeed, reliable and *fast* photon production underpins many of the target applications for this work. The use of a finite-time pulse always reduces the output probability relative to the adiabatic result; the important consideration for the validity of our approach (i.e., the separation of the cavity parameters from the driving pulse) is whether the ion-cavity system is likely to be driven rapidly enough to deviate significantly from the infinite-time result.

The validity of this assumption varies depending upon many experimental factors, including the parameters  $g_0$ ,  $\kappa$ , and  $\gamma$ , and the characteristic time scale and exact pulse shape of the driving pulse used. However, for typical smooth pulse profiles, the most important criterion for determining the output probability is the dimensionless quantity  $\kappa\tau$ , where  $\tau$  is the width of the finite-time window in which the photon must be produced. To test this, we simulate the integrated output probability versus  $\tau$  for a variety of  $\Lambda$ -type systems driven with a sine-squared pulse amplitude (as used in Ref. [58]), with the peak amplitude and pulse width chosen to maximize the output for each  $\tau$ .  $\Lambda$  systems are chosen with a range of cooperativities  $C = \{0.1, 1, 10\}$  and cavity-to-atomic decay-rate ratio  $\kappa/\gamma = \{0.1, 1, 10\}$  in order to account for the variety of cavity parameters in current microcavity experiments. The results, presented in Fig. 9, show that the output approximates the adiabatic limit for  $\kappa\tau \gtrsim 10$ . Therefore, as long as the photon-production window is several times longer than the characteristic decay time of the cavity, the achievable output is very similar to the adiabatic limit and thus the latter provides an adequate approximation while optimizing the cavity design.

The specific systems considered in this work cover cavity decay rates where  $20 \text{ MHz} \lesssim \kappa/2\pi \lesssim 300 \text{ MHz}$ , which satisfies the adiabaticity condition for

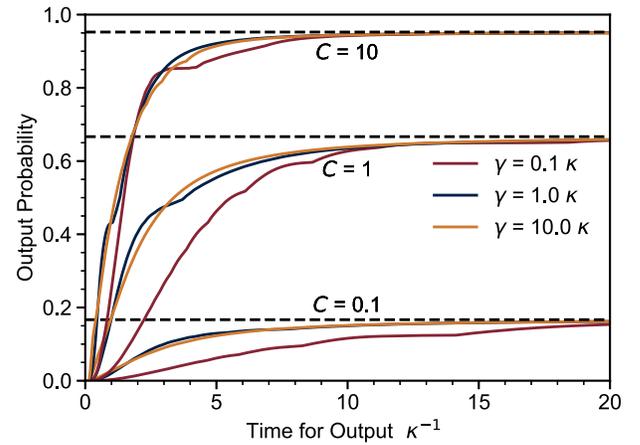


FIG. 9. A plot of the finite-time photon outputs for a range of ion-cavity systems. The evolution of each system is simulated with a selection of pulses the amplitude of which follows a sine-squared shape and is always on resonance with the  $|0\rangle \rightarrow |e\rangle$  transition. At each time point, the largest value that the output probability takes over the range of simulated pulses is recorded as an example of a lower bound on the achievable output. The systems simulated form three groups with cooperativities of  $C = 0.1$ ,  $C = 1.0$ , and  $C = 10.0$ , where the adiabatic limit  $P_{\text{gen}} = 2C/(2C + 1)$ , valid for slow driving or an infinite output time window, is indicated with a dashed black line. Within each set,  $\gamma$  takes three different values to compare across a range of systems; note that for system configurations close to the  $\gamma = 0.1\kappa$  case,  $\kappa$  is generally high and using a longer pulse  $\kappa\tau \gg 10$  is unlikely to present a bottleneck.

$500 \text{ ns} \gtrsim \tau \gtrsim 33 \text{ ns}$ . These times are smaller than or comparable to other processes required during each photon-production attempt, such as optical pumping for state preparation, optical modulator latency, and other control-system-related overheads and, in larger networks, the photon propagation time. We therefore consider that achieving near-adiabatic performance is possible for all the systems we consider, without a significant reduction in achievable repetition rate, although we note that for the lowest- $\kappa$  cavities that we consider, these effects are certainly worthy of careful consideration if using repetition rates approaching 1 MHz. In these cases, it should also be noted that the values presented here represent a lower bound on the output achievable for a pulse of this length, given that the pulse that we use has a predefined shape and frequency; carefully designing the driving pulse can help increase the output probability [59] for lower  $\kappa\tau$ , though the improvement is typically modest over the functions used here.

- [1] P. Hamilton, M. Jaffe, J. M. Brown, L. Maisenbacher, B. Estey, and H. Müller, Atom Interferometry in an Optical Cavity, *Phys. Rev. Lett.* **114**, 100405 (2015).

- [2] A. J. Park, J. Trautmann, N. Šantić, V. Klüsener, A. Heinz, I. Bloch, and S. Blatt, Cavity-enhanced optical lattices for scaling neutral atom quantum technologies to higher qubit numbers (2022).
- [3] B. Braverman, A. Kawasaki, E. Pedrozo-Peñafiel, S. Colombo, C. Shu, Z. Li, E. Mendez, M. Yamoa, L. Salvi, D. Akamatsu, Y. Xiao, and V. Vuletić, Near-Unitary Spin Squeezing in  $^{171}\text{Yb}$ , *Phys. Rev. Lett.* **122**, 223203 (2019).
- [4] E. Pedrozo-Peñafiel, S. Colombo, C. Shu, A. F. Adiyatullin, Z. Li, E. Mendez, B. Braverman, A. Kawasaki, D. Akamatsu, Y. Xiao, and V. Vuletić, Entanglement on an optical atomic-clock transition, *Nature* **588**, 414 (2020).
- [5] J. A. Muniz, D. J. Young, J. R. K. Cline, and J. K. Thompson, Cavity-QED measurements of the  $^{87}\text{Sr}$  millihertz optical clock transition and determination of its natural linewidth, *Phys. Rev. Res.* **3**, 023152 (2021).
- [6] S. J. van Enk, H. J. Kimble, and H. Mabuchi, Quantum information processing in cavity-QED, *Quantum Inf. Process.* **3**, 75 (2004).
- [7] L.-M. Duan and H. J. Kimble, Scalable Photonic Quantum Computation through Cavity-Assisted Interactions, *Phys. Rev. Lett.* **92**, 127902 (2004).
- [8] J. Cho and H.-W. Lee, Generation of Atomic Cluster States through the Cavity Input-Output Process, *Phys. Rev. Lett.* **95**, 160501 (2005).
- [9] W. J. Munro, A. M. Stephens, S. J. Devitt, K. A. Harrison, and K. Nemoto, Quantum communication without the necessity of quantum memories, *Nat. Photonics* **6**, 777 (2012).
- [10] A. Reiserer, N. Kalb, G. Rempe, and S. Ritter, A quantum gate between a flying optical photon and a single trapped atom, *Nature* **508**, 237 (2014).
- [11] A. Kuhn, M. Hennrich, and G. Rempe, Deterministic Single-Photon Source for Distributed Quantum Networking, *Phys. Rev. Lett.* **89**, 067901 (2002).
- [12] D. L. Moehring, P. Maunz, S. Olmschenk, K. C. Younge, D. N. Matsukevich, L.-M. Duan, and C. Monroe, Entanglement of single-atom quantum bits at a distance, *Nature* **449**, 68 (2007).
- [13] 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).
- [14] T. E. Northup and R. Blatt, Quantum information transfer using photons, *Nat. Photonics* **8**, 356 (2014).
- [15] P. Kobel, M. Breyer, and M. Köhl, Deterministic spin-photon entanglement from a trapped ion in a fiber Fabry-Pérot cavity, *npj Quantum Inf.* **7**, 6 (2021).
- [16] 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).
- [17] C. Maurer, C. Becher, C. Russo, J. Eschner, and R. Blatt, A single-photon source based on a single  $\text{Ca}^+$  ion, *New J. Phys.* **6**, 94 (2004).
- [18] 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).
- [19] M. Hijlkema, B. Weber, H. P. Specht, S. C. Webster, A. Kuhn, and G. Rempe, A single-photon server with just one atom, *Nat. Phys.* **3**, 253 (2007).
- [20] H. G. Barros, A. Stute, T. E. Northup, C. Russo, P. O. Schmidt, and R. Blatt, Deterministic single-photon source from a single ion, *New J. Phys.* **11**, 103004 (2009).
- [21] J. D. Sterk, L. Luo, T. A. Manning, P. Maunz, and C. Monroe, Photon collection from a trapped ion-cavity system, *Phys. Rev. A* **85**, 062308 (2012).
- [22] M. Mücke, J. Bochmann, C. Hahn, A. Neuzner, C. Nölleke, A. Reiserer, G. Rempe, and S. Ritter, Generation of single photons from an atom-cavity system, *Phys. Rev. A* **87**, 063805 (2013).
- [23] T. Walker, S. V. Kashanian, T. Ward, and M. Keller, Improving the indistinguishability of single photons from an ion-cavity system, *Phys. Rev. A* **102**, 032616 (2020).
- [24] H. J. Kimble, The quantum Internet, *Nature* **453**, 1023 (2008).
- [25] A. Reiserer and G. Rempe, Cavity-based quantum networks with single atoms and optical photons, *Rev. Mod. Phys.* **87**, 1379 (2015).
- [26] H. Takahashi, E. Kassa, C. Christoforou, and M. Keller, Strong Coupling of a Single Ion to an Optical Cavity, *Phys. Rev. Lett.* **124**, 013602 (2020).
- [27] 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).
- [28] N. Podoliak, H. Takahashi, M. Keller, and P. Horak, Comparative Numerical Studies of Ion Traps with Integrated Optical Cavities, *Phys. Rev. Appl.* **6**, 044008 (2016).
- [29] C. Monroe and J. Kim, Scaling the ion trap quantum processor, *Science* **339**, 1164 (2013).
- [30] N. H. Nickerson, J. F. Fitzsimons, and S. C. Benjamin, Freely Scalable Quantum Technologies Using Cells of 5-to-50 Qubits with Very Lossy and Noisy Photonic Links, *Phys. Rev. X* **4**, 041041 (2014).
- [31] 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).
- [32] J. Bochmann, M. Mücke, G. Langfahl-Klabes, C. Erbel, B. Weber, H. P. Specht, D. L. Moehring, and G. Rempe, Fast Excitation and Photon Emission of a Single-Atom-Cavity System, *Phys. Rev. Lett.* **101**, 223601 (2008).
- [33] G. S. Vasilev, D. Ljunggren, and A. Kuhn, Single photons made-to-measure, *New J. Phys.* **12**, 063024 (2010).
- [34] B. W. Shore, Picturing stimulated Raman adiabatic passage: A STIRAP tutorial, *Adv. Opt. Photonics* **9**, 563 (2017).
- [35] K. Bergmann *et al.*, Roadmap on STIRAP applications, *J. Phys. B: At., Mol. Opt. Phys.* **52**, 202001 (2019).
- [36] P. B. R. Nisbet-Jones, J. Dille, D. Ljunggren, and A. Kuhn, Highly efficient source for indistinguishable single photons of controlled shape, *New J. Phys.* **13**, 103036 (2011).
- [37] V. Krutyanskiy, M. Meraner, J. Schupp, V. Krcmarsky, H. Hainzer, and B. P. Lanyon, Light-matter entanglement over 50 km of optical fibre, *Npj Quantum Inf.* **5**, 72 (2019).

- [38] L. Luo, D. Hayes, T. Manning, D. Matsukevich, P. Maunz, S. Olmschenk, J. Sterk, and C. Monroe, Protocols and techniques for a scalable atom-photon quantum network, *Fortschritte der Physik* **57**, 1133 (2009).
- [39] A. Stute, B. Casabone, P. Schindler, T. Monz, P. O. Schmidt, B. Brandstätter, T. E. Northup, and R. Blatt, Tunable ion-photon entanglement in an optical cavity, *Nature* **485**, 482 (2012).
- [40] P. B. R. Nisbet-Jones, J. Dilley, A. Holleczeck, O. Barter, and A. Kuhn, Photonic qubits, qutrits and ququads accurately prepared and delivered on demand, *New J. Phys.* **15**, 053007 (2013).
- [41] S. Gao, Ph.D. thesis, University of Oxford (2020).
- [42] For example, in the case of a cavity with optical axis near parallel to the magnetic field, driving a  $\pi$  transition.
- [43] Hard clipping by an aperture located on the cavity axis away from the mirrors has a similar effect, with the dependence on the diameter scaled according to the expected mode diameter at the aperture location.
- [44] J. M. Bennett, Recent developments in surface roughness characterization, *Meas. Sci. Technol.* **3**, 1119 (1992).
- [45] K. Ott, S. Garcia, R. Kohlhaas, K. Schüppert, P. Rosenbusch, R. Long, and J. Reichel, Millimeter-long fiber Fabry-Pérot cavities, *Opt. Express* **24**, 9839 (2016).
- [46] 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).
- [47] D. P. Nadlinger, P. Drmota, B. C. Nichol, G. Araneda, D. Main, R. Srinivas, D. M. Lucas, C. J. Ballance, K. Ivanov, E. Y.-Z. Tan, P. Sekatski, R. L. Urbanke, R. Renner, N. Sangouard, and J.-D. Bancal, Experimental quantum key distribution certified by Bell's theorem, *Nature* **607**, 682 (2022).
- [48] R. Nigmatullin, C. J. Ballance, N. de Beaudrap, and S. C. Benjamin, Minimally complex ion traps as modules for quantum communication and computing, *New J. Phys.* **18**, 103028 (2016).
- [49] T. H. Doherty, Ph.D. thesis, University of Oxford (2021).
- [50] G. Vittorini, D. Hucul, I. V. Inlek, C. Crocker, and C. Monroe, Entanglement of distinguishable quantum memories, *Phys. Rev. A* **90**, 040302 (2014).
- [51] E. Kassa, W. Hughes, S. Gao, and J. F. Goodwin, Effects of cavity birefringence on remote entanglement generation, Arxiv e-prints **2008**, 11712 (2020).
- [52] 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).
- [53] D. Clarke and P. Horak, in *Quantum Technologies 2018*, edited by J. Stuhler, A. J. Shields, and M. J. Padgett, *Proc. International Society for Optics and Photonics*, Vol. 10674 (SPIE, Bellingham, WA, 2018).
- [54] D. Hunger, T. Steinmetz, Y. Colombe, C. Deutsch, T. W. Hänsch, and J. Reichel, A fiber Fabry-Pérot cavity with high finesse, *New J. Phys.* **12**, 065038 (2010).
- [55] In the rarer case of  $R \gg L_{\min}$  combined with small  $D$  and/or large  $M$ , minimizing  $L$  can again be the best strategy; however, overall performance in this regime tends to be poor and we do not consider it further.
- [56] The code used in this work is available at <https://doi.org/10.5281/zenodo.7020047>.
- [57] <https://doi.org/10.5287/bodleian:o8jV7bYjN>.
- [58] P. K. Molony, P. D. Gregory, A. Kumar, C. R. Le-Sueur, J. M. Hutson, and S. L. Cornish, Production of ultracold  $^{87}\text{Rb}^{133}\text{Cs}$  in the absolute ground state: Complete characterisation of the stimulated Raman adiabatic passage transfer, *ChemPhysChem* **17**, 3811 (2016).
- [59] G. S. Vasilev, A. Kuhn, and N. V. Vitanov, Optimum pulse shapes for stimulated Raman adiabatic passage (2009).