An investigation towards single-emitter hybrid photonic devices

Palstra, I.M.

Citation for published version (APA):
Palstra, I. M. (2021). *An investigation towards single-emitter hybrid photonic devices*.

An investigation towards single-emitter hybrid photonic devices

Isabelle M. Palstra

An investigation towards single-emitter hybrid photonic devices

Isabelle Maria Palstra

An investigation towards single-emitter hybrid photonic devices

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad van doctor

aan de Universiteit van Amsterdam

op gezag van de Rector Magnificus

prof. dr. ir. K.I.J. Maex

ten overstaan van een door het College voor Promoties ingestelde commissie,

in het openbaar te verdedigen in de Agnietenkapel

op maandag 27 september 2021, te 14.00 uur

door

Isabelle Maria Palstra

geboren te Summit

When things go wrong, as they sometimes will,

When the road you're trudging seems all uphill,

When the funds are low and the debts are high,

And you want to smile but you have to sigh.

When all is pressing you down a bit-

Rest if you must, but don't you quit.

Success is failure turned inside out,

The silver line on the clouds of doubt,

And you can never tell how close you are,

It may be near when it seems far.

So stick with the fight when you're the hardest hit -

It's when things go wrong that you must not quit.

- John Greenleaf Whittier, 1892

CONTENTS

1

INTRODUCTION: WHERE LIGHT AND MATTER MEET

*"The White Rabbit put on his spectacles. 'Where shall I begin, please your Majesty?' he asked.
'Begin at the beginning,' the King said gravely, 'and go on till you come to the end: then stop."*

— Lewis Carroll, *Alice's Adventures in Wonderland*, 1865

THIS thesis concerns the manipulation of light with matter. When it comes to manipulating light, there are a great number of methods to choose from that are evident in our day-to-day life. Applying a coat of paint to a newly finished project is, in a sense, a form of manipulating light, as this determines which wavelengths of light the object will absorb, which colors it will transmit, and which colors will be reflected after scattering [1]. The central importance of color and light is also evident from numerous examples in nature. For instance, the color of plants is essential for their energy balance and photosynthesis. Animals display the brightest colors to attract a mate or the subtlest ones to hide from predators [2–4]. Hummingbirds [5], beetles [6], butterflies [7], and even fruits [8] create a splendor of iridescent hues and vibrant patterns by growing structures on the scale of visible wavelengths, while corals, fish, and scorpions glow eerily under UV illumination [9].

Since the dawn of civilization, humans have joined in this competition of color. With the hue of our clothes and other possessions, we send signals to each other, communicating our mood, intention, status, and personality. In more recent times, research has expanded our knowledge about the nature of light and electromagnetic waves more generally. Indeed, with the development of the radio just around the turn of the 20th century, humans entered the era of telecommunication, first using light that our human eyes can't observe. In 1909, Guglielmo Marconi and Karl Ferdinand Braun were jointly awarded a Nobel prize "in recognition of their contributions to the development of wireless telegraphy" [10]. Since that time, with a rapidly accelerating pace, we gained an aptitude for worldwide and nearly instant communication that nowadays speeds through endless lengths† of glass fiber optical cables spanning our planet, allowing us to communicate with people across the world with imperceptibly small delays. Charles Kuen Kao was awarded a Nobel Prize "for groundbreaking achievements concerning the transmission of light in fibers for optical communication" [11], one hundred years after the Nobel awarded to Marconi and Braun. Our current communications network is made possible by a wealth of knowledge and skills we have acquired that allow us to manipulate how light and matter interact. The rise of fiber-optical communication is furthermore fueling the emergence of many data-intensive devices. Many of these themselves leverage optics and light-matter interaction, such as optical imaging and sensing [12–14]. The data that these applications produce in turn drives the need for data transport and data processing technologies that go well beyond signal transport by fibers.

At the time of writing, people are isolating themselves from each other in order to stem a global pandemic. Physical closeness, the most natural way for humans to connect, is kept to a bare minimum. In this time, light-based technology in many forms is helping to bring people together.

1.1. Single photons and single quantum emitters

Within the broad field of light-matter interaction, this thesis more specifically deals with light-matter interaction on the level of single quanta of light, i.e., single photons and single quantum systems. The interaction of single photons with single emitters,

---

†Submarine fiber cables are the backbone of the internet and reportedly span 1.2 million kilometers. Nowadays fibers go to the home, and ca 500 million kilometers is installed yearly. Source: Fiber Optic Cable Market - Growth, Trends, and Forecasts (2020 - 2025)

such as atoms, quantum dots, and molecules, is pursued worldwide for a variety of applications, which range from quantum information processing [15–17], to spectroscopy and microscopy of biophysical systems [18, 19]. Some of the most pressing questions in the physics of how to make single photons and single emitters interact are posed by the domain of information processing. Our current computers are electronic, and use conductive transistors to do calculations with electrons. Classical photonic computing has been proposed as a technology in which photons and their interactions would be used to compute [20, 21]. This would allow calculations to be done at the speed of light and could harness the very high bandwidth allowed by the PHz carrier frequency (versus GHz for electrons). Beyond classical computing with light to replace computing with electrons, quantum computing promises an exponential speedup of certain algorithms through operation on entangled quantum degrees of freedom [22–25]. In that domain, photons are expected to play a main role because they could either be used to perform quantum calculations directly or instead they would be needed to transport quantum information, with low loss, from one location to another, for instance in form of streams of single photons [26].

One of the first main obstacles for such technologies is that the interaction of light and matter is not, by itself, a strong interaction. Single emitters, i.e., quantum systems that can absorb a single pump photon, retain that energy for a characteristic time, and then emit the energy again in the form of another photon, are in principle readily available (See Section 1.3). However, a typical single emitter at room temperature, such as a dye molecule in the focus of a microscope objective, only absorbs of order 1 in 10^5 incoming photons [27]. Nanophotonics is a research field that seeks to realize optical microresonators and antennas to boost such weak interactions to unit efficiency [28]. These systems are used to both retain the light for a longer period of time around an emitter, as well as squeeze it into a smaller volume, which both boost interaction strength, as will be discussed below. Optical resonators serve both to enhance the absorption of incoming light and to accelerate the emission of single photons in preferred emission channels.

A second main problem that light-based technology faces is that, much unlike electrons, photons do not have any *direct* mutual interaction. Hence such interaction needs to happen *indirectly* via photonically active media. Making photons interact either requires classical nonlinear optical media or can be achieved in the regime of strong coupling of emitters to optical modes, which in turn requires strongly resonant nanophotonic environments. Through this mechanism, photons can be made to interact with each other, which paves the road towards quantum technologies. Beyond quantum technologies, the many advances in nanophotonics to harness single emitters also finds applications in biosensing [29, 30], where cavity and plasmonic resonances are harnessed for ultra-sensitive detection, such as improvements in blood glucose monitoring [31]. Other applications include Raman spectroscopy [32, 33], improvements in solar cells [34], and solid-state lighting [35, 36].

In the following parts, we will discuss the requirements of optical cavities and antennas on the one hand, and of single emitters on the other.

1.2. Optical cavities and antennas

1.2.1. Cavity-emitter coupling

When a single emitter, such as a dye or a quantum dot, is coupled to an optical mode of a nano- or microstructure, the altered electromagnetic environment can cause changes in the behavior of this emitter [37–39]. Starting with the most general case, we consider the transition of an emitter from an excited state $|e\rangle$ to a ground state $|g\rangle$. This transition is governed by the interaction Hamiltonian, which describes the interaction between the emitter and the emitted light. It is given by $\mathcal{H} = -\hat{\boldsymbol{\mu}} \cdot \hat{\boldsymbol{E}}$, where $\hat{\boldsymbol{\mu}}$ and $\hat{\boldsymbol{E}}$ are the operators for the transition dipole moment and the electric field operator, respectively. For many applications, it is desirable to couple the emitter in question to an optical mode. This can be achieved by bringing the emitter in close proximity to a structure that affects the optical modes available to the light. For optimal control, the structure can often be tuned to allow only a single optical mode. If the initial state $|i\rangle$ is that of an empty optical mode that transitions to the final state $|f\rangle$ when a photon radiates from an emitter, it holds that $\langle i|\hat{\mathcal{H}}|f\rangle = \hbar g$, where g is the emitter-field coupling rate. This is the rate at which the emitter and the optical mode exchange energy, given by

$$g = \mu\sqrt{\omega/2\hbar V \epsilon_0 \epsilon}, \tag{1.1}$$

where $\mu = |\langle g|\hat{\boldsymbol{\mu}}|e\rangle|$ is the dipole moment for the emitter transition, ω is the angular frequency of the field, V is the mode volume of the optical mode, ϵ_0 is the vacuum permittivity, and ϵ is the relative permittivity of the material surrounding the emitter. Eq. (1.1) shows that the coupling constant depends on two different contributions. Firstly, the *internal* electronic properties of the emitter itself are captured in the transition dipole moment μ. Secondly, the photonic properties *external* to the emitter are captured in the mode confinement V. This shows how the coupling rate can be tuned through the choice of emitter and the mode volume of the optical mode.

Excepting a peculiar family of recently discovered dark modes (bound states in the continuum [40]), optical modes generally outcouple to free space, which causes light from the mode to be lost. The rate at which light escapes from the optical mode to free space is given by κ. Providing that emitter losses are low, comparing the coupling and loss rates of a mode to each other shows that there are two distinct regimes of light-matter interaction, namely weak and strong coupling.

Weak coupling

If $2g < \kappa$, the emitter is weakly coupled to the optical mode, and the energy of the light from the emitter is lost more quickly by the optical mode than it is exchanged between the mode and the emitter. The onset of the strong coupling regime is where $g > 2\kappa$, where the energy is exchanged between the optical mode and the emitter faster than it is lost. Strong coupling of single emitters is an essential ingredient in quantum cavity electrodynamics (CQED), as it allows for coherent control over the emitter. Instead, the weak coupling limit is the operation regime of choice for single-photon sources, where the photonic environment is used to accelerate spontaneous emission. In the weak coupling limit, the decay rate of the emitter is dictated by Fermi's golden rule [41]

$$\Gamma = \sum_f \frac{2\pi}{\hbar} \left|\left\langle i|\hat{\mathcal{H}}|f\right\rangle\right|^2 \delta(E_i - E_f), \tag{1.2}$$

where $\langle i|$ and $|f\rangle$ are the initial and final state, respectively, $\rho(E_f)$ is the density of states at the energy E_f of the final state. $\hat{\mathcal{H}} = -\hat{\boldsymbol{\mu}} \cdot \hat{\boldsymbol{E}}$ is the interaction Hamiltonian, with $\hat{\boldsymbol{\mu}}$ and $\hat{\boldsymbol{E}}$ the operator for the dipole moment of the optical transition, and the electric field, respectively. Fermi's Golden Rule ccan be rewritten to

$$\Gamma = \frac{\pi\omega}{3\hbar\epsilon_0} |\mu|^2 \rho_\mu(\omega, \boldsymbol{r}), \tag{1.3}$$

where $\rho_\mu(\omega, \boldsymbol{r})$ is the local density of states (LDOS), which gives the number of states per unit volume and unit frequency, at frequency ω, available to an emitter at position $\boldsymbol{r}$, oriented along $\boldsymbol{\mu}$. The concept of optical density of states for light was introduced in the 1990's by Sprik et al. [42], although the fact that electromagnetic environments can modify spontaneous emission through its mode structure was already known since seminal experiments by Drexhage in the 1970s [43, 44].

The emission enhancement of an emitter coupled to an optical mode is governed by the ratio of the LDOS at the position of that emitter and the emitter decay rate in a homogeneous medium γ_{hom}, where $\rho_{\text{hom}} = \omega^2 n/(\pi^2 c^3)$. At the resonance of the optical mode in question, the LDOS is given by $\rho_{\text{c}} = 6/\pi\kappa V$, and the decay rate becomes $\Gamma_{\text{c}} = 4g^2/\kappa$. This is used to find the Purcell factor

$$F_{\text{P}} = \frac{\Gamma_{\text{c}}}{\gamma_{\text{hom}}} = \frac{4g^2}{\gamma_{\text{hom}}\kappa} = \frac{3}{4\pi^2} \frac{Q}{V} \left(\frac{\lambda}{n}\right)^3, \tag{1.4}$$

with which one can characterize the enhancement of emission by the emitter w.r.t the emission in a homogeneous background. Here, Q is the quality factor of the optical mode, given by $Q = \omega/\kappa$, which dictates the linewidth of the optical mode, and (λ/n) is the wavelength of the optical mode in the surrounding medium. It was first stated by Edward Mills Purcell in 1946 [45]. The first observed LDOS control of spontaneous emission was not reported for cavities, but instead for emitters near reflective interfaces by Karl Drexhage. He placed emitters at a well-controlled distance from a metal mirror by preparing thin layers on the order of the wavelength of light in thickness. Through constructive and destructive interference, the modes at a metallic interface form standing waves, leading to a characteristic oscillatory LDOS as a function of distance to the mirror. An alternative viewpoint on this LDOS effect is that the emitter interferes with its own mirror image which alters the LDOS [43, 44]. At present, it has both been possible to suppress spontaneous emission (ca 20—70-fold slow down of emission in photonic band gap [46]), to accelerate it up to c.a. 30-fold in dielectric microcavities by the Purcell effect [47–52], and to accelerate it ca. 500 fold in plasmonic nano-junction systems [53].

Strong coupling

If the coupling rate between an emitter and optical mode increases to $2g > \kappa$, excitations are exchanged between the emitter and the mode faster than it is lost to free space. This regime is called strong coupling. In the above section, we have neglected losses by the emitter itself, assuming these to be small with respect to κ. Including these in the strong coupling condition, it reads $2g > \kappa + \gamma + 2\gamma^*$, where γ and γ^* are the emitter's decay rate and dephasing rate, respectively [54, 55].

In the strong coupling regime, Fermi's Golden Rule no longer applies [56], as spontaneous emission is no longer irreversible. Instead, after excitation of an emitter,

vacuum Rabi oscillations can be observed, signaling the coherent evolution of the light-matter system [57]. Alternatively, the strong coupling between an emitter and an optical mode can be seen in the splitting of modes in emission spectra [58]. This splitting is the hallmark of mode hybridization, where the eigenstates of the system are a hybridization of light and matter excitation. These phenomena were first observed for single atoms in highly excited states (Rydberg atoms) in interaction with mm-wave radiation by the group of Haroche, who received the Nobel prize in 2012 [59]. Generally, outside the strong coupling regime, photons do not interact with each other. However, the hybridized eigenmodes of strong coupling, also known as dressed states, can be used as a pathway to allow for photon-photon interactions. Since photon-photon interactions are an essential element of quantum photonic networks, it is unsurprising that since the advent of CQED, it has been a holy grail to push single-emitter strong coupling from the domain of atomic physics to the domain of solid-state integrated nanotechnology. This domain can currently be reached at cryogenic temperatures in III-V resonators. Recently reports have appeared that claim room temperature strong coupling of single emitters in plasmonic nanojunctions [60].

1.2.2. Optical cavities

The crucial parameters for light-matter interaction are the electromagnetic mode confinement V that determines the coupling strength g, and the cavity quality factor Q that determines the cooperativity, or Purcell factor. When it comes to creating the desired optical modes, a common approach is designing and fabricating optical cavities. These can be as simple in design as two opposing mirrors forming a Fabry-Pérot etalon, or more elaborate, such as ring resonators [61, 62], which come in variations of microdisks, - toroids, and -spheres. Photonic crystals utilize a periodic structure in the dielectric to create a band gap. A defect in the crystal can then function as another type of cavity, as demonstrated for 2D photonic crystal cavities [63, 64] and photonic crystal micropillars [65].

Microcavities are commonly used in applications for light trapping [66]. As many dielectrics have a frequency window with low losses, impressively high Q factors can be achieved with these structures, with experimental values up to 10^7 being reported for photonic crystals [67], and even much higher, up to 10^9 for larger cavities such as microtoroids and microspheres [68, 69]. The high Q allows the light to be trapped for many oscillation cycles, which corresponds to a long interaction time with the emitter. This makes high-Q cavities a common approach for single-emitter strong coupling. However, they suffer from a major drawback. The confinement in these cavities is bounded by the diffraction limit. Thus, in most cases dielectric cavities can't achieve mode volumes smaller than $\sim (\lambda/n)^3$ [70], although there have been reports of some cavity designs that push this limitation by an order of magnitude by employing narrow slits between high refractive index materials [71–73].

Emitter placement is crucial when it comes to achieving high coupling rates between cavity and emitter. Since optical modes decay evanescently away from the cavity, placing the emitter as close as possible to the field maximum is paramount. In particular, there have been demonstrations of micropillars [74–76] and 2D photonic crystal cavity systems [63, 67, 77–81] fabricated from materials that contain intrinsic color centers or embedded quantum dots, such as InGaAs or GaN . Typically, for these systems, the emitters of the system need to be located inside of it and the cavity fabricated around it.

Figure 1.1: Examples of common dielectric optical microcavities, with photonic modes illustrated in red. (A) a Fabry-Pérot cavity, (B) a microtoroid, (C) a microdisk, (D) a microsphere, (E) a planar photonic crystal cavity, and (F) a photonic crystal micropillar. The red color denotes the region where the light is trapped.

Alternatively, there have been efforts to place single emitters at a high-LDOS location of a pre-existing cavity using atomic force microscopy (AFM) [82] and electron beam (e-beam) lithography [83–85]. Due to the unresolved issue of precisely placing emitters in cavities, measured Purcell factors (measured values up to $F_P \approx 20$) have fallen short of theoretical expectations [48, 49, 51, 52].

1.2.3. Optical antennas

In terms of merits, optical antennas are the precise opposite of optical cavities [86]. Where optical cavities are usually made of dielectrics with high refractive indices, optical antennas are typically made of metals, which have a negative dielectric constant. Their metallic properties have a number of implications. Incoming light can excite resonant oscillations of the free carries of the metal to oscillate in the antenna, a phenomenon known as plasmon resonance [87]. This light can be lost either in the form of heat of the antenna due to Ohmic resistance inside the antenna or to its surroundings as a photon. Whereas dielectric optical cavities, being optimized for their high quality factors, intrinsically avoid absorption altogether, and need 10^3 to 10^9 optical cycles to couple light to free space, optical antennas lose their excitation within 10 to 40 optical cycles. Cavities poorly couple to free space, whereas antennas are excellent far-field transducers. Plasmon antennas will often behave like induced electric dipole resonators, which are well-known for their strong radiation to the far-field. Coupled with the inherent losses of metals, optical antennas are severely limited in their quality factor. At typically targeted values $Q \sim 20$ half the excitation is lost as heat, and half as far-field radiation.

The plasmonic properties and negative dielectric constants of optical antennas have a salient benefit, however. They allow for light to be squeezed into much smaller mode volumes than their dielectric cavity counterparts. Antennas with highly sub-

Figure 1.2: Examples of optical antennas, where the maximum of the plasmonic mode is shown in red. Emittes placed here will have optimum coupling to the antenna. (A) a phased array antnna, specifically a Yagi-Uda, (B) and (C) show dipole antennas, where (C) is a a dimer anteanna. (D) shows an example of patch antenna: in this case, a nanocube on a metallic surface.

wavelength gaps are typical examples of this [53, 88, 89]. It is important to note here that the usual definition of mode volume, which is applicable to high-Q cavities, does not apply to low-Q plasmonic antennas. A more complete discussion can be found in [90–93].

Many different designs of optical antennas have been proposed and utilized over the years. They roughly fall within the following categories: Phased-array antennas [83, 94], dipole antennas [95], and patch antennas [60, 96, 97]. Of these, the phased-array antennas have the largest mode volume, and patch antennas the smallest, owing to their capacity to achieve extremely narrow gaps. A number of examples are shown in Fig. 1.2. The low Q paired with a small V of optical antennas allows for a strong but short-lived interaction with an appropriately placed emitter, with the highest LDOS achieved when an emitter is brought close to the metal interface. Patch plasmonic antennas realized by placing monocrystalline nanocube and nanosphere particles on ultrasmooth metallic films with just a nanometer thick spacer as separation have been claimed to reproducibly provide mode volumes as small as $(\lambda/n)^3/10^5$. Observed Purcell enhancements are of order 300 to 500 times [53].

1.2.4. Optical hybrids

In recent years, several groups have suggested to realize plasmonic-photonic hybrid resonators with the goal to achieve structures that combine the merits and remove the drawbacks, of plasmonics on the one hand and dielectric cavities on the other hand. In other words, dielectric cavities are limited to poor confinement V, but allow flexible control of Q. Conversely, optical antennas can achieve atomic-scale V, but are severely limited in Q. Hybrids may combine subwavelength confinement and high Q. The first experimental demonstration of optical hybrids was reported by Won-Tae Kim et al. in 1999 in the context of Raman scattering [98]. Apart from their use in Raman scattering [99–101], over the past two decades, many more designs of optical hybrids have been proposed and realised, such as whispering-gallery-mode hybrids [102–108], photonic crystal cavities [109], sensing [110], optical trapping [111–114], nanolasers [115, 116], and near-to-far field coupling [117–121].

Recently, it was shown by Doeleman et al. that optical hybrids can be modeled as a system of coupled harmonic oscillators to derive the LDOS in a hybrid system on basis

Figure 1.3: (A) The homogeneous, uncoupled antenna polarizability (blue), and the cavity response function (green). (B) The hybrid polarizability, when the two are coupled together. We show the bare antenna polarizability in the background, as a comparison.

of the resonance of the constituents [122]. We will discuss this model in more depth in Chapter 2, but in essence, this model considers both the antenna and the cavity to be harmonic oscillators with a Lorentzian response as a function of angular frequency ω, which is proportional to

$$\frac{1}{\omega_0^2 - \omega^2 - i\omega\gamma},$$

(1.5)

where ω_0 is the resonance frequency and γ is the loss rate of the cavity or antenna. For the cavity, this loss rate is typically low, much smaller than the resonance frequency. For the antenna, this loss rate is typically much larger, with $\omega_0/\gamma \sim 10$. This is illustrated in Fig. 1.3A, where we plot the antenna and cavity response functions. We see that the high-Q cavity has a narrow response, and the low-Q antenna has a broad one. If the two are coupled, the resulting combined response will have a Fano lineshape, as shown in Fig. 1.3B. This lineshape shows that, depending on the frequency, the hybrid system can have a more cavity-like response, or a more antenna-like response. In most cases, it is the cavity-like response that is of interest. It presents a resonance with a quality factor that is close to that of the cavity (usually spoiled to be lower by a factor 2 – 10 by the perturber), and at the same time can provide an optical mode volume that can be 1 to 2 orders of magnitude below that of the cavity. Similarly to the polarizability, the LDOS will take on a Fano lineshape.

Qualitatively, this formalism predicts that the LDOS that quantifies spontaneous emission enhancement may indeed show high-Q, yet low V resonances by hybridization of antenna and cavity, which expresses as LDOS Fano resonances. Evidencing such Fano lines is highly challenging. Intensity-based measurements of enhancements are challenging because both the LDOS *and* the directionality of emitted light, i.e., the collection efficiency of photons, simultaneously display Fano resonances. Decay rate measurements would avoid this, but are hampered by the fact that matching the spectra of emitters and LDOS resonances is challenging. Finally, realising hybrids with precisely placed emitters is a hard nanofabrication challenge. Due to these difficulties,

only a few demonstrations of LDOS measurements have been reported so far, with experimental LDOS enhancements up to a few factors of 10 reported [82, 85, 123].

1.3. Single photon emitters

In the following section we will discuss basic properties of single emitters, and discuss the main single emitter systems that are currently studied in the solid state.

1.3.1. Single emitters as two-level systems

The textbook example to demonstrate the basic physics of a single emitter is to consider a single atom, even abstracted to a single two-level system [27]. The basic idea is that a two level system with ground state $|g\rangle$ and excited state $|e\rangle$ split by an energy $\hbar\omega$ may transition between the two states by absorbing and emitting single photons of frequency ω. After excitation, according to Fermi's Golden Rule, the atom will typically remain in the excited state for some characteristic time before re-emitting another photon at a rate of γ_0. For a two-level system, the emission rate, the oscillator strength, and the transition dipole moment are equivalent indicators of the strength of the transition. Indeed, the strength of the transition is quantified by the emitter's oscillator strength

$$f = \frac{2\mu^2 m_e \omega}{\hbar e^2} \tag{1.6}$$

where m_e and e are the electron mass and charge, respectively, and where the transition dipole moment can be extracted from the decay rate according to

$$\mu = \sqrt{\gamma_0 \frac{3\pi\epsilon_0 \hbar c^3}{n\omega^3}}. \tag{1.7}$$

The spectrum of emitted light in this system is Fourier-limited, meaning it is Lorentzian with width set by γ_0.

1.3.2. Single emitters as multi-level systems

Single atoms may be near-perfect two-level systems, but they are far from practical to implement, as one requires vacuum systems and atomic physics laser cooling equipment to hold atoms still in (macroscopic) cavities. Hence solid-state single emitter systems are of large interest. However, all solid-state emitter systems have significant drawbacks with respect to a perfect two-level system due to the presence of a complex and dynamically changing environment. An excellent qualitative understanding of emitter physics beyond the two-level system is presented by the so-called Jabłoński diagram, developed originally for molecular fluorophores by Aleksander Jabłoński in 1933 [124]. An example of this is shown in Fig. 1.4A. The main idea is that each of the electronic ground and excited states, typically separated by a few eV in energy, are in fact manifolds of more closely spaced (10s of meV) (ro)vibrational levels. At room temperature, molecules typically reside in the vibrational ground, or at most, first excited, state. By exciting with by pump light, it is possible to access the vibrationally excited states in the first electronic excited state manifold. After excitation, the vibrations typically cool within picoseconds, leaving the emitter in the vibrational ground state of the first electronic excited state to decay back down. This decay is

Figure 1.4: (A) Jablonski diagram of different photonic interactions of a multi-level system, showing the initial excitation (blue), the intraband relaxation (red), and fluorescence from the excited to the ground state (green). Also shown is the phosphoerescence pathway via the triplet state (orange). (B) The absorption (blue) and emission (green) of Rhodamine 6G. The shift between the two peaks is called the Stokes shift.

then possible to a multitude of the vibrationally excited electronic ground state levels. In this picture, molecules have wide absorption bands, extending to higher frequency from the zero-phonon transition, while fluorescence likewise shows a broad spectrum, but extended towards the red, not the blue. The fact that emitted photons are almost inevitably red-shifted relative to the pump due to the loss of excitation into motion, is expressed in the Stokes shift, named after George Gabriel Stokes [125]. An example is shown in Fig. 1.4B, where we show the absorption and emission spectra of a commonly used organic fluorophore, namely Rhodamine 6G [126].

The width of the emission line of emitters at room temperature is thus greatly different from that of perfect two-level systems. Due to the uncertainty principle, the lower limit on the width of the emission spectrum of a single emitter is simply the fluorescence decay rate $\Delta\omega = \gamma_0$, also known as homogeneous linewidth. An emission line that has this width is called lifetime limited. Since typical emitter decay rates are nanoseconds, homogeneous linewidths are in the MHz to GHz range [27, 127]. Instead for room-temperature emitters in the solid states, spectra have linewidths of 20 to 100 nm (10s of THz), due to the vibrational structure of emitters expressed in Jabłońskis' diagram, which is further affected by a variety of thermal broadening effects. Thereby the inhomogeneous linewidth of emitters far exceeds the homogeneous linewidth. The linewidth of an emitter is closely connected to another quantity known as dephasing [128]. The dephasing rate is the sum of all decay processes, both radiative and non-radiative, and is the rate at with a quantum state of a given emitter is lost. For quantum optics, it is favorable for the dephasing rate to be as low as possible, as this allows for time for the emitter to couple to the optical mode for longer, facilitating interaction.

1.3.3. Techniques to probe single photon emitters

Since in this thesis we will present measurements on single emitters, we discuss the most common techniques to probe the behavior of single emitters. We discuss these

Figure 1.5: An schematic of the setup used in this work, which is an example of a typical TCSPC measurement setup. It shows the pump laser, the reference diode, and the emitters under study. The red-shifted light is separated from the pump using a dichroic mirror and cleanup filters. Using a flip mirror, the emission can be directed towards either the spectrometer or the photodiodes.

techniques using as illustration the main instrument used in our own work shown in Fig. 1.5, which is a confocal microscopy setup that performs time-correlated single-photon measurements (TCSPC) and spectroscopy. In these measurements, an emitter is typically pumped with a blue or green pulsed laser that provides a train of picosecond pulses at MHz repetition rates. As the fluorescence light is red-shifted with respect to the pump light, it can be separated from the pump light by a dichroic mirror and further cleaned up using emission filters, which are engineered to have maximum transmission of the emission light, while maximally blocking the pump frequencies. By combinations of filters, rejection contrasts of $> 10^8$ can be reached. This technique allows the low light levels from single-photon emitters, typically 10^3-10^6 photons/second, to be distinguished against the pump light, which is often upwards of 10^{12} photons/s.

Using a flip mirror, the emission can be sent to a spectrometer for spectral analysis. Present-day spectrometers are equipped with CCD or CMOS cameras that can have near-unity quantum efficiency ($> 95\%$ probability of converting an impinging photon to a photoelectron can be reached) and very low read noise. Alternatively, photons are counted on avalanche photodiodes (APDs). These also are available at near-unity quantum efficiencies, and when operated in Geiger counting mode can report accurately on the arrival time of each detected photon, with a time resolution of several 10's to a few 100 picoseconds, though timing accuracies of >3 ps have been reported using superconducting nanowire detectors [129]. With accurate photon arrival information, the decay rate of an emitter can accurately be determined. When using a pulsed laser, correlating the photon arrival times with the firing times of the laser, a histogram can be built of photon delay times. For usual emitters, the probability of photon emission follows an exponential decay law as a function of time. It is given by

$$P = \mathrm{e}^{-\gamma t} = \mathrm{e}^{-t/\tau}, \tag{1.8}$$

where γ is the total decay rate of the emitter, which is the sum of radiative- and

non-radiative decay rates. For a lifetime measurement, only a single APD is needed. However, single emitter microscopes typically use a set of two detectors, separated by a 50/50 beam splitter. This configuration is known as a Hanbury-Brown Twiss setup and is required to assess the single-photon purity of an emitter can be determined. For a single photon emitter, photons are expected to exhibit 'antibunching', meaning that two photons are never emitted simultaneously. Measuring this antibunching on a single APD is technically impossible due to the dead times of APD detectors. An APD is essentially a photodiode biased at a voltage close to breakdown, such that a single photon can trigger an electron avalanche. After it fires, it takes several tens of nanoseconds to build up to this voltage again, during which no photon events can be registered. Using two APDs overcomes the fundamental limitations of a single APD for measuring the temporal correlation between photons for times shorter than the APD dead time. The single-photon purity is found from the time-correlation between the photon event data streams on the two APDs, known as $g^{(2)}(\tau)$. The threshold for antibunching and single-photon emission is $g^{(2)}(0) < \frac{1}{2}g^{(2)}(\tau_L)$, where τ_L is the repetition time of the laser. The best single-photon emitters, with the highest purity, have $g^{(2)}(0) = 0$. Using this criterion, emitters can be selected based on single-photon purity.

1.3.4. A portfolio of emitters

In this section we briefly summarize salient properties of solid-state emitter systems. This thesis is motivated by the desire to establish the potential of hybrid plasmonic-photonic resonators for controlling light-matter interaction, preferably with emitters at room-temperature. We will discuss the properties of main categories of emitters, as well as their benefits and disadvantages for their integration with hybrid plasmonic-photonic resonators.

Color centers

Color centers are atomic-scale defects, typically in large band gap semiconductors. They can be made of inclusions, vacancies, or exchanges of atom positions. The most well-known example is the nitrogen-vacancy center (NV-centre) in diamond, which consists of a vacancy in the diamond lattice adjacent to a nitrogen atom [41]. By virtue of the wide band gap semiconductor surroundings, such a color center can have atom-like electronic states that provide for highly reproducible, stable, and sometimes even narrow linewidth emission. For instance, the NV center in diamond, when filtered on the zero phonon line, can be a source of indistinguishable photons in quantum optics experiments [130, 131], and has proven almost indefinitely photostable under excitation [132–134]. Further, the main property driving the fame of the NV center is that the emission provides a handle to read out and control electron spin. A drawback to color centers is that they tend to have significant spectral broadening due to phonons (lattice excitations of the surrounding medium). Even at low temperature the Fourier-limited zero-phonon-line of the NV center only contains a few percent of the total emitted intensity. This limitation is partially overcome by, for instance, the silicon-vacancy in diamond, and is further amenable to photonic engineering by embedding a color center in optical cavities that provide a high Purcell factor at the zero-phonon line. While major strides have been taken in this regard [135–138], color centers are almost exclusively amenable to monolithic integration: in bulk diamond color centers have

Figure 1.6: Examples of single-photon emitters. (A) A schematic of an NV^- center, where two atoms from a diamond lattice are displaced, and a single Si atom takes their place. (B) The base structure of Rhodamine, a family of fluorescent molecules. (C) A cartoon of a quantum dot.

excellent properties, but once they are within 10s of nanometer distance from surfaces, they tend to become unstable and lose quantum efficiency. While nanodiamonds with NV centers are an effective field of study, they have not shown emissive qualities that compare well with those of NV centers in bulk diamond [139, 140]. Integration of such nanodiamonds in nanostructures has been demonstrated with the very challenging technique of AFM-based nanomanipulation [141].

Organic molecules
Organic molecules, particularly those with conjugated carbon chains, and polycyclic aromatic molecules, can show a strong absorption, usually in the visible or UV range, and can, depending on their particular chemical composition, be controlled to emit photons throughout the visible and near-infrared regime [41]. Their nanosecond lifetimes and high quantum yields are furthermore of large use for making bright emissive nanosystems. A wide variety of organic emitter families have been realized, optimizing for solubility in a variety of solvents, and for excellent emission quantum efficiency, particularly in solvated settings and in polymer matrices [142–144]. Due to the ease of solution based handling they have been an emitter of choice for deposition on plasmonic systems. However, they are intrinsically very hard to deterministically and selectively position in desired locations. Molecular dyes are particularly sensitive to a process called bleaching. This generally occurs when a molecule is exposed to a strong electromagnetic field, such as a pump laser, which can cause irreversible structural changes to the molecule, effectively killing their luminescent properties due to photo-activated chemical reactions with the surrounding.

As dye molecules are structurally identical, it would seem that they would offer a perfectly reproducible emitter. However, their properties are highly sensitive to the electronic and vibrational structure of their surroundings [143]. Some molecular dyes are known to have markedly improved behavior when embedded in very specific crystalline molecular host matrices. An example is dibenzoterrylene (DBT) in anthracene. This system, and a few similar molecule-host combinations become near ideal quantum systems at liquid helium temperatures, with Fourier-limited spectra [144–146]. Integration of these molecules with nanophotonic structures is a main challenge, first due to the challenge of aligned nanofabrication, and second because it appears the optical properties of DBT are adversely affected by proximity to interfaces. A recent

breakthrough by Pazzagli et. al is the development of nanoparticles with DBT that show advantageous emission properties [147].

II-VI semiconductor quantum dots
Colloidal quantum dots are typically 3 to 10 nm size crystallites of a semiconductor material. As the size is below the exciton Bohr radius, the excitons in them can experience quantum confinement effects, with luminescent discrete levels that can be tuned by quantum dot size. In analogy to the seminal 'particle in a box' model, the diameter of a quantum dot will determine the discrete exciton energy states. Thereby the size of quantum dots tunes their emission. Noteably CdS and CdSe quantum dots can show strong absorption cross sections for photoexcitation, and near-unit quantum efficiency emission properties in the visible regime, with nanosecond emitter lifetimes. Furthermore, quantum dots can show sustained photon count rates over 10^7 per second [148].

Instrumental for good performance of such quantum dots is control of the surface, typically obtained by growing multiple shells of larger-band gap materials over the core. Such shells reduce the susceptibility of quantum dots for intermittent emission behavior. Intermittency is also known as blinking, and appears as a random switching on and off of emission. While also organic fluorophores are known to blink, blinking behavior has gained major interest in the scientific community for II-VI quantum dots [149–153]. A number of different mechanisms for blinking have been proposed, and it is generally agreed upon that even a single atomic defect, and exchange of even a single charge carrier of the quantum dot with its environment, can cause a II-VI quantum dot to switch from bright to dark.

As regards integration of colloidal quantum dots with nanophotonic devices, a major benefit is that quantum dots can be functionalized through ligand exchange to have organic ligands that can selectively bind to prefunctionalized sites. This feature has recently been leveraged for one of the first demonstrations of light-matter coupling control in hybrid plasmonic-photonic systems [85]. Unfortunately II-VI quantum dots have disadvantageous dephasing properties, meaning that they have strong inhomogeneous broadening at both room and cryogenic temperatures [154, 155].

Perovskite quantum dots
Very recently a new type of quantum dot has gained considerable attention for its remarkable properties. This category is that of lead halide perovskites, a highly ionic material with a chemical composition of $MAPbX_3$ or $CsPbX_3$, where MA stands for the methylamonium cation, and X stands for either Cl, Br, I, or a combination thereof. Organic-inorganic perovskites, as well as all-inorganic lead halide perovskite materials are of large interest for emissive devices, as well as for photovoltaics. In 2015 Protesescu [156] reported the first realization of cesium lead halide perovskite quantum dots. Emission properties can be tuned both through the size and the chemical composition, which together provides coverage of the entire visible range. Dots containing Cl^- on the anion locations have the highest energy emission, as short as $400\,nm$. Bromide mixtures span the region from $460 - 520\,nm$, where I^- ions at the anion position yield emissions around $600 - 700\,nm$. To further tune the emission frequency, the anions can be mixed and the size of the particles can be controlled.

Lead halide perovskites have been reported to boast high quantum yields, reduced photobleaching, high spectral purity, and a wide color gamut, which makes them

an interesting candidate for single-photon emitters [156–159]. Moreover, lead halide perovskites are generally understood to be highly defect-tolerant, which may be expected to combat quantum dot intermittency. However, as they are made from ionic materials, these materials are sensitive to ambient humidity, which can cause their rapid degradation when left exposed. Additionally, they lack the surface functionalization developed for more traditional quantum dots, and thus have a reduced affinity to surfaces, making them difficult to place inside a plasmonic or photonic cavity system.

1.4. This thesis

One of the main goals of cavity quantum electrodynamics is gaining maximum control over an emitter placed inside an optical resonator. In order to achieve this high level of control, there are three main building blocks that need to be completed. Firstly, an optical system is required that has an appropriate quality factor, mode volume, and Purcell factor, such that the light from the emitter can be adequately funneled into the cavity. For this, a cavity with a similar linewidth to the emitter is preferable, and small mode volumes are needed to confine the light and achieve a high Purcell factor. Secondly, a high level of control over the emitter is needed, as well as an understanding of its behavior. It is necessary to use the best possible tools to analyze the properties of emitters, so that an emitter can be selected that has a high single-photon purity, as well as the appropriate emission dynamics. The final challenge then lies in combining the previous two, where a single emitter is placed accurately inside an optical resonator.

This thesis was conceived to follow these three steps, and was motivated by the proposition that hybrid plasmonic-photonic resonators could be ideal candidates for matching resonators to emitters at even room temperature, or at least liquid nitrogen, instead of liquid helium temperatures. Realizing this promise would be an important breakthrough for nanoscale quantum optics. The main focus of Part I of this thesis is the design and fabrication of nanophotonic hybrids with high performance, i.e., with high confinement, quality factor, and light-matter interaction strength. In parallel, we have explored the physics of a particularly promising emerging solid-state emitter system, which is reported as Part II of this thesis. Developing each of these parallel threads has involved large challenges, and the full route from start to the ultimate goal of incorporation of these emitters inside the nanophotonic resonators is not commensurate with a four-year research project. However, this work reports on main ingredients to build towards this goal.

Part I: Plasmonic-photonic hybrids and their performance limits

Chapter 2 discusses the properties of photonic crystal nanobeam cavities when coupled to plasmonic dimer gap antennas. We show that the combined system has a number of advantages over the individual constituents, and also over using single plasmonic dipole antennas without gap. We present generalized design rules for optical hybrids, as well as the requirements of the optical system as well as the emitters for a set of applications. The discussion includes strong coupling, weak coupling bright single-photon sources, and single-photon indistinguishability.

In Chapter 3 we describe the experimental realization of the systems discussed in Chapter 1. We focus on a modified photonic crystal nanobeam design as compared

to Chapter 2 that can be fabricated from high quality Si_3N_4 on glass, instead of being a suspended nanobeam in air. We report on the complex multistep electron beam lithography fabrication procedure of the nanobeams as well as of gold nanogap antennas, and describe the full nanofabrication process to create hybrid resonators in which nanobeams and antennas are precisely aligned to each other.

Chapter 4 discusses optical experiments on hybrid resonators composed of Si_3N_4 nanobeams on glass combined with gold gap nanoantennas. We use dark field spectroscopy to quantify the antenna resonances, and developed a tuneable diode laser spectroscopy method to identify nanobeam resonances. For the hybrid systems we examine the cavity quality factors and resonance frequencies as a function of the antenna size and tuning, evidencing the effect of hybridization through perturbation of the cavity mode. To gain access to local density of states modifications we use a Raman-active species as reporter. Surface Enhanced Raman Scattering (SERS) signals allow to decouple pump field enhancements and LDOS effects. For both types of effects we identify Fano lineshapes.

Part II: Intermittent behavior and single perovskite quantum dot emitters

In Chapter 5 we discuss a method to characterize the intermittency properties of single emitters. We create a toolbox that uses Monte Carlo methods to simulate multi-level, intermittent single emitter behavior with a wide array of properties. The same toolbox uses changepoint analysis, a Bayesian clustering algorithm, and correlation calculations to analyze both simulated and measured data. We use the simulated emitters to benchmark the analysis and determine the boundaries within which conclusions can be drawn from the available data.

In Chapter 6 we apply changepoint analysis to a measured set of $CsPbBr_3$ quantum dots. We find that the emitters under study exhibit a large number of gray states, where we find an inverse correlation between brightness and decay rate. Additionally, we use the segments found through changepoint analysis to investigate memory effects and aging behavior of these emitters.

I

Plasmonic-photonic hybrids: Their performance and their limits

2

Hybrid cavity-antenna systems for quantum optics outside the cryostat?

> "'Take some more tea,' the March Hare said to Alice very earnestly.
> 'I've had nothing yet,' Alice replied in an offended tone, 'so I can't take more.'
> 'You mean you can't take less,' said the Hatter: 'it's very easy to take more than nothing'"
> — Lewis Carroll, *Alice's Adventures in Wonderland*, 1865

2.1. Introduction

MICROCAVITIES are a key building block for all branches of optics, and over the last thirty years, their development has been a mainstay of micro- and nanoscale optics research efforts. For instance, vertical-cavity surface-emitting lasers in active III-V semiconductor systems are a key technology for optical interconnects in information processing in any data center. At the same time, the narrow spectral lines of microcavities are key for label-free sensing down to the level of single proteins [160], and for metrology of distances down to the picometer scale [161]. In quantum optics, microcavities are particularly sought after for their ability to turn intrinsically slow and isotropic emitters into directional and fast single-photon guns [27, 162, 163] and even to bring quantum emitters into strong coupling regimes where spontaneous emission is replaced by quantum entanglement of light and matter [56]. Notwithstanding the large diversity of microcavity designs, generally they are characterized by two figures of merit: the first relates to *temporal* and the second to *spatial* confinement. The quality factor Q measures the time duration for which light is stored in the resonator in units of optical cycles, while the mode volume V is a measure for how tightly light is confined in three dimensions. For actual applications, generally algebraic combinations of Q and V determine performance. For instance, the most well-known figure of merit for a cavity is undoubtedly the Purcell factor for spontaneous emission enhancement in a cavity, which reads [45]

$$F_{\mathrm{P}} = \frac{3}{4\pi^2} \frac{Q}{\tilde{V}}.$$ (2.1)

where $\tilde{V} = V/(\lambda/n)^3$ is the mode volume expressed in units of wavelength cubed in the medium of interest. This factor quantifies the Local Optical Density of States (LDOS) at resonance [41, 42], and is thereby fundamental for many light-matter interactions, such as creating desirable single-photon sources [27] or building sensors for analytes that have optimum sensitivity [69, 160]. However, for other applications, ranging from spectral filters to optical memories, building lasers [164], enhancing nonlinear optical effects [165], and cavity QED [28, 56], other metrics apply. For instance, in cavity QED the so-called regime of strong coupling is sensitive to $Q/\sqrt{\tilde{V}}$ [56].

Approximately 15 years ago, plasmonics was proposed as the solution for the main perceived drawback of classical dielectric microcavities: whether in form of microdisks, micropillars, or photonic crystals, dielectric cavities are limited in achievable confinement to approximately the diffraction limit, meaning that target performances tend to require minimum quality factors (typically $Q > 10^4$). In contrast, plasmonic resonators have stellar confinement, but *exceptionally* poor Q in the order of 10. Indeed, plasmon nanoantenna resonators have been reported that provide measured Purcell factor up to 10^3 [96, 166, 167], and recent claims are that plasmon antennas allow strong coupling with single emitters at room temperature, as opposed to at cryogenic temperatures as achieved in microcavities [60, 89]. Figure 2.1 shows quality factors and mode volumes of dielectric microcavities and plasmonic antennas that are at the state of the art. The striking observation is that even if similar Purcell factors are possible (constant F indicated by diagonals in the diagram), there is a huge gap between nano-antennas and microcavities.

Apparently, reaching *intermediate* $Q/\tilde{V}$-values, where one trades in part of the plasmonic confinement in favor of higher Q is extremely difficult. This is unfortunate,
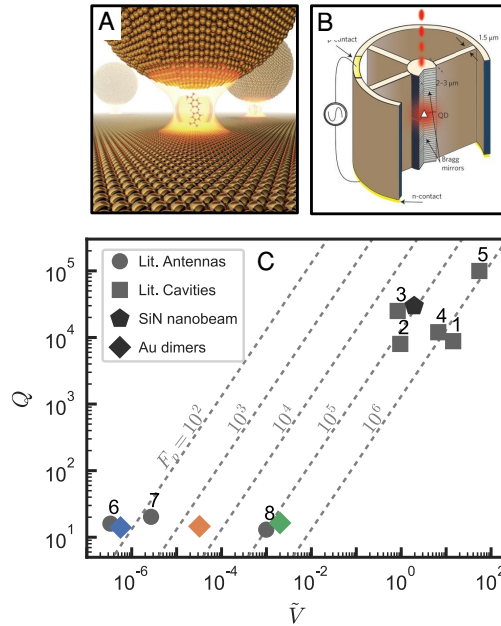
Figure 2.1: Reported quality factors and mode volumes for dielectric cavity and plasmonic antenna systems. (A, B) Renderings of state of the art plasmonic [60] (A) and dielectric [168] (B) microcavities. (C) Q and $\tilde{V}$ values for state-of-the art literature cavities (1: [65], 2: [63], 3: [64], 4: [61], 5: [62]) and antennas (6: [60], 7: [89], 8: [96]). Also shown are values for the cavity and antennas used in this Chapter, shown in black — a silicon nitride (SiN) nanobeam cavity and gold dimer antennas with dimer gaps of 1 (blue), 5 (green) and 25 nm (red). Dashed lines indicate lines of constant Purcell factor F_P. This image is reproduced from our paper Ref. [169] under CC-BY license. To provide image credits for panels (A) and (B) we here quote literally the image credits listed for them in Ref. [169]: *Panel (A): Reprinted by permission from Springer Nature [60], Copyright 2016. Panel (B): Reprinted by permission from Springer Nature [168], Copyright 2016.*

as plasmonic quality factors ($Q = 10$), and the fact that they can not be controlled at will, can hardly be classified as a practical proposition for many envisioned applications. Moreover, the extremely low antenna mode volumes require exceptional control over the spatial alignment of the emitter and antenna. On the other hand, the extremely narrow linewidths of high-Q cavities make it difficult to couple to luminescent materials, which generally have much broader linewidths unless one works at cryogenic temperatures. Working in the cryostat, however, is not an ideal solution to these problems. In addition to high cost and reduced collection efficiency, matching narrow linewidths of different emitters and devices becomes highly challenging at high Q.

Recently, several groups, including our own, suggested that hybrid plasmonic-dielectric resonators can be constructed [82, 100, 102, 104, 105, 108–110, 115, 122, 170–176], raising the idea that exactly this trade-off between confinement and Q can

be reached. In this work we present a survey of the performance that should be available with hybrids if one assumes access to state-of-the-art building block cavities and antennas. To this end we discuss full-wave calculations on actually envisioned combinations of constituents, and on basis of a simple model, propose and benchmark a set of crucial design rules of thumb. Having mapped out that one can in principle indeed construct hybrids of even better $Q/\tilde{V}$ than the constituents, yet at essentially any intermediate $\tilde{V}$, we critically examine if this is of any actual use towards several applications, such as strong coupling with promising quantum emitters, bright single-photon sources, as well as high-speed light-emitting diodes, low-threshold lasers, sensing and vibrational spectroscopy. We have to conclude that room-temperature strong coupling will be as difficult to achieve with hybrids as it is with plasmon antennas alone, although once you achieve it, you have full freedom of choice over linewidth. At the same time, we conclude that hybrids are unique for their very high Purcell factors at any Q, even if their confinement is not as good as that in the very best plasmon antenna. This characteristic may offer a pathway to single-emitter strong coupling at liquid nitrogen temperatures with many different types of emitters, and to bright, low-jitter single-photon sources that might reach indistinguishability yet even operate at room temperature. If these findings would be turned into actual reality in the laboratory, this could be of large practical importance given that one could finally pass the first litmus test of "practicability" that many of the now available, highly impressive solid-state quantum light sources fail, namely that no liquid helium temperature is required.

2.2. A model system of dimer antennas coupled to a nanobeam cavity

To assess the potential of hybrid plasmonic-dielectric resonators we first numerically explore a model system using full-wave finite-element simulations. The model system is designed to overcome the main limitations of previous designs in our group [122], which was strictly confined to gapless single-particle antennas that intrinsically had quite poor LDOS enhancement characteristics, and which focused on cavities of intrinsically large mode volume. Here instead we obtain Q, $\tilde{V}$, and the LDOS of dielectric-plasmonic hybrids consisting of a silicon nitride (Si_3N_4) nanobeam cavity and a family of gold dimer antennas. The nanobeam cavity is among the smallest mode volume high-Q cavities achievable in the near-infrared, while the gold dimer antennas offer LDOS enhancements in their gap that goes well beyond the enhancements possible with single-particle nano-antennas.

In the following sections, we first describe the separate components, followed by a discussion of the merits of the hybrid system.

2.2.1. The bare cavity and antenna

For the cavity, we focus on a Si_3N_4 photonic crystal nanobeam with a design inspired by Deotare *et al.* [177]. As we will show in Section 2.3, hybrid designs have best performance if the constituent cavity and antenna themselves offer high Q and low mode volume $\tilde{V}$. Photonic crystal nanobeam cavities have shown to be near-ideal cavity systems, with high confinements and quality factors [177–179], where Si_3N_4
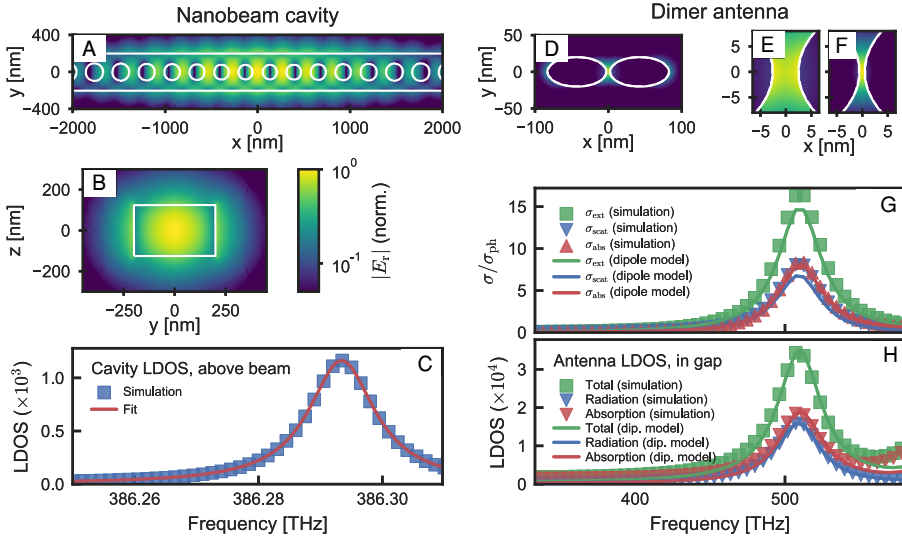
Figure 2.2: Properties of the separate components of the nanobeam cavity and the dimer antenna. (A) x, y crosscut of the cavity design and of the field distribution in the cavity. We see that the field is confined in the central defect of the photonic crystal. (B) y, z crosscut in the center of the nanobeam cavity, showing that most of the field is centered in the high-index Si_3N_4. (C) The LDOS of the cavity system at 25 nm above the beam. It has a maximum value of 1.16×10^3 and a width of 0.013 THz, which gives a Q of 3.0×10^4. (D) a crosscut of the plasmonic antenna. Clearly, most of the field is concentrated in the 5 nm gap between the metal particles. (E), (F) magnified images of the field profiles of the antennas with a gap separation of 5 and 1 nm, respectively. Note that all field profiles in (A-B) and (D-F) are normalized to their maxima and shown on a logarithmic scale. (G) Extinction, scattering and absorption cross sections σ_{ext}, σ_{scat} and σ_{abs}, respectively, of an antenna with $L = 80$ nm and a gap of 5 nm, showing the increase of cross sections at resonance. Cross-sections are normalized to the physical cross section σ_{ph}. We compare values obtained directly from simulations and indirectly via a dipole model, showing good agreement. (H) The LDOS in the center of the antenna gap (same antenna as in (G)). It shows a peak at the dipolar resonance. Again, we find good agreement between our dipole model and the numerical simulations. Deviation at the highest frequencies is due to the onset of a multipolar resonance.

is a commonly used material for dielectric cavities designed to operate in the visible to near-infrared due to its low material losses and moderately high refractive index ($n \approx 2$). Our nanobeam cavity consists of a rectangular beam that is 400 nm wide and 250 nm thick. Two sets of 6 cylindrical holes with a pitch of 300 nm and hole radii of 100 nm form photonic-crystal mirrors. Between these, a cavity is formed by a taper consisting of five holes, in which pitch and hole size are reduced linearly to 260 and 77 nm, respectively, towards the center of the cavity. The spacing between the edges of the

25

two central holes is 52 nm. The central part of the cavity design is shown in Fig. 2.2A, B.

Using the eigenmode solver of COMSOL Multiphysics, we calculate the eigenfrequency, Q and field distribution of the fundamental mode, which is confined in the center of the beam Fig. 2.2A. We calculate the effective mode volume $\tilde{V}_c$ for the cavity as

$$\tilde{V}_c = \frac{\int \varepsilon(\vec{r}) |\vec{E}(\vec{r})|^2 d^3\vec{r}}{\varepsilon(\vec{r}_e) |\vec{E}(\vec{r}_e)|^2} \left(\frac{n(\vec{r}_e)}{\lambda_{vac}} \right)^3 , \tag{2.2}$$

where ε is the dielectric constant, n is the refractive index, $\vec{E}$ is the electric field amplitude, λ_{vac} is the wavelength in vacuum, and $\vec{r}_e$ is the location of the emitter. This is the textbook definition for mode volume (ignoring the problems with it pointed out, resp. resolved in [90] and [92]), barring the fact that we evaluate the mode volume felt by an emitter centered at 25 nm *above the surface* of the beam, at position $\vec{r}_e$. This is significantly larger than the traditional mode volume that is referenced to the mode maximum, but appropriate for our envisioned antenna and emitter placement, since placing an antenna in the center of such a nanobeam cavity is not feasible. We find $\tilde{V}_c \approx 2$ which is not that far from the diffraction limit ($\tilde{V} = 1/2^3$) if one considers the disadvantageous placement. These values are similar to experimental and theoretical values reported previously in literature [177–179], as shown by the hexagon in Fig. 2.1. Though methods based on further slots taken out of the cavity have been proposed to reduce the mode volume even well below the diffraction limit [71, 73, 180], we will show that this does not significantly improve the performance of hybrids, which is why we choose the simpler cavity design. As verification of our eigenmode calculations, we have also performed driven dipole calculations to determine the LDOS spectrum, as shown in Fig. 2.2C for a dipole above the central defect of the beam, 25 nm from the Si_3N_4. We find a maximum LDOS of 1.16×10^3 at 386 THz and a linewidth of 0.013 THz, which corresponds to $Q = 3.0 \times 10^4$ and $\tilde{V}_c = 1.9$. Throughout this work, quoted LDOS values are normalized to the LDOS in vacuum at the same frequency, and LDOS is understood to mean the sum of radiative and nonradiative effects (as contained in the imaginary part of the dyadic Green function [41]).

Next, we turn to our model antenna system, for which we use a family of gold dimer ellipsoid antennas. Currently, the smallest mode volumes in plasmonics, reported to be as low as $\lambda^3/10^6$ [60, 88] are achieved not in dimer gap antennas but in metal-insulator-metal (MIM) structures composed of nanoparticles on smooth metal films, separated by a dielectric spacer [54, 93, 181–183]. However, due to the peculiar metal-film geometry and the dominant out-of-plane polarization characteristics, these MIM antennas appear less amenable to hybrid integration with a cavity than, e.g., nanorod and bow-tie antennas. In this chapter, we choose dimer ellipsoid antennas for simplicity, ease of tunability, and high LDOS in the gap. Their performance is similar to bow-tie antennas at equal gap sizes (see supplement). We study dimers with widths of 40 nm, lengths L of 60, 80, and 100 nm (for a single ellipsoid), and gaps varying between 1 and 25 nm. The length controls the antenna resonance frequency, and the antenna scattering strength. At the same time, tuning the gap tunes the LDOS enhancement at the antenna center. Fig. 2.2D shows a crosscut of the field for one example case, with $L = 80$ nm and a gap of 5 nm. As expected, most of the field is concentrated in the gap between the antennas (magnified image Fig. 2.2E). This confinement increases with decreasing gap size, as shown in Fig. 2.2F for a gap of 1 nm.

As with the nanobeam, we perform finite elements simulations on the antenna, now driving it with an incident plane wave polarized along its long axis. We can directly retrieve antenna scattering, extinction and absorption cross-sections $\sigma_{\text{scat}}, \sigma_{\text{ext}}, \sigma_{\text{abs}}$, respectively, as shown by the data points in Fig. 2.2G. We observe a resonance corresponding to an electric dipole mode, with an albedo (defined as $A = \sigma_{\text{scat}}/\sigma_{\text{ext}}$) of roughly 50%. A comparison to a dipole model, shown by the lines in Fig. 2.2G, H, is discussed in Section 2.3. A simulation with a point source at the center of the antenna gap reveals a significant LDOS, peaking at 3.4×10^4 at an apparent quality factor of $Q = 14$, as shown in Fig. 2.2h. Increasing antenna length L causes a redshift of this dipole mode and a slight increase of albedo (due to increased volume). As expected, decreasing the gap enhances both radiative and absorptive LDOS in a roughly equal manner. From these simulations we can retrieve antenna quality factor Q_a and mode volume $\tilde{V}_a$, which are shown by the colored markers in Fig. 2.1 for gaps of 1, 5 and 25 nm. Note that we use the term 'mode volume' here *not* as an endorsement of the validity of this concept per se for plasmonics [90], and the term *neither* indicates that we employed a formula similar to Eq. (2.2) *nor* that we deployed a quasi-normal mode formalism [91, 92, 175]. Instead, we obtain antenna mode volume by *inversion* of Eq. (2.1), and use it as a metric for how high the LDOS enhancement is on resonance, given the antenna quality factor Q_a.

2.2.2. The hybrid systems

To determine the properties of hybrid systems, we perform simulations of our nanobeam cavity with a gold ellipsoid dimer placed on top of the beam (see Fig. 2.3A). The dimer is placed above the center of the beam (antenna gap center is 25 nm above the surface of the beam). The long axis is aligned in the y-direction, matching the cavity mode polarization. Ellipsoid length is varied from 60 to 100 nm, and gap size from 1 to 30 nm. The source is placed at the center of the antenna gap, matching the source positions of both the bare antenna and cavity simulations. As an implementation note, in COMSOL we ensured that all calculations (for bare constituents and hybrid, driven and eigenmode calculations) use the very same mesh, where we cycle through the distinct structures by setting material constants appropriately. This approach safeguards against small shifts in frequency and Q that can occur as function of mesh and geometry truncation in COMSOL. Figure 2.3B and C show field profiles of the hybrid mode, obtained from eigenmode calculations (without source) for antennas with 40 and 80 nm short and long axes, and a 5 nm gap. In stark contrast with the bare cavity mode shown in Fig. 2.2A and B, the hybrid mode is strongly concentrated around the antenna. Nonetheless, the mode Q remains high (order 10^3) and the waveguide crosscut also clearly shows energy density inside the beam. These characteristics indicate hybridization of cavity and antenna. Figure 2.3D shows the LDOS spectra of a hybrid system with an antenna length of 80 nm and gap size of 5 nm. We observe an LDOS peak of 7.5×10^4. This is a remarkable 64-fold increase over the bare cavity LDOS, and a 2.2-fold increase over that of the antenna at resonance. Moreover, the lineshape is no longer Lorentzian but slightly asymmetric; such a Fano lineshape is characteristic of interference between a narrow resonance and a broad background [184], and has been predicted by several groups to occur in hybrid cavity-antenna system LDOS [82, 122, 171, 175, 176]. Strong enhancements of the LDOS, as compared to the bare components, has also been reported in theoretical
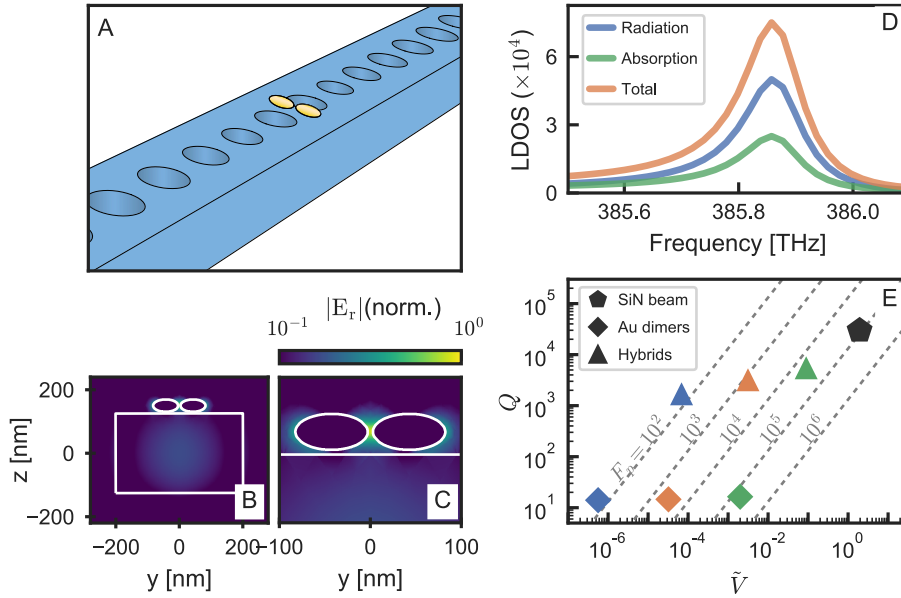
Figure 2.3: Properties of the hybrid system calculated with numerical simulations. (A) Sketch of the hybrid system, with a gold dimer placed just above the center of a nanobeam cavity. (B) Crosscuts of the field profiles of the hybrid at $x = 0$, with a gap of 5 nm. Contrary to the bare cavity shown in Fig. 2.2B, here the field is strongly localized at the antenna. (C) Magnified image of (B), showing that locally the field resembles the bare antenna mode profile in Fig. 2.2D. Fields in (B-C) are normalized to their maxima and shown on a logarithmic scale. (D) Numerically calculated radiative, absorptive and total LDOS of the hybrid. LDOS shows Fano lineshapes and we find a maximum total LDOS of 7.5×10^4 at a Q of 3000, yielding $\tilde{V}_\mathrm{H} = 3 \times 10^{-3}$. (E) Comparing a selection of hybrid systems with the bare cavity and antennas for different gap sizes. Here we show gaps of 25 nm (red), 5 nm (green), and 1 nm (blue). Hybrid Q and $\tilde{V}$ lie in between those of their constituents, and a decrease in $\tilde{V}_\mathrm{a}$ leads to a similar decrease in hybrid $\tilde{V}$.

studies of other hybrid systems [122, 175, 176]. LDOS can be further increased in our antenna-cavity systems using several tuning mechanisms (see next section).

We now compare this hybrid system in terms of the resultant Q and $\tilde{V}$ with the bare cavity and antenna. Hybrid and antenna Q are obtained through a fit with a Fano or Lorentzian lineshape, respectively. Mode volumes $\tilde{V}$ are again obtained through the peak LDOS and inversion of Eq. (2.1). A selection of the obtained values for Q and $\tilde{V}$ for our cavity, bare antennas and hybrids are shown in the phase diagram in Fig. 2.3E. Clearly, the hybrids appear right in the intermediate regime, with Q and $\tilde{V}$ between those of the cavity and antenna. Moreover, we see that hybrid systems always have higher LDOS than their individual constituents. This highlights the great potential of hybrids for practical devices that operate at intermediate Q, while maintaining

high LDOS. We note that a large host of simulations of different antenna geometries (systematically varying dimer width, length and gap size) and materials (silver and gold) all show similar behavior: hybrids show enhanced LDOS as compared to cavity and antenna, Fano lineshapes, and a quality factor in between that of the cavity and antenna.

A salient feature of the data in Fig. 2.3E is the proportional scaling of the hybrid mode volume with the antenna mode volume. As antenna mode volumes are reduced by narrowing the gap, the hybrid mode volume reduces equally. This raises the question: What determines hybrid Q and $\tilde{V}$, and what possibilities do we have to optimize these parameters? In the following section, we employ an analytical model to better understand the effects of cavity and antenna on the properties of the resulting hybrids.

2.3. Analysis of hybrids in the coupled harmonic oscillator model

Different applications of resonators in classical and quantum optics will place different requirements on resonator frequency, Q and $\tilde{V}$. To design hybrids that meet such requirements, it is important to understand how the hybrid's properties depend on those of its constituents. Here we study these properties using an analytical coupled oscillator model [122]. We refer to [122] and to Chapter 2 of [185] for a complete treatment, but recapitulate the main result here. This model is generally valid for any cavity or antenna geometry, provided that

1. the antenna can be treated as a dipolar scatterer, with no higher order multipole contributions in scattering,

2. far-field radiation overlap between cavity and antenna can be neglected, meaning that *interference* between far-field loss of the antenna and cavity can be neglected (if this is not the case, we refer to Ref. [186] for the unconventional resulting physics).

Furthermore, we assume that the near-field gap effects that imbue the antenna with a large local enhancement of LDOS as compared to a simple dipole picture, can be lumped into a prefactor that does not change with environment (see appendix). Although we expect that our systems fulfill both conditions, the purpose of this section is not to propose a perfect model, but instead to provide simple 'rules of thumb' for hybrid system design. The strength of the coupled oscillator model is that it can predict the properties of a hybrid system, given those of the individual constituents.

2.3.1. Coupled harmonic oscillator model

In the model, the antenna response is abstracted to that of a point dipole with a Lorentzian resonance in its polarizability. In the time domain, this corresponds to likening an antenna to a harmonic oscillator with charge q and mass m at resonance frequency ω_0. The coupled equations of motion describing a coupled antenna-cavity system, driven by a point dipole (i.e., classical representation of a quantum emitter) p_{dr}

at frequency ω are

$$\left(\omega_a^2 - \omega^2 - i\omega\gamma_a\right) p - \beta E_c = \beta G_{bg} p_{dr}, \tag{2.3}$$

$$\left(\omega_c^2 - \omega^2 - i\omega\kappa\right) E_c - \frac{\omega^2}{\epsilon_0 \epsilon V_c} p = \frac{\omega^2}{\epsilon_0 \epsilon V_c} p_{dr}. \tag{2.4}$$

where the induced antenna dipole moment p and the excitation of the cavity mode field E_c are the free variables, ω_a (ω_c) and γ_a (κ) describes antenna (cavity) resonance frequency and damping rate, respectively, β is antenna oscillator strength, which can be derived from the antenna scattering cross sections. For a Drude-metal sphere of volume V_{ant}, $\beta = 3V_{ant}\epsilon_0\omega_a^2$. G_{bg} is the background Green's function determining antenna-source coupling strength, V_c is cavity mode volume and $\epsilon = \epsilon(\vec{r}_a)$ is the relative permittivity at the antenna location. Importantly, $\gamma_a = \gamma_i + \gamma_r$ includes both Ohmic damping γ_i and radiative damping γ_r, ensuring that the model is self-consistent and applicable both to strongly and weakly scattering antennas [187].

At first we will consider the above equations of motion as uncoupled. That is, we consider a standalone cavity and antenna. The antenna polarizability as a function of frequency ω becomes

$$\alpha_{hom} = \frac{\beta}{\omega_0^2 - \omega^2 - i\omega\gamma_a}, \tag{2.5}$$

where γ_a is the total damping rate of the antenna. The damping rate γ_a combines Ohmic damping, and radiation damping into free space. We note that in the classical Lorentz-oscillator model a charge q and mass m appear, where $q^2/m = \beta$.

Similarly, we can consider the behavior of the cavity in absence of an antenna. If the field of a cavity near its resonance is dominated by just one of its orthogonal modes, which is generally applicable in high-Q cavities, we find that $E_c = \chi_{hom} p_{dr}$. Through similar methods as applied to the antenna, the bare cavity response becomes

$$\chi_{hom} = \frac{1}{\epsilon_0 \epsilon \tilde{V}_c} \frac{\omega^2}{\omega_c^2 - \omega^2 - i\omega\gamma_c}, \tag{2.6}$$

where $\tilde{V}_c$ is the normalized effective mode volume of the cavity, and γ_c is the cavity loss rate. This cavity susceptibility presents the response for point excitation at the point in space where for the hybrid the antenna will be placed. Thus it it encodes for the Lorentzian LDOS spectrum.

Next, one can consider the antenna to be placed in the cavity, with interactions through a multiple scattering series between cavity mode and plasmon antenna. By coupling the two systems and taking into account that the scattered fields act as additional driving terms, the hybrid polarizability and response functions can be found as

$$\alpha_H = \alpha_{hom}(1 - \alpha_{hom}\chi_{hom})^{-1} \tag{2.7}$$

$$\chi_H = \chi_{hom}(1 - \alpha_{hom}\chi_{hom})^{-1}. \tag{2.8}$$

This result signifies that the antenna, and the cavity, perturb each other. For instance, Eq. (2.8) signifies that the cavity response in presence of the antenna will be shifted in frequency and deteriorate in Q due to the introduction of the lossy antenna. Conversely, Eq. (2.7) claims that the antenna polarizability, and hence scattering properties, will be modified by the cavity resonance.

To find LDOS, we calculate the power emitted by the source dipole, given as [41]

$$P_{\mathrm{dr}} = \frac{\omega}{2} \operatorname{Im} \left\{ p_{\mathrm{dr}}^* E_{\mathrm{tot}} \right\}, \tag{2.9}$$

where E_{tot} is the total field at the source position, which can be found by solving the equations of motion Eqs. (2.3) and (2.4) and which consists of several multiple-scattering contributions. Normalizing P_{dr} to the power emitted by the source in vacuum yields the total LDOS

$$\mathrm{LDOS_H} = 1 + \frac{6\pi\epsilon_0 c^3}{\omega^3 n} \operatorname{Im}\left[\alpha_{\mathrm{H}} G_{\mathrm{bg}}^2 + 2 G_{\mathrm{bg}} \alpha_{\mathrm{H}} \chi_{\mathrm{hom}} + \chi_{\mathrm{H}}\right]. \tag{2.10}$$

In this expression, the very first term is just the contribution of free space modes to the emission rate. The other three terms arise from coupling of the emitter to the antenna-cavity system. Of these, the first term indicates acceleration of the emission due to coupling of the emitter to the antenna, however taking into account that the antenna response is dressed by the cavity. Similarly, the very last term indicates the Purcell enhancement experienced by the emitter due purely to the perturbed (shifted and spoiled by the antenna) cavity. Finally, the middle term signifies interferences between these pathways. It is also of interest to consider the limiting cases of no antenna-cavity coupling. For an antenna alone, $\chi_{\mathrm{hom}} = 0$, in which case the antenna is solely responsible for the LDOS enhancement. Indeed, if the drive dipole is in the hotspot of the antenna, and polarized in line with the antenna's dipole, the electric field exerted by the point source on the antenna equals $G_{\mathrm{bg}} p_{\mathrm{dr}}$, leading to an induced antenna dipole $\alpha_{\mathrm{hom}} G_{\mathrm{bg}} p_{\mathrm{dr}}$, and thereby a scattered field returning to the drive dipole that has strength $G_{\mathrm{bg}} \alpha_{\mathrm{hom}} G_{\mathrm{bg}}$. Similarly, in absence of the antenna, the cavity susceptibility χ_{H} reduces to χ_{hom}, and the LDOS expression by construction reduces to the bare cavity Purcell factor. From an operational perspective, we can use this correspondence to match full wave numerical simulation output to analytical model input parameters. For the antenna this means that α_{hom} is extracted from far field scattering cross sections, and subsequently G_{bg} is adjusted to match calculated antenna LDOS. A decrease in antenna $\tilde{V}$ due to, for instance, narrow gaps in dipole antennas, is captured mainly by an increase in G_{bg}.

We fit the simulated cavity LDOS to retrieve cavity resonance frequency ω_{c}, quality factor Q_{c} and mode volume V_{c}, as shown in Fig. 2.2C. From the antenna simulations, we directly obtain antenna polarizability $\alpha(\omega)$ (along the antenna long axis) by integrating the polarization currents in the antenna under plane wave driving. For a dipolar particle in a homogeneous medium, the polarizability α should relate directly to σ_{scat} and σ_{ext} [188], so we can compare 'dipolar' cross-sections (from retrieved polarizability) to the directly obtained cross-sections. Fig. 2.2G shows good agreement, indicating that our antennas are indeed dominated by electric dipole resonances.

To verify the fidelity of this retrieval, we compare analytical expressions for LDOS to simulation data (Fig. 2.2h), which show good agreement. For further details on the model and the retrieval of cavity and antenna parameters, we refer to the supplement. The interpretation of Eq. (2.10) is, crudely speaking, that the LDOS in the hybrid is that of vacuum, plus three contributions. The first ($\propto \operatorname{Im}\left\{\alpha_{\mathrm{H}} G_{\mathrm{bg}}^2\right\}$) is the contribution of just an emitter coupled to a polarizable antenna, with the caveat that the antenna polarizability is modified by the cavity. Conversely, the last term is exactly the LDOS one would expect from an emitter coupled to just a cavity, but with the caveat that the

cavity is perturbed by the antenna. Thus χ_H accounts for the change in frequency and Q predicted by cavity perturbation theory [189, 190]. Finally, the middle term contains interferences between antenna and cavity contributions.

2.3.2. Hybrid system design rules

We now study the influence of four parameters on the hybrid system properties - the antenna-cavity detuning, antenna mode volume V_a and cavity Q and V_c. This leads to four design rules for a hybrid system, which we discuss below.

    Rule of thumb I: decreasing antenna-cavity detuning decreases hybrid Q and $\tilde{V}$, at roughly equal LDOS. Figure 2.4A shows hybrid LDOS spectra for four hybrids, each with different cavity frequency ω_c. Each spectrum shows a broad peak at the bare antenna resonance and a narrow peak close to the bare cavity resonance. The width of this hybrid resonance, however, varies greatly with detuning. Figure 2.4B shows this dependency of Q_H (the hybrid Q) on detuning more explicitly. Far red-detuned from antenna resonance, Q_H approaches the bare cavity Q, yet near resonance, we see a decrease in Q_H of more than 2 orders of magnitude. The basic effect at work is captured by cavity perturbation theory, which states that the complex resonance frequency of a cavity will shift by $\Delta\omega = -\omega\alpha/(\epsilon_0\epsilon V_c)$. This implies a strong perturbation of Q_c only near antenna resonance (large imaginary α). Remarkably, the hybrid mode volume $\tilde{V}_H$ experiences a similar trend as Q_H, thus keeping peak LDOS (i.e. $Q_H/\tilde{V}_H$) roughly constant. This is shown by the envelope function in Fig. 2.4A, which describes peak LDOS for varying detuning. While Q and V vary orders of magnitude, LDOS varies only by a factor ~ 4. Letting cavity frequency ω_c vary over the spectrum in Fig. 2.4A while keeping the antenna constant results in $Q - \tilde{V}$ curves as shown in Fig. 2.4C. Just by changing detuning, the hybrid Q and $\tilde{V}$ can be changed over orders of magnitude at roughly constant LDOS.

    Rule of thumb II: Better antennas make better hybrids. Figure 2.4C also displays the influence of the antenna mode volume. Decreasing the dimer gap size leads to significantly 'better' antennas, meaning higher LDOS in the gap, or equivalently lower antenna mode volume. Fig. 2.4C shows antennas and hybrids with gaps of 1, 5, and 25 nm. As antenna mode volume (i.e. gap size) is decreased, hybrid mode volumes decrease proportionally. This reflects the fact that hybrids enjoy the benefits of strong local antenna hotspots in the same manner as a bare antenna does. In other words, the emitter-antenna coupling (captured here in G_{bg}) is not affected by the photonic environment of the antenna (i.e. the presence of a cavity). Indeed, from Eq. (2.10) we see that an increase of G_{bg} leads to an increase of LDOS$_H$ as well.

    Rule of thumb III: Decreasing cavity mode volume decreases both hybrid Q and $\tilde{V}$, while keeping LDOS fixed. Indeed, Fig. 2.4D shows that a 10-fold decrease in cavity mode volume $\tilde{V}_c$ simply shifts the hybrid parameters along the diagonal lines of constant LDOS. This is best understood by considering the expressions for Q_H and $\tilde{V}_H$, given as [122]

$$Q_H = \omega_c \left(\kappa_c + \frac{\omega_c}{\epsilon_0\epsilon V_c} \text{Im}\{\alpha\} \right)^{-1} \tag{2.11}$$

$$\tilde{V}_H = \frac{\tilde{V}_c}{|1 + \alpha G_{bg}|^2}, \tag{2.12}$$

where $\kappa_c = \omega_c/Q_c$ is the bare cavity loss rate, ϵ_0 is the vacuum permittivity and ϵ
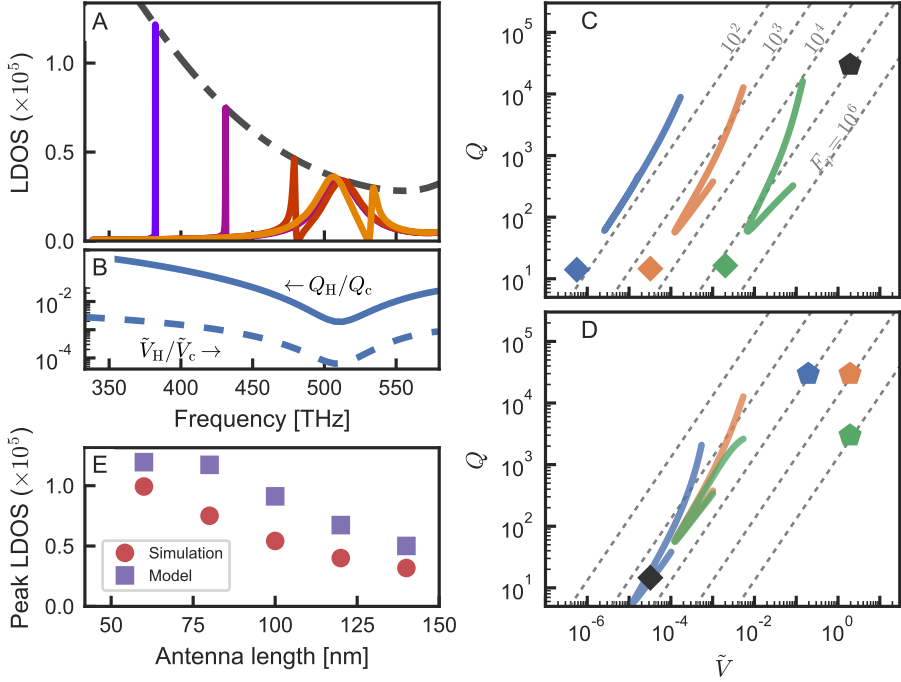
Figure 2.4: Hybrid design rules. (A) Examples of LDOS spectra for four hybrids with different cavity resonance frequencies ω_c. We fix Q_c and $\tilde{V}_c$ to those of our nanobeam cavity, and let ω_c change. The antenna is a gold dimer with ellipsoid length $L = 80$ nm and a gap of 5 nm. The dash-dotted grey line indicates the envelope function describing hybrid peak LDOS as a function of ω_c, and is given by $3/(4\pi^2)Q_H/\tilde{V}_H$, with Q_H and $\tilde{V}_H$ given in Eq. (2.11) and Eq. (2.12), respectively. (B) Hybrid quality factor Q_H and mode volume $\tilde{V}_H$ relative to cavity values, as a function of ω_c. We see a dramatic decrease of both Q_H and $\tilde{V}_H$ near antenna resonance. (C) The effect of changing detuning and antenna mode volume on Q_H and $\tilde{V}_H$. Combining a nanobeam cavity (black marker) with antennas of $L = 80$ nm and gaps of 1 (blue marker), 5 (green marker) and 25 nm (red marker), while letting ω_c vary over the spectral range shown in (A-B), leads to hybrids with Q_H and $\tilde{V}_H$ shown by the full curves (color corresponding to the antenna used). We see that Q_H and $\tilde{V}_H$ are tunable through cavity-antenna detuning, and that $\tilde{V}_H$ scales with antenna $\tilde{V}$. The dashed lines in (C) and (D) indicate constant F_P. (D) The effect of cavity Q and $\tilde{V}$. Similarly to (C), we combine an antenna with $L = 80$ nm and a gap of 5 nm (black marker) with either one of three nanobeam cavities, with Q_c and $\tilde{V}_c$ as simulated (green marker), with a 10-fold reduced Q_c (red marker) or a 10-fold reduced $\tilde{V}_c$ (blue marker). Hybrid parameters are indicated by the curves in corresponding color. We see that a reduced $\tilde{V}_c$ leads to an equal reduction in Q_H and $\tilde{V}_H$, and that Q_c only matters when Q_H approaches it. (E) Comparison between peak LDOS in hybrid systems of varying antenna length as obtained by finite element simulations of the full hybrid system (purple) and by the coupled oscillator model (yellow). Despite deviations, the model predicts the correct trend and order of magnitude.

33

the relative permittivity of the antenna host medium. As long as the hybrid Q_H is dominated by antenna losses (i.e. the second term in Eq. (2.11)), both Q_H and $\tilde{V}_H$ are proportional to cavity mode volume $\tilde{V}_c$. Thus, decreasing $\tilde{V}_c$ decreases Q_H and $\tilde{V}_H$ equally, keeping LDOS constant. This behavior breaks down when cavity losses become dominant, which happens for large detuning (small Im$\{\alpha\}$), and of course in bad cavities (low Q_c or large $\tilde{V}_c$ to start with).

Rule of thumb IV: bare cavity Q is irrelevant unless the hybrid Q approaches it. As shown in Fig. 2.4D, changing cavity quality factor has little influence on hybrid properties. Only when hybrid Q approaches that of the bare cavity, it is possible to gain any performance in Purcell factor by increasing the bare cavity Q. This is again well understood from Eqs. (2.11) and (2.12), which show that cavity losses do not affect $\tilde{V}_H$, and affect Q_H only when antenna losses are so small that cavity losses are dominant.

The strength of the coupled oscillator model used here to find these rules lies in its simplicity, even if they are not quantitatively accurate rules of thumb (only passably accurate on a log-log scale). For example, Fig. 2.4E shows a comparison between peak LDOS in hybrid systems obtained directly from simulations of hybrids, and predicted by the model. We observe that the model correctly predicts the order of magnitude and the trend in peak LDOS as antenna length is varied, although exact values deviate by up to a factor of two. This deviation is in fact easily solvable even within the analytical model. It is mainly caused by the fact that we retrieved antenna parameters in *complete absence* of the Si_3N_4, whereas in fact the nitride substrate induces an antenna redshift that is completely unrelated to the cavity resonance per se. Including this nonresonant shift largely resolves the discrepancy [122]. The rules of thumb discussed above can be used to understand the requirements for the components of a hybrid for different applications, which we will discuss in the following section.

2.4. Applications

The promise of hybrid plasmonic-photonic resonators is to provide resonances with linewidths intermediate between constituent antenna and resonator, and mode volumes that can not be reached by microcavities alone. We have now established by full-wave simulation that in principle a family of nanobeam-dimer antenna hybrids indeed allows deep subwavelength confinement, high Q and high Purcell factor, and furthermore proposed four rules for the design of resonators with quality factors and mode volumes anywhere between those of the bare cavities and antennas. Exactly which combination of hybrid quality factor and mode volume is desirable depends on the exact application one targets. In the following section, we focus on select applications in quantum optics, and discuss how hybrid systems can benefit these.

2.4.1. Single-emitter single-photon strong coupling

Strong coupling between an optical (cavity) mode and a single emitter has long been pursued in the field of quantum optics. It is a cornerstone of cavity QED, as recognized by the 2012 Nobel prize for its realization in atomic physics [191], and is hotly pursued for on-chip quantum information processing with photons and matter in the benchmark material system of III-V semiconductors and quantum dots at liquid helium temperatures [28, 61, 63, 65]. A major promise of plasmonic antennas has been to provide room-temperature strong coupling of single emitters and light, using the exceptionally

tight confinement and concomitantly large single-photon field strength to overcome the poor linewidth of emitters at room temperature. This regime has been claimed to be reached in select plasmonic nanogap antennas [60, 88, 89] that feature single-digit or sub-nanometer gaps.

In strong coupling, the emitter-cavity coupling is sufficiently strong for energy to be coherently exchanged between the emitter and the cavity before either the photon or the coherence of the emitter is lost. Experimentally, the signature is typically observed either by a spectral splitting in the frequency domain (vacuum Rabi splitting for a single emitter in a cavity) or by Rabi oscillations in the time domain. Strong coupling offers a route to create effective interactions between single photons, as the reflection of the cavity becomes different for, e.g., single and two photon states due to the saturable absorption of the emitter, and the nonlinearity of the Jaynes-Cummings ladder [165].

The onset of the strong coupling regime occurs where the coupling rate between an emitter and a cavity system exceeds the sum of the loss rates [54, 55]

$$2g > \kappa + \gamma_e, \tag{2.13}$$

where $\kappa = \omega/Q$ is the loss rate of the cavity, γ_e is the full linewidth of the emitter (including dephasing). The coupling rate g between emitter and the cavity is given by

$$4g^2 = \kappa\gamma_0 F_P = \gamma_0 \frac{3}{4\pi^2} \frac{\omega}{\tilde{V}}, \tag{2.14}$$

where ω is the frequency of the emitter transition and γ_0 is its radiative decay rate. This decay rate differs from the total decay rate γ_{tot}, which is the quantity usually obtained in lifetime measurements, by a factor which is the quantum efficiency QE, i.e. $\gamma_0 = \gamma_{tot} \times$ QE. The radiative lifetime and emission frequency are related to the emitter's oscillator f strength via [39, 54]

$$f = \frac{2\mu^2 m_e \omega}{\hbar e^2}, \tag{2.15}$$

where m_e and e are the electron mass and charge, respectively, and

$$\mu = \sqrt{\gamma_0 \frac{3\pi\varepsilon_0 \hbar c^3}{n\omega^3}} \tag{2.16}$$

is the emitter's transition dipole moment. Here, n is the refractive index of the medium embedding the emitter. Using these equations, we can determine conditions that a resonator needs to satisfy for strong coupling if the emitters emission frequency, linewidth, and radiative lifetime or oscillator strength are known.

At Room temperature

A starting point for the discussion is that any given emitter determines a characteristic curve in the $Q - \tilde{V}$ diagram, above which strong coupling is achieved. The condition for strong coupling in Eq. (2.13) yields a required minimum Q for an optical cavity at a given $\tilde{V}$ for given emitter properties set by

$$Q > \frac{\omega}{\sqrt{\frac{3\omega\gamma_0}{4\pi^2 \tilde{V}} - \gamma_e}}. \tag{2.17}$$

Emitter	$\omega/2\pi$ [THz]	$\gamma_e/2\pi$ [THz] 300K	$\gamma_e/2\pi$ [GHz] 77 K	$\gamma_{tot}/2\pi$ [MHz]	n	QY	f
DBT [144, 192]	380	38	2	208	1.59	0.24	5.5
Rh6G [193, 194]	535	47		172	1.5	0.98	10
Methylene blue [60, 195]	490	20		7.7	1.4	0.03	0.02
Lum. F Red 305 [193, 194]	490	81		126	1.5	0.91	8.1
CsPbX3 qdots [156]	545	64		67	1.5	0.7	2.7
CdSe/ZnSe qdots [151, 196–198]	500	26	1000	61	1.5	0.8	3.3
GaN qdots [199, 200]	1050	1.09	970	3330	3.4	0.98	22
InGaAs qdots [201]	307	0.96	12	1000	3.4	1	79
SiV [202]	405	0.4	120	1000	2.4	0.05	3.2

Table 2.1: Selection of emitters at room- and liquid nitrogen temperatures. Because in our calculations the emitter is assumed to be in vacuum, here we correct for the index of the surrounding medium (and for the quantum yield) when calculating oscillator strength of the emitter in vacuum, such that $\gamma_0 = \gamma_{\text{tot}} \times \text{QE}/n$.

This condition is plotted for several emitters at room temperature in Fig. 2.5A, that poignantly visualizes the different roles of the linewidth and the radiative lifetime of the emitter. The curves are typified by an inflection point at a critical combination of mode volume and Q. The emitter's radiative lifetime determines the left-right position of the curve: a faster radiative decay is equivalent to a larger oscillator strength, which shifts the curve to higher mode volumes, relaxing the strong coupling condition. At mode volumes $\tilde{V}$ below the inflection point, the dominant inhibiting factor for strong coupling is if light is lost before a Rabi oscillation is completed. In this limit, strong coupling can be achieved by increasing the quality factor of the cavity. If one traces the curve to higher $\tilde{V}$, increasing the Q to match, at some point the emitter dephasing rate becomes the limiting factor. From this point onward, no matter how much the Q is improved, the system can not be brought to strong coupling. The only way out would be to reduce the mode volume, or alternatively to cool the emitter in order to reduce dephasing.

We have made an inventory of promising efficient single photon emitters in literature, taken from a pool of outstanding organic quantum emitters [60, 144, 192–195], II-VI and III-V/III-N semiconductor quantum dots [151, 156, 196–201] and color centers [202]. The relevant properties of the emitters shown here are given in Table 2.1. For almost all single photon emitters at room temperature, the linewidths are of order 20 to 50 THz. As such, we see that at room temperature, the minimum required mode volumes are invariably between $10^{-5}(\lambda/n)^3$ and $10^{-7}(\lambda/n)^3$, irrespective of the Q that could be achieved. For reference, those systems for which room temperature strong coupling have been claimed have $Q\sim20$, and claimed mode volumes from $\sim10^{-7}(\lambda/n)^3$ to $10^{-6})(\lambda/n)^3$, just sufficient for strong coupling. The exceptional emitters in this diagram are the silicon-vacancy (Si-V) defect center in diamond and GaN and InGaAs quantum dots, which retain narrow zero-phonon lines at room temperature [199–202]. Unfortunately these exceptions are also all but impossible to embed in a gap of a few nm, making them quite unusable for room-temperature strong coupling in a hybrid structure.

To assess if hybrid resonators will facilitate strong coupling, Fig. 2.5 shows the $Q, \tilde{V}$ values for the separate components and the curves for hybrid performance, for the three hybrids from Fig. 2.4C. With the exception of SiV and InGaAs quantum dots, the hybrid curves do not reach the SC regime. At the smallest antenna gap size of 1 nm,
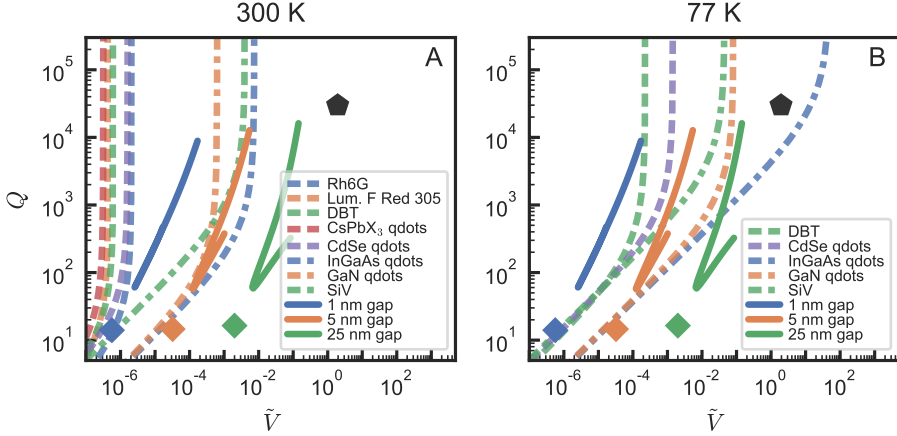
Figure 2.5: Strong coupling conditions for different emitters at room- and liquid nitrogen temperatures (above dashed lines). We find a threshold for the mode volume above which strong coupling becomes impossible. Upon cooling the emitters to liquid nitrogen temperatures, this threshold is relaxed by several orders of magnitude. We show the same cavity (black marker), antennas (colored diamonds) and hybrids (full colored lines) as in Fig. 2.4C. At room temperature (A), the Q, $\tilde{V}$ are insufficient for strong coupling with most emitters. At 77 K (B), the reduced dephasing of the emitters relaxes the strong coupling condition such that the hybrids can reach strong coupling where the components can't.

the antenna could reach SC with a room-temperature dibenzoterrylene (DBT) molecule by itself, but the hybrid will not. This illustrates the difficulty of strong coupling with hybrids at room temperature. Compared to the bare antenna, hybrids gain in Q but unfortunately also in $\tilde{V}$. This means that hybrids are advantageous for high Purcell factors at any Q. At the same time, the peculiar upswing of the strong coupling curves at a critical mode volume means that hybrids are not advantageous for strong coupling at room temperature compared to the bare antenna. Antennas that will allow a hybrid system to reach strong coupling with a single emitter at room temperature will usually be able to reach it without help of the cavity. Thus, while there could still be a benefit to have linewidth control, hybrids do not provide any alleviation of the nanofabrication problem involved in reaching room-temperature strong coupling. Similar conclusions can be gleaned from the work of Gurlek *et al.* [176]. They show that strong coupling could be possible in a tuneable-mirror cavity hybridized with a nanocone antenna, but only by virtue of the huge LDOS boost that the nanocone by itself provides.

At 77 K
When the temperature is reduced from room temperature, emitters become significantly more well-behaved. Linewidths narrow as dephasing is reduced, while radiative lifetimes remain mostly unchanged [200–204]. A main technological advantage would be if liquid helium temperatures could be avoided for a platform based on single

emitter strong coupling, ideally reaching out to temperatures achievable with Peltier coolers (-100°C), or at least no colder than liquid nitrogen. Fig. 2.5 shows that the condition on the mode volume is relaxed significantly at these temperatures. Here, the hybrid systems show a distinct improvement on the separate components, with strong coupling achievable for a variety of emitters, even in cases where the individual constituent cavity and antennas alone do not suffice. Moreover, it is possible for hybrid systems with gaps as large as 5 nm to reach strong coupling with nearly all emitters shown here, significantly outperforming its components. Hence, hybrids offer a practical route to strong coupling at 77 K — compared to cavities, they offer a larger choice in emitters and alleviate the demands on spectral alignment by operating at lower Q. Compared to antennas, spatial alignment criteria are relaxed since larger gap sizes (mode volumes) can be used. For this, one pays the price of a more complex, multi-step fabrication procedure.

2.4.2. Multiple-emitter strong coupling

Though the strong coupling of single quantum emitters to a cavity mode is generally seen as the main path towards quantum optics on basis of single-photon nonlinearities, currently many efforts are also put into achieving strong coupling with many emitters [205–210]. Strong coupling with multiple emitters is easier to achieve owing to the fact that the oscillator strength of an ensemble of N emitters effectively scales with $\sqrt{N}$, which facilitates strong coupling at larger mode volumes and lower quality factors. Fig. 2.6 shows the strong coupling condition for 1, 10^2, 10^4 emitters. These emitters have a 100 MHz radiative decay rate γ_0 (typical for organic molecules) and emit at 400 THz with a 1 THz linewidth, which lies between that of room temperature and liquid nitrogen temperature DBT molecules (see Table 2.1). The figure illustrates the benefit of an increased oscillator strength, where every factor of 100 emitters shifts the inflection point of the curves by an order of magnitude.

Strong coupling with ensembles of molecules is emerging as a topic of interest in several different fields. A main reason comes from the field of cavity exciton-polariton physics, where it is realized that intrinsically non-interacting photons become strongly interacting particles when hybridized with excitons into exciton polaritons [211, 212]. This is achieved in the multi-emitter strong coupling case. The resulting interacting quasiparticles are interesting as a realization of quantum superfluids that show phase transitions like condensation [213], superfluidity [214, 215], long range coherence [216, 217], and nonlinear states [218]. Coupled arrays of cavity exciton-polariton systems could form the basis of quantum simulators [22, 219], a topic that is being pursued in organic and semiconductor exciton systems, as well as novel materials like perovskites [220, 221] and 2D transition metal dichalcogenides [222, 223]. From an optical point of view, these systems could mean a new venue in which to study quantum light sources, such as polariton lasers, and super- and subradiance phenomena. From a more matter-oriented point of view, Hutchison *et al.* [224, 225] have pioneered the notion that the coupling of collective molecular resonances to an optical mode can alter chemical reaction energy landscapes, work functions, phase transitions, and electronic transport. Experiments show that ensembles of molecules coupled to optical cavities allow for optical and reversible switching between the weak- and ultrastrong coupling regime, with Rabi splittings that approach the molecular transition energy, and with tangible effects on chemical rate constants [205, 212, 224–228]. This line of research

Figure 2.6: Strong coupling conditions for ensembles of emitters, at different emitter number N (above dashed lines). We use emitters with $\gamma_0/2\pi = 100$ MHz, $\gamma_e/2\pi = 1$ THz, and $\omega/2\pi = 400$ THz. The oscillator strength scales with $\sqrt{N}$, and a factor of 100 in N shifts the condition for strong coupling by an order of magnitude in $\tilde{V}$. Note that even a $20 \times 20 \times 1$ nm^3 gap could in principle fit over 10^2 molecular emitters. We show the same cavity (black marker), antennas (colored diamonds) and hybrids (full colored lines) as in Fig. 2.4C and Fig. 2.5 for comparison.

also extends to vibrational spectroscopy, in an emerging field coined (cavity-enhanced) molecular optomechanics. For example Refs.[208, 226] experimentally and theoretically consider scenarios of collective strong coupling of molecular vibrations to an (infrared) cavity resonance, arguing that a macroscopic coherent superposition of molecular vibrations arises that behaves as a single mechanical oscillator. In a related context, researchers pursue Raman phenomena in plasmonic picocavities at the limit of strong coupling [88, 229]. An excellent review has been published by Flick et al. [230].

Hybrid plasmonic photonic resonators could provide a new venue for few/multi-emitter strong coupling. Currently, these type of room temperature strong coupling experiments have used extended plasmonic lattices in order to obtain higher Q (order 10^2 higher than plasmonics alone provide), requiring of order 10^7 molecules per plasmon antenna [207]. In microcavity systems that displayed strong coupling with organic fluorophores [231], the number of dye molecules per cubic wavelength of device volume was similar within one order of magnitude. For room-temperature organic fluorophores, hybrids have no advantage over pure antenna systems, as discussed in Section 2.4.1. However, at slightly lower temperatures between liquid nitrogen and room temperature, perhaps even around those provided by Peltier coolers, hybrids do offer new opportunities. Figure 2.6 shows that there, hybrid plasmonic-photonic resonators could facilitate strong coupling of few-emitter ensembles of organic molecules as small as 5-10, and upwards, at quality factors that are 100 to a few thousand. This would thus allow to very controllably study cooperative phenomena in few-emitter systems. As the dominant loss-channel of hybrids can be through the cavity

39

input-output channels, one could really envision making individual 'nodes' that are waveguide addressable, and that could be made to interact through integrated optics networks, in vein of quantum simulator demands. This should be contrasted to the extended microcavity and plasmon array systems studied in literature. Also, one could envision creating interacting strongly coupled nodes by hybridizing a single cavity with *multiple* antennas, each coupled to a patch of molecular matter.

2.4.3. Single photon sources: time jitter, brightness and indistinguishability

Single photon sources, essential for photonic quantum networks [27, 232], should fulfill a number of conditions: Applications such as quantum key distribution require high repetition rates and low timing jitter, meaning that there should be low uncertainty in when the photon is emitted [233]. This is achieved by having a short lifetime, i.e. placing the emitter in a high Purcell factor resonator. It also requires that the cavity and emitter should not be in the strong coupling regime, as this would increase jitter due to Rabi oscillations. Clearly, hybrids are excellent candidates for single-photon sources due to their high achievable Purcell factors that exceed those of the individual components. Moreover, as we have seen in Section 2.3, through cavity-antenna detuning the linewidth can be chosen to match that of the emitter, while keeping roughly the same Purcell factor. This facilitates the coupling to emitters at non-cryogenic temperatures.

Any quantum optical process requiring single-photon sources benefits from high source brightness, meaning that the source should produce as many (single) photons per second as possible [27, 96, 233]. This relates again to the lifetime, as the repetition rate of the source can never be higher than the inverse lifetime, but also to radiative efficiency. Hence, if resonators are used to decrease emitter lifetimes, these resonators should not be too lossy. Again, hybrid systems have an advantage over only-plasmon antennas, because changing the resonator linewidth also changes the distribution of energy over the cavity and the (lossy) antenna. Hence, non-radiative antenna losses can be mitigated by going to the red-detuned regime where a large fraction of the energy exits the system via radiative cavity losses. In fact, it was shown that hybrid systems can show highly efficient power extraction into a single mode waveguide, while keeping LDOS high [122, 234].

Applications that rely on the interference between two single photons to create an effective photon-photon interaction, such as several schemes for quantum computation and communication [235, 236] or boson sampling [237, 238], also require these photons to be indistinguishable [168, 239]. This implies firstly excellent control over the polarization and optical mode that the photon is emitted into, and secondly, emission spectra that are (close to) Fourier-transform limited [27], i.e. no strong dephasing. Assuming that photons are always emitted into the same optical mode with the same polarization, the photon indistinguishability I produced by an emitter with a Purcell-enhanced radiative rate of $F_P\gamma_0$ and a total linewidth γ_e (in absence of the cavity, including dephasing and radiation into the background medium) is given as [39]

$$I \approx \frac{F_P\gamma_0}{F_P\gamma_0 + \gamma_e}. \qquad (2.18)$$

This implies that in presence of dephasing, a minimum Purcell factor is required to achieve the desired indistinguishability. Moreover, to fully benefit from the Purcell
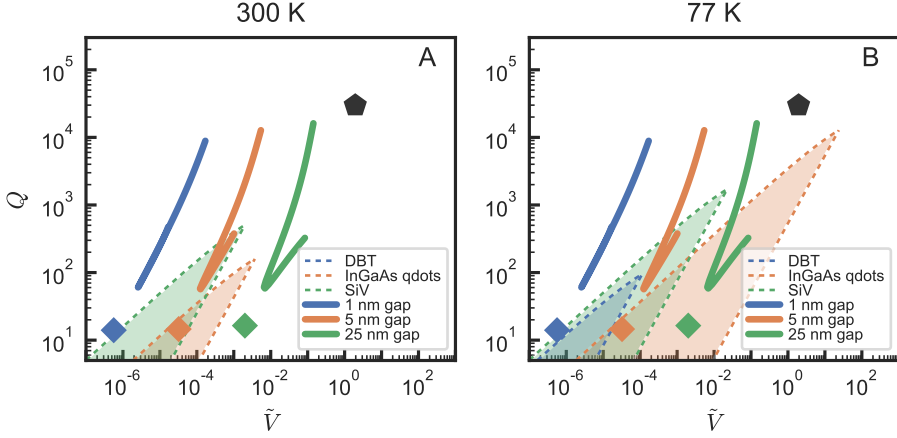
Figure 2.7: Conditions for indistinguishable photon sources, for different emitters at 300 K (A) and 77 K (B). We show Q and $\tilde{V}$ of the same cavity, antennas and resulting hybrids as shown in Fig. 2.4C and Fig. 2.5 (full curves and markers). Furthermore, for three emitters the shaded regions indicate the range of resonator Q and $\tilde{V}$ for which a single-photon source with indistinguishability $I \geq 50\%$ can be attained, i.e. Eqs. (2.18) and (2.19) are satisfied. Note that for DBT at 300 K, this region falls just outside the plot range. We show the same cavity (black marker), antennas (colored diamonds) and hybrids (full colored lines) as in Fig. 2.4C, Fig. 2.5, and Fig. 2.6.

enhancement provided by a resonator, one has to be in the 'good emitter limit' [240], meaning that the total emitter linewidth in presence of the cavity should not exceed the cavity linewidth, i.e. *

$$\gamma_e + F_P \gamma_0 \leq \kappa. \tag{2.19}$$

This condition also ensures that the device operates in the weak-coupling regime. Together, Eqs. (2.18) and (2.19) define a region in the $Q - \tilde{V}$ diagram, specific for each emitter, in which good indistinguishability can be achieved. Figure 2.7 shows these regions for a selection of emitters and a minimum indistinguishability of 50%, at room temperature and at 77K. We notice that only a narrow region is available for each emitter. High-Q cavities often fall outside this region, mainly because these do not satisfy Eq. (2.19). Antennas appear a more natural choice (at these temperatures), yet in practice, there is often another constraint to consider. Most emitters support multiple emission lines or phonon side bands, often close in frequency to the emission line of interest. To avoid also enhancing these lines (thus decreasing indistinguishability), one typically tries to match the resonator Q as closely as possible to width of this emission line. This implies that often, only the top parts of the shaded regions in Fig. 2.7 are useful. This shows that, to make a good single-photon source at high temperatures, good control over the exact Q and V of the resonator is essential. This is, of course,

---

*Note that one could also operate in the bad emitter limit, but this typically comes at the cost of source brightness, as not all photons are emitted into the (cavity) bandwidth of interest.

exactly what hybrids provide.

2.4.4. High-speed LEDs, lasers and label-free particle sensors

For a large number of applications besides spontaneous emission control, key figures of merit directly depend on the Purcell factor F_P. An obvious example is the modulation speed of a light-emitting diode (LED). LEDs could play a role in optical circuits, which have been proposed [241, 242] to replace electrical interconnects on a microprocessor. Owing to their many advantages compared to lasers, including energy-efficient operation, low cost, and high reliability, LED's are attractive as light sources for such interconnects. However, LED switching speeds are currently limited to ~100 MHz, approximately 2 orders of magnitude slower than a typical solid-state laser [242]. Since the switching rate of an LED is ultimately limited by the excited state lifetime of the carriers — although in practice also other limiting factors such as device capacitance may play a role — this rate scales proportionally with F_P. Moreover, enhancing the spontaneous emission rate additionally provides control over where the light is going [28], which enables, for example, directional emission or efficient collection of the light from such LEDs.

Purcell enhancements can also benefit the development of small, low-threshold lasers. Spontaneous emission and stimulated emission are intimately linked through the Einstein coefficients. It is therefore not surprising that, to first order, the pump power required to reach the lasing threshold is proportional to V/Q, with Q and V the quality factor and mode volume of the laser cavity mode [164]. Hence, large F_P decreases the minimal operation power of a laser, which can lead to a reduction of energy usage in optical communication [164, 242]. We note that in practice, the precise threshold power of a laser also depends on other (geometrical) parameters [243].

Besides influencing emission processes, the Purcell factor also plays a role in label-free optical particle detection. Optical resonators can be used to sense small particles such as single viruses or molecules [69, 160] through the fact that their resonance shifts when polarizable objects are placed in their near-field. Generic schemes to measure this shift convert the resonance shift in an intensity change for a narrow band laser tuned to the resonance edge in transmission or scattering. The detection sensitivity is determined by $\Delta\omega/\kappa$, i.e., by the resonator lineshift $\Delta\omega$ induced by the particle compared to the resonator linewidth. Cavity perturbation theory [190] states that, for a single small particle of polarizability α_δ perturbing a resonator, $\Delta\omega \propto -\alpha_\delta/V$. As a consequence, sensitivity is directly proportional to the Purcell factor.

LEDs, lasers, and particle sensors can clearly benefit from hybrid cavity-antenna systems, since these can achieve extremely large LDOS. In fact, hybrids have already been explored experimentally for nano-scale lasers and single-particle sensors. First hybrid lasers were demonstrated using a bow-tie antenna on a photonic crystal cavity [115, 173], however these did not harness the full potential of the device since the gain medium was embedded in the photonic crystal, far from the antenna mode maximum. As particle sensors, hybrids have been experimentally studied extensively, particularly high-Q whispering-gallery-mode cavities dressed by plasmonic antennas [102, 104, 109, 110, 170], with notable achievements including the detection of single ions in solution [106]. Note that, while hybrids are excellent as single-particle sensors, they prove less effective as bulk refractive index sensors [105].

2.4.5. Raman spectroscopy

Raman spectroscopy is based on the conversion of pump light to light at slightly up- or down-shifted frequency through an interaction with vibrations in a target material [244]. It is widely used to identify materials, as each material has its own unique spectroscopic vibrational fingerprint. Many efforts have been invested in increasing the poor efficiency of Raman processes, for example by harnessing the field confinement near a metallic surface or plasmonic antenna to boost pump intensity, and to boost the emission rate at the Stokes/anti-Stokes line [245–247]. To zeroth order, Raman signals are usually considered to scale as $\eta_R \propto |E/E_0|^4$ (fourth power of 'field enhancement') [229, 244], yet Raman scattering is a two-frequency process, and efficiency factorizes as the product of pump field enhancement at frequency ω_1, i.e., $|E(\omega_1)/E_0|^2$, and the LDOS at the shifted frequency ω_2 [175]. While in plasmonics resonances are so broad that pump-field enhancement and LDOS contributions are often near-identical, in microcavities one can separately control pump effect and Raman emission [100, 248, 249]. Hybrid cavity-antenna resonators could provide a unique venue here. Since the hot spot is pulled out of the cavity and into the antenna gap, it is directly available for the Raman-active species under investigation. At the same time, the hybridization of resonances, especially when considering the possibility of engaging several cavity modes and one antenna, could allow to independently structure the enhancement factors at pump, Stokes, and anti-Stokes frequency, judiciously matching resonances and their linewidths. This could be an especially exciting direction in the molecular optomechanics paradigm proposed by Roelli *et al.* [88, 229, 250]. We will revisit Raman spectroscopy in Chapter 4

2.5. Conclusion

In this Chapter, we have quantitatively assessed the merits of hybrid plasmonic-dielectric cavity-antenna systems, focusing on the achievable trade-off in confinement and quality factor, and the merits for diverse applications. For this, we have performed a systematic survey of performance metrics achievable in hybrids composed of a state-of-the-art high-Q nanobeam cavity and a family of plasmon-gap antennas, where the gap size tunes the bare antenna LDOS. Full-wave simulations and a simple analytical model all confirm that a large freedom over Q, on par with those of microcavities, and $\tilde{V}$, deeply subwavelength, is possible within a set of four 'rules of thumb'. These are (1) that antenna-cavity detuning controls hybrid Q, at hybrid $Q/\tilde{V}$ that remains on par with the peak antenna LDOS, (2) better antennas in terms of LDOS, make better hybrids (3) more cavity confinement helps more confinement in the resulting hybrids, again at constant $Q/\tilde{V}$, (4) the cavity Q is quite irrelevant unless one targets hybrids with similarly high Q. By themselves these 'rules of thumb' are approximate - i.e., they are accurate on log-log plots, and give a general guidelines rather than an exact quantification. With parameters extracted from full-wave simulations for the individual components, the analytical model makes them quantitative for hybrids.

In excellent agreement with previous reports of several groups [82, 108, 122, 170, 172, 174, 176, 251], hybrids can outperform the individual constituents in terms of Purcell factor, and can do so at any Q that bridges the gap between antenna and high-Q cavity. Detuning allows one to choose Q while keeping F_P almost constant. That this performance is available is a remarkable result in itself, as the mechanism

by which hybrids operate is through delicate interferences, and not through, e.g., an incoherent addition or multiplication of metrics. These interferences are directly visible in calculated LDOS lineshapes (Fig. 2.4A) that show an entire family of Fano lines going from LDOS enhancement to transparency. Notably, achieving a narrow LDOS peak at plasmonic LDOS-values is best *not* achieved by choosing plasmon-antenna and cavity *both* on resonance with the emitter.

The main purpose of this chapter is to critically assess not only the accessible performance metrics, but also if they are of use for challenges set by promises of, for instance, the plasmonic quantum optics community. The fact that phenomenally high Purcell factors at tunable quality factor are available, is a big advantage for those applications that require high Purcell factor, yet not strong coupling. These are for instance the development of room-temperature ultra-bright sources of single photons on demand, where Q-control is a crucial parameter for photon indistinguishability. Moreover, the fact that hybrids can be designed to have all their loss through cavity loss-channels, such as critically coupled waveguides, helps photon collection efficiency exceed the values achieved in ultra-high Purcell factor nano-antennas so far [96, 166].

Intuitively, one might think that the fact that *very* high Purcell factors are in reach also widens the prospects for room-temperature quantum strong coupling with single emitters, a feat so far claimed only to occur in select plasmon antennas with single-digit nanometric gaps [60, 89]. However, for this scenario, the huge dephasing rates reported in literature for actual emitters at room temperature mean that extremely small mode volumes are necessary for strong coupling to a single emitter, regardless of Q. As a consequence, hybrid cavity-antenna structures that excel at Q, but only at moderately subwavelength confinement, can not provide strong coupling conditions for any emitter at room temperature except maybe the SiV color center in diamond. Moreover, the geometrical requirements in terms of the ultra-narrow gaps that are required for strong coupling are in no way alleviated by the hybrid structure. These findings rationalize reports by Dezfouli *et al.* [178] and Gurlek *et al.* [176] that pointed out hybrids as highly promising for strong coupling, but for that needed exotic antenna shapes or gaps.

A useful niche could be in reaching strong coupling at liquid nitrogen temperatures, a regime that is significantly less demanding for real-life applications than the current liquid helium conditions of solid-state quantum optics. In this temperature regime, dephasing decreases and higher mode volumes are allowed. Consequently, strong coupling is possible with a host of different emitter choices in hybrids at relaxed fabrication conditions, where neither the antenna nor the cavity alone would suffice, and where hybrids offer a large flexibility in choosing Q. Thus, hybrids can offer strong coupling at both relaxed spectral alignment criteria (lower Q), as compared to the cavity, and relaxed spatial alignment and fabrication criteria (smaller gaps) as compared to the antenna. Finally, hybrids can also have applications outside single-emitter optics. These include few/many-emitter strong coupling in vein of work on the interface of chemistry and polaritonics, high-speed LEDs, low-threshold nano-scale lasers, and single-particle sensing.

APPENDICES

2.A. Comparison between ellipsoid dimers and bow-tie antennas

In this appendix we show a comparison between the quality factors and LDOS enhancements (parametrized through antenna mode volumes) of ellipsoid antennas and bow-tie antennas. As mentioned in the main text, bow-tie antennas form a good alternative to ellipsoid dimers, as these can be lithographically fabricated (albeit not with gaps below ~ 5 nm) [252]. Here we argue that the results discussed in the main text for ellipsoid dimers hold equally well for bow-tie antennas with similar gaps, since these can achieve similar figures of merit.
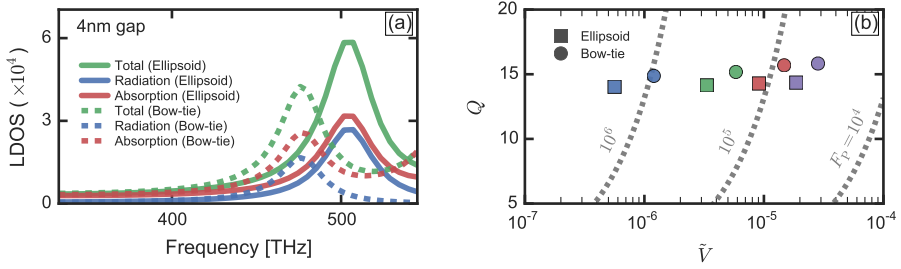
Figure 2.A.1: Comparison of LDOS spectra, Q and $\tilde{V}$ between ellipsoid dimers and bow-tie antennas. **(a)** LDOS spectra of an ellipsoid dimer and a bow-tie antenna, both with a gap of 4 nm. We see similar linewidths and peak heights, even if exact resonance frequencies are different. **(b)** Comparison of Q and $\tilde{V}$ for ellipsoid dimers and bow-tie antennas with gaps of 1 (blue), 2 (green), 3 (red) and 4 nm (purple). We see the same scaling with gap size for both antennas. Quality factors are strongly alike and mode volumes differ by up to a factor ~ 2. For all simulations, the point source was oriented along the antenna long axis and situated in the center of the antenna gap.

To verify the similarity between bow-ties and ellipsoid dimers, we perform finite element simulations of both. Ellipsoid dimer simulations are described in the main text, and here we consider only the antennas with an ellipsoid length of 80 nm and gaps of 1-4nm. We design bow-tie antennas consisting of two equilateral triangles of 40 nm height and 70 nm side length, with the tips facing each other and seperated by gaps of 1-4 nm. All edges and points are rounded with rounding radius of 2.5 nm, preventing unphysical hot spots due to sharp edges. A point source is placed at the center of the bow-tie gap, oriented along the antenna long axis.

Figure 2.A.1a shows a comparison between the LDOS spectra of an ellipsoid dimer and a bow-tie with gaps of 4 nm. Resonance frequencies and albedo (ratio of scattering to total LDOS) are different, but these can be tuned through, e.g., height and size. The important point is that both achieve similar linewidths and peak LDOS, which is what matters most for the hybrid performance. Figure 2.A.1b shows Q and $\tilde{V}$ for ellipsoids and bow-ties with gap sizes between 1 and 4 nm. Again, Q is almost identical and,

although there are differences in the exact mode volumes due to the difference in peak LDOS, ellipsoid and bow-tie mode $\tilde{V}$ have the same order of magnitude and show the same scaling with gap size. These results show that hybrids composed of bow-tie antennas coupled to cavities should achieve similar figures of merit as demonstrated in the main text for ellipsoid dimers.

2.B. Predicting hybrid performance using a coupled oscillator model

Here we explain how we predict hybrid LDOS, Q and $\tilde{V}$ using a coupled oscillator model and parameters retrieved from finite-element simulations of the bare cavity and antenna. The LDOS in a hybrid system is given by Eq. (2.10) in the main text. To evaluate this equation, for example to produce spectra as in Fig. 4a of the main text, we need the following information: cavity resonance frequency ω_c, linewidth κ and effective mode volume V_c, as well as antenna polarizability α and the background Green's function G_{bg} determining antenna-source coupling strength. Note that, while Eq. (2.3) explicitly assumes a Lorentzian response for the bare antenna, this is not a necessary condition for Eq. (2.10) in the main text to hold. If α can be retrieved directly, knowledge of ω_a, γ_a and β individually (which determine α for a Lorentzian dipole antenna) is not required.

Cavity parameters are obtained from a finite-element simulation of the cavity, driven by point dipole 25 nm above the top surface at the center of the nanobeam, oriented in the y-direction (in the frame of Fig. 2a in the main text). We calculate emitted power and normalize to the power emitted in vacuum to obtain the LDOS spectrum shown in Fig. 2c of the main text. The coupled oscillator predicts a bare cavity LDOS spectrum given as [122]

$$\text{LDOS}_c = 1 + \frac{6\pi\varepsilon_0 c^3}{\omega^3 n} \, \text{Im}\{\chi_c\}, \tag{2.B.1}$$

that is, just the last term in Eq. (2.10) in the main text, with hybridized cavity response χ_H replaced by the bare cavity Lorentzian response function

$$\chi_c = \frac{1}{\epsilon_0 \epsilon V_c} \frac{\omega^2}{\omega_c^2 - \omega^2 - i\omega\kappa}. \tag{2.B.2}$$

Equation (2.B.1) is fitted to the simulated LDOS spectrum (see Fig. 2c of the main text) to obtain ω_c, κ and V_c.

To get antenna parameters, finite element simulations of a bare antenna could be combined with a similar fit as for the cavity, as was also done in earlier work [122]. Here, however, we choose not to use a fit but instead a direct retrieval method. This largely circumvents problems due to non-dipolar contributions to the antenna LDOS, such as non-resonant quenching at very small gaps or multipolar resonances at high frequencies. Polarizability α is thus directly calculated by integrating polarization currents in an antenna driven by an external plane-wave, polarized along the antenna long axis. We then perform another simulation of the same antenna, now driven by an x-oriented point dipole at the center of the dimer gap (in the frame of Fig. 2d in the main text), producing LDOS spectra as shown in Fig. 2h of the main text. The coupled

oscillator predicts a bare antenna radiative LDOS spectrum given as [122]

$$\text{LDOS}_{\text{r,a}} = |1 + \alpha G_{\text{bg}}|^2. \qquad (2.\text{B}.3)$$

If we assume that G_{bg} is real and positive (i.e. no phase delay between the source dipole moment and the field driving the antenna), which is a good approximation for a source positioned in the hotspot of an antenna dimer, we can invert Eq. (2.B.3) to get G_{bg}, given α and the simulated $\text{LDOS}_{\text{r,a}}$. To verify this procedure, we use the retrieved α and G_{bg} to predict total and absorptive antenna LDOS, given as

$$\text{LDOS}_{\text{tot,a}} = \text{LDOS}_{\text{r,a}}/A, \qquad (2.\text{B}.4)$$
$$\text{LDOS}_{\text{abs,a}} = \text{LDOS}_{\text{tot,a}} - \text{LDOS}_{\text{r,a}}, \qquad (2.\text{B}.5)$$

respectively, with $A = \sigma_{\text{scat}}/\sigma_{\text{ext}}$ the antenna albedo that we obtain from the antenna cross-sections given by the dipole model (lines in Fig. 2g of the main text). Fig. 2h in the main text compares the simulated LDOS data to these expressions, showing good agreement. Note that we could have also retrieved G_{bg} by inverting an expression for the *total* antenna LDOS, which is similar to the first term in Eq. (2.10) of the main text. However, by using the radiative LDOS instead we are less sensitive to quenching and higher-order multipole resonances, which mainly affect $\text{LDOS}_{\text{abs,a}}$ and not $\text{LDOS}_{\text{r,a}}$.

With cavity and antenna parameters known, we can calculate LDOS spectra using Eq. (2.10) in the main text. Moreover, for several figures in the main text, including the various Q-$\tilde{V}$ diagrams, we determine hybrid Q and $\tilde{V}$. In that case, we use Eqs. 2.11, 2.12 in the main text.

3

FABRICATION OF NANOBEAM-ANTENNA HYBRID RESONATORS

> *'Explain all that,' said the Mock Turtle.*
> *'No, no! The adventures first,' said the Gryphon in an impatient tone.*
> *'Explanations take such a dreadful time.'*
> — Lewis Carroll, *Alice's Adventures in Wonderland*, 1865

3.1. Introduction

O PTICAL resonators that are hybrids of plasmonic antennas and photonic micro-cavities offer a wide range of possibilities that give them significant advantages over single-element systems of cavities or plasmonics for light-matter interaction, as reviewed in Chapter 2. The main two benefits are a very broad tunability of quality factor (Q) and at the same time the ability to reach exceptionally high local density of optical states (LDOS) due to the plasmonic confinement of field to subwavelength dimensions [253, 254]. Due to this, they are promising systems for fundamental studies into single- or multi-emitter strong coupling [205, 211, 212], innovative light sources [115, 123], or single-particle (bio)sensors [104, 105, 110]. Juxtaposed to these potential advantages of optical hybrids stand significant difficulties associated with their fabrication and utilization. Experimental realization of these systems by top-down methods often requires a sequence of mutually aligned (electron beam) lithography steps, which places limitations on the design space. Current e-beam lithography is limited to features of $\gtrsim 10$ nm [255], standing in the way of reaching ultranarrow gaps between metals. This is a main limitation for lithographically realized plasmon antennas such as bow ties [256], and equally applies to reaching the low mode volumes and high LDOS favorable for hybrids foreseen in Chapter 2. Next, both the performance of plasmonic antennas with gaps and of photonic crystal nanobeams so sensitively depend on geometry that a large challenge is posed by random variations between fabricated devices. For instance, the required alignment accuracy is at or beyond the limits of alignment in present-day multi-step electron beam lithography.

For these reasons very few lithographically realized hybrid plasmonic-photonic resonators have been reported. Notably, the only system in which a full comparison of antenna properties, cavity properties, and LDOS, have been attempted are Si_3N_4 microdisk resonators with simple nanorod antennas, functionalized with quantum dots [85]. Beyond these, the most successful reported hybrids have used other methods of antenna-cavity alignment. Gross et al. used a metallic tip as an antenna, which is brought close to the cavity by the use of an atomic force microscope [89]. The groups of Vollmer and of Goldsmith used microsphere cavities in solution, with randomly deposited metal nanoparticles [82, 106–108, 115, 121, 257, 258], Mukherjee et al., fabricated Si_3N_4 nanobeam cavities using a focused ion beam (FIB) underneath spin-coated silver nanoparticles [259], and a hybrid of a plasmon antenna and a photonic crystal cavity using a multi-step lithography process was reported by Zhang et al [115].

In this chapter, we report on the fabrication of optical hybrids made from Si_3N_4 photonic crystal nanobeams and gold dimer nanoantennas with a narrow gap, a similar design as discussed in Chapter 2. In that chapter, we discussed the properties of a freely suspended nanobeam crystal in vacuum. However, since free-hanging beams are notoriously challenging to fabricate, we opt instead to fabricate a hybrid immersed in a dielectric medium of lower index. This host medium consists of glass supporting the nanobeam, and a polymer top layer. A sketch of this system is shown in Fig. 3.1. This chapter is organized as follows. In section Section 3.2 we discuss our method for fabricating Si_3N_4 nanobeam photonic crystal cavities using electron beam lithography. In Section 3.3 we discuss our method to define nano-antennas in a second lithography step that is carefully aligned to the first. In Section 3.4 we provide an overview of the fabricated structures in terms of quality and yield on basis of scanning electron microscopy imaging.

Figure 3.1: Sketch of a photonic crystal nanobeam with a gold dimer antenna. We will fabricate this structure from Si_3N_4 effectively immersed in glass (in fact, on glass and covered with a polymer).

3.2. Nanobeam fabrication

In this section, we discuss the processes used to fabricate photonic crystal nanobeam hybrid systems. The process to fabricate the beams is sketched in Fig. 3.2.

Design

As fabrication of thick layers of Si_3N_4 often yields high internal mechanical stress that can cause cracking, we opt to alter our design from Chapter 2 to use a more conventional 210 nm Si_3N_4 thickness. Also, to avoid the challenge of fabricating suspended structures, we focus on nanobeams on a glass support, covered by PMMA to match the glass' refractive index. These choices are commensurate with available wafers of ultra-low loss Si_3N_4 on silicon oxide produced by LioniX BV. Using COMSOL to alter the design of Chapter 2, we converge on the following structure. We use a 440 nm wide beam, which has two sets of 20 elliptical holes with long axis 153 nm, short axis 51 nm, and pitch 143 nm form photonic crystals on either side of a cavity. This cavity is formed by two sets of 10 holes, where the long axis and pitch taper down linearly to 51 nm and 228 nm, respectively. The two center holes at the center of the cavity have a spacing of 143 nm. To couple the light in and out of the system, we taper out the width of the waveguide to 2.5 µm, followed by a grating. The gratings serve for out-of-plane incoupling and outcoupling of light, and each consist of 6 grating lines of Si_3N_4 250 nm wide, with a pitch of 500 nm. Using COMSOL Multiphysics 5.2 finite element simulations, we calculate the eigenmodes of this system, which predicts the resonance at 784 nm, with $Q \sim 3500$ and $V \sim 7(\lambda/n)^3$. This is the effective mode volume calculated 30 nm above the surface of the beam.

Wafer choice

The substrates are wafers produced by LioniX International. They consist of a silicon wafer as structural support layer $\sim$1.5 mm thick, with 8 µm of glass SiO_2 topped with a 210 nm layer of Si_3N_4, deposited using low-pressure chemical vapor deposition (LPCVD). The glass under the Si_3N_4 provides a layer with a lower refractive index than that of the Si_3N_4, with $n = 1.47$ and $n\sim2$, respectively.

Figure 3.2: A schematic of the process to fabricate nanobeams. (A) The sample is cleaned and a monolayer of HDMS and $430\,\text{nm}$ of CSAR are spin coated on top of the Si_3N_4. (B) The shape of the beams is exposed using e-beam lithography, and removed using a set of developing solvents. (C) A Si_3N_4 etch is used to transfer the pattern from the resist into the Si_3N_4 itself.

Substrate preparation

The wafer is cut into $12\,\text{mm} \times 12\,\text{mm}$ pieces. When sawing, the Si_3N_4 side of the wafer is glued to a plastic protective coating. To remove this, as well as inorganic particles and dust, the samples are placed in a beaker of acetone for 10 minutes and sonicated. After that, they are rinsed twice in isopropanol and dried with a nitrogen gun. Next, the samples are placed in a beaker of base piranha solution, prepared from $50\,\text{mL}$ of water, and $10\,\text{mL}$ of each ammonium hydroxide and hydrogen peroxide, heated to $75\,°\text{C}$, and kept for 15 minutes. To remove residue of the piranha solution, the samples are rinsed in water twice, once in isopropanol, and dried with a nitrogen gun. To functionalize the surface, the samples are exposed to an oxygen plasma (chamber pressure $0.25\,\text{mbar}$, $180\,\text{W}$, with $30\,\text{s}$ N_2 flush before venting) for 5 minutes, in a Diener Electronic Pico QR-200-PCCE (Pico).

Spin coating of electron beam resist

Prior to spin-coating e-beam resist an adhesion layer of hexamethyldisilizane (HDMS) is applied to the surface of the Si_3N_4 (spinning for 4000 rpm for 35 seconds). To dry the layer completely, the samples are put on a hot plate for 1 minute at $150\,°\text{C}$. After the HDMS, a layer of positive e-beam resist (CSAR 6200.13, Allresist GmbH, nominal resolution $\sim10\,\text{nm}$) is applied (~3500 rpm), with a target thickness of $430\,\text{nm}$, and baked for 3 minutes for $150\,°\text{C}$. Finally, a conductive coating (Electra 92, 5091.02, Allresist GmbH) is applied to the top by spinning at 2000 rpm for 60 seconds. This prevents charging of the all-dielectric stack on the Si_3N_4 while writing the structures. In order to keep the back of the samples clean, we use a piece of adhesive tape (either one of 1008R or 1009R Silicone-Free Blue Adhesive Plastic Film) on the back while spin coating, which peels off easily and doesn't leave residue. This facilitates keeping the back of the small samples free of resist spills, which could cause the samples to not lie flat in during the e-beam lithography step.

Electron beam lithography

To "write" our structures, we use a Raith Voyager e-beam lithography system with a $50\,\text{kV}$ thermal field column emission. In order to achieve high resolution, we operate at a beam current of $\sim0.130\,\text{nA}$, with a $5\,\text{nm}$ step size. The optimal dose for our process was found to be $\sim135\,\text{µC/cm}^2$. Great care has to be taken to align the system. This is especially important, as the antenna placement in the second lithography step needs

to be within ca. ~ 10 nm. This second alignment rests on the precise identification in the second lithography step of alignment markers written in the first step. As the system is more sensitive to misalignment of structures relative to alignment markers when larger write fields are used, we use a $200\,\mu m \times 200\,\mu m$ write field, whereas the system's maximum write field size is $500\,\mu m \times 500\,\mu m$. To enable alignment in the second lithography step, cross-shaped markers $200\,\mu m$ across are placed about $3\,mm$ from the corners of the sample. In addition, smaller alignment markers, again shaped like crosses, are written in the four corners of each write field.

Processing the resist after exposure

Before resist development, the Electra is washed off by rinsing the samples in water twice for 15 seconds each, followed by isopropanol. Subsequently, the CSAR is developed in pentyl acetate for 2 minutes, gently moving the sample through the solution to ensure a light flow of the developer around the sample. The residue left after developing is removed by transferring the sample into a beaker with o-Xylene for 7 seconds, then into a mixture of 9 parts Methyl isobutyl ketone and 1 part isopropanol for 15 seconds to stop the development. To remove all development solvents, the sample is rinsed in isopropanol twice and dried with a nitrogen gun.

Etching

At this point, the shapes drawn with e-beam lithography are patterned into the CSAR. To transfer this pattern into the Si_3N_4, we use CSAR as an etch mask for plasma etching into Si_3N_4 (Oxford Instruments Plasma Technologies Plasmalab 100 Cobra). We use a gas flow of $16\,sccm$ of SF_6, and $80\,sccm$ of CHF_3 for 100 seconds. We cool the table to $0\,°C$, with a small drop of Fomblin oil to enhance thermal contact between the sample and a Si carrier wafer. We use a reactive ion etch (RIE) forward power of $50\,W$, and an inductively coupled plasma (ICP) power of $500\,W$, at a pressure of $9\,mTorr$. We have found that the etch rate of Si_3N_4 using this recipe is $130\,nm/\min$, while the CSAR etches away at $240\,nm/\min$. Before etching the sample, we run the built-in O_2 cleaning program to ensure the chamber is clean. Subsequently, we run the process described above for 5 minutes in order to flush the system of other gases and to ensure consistency between fabrication sessions. After etching, the back of the sample is cleaned of the Fomblin oil with an acetone-soaked swab. As the CSAR that is still on the sample hardens considerably during the plasma etching, it is removed by O_2 ashing (Sinvacon Diener Pico).

3.3. Antenna fabrication

Similarly to the design in Chapter 2, a dimer antenna is placed on the beam, at the center of the cavity. In order to have these antennas slightly blue detuned with respect to the cavity resonance, we make antennas with a full width of $40\,nm$ and full length of each part of the dimer ranging from $50\,nm$ to $80\,nm$ in steps of $10\,nm$. The gap between the antennas is $20\,nm$, close to the smallest spacing achievable with our equipment. As seen in Chapter 2, the achievable F_P of a hybrid is strongly dependent on the mode volume of the antenna, where a small mode volume can be achieved through a very narrow gap in a dimer antenna. A sketch of the process is shown in Fig. 3.3
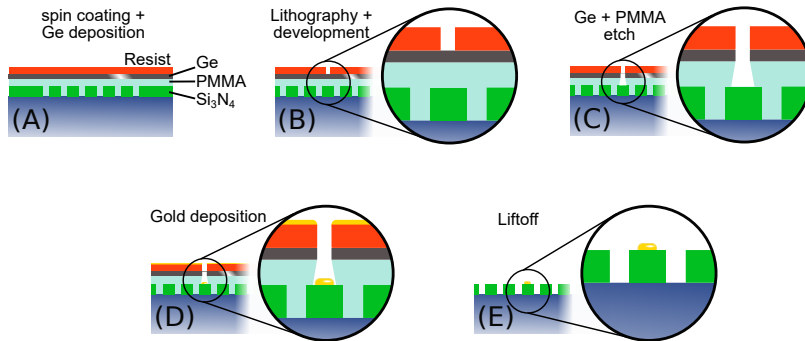
Figure 3.3: A sketch of the process to fabricate antennas on top of existing nanobeams. (A) shows the stack of resists and Ge hard mask that is applied over the beams. In (B), the lithography field is aligned using the alignment markers. The top layer of resist is exposed at the location of the antennas and removed by a liquid developing agent. (C) shows the resist stack is now etched down to the Si_3N_4 underneath. When $40\,nm$ of gold is evaporated on top, an antenna is formed on the Si_3N_4, shown in (D). In (E) warm acetone is used to dissolve the PMMA and lift off the Ge hard mask and the remaining CSAR, leaving the antenna in place on the beam.

Resist preparation

In order to fabricate these antennas and achieve a gap of only $20\,nm$, we require the highest lithographic resolution achievable in a lift-off process. In order to achieve this, we use a stack of layers, rather than a single resist layer (Fig. 3.2A). We first spin a $\sim 200\,nm$ layer of PMMA (Sigma-Aldrich, 4% solid weight in anisole) on top of the beams. Over that, an $18\,nm$ layer of germanium is deposited via thermal evaporation in a Polyteknik Flextura M508 E. Lastly, a $60\,nm$ layer of positive e-beam resist (CSAR 6200.09, Allresist GmbH, nominal resolution $\sim 10\,nm$), diluted in anisole in a 1:1 ratio, with spin speed ~ 3500 rpm. Compared to the process used for nanobeam lithography, the use of a different type of CSAR and the dilution in anisole yields a much thinner layer of resist. The philosophy behind this material stack is that only the CSAR layer functions as an e-beam resist, ultimately determining the resolution. Patterns in CSAR are transferred into the Ge layer, which subsequently acts as a mask for etching into the PMMA, in which a large undercut is created. The large undercut facilitates liftoff to produce gold nanostructures, where the gold is evaporated onto the sample through the Ge mask. This configuration of layers allows us to fabricate the $20\,nm$ gap in the dimer antenna, which would be inaccessible with a single-layer configuration.

E-beam lithography for antennas

In order to allow for a liftoff process, the PMMA *on top* of the nanobeams needs to be at least 3 times the thickness of the targeted material of the antennas, i.e., gold. Due to the topography of nanobeam samples, the thickness of spin-coated PMMA on nanobeams is different from the thickness on flat glass. To verify the thickness of the PMMA on the nanobeams, we used a focused ion beam (FIB) to make a crosscut of a nanobeam and

Figure 3.4: SEM of a crosscut of a nanobeam on glass. The beam is coated with PMMA. We see here that spin coating 250-300 nm on the sample will result in a ~250 nm thick layer of PMMA on top of the beam. Over the PMMA are an 18 nm layer of Ge and a 60 nm layer of CSAR.

the resist stack (Fig. 3.4). For the lithography progress, no Elektra conductive polymer is required as we benefit from the conductive properties of the germanium layer. We again used the Raith Voyager system with a beam current of ~0.130 nA and a step size of 5 nm. The electron dose is $150\,\mu C/cm^2$. To ensure good electronic conduction, we use two metal clamps, at opposite sides of the sample, to fix it to the sample holder.

To align antennas to cavities, the e-beam system must be meticulously aligned. We start by performing a manual 3-point alignment on the large markers located at the four sample corners. As the glass, Si_3N_4 layer, and the PMMA are not conductive, the contrast between the marker and the background is quite low. Nonetheless, it is possible to locate these markers. This allows us to coarsely align the system within ~0.5 µm. We note that accurate alignment in this step is crucial for the accurate execution of the automatic alignment procedure in the next step. The second, more precise alignment is performed on the smaller alignment markers, which are placed in the four corners of every write field, each 10 µm across. The geometry of the markers relative to the nanobeams is shown in Fig. 3.5. We use a built-in alignment procedure of the Raith Voyager which, once set up, is fully automated. The procedure relies on SEM imaging across two arms of each alignment marker. For a total of 8 scans the midpoint of each arm is determined through an edge detection algorithm. This gives sufficient accuracy to consistently place the antennas in the center of the cavity.

Development and etching

For resist development after e-beam lithography, we use the same sequence of steps as for nanobeam lithography, but with slightly different timing to accommodate for the different layer thickness. The CSAR is developed in pentyl acetate for 1 minute, the residue is removed in o-Xylene for 6 seconds, and the development is stopped in the mixture of 9 parts Methyl isobutyl ketone and 1 part isopropanol for 15 seconds. Next, we etch the CSAR pattern into the Ge layer using plasma etching (Oxford Instruments Plasma Technologies Plasmalab 80 Plus). The germanium is etched for 60 seconds with 12.5 sccm of SF_6 and 2.5 sccm of O_2. The process pressure is 8 mTorr, and the RF

Figure 3.5: A full layout of a single $200\,\mu\text{m} \times 200\,\mu\text{m}$ write field. The insets show an enlarged alignment marker, a beam, and the middle section of the beam, including the dimer antenna. In green are all areas written by the electron beam in the first lithography step, the placement of the automatic alignment scans are given in red, and the antenna placement is given in red. The diameter of the smallest holes is $100\,\text{nm}$.

forward power is 30 W. After the Ge etch, we immediately etch the PMMA using an O_2 descum recipe, for 5 minutes, with 25 sccm of O_2, at a pressure of 30 mTorr and an RF forward power is 50 W.

Gold deposition and lift off

We use e-beam evaporation to deposit 40 nm of gold at 0.05 nm/s in a Polyteknik Flextura M508 E layer deposition tool. Next we perform lift off in a beaker of acetone heated to 50 °C, and covered with aluminum foil to reduce evaporation. The samples are scratched around the structures to provide access for the acetone, and placed upright in the solvent for a few minutes until the gold starts to peel off. To let the PMMA dissolve and to remove the resist stack, we gently agitate the solvent using a glass syringe. When the gold-resist stack has sufficiently separated from the sample, we transfer the sample to a beaker of isopropanol. After a second rinse in isopropanol, we gently dry the sample with a nitrogen gun. We do not allow the acetone to dry on the sample, as it leaves residues.

To complete the hybrid, we spin coat a thick layer (500-800 nm of PMMA) over the top of all the nanobeams, including the beams that have no antenna. We then bake the sample for 5 minutes at 150 °C. As the PMMA has a refractive index very close to that of glass, this ensures that the beam is encased in a mostly homogeneous medium. Without this step, the higher index of the glass substrate would draw the optical mode away from the antenna, resulting in a much poorer coupling between the cavity and the antenna.

Antennas on glass

For reference scattering measurements we have also fabricated arrays of antennas on glass, instead of on nanobeams. The recipe to fabricate antennas on glass is very similar to that of antennas on Si_3N_4 beams. However, there are a few small differences. First, to obtain a similar PMMA thickness as on the nanobeams, we spin a nominally 150 nm thick layer. During etching, instead of etching the PMMA for 5 minutes, we only etch for 2 minutes. The reason is that the thickness of PMMA on glass is much more controlled than the thickness on top of pre-existing beams. We have verified that we have no problems with over-etching using the longer PMMA etch times in hybrids.

3.4. Fabrication results

In order to make clear scanning electron micrograph images, we deposit a thin layer of chrome, around 4 nm on the samples to make the surface electrically conductive. Without this chrome layer, the fully dielectric material stacks would have made imaging all but impossible. While for the final optical experiment the structures are covered by PMMA, naturally, we ensure for SEM imaging that there is no PMMA on top of the hybrids before applying chrome. To this end we deposit Cr directly on a freshly fabricated sample. If samples are first used in optical experiments, we remove the PMMA with acetone, with an isopropanol rinse, and blow dry. The layer of chrome can be removed by dipping the sample in a chromium etchant (Sigma-Aldrich 651826) and rinsing in water. This etches Cr at a rate of 4 nm/s, and should leave Si_3N_4 and gold untouched. However, we have refrained from doing measurements on samples that have had Cr evaporated and removed.

A scanning electron microscopy (SEM) image of bare nanobeams is shown in Fig. 3.6. Fig. 3.6A shows an example of a full nanobeam system including the incoupling and outcoupling structures. Since we use a positive resist, we need to expose a rectangle around the beam itself. Fig. 3.6B and C show details of the nanobeam cavity. The holes in our beam are subject to some roughness, which will likely negatively impact the devices' quality factor. This leads to a variation in hole diameter of about 6 nm. The sides of the beams, however, are somewhat more smooth, with variations of about 3 nm. We note that proximity effects during the e-beam process causes the holes closest to the center of the cavity to be slightly smaller, and the holes further away, larger. In future designs, this can be corrected for.

Figure 3.6E and F show SEM images of 110 nm gold dimer antennas on a glass substrate, where we have deposited 4 nm of Cr over the entire system to facilitate conduction. We see that, using this method, we can reproducibly fabricate gold dimer antennas with gap sizes of 20 nm. We find that antenna roughness is the main contributor toward antenna length and gap size variation, causing a 3 nm variation. The smallest gaps observed were 15 nm, the largest 27 nm.

Finally, Fig. 3.6G shows an example of a fabricated hybrid. We can clearly see the gold antenna placed with an accuracy of a few nm, as well as achieving a narrow gap of 20 nm. We have observed that accurate placement of the antenna is not always successful, likely due to the challenges of imaging them with the electron beam in the second lithography step. The alignment markers, discussed in Section 3.3, present a low contrast for the electron beam, since they are only holes in Si_3N_4 on a substrate of glass, covered in PMMA, which are all dielectrics. Further, the markers are imaged through a layer of germanium. We find that a remarkably small variation in antenna placement of about 10-20 nm can be achieved. However, in rare instances, the alignment procedure wasn't able to find the alignment markers, and antennas are up to 100 nm away from their target location.

3.5. Conclusion

In this Chapter, we have shown a two-step lithography process to fabricate Si_3N_4 photonic crystal nanobeams with Au dimer antennas. In the first step, we use electron beam lithography to define photonic crystal nanobeams and alignment markers in Si_3N_4. In a second step we perform electron-beam lithography and lift off to create gap antennas that are aligned to the nanobeams. A main inconvenience for the antenna placement is the poor contrast in electron microscopy of the buried, dielectric, alignment markers. Nonetheless we are able to place the antennas on top of the nanobeam cavities with quite high accuracy. The cavities were made of Si_3N_4 on a glass substrate, and are designed to be covered by PMMA to symmetrize the refractive index environment. The antennas are fabricated on top of the existing beams using a 3-layer resist stack containing a hard mask of germanium. This allows us to fabricate dimer antennas with a gap as small as 20 nm, which we can place on top of nanobeams of $\sim$500 nm width in the $\lesssim$100 nm space between two air holes.

In order to reduce surface roughness of the Si_3N_4 nanobeams, alternative development and etching recipes could be investigated. To facilitate the lithography alignment, an inversion of the fabrication process could be explored. If the gold antennas could be deposited first, and the cavities could be fabricated underneath them in a second step,

this would allow for high-contrast gold alignment markers. Higher contrast markers have the potential to lead to improved antenna-cavity alignment due to the higher fidelity with which their position can be extracted from electron microscopy.

Figure 3.6: SEM images of the fabricated systems. (A) A full nanobeam. (B), (C) close-ups of the beam's cavity. (D) A detail of the outcoupling grating. (E), (F) show two details of antennas on glass, showing their homogeneity in both length and gap size. (G) shows an example of an optical hybrid, with a gold antenna placed in the center of the optical cavity.

4

OPTICAL CHARACTERIZATION OF HYBRID NANOBEAM-ANTENNA RESONANCES

"Curiouser and curiouser!"
— Lewis Carroll, *Alice's Adventures in Wonderland*, 1865

4.1. Introduction

THE optical characterization of hybrid plasmonic-photonic resonances ideally follows
a sequence of steps. First, one should investigate the resonant properties of the in-
dividual components to ascertain at least the resonance frequencies and quality factors
of the nanobeams and the antennas and ideally also mode volumes and scattering
strengths. The next step is to perform mode spectroscopy on the hybrid system to
evidence if the plasmonic antenna and photonic cavity mode couple according to
theoretical expectations for mode hybridization, such as those outlined in Chapter 2. A
telltale signature of such coupling would for instance be the observation of perturbed
cavity mode frequencies and linewidths. Finally, any experiment to probe merits for
light-matter interaction ultimately requires introducing a third element into the system:
active matter located in the antenna gap to report on local density of optical states
(LDOS). This chapter applies this full sequence of steps to the photonic nanobeam-gap
antenna hybrids developed in Chapter 3. It is an agenda of steps that faces significant
practical obstacles.

A first significant obstacle comes from fabrication disorder. In order to accurately
assess the properties of a hybrid on basis of the qualities of the bare constituent
components, the components should be identical in both the hybrid and individually. In
practice, inevitable imprecisions in the fabrication process of each constituent and in the
relative alignment between them, make quantitative comparisons difficult. For instance,
a significant body of literature exists on cavity perturbation by nanoscale resonantly
scattering objects [186, 189, 190, 260, 260–262], but to our knowledge only a single report
compares perturbation across many lithographically prepared devices as a function of
cavity mode volume, antenna scattering strength, and detuning [85]. The reason is
that the variation in resonance frequency and quality factor between devices generally
exceeds the perturbation due to the plasmon antenna or scattering element. Mapping
perturbations therefore has often used near-field scanning microscopy, which verifies
the spatial dependence of scatterer-cavity coupling, but not the absolute coupling
strength. Outside the realm of lithographic devices, an elegant approach to track
quantitatively the coupling between scatterers and cavities was reported by Heylman
et al. [263]. They used photothermal effects to correlate particle cross-sections and
cavity perturbation.

Beyond the challenge of quantitatively understanding the hybrid mode resonances,
a further challenge arises when one desires to probe LDOS. Firstly, this would require
a third aligned fabrication step in order to selectively position an active species, such
as fluorescent molecules or quantum dots, in the antenna gap [83, 85]. Secondly, these
species should be matched in spectrum (peak emission wavelength and linewidth)
to the hybrid resonances of interest. Finally, one faces the usual challenge in LDOS
measurements on basis of fluorescence of disentangling pump enhancement, quantum
efficiency modifications, and emission directivity contributions, as reviewed in [264].
Regarding the assembly, important steps in literature have used AFM assembly of
NV centers and quantum dots with a bow-tie photonic crystal cavity [265]. Other
works [85] have pursued lithographic assembly of II-VI quantum dots near antennas on
microdisk cavities, following recipes originally developed by Curto et al. for coupling
quantum dots and Yagi-Uda plasmon antennas [83]. The emitters involved in these
assemblies are intrinsically not ideal to probe LDOS in the sense that the emission
linewidth is not well matched to hybrid resonator systems. NV centers and single

colloidal quantum dots [131, 134, 266, 267] generally show broad spectra (several tens of nanometer spectral width), as compared to the spectral width of hybrid resonances (1 to 5 nm). Conversely, for narrowband emitters, such as organic molecules in crystalline hosts [127], if assembly is even possible, a new issue arises in form of the need for spectral tuning of the emitter and hybrid.

In this Chapter, we report on mode spectroscopy of bare nanobeams, antennas, and hybrids, and provide evidence for hybridization in form of observed cavity perturbations. Next, we present preliminary experiments that use a Raman active species deposited on the metallic antennas to evidence pump-field enhancements and LDOS enhancements. The advantage of Raman active species is that vibrational lines are quite well matched in linewidth to the Q of our hybrid systems. Moreover, unlike fluorescence lines that can generally not be tuned at will to match hybrid resonances, Raman lines can simply be tuned through a photonic resonance by tuning the pump laser. As we will show, it is possible to separate near-field enhancement effects acting on excitation light from LDOS effects. This Chapter is structured as follows. Section 4.2 discusses some theoretical concepts relevant for probing the behavior of the hybrids. Section 4.3 presents methods and results for spectroscopy of the bare nanobeams. Section 4.4 reports characterization of the antennas realized in Section 3.3, on basis of dark field microscopy. In Section 4.5 we put the components together and present mode spectroscopy of hybrids. Lastly, in Section 4.6 we present results from Raman spectroscopy performed on these hybrids.

4.2. Theory of optical signatures of hybrids

In this Section, we will discuss some theoretical background required to understand the optical signatures of hybrid resonators. First, we will discuss how the presence of an optical antenna influences the behavior of an existing cavity mode. Secondly, we will discuss Raman spectroscopy, which is a powerful tool for characterization of hybrid resonators, which can be used to probe both the enhancement of the pump field in the presence of an antenna, and to probe the local density of optical states (LDOS).

4.2.1. Hybrid resonance properties

The main signature of hybridization that we have access to through interrogation of the mode structure of a cavity is the perturbation of the cavity at hand by the presence of an antenna. This perturbation expresses as a change in quality factor and in resonance frequency. We use the language of Chapter 2 and refer to Fig. 4.1 for the scenario at hand. Provided the hybrid nanophotonic system resides in the good cavity limit, where the loss rate is much smaller than the resonance frequency, an antenna can be treated as a small perturbation on a cavity mode. The antenna's inherent polarizability will perturb the cavity mode, following

$$\frac{\delta\omega}{\omega_{\mathrm{c}}} = -\frac{\alpha(\omega_{\mathrm{c}})}{V_{\mathrm{cav}}} \tag{4.1}$$

where $\delta\omega$ is a change in the *complex* cavity resonance, ω_{c} is the bare cavity resonance frequency, V_{cav} is the cavity mode volume, and α is the antenna polarizability† given

---

†Note that in this section, we will use response functions and polarizablities that have units of m³. To retrieve the polarizability in SI units, we must multiply by $4\pi\epsilon_0$

Figure 4.1: (A) and (B) show the real and imaginary parts of an antenna polarizability, respectively. The highlight shows a range of values of ω_c for experiments that focus on antennas that are blue-detuned from cavity resonances.

by

$$\alpha = \frac{3 V_{\text{ant}} \epsilon}{4\pi} \frac{\omega_{\text{ant}}^2}{\omega_{\text{ant}}^2 - \omega^2 - i\gamma\omega}. \tag{4.2}$$

Here, V_{ant} is the antenna mode volume, ω_{ant} is the real antenna resonance frequency, and γ is the antenna loss rate. The real and imaginary parts of the antenna polarizability that is typically expected for a plasmonic scatterer is shown in Fig. 4.1A and B, respectively. This expression for the perturbation of a cavity traces to work by Bethe and Waldron [189, 190], has been tested mainly by using NSOM tips to perturb cavities [260–262], and has also been invoked to rationalize whispering gallery mode biosensing performance by Vollmer and Arnold [69]. The formalism in Chapter 2 generates the perturbation theory result through the hybrid cavity susceptibility in Eq. (2.8).

As Eq. (4.1) describes a complex-valued resonance shift, both a frequency shift and a linewidth change are predicted. The real part of α will determine the shift, and the imaginary part will determine the linewidth change. As we can see in Fig. 4.1A, the direction of the lineshift depends on the relative values of ω_c and ω_{ant}, i.e., on the detuning of the cavity relative to the antenna resonance. If $\omega_c > \omega_{\text{ant}}$, the cavity mode will experience a redshift in the presence of an antenna. Conversely, if $\omega_c < \omega_{\text{ant}}$, the cavity mode will experience a blueshift. At zero detuning, the resonance shift vanishes. Instead, the linewidth follows the Lorentzian lineshape of $\text{Im}(\alpha)$. As $\text{Im}(\alpha)$ is always positive, the resonance can only broaden. This deterioration of cavity Q is maximized when $\omega_c = \omega_{\text{ant}}$. To obtain a numerical estimate that places an upper bound on the expected resonance perturbation we use the dynamic polarizability of a plasmonic dipole scatterer, obtained by including radiation damping into Frölich's model for a metallic sphere [187]. We estimate the polarizability of the antennas used in this work to be of order $(5 \times 10^{-22} + 3i \times 10^{-22})\,\text{m}^3$. This corresponds to an expected line shift and broadening of about 0.2%, which amounts to about $1.5\,\text{nm}$.

As we have discussed in Chapter 2, the highest hybrid LDOS is achievable with antennas that are blue detuned relative to the cavity. Therefore the design space covered

in this chapter focuses on antennas which are blue-detuned relative to the cavity mode. We will examine hybrids with different sizes of beams and antennas, thus varying the resonance frequencies with respect to each other within this regime of blue detuning. The range of cavity frequencies that we have access to is limited by the tunable laser with which we interrogate the systems (ca. 765 to 810 nm). Within the accessible design space, firstly, we expect to see a redshift of the hybrid resonance relative to the bare beam resonance. As we increase the antenna size, their polarizability shifts to the red and grows in magnitude, meaning that we expect to see the hybrid resonance linewidth broaden (Fig. 4.1, tracing the curve from the blue-shaded domain towards resonance).

4.2.2. Raman spectroscopy for hybrids

We will use surface-enhanced Raman scattering (SERS) as a reporter for LDOS in the hybrid resonators. When light impinges upon a molecule, there is a possibility not only for absorption but also for elastic and inelastic scattering. The inelastic scattering process is known as Raman scattering, named after the discoverer, the Indian physicist Sir Chandrasekhara Venkata Raman. The Raman scattering process couples light and molecular motion, such that scattered light can either excite a molecular vibration or annihilate one. The characteristic resulting spectral lines are called Stokes scattering when photon energy is lost to the molecule (excitation of a vibration, scattered photons are redshifted by the vibrational energy), and anti-Stokes scattering when the scattered
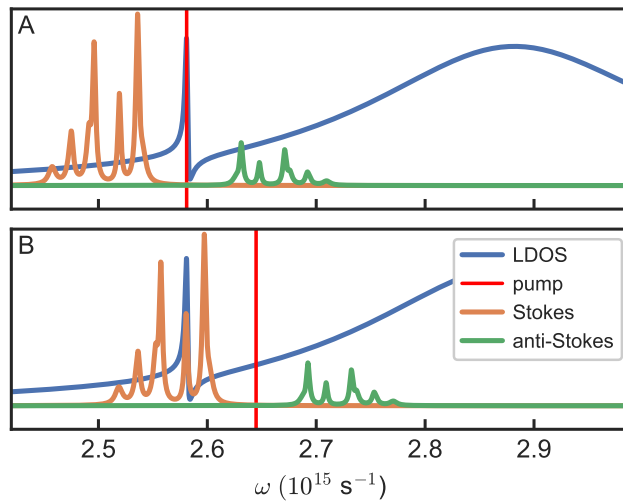
Figure 4.2: Sketch of a fictitious Raman spectrum with Stokes (orange) and anti-Stokes (green) lines on either side of the pump laser (red), overlayed with a hybrid LDOS spectrum (blue). (A) The pump light is coincident with the wavelength of the hybrid resonance. This will lead to a pump enhancement, where all Raman lines are enhanced equally in intensity. (B) The pump laser is blue-shifted such that one of the Raman lines overlaps with the hybrid line. This line will be more strongly enhanced than the others, due to LDOS enhancement.

photon energy is augmented by one quantum of vibrational energy [244]. At room temperature, vibrations in molecules are only weakly populated, and Stokes scattering dominates.

The cross-section for Raman scattering is weak. It has long been realized that by placing the molecules in electromagnetic hot spots, such as the enhanced near fields of tightly confined plasmonic modes, the intensity of both the Stokes and the anti-Stokes lines can be greatly enhanced. Since the first reports of plasmon-enhanced Raman scattering scattering [268, 269], surface-enhanced Raman scattering (SERS) has been a topic of strong interest and hot debate. SERS enhancement factors are now understood to be in part due to chemical effects, i.e., structural modifications of analyte molecules in proximity to metal interfaces, and in part due to electrodynamic effects [100, 245–249]. Furthermore, although colloquially SERS is often stated to scale as the fourth power of field enhancement $|E|^4$, it is in fact well understood that the enhancement is due to the product of two factors. The first of these is a boost of the Raman signal through pump field enhancement E_{pump}, whereby the local electric field strength E_{pump} at the driving frequency of the pump laser can far exceed the field strength of the incident beam in free space due to coupling to plasmonic resonances. The second contributor is LDOS enhancement at the Raman-scattered frequency. In a classical oscillator picture for Raman scattering, Raman scattering is the radiated output of an induced electromagnetic dipole moment oscillating at the Raman-shifted frequency, and with a magnitude proportional to the incident optical pump field, and the amplitude of the vibration. Since classically the radiated output of an oscillating dipole is proportional to LDOS, the Raman signal can be enhanced by LDOS enhancements at the Raman-shifted frequencies [100, 245–249]. The ability to separately measure pump enhancement and LDOS enhancement makes Raman scattering a highly sensitive probe to map out the physics of nanophotonics systems [270]. Fig. 4.2 sketches the use of Raman spectroscopy in this work, using as illustration a fictitious Raman spectrum with both Stokes and anti-Stokes lines, and a hybrid LDOS curve with its characteristic Fano lineshape. Importantly, typical Raman lines will have a spectral *width* comparable to or smaller than hybrid resonance linewidths. At the same time, Raman shifts well exceed the hybrid resonance linewidth. This makes it possible to separately examine pump enhancement effects and LDOS effects addressing individual vibrational lines. In Fig. 4.2A we show a configuration tuned to obtain pump field enhancement at the hybrid mode. Since all Raman lines in the spectrum are driven by the same pump field, pump field enhancement will affect all Raman peaks equally. Tuning of the pump laser (red line) around the hybrid resonance (blue Fano resonance) will cause all vibrational lines to show the same intensity variation. In Fig. 4.2B we show another configuration, with the pump laser blue detuned by an amount equal to the vibrational energy of one the Raman lines. In this configuration, the pump does not show a narrow band enhancement, although it benefits from the SERS enhancement that the antenna would also provide in isolation. However, in this configuration the LDOS is enhanced at just one of the Raman Stokes lines. When slightly tuning the pump laser, one will observe that most Stokes line will show no significant intensity variation. Only the Stokes line that is swept over the LDOS resonance will be subject to a strong change. Utilizing the different peaks in a Raman spectrum, together with judicious tuning of structure and pump laser creates a powerful tool to map out LDOS variations, and to distinguish them from pump field enhancements.

4.3. Characterization of bare nanobeam cavities

We have used two methods to characterize the resonances of photonic crystal nanobeam cavities. The main method is based on a scanning narrowband laser diode to perform excitation spectroscopy. The advantage of this approach is that it can probe resonant transmission through devices with a spectral resolution much better than the resonant linewidths of the devices at hand. The drawback of this approach is that it is tedious to perform, particularly if one has not first ascertained the approximate resonant frequencies of devices. For this reason, we use photoluminescence (PL) spectroscopy imaging for initial characterization of the bare beam resonances. The merit of this method, as opposed to excitation spectroscopy using a tuneable wavelength laser, is that a single measurement reveals optical modes over a wide wavelength range, allowing rapid assessment of the resonances of many beams. We discuss both methods in this section.

4.3.1. Bare beam photoluminescence measurements

In order to find the resonance frequencies of the fabricated nanobeams that we reported on in Chapter 3, we use photoluminescence of Si_3N_4 generated by green excitation light. The setup is described in Fig. 1.5. We use a 532 nm green pulsed laser (Time Bandwidth, Lynx) to probe the beams without antennas. We align the system in a confocal reflection mode, such that the laser spot is aimed at the center of a single nanobeam, and light is collected with the same objective (Nikon Plan Fluor 100x, 0.9 NA). When pumped with sufficient intensity, Si_3N_4 has a spectrally broad fluorescent response, which will be subject to a small enhancement on resonance with the cavity. This results in a small peak on the broad PL signal. We separate the reflected laser light from the PL using a dichroic mirror (Semrock LPD01-532Rs) and an emission filter (Semrock BLP01-532R-25), with optical densities of 2.8 and 6.4 at 532 nm, respectively. The total laser power is 0.38 mW, and we excite the center of the cavity for 30 seconds. The PL is spectrally analyzed by an Acton Research Corporation SpectraPro-2300i spectrometer using a grating with 300 lines/mm, and a Princeton Instrument Pixis 100b back-illuminated Si CCD camera, with an integration time of 30 seconds. We show the results in Fig. 4.3. For clarity, we subtracted the broad background spectrum and vertically offset the PL spectra with respect to each other. The wavelengths of the cavity modes are clearly visible as small peaks. We find two modes for each investigated beam, which we will call the first (reddest) and second (bluest) modes. We engineered the structures such that the first mode falls within the scanning range of the laser used in Section 4.3.2. The second mode falls outside this range, and we will not investigate it further.

In order to obtain a range of nanobeam resonances, we have fabricated a number of different sizes of beams, as indicated by the scaling factor on the right-hand side of Fig. 4.3 according to the procedures reported in Chapter 3. While the beams after fabrication are on glass and in air, we apply a thick covering layer of PMMA (500 — 800 nm, precise thickness not a critical parameter). This symmetrizes to a large degree the refractive index environment around the beam, and as per the COMSOL design reported in Chapter 3, should thereby provide reasonably high Q confined resonances, despite the overall modest index contrast. The size scaling is applied such that the resonance of interest (the reddest mode) sweeps through the wavelength range available to us in experiments that use our narrowband tunable diode laser, and is

Figure 4.3: Results from PL measurements on nanobeam cavities without antennas. A single measurement is shown for each beam size, which is scaled up in 0.5% increments. The plots have their common background subtracted and are offset by 280 counts with respect to each other, to improve legibility.

applied to all parameters of the beam (width, hole diameter, etc.) except the beam height, which is fixed by the thickness of the Si_3N_4. The scaling is not applied to the gratings at either end of the beam, which are kept identical for every device. The range of beam sizes yields resonance wavelengths of 777.4, 781.0, 784.5, 789.1, 791.1, 793.9, 797.7, 802.0, and 803.9 nm for scaling factors from 1.065 to 1.105 in increments of 0.005. This shows that for this range this scaling successfully shifts the resonance frequencies of the devices.

4.3.2. Bare beam transmission measurements

In order to more accurately probe the resonance frequencies and linewidths of the nanobeams, next we measure the transmission signal of each beam. A schematic of the setup is shown in Fig. 4.4. Using an IR LED, and a 3D translation stage, a desired beam can be brought into focus on a simple camera (Basler Ace acA1920-40 µm). With each beam less than 50 µm in length, a single beam fits into the field of view afforded by our microscope objective (Olympus MPlan IR, 100x, 0.95 NA). The LED light is for alignment purposes only, and is turned off for the measurement itself. Next, we align light from a narrowband tunable diode laser (Toptica DL-Pro 780), coupled to free space with a fiber collimator (60FC-4-M10-02, Schäfter+Kirchhoff) onto one grating coupler of the nanobeam under study, at a power of 44.6 µW. A set of Galvanometer mirrors (Thorlabs GVSM002/M) allows us to accurately aim the laser spot through the microscope objective (Olympus MPlan IR, 100X, NA= 0.95) and at the incoupling grating. Using the same objective, the light from the other end of the nanobeam is collected and sent to the camera. A beam block is placed in front of half the camera chip

Figure 4.4: A sketch of the experimental setup used to measure the transmission signal of the nanobeams. An IR LED illuminates the sample, allowing the desired structure to be brought into focus. The light from a tunable laser is led through a fiber polarization controller, and coupled to free space with a 10 mm fiber collimator. A set of galvanometer scanning mirrors controls the location of the beam through the objective onto the sample. A part of the light is directly reflected, and part is transmitted through the beam and collected by the same objective. This light is sent to a camera using a set of beam splitters, where the part of the laser light that is directly reflected is blocked to protect the camera. The spectrometer allows for Raman spectroscopy.

to shield it from being saturated by the reflected laser light coming off the incoupling grating.

The laser power is stabilized using a Thorlabs PM100D power meter console, which is connected in turn to an analog input port of the laser controller. An example of a beam with scaling factor 1.075 is shown in Fig. 4.5. The top image shows an image of the beam as it is seen with the IR LED. The red circle is centered around the middle of the cavity. We only show this part of the beam, as the rest is shielded to protect the camera. The lower image shows the same beam with no light from the IR LED, now illuminated at 780 nm by the laser on the incoupling grating on the left-hand side of the image (obscured by the beam block). We clearly see that a significant amount of light couples out of the grating on the right-hand side of the image, indicating transmission of laser light through the nanobeam cavity. Additionally there is a low but noticeable background of light scattering from the general area of the nanobeam, mainly from the area around the cavity where the hole size and pitch tapers down. We see this consistently throughout our measurements.

We sweep the laser in steps of ≤ 0.1 nm across the resonance wavelengths as found in Fig. 4.3, taking a single snapshot with the camera for every step, with a 20 ms integration time, waiting 1 s between measurements to allow the laser to stabilize.

69

Figure 4.5: An example of a bare beam, (A) in white light illumination and (A) when the incoupling grating is illuminated at 780 nm. We see some light at the location of the cavity center, but the majority scatters out at the outcoupling grating. The location of the center of the cavity, where the antenna is in the case of hybrids, is circled in red.

Fig. 4.6 shows the integrated counts on the camera from the outcoupling grating for a set of nanobeam sizes. As in Fig. 4.3, the resonance wavelength redshifts as the size of the beam is increased. We use these curves to extract both the resonance wavelength λ_0 and the full width at half maximum (FWHM), from which we obtain the quality factor. The behavior of all 32 measured bare nanobeams, across 8 different sizes, is summarized in Fig. 4.7. Fig. 4.7A shows the resonance wavelengths, where once again the linear increase in resonance is shown with increasing beam size, with the exception of $f = 1.080$. For these bare beams, there is quite a small dispersion, with the standard deviation in wavelength smaller than 0.6 nm for each set of beams. Fig. 4.7B shows the quality factors of all these bare nanobeams. The average of these quality factors is around 550±220, at an average linewidth of 1.4 nm. This broad distribution in Q shows there is quite some deviation between devices. We attribute this spread to the effects of fluctuations in hole size and shape, and surface roughness of the beams, which is created during the fabrication process. The best Q-factor devices approach $Q = 1000$.

4.4. Plasmon antenna scattering resonances

As discussed in Chapter 2, favorable hybrid properties can be achieved by combining a cavity with an antenna that has a relative blueshift. We fabricated a range of dimer antennas on a glass substrate, where the full length of a single antenna ranges from 50 nm to 110 nm, with a gap size of 20 nm. In order to characterize the antennas, we use a dark field microscope, shown in Fig. 4.8A. The home-built microscope is based on a Nikon LV-UEPI2 epi illumination unit, with a 100 W halogen light source. The dark field optics of this illuminator reflect the light through a collar around the lenses of the objective (Nikon CFI Plan Fluor BD, 20X, NA 0.45), which then reflects the light towards the sample at a glancing angle. The design of the objective allows only light scattered by the sample to be collected by the lenses in the objective, while the specular

Figure 4.6: The collected light from the gratings for a set of bare nanobeams in the full size range under study in this thesis. We see that the scaling factor defines the resonance of the nanobeam cavity, with larger beams shifted towards the red. We attribute the relatively broad spread in FWHM to fabrication error.

Figure 4.7: The summary of all the measured bare beams. (A) The resonance as a function of the scaling factor, where we see the resonance wavelength of the cavities shift to the red with increasing size. (B) The quality factor as a function of beam scaling factor, where we find quite a broad spread.

reflection is blocked. The dark field optics pass the scattered light back through the illuminator. Behind the illuminator it is focused by a tube lens (Thorlabs AC254-200A). A beam splitter (Thorlabs BS022) projects 30% of the light on a camera (Imaging Source DFK 21AU04), while the other 70% is projected on an image plane in which a fiber end is positioned. The 200 µm core of the fiber (Thorlabs M36L01) delimits a collection area such that only light scattered by one field of antennas is guided towards the spectrometer (Avantes 2048TEC-2-USB2, grating-based mini-spectrometer with a cooled Si CCD camera detector).

The resulting dark field spectra are shown in Fig. 4.8B and C, where we have probed antennas on glass, without a coating of PMMA, and with 400 nm of PMMA, respectively. We see that the antenna resonance shifts to the red with increasing antenna length, from approximately 620 nm to 720 nm when the antennas are exposed to air.
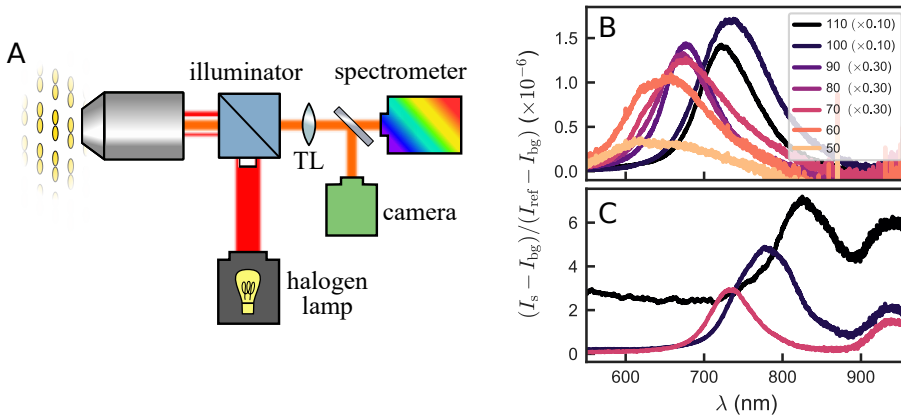
Figure 4.8: (A) A sketch of the dark field setup used. White light is focused onto an array of antennas with an objective. An illuminator unit blocks all but the scattered light, with is focused onto a spectrometer or a camera using a tube lens (TL). (B) The summary of the measured bare antennas with no PMMA top layer, with lengths of $50\,\mathrm{nm}$ to $110\,\mathrm{nm}$. We see that, as a function of increasing length of the antennas, the scattering intensity both increases and redshifts, from $620\,\mathrm{nm}$ to $720\,\mathrm{nm}$. (C) shows similar measurements, where a top layer of $400\,\mathrm{nm}$ of PMMA has been applied over the antennas, resulting in a red shift of the resonance wavelengths of $\sim10\%$.

When the PMMA top layer is applied, we see a redshift of the resonance wavelengths of $\sim10\%$. The resonance generally appears stronger for larger antennas, and presents a quality factor of order 10, commensurate with expected values for plasmon antennas with significant radiative damping. These findings allow us to study the effect of the detuning between the beam and the antenna, i.e. between both the constituents of the hybrid, on the behavior of the hybrid.

4.5. Transmission spectroscopy of hybrid plasmonic-photonic resonators

In this Section, we report on spectroscopy of nanobeams coupled to antennas. We use the same tuneable diode-laser based transmission spectroscopy method as for the bare nanobeams (Section 4.3.2) also to interrogate hybrid resonators. These hybrids are fabricated in the same size range as in the above section, with the difference that each beam has an antenna dimer located at the center of the cavity. In order to optimize antenna-cavity detuning, we use gold antennas where each part of the dimer has a full length of 50, 60, 70, and 80 nm, and the gap between the two is $20\,\mathrm{nm}$, of which the properties were discussed in Section 4.4. We show an example of a transmission measurement in Fig. 4.9. As in Fig. 4.5, this beam has scaling factor $f = 1.075$, with a bare beam resonance wavelength of $781\,\mathrm{nm}$. This nanobeam has a dimer antenna with

Figure 4.9: An example of a hybrid, (A) in white light illumination, where a red circle marks the location of the antenna and (B) with the incoupling grating illuminated at 780 nm. We see a high intensity of scattered light from the antenna at the cavity position. The color scale is cropped to make the light from the outcoupling grating visible.

a length of 80 nm placed on top of the cavity. In the top panel, we again see an image of the beam under wide-field LED illumination, where the antenna is visible as a dark dot on the beam on the left of the image. In the bottom panel, the light is provided by the tunable laser at 780 nm, aimed at the incoupling grating located on the left-hand side of the image (obscured due to beam block). The most salient feature is the strong light emission emanating from the antenna, which far outshines the light coupled out from the grating on the right-hand side. This observation clearly shows that it is possible to drive the gold dimer antenna through the nanobeam structure, implying that there is optical coupling between the optical cavity mode and the gold dimer antenna plasmon resonance.

As with the bare nanobeams, we transmit light through the nanobeam hybrids, mapping the intensity of scattered light as a function of the excitation wavelength. We examine both the intensity scattered by the outcoupling gratings, and the light that is scattered to free space by the antennas. For both signals we assess scattered intensity by integrating camera counts over a small region of interest. We show an example of the results, for a family of hybrid resonators with scaling factor $f = 1.070$ in Fig. 4.10 combined with different antenna sizes, where Fig. 4.10A shows the resulting integrated counts from the grating ($3\,\mu m \times 3.5\,\mu m$ integration area), and Fig. 4.10B shows the integrated counts from the antennas ($1.5\,\mu m \times 1.5\,\mu m$ integration area). The resonance of the bare beams as found on the grating is plotted in a vertical dotted line. For each antenna size, we provide two example data sets (solid and dashed lines of same color) to illustrate the device-to-device variability.

Firstly, we see that generally, there is a clear resonance with 1 to 2 nm width for all devices under study. In addition, the signal from the grating coupler is invariably stronger, with less background, than its counterpart retrieved from the antenna, while the wavelengths at which the signals peak shows a strong correlation. Finally, we see that most hybrid modes are redshifted with respect to the bare cavity mode by of order 1 to 2 nm. The fact that we systematically see a redshift of the hybrid resonances for the

Figure 4.10: An example of the signal from a set of nanobeam hybrids. (A)
and (B) show the total integrated counts collected over the grating and over the
antenna, respectively. We plot the curves of 4 different antenna sizes, and plot 2
hybrids for each antenna size, where half of the curves are dashed to distinguish
them. The black dotted line shows the mean of the resonance measured for the
bare nanobeams of this size (f=1.070).

blue-detuned antennas aligns with the expectations from perturbation theory discussed
in Section 3.1. At the same time we note that different realizations of nominally identical
hybrids with the same antenna length show a variability in resonance wavelength by up
to 2 nm. In other words, the device-to-device variability due to geometrical variations in
nanobeam geometry (bare cavity frequency), antenna size and gap (antenna resonance)
and antenna-to-cavity alignment quality cause spectral differences between devices that
are of the same order as the expected perturbation.

In order to survey more general trends in the data, we summarize the results of 256
hybrids and 32 bare beams. Fig. 4.11 shows the resonant wavelengths found for every
hybrid, and compares them with the resonances of the bare nanobeams. In Fig. 4.12
we show a similar graph, but for the inverse quality factor of the measured nanobeams
and hybrids as a function of antenna size, again with the systems organized in panels
separated by beam size. For the resonant wavelengths, the dominant feature is the
resonance shift with nanobeam size parameters, as also obtained for the bare beams.
For some families of hybrids (trend strongest for f=1.095) there is some evidence for
the resonance wavelengths to show a trend, i.e. a shift towards longer wavelengths for
longer antennas. Such a trend is consistent with perturbation theory Eq. (4.1) according
to which larger antennas, as long as they are blue-detuned from the cavity, cause larger
redshifts. However, this trend is only weakly visible in the data, likely because it is of
the same order as the device variability. A clearer picture emerges from the inverse

Figure 4.11: A summary of 288 measured nanobeam systems, where the observed resonance wavelength is plotted against the size of the antenna for the 8 different beam sizes. In green, at $0\,\text{nm}$ antenna length, the bare nanobeam resonances are plotted, in purple the hybrid resonances.

quality factor. The hybrid systems have an average Q of 400 ± 185, clearly below the average Q of the bare nanobeams. There is a clear trend toward lower Q as a function of increasing antenna size. For antennas with full lengths of 50, 60, 70, and 80 nm, we find mean quality factors of 460 ± 170, 440 ± 170, 340 ± 140, and 250 ± 110, respectively. This corresponds to a change in linewidth from $1.4\,\text{nm}$ for the bare beams to $1.7-3\,\text{nm}$ for the hybrids. This trend is qualitatively commensurate with the prediction of perturbation theory that the cavity linewidth increases in proportion to the antenna polarizability, which increases with antenna size. Quantitatively the Q-change is of the same order, or arguably even somewhat larger than expected from perturbation theory. Here it should be noted that the polarizability that enters perturbation theory is difficult to estimate precisely, while furthermore perturbation theory disregards multipolar near-field corrections to antenna response. Such corrections could be important for antennas directly placed on top of the Si_3N_4 - PMMA interface.

4.6. Probing optical hybrids with Raman spectroscopy

In order to finally investigate the hybrids with Raman spectroscopy, we remove the previously applied PMMA from the hybrids by soaking the sample in room temperature acetone for about 10 seconds, followed by a rinse in ethanol, and blowing the sample dry

Figure 4.12: A summary of 288 measured nanobeam systems, where the observed inverse Q is plotted against the size of the antenna for the 8 different beam sizes. In green, at $0\,\mathrm{nm}$ antenna length, the inverse Q of the bare nanobeam are plotted, in purple those of the hybrids.

with an N_2 gun. To remove further traces of organic materials, the sample was exposed to a gentle oxygen plasma (chamber pressure $0.25\,\mathrm{mbar}$, $180\,\mathrm{W}$, with $30\,\mathrm{s}$ N_2 flush before venting) for 1 minute, in a Diener Electronic Pico QR-200-PCCE (Pico). We then soak the sample in a biphenyl-4-thiol (BPT) solution for 12 hours ($1\,\mathrm{mmol}$ in ethanol). The BPT has an affinity to adhere to the gold of the antennas in form of self-assembled monolayers, and is therefore a standard analyte for SERS experiments. It has no affinity to the other materials at hand, i.e. Si, Si_3N_4 or SiO_2 [271]. The sample is rinsed in ethanol and dried with an N_2 gun. Lastly, $400\,\mathrm{nm}$ of PMMA (Sigma-Aldrich, 4% solid weight in anisole) is spin-coated over the top of all hybrids, and baked at $160\,°\mathrm{C}$ for 4 minutes, in order to restore the index-matching environment around the beam. In order to probe Raman spectra, an additional branch in the setup for hybrid nanobeam spectroscopy (Fig. 4.4) is implemented. Signal collected by the microscope objective (Olympus MPlan IR, 100X, NA= 0.95) can be routed to a fiber-coupled spectrometer (Andor Shamrock A-SR-303i-B-SIL) equipped with a cooled silicon CCD camera (Andor iVac A-DR324B-FI). Excitation is performed with a narrowband tunable diode laser (Toptica DL-Pro 780), spectrally cleaned with a pair of bandpass filters (Semrock TBp01-790/12) before being inserted into the setup. Prior to insertion of detected light into the collection fiber, the Rayleigh-scattered light is filtered out with a set to two notch filters (both Thorlabs NF785-33). This configuration allows for the detection of both Stokes and anti-Stokes

signal, with the tuning of the pump frequency over a relevant (tens of nm) range.

Prior to measuring hybrid nanobeam systems, we verified the SERS spectroscopy for antenna arrays on glass, and for isolated antennas on Si_3N_4 in PMMA that are not coupled to a nanobeam resonance. The main findings are, firstly, that the SERS signal from BPT remains unaffected upon application of PMMA, and secondly, that the signal arises uniquely from SERS enhancement right at the antenna. Here we only show representative results for SERS of BPT at a single antenna on Si_3N_4, and in PMMA, comparing different pump polarizations. Fig. 4.13A shows a full set of Raman spectra for polarization parallel to the long axis of the antenna, where the pump laser is scanned from 765 to 779 nm, taking a full spectrum at every step. We see a number of Raman lines clearly, which shift as the pump laser is swept (remnant of pump laser weakly visible, scanning through the notch filter pass band). This result demonstrates the capability of our setup to detect Raman lines while scanning the excitation wavelength. We will use this to perform tuning of either the pump wavelength or a Raman line to hybrid resonances. Fig. 4.13B shows a Raman spectrum for pump polarization perpendicular to the long axis of the antenna, while Fig. 4.13C shows the spectrum for the same structure with the pump light along the long antenna axis. Since the pump wavelength (775 nm) is close to the long-axis resonance, and not the short-axis antenna resonance, we only expect SERS enhancement for the configuration in panel B. In both panels, we see a clear Raman signature from the deeply buried silicon substrate underneath the oxide, which has a well-known Raman line at $516\,cm^{-1}$. For the perpendicular polarization case in Fig. 4.13B there is essentially no other signal, where we note that it was taken with an integration time of 1200 s, as opposed to the parallel polarization case where we required only 20 s (Fig. 4.13C). This caused the spectrometer to oversaturate at the Si peak, making it appear lower in Fig. 4.13B The Raman lines for parallel polarization are characteristic of BPT. In this chapter, we will focus on the peaks at 473, 1073, and $1576\,cm^{-1}$. Clearly, the Raman signal without the enhancement from the antennas is weaker by at least two orders of magnitude compared to the SERS signal when driving the long antenna axis. We can thus conclude that any measured Raman signal on hybrids will originate uniquely from BPT molecules in the close vicinity of the antennas.

4.6.1. SERS in hybrids - pump enhancement

We now turn to the use of SERS as a method to probe hybrid resonances. We exploit both of the scenarios discussed in Section 4.2.2. First, we report on a nanobeam with a resonance well within the wavelength range of our pump laser, with scaling factor f=1.065. As a consistency check to identify the hybrid resonance frequency, we first couple in to the grating of a bare nanobeam hybrid with a fiber coupled white light source, and collect the light from the opposite end onto our spectrometer. As shown in Fig. 4.14A, we find a resonance consistent with our previous, more finely sampled transmission measurements, with $Q = 323$. This result indicates that the structures are not affected by the removal and redeposition of PMMA. We then move to a hybrid with the same size, but with an antenna with length 80 nm. Figure Fig. 4.14B shows Raman spectra for free space excitation directly focused on the antenna, and with signal collection from the same location and through the same objective. We sweep the laser across the resonance of the cavity, from 765 to 780 nm, while keeping the power constant at $400\,\mu W$. For each wavelength step, we integrate the spectrum for 20 seconds. For

Figure 4.13: Raman spectra taken from the same antenna on a nanobeam, for two different polarizations of the pump laser. (A) An image of 44 Raman spectra taken, sweeping the pump laser from 765 to 779 nm, with the polarization along the long axis of the antennas. A number of Raman peaks can be seen, as well as a dark band where we filter the pump laser. (B) A spectrum where the pump polarization is perpendicular to the long axes of the dimer antenna, integrated for 1200 seconds. We see a clear silicon signature at $516\,\mathrm{cm}^{-1}$. We don't see any of the characteristic Raman peaks of the BPT. (C) We use the same hybrid, but rotate the polarization of the pump so it in line with the long axis of the antenna. Here, we use a $20\,\mathrm{s}$ integration time. Here, we again see the Si peak, now accompanied by a number of Raman peaks of the BPT at 473, 1073, and $1576\,\mathrm{cm}^{-1}$. (B) and (C) were taken with a pump laser at 775 nm.

each spectrum, we integrate the counts underneath the Raman peaks at 473, 1073, and $1576\,\mathrm{cm}^{-1}$ around the maximum of the peak, over a range of $-20\,\mathrm{cm}^{-1}$ to $+20\,\mathrm{cm}^{-1}$, completely encompassing the peaks. Curves in Fig. 4.14B correspond to the signals under the peaks at 473, 1073, and $1576\,\mathrm{cm}^{-1}$, respectively (blue, green, and red line colors). The signals are normalized by the $520\,\mathrm{cm}^{-1}$ Raman line from the underlying substrate, which we argue can be used as an intensity reference. We see that, on this hybrid, we have a clear signal from each of our selected Raman lines. As we are exciting at the resonance wavelength of the hybrid, all the Raman peaks are detuned from the maximum in LDOS. However, the pump frequency is directly swept through the hybrid resonance Thus, we expect an identical signature from the hybrid resonance in all Raman lines, as they experience the same Fano-line in pump enhancement, and an

Figure 4.14: Properties of a nanobeam hybrid, with scaling factor f=1.065, measured for Raman spectroscopy. (A) a transmission spectrum of a bare beam, using white light that is coupled in and out from the gratings. (B,C) Spectra on a nanobeam hybrid with a dimer antenna with length 80 nm performed by scanning a narrowband pump laser across the resonance of the hybrid. Light is collected from a spectrometer, and integrated underneath the three Raman peaks, resulting in the blue, red, and green curves. (B) is measured by pumping and collecting directly on the laser, and (C) is measured by coupling in through the grating and collecting at the antenna. We see that this method of measuring causes the signal to be filtered by the cavity itself. (B, C) show that all Raman peaks follow the same trend, demonstrating that this signal is a result of pump enhancement on the BPT. The dashed black curve shows the results from the CHO model, showing excellent agreement with the data.

unstructured LDOS. This situation corresponds to the case outlined in Section 4.2.2A.
We find that the curves all follow a Fano lineshape, the signature of hybridization of
the plasmonic and cavity resonance on the local pump field.

Fig. 4.14C shows SERS results on the same hybrid, but interrogated through a
different input channel as now we excite the antenna through the waveguide and
cavity. To this end, we align the pump laser to one of the gratings, but still collect
light from the location of the antenna. We find that this method gives a sufficiently
strong waveguide-addressed SERS signal to show a clear signal. When we sweep the
laser frequency, we again find that all Raman lines follow a very similar curve as a
function of wavelength. However, instead of a Fano lineshape, the SERS enhancement
now traces an almost Lorentzian resonance. This observation can be rationalized from
the fact that again the SERS enhancement is due to a strong spectral modification of
the pump enhancement, yet an unstructured LDOS at the vibrational lines. When
pumping through the waveguide, the only condition at which the pump light can
reach the antenna is when the pump light is resonant with the perturbed cavity. The
photonic crystal cavity essentially acts as a strong spectral filter, which will only permit
light to pass when it is on resonance with the cavity. The broad plasmon resonance
that was responsible for the destructive interference in the Fano line does not strongly
contribute in this scenario.

4.6.2. SERS in hybrids - LDOS enhancement

We now turn to probing selective LDOS enhancement in hybrid systems, as opposed
to pump enhancement. To this end, we select a different cavity with a lower resonance
frequency, with f=1.105. In contrast to the previous section, we do not sweep the pump
laser across the hybrid resonance. Instead, we sweep the laser through a wavelength
interval blue detuned relative to the hybrid resonance by about 25 nm, such that a
single Stokes-shifted Raman peak tunes across the hybrid resonance. This realizes the
situation shown in Fig. 4.2B.

Once again we used a white light transmission spectrum on a bare nanobeam
(shown in Fig. 4.15A) to find the resonance frequency, establishing $Q \sim 640$ at 799.5 nm.
After this, we once again probe the hybrid under study using SERS, using the same
method for light collection and data processing as described in the previous section,
now sweeping the BPT Raman line at $473 \, \text{cm}^{-1}$ across the hybrid resonance. This
corresponds to sweeping the laser from 765 nm to 779 nm. Again, we interrogated
structures both fully from free space, and through the waveguide. Fig. 4.15B shows
measurements taken from free space by pumping on and collecting from the antenna,
taking Raman spectra at all wavelengths, and integrating counts under each peak. In
stark contrast to the earlier results with pump light resonant with the hybrid, now just
one Raman line shows a striking Fano lineshape, while the other Raman lines follow a
gently sloping shoulder. The fact that only the $473 \, \text{cm}^{-1}$ line that is swept through the
resonance shows a Fano-line, and not the other lines, is consistent with the notion that
the pump field is not strongly modified by the cavity mode, and that the hybrid LDOS
shows a Fano line that acts only on the vibration that it is resonant with. The gently
sloping shoulder in the other lines we attribute to the typical SERS enhancement of the
bare antenna, which diminishes as one tunes away from resonance.

Finally, Fig. 4.15C shows waveguide addressed SERS in this system. As the photonic
crystal cavity prohibits any optical transport between the grating coupler and the
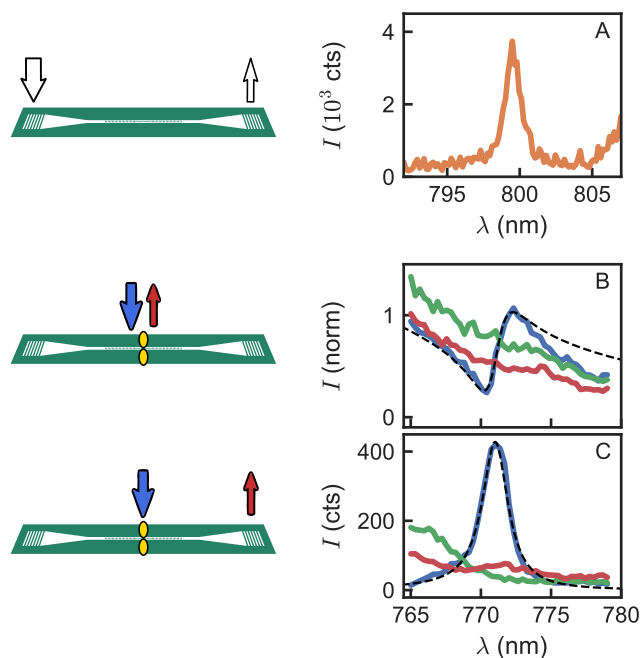
Figure 4.15: Properties of nanobeam hybrid systems, with scaling factor f=1.105, measured for Raman spectroscopy. (A) a transmission spectrum of a bare beam, using white light that is coupled in and out from the gratings. (B,C) Spectra on a nanobeam hybrid with a dimer antenna with length $80\,\text{nm}$ performed by scanning a narrowband pump such that the Raman line at $471\,\text{cm}^{-1}$ is scanned across the cavity resonance. Light is collected from a spectrometer, and integrated underneath the three Raman peaks, resulting in the blue, red, and green curves. (B) is measured by pumping and collecting directly on the antenna. In this image, we normalize by the Si line, which we take as a reference. (C) is measured by coupling in through the grating and collecting at the antenna. We again see that the signal is filtered by the cavity. (B, C) show that all Raman peaks follow the same trend, demonstrating that this signal is a result of pump enhancement on the BPT. The dashed black curve shows the results from the CHO model, showing good agreement with the data.

antenna except on cavity resonance, we now offer the pump light from free space, and collect light from the grating, in a reverse configuration of Fig. 4.14C. When aligning the $473\,\mathrm{cm}^{-1}$ Stokes-shifted line to the hybrid resonance, one would expect Raman scattered light to be transmitted through the photonic crystal nanobeam cavity, to be collected at the grating. Similarly to Fig. 4.15B we indeed see that the curve from the $473\,\mathrm{cm}^{-1}$ line shows strikingly different behavior from the other two. Whereas the 1073 and $1576\,\mathrm{cm}^{-1}$ lines simply follow a featureless shoulder, the $473\,\mathrm{cm}^{-1}$ line shows a strong Lorentzian resonance. This observation is consistent with the fact that there is no strong modification in the pump field as the laser is swept, and that the hybrid resonance provides the LDOS and the collection efficiency into the waveguide required to enhance and efficiently collect the Raman emission.

4.7. Coupled oscillator model for Raman scattering

Recently, Shlesinger et al. proposed a model [272] that extends the reach and scope of the coupled harmonic oscillator model discussed in Chapter 2 by adding terms that couple the cavity to input and output waveguides. In brief, for this model it is useful to cast the equations of motion in Eq. (2.4) into a symmetrized form that maps onto coupled mode theory [273]

$$\begin{cases} -i\left(\omega - \omega_a + i\gamma_a/2\right)a - iKc = \sqrt{\eta_{a,in}\gamma_a}\,s_{a,in} \\ -i\left(\omega - \omega_c + i\kappa/2\right)c - iKa = \sqrt{\eta_{a,in}\kappa}\,s_{c,in}. \end{cases} \quad (4.3)$$

Here, the symbols c, a label the cavity, resp. antenna, the second order order frequency dependence has been changed to first order by assuming $\omega_{a,c}^2 - \omega^2 \approx 2\omega(\omega_{a,c} - \omega)$ (limit of small detuning), and the quantities have been rescaled for symmetrization as

$$a \equiv \frac{\omega}{\sqrt{2\beta}}p, \quad K \equiv \sqrt{\frac{\beta}{2}}\frac{\tilde{E}_c}{2} = \frac{1}{2}\sqrt{\frac{\beta}{\epsilon_0\epsilon V_c}} \quad \text{and} \quad \sqrt{\eta_{a,in}\gamma_a}\,s_{a,in} \equiv \frac{i}{2}\sqrt{\frac{\beta_a}{2}}E_{ext}, \quad (4.4)$$

where V_c is the mode volume of the cavity mode at the position of the antenna. Compared to Eq. (2.4) the main change lies in the terms on the right-hand side, where $s_{a,c,in}$ quantify resp. the waveguide, and free space input-output channels. The $\eta_{a,c}$ parameters are coupling efficiencies that measure which fraction of antenna loss radiates into our objective collection NA, resp. which fraction of cavity loss is into the waveguide. Assuming coupling efficiencies $\eta_{a,c}$ this model can be elaborated to predict pump field enhancements, LDOS enhancements, and extraction efficiencies of SERS signal in waveguide and free space channels [272]. We compare our measurement results with results from this model. The input parameters for the model were chosen based on a number of sources. Previous transmission measurements yielded $Q = 520$ at $389\,\mathrm{THz}$ for this cavity, while dark field measurements yielded $Q = 10$ at $403\,\mathrm{THz}$ for the antenna (which pins the magnitude of the polarizability at $\beta = 0.1\,\mathrm{C}^2/\mathrm{kg}$, given that the antenna linewidth is limited by radiation damping). Full-wave COMSOL simulations yielded a mode volume of $8\,(\lambda/n)^3$ for the cavity and $10^{-4}\,(\lambda/n)^3$ for the antenna. We insert these values into the model, together with estimated coupling efficiencies for the cavity to waveguide coupling set at $\eta_c = 0.025$, a pump efficiency to the antenna of $\eta_{a,in} = 0.07$, and from antenna to our collection $\eta_{a,out} = 0.35$, at NA= 0.9. It should be noted that *constant* values for $\eta_{a,c}$ give rise to frequency-dependent

Fano structures in the extraction efficiency of SERS signal into the waveguides and into free space [122, 272].

The curves predicted by this model are shown in the black dotted curves in Fig. 4.14B, C, for the case of pump enhancement signatures. We find that our measurements can be well parametrized by this modified CHO model. We have also applied this model to the case of LDOS enhancement. As the coupling gratings and the antenna are unchanged, we use the same model parameters as above for the antenna and coupling properties, and only make changes to the nanobeam cavity parameters. We now use $Q = 500$ at 375 THz, with a slightly larger mode volume of $\tilde{V} = 11 \, (\lambda/n)^3$. The results of this model are shown in the black dashed lines in Fig. 4.15B and C, for the curve extracted from the $473 \, \mathrm{cm}^{-1}$ line. We again find a good agreement for the lineshapes in both observables. While one must be very careful when interpreting these results, given the coarse efficiency estimates, the fact that the model matches the lineshapes in all four configurations simultaneously is remarkable. This indicates that the CHO model provides a good understanding of the physics of pump enhancement, LDOS enhancement, and extraction efficiency. As long as the coupling constants for the grating couplers and antenna are unknown, measuring an absolute LDOS enhancement is out of reach. However, from the CHO model, it is possible to estimate the *ratio* of the LDOS enhancements of the bare antenna and the hybrid system. We find that hybridizing the antennas at hand with a nanobeam cavity reduces the LDOS by a factor of 3.34 at the dip of the Fano, and increases it by a factor of 1.35 at its peak. Much more substantial suppressions and enhancements are predicted by this model for geometries that use cavities with smaller mode volumes as starting point, and through optimization of the tuning of antenna and cavity relative to each other.

4.8. Conclusion

To conclude, we have reported the optical characterization of Si_3N_4 photonic crystal cavity hybrids with a narrow-gap dimer antenna, fabricated through a 2-step lithography process. We demonstrated that we can tune the resonance of the cavities and antennas with respect to each other by scaling their sizes, allowing us to detune the two components with respect to each other. The hybrids can be addressed through input and output waveguides to which we couple via a set of gratings. When doing so, we see a clear optical signal from the antenna. We observe that the presence of an antenna on a nanobeam cavity causes the resonance to redshift with respect to that of the bare cavity by $\sim 1.2 \, \mathrm{nm}$, and to broaden by $\sim 1 \, \mathrm{nm}$ on average. Stronger effects are observed for larger antennas, which is in line with predictions from perturbation theory. We see that almost all devices are functional and follow this trend, indicating a high success rate in the fabrication of the beam, the antenna, and the antenna placement.

In the second part of this chapter, we have shown that we can use SERS as a means to probe the hybrid properties in terms of field enhancement and LDOS. The SERS signal is highly selective for molecules that probe LDOS and field enhancement right at the antenna position, as it is only visible when addressing the antennas in the correct polarization. We have shown that we can effectively probe the properties of the nanobeam hybrids with SERS, and can make a clear distinction between SERS signal that is mainly enhanced through the pump, and enhancements originating from the narrow LDOS line of the hybrid resonance. There is good qualitative agreement

between the measured SERS signal and the signal predicted in a coupled harmonic oscillator model. While it is difficult to convert these observations into absolute field and LDOS enhancements, these observations validate our model understanding of hybrid resonators as developed in Chapter 2. Finally, we note that instead of viewing SERS as a probe of a hybrid nanostructures, there is also a strong interest in developing hybrids to leverage SERS. In particular, the paradigm of molecular optomechanics suggests that sideband-resolved SERS, wherein optical resonances are used that are narrower than the Raman shift, can be instrumental for applications of SERS beyond spectroscopy. These include transduction of THz and IR to visible light, cooling and parametric driving of vibrations, and vibrational strong coupling for chemistry. Our work evidences important steps to this goal, including sideband resolved SERS, and waveguide addressing of sideband resolved SERS hot spots. Important next steps will be to benchmark and improve SERS enhancement factors, and to realize multi-resonant systems in which *both* the pump *and* vibrational lines are controlled through on-chip pathways.

Towards the goal of improving Purcell enhancement factors, we note that the low index contrast between the Si_3N_4 and the surrounding media of PMMA and SiO_2 limit the confinement that can be achieved by such cavities. Free-standing nanobeams have the potential to achieve higher Q and smaller V, resulting in higher-quality hybrids. Likewise, the antenna properties can be enhanced by improvements in design and fabrication. This work demonstrates that Si_3N_4 photonic crystal nanobeam hybrids have the potential to harness light-matter interaction, perhaps to the point of achieving strong coupling. Especially if we could improve confinement and Q by working with suspended instead of immersed nanobeams and if the gap could be reduced by a factor ~ 2, these systems could contribute towards spanning the gap between photonic microcavities and plasmonic nanoantennas. This regime has a number of useful applications, including few/many-emitter strong coupling and single-particle sensing.

II

INTERMITTENT BEHAVIOR AND SINGLE PEROVSKITE QUANTUM DOT EMITTERS

5

A python toolbox for unbiased statistical analysis of fluorescence intermittency of multi-level emitters

"'Would you tell me, please, which way I ought to go from here?'
'That depends a good deal on where you want to get to,' said the Cat.
'I don't much care where -' said Alice.
'Then it doesn't matter which way you go,' said the Cat.
'- so long as I get SOMEWHERE,' Alice added as an explanation.
'Oh, you're sure to do that,' said the Cat, 'if you only walk long enough."
— Lewis Carroll, *Alice's Adventures in Wonderland*, 1865

5.1. Introduction

S INCE the seminal first observation of single molecule emitters in fluorescence microscopy three decades ago [274], single quantum emitter photophysics has taken center stage in a large body of research. On one hand, single quantum emitters as single-photon sources [27] are held to be an essential part of quantum communication networks and are deemed essential for building optically addressed and cavity-QED based quantum computing nodes [26]. This has particularly spurred research in III-V semiconductor quantum dots [28, 168], color centers in diamond, silicon carbide, and 2D materials [275–277], and organic molecules at low temperature [127]. On the other hand, classical applications of ensembles of emitters for displays, lighting, lasers, and as microscopy-tags drive the continuous development of new types of emitters, such as II-VI self-assembled quantum dots 20 years ago [267, 278, 279], and inorganic perovskite quantum dots just recently [156, 157, 280–285]. For all these systems, understanding the photophysics on the single emitter level is instrumental, whether the intended use is at the single or ensemble level. A common challenge for almost all types of emitters is that they exhibit intermittency, also known as blinking [286, 287]. Under constant pumping emitters switch, seemingly at random, between brighter and dimmer states, often corresponding with higher and low quantum yield (QY), and different fluorescence decay rates. Frequently the switching behavior also shows peculiar, power-law distributed, random distributions of durations of events. Determining the mechanism through which emitters blink, i.e., the origin of the involved states, the power-law distribution of residence times, and the cause of switching, have been the topic of a large number of studies particularly for II-VI quantum dots as recently reviewed by Efros and Nesbitt. [288]. Recent studies on inorganic perovskite quantum dots uncover intermittency behavior that does not fit common models for intermittency in their II-VI counterparts [157, 280–285].

In order to quantify intermittent behavior, the simplest and most commonly employed method is to subdivide a measurement stream of individual photon-arrival times into short bins of a few ms, to calculate the intensity (in counts/second) of each bin. Every bin can then be assigned to a state (on, off or grey) according to its brightness so that on/off times as well as intensity levels can be defined and analyzed [286]. For pulsed laser excitation, also quasi-instantaneous fluorescence decay rates can be obtained [266, 289]. However, it is well known that this method of binning time streams and histogramming binned intensities causes detrimental artefacts [286, 290–296]. Retrieved parameters of the quantum dot behavior often exhibit a dependency on the choice of the bin width, which affect estimates of switching time distributions and power laws, and also the objective assignment of intensities to intrinsic levels. Narrower bin widths in principle allow better resolution, but run into shot noise limits, while conversely choosing larger bins suppresses noise, but will render the analysis blind to fast events.

To overcome these issues, Watkins and Yang [290] proposed changepoint analysis (CPA) as a Bayesian statistics approach for the unbiased determination of switching times that is optimal, i.e., gives the best performance given the constraints of shot noise in the data. Changepoint analysis and clustering is one example of Bayesian inference methods to determine the transitions and underlying levels in single photon trajectories. In the domain of high-throughput single-molecule analysis [297] many methods to process single photon trajectories have appeared, that one can classify as supervised
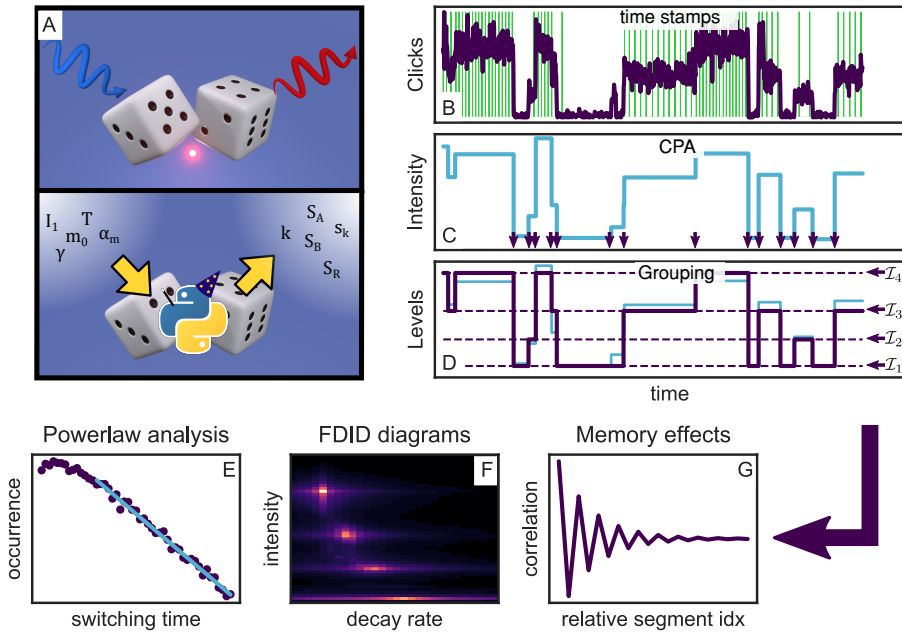
Figure 5.1: A schematic overview of the working of the toolbox. (A) An illustration of the two methods available to obtain TCSPC data, either by photoluminescence TCSPC measurements of a single emitter (top panel) or through simulation of a single emitter (bottom panel). The latter is provided in the toolbox. Both will result in a stream of timestamps that can then be further analyzed by the toolkit. The simulation part of the toolbox simulates dots of m_0 levels, with associated count rates and fluorescence decay rates I_m and γ_m, with the dot visiting levels in random order, and with residence times for each segment chosen drawn according to specified power law exponents α_m. The output simulated data consists of time stamped photon arrivals over a total time span T, where for each photon $k = 1, 2, 3 \ldots$, a time stamp s_k is recorded. These time stamps are randomly distributed over two detector channels S_A, S_B. The delays of each of the photon time stamps relative to the time stamp in the third channel S_R, representing the periodic pump laser pulse train, is chosen in accordance with the set emitter decay rate. (B) Starting out with this stream of photon events, using CPA (C), the changepoints are found, and with these, the instantaneous intensities. (D) Subsequently, these events are grouped in order to find the most likely underlying intensity levels between the behavior. Following this, a number of analyses can be done, such as (E) ascertaining whether the time between switching events is power law distributed, (F) the visible separation of states in FDID/FLID (fluorescent decay rate intensity diagrams, resp. fluorescent lifetime intensity diagrams), and (G) the presence of memory in the switching behavior. Here we show an example of a simulated, 4 level emitter, with simulation parameters chosen for clarity of illustration.

learning methods with a priori model assumptions on one hand, and unsupervised approaches on the other hand. Prominent are so-called hidden Markov Model (HMM) methods [298] that view photon data streams as experimentally measured output of transitions between hidden transition states. Bayesian inference can then estimate parameters such as transition probabilities if one a priori postulates the number of levels and the allowed transitions. As this underlying model is a priori often not known, one can apply HMM with different possible models, and rank them according to probabilistic criteria, such as the Bayesian Information Criterion. Also, the requirement for a priori known models is relaxed in so-called aggregated Markov models [299], and non-Markov memory kernel models [300].

Juxtaposed to such supervised analysis methods are unsupervised approaches. Such methods apply changepoint analysis to partition data into time segments between jumps, and subsequent clustering of intensity levels. The CPA method pioneered by Watkins and Yang is essentially such a combination of changepoint detection and hierarchical agglomerative clustering, using the Bayesian Information Criterion to determine the best clustering of measured intensities in distinct levels (states) with as sole assumption that intensity in each segment of time wherein the emitter is in a given level, the counts are Poisson distributed [290, 293]. A main drawback is that particularly the clustering is slow. We refer to Ref. [301] for recent developments in machine learning to mitigate this problem. A main advantage of CPA is that no underlying model is required, and that the data is segmented and clustered to the level that the data allows, given that the data is Poisson distributed in intensity, and given a required confidence level stipulated by the user. Variations on CP for other types of noise, such as Gaussian noise, have also appeared [302, 303]. Despite the well-documented superior performance over binning of photon counting data, in the domain of single photon counting data from quantum dots only very few groups have adopted these methods [152, 282, 290, 292, 293, 296, 304].

In this chapter, we provide, benchmark, and document a Python toolbox for changepoint analysis, state clustering, and analysis of fluorescence-intensity-decay rate correlations that is posted on GitHub [305]. A main motivation lies in the emergence of new quantum emitter systems with complex photophysics. While II-VI quantum dots for which CPA was originally developed, are generally understood to switch between just two or three states, the problem of accurate analysis of intermittency is gaining in prominence with the advent of novel emitters, such as perovskite quantum dots which appear to switch not between just two, but instead a multitude of states [157, 280–285]. There is hence a large need for a toolbox that provides unbiased, model-free analysis of photon counting data, for which reason we provide a CPA implementation and benchmark it for complex multilevel emitters. Our toolbox is both applicable to real quantum dot data, and valuable as a testbed for both testing models and analysis techniques on synthetic, i.e., numerically generated data. Indeed, the code toolbox includes code to generate numerically random photon time arrival data streams for 'synthetic' quantum dots that jump between an arbitrary set of intensity levels and decay rates, with jump time statistics and photon budgets that can be set by the user.

As results we provide benchmarks on the performance of CPA for detecting change points in function of the number of intensity levels and total photon budget, and we explore the limits to the number of distinct states that the clustering analysis can reliably separate. Moreover, the toolbox allows to us test the accuracy of jump time statistics, such as power law statistics, for such multilevel dots. Finally, the

test suite also allows to benchmark the accuracy of fluorescence decay model fitting with maximum likelihood estimation, and we discuss the construction of fluorescence-intensity-decay rate correlations from CPA-partitioned data. The applicability of the toolbox to experimental data will be illustrated in Chapter 6. We note that some other approaches such as HMM methods may be more suited for processes where more knowledge on the underlying physical processes is available. In contrast, our toolbox is ideal for cases in which one wants to make no a priori assumption on the physical mechanism behind intermittency. Furthermore, the CPA method in the toolbox operates on the finest level of information available in photon counting, i.e. the distribution of individual photon arrival times, as opposed to methods that are optimized to work on camera frame data [297]. This Chapter is structured as follows. In the Methods section we summarize the Bayesian statistics tools we implemented to analyze all aspects of our data. Next, we benchmark the performance of changepoint analysis to pinpoint intensity jumps, and of level clustering to identify the number of levels between which a dot switches on the basis of Monte Carlo simulations. Next we present considerations on the dependence of on-off time distributions, decay rate fit, and so-called 'fluorescence decay-rate versus intensity diagrams'(FDIDs) on count rates.

5.2. Methods

In this Section, we present all the methods implemented in our Python toolbox, as well as the methods for benchmarking them. Benchmark results are presented in the results section. We refer to the supporting information of [306] for a full manual to the code and to our GitHub repository [305] the code itself.

5.2.1. Changepoint analysis

First, we summarize changepoint analysis (or CPA), a Bayesian statistics method for the unbiased determination of jumps or 'changepoints' in time traces of discrete events [150, 152, 153, 282, 290, 292, 293, 296, 303, 307–309]. Bayesian statistics is a paradigm that reverses the usual standpoint of probability theory. Usual probability theory views a data set as a *random draw* from a probability distribution, given a hypothesis on the parameters of the underlying physical process. In this framework, one can calculate the likelihood of drawing the specific measured data set. Bayesian statistics, on the other hand, compares the likelihood of distinct hypotheses, *given* a measured data set and assumptions on the underlying measurement noise.

We consider time-tagged single photon counting data consisting of an ordered list of measured photon arrival times s_k, collected over a measurement time T. For a single emitter with no memory that emits at a count rate of N photons in a time T, the waiting times - i.e., the times between photon arrivals - are exponentially distributed with waiting time $\tau_w = T/N$. In order to determine whether there is a changepoint in some segment q, CPA compares the likelihood of two distinct hypotheses, (1) there is a jump in emission intensity (i.e. the average waiting time τ_w jumping from some value to another) against (2) there is the same intensity throughout the measurement interval. When testing for a jump at photon detection event k at time s_k in this trajectory q with time duration T_q containing N_q photon events, this leads to a log-likelihood ratio, or

'Bayes factor' [290, 292, 293]

$$\mathcal{L}_k = 2k \ln \frac{k}{V_k} + 2(N_q - k) \ln \frac{N_q - k}{1 - V_k} - 2N_q \ln N_q, \tag{5.1}$$

where $V_k = s_k/T_q$. Derivation of this log-likelihood ratio involves several steps. First, it incorporates the assumption that in between jumps, the waiting time between photons is exponentially distributed, on basis of which one can assess the likelihood of measuring the given data set for a given hypothesis on the exponential waiting time τ_w. Second, it uses maximally non-informative priors for $\mathcal{L}_k$ to compare the hypothesis of presence versus absence of a changepoint without further restrictive assumptions on the involved intensity levels.

It should be noted that there are other ways to arrive at the same log-likelihood ratio test. One alternative starting point is a binary time series in which there is an underlying uniform and small probability distribution of photon detection per bin (e.g., imagining the time axis binned in by the timing card resolution (of order 0.1 ns for typical hardware)) [304]. Such a uniform distribution would emerge as a direct consequence of exponential waiting time distributions. In this case one should start from a binomial distribution and ultimately arrives at the same formula after application of Stirling's formula. Another starting point is CPA applied to binned data with wider bins with multiple counts, i.e. to series of Poisson distributed intensities instead of discrete events [290, 293]. However, the binning would introduce an undesirable time scale through the chosen bin width. Of these three methods, working with photon arrival times is the most data-efficient approach and introduces no artificial partitioning whatsoever. We refer to Ref. [293] for a derivation of the log-likelihood ratio in all these three scenarios, which includes a precise description of the use of maximally non-informative priors.

Following Watkins and Yang [290] and Ensign [293], the most likely location of a changepoint, if any, is at the k that maximizes the Bayes factor $\mathcal{L}_k$. The hypothesis that this most likely changepoint is indeed a real event is accepted if $\mathcal{L}_k$ exceeds a critical threshold value for $\mathcal{L}_k$ or 'skepticism'. This value is chosen to balance false positives against missed events. A full data set is partitioned recursively, i.e., by recursively checking if data sets between two accepted changepoints themselves contain further changepoints. This results in a division of the data set into segments, each of which starts and ends at an accepted changepoint, and with the level of skepticism as stop criterion for the recursion. The resulting segmentation provides the most likely description of data as consisting of segments within which the intensity is constant, given the value chosen for the degree of 'skepticism', and given the amount of data collected. Since the algorithm works with the list of individual photon arrival times this segmentation entails no arbitrary partitioning. An accepted rule of thumb is that if the Bayes factor $\mathcal{L}_k$ exceeds a 'skepticism' value of just between 1-3, the evidence for a changepoint is highly ambiguous, while values in the range 7-10 are deemed strong evidence. The toolbox is supplied with a default value of skepticism of 8, set following the analysis of [290] and [293]. The reader is warned that for a given photophysics scenario (intensity levels, segment duration statistics) it is advisable to set the level of 'skepticism' on basis of simulations, in order to optimize the trade-off between missing changepoints altogether (false negatives) and precision (avoiding false positives). Our result section provides an example of such an optimization.

5.2.2. Clustering

Changepoint analysis splits the data into segments separated by jumps (a list of Q jumps delineate $Q - 1$ segments). One can now ask what the statistical properties are of the segmentation, i.e., what the statistics are of the length of segments, the intensity levels most likely corresponding to the segments, and the fluorescence decay times associated with the segments. For instance, it is a nontrivial question how many distinct constant intensity levels, or states, m_r actually underlie the $N - 1$ found segments, with intensities $I_1 \ldots I_{Q-1}$. To answer this question Watkins and Yang [290] proposed a clustering approach. The recent work of Li and Yang [303] provides a detailed explanation of the reasoning involved, though quoting results for Gaussian instead of Poissonian distributed data. The idea is that with the $Q - 1$ found segments, each with their associated recorded intensities I_q one can use expectation maximization to calculate, for a hypothesized and fixed number of levels n_G what the most likely underlying intensity levels $\mathcal{I}_m$ are (with $m \in 1 \ldots n_G$), and how probable it is that each segment is ascribed to a given level (probability p_{mq}). Subsequently, Bayesian inference is used to establish what the most likely number of levels (i.e. states) m_r and associated intensities $\mathcal{I}_m$, with $m \in \{1 \ldots m_r\}$ is that describes the data.

Following Ref. [290, 303], the expectation minimization in our toolbox is implemented as an iterative algorithm started by a first guess of the segmentation. This guess is obtained by a hierarchical clustering of $Q - 1$ segments in $m = 1, 2, \ldots Q - 1$ levels that proceeds recursively. In each step, it identifies the two segments in the list with the most similar intensity levels as belonging to the same level. This provides an initial clustering of the measured data in any number $m = 1, 2, \ldots Q - 1$ of levels. For the expectation maximization, the idea is to simultaneously and iteratively optimize the probability p_{mq} for segment q to belong to the mth level, as well as an estimate of the intensities of these levels $\mathcal{I}_m$. In each iteration, the intensities of all levels are estimated from the level assignment from $p_{m,q}$. Following this, the probability distribution p_{mq} is updated to redistribute the segments over the levels. In this calculation, it is important to understand the type of noise statistics the data obeys. In the case of single-photon measurements, and for the purpose of this discussion, the intensities are Poisson distributed. The iteration is repeated until p_{mq} converges (practically also capped by a maximum number of iterations). The final outcome is a most likely assignment of the measured segments into n_G levels. Next for each value of nG one assesses the 'Bayesian information criterion' (BIC). This criterion is a measure for how good the description of the segmented intensity trace is with n_G intensity levels given the assumption of Poisson counting statistics for each fixed intensity level. Beyond a mere 'goodness of fit' metric that would simply improve with improved number of parameters available to describe the data, this metric is penalized for the number of parameters to avoid overfitting. For Poisson distributed data the criterion is derived in [290] as

$$\text{BIC} = 2\mathcal{L}_{EM} - (2n_G - 1)\ln Q - Q \ln N \qquad (5.2)$$

where Q again is the number of change points detected, n_G is the number of available levels. The term $\mathcal{L}_{EM}$ is the log-likelihood function optimized in the expectation maximization step, i.e $\mathcal{L}_{EM} = \sum_q \sum_{m=1}^{n_G} p_{mq} \ln[p_m \mathcal{P}(I_q; \mathcal{I}_m)]$ with $\mathcal{P}(x; \lambda)$ the Poisson probability function at mean λ, p_m the probability of drawing level m. The second term in the BIC is the term penalizing the BIC for overfitting. The accepted best description of an emitter in n_G-levels is taken to be at the value of n_G where the BIC peaks.

5.2.3. Intensity cross/autocorrelation and maximum likelihood lifetime fitting

Many single photon counting experiments are set up with pulsed laser excitation for fluorescence decay rate measurements, and with multiple detectors to collect intensity autocorrelations (e.g., to verify antibunching in $g^{(2)}(\tau)$ for time intervals τ comparable to the fluorescence decay rate, and shorter than the commonly longer detector dead time). In a typical absolute-time tagging set up, this results in multiple data streams S_A, S_B, S_R of time stamps corresponding to the detection events on each detector, and the concomitant laser pulses that created them, respectively. Our Python toolbox contains an implementation of the correlation algorithm of Wahl et al. [310] that operates on timestamp series, and returns, for any combination of channels S_1, S_2 ($1, 2 \in \{A,B,R\}$), the cross correlation $C(\tau)\Delta\tau$, i.e., the number of events in the time series S_1 and the time series S_2 that coincided when shifted over τ, within a precision $\Delta\tau$.

Cross-correlating detected photons and laser arrival times, taking $\Delta\tau$ to be the binning precision of the counting electronics and the range of τ equal to the laser pulse repetition rate, returns a histogram of the delay times between photon detection events and laser pulses. To obtain $g^{(2)}(\tau)$ to investigate antibunching, streams of photon events from two detectors in a Hanbury-Brown Twiss set up are cross-correlated. $\Delta\tau$ is taken to be the binning precision of the counting electronics and the sampled range of τ as an interval is taken symmetrically around $\tau = 0$ and several times the laser pulse interval. Finally auto- or cross correlating detector streams over τ-ranges from nanoseconds to seconds, coarsening both τ and $\Delta\tau$ to obtain equidistant sampling on a logarithmic time axis, results in long time intensity autocorrelations of use in intermittency analysis [295]. Our toolbox also provides this logarithmic time-step coarsening version of the correlation algorithm of Wahl et al. [310]

Of particular interest for intermittent single emitters is the analysis of fluorescence decay rates in short segments of data as identified by changepoint analysis, that may be so short as to contain only 20 to 1000 photons. For each of the photon detection events in a single CPA segment, cross-correlation with the laser pulse train yields a histogram of the N_q photons in segment q. In each of the bins (with width $\Delta\tau$) the photon counts are expected to be Poisson distributed. Therefore the optimum fit procedure to extract decay rates employs the Maximum Likelihood Estimate procedure for Poisson distributed data, as described by Bajzer et al. [311]. In brief, for a decay trace sampled at time points τ_i relative to the laser excitation, with counts per bin $D(\tau_i)$, the merit function reads

$$M = - \sum_{\text{all data bins } i} \{D(\tau_i) \log [F_A(\tau_i)] - F_A(\tau_i)\}. \tag{5.3}$$

Assuming a chosen fit function $F_A(t)$ the parameter set A that minimizes this merit function provides the parameter values that most likely correspond to the data. The estimated errors in these parameters then follow from the diagonal elements of the inverse of the Hessian of M relative to the parameters A. Importantly, the fact that the Poisson distribution is tied to absolute numbers of counts, implies that this approach requires that the data is *neither* scaled nor background subtracted. Instead, the background should be part of the fit function either as a free parameter or a known constant. Furthermore, it should be noted that time bins with zero counts are as informative to the fit as non-empty ones, and should not be left out.

5.2.4. Generating synthetic quantum dot data

To benchmark the CPA and clustering method and to test its limits, our toolbox provides an example routine to generate artificial data mimicking quantum dot intermittency. To obtain mimicked quantum dot data, we first choose a number m_0 of intensity levels $\mathcal{I}_{0,m}$ between which we assume the dot to switch. Next, we generate switching times for each of the states. In this work, we choose all switching times from a power law distribution. For benchmark purposes, we will present results with power law exponent $\alpha = 1.5$, though any exponent can be set in the code. On the assumption that intensity levels appear in a random and uncorrelated order, this segments the time axis into a list of switching events $T_{0,j}$, $j = 1, 2, \ldots m_0$, where for each segment we randomly assign one of the nominal intensities $\mathcal{I}_{0,m}$. Next, to mimic a pulsed excitation experiment, we imagine each of these segments to be subdivided in intervals of length τ_L equivalent to a laser repetition rate ($\tau_L = 100$ ns in the examples in this work). We assign each of these intervals to be populated with one photon at probability $p_m = \mathcal{I}_{0,m}\tau_L$. This ensures that the number of photons N_q in every segment is drawn from a Poisson distribution at mean $T_{0,q}\mathcal{I}_{0,m}$. By removal of all empty bins, the binary list is translated into a list $S = (t_1, t_2, t_3, \ldots)$ of photon arrival time stamps at resolution τ_L to which one can directly apply CPA to attempt a retrieval of switching times, and apply clustering to retrieve the number of states.

To also enable fluorescence decay rate analysis, we further refine the photon arrival time list. Recalling that we have generated switching events $T_{0,q}$ between intensity states $\mathcal{I}_{0,m}$, we now also assume fluorescence decay rates $\gamma_{0,m}$. As each segment q, was already chosen to correspond to some level m_q, we now impose decay rate γ_{0,m_q} on the photon arrival times. To do so, for each of the photon events $k = 1 \ldots N_q$ already generated at resolution τ_L we now randomly draw a delay time Δ_k relative to its exciting laser from an exponential distribution characterized by rate γ_{0,m_q}. To mimic the behavior of typical TPSPC counting equipment, the delay time is discretized at a finite time resolution $\Delta\tau$ (in this work chosen as 165 ps to match the hardware in the provided example experimental data measured in our lab (Becker & Hickl DPC-230), though of course in the toolbox the value can be set to match that of any TCSPC card vendor). Testing of fluorescence decay trace fitting can operate directly on the generated list of delay times, or alternatively one can synthesize a TCSPC experiment by re-assigning S to represent laser-pulse arrival times S_R, and defining photon arrival times as the events in S_A, S_B each shifted by its delay time, i.e. $S_X = (t_1 + \Delta_1, t_1 + \Delta_2, \ldots)$, with X $\in$ A,B. Cross correlation of S_R and S_X returns the delay time list. We note that although our work does not focus on antibunching, our quantum dot simulation routine provides data distributed over two detector channels, where emission events antibunch, while an uncorrelated background noise level of the detectors can also be set.

5.2.5. Practical implementation

We have implemented the toolbox ingredients in Python 3.8. As timestamp data can be substantial in size, we use the 'parquet' binary format to store timestamps as 64-bit integers. Processing and plotting the data is dependent on Pythons' standard libraries numpy, matplotlib, while we use Numba, a just-in-time compiler, to accelerate the time-stamp correlation algorithms. An example script to generate synthetic data, and to run the entire workflow on simulated data is provided. We refer to the supporting information for a guide to the practical implementation and use of our toolbox. The

toolbox comes also with example experimental data on single $CsPbBr_3$ quantum dots from the experiments in Chapter 6. Table 5.1 list scaling and performance metrics for the algorithms contained in the toolbox.

5.3. Results

The remainder of this work is devoted to presenting benchmarks of the provided methods. Benchmarks for emitters with 'binary' switching, i.e., two well-separated intensity levels as is typical for II-VI quantum dots, have already been presented in literature [290, 296, 304]. However, emitters under current study, such as perovskite quantum dots appear to have a multitude, or perhaps even a continuum, of intensity levels. Our tests hence focus on determining the performance of CPA and level clustering for many-level single photon emitters.

5.3.1. Precision of identifying individual changepoints

Figure 5.2(A) and (B) show examples of CPA analysis applied to a simulated quantum dot with a single jump in its behavior, from an intensity level 10^4 to 10^3 cts/s resp. from 4.5×10^2 to 2.25×10^2, with ~ 900 and ~ 150 photons left and right of the changepoint, respectively. Purely for visualization purposes, the data is plotted in a binned format, as the analysis itself does not make use of any binning. Alongside the binned intensity trace, we also show the log-likelihood ratio $\mathcal{L}_k$. In both cases, the log-likelihood ratio clearly peaks at or close to the point where there is a changepoint in the data. Since the Bayes factor is actually a *logarithmic* measure for the comparison of hypotheses, the algorithm indeed identifies the changepoint with high probability and to within just a few photon events, even where the jump is far smaller than the

Algorithm	Scaling	Timing
CPA	$\mathcal{O}(N\sqrt{N})$	Ca. 10 s for $N = 10^6$
Grouping	Initial clustering dominates over iterative algorithm	
Initialization∗	$\mathcal{O}(Q^2)$	0.5 s for $Q = 3 \times 10^2$
Iterative optimization†	Overhead dominated	0.75 ms per iteration ($Q < 500$)
$g^{(2)}$‡	$\mathcal{O}(Nn_{\text{plot points}})$	8 ms per plot point at $N = 10^6$
Long-time autocorr.	$\mathcal{O}(Nn_{\text{cascade}})$	Ca. 2s for full plot

∗ For large Q this is accelerated by first clustering subsets, merging, and continuing clustering
† Typically 5 to 10 iterations are required when $n_G \approx m$.
‡ Numba JIT acceleration assuming int64 provides over 2 orders of magnitude accelation. A typical $g^{(2)}$ plot has ca. $n_{\text{plot points}}) = 2000$, and hence requires 10 to 20 seconds to evaluate.
  Essentially repeating $g^{(2)}$ and logarithmic coarsening every n_{cascade} points.

Table 5.1: Algorithm scaling and computation time as function of the total number of photons N and total number of changepoints Q. Timings were obtained on a standard desktop [Intel I7 4790 at 3.6 GHz, with 16 Gb of DDR3 RAM), and are obtained on basis of 5×100 photon trajectories (100 independent draws for 5 different trajectory record lengths with from ca. 10^4 to 3×10^6 photon events).
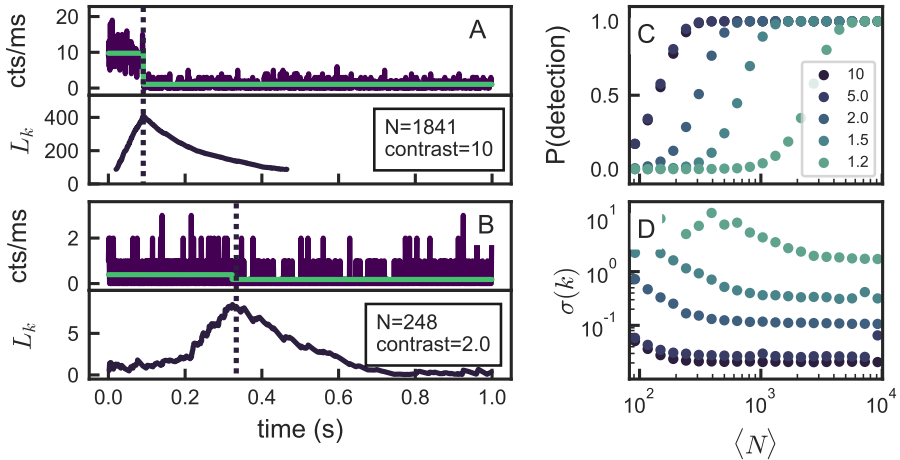
Figure 5.2: Demonstration of changepoint detection applied to a synthesized data set with a single changepoint, with equal photon counts before and after the changepoint. In (A) and (B) the contrast between intensities is a factor of 10 and 2, respectively, while the total photon budget is approximately 2000 and 300. The bottom panels show the log-likelihood ratio test, which clearly peaks at the changepoint in both cases. The y-axis unit cts/ms stands for counts per millisecond. The robustness of the method is demonstrated in (C) and (D), where we show the likelihood of detecting a changepoint in such a series for different intensity contrasts, and the variance of the found times, respectively, as a function of the total photon numbers. To gather accurate statistics, 10^4 photon traces were generated for each data point. The data is plotted in a binned fashion (1 ms bins) for visualization purposes only.

shot noise in the binned representation in the plot, at a relative intensity contrast of just a factor of 2. Generally, the probability with which the algorithm identifies or misses the changepoint is dependent on the total number of photons recorded both before and after the changepoint, and on the contrast in intensities, consistent with the findings of Watkins, and Ensign [290].

5.3.2. Intermittency and on-off time histograms

To identify the limits of CPA [293, 296] we consider the feasibility of identifying changepoints of contrast I_2/I_1 as a function of the total number of photons in the time record. The results are shown in Figure 5.2C for the likelihood of detecting a changepoint, and Figure 5.2D for the error in identifying the precise event k at which the changepoint that is identified occurred. Here, we only consider the case where there are roughly an equal number of photon events before and after the changepoint. This data is obtained by simulating 10^4 switching events of the type as shown in Fig 5.2(A,B) for each contrast and mean photon count shown. The range of contrasts is chosen commensurate with reported on-off contrasts for typical quantum dots in

Figure 5.3: (A) Typical time trace of a simulated quantum dot. The intensity
duty cycle switches between 0.5×10^4 and 2×10^4 counts/s. It shows an
on/off input duty cycle generated with a power-law distribution (orange), the
duty cycle with Poissonian noise (purple), and the retrieved duty cycle (green).
Overall, the original intensities and lifetimes (B) are retrieved well.

literature, which generally fall in the 1.5 to 5-fold contrast range. At high intensity
contrast, exceeding a factor 5, a total photon count as low as 300 is enough for near-
unity detection. Moreover, for sufficiently high photon count left and right of the
changepoint, even very small changes in intensity have a high likelihood of being
accurately detected, even if in binned data representations the jump is not visible
within the shot noise. Figure 5.2(D) provides a metric for the accuracy within which
changepoints are pinpointed. Changepoint analysis returns the most likely photon
event k in which the jump occurred, which in our analysis can be compared to the
actual photon event index k_0 at which we programmed the Monte Carlo simulation to
show a jump. Figure 5.2d reports the mean error ($\sqrt{\text{var}(k - k_0)}$) as a metric of accuracy.
At jump contrasts above a factor 2, changepoints are identified to within an accuracy of
almost one photon event even with just 10^2 photon counts in the total event record. At
very small contrasts, the error in determining the location of a changepoint is generally
on the level of one or two photon events, only worsening when there are fewer than
200 counts. This observation highlights the fact that if the photon record has just a few
counts in total, the error in estimating the count rate before and after the jump becomes
comparable to the magnitude of the jump.

As next step in our Monte Carlo benchmarking we turn to time series with many,
instead of single, jumps. Fig. 5.3A shows a representative example for a simulated
intermittent quantum dot with two states, assuming a contrast ratio between states
of 2×10^4 and 5×10^3 s^{-1}. We generally observe that the recursive CPA algorithm
accurately identifies switching events, barring a number of missed events of very short
duration.

From the CPA analysis we retrieve the time duration between switching events.
Figure 5.4A shows a histogram of time durations, plotted as a probability density
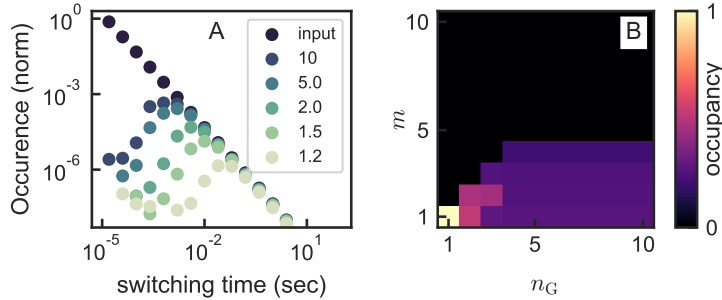
Figure 5.4: (A) A histogram of the number of switching events as a function of their duration. Each data point represents 10.000, 10 s power-law distributed time traces. The input shows the initial power-law distribution, the lighter colors show the number of retrieved changepoints, after applying Poisson noise and CPA, at different contrasts, with $I_1 = 10^5$. counts/s. We can see that even at low contrast, events with long times between switching are retrieved, but each contrast has a characteristic duration below which changepoints can't be accurately retrieved. This puts a fundamental limit on the information that can be extracted from a given data set. (B) An occupancy diagram illustrating the behavior of the clustering algorithm for a system with m_0=4 intensity levels. The color scale indicates the amount of time spent in each state m_i after the assignment of states for a given number of available states n_G. We see that when $n_G > m_0$, effectively all segments are distributed across only $n_G \leq m_0$ intensity levels.

function obtained from a whole series of Monte Carlo simulated time traces of varying contrast between states (see legend). Notably, if we simulate quantum dots that have switching times that are power-law distributed, the retrieved distribution indeed follows the assumed power-law, particularly for long times. At shorter times, the histogram remains significantly below the power law, particularly at low intensity ratios between the two assumed states. This indicates that CPA misses fast switching events, and is consistent with the observation from Figure 5.2C that a minimum photon count is required to observe switching events of a given contrast. As a rule of thumb, usual II-VI colloidal quantum dots have a contrast between dark and bright states of around 5, meaning that of order 200 photons are required to detect a change point with near-certainty. At the assumed count rates ($2 \times 10^4 \, \mathrm{s}^{-1}$ for the bright state) this means one expects CPA to fail for switching times below 10 ms, where the on-off time histogram indeed shows a distinct roll-off. This result suggests one should interpret on-off time histograms from changepoint detection with care: one can generally rely on the long-time tail, but should determine the shortest time scale below which the histogram is meaningless on basis of the intensity levels present in the data.

5.3.3. Error analysis for trajectories with multiple jumps

The changepoint analysis results in Figure 5.2(C) essentially quantifies the algorithm performance in terms of the fraction of correctly identified change points (true positives) for traces with a single step in intensity, as a function of contrast and photon budget. Actual single emitter photon trajectories have a plethora of steps, where CPA is mainly likely to miss short segments as only those changepoints are accepted for which the evidence in the data is compelling, relative to the shot noise in it. Indeed, the short-time roll-off in Figure 5.4A highlights exactly this tendency of CPA to under-report on closely-spaced changepoints (false negative rates high for short segments). The level of skepticism set as parameter for running CPA sets the overall accuracy of the algorithm, essentially trading off the rates of false positives, and false negatives. When using the toolbox for a particular photophysical scenario, the reader is recommended to study the error rates as a function of skepticism. To demonstrate that type of study, here we report on algorithm performance as a function of skepticism using the error metrics *accuracy*, *precision* and *recall*. To this end, we generate synthetic data and match the list of nominal changepoints and retrieved changepoints to determine the rate TP of true positives, the rate FP (false positives) of detected transitions for which no transition was actually present, and the rate FN of false negatives, in which a true transition is not detected by CPA. The standard definition for the error metrics reads [298, 301]

$$\text{accuracy} = \frac{\text{TP}}{\text{TP+FN+FP}} \tag{5.4}$$

$$\text{precision} = \frac{\text{TP}}{\text{TP+FP}} \tag{5.5}$$

$$\text{recall} = \frac{\text{TP}}{\text{TP+FN}}. \tag{5.6}$$

The accuracy benchmarks overall performance, whereas precision measures the false positive error rate, and recall quantifies the false negative rate. Since firstly changepoint detection is not accurate to the level of a single photon arrival time, and secondly the set of stored nominal switching times in our toolbox may fall in between synthesized

Figure 5.5: Accuracy (A), precision (B), and recall (C) for changepoint analysis as a function of tolerance in milliseconds for which nominal and detected changepoints are matched as equal, and as a function of the level of "skepticism" which the Bayes factor needs to exceed for a changepoint to be accepted.

Figure 5.6: Examples of the BIC for simulated quantum dots, with $m_0 = 2, 3, 4$. The most likely number of levels is indicated by a small peak in L_k. In the inset, we show that the distributions indeed peak at their respective m_0. It should be noted that the BIC shows a very sharp rise, and then from m_0 onward appears almost flat. There is in fact a shallow downward slope. Here one should keep in mind that the BIC is a logarithmic metric. On a linear scale the maximum is significant.

photon events, such a comparison requires a tolerance range to be meaningful. Figure 5.5 presents the algorithm performance as a function of the level of skepticism (vertical axis), and as a function of the tolerance range within which change points are accepted as true positives, measured in milliseconds. The results are for a more challenging case than a two-level dot, namely a 4 level system with mean count rate 5×10^4 counts per second, and 4 equidistant intensity levels (2, 4, 6 and 8×10^4 counts per second), and power law distributed segment residence times (exponent 1.5, with shortest residence time of 10 ms in a segment). Presented results are obtained from 200 photon trajectories with on average 5×10^5 photons and 10^2 changepoints each.

If the time axis for the tolerance is chosen as short as the inverse mean count rate, the apparent algorithm precision is low, indicating that changepoints are generally found close to, but not quite at, the moment where the switching event occurs. At tolerances of 2 to 5 ms (containing of order 50-250 photons typically at the given rate, and for the various assumed intensity levels) the error rate saturates at above 90%. The *accuracy* for this example peaks at a skepticism of ca. 7.0 (Figure 5.5A). At higher levels of skepticism, the *precision* increases, i.e., the number of false positives reduces further (Figure 5.5B). However, this is at the expense of recall, i.e., the number of missed changed points. The false negatives rate decreases only with skepticism lowered to below 10 (Figure 5.5C).

5.3.4. Performance of level clustering

Next, we consider the performance of the grouping algorithm applied to the segmentation of simulated time traces. For reference, Figure 5.6 shows the Bayesian Information Criterion versus n_G for the example of simulated dots with $m_0 = 2, 3, 4$ intensity levels. Generally, the BIC rapidly rises as n_G approaches the actual number of levels with which the data was simulated, and gently decreases once n_G exceeds the actual number of
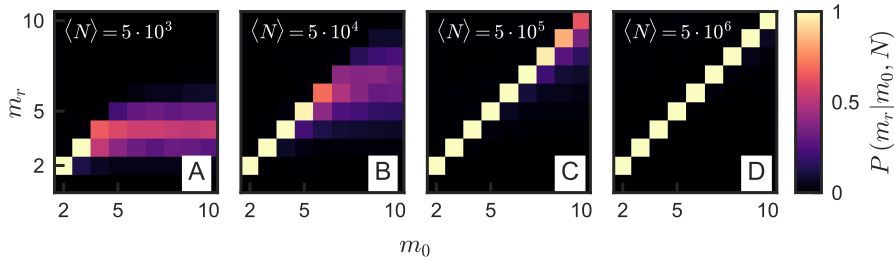
Figure 5.7: Likelihood P of finding m_r levels for a simulated system, given m_0 initial intensity levels and $\langle N \rangle$ the mean number of photon counts. We see that the photon budget plays a defining role in the total number of states that one can reliably resolve. At low photon budgets, the number of levels is systematically underestimated, whereas at high photon budgets, $P(m_i = m_0)$ remains high even at high m_0. We see that m_r is never overestimated.

levels in the data m_0. The fast rise indicates that within the assumption of Poisson distributed intensities, the data can not at all be described by fewer than m_0 levels. The slow decrease is due to the penalization of the BIC by the number of fit parameters. Since the BIC criterion actually relates to the *logarithm* of the probability with which n_G states are the appropriate description of the data, even an apparently gentle maximum in BIC actually coincides with an accurate, unique determination of m_0.

To gauge the accuracy of the retrieval of the number of states for multi-state quantum dots, we simulated quantum dot data with power-law distributed ($\alpha = 1.5$) switching events, assuming switching between from $\tilde{m}_0 = 2$ to 10 equally likely levels, where we assumed intensity levels to be assigned to segments randomly, and where we assumed levels for simplicity evenly spaced from dark to bright. Lastly, all segments were reassigned an intensity according to Poisson statistics. In other words, we introduced shot noise.

5.3.5. Accuracy of decay rates versus photon budget

For many random realizations with different m_0 and $\langle N \rangle$ we determine the most likely number of states m_r (BIC(m_r) = max(BIC)) according to the clustering algorithm, and construct histograms of outcomes. The total photon budget is set by the product of assumed record length and the mean count rates of the different levels. The outcome of these calculations are shown in Figure 5.7(A-D), where each panel corresponds to a different photon budget. A plot with only diagonal entries signifies that the number of levels retrieved by the clustering algorithm always correspond to the number of levels assumed, so $m_r = m_0$. At high photon budgets (Figures 5.7C, D), the retrieval of the number of states is indeed robust, even for simulated dots that switch between as many as 10 intensity levels. At low total photon budgets (Figures 5.7A, B), we see that m_r is often underestimated. This signifies that there is high uncertainty due to shot noise in the assigned intensity levels, so that levels can not be discriminated within the photon budget. It is remarkable that at photon budgets of 10^6 photons, as many as 10 intensity levels can be robustly discerned even though the smallest contrast

Figure 5.8: (A) Two examples of decay traces ($\gamma = 0.1, 1\,\mathrm{ns}^{-1}$) with low photon count ($N = 300$) fitted to a single-exponential decay. (B) The standard error of the fitted decay rate w.r.t. the input decay rate as a function of the total photon count, for different decay rates. Each data point is the average of 10^3 simulated decay traces.

between levels is as small as 10% in intensity in view of Fig. 5.2, where it is evident that detecting changepoints for small intensity jumps is difficult. Here, however, one should realize that in contrast to Fig. 5.2C, where single small-contrast changepoints are studied, here many levels are visited in random order. Thus the correct detection of small level differences is not reliant on the detection of small changepoint contrasts, but on having sufficient photon statistics to resolve the segment count rates of already identified segments. As a secondary metric, additional to the BIC, is the occupancy of the levels assigned by the clustering algorithm. The clustering algorithm assigns to each data segment the most likely intensity level that it was drawn from. Occupancy is a metric for how often each of the n_G levels available to the algorithm is actually visited by the measured intensity sequence. We find that if one allows the clustering algorithm to use more levels than originally used to synthesize the quantum dot data ($n_G > m_0$), the additional levels take essentially no occupancy. We show this in Figure 5.4B for an example of a dot assumed to have four intensity levels with a total photon count of 5×10^5. As soon as additional states ($m \geq 5$) are offered to the grouping algorithm, these additional states take no occupancy and do not change the distribution of segments over the states found at the correct m. Thus we confirm again that the grouping algorithm does not over-estimate the number of states.

Figure 5.8(A) shows examples of fitted simulated data for slow and fast decays, as examples of the Monte Carlo simulations we have performed to benchmark the accuracy of decay rate fitting in function of photon budget, and decay rate (panel (B)). We find that the error in γ very roughly follows roughly a power law with an exponent of $0.9 - 1.1$, with slower decay rates showing higher errors. Consistent with Ref. [312] we find by Monte Carlo simulation that one requires approx. 200 (50) counts to obtain an error below 10% (20%) in decay rate if one fits mono-exponential decay with free parameters. A problem intrinsic to the use of CPA is that fast switching events may be missed, leading to an averaging of short time intervals with others. This leads to decay traces that are in fact not attributable to a single exponential decay.

5.3.6. FDID diagrams

Correlative diagrams that plot correlations between intensity levels and fluorescence
lifetime [266, 289, 296] form a powerful visualization of quantum dot photophysics.
Our toolbox contains code to generate both fluorescence decay rate intensity diagrams
(FDIDs), and fluorescence lifetime intensity diagrams (FLIDs). The considerations in
this section hold equally for FDID and FLID diagrams, although the provided example
is for a FDID analysis. Essential to the construction of FDID/FLID diagrams is that for
each detected photon also the delay time to the laser pulse that generated it is known so
that decay rates can be fitted even to short segments of a time trace segmented by CPA.
Here we discuss the construction of FDID diagrams derived from CPA-analysis, again
illustrated by examining simulated data for a quantum dot that switches between two
states of distinct intensity and lifetime. FDID diagrams are conventionally constructed
from time-binned data, where it is interpreted as a simple histogram in which each
time-bin contributes a single histogram count to one single intensity-decay rate bin.
It is not trivial to extend this notion to CPA-segmented data since CPA segments
intrinsically have very different time durations instead of having equal width as in
conventional time-binning. We propose two modifications to the construction of a
FDID as a histogram. First, instead of representing FDID entries as a single binary
entry in just one histogram bin (one time segment contributes one count to a single
pixel in a FDID), we propose to incorporate the uncertainty in intensity and decay rate
that is associated with each time segment. To this end, each segment contributes to
the FDID diagram according to a 2D Gaussian function centered at the CPA-segment
decay rate and intensity (total counts C_j in segment j divided by segment duration T_j),
where the width is given by the fit error in the decay rate and the shot noise error in
the segment $\sqrt{C_j}/T_j$. If one would apply this logic to regular time-binned data, giving
each Gaussian contributor the same integrated weight, one obtains a diagram similar
to a regular FDID histogram except that the results is smooth and with less dependency
on a chosen histogram bin width. Instead, the feature size in FDID represents the actual
uncertainties in intensity and rate.

  As a second modification, we propose to reconsider the *weights* of the Gaussians
— i.e., the *integrated* contribution to each entry in the FDID. For time-binned data, one
assigns each segment equal weight so that equal lengths of time contribute equal weight.
Since CPA results in segments of unequal length, several choices for constructing
FDID diagrams are possible, which to our knowledge have not been discussed in
CPA literature. Giving equal weight to each CPA fragment will lead to FDID diagrams
from those obtained from binned data, since effectively long time segments are then
underrepresented compared to short segments. This leads to under-representation
of states with steeper powerlaw distributions in their switching times. The direct
equivalent to regular FDID-weighting for CPA-segmented data is that a segment of
duration T_i has a weight proportional to T_i (henceforth 'duration-weighted FDIDs').
Alternatively one could argue that since *time-averaged* intensity and fluorescence decay
traces are rather set by the contribution in emitted photons, one could instead use the
total number of photons C_i in each segment as weight (henceforth 'count-weighted
FDID').

  It has been established in a multitude of studies of II-VI quantum dots that the
distribution of on/off times follows a power-law (exponents $\alpha_{on,off}$) truncated by an
exponential with specific time τ_l, giving the distribution [286–288] $t^{-\alpha}e^{-t/\tau_l}$. We

Figure 5.9: Four FDID diagrams of a simulated bimodal quantum dot with $(I, \gamma) = (2 \times 10^4 \text{ s}^{-1}, 0.1 \text{ ns}^{-1})$ and $(0.5 \times 10^4 \text{ s}^{-1}, 0.4 \text{ ns}^{-1})$. The color scale indicates the height of the 2D histogram. In each diagram, we apply a different cutoff time T_c to the simulated powerlaw at ∞, 10 s, 1 s, and 0.1 s. We find that a shorter cutoff causes an increasingly strong smearing effect between the two states.

analyze Monte Carlo simulated FDIDs to establish if there are conditions of the truncation time under which a two-state quantum dot would appear not as a bimodal distribution. Figure 5.9 shows duration-weighted FDIDs for simulated quantum dot data. For panel (A) we consider a two-state quantum dot with power-law distributed switching times. In panels (B-D) we show a quantum dot simulated with similar parameters, but with the maximum duration of the segments $T_c = 10$, 1, 0.1 seconds. Evidently, the bi-modal nature of the quantum dot is faithfully represented by the FDID diagram constructed through CPA. This remains true also for power laws with a long time truncation ($t^{-\alpha} e^{-t/\tau_l}$), unless truncation time τ_l are as short as 20 ms, so that there are no segments with over approximately 10^2 counts. This limit of our benchmarking space zooms in on the regime where CPA intrinsically fails (Fig. 5.2(C)). In this regime, the originally assumed bimodal quantum dot behavior no longer results in a bi-modal FDID. Instead a significant broadening is evident. We can conclude that for most realistic quantum dot systems, CPA-generated FDIDs will not suffer from this artificial broadening artifact.

5.4. Conclusion

In this work, we have provided a Python toolbox for change-point analysis, and for determining the most likely intensity level assignment for intermittent multilevel emitters. We investigated the limits of changepoint analysis and clustering as fundamentally set by photon budget, and for the case of many state emitters. We have shown that for long switching times, the typical power-law behavior of many quantum emitters can be accurately retrieved. We also show that in the case of many-state emitters, the number of intensity levels can be retrieved with high fidelity, provided the photon count is high enough. At low photon counts, the number of states is systematically underestimated. This shows in which way the photon budget puts a fundamental limit on the amount of information that can be retrieved from a given TCSPC data set. We show that the photon budget also poses a limitation on the accuracy at which the slope of a

single-exponential decay can be retrieved. Additionally, we investigate the commonly
used intensity-decay time diagrams. We show that with CPA, a two-state simulated
quantum dot is well-represented in an FDID, but when a cutoff is introduced to match
that commonly found in literature, the states become increasingly poorly defined in
an FDID representation. While the Bayesian inference algorithms in this toolbox were
reported earlier for application to quantum dots with just two or three intensity states,
this toolbox and the provided benchmarks point at the applicability even to emitters
that jump between many closely spaced intensity levels, which will, in our view, be
of large practical use for many workers analyzing the complex photophysics of, e.g.,
perovskite quantum dots. Also, the toolbox can be used for theory development,
following a workflow in which hypotheses are cast in synthetic photon counting data,
which in turn can be subjected to the CPA analysis suite, to assess how hypothesized
mechanisms express in observables, and in how far they are testable. The limit of this
testability generally depends on an interplay of total photon budgets, residence time
in each level, intensity contrast between levels, and segment durations. For a given
photophysics scenario, the user can easily deploy the toolbox to directly assess data
segmentation in terms of accuracy, precision, and recall error rates, in dependence
of the level of skepticism that the user wishes to apply in order to accept assertions
regarding the segmentation of data in segments and intensity levels. These error rates,
and hence the testability of a hypothesized photophysics scenario, are ultimately limited
by the evidence in the counting statistics, and not the segmentation algorithm.

The full code base used for this paper has been made freely available on
GitHub [305].

APPENDICES

5.A. Processing toolbox - overview

I$_{\text{N}}$ this supplement, we will briefly discuss the different uses of the Python toolbox for unbiased statistical analysis. For a full discussion, we refer to the supplementary material of [306]. The toolkit accepts both data taken in a TCSPC measurement and data created in Monte-Carlo (MC) simulations. We will first discuss the data processing of photon time stamp files, and subsequently the provided routine to simulate m_0-level single photon emitters. The full toolkit is available for use on GitHub [305].

Figure 5.A.1 shows the different parts of the data analysis tools provided in the toolbox, and describes a workflow to perform a full analysis. The diagram shows main functions as dark blue ellipses, input and output files as green boxes, with blue frames indicating groupings or related tasks the code performs.
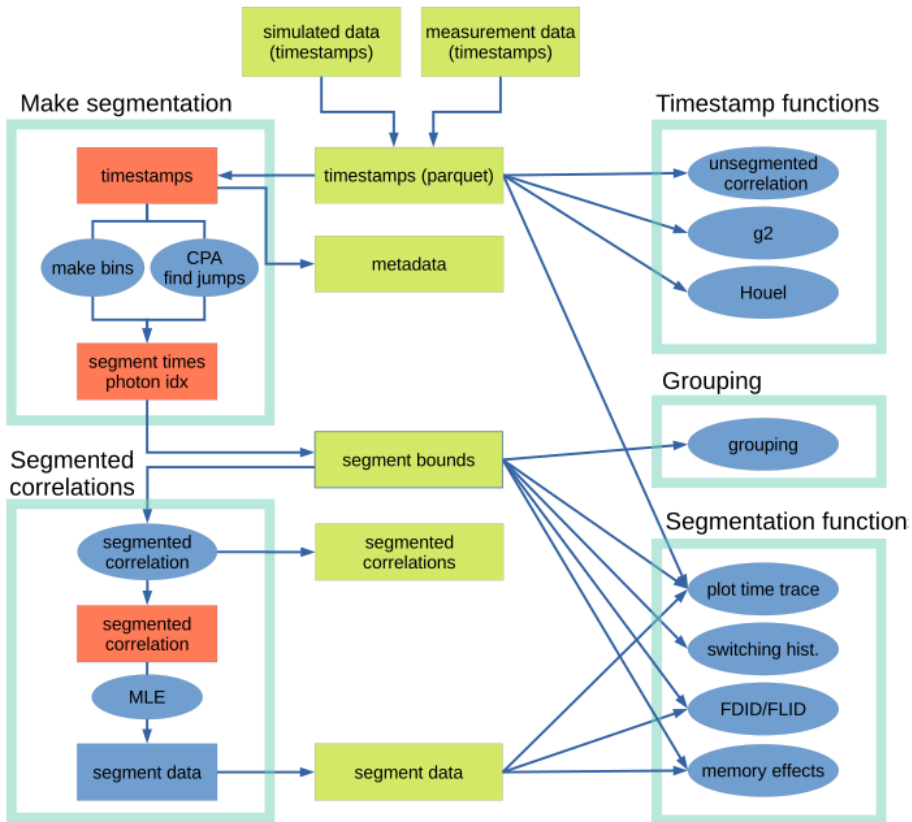
Figure 5.A.1: Flow chart of the processing toolbox.

107

The starting point for the code is a provided set of three lists of timestamps, which represent three channels, i.e., two photon counting channels, and a pulsed laser reference record. These timestamps are provided as sorted lists of 64-bit integers through binary "parquet" files. The data is expected to be in "reverse start-stop" format as usual for TCSPC experiment, meaning that the pulsed laser reference trace contains time stamps for the laser pulses immediately after each photon event.

The processing code has the following functionality. First, the data, either measured or simulated, is segmented into intervals using either changepoint analysis, or with fixed-width binning, shown in the "Make segmentation" block in Fig. 5.A.1. This returns the jump times at which the emitter at hand jumps in intensity, and between which the dot is constant in intensity, within the CPA threshold. It also yields the index of the corresponding photon event, from which intensities can be extracted. With a fixed bin width, it will return the times and photon indices at the bins' edges.

Next, for each segment, the detector and photon channels are time correlated (block "segmented correlations") to obtain fluorescence decay trace photon counting histograms per segment, to which a decay rate is fitted with the maximum likelihood estimator (MLE). These processes lead to a set of intermediate, ASCII-formatted files that contain the segmented jump times, intensities and decay rates.

The reminder of the routines are found on the right hand side of the diagram. The top block "timestamp functions" acts not on the segments, but on the timestamp list as a whole. They convert the timestamp list into time-correlations between detectors, and detector and laser, across the whole data set, to provide a decay trace, $g^{(2)}(\tau)$, the full decay trace, and the long-term autocorrelation trace, discussed in Section 5.2.3 and Section 6.4.3.

The remainder of the blocks post-processes the segmented data. This includes on one hand *grouping* according to Bayesian analysis, and on the other hand plotting of the segmented data as time traces, FDID diagrams, and switching time histograms.

5.B. Simulating N-level emitters

We provide a routine to generate time-tagged data of an m_0-level single emitter for Monte Carlo simulations testing the toolbox, or for testing how assumed photophysics would appear in experiments. This routine is independent from the processing toolbox. It produces timestamps in parquet file format that, together with an appropriately specified preamble files, can be processed in exactly the same manner as the example experiment data set.

The steps followed by the routine are shown schematically in Fig. 5.B.1. The input parameters are shown in orange. The first is the number of desired levels, and the others are their exponents for powerlaw switching, their intensities, and their decay rates. Lastly, it takes some level of background noise. We will walk through the routine in order of its operations. First, a random selection is done, to create a list of levels. This will provide the which-level information, used in the subsequent steps. The second step is determining the length of each segment, bounded by switching events. This is done by sampling a powerlaw with the chosen exponents. When the length of each segment is known, they can be filled with photon events. This is done by drawing from a uniform distribution, to create a constant intensity for the duration of each segment. This step only defines for which excitation events a photon event is registered. So for the next step, the routine determines the delay w.r.t. the excitation event. This is done by drawing from an exponential decay with the appropriate decay rate for each segment.
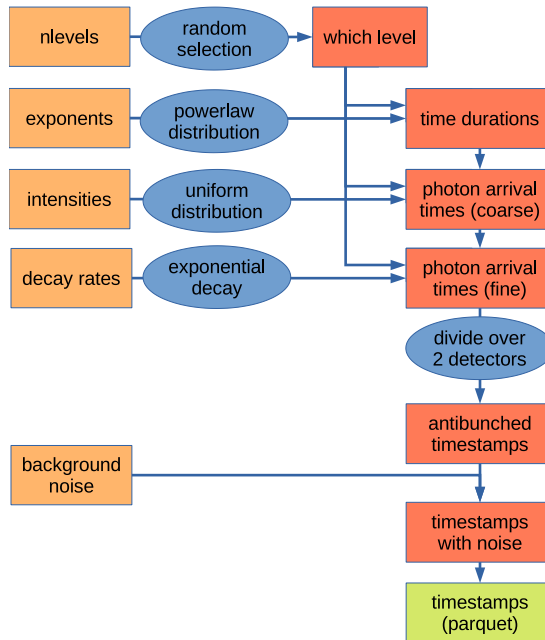
Figure 5.B.1: Flow chart of the routine to simulate multi-level emitters.

At this point, all photon events which are not background have been defined. In order to simulate antibunching, the events are divided randomly into two channels, which are equivalent to two detectors. Now the background noise can be added, the timestamps of the simulated photon events are written to the binary parquet files, and can be analyzed using the main routine. The philosophy of the code is that it can interchangeably run on experimental and simulated data simply by setting a Boolean switch in the code.

6

Intermittency of CsPbBr$_3$ perovskite quantum dots analyzed by an unbiased statistical analysis

"Everything's got a moral, if only you can find it."
— Lewis Carroll, *Alice's Adventures in Wonderland*, 1865

6.1. Introduction

Cᴇsɪᴜᴍ lead halide perovskite nanocrystals, introduced in a seminal paper by Prote-
sescu et al. [156] have emerged as highly attractive quantum dots, with advanta-
geous properties in comparison to traditional colloidal II-VI semiconductor quantum
dots. These include very large photon absorption cross sections [159], a wide degree of
tunability by both size and halide (Br, I, Cl) composition [156], and reportedly a very
high luminescence quantum yield without the need of protecting the nanodot core
with epitaxial shells, as is required for CdSe quantum dots [157, 158]. Furthermore,
inorganic halide perovskite materials generally show an exceptionally high tolerance
to defects [313]. Owing to these properties perovskite nanocrystals are intensively
pursued as solar cell materials [314], as emitters for LEDs, display technologies and
lasers [156, 315, 316], and could be interesting as single photon sources. For the
purpose of single photon sources, emitters need to satisfy a variety of requirements
beyond brightness, tunability and high quantum efficiency, which includes single-
photon purity, tight constraints on inhomogeneous spectral broadening, and stability
in spectrum, decay rate and intensity [27].

Perovskite nanocrystals unfortunately follow the almost universally valid rule that
solid-state single emitters at room temperature show intermittency [157, 280–285]. In
the field of II-VI quantum dots, intermittency has been studied for over two decades,
with the aim of identifying the nature of the usually two or three distinct bright, dark,
and gray states, and the mechanism by which switching occurs, by analysis of the
apparently discrete switching events between dark and bright states [288, 317], and
concomitant jumps in spectrum and lifetime. For instance, for II-VI quantum dots a
popular model (reviewed in Ref. [318]) is the charging/discharging model whereby
quantum dots turn from bright to dark upon acquiring a single charge.

Many efforts have been made to explain the typically power-law distributed res-
idence times for on and off states, for instance through hypothesized mechanisms
by which charges are exchanged with the environment [286, 287, 317, 319]. In this
respect another powerful model is the so-called Multiple Recombination Center (MRC)
model proposed by Frantsuzov et al. in 2009 [320, 321], which argues that the wide
distribution of on/off times underlying binary blinking is due to typically of order 10
available recombination centers. This model furthermore is applicable to a wide array
of systems such as quantum dots, rods and wires, as it can explain also qualitatively
different intermittency behavior, such as systems that do not show two but multiple
intensity levels, in function of assumed underlying recombination center physics [322].

For perovskite nanocrystals several groups studied intermittency [157, 280–285] and
found quite different physics. A set of works observe that perovskite quantum dots
do not show bimodal behavior, like II-VI quantum dots do, but instead a continuous
distribution of states between which they switch [280, 281, 283]. These observations are
difficult to rationalize in a charging-discharging model, but can be described within
the MRC model of Frantsuzov et al. [320, 322], as pointed out for CsPbBr₃ dots by Li et
al. [281]. Within this model activation of individual recombination centers can provide
a wide distribution over intensity and rate. An important observation consistent with
the MRC model is a linear dependence between emitted intensity and fluorescence
lifetime. Further evidence for multiple recombination center physics in the context of
perovskite PL has been reported in the context of emitting perovskite microparticles
that show no quantum confinement but nonetheless blink, for instance in a recent report

by Merdasa et al. that evidences extremely efficient dynamic quenching sites that can appear and disappear [323]. In contrast to Refs. [280, 281, 283], another group has analyzed intermittency on basis of changepoint analysis and cluster analysis, which are Bayesian inference tools for the unbiased estimate of the number of states, reporting that just of order 2-4 states are involved instead of a continuum [282]. Finally, a recent study points at memory effects in intermittency, visible in that work as correlations between subsequent dwell times in the brightest state [285]. These reported memory effects for perovskite dots are similar to those observed over 15 years ago for II-VI dots by Stefani and coworkers [324], which were explained by the MRC model [321].

Intermittency analysis is a field known to be fraught by statistical bias in analysis methods [319], primarily due to binning of data prior to analysis. This is a recognized problem already for interpreting data from bimodal dots. These artifacts may be even more severe for multilevel dots. In this work we report a study of cesium-lead-bromide nanocrystal intermittency, analyzing the photon statistics of a large number of dots using unbiased Bayesian statistics analysis tools, tracing brightness and fluorescence lifetimes simultaneously, and screening for memory effects. These Bayesian statistics methods were first developed by Watkins and Yang [290, 292, 293] and have since been applied in a small set of papers to two/few-level II-VI dots, and in one recent work to $CsPbBr_3$ dots [282]. Our implementation is through a freely available Python-based analysis toolbox [306], which we have specifically benchmarked by Monte Carlo methods for application to highly multi-state, instead of bi-modal, systems. In this work, a first main purpose is to obtain statistically unbiased estimates, or at least lower bounds, for the number of dark/gray states of perovskite quantum dots from a large number of single dot measurements. Our conclusions solidly support [280, 281, 283], but not Ref. [282], since we find blinking between a single well-defined bright state and a continuum — or at least over 10 — gray/darker states. These are findings that fall within the class of phenomena explainable by the MRC model [281, 320, 322]. Next, our purpose is to screen for memory effects in residence times, intensity levels and decay rate sequences in data that has been separated into segments by unbiased changepoint analysis, thereby extending Ref. [285], which did not leverage the benefit of CPA analysis. We find no evidence for memory in residence times, but do find that a substantial fraction of dots tend to switch back and forth repeatedly between the quite uniquely defined bright state and the band of gray states, instead of jumping through all states in an uncorrelated random fashion.

6.2. Experimental methods

To introduce our measurement protocol and the photophysics of the $CsPbBr_3$ dots at hand we first present in Figure 6.1 the typical behavior of a $CsPbBr_3$ quantum dot, as analyzed with the standard approach of plotting time binned data. We prepare quantum dots according to a modified literature report [156, 325].

Preparation of cesium oleate. We load 0.814 g of Cs_2CO_3 into a 100 mL 3-neck flask along with 40 ml of octadecene (ODE) and 2.5 ml of oleic acid (OA) and dry this for 1 hour at 120 °C. This is then heated under an N_2 atmosphere to 150 °C until all Cs_2CO_3 has reacted with OA. To prepare for the next step, we preheat the resulting cesium oleate to 100 ∘C before injection. This is necessary as it precipitates out from ODE at room temperature.
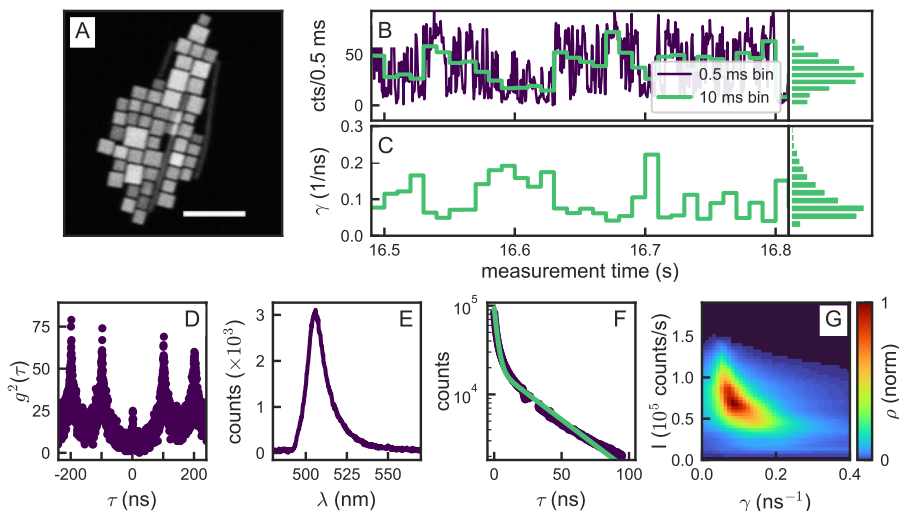
Figure 6.1: Properties of a CsPbBr$_3$ quantum dot. (A) a SEM image showing a cluster of CsPbBr$_3$ quantum dots. The scale bar is 100 nm. A time trace of the (B) intensity and (C) fluorescence decay rate of a typical quantum dot when split into bins of 10 ms (green). We find a single peak in both the intensities and lifetimes around 60 counts/ms and 0.05 ns^{-1}, respectively. For visualization purposes, we also show the photon events binned into 0.5 ms bins (purple). (D) The $g^2(\tau)$ of this qdot. The dots used in this analysis were selected for having $g^2(0) < 0.5 \cdot g^2(100\text{ns})$. (E) The PL spectrum of this qdot when exciting at 450 nm. We find a peak in the emission at 505 nm. (F) The decay trace of all the photon events combined. We have excluded an electronic artifact between 20 and 30 ns. We have a reasonable fit to a bi-exponential decay with rates of $\gamma_1, \gamma_2 = 0.43, 0.03$ ns^{-1}, respectively. (G) The FDID diagram of this dot. We see a main peak at $I, \gamma = 0.7 \times 10^5$ cts/s and 0.09 ns^{-1}

Synthesis of CsPbBr$_3$ nanocubes. We load 0.188 mmol of PbBr$_2$ in 5 ml of ODE, 0.5 ml of oleylamine and 0.5 ml of OA into a three-neck round bottom flask and dry this under vacuum at 120 °C for an hour, after which the reaction atmosphere is made inert by flushing the flask with N$_2$. After complete solubilization of PbBr$_2$, the temperature is raised to 200 °C and 0.4 ml of the preheated cesium oleate is injected into the three-neck flask. After the injection, the color of the solution turns from colorless to greenish yellow indicating the formation of perovskite cubes. Then we lower the temperature to 160 °C and anneal the solution at that temperature for 10 min to get uniform size dispersion of the cubes. After that, we cool down the solution using an ice water bath for further use.

Isolation and purification of CsPbBr$_3$ cubes. After the synthesis, we centrifuge our solution twice to collect the cubes. First, we take 1 ml from the stock solution just after the synthesis and centrifuge at 8000 rpm for 20 min to collect all CsPbBr$_3$ particles from the solution. We discard the supernatant, gently wash the inner wall of the tube

using tissue paper and add 2 ml of toluene to disperse the $CsPbBr_3$ solid. The second step of centrifugation is run at 2000 rpm for 5 min to get rid of all the particles that are too large. In the supernatant, we have 2 ml of toluene containing $CsPbBr_3$ nanocubes having a size distribution around 10-15 nm. As the scanning electron micrograph in Figure 6.1(A) shows our quantum dots are essentially cubic in shape.

Before the measurement, about 400 μL of the solution is spin-coated at 1000 rpm on glass coverslips that had been cleaned in a base piranha solution. In order to protect the quantum dots from moisture in the air, the quantum dots were covered by a layer of PMMA (8% solid weight in anisole), by spin-coating for 60 s at 4000 rpm. Quantum dots stored in solution were found to be unchanged in their properties over $\gtrsim$ 6 months. For the optical experiments we prepared microscope slides with samples from solutions no more than 1 month old, and then performed microscopy on a given substrate within a time span of 7 days. We found no differences between data taken directly, and data taken after 7 days.

Single emitter microscope. For optical characterization and measurements, we use the inverted optical microscope shown in Fig. 1.5 to confocally pump the dots at 450 nm (LDH-P-C-450B pulsed laser, PicoQuant) at 10 MHz repetition rate of $<$ 70 ps pulses, with 90 nW inserted into the microscope. An oil objective (Nikon Plan APO VC, NA=1.4) focuses the pump laser onto the sample and collects the fluorescence. The excitation provides similar pulse energy density as in [282]at the lowest energy density $\langle N \rangle \ll 1$ quoted in that work. With the estimated efficiency of our setup, the excitation probability per optical pulse is estimated at $<$0.1 from the count rate. The fluorescence from the sample is directed either to a camera (PCO.edge 4.2, PCO AG), a spectrometer (PI Acton SP2300) or two fiber-coupled avalanche photodiodes (APDs) (SPCM-AQRH-14, Excelitas) in a Hanbury-Brown & Twiss configuration. The APDs are coupled to a photon correlator (Becker & Hickl DPC-230) that measures the absolute photon arrival times.

Measurement protocol. Using the camera and wide-field pump illumination, we select an emitter that appears to be diffraction-limited. After driving it to the laser spot, we do a time-correlated single-photon counting (TCSPC) measurement to collect photon arrival times. To calculate the photon correlations we use a home-built TCSPC toolkit that utilizes the algorithm developed by Wahl et al. [310] to calculate $g^2(\tau)$ and the lifetimes for the different emitters and for the individual CPA segments. From $g^2(\tau)$ we select the emitters with a strong anti-bunching signal (normalized $g^2(\tau = 0) < 0.5$) to ensure single quantum emitter behavior. Of the 75 dots measured, 40 passed this test. We note that within those 40 dots we found no systematic correlations between any of the variables (brightness, decay rates, apparent number of levels, residence time power law exponent) and the normalized $g^2(\tau = 0)$ value. Our TCSPC measurements are taken over 120 seconds of acquisition time. We note that in our decay traces taken using a Becker-Hickl DPC 230 photon-counting and correlator card in reverse start-stop configuration, a small time interval centered at around 30 ns is subject to an electronic artifact which we attribute to a ringing in the DPC-230 TDS timing chip. Therefore we exclude this time interval for decay rate fitting.

Initial characterization. Figure 6.1 presents initial characterization of an exemplary single dot on basis of standard time-binned analysis, where the data is sliced in 10 ms long segments, to each of which intensity and decay rate is fitted. Throughout this work we consider photon counting data, in which absolute time-stamps are collected with 0.165 ns resolution for all collected photons and concomitant excitation laser pulses,

on two avalanche photodiodes (APDs) in a Hanbury-Brown and Twiss configuration. This allows to construct *a posteriori* from one single data set the intensity, fluorescence decay rate, and the $g^{(2)}(\tau)$ photon-photon correlation. In our optical measurements we post select all single nanoparticles on basis of photon antibunching ($g^{(2)}(\tau = 0) < 0.5$). For the example at hand, the selected emitter shows clear intermittency in intensity and decay rate (panels Fig. 6.1(B,C) discussed further below), while Fig. 6.1(D) shows a marked anti-bunching at zero time delay in the $g^{(2)}(\tau)$ that is constructed from the full photon record. The quantum dot in Fig. 6.1(E) shows a time-averaged emission spectrum that peaks at around 505 nm and has a spectral FWHM of 20 nm, which is consistent with reports by Protesescu et al. [156], and together with the antibunching photon statistics points at quantum confinement. The time-integrated fluorescence decay trace (Fig. 6.1F) is markedly non-single exponential. Fitted to a double exponential decay we find decay rates of $\gamma_1, \gamma_2 = 0.43, 0.03$ ns^{-1}. We must note, however, that a double exponential is often not sufficient to fit these emitters, and typical decay rates for our dots range from 0.05 to 0.9 ns^{-1}. At these decay rates, the fastest decay rate component of the quantum dots generally spans at least 10 timing card bin widths.

Figure 6.1(B, C) shows just a fraction of the intensity and decay rate time trace, plotted according to the common practice of partitioning the single photon data stream in bins. The fluorescence decay rate for each bin is obtained by fitting data within each 10 ms bin to a single exponential decay law employing a maximum likelihood estimator method that is appropriate for Poissonian statistics [311]. As expected from prior reports on single perovskite nanocrystal blinking [157, 280–285], the intensity and decay rate time trace show clear evidence for intermittency. The intensity varies from essentially zero to 150 counts per ms. Figure 6.1(B, right panel) shows a histogram of intensities, binned over the entire time trace (for all dots in this work, 120 s, or till bleaching occurred). The histogram shows a broad distribution of intensities with the most frequent intensities around 60 cts/ms. This is in contrast with the typical bi- or trimodal physics of II-VI quantum dots, which usually show distinct bright and dark states [149, 286–288, 317, 318]. However, the width of the peak well exceeds the Poisson variance expected at these count rates, suggesting that there are many intensity levels. The decay rate histogram also displays intermittent behavior, in step with the intensity blinking. The most frequent decay rate is around 0.07 ns^{-1}. Fig. 6.1G displays a *Fluorescence Decay Rate Intensity Diagram* [FDID], a 2D histogram displaying the frequency of occurrence of intensity-decay rate combinations. This type of visualization was first introduced by [266, 289, 296] to identify *correlations* between intensity and fluorescence decay rate (FLIDs in those works, using lifetime instead of decay rate). For II-VI quantum dots, FDID diagrams typically separate out bright and slowly decaying states from dark, quickly decaying states [266, 289]. Instead, for the perovskite quantum dot at hand, the FDID diagram presents a broad distribution with a long tail towards dim states with a fast decay.

The picture that emerges from Fig. 6.1 is consistent with recent observations of several groups [280, 281, 283], showing a continuous distribution of dark, gray states. This should be contrasted to typical II-VI quantum dot behavior in which blinking usually involves just two or three apparent intensity levels, and also the recent report by [282] on very similar CsPbBr$_3$ dots, but taken under very low repetition rate excitation conditions (fs pulses at very low repetition rate, as opposed to picosecond pulses at ≥ 10 MHz — at similar $\langle N \rangle < 0.1$).

6.3. Computational methods

Since extreme caution is warranted when scrutinizing photon counting statistics to determine quantitative intermittency metrics due to artifacts of binning [290, 292–294, 296], we proceed to analyze the data of a large number of dots with state-of-the-art bias-free statistical analysis to determine a lower bound to the number of involved states, and the switching dynamics and memory effects therein. We apply tools of Bayesian statistics, specifically, changepoint analysis (CPA), discussed in Chapter 5, to partition the data in segments separated by switching events, and level-clustering to determine (a lower bound to) the number of states, as a rigorous and bias-free approach to investigate the intermittency of quantum dots. These tools were first proposed by Watkins and Yang [290], and later also used and extended in the context of quantum dot intermittency by [150, 152, 153, 282, 292, 293, 296, 307–309]. We have discussed changepoint analysis and clustering extensively in Chapter 5, but will summarize the salient features here.

CPA performs segmentation of the time record of single photon counting events into intervals within which the count rate is most likely a constant value, delineated by switching events or 'changepoints' at which the count rate changes, in as far as can be judged given the shot noise in the data. Since CPA works on a full time series with many jumps by finding a single jump at a time, and successively subdividing the time stream until segments with no further jumps are found, the ultimate performance is ultimately set by how well CPA can pinpoint in the last stage of the subdivision single jumps in short fragments of the photon stream. For significant intensity contrasts CPA detects changepoint in very short fragments (e.g., to accurately resolve a jump with a 5-fold count rate contrast, a record of just 200 photons suffice), with single-photon event accuracy. Smaller jumps are missed unless fragments are longer (e.g., factors 1.5 contrast jumps require fragments of ca. 10^3 photons for near sure ($> 90\%$) detection. At typical *practical* count rates of 10^5 counts/s, this means that switching events further apart than 10 ms are accurately identified as long as jump contrasts exceed a factor 1.5 (ca. 100 ms for contrasts as small as 1.2). Switching events that are even closer in time are missed by CPA. This is intrinsic to the photon budget, i.e., the ultimate information content in the discrete event time stream fundamentally does not allow pinpointing even more closely spaced switching evens.

After dividing the time trace into segments spaced by changepoints, one is left with sequences containing the residence times T_q for each segment, photon counts N_q and instantaneous segment intensities ($I_q = N_q/T_q$), as well as decay rates γ_q, obtained by maximum likelihood fitting of the decay trace from each segment to a single exponential decay. The question how many actual intensity levels most likely underlie the measured noisy sequence $\mathcal{I}_{m_r}$ can be determined using Watkins & Yang's clustering algorithm [290]. While Watkins and Yang considered Poisson distributed noise, as in this work, we recommend also the work of Li and Yang [303] as a very clear explanation of the method, though applied to Gaussian distributed noise. The idea is that expectation-maximization is used to group the most similar segments together into n_G intensity levels, where $n_G = 1, 2, 3, \ldots$. After this, the most likely number of levels, $n_G = m_r$, required to describe the data, given that photon counts are Poisson distributed, can be determined by a so-called Bayesian Information Criterion (BIC) [290, 303]. We have extensively verified by Monte Carlo simulations the performance of CPA and level clustering for dots with *many* assumed discrete intensity levels in Chapter 5 In

brief, at small photon budgets in a total time series, only a few levels can be detected, but conversely at the total photon budgets in this work, exceeding $5 \cdot 10^6$ events, clustering has a $> 95\%$ success rate in pinpointing the exact number of levels in dots with at least 10 assumed intensity levels. Moreover, for photon budgets that are too small to pinpoint all levels exactly (e.g., at 10^4 counts in a total measurement record, only up to 4 levels can be accurately discerned), clustering always returns a *lower bound* for the actual number of intensity states.

6.4. Results and discussion

6.4.1. Changepoint analysis and FDID diagrams

We have applied the unbiased CPA analysis and Bayesian inference tools to data from 40 single CsPbBr₃ quantum dots. We first discuss an exemplary single dot as example, and then discuss statistics over many single dots. The example dot is identical to the one considered in Figure 6.1 and refer to the supporting information for results on all dots. In Figure 6.2A we see that CPA is able to accurately follow the intensity trace of a typical CsPbBr₃ quantum dot. We show only a section of the total measurement for clarity, and strictly for plotting purposes only, binned the photon arrival times in 0.5 ms intervals. Note that this binning is only for visualization, and does not enter the CPA algorithm. Figure 6.2B displays the fitted decay rates for the same selected time interval, obtained by fitting each of the identified segments. The right-hand panels of Figures 6.2A and B show histograms of intensity and lifetime as accumulated over the full time trace. It should be noted that these histograms are intrinsically different from those in Figure 6.1 for two reasons. First, binned data has entries from bins containing jumps, leading to a smearing of the histogram. Second, since histogramming of segment values I_q is agnostic to segment *duration*, events are differently weighted. Thus the histogram of intensities now shows a bimodal distribution. The histogram of the decay rates still exhibits only a single peak at ca. $0.05 \, \text{ns}^{-1}$.

Next, we construct correlation diagrams of fluorescent decay rate versus intensity (FDIDs) from CPA data. Customarily FDIDs are 2D histograms of intensity and decay rate as extracted from equally long time bins in binned data. As the length of segments found by CPA can vary over many orders of magnitude, an important question is with what weight a given segment should contribute to a CPA-derived FDID. A first approach is to give all segments an equal contribution to the FDID, which emphasizes the probability for a dot to jump to a given intensity-decay rate combination. Alternatively, one could weight the contribution of each segment to the histogram by the number of counts it contributes. This histogram hence emphasizes those entries that contribute the most to the time-integrated observed photon flux. Lastly, if one uses the segment durations as weights for the contribution of segments to the FDID one obtains an FDID closest in interpretation to the conventional FDID diagram, which presents the probability density for being in a certain state at a given time. Figures 6.2C, D and E provide all three visualizations. The data shows variations in intensity levels over approximately a factor 10, with concomitant decay rates also varying over an order of magnitude. Overall, all diagrams suggest an inverse dependence qualitatively consistent with the notion that the dots experience a fixed radiative rate, yet a dynamic variation in the number of available non-radiative decay channels, that make the dot both darker and faster emitting. This inverse dependence was also observed
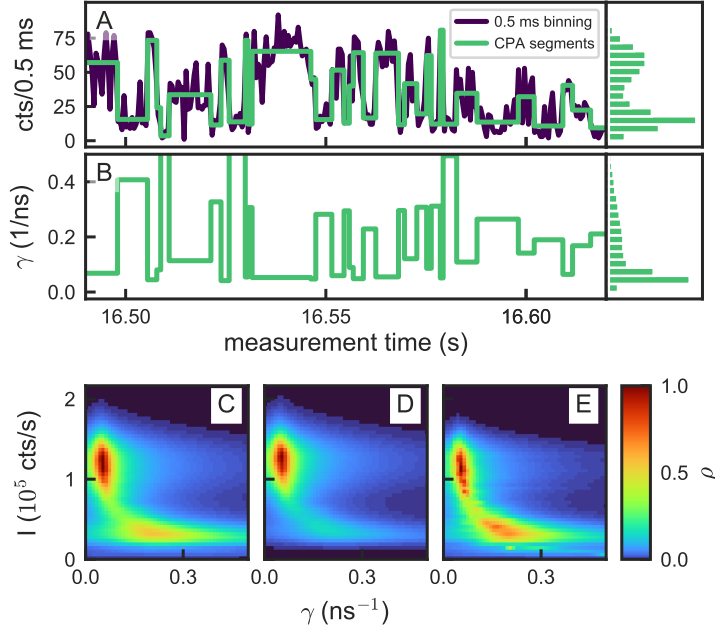
Figure 6.2: (A) An example of the intensity time trace of a measured quantum dot (purple, binned in 0.5 ms bins for visualization purposes), and the intensity segments found by CPA (green). In the lower panel (B) the lifetime for the found CPA segments is shown. On the right are histograms of the occurrences of the intensities for both treatments with segments weighted by their duration. (C-E) Three FDID plots weighting each CPA segment (C) equally, (D) by their number of counts, (E) by their duration. The choice of weights puts emphasis on different parts of the intensity-decay rate diagram, as they report on differently defined probability density functions.

for perovskite dots by [281], and can be explained by the MRC model [320, 322]. The unweighted and photon count weighted FDIDs show a peak at similar intensity and decay rate at $(\gamma, I) = (0.06 \text{ ns}^{-1}, 12 \times 10^4 \text{ s}^{-1})$, indicative of the most frequently occurring intensity/rate combination that is simultaneously the apparent bright state. The different FDID weightings emphasize different aspects of the data. For instance, weighting by counts highlights mainly the emissive states and underrepresents the long tail of darker state, with respect to the other weighting approaches. This qualitative difference can result in a quantitative difference for extracted parameters, such as the apparently most frequently occurring combination of intensity and decay rate.

For essentially all dots (see supporting information), FDIDs are much like the example shown in Figure 6.2, showing a slow decaying bright state with a long tail towards both lower intensity and faster decay. In fact, we can collapse the FDIDs of all 40 dots onto each other by summing histograms (no weighting by, e.g., segment

119

Figure 6.3: FDID of all 40 single dots, obtained by summing single dot FDIDs for which the segment intensities were normalized to the mean intensity. A simple histogramming was used (no specific weighting of entries by duration or counts). Overplotted is a parametric curve of the form $(\gamma_r + \gamma_{nr}, B + I_0\gamma_r/(\gamma_r + \gamma_{nr}))$ with as input a fixed value γ_r, and a background $B = 0.06I_0$, with I_0 adjusted to match the peak in the FDID, and γ_{nr} scanned.

duration) of *normalized* intensity $I/\langle I\rangle$ versus γ, which further underlines this generic behavior, see Figure 6.3). An appealing explanation for the observed dynamics is if the perovskite dots are characterized by always emitting from one unique bright state that is efficient and has a slow rate of decay γ_r [labeled as radiative decay rate], while suffering fluctuations in both brightness and rate through jumps in a nonradiative rate γ_{nr}, as in the MRC model [281, 320, 322]. In this picture, one would expect the FDID feature to be parametrizable as $I \propto B + I_0\gamma_r/(\gamma_r + \gamma_{nr})$. The feature in the collapsed FDID plot can indeed be reasonably parametrized as such a hyperbola. This parametrization is consistent with Ref. [281] in which a linear relation between intensity and fluorescence lifetime was reported. The required radiative decay rate for the bright state is $\gamma_r \sim 0.075$ ns^{-1}, while the parametrization requires a residual background $B = 0.06I_0$. This residual background is not attributable to set up background or substrate fluorescence, suggesting a weak, slow luminescence component from the dots themselves. Moreover, we note that the FDID feature clearly has a somewhat stronger curvature than the hyperbolic parametrization (steepness of feature at $\gamma < 0.1$ ns^{-1}, and $I/\langle I\rangle > 1.0$).

6.4.2. Clustering analysis

The FDID diagrams at hand qualitatively support the continuous distribution of states also observed by Refs. [280, 281, 283]. As quantification of the number of states involved we perform clustering analysis [290, 292], discussed in Chapter 5 to estimate the most likely number of intensity states describing the data on basis of Bayesian inference. A plot of the Bayesian Information Criterion as a function of the number of levels n_G allowed for describing the data of the specific example dot at hand is shown in Figure 6.4A. Strikingly, the BIC does not exhibit any maximum in the range $n_G = 1 \ldots 5$, but at $n_G = 13$. Recalling that the BIC criterion in clustering analysis for multistate dots at finite budget generally reports a lower bound, this finding indicates that the data for this dot requires at least as many levels to be accurately described, if a discrete level model is at all appropriate.

Similar conclusions can be drawn from Figure 6.4B. We have found in Monte Carlo simulations that if one allows the level clustering algorithm to find the best description of intensity traces in n_G levels for dots that in fact have just $m < n_G$ levels, then the returned description of the data utilizes just m levels, with the remaining levels having zero occupancy in the best description of the data returned by the algorithm. Figure 6.4B shows the occupancy assigned by the clustering algorithm for our measured quantum dot as a function of the number of states offered to the algorithm for describing the segmented intensity trace. Each additional state offered to the clustering algorithm is in fact used by the algorithm, whereas Monte Carlo simulations, discussed in Chapter 5, have shown that at the photon budgets involved (5.5×10^6 photons) the clustering algorithm generally does not assign occupancy to more than m levels to simulated m-levels dots. The occupancy diagram hence confirms the conclusion from the BIC criterion that the dot at hand requires many levels, or even a continuous set of levels, to be described.

For all 40 dots, we extracted wavelength, brightness, and performed the same CPA and clustering analysis as for the example dot. Moreover, we examined segment duration statistics for power law exponents. The supporting information contains a detailed graphical overview of the CPA results for each of the 40 dots, while summarized results are shown in Figure 6.5. Figure 6.5A shows that the dots have a low dispersion in peak emission wavelength, with emission between 500 and 510 nm. All considered dots offered between 2 and 8×10^6 photon events (Figure 6.5B) for analysis (120 seconds collection time, or until photobleaching). The mean intensity per measured dot (histogram Figure 6.5C) is typically in the range from 15×10^3 and 80×10^3 cts/s, with one single dot as bright as 110×10^3 counts/s. According to the Monte Carlo analysis in, as discussed in Chapter 5 the total collected photon count for all dots therefore provides a sufficient photon budget to differentiate with high certainty at least up to 10 states. We can thus with confidence exclude that intermittency in these perovskite quantum dots involves switching between just two or three states as in usual quantum dots. Instead, any physical picture that invokes a set of m discrete levels requires a description in upwards of $m = 10$ levels. In how far further distinctions between > 10 discrete levels, or instead a continuous band can be made on basis of data is fundamentally limited by the finite photon budget that can be extracted from a single emitter. This quantification matches the observation in Ref. [280, 281, 283, 285] (based on examining time-binned FDID diagrams). The main other work that applied CPA tools to perovskite dot by Gibson et al. [282], however, arrived at an estimate $m_r = 2.6$, which is at variance with
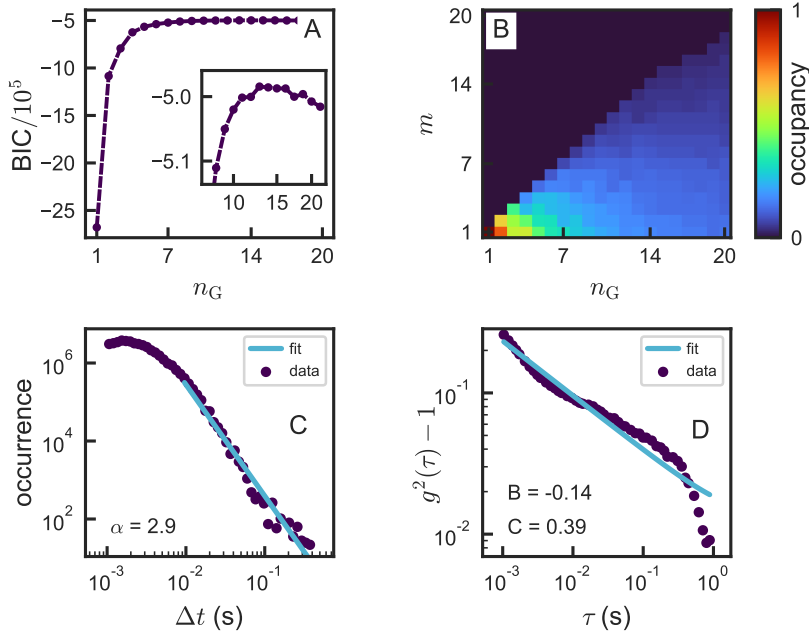
Figure 6.4: (A) BIC criterion for level clustering analysis of a single CsPbBr₃ quantum dot. We see that the BIC of this dot peaks at $n_G = 13$. (B) The occupancy diagram of the same quantum dot. The number of occupied states keeps growing with the number of available states, saturating around $n_G = 15$. (C) Histogram showing the durations of CPA segments of the CsPbBr₃ quantum dots as scatter plot. For this dot we fit (line) a powerlaw tail with an exponent of $\alpha = 2.9$. (D) Long-term autocorrelation trace of a single CsPbBr₃ quantum dot. Quoted coefficients B, C refer to parameters in $At^{-C}\exp(-Bt)$ fitted (line) to the data (scatter plot).

our findings as well as with Ref. [280, 281, 283, 285].

This difference may be attributable to the different excitation conditions that are unique to Gibson et al. [282] relative to all other works. Gibson et al. [282] report that the lower excitation duty cycle resulting from both lower repetition rate (sub-MHz) and shorter pulse (order 0.1 ps versus 10 ps) excitation, promotes photostability. We note that from a purely experimental point of view, this benefit is not immediately clear to us, at least not when expressing photostability in number of excitation cycles as we observe dots for 2 minutes at 10 Mhz repetition rate, versus 10 minutes at 0.3 MHz in Ref. [282], at similar count rates per excitation pulse. Among possible explanations, we can exclude effects purely due to thermal load: according to established thermal analysis of nanoparticles under pulsed excitation [326] a nanocrystal and its environment cool down within nanoseconds after excitation, meaning that although our work and Refs. [280, 281, 283, 285] use higher repetition rates (up to 20 MHz), there is no ground

to believe that heating effects build up more strongly than in Ref. [282]. Regarding electronic processes, several works recently claimed that intermittency in perovskite dots arises not from one, but from several competing mechanisms including non-radiative bandgap carrier recombination, trion-mediated recombination, and hot carrier blinking [327, 328]. There is a wide range of involved time constants, some of which are hypothesized to be slower than the typical MHz laser repetition rates. For instance Ref. [328] argues that there is evidence for shallow trap states with long lifetimes (> 250 ns), and some reports claim microsecond timescale delayed emission for lead halide perovskite quantum dots [329–331] which is hypothesized to originate from carrier trapping/detrapping between the band edge state and energetically shallow structural disorder states. We note that this means that laser repetition rate is ideally a variable in experiments. However it is not trivial to extend CPA studies to deep sub-MHz repetition rates as the concomitant fall in overall count rate means that the tail of dark states will become comparable in strength to the fixed background of the single photon detectors (which contribute of order 250 counts/s in our work, summing over both Excelitas detectors).

6.4.3. Residence times

In Figure 6.4C we show a histogram of the segment lengths found by CPA. In other works, on-states and off-states are often separated explicitly by thresholding following which on-times and off-times are separately analyzed, for instance to ascertain the almost universally observed power-law dependencies and their exponents. In the case of our CsPbBr$_3$ quantum dots a level assignment in on and off states is not obvious. Therefore we simply combine *all* segment lengths irrespective of intensity level in a single histogram. These switching times are power-law distributed, at least from minimum time durations of 10 ms onwards. The short-time roll-off is consistent with the limitations of the information content of the discrete photon event data stream: for segments shorter than 50 photons or so, even if physically there would be a jump, the photon number would not suffice to resolve it. Thus the roll-off does not exclude that power-law behavior also occurs for shorter times, but instead signifies that the testability of such a hypothesis is fundamentally limited. Fitting the power law $t^{-\alpha}$ for time > 10 ms indicates a power-law exponent of $\alpha = 2.9 \pm 0.1$.

The peculiar segment-duration power law statistics with exponent $\alpha \sim 2.5$ of our example dot also extends to the full ensemble. Figure 6.5E shows the distribution of power-law exponents that were fitted to the tail of the switching time histograms. We find a broad distribution of power-law exponents ranging from 1.5 to 3.0, with the bulk of the dots showing exponents in the range of $\alpha = 2.0$ to 3.7. These values are significantly higher than the values found for many semiconductor quantum dots, which generally are close to 1.5 [286]. Also, these values are significantly higher than the exponents reported for on-times of CsPbBr$_3$ dots extracted from intensity-thresholded time-binned data. We note that one can (somewhat arbitrarily) threshold CPA-segmented data in an attempt to isolate 'on-times' for the bright state from the 'residence times' associated with the long dark/gray tail of states. Doing so with thresholds $I/\langle I \rangle > 1.3$ (on-state) and $I/\langle I \rangle < 0.7$ (tail of gray/dark states) estimated from Fig. 6.3, resulted in residence-time histograms for on- and off-times with similarly high slopes as we obtain for the full set. We thus find no support in our data for power laws generally being close to 1.5 or even below, as reported in other recent
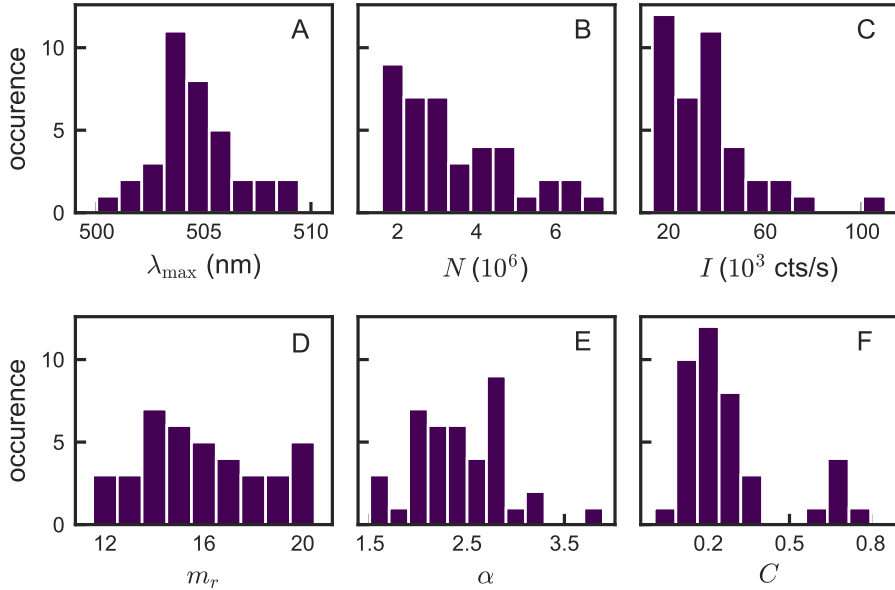
Figure 6.5: Summary of the behavior of the 40 measured single quantum dots.
We show the distribution of found (A) peak wavelengths, (B) total photon count,
(C) intensities, (D) most likely number of states, (E) powerlaw exponents of the
switching time α, and (F) the power-law exponent of the autocorrelation C.

reports [282, 285]. We note that apart from the methodological difference of not working
with binned thresholded data but with CPA analysis, also the selection of dots reported
on may matter. In this work we report on all dots identified as single photon emitters
by their $g^{(2)}(0)$. Instead in Ref. [285] dots are reported to have been selected as those for
which inspection of binned time traces suggested the most apparent contrast between
bright and dim states, qualitatively appearing closest to bimodal behavior. According
to our analysis of FDIDs and in light of the MRC model, this post-selection may not
single out the most representative dots.

An alternative approach to quantifying blinking statistics and power law exponents
that requires neither thresholding binned data nor CPA is to simply determine intensity
autocorrelation functions $g^{(2)}$ for time scales from ms to seconds, as proposed by Houel
et al. [295]. According to Houel et al. [295] the normalized autocorrelation minus 1 may
be fit with the equation $At^{-C}\exp(-Bt)$. Figure 6.4(D) show such an analysis for the
exemplary dot at hand, for which we find a reasonable fit with $C = 0.39$. As Figure 6.5F
shows, across our collection of dots we generally fit exponents C in the range 0.10 to
0.75 to intensity autocorrelation traces. We note that the relation $\alpha = 2 - C$ put forward
by Houel et al. [295] is only expected to hold for two-state quantum dot, and C does
not relate directly to α for quantum dots in which more than two states are at play.

6.4.4. Memory effects, aging and correlations in CPA sequences

Finally, we examine the dots for aging and memory effects, leveraging the fact that CPA gives an unbiased data segmentation into segments $n = 1 \ldots N$ that are classified by segment duration $T_1, T_2, \ldots$, intensity in counts/sec $I_1, I_2, \ldots$, and decay rate $\gamma_1, \gamma_2, \ldots$ that is established without any distorting temporal binning. Memory effects were first studied by Stefani et al. [324] for II-VI quantum dots, and later for perovskite dots in [285], in both systems evidencing memory effects in on/off times. We present results again for the same example dot as in Fig. 6.1 in Figures 6.6 and 6.7. With regard to aging, one can ask if over the full measurement time in which a dot undergoes of order 10^8 excitation cycles, the distribution of segment duration, intensity, and decay rate show any sign of change. To this end, we subdivide the total measurement period (e.g., Fig.6.6(A-C), total measurement time 60 s for this dot) in 100 slices that are equal length in terms of wall-clock time, and examine the evolution of histograms of I_n, γ_n and T_n for these short measurement intervals as a function of their occurrence in the measurement time. As the residence times are very widely distributed, we plot histograms of $\log_{10} T_q$, with q the index of the segments. There is no evidence that any of these observables change their statistical distribution over the time of the measurement. While Fig.6.6(A-C) shows an example for just one dot, this conclusion holds for all dots in our measurement sets, with the caveat that for some dots drifts in microscope focus caused a small gradual downward drift in intensity. We observed no photobrightening of dots during the experiment.

Clustering allows us to ask questions that are not accessible with simple binning of data, as we can examine the data sets for correlations between parameters and between subsequent segments. In terms of cross-correlating different observables, beyond FDIDs that correlate intensity and decay rate, one can also examine correlations between intensity and segment duration, and between decay rate and segment duration. Histogramming the clustered data to screen for such correlations (Fig.6.6(D,E)) show that both the distribution of intensities and of decay rates are uncorrelated, or only very weakly correlated, with the segment duration. In other words, we find no evidence that within the distribution of states between which the dot switches, some states have different residence time distributions than others.

Memory effects [285, 324] should appear as correlations in the values for any given observable in subsequent segments, i.e. in conditional probabilities that quantify what the probability $P_{\Delta n}(A|B)$ is that a chosen observable to obtain a value A is given that it had a value B in the previous segment ($\Delta n = 1$), or generally counting Δn events further back into the history of previous segments. Figure 6.7 shows such conditional probabilities for $\Delta n = 1$ (panels A-C) and $\Delta n = 2$ (panels D-F) , for intensity (panels A,D), decay rate (panels B,E) and segment duration (panels C,F). These diagrams are obtained by applying a simple 2D histogramming approach, listing the value of B as $x-$axis, the value of A as $y-$axis, and normalizing the sum of each of the columns to obtain a conditional probability. We note that this approach means that at the extremes of the histograms (far left, and far right), there are few events to normalize to, leading to large uncertainty. When screening for memory in intensities, it's important to consider that the CPA algorithm itself *selects* for intensity jumps. Due to this the intensity after one jump ($\Delta n = 1$) is *a priori* very unlikely to achieve a similar value, which leads to a near-zero conditional probability at the diagonal of Figure 6.7A. Nonetheless, the distinct features in the diagram at $\Delta n = 2$ (Fig. 6.7D) do suggest that the dot generally

125

Figure 6.6: Analysis of (absence of) aging during photocycling of a single perovskite quantum dots, histogramming intensity (A), decay rate (B) and segment duration (C) in slices of 0.9 sec for a total measurement of 90 sec. (D, E) correlation histograms of intensity versus segment duration, and decay rate versus segment duration, evidencing that these are uncorrelated quantities.

alternates repeatedly back and forth between a bright and a dark state. More telling than diagrams for intensity are those for decay rate. They show that if, in a given step the decay rate is low (slow, bright feature in FDID at $< 0.1 \, \mathrm{ns}^{-1}$), then in the subsequent step the decay rate is usually fast, yet widely distributed from 0.1 to 1 ns^{-1}, and vice versa from any of the fast decaying states, the dot is likely to jump to the quite narrowly defined slow rate of the bright state. If one considers the conditional rate at $\Delta n = 2$, the conclusion is that if the dot is in the bright state with its slow decay rate at a given step, then likely after two jumps it comes back to this bright, slowly decaying state. If, however, the rate was fast anywhere in the interval from 0.1 to 1 ns^{-1}, then after an excursion to the slow rate at $\Delta n = 1$, the dot likely in the second step again takes on a fast rate in the interval from 0.1 to 1 ns^{-1} but without a particularly clear preference for any value in that wide interval. Finally we note that there is no indication in our data that subsequent residence times ($P(T_{m+\Delta n}|T_m)$ show any memory (Fig. 6.7(c,f), showing result for $\log_{10} T_m$). Thus our data do not confirm the observation of Hou et al. [285] that there are memory effects in subsequent on-off times. Those memory effects mirror the mirror effects observed by Stefani et al. in 2005 for II-VI quantum dots [324], and indeed the MRC model [321] predicts memory effects in subsequent on/off durations. We note that the analysis in these previous works is contingent on thresholding to define on-off states and times, a process counter to the findings of CPA analysis that there are not simply two intensity levels. Moreover, we note that in this work we indiscriminately report on all dots we identified as single photon emitters by their $g^{(2)}(0)$, instead of post selecting those that qualitatively appear closest to bimodal behavior as in Ref. [285]. The fact that the very definition of on-off time is unclear for these quantum dots rather defies analysis of memory in these quantities in the terms used by [285, 321, 324]. Since it appears that the dots at hand switch between
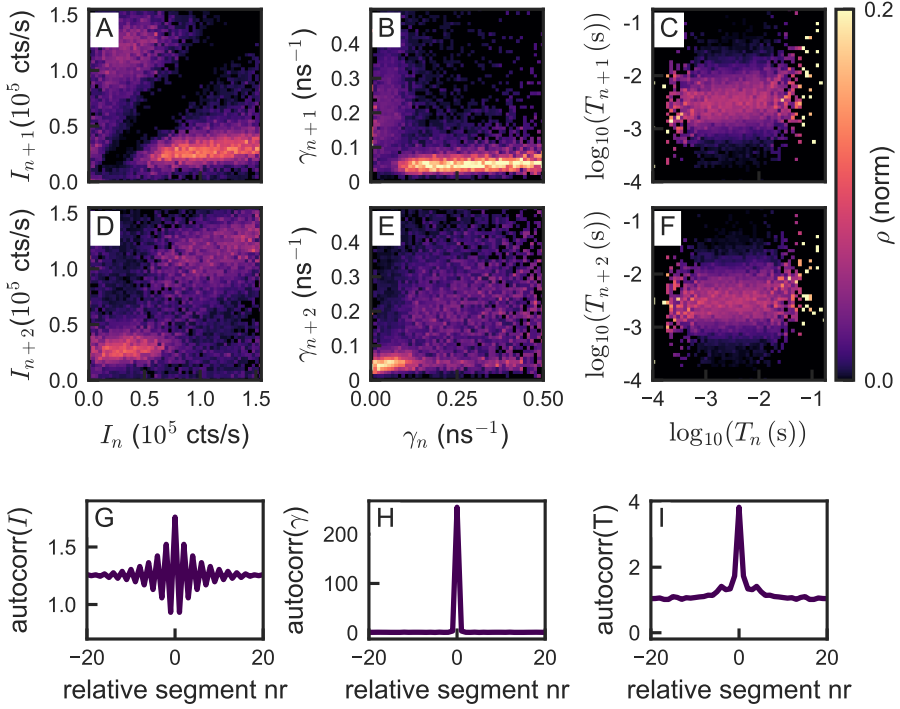
Figure 6.7: (A-F) Conditional probability of observing a value for intensity (A, D), rate (B, E) or segment duration (C, F), given the value of the same observable one or two steps earlier, respectively. (G,H,I) Normalized autocorrelation (difference from 1) of the sequence of intensity, decay rate and segment durations. This data is for the same single dot as considered in Fig. 6.6.

a reasonably unique bright state and the entire tail of dark gray states suggests to define on-times, as selected from CPA by thresholding at circa $I/\langle I \rangle > 1.3$. With this approach, we found no memory effects for the sequence of on-times.

One could speculate that the information gleaned from such conditional probability diagrams could be advantageously condensed in autocorrelations of the traces I_q, γ_q and T_q. We plot normalized autocorrelation traces $G(\Delta q) - 1$ where for any sequence $H \in I_q, \gamma_q, T_q$, one defines

$$G(\Delta q) = \langle H(q)H(q + \Delta q)\rangle/\langle H(q)\rangle^2$$

(where $\langle . \rangle$ denotes the mean over q are all segment indices and Δq is $1, 2, \ldots$), so that at long times $G_m - 1$ vanishes.

In Fig. 6.7(G-I) we plot $G(\Delta q) - 1$ for intensity, rate and segment duration. Such segment-autocorrelations are distinct from, e.g., the usual intensity autocorrelation traces that one might examine to determine blinking power laws, since here one autocorrelates subsequent intensity segments *without any regard for their time duration*.

For a conventional two-level dot, the autocorrelation trace I_q would oscillate with large contrast up to very large q. Instead, we find that the dot at hand shows an oscillation with a distinct contrast in the intensity segment autocorrelation contrast for up to 5—10 cycles. In the normalized autocorrelation for decay rates, the memory is far less evident. We attribute this not to a lack of memory, but note that if a dot switches between a state of well-defined slow rate, and an array of states with highly distributed fast rates, then upon averaging the wide distribution of fast rates washes out any autocorrelation signature. Finally, the residence times, which we already found to be uncorrelated between subsequent jumps, show no autocorrelation signature for $z \neq 0$. A similar behavior to that shown in Figure 6.7G was observed for circa 30% of the dots studied, with other dots showing no clear intensity autocorrelation.

6.5. Conclusion

To conclude, we have reported on intermittency properties of a large number of CsPbBr$_3$ quantum dots on basis of a Bayesian inference data analysis. This approach works with raw, unbinned, photon counting data streams and thereby avoids artifacts commonly associated with the analysis of time binned data. We find that dots have in addition to their bright emissive state a tail of gray states that qualitatively appears continuous in FDID diagrams, and that according to clustering analysis requires at least 10 to 20 levels to describe, if a discrete-level description would be appropriate. Thereby our work provides a confirmation of claims in earlier works [280, 281, 283] under similar excitation conditions, with the distinction that we do not use time binned data but rigorously exploit all the information in the data stream to the level that its intrinsic noise allows. We note that the same type of dots have displayed a different behavior, indicative of 2 to 3 levels, in Ref. [282]. Since that work uses almost identical Bayesian inference methods, we conclude that this distinction is really due to the different physical realization. Alongside possible differences in sample preparation, we note that Ref. [282] also stands out from all other reports due to its quite different excitation conditions, particularly using shorter pulses and significantly lower pulse repetition rates. While Ref. [282] states that this choice improves photostability, when expressed in number of excitation cycles, our experiment is not actually at a disadvantage in terms of photostability since we follow dots for 2 minutes at 10 MHz repetition rate, versus for 10 minutes at 0.3 MHz. Also, our estimates exclude the idea that higher pulse repetition rates but at similar per pulse excitation densities, would cause a more significant heating of the dot that would explain thermally activated modifications since nanoparticles under pulsed excitation loose their energy to the environment in nanoseconds [326]. A possible explanation might lie in the fact that perovskite quantum dots have been reported to have slow-time constant electronic processes, such as delayed exciton emission (microsecond time scales) [329–331], and shallow trap states with lifetimes exceeding 250 ns [327, 328]. These observations imply that there are photophysical processes that may be involved in blinking and flickering, and that may not fully relax at higher laser repetition rates. Finally, a caveat on experimental limitations in the effort to determine dim intensity levels is that, even if the physics is unchanged, lower repetition rate experiments are less likely to identify many gray /dark states once the dark state count rates approach detector background levels. In our setup, the combined dark count rate of both detectors is of order 200—250 counts

per second, meaning that the darker levels would be comparable in count rate if we would reduce the excitation rate 30-fold.

Overall, our results support [280, 281, 283], and as in the first report proposing the validity of the MRC model [320, 322] for perovskite dots [281], we find that the tail of gray states display an inverse correlation between intensity and rate, suggesting that the dots have a unique bright state with a given decay rate, to which random activation of recombination centers add nonradiative decay channels. However, we note that this observation merits further refinement of models: while plotting intensity versus lifetime may point at strict proportionality, plotting rate instead of lifetime accentuates deviations, notably a deviation in curvature of our data relative to inverse proportional dependence. Finally, we have analyzed correlations in the measured CPA-segmented sequences of intensity levels, decay rates, and segment lengths. We find no evidence for aging, i.e., gradual shifts in e.g., decay rate or blinking dynamics during photocycling of dots through 10^6 to 10^7 detected photons (i.e., well over 10^8 cycles). Also, our data indicate that residence times are not correlated to the state that a dot is in. The residence times can be fiducially extracted for a limited time dynamic range from ca. 5-15 ms, limited intrinsically by count rate, to ca. 10 s, limited by the length of the photon record. We note that in residence time histograms determined by CPA, according to Monte Carlo simulations at long times the analysis fiducially reports on power laws without introducing artifacts, such as apparent long-time roll-offs. The exponents that we find are in the range from 1.5 to 3.0, which appears high compared to the near-universal value of 1.5 observed for II-VI single photon sources. In the domain of CsPbBr$_3$ dots, reports have appeared of even lower exponents (down to 1.2) [282, 285] with exponential roll-offs at times ~ 0.1 s that cause a steepening of the residence time histogram at longer times. We note that although exponential roll-off certainly steepens slopes in the residence time histogram, our histograms do not point at exponential, but at high-exponent power law behavior.

Regarding memory, we found a distinct memory effect in intensity and rate in the sense that dots appear to switch between a quite unique bright state with a slow decay rate that is evident as the bright pocket in FDIDs, and the entire tail of dim states in the FDID. Moreover, the dots do not appear to return preferentially to this dim state, but explore the entire tail anew at each transition from the bright state. This is an important refinement on the MRC model which in itself leaves open if dots return at all to the bright state before choosing another dim state, and which does not specify if dots make repeated visits to the same dim state or not. In terms of analysis, this memory is only partially visible in autocorrelation traces of sequences of CPA intensity and rates, as the dim states are so widely distributed. The averaging involved in evaluating autocorrelations washes out some memory effects that do appear more clearly in conditional probability histograms reporting on subsequent jumps. Finally, we found no evidence in our data set for the apparent memory in residence times (segment lengths T_q) reported by Hou et al. [285] for on-times.

In our view, this rich data set will stimulate further theory development in the domain of inorganic quantum dot intermittency. Compared to the case of II-VI quantum dots, a host of different effects could be at play in perovskite quantum dots. For instance, vacancy concentrations in perovskites are orders of magnitude higher than those in II-VI materials, and vacancies are highly mobile, which may affect photoluminescence [325]. Also, halide perovskites are known to undergo reversible surface (photo)chemical reactions. Given the role of surface defects in blinking (as understood for II-VI dots), this

may be highly important for perovskites. Blinking studies in different environmental gases could elucidate this [332]. Also, one could speculate that the strong polaronic effects in perovskites affect blinking, through involvement with the screening of trapped charges [333]. Finally, in terms of electronic structure perovskite materials are different from II-VI dots not just in weak versus strong confinement, but also in having strongly anharmonic potentials, near-equal hole and electron effective masses, and a band structure that causes defect levels to be at shallow trap levels, instead of deep trap levels (see Ref. [327–331] for possible relations to intermittency).

Summary

> *"'Speak English!' said the Eaglet. 'I don't know the meaning of half those long words,*
> *and what's more, I don't believe you do either!'"*
> — Lewis Carroll, *Alice's Adventures in Wonderland*, 1865

LIGHT drives human life. As a species, sight is by far the dominant sense we use to navigate the world. With the advent of radio, optical fibers, and nanofabrication on basis of optical lithography, our uses and dependency on light has only increased. It is used on a global scale for optical fiber communication, which enables the World Wide Web. Ever improving techniques in nanofabrication enable us to carry smartphones in our pockets that are more powerful than the computers that brought humans to the moon. The uses of light-based technology are immensely diverse, from using luminescent markers to study the life cycle of single cells, tracking minute chemical changes in blood, to improving solar cells. Light is used in applications on scales as small as visualizing single atoms and molecules, and as large as mapping out our intergalactic neighborhood.

In this thesis, we study the interaction of light and matter at the nanoscale, and we aim to bring together a number of different components. The first of these concerns the trapping of light in so-called optical hybrid resonators. Such hybrids bring together different methods of containing light, where the goal is to hold on to the light as long as possible, while simultaneously squeezing it into the smallest possible volume. The other building block we investigate is that of single emitters. These are systems that absorb and emit only a single photon at a time, making them excellent communicators for small-scale electronic information. The second half of this thesis is dedicated to studying a particular type of these emitters, as well as how to best investigate their behavior.

When it comes to light trapping, there are two main approaches. The first is that of optical microcavities. This is a class of systems, often made from transparent materials, that can trap the light for very long times, while it reverberates many hundreds or even thousands of times, giving them what is called a high quality factor (Q). The other approach is that of optical antennas. These are made of metallic materials, often gold or silver. Just as a radio antenna has a strong interaction with radio waves, an optical antenna resonates in tune with visible light. These systems are capable of squeezing the light into volumes much smaller than the wavelength of that light for a short period of time, usually only about ten cycles.

Combining an antenna and a cavity creates an optical hybrid, which can be used to combine the positive qualities of either component. In Chapter 2 we examine how small a volume the light can be squeezed into, and how long it can be retained in these

hybrid systems. These characteristics together determine the strength of light-matter interaction through a figure of merit that is known as the "Purcell factor". On basis of the achievable Purcell factors we examine the merits of hybrids for various applications in classical and quantum optics, particularly asking if hybrids can avoid the need to cool to temperatures close to absolute zero, which is $-273.15\,°C$.

The approach of Chapter 2 is to first derive design rules for hybrid resonances from a simple analytical model. These rules are checked against full-wave simulations of hybrids composed of nanobeam cavities and gap antennas. Specifically, we study nanobeam optical hybrids, made of silicon nitride (Si_3N_4) and gold. Si_3N_4 is, just like glass, transparent to visible light. Our optical hybrid design is somewhat similar to an optical fiber to guide light, together with periodic variation of refractive index to reflect light by constructive interference, much like the reflective properties of opals and mother of pearl. The nanobeam design consists of a rectangular beam of Si_3N_4 with holes drilled into it at regular intervals. These holes reflect the light, creating a mirror. In the middle of the beam, we change the size and distance between the holes, which creates an optical cavity between two of these mirrors, where the light can be trapped. On top of this nanobeam cavity, we place an optical antenna, which has the shape of two ellipses with the pointy ends nearly touching. This traps the light mostly in that narrow gap. Combining the beam and the antenna creates the optical hybrid.

We find that hybrids can outperform the plasmonic and cavity constituents by combining a bit of the best of both methods of trapping. Additionally, we find that by changing the resonance frequency of the antenna and cavity with respect to each other, we can trade off the confinement of the light against the duration of the light trapping, while keeping their ratio roughly the same. This detuning of the components allows us to span the full range of parameters in between the two components.

In Chapter 3 and Chapter 4 we report on our experiments to realize and optically interrogate our design of optical hybrids. We modified the design with respect to Chapter 2 slightly to a system that is easier to fabricate. Instead of beams suspended in air, we work with a design that works for beams lying on glass. In Chapter 3 we describe the process to physically realize these systems using nanofabrication techniques. We start with a thin substrate of silicon that acts as a mechanical support layer. On top of that there is a thin layer of glass, and an even thinner layer of Si_3N_4, only $200\,nm$ thick, which is over 100 times thinner than a human hair. Since Si_3N_4 has a higher refractive index than glass, it can support guided modes. Using a tightly focused electron beam, we write the structures of these nanobeams on a polymer, which are then transferred into the Si_3N_4 using a high-power plasma. Using a different polymer, we inscribe the antennas over the existing beams, again with an electron beam. This forms a mold that allows us to evaporate gold onto the cavities in the shape of our antennas, resulting in two ellipsoidal nanoparticles with a $20\,nm$ gap between the pointy ends. With proper alignment of the cavity and antenna, we create an optical hybrid.

In Chapter 4 we study the behavior of the components of these hybrids independently, and subsequently the behavior of the hybrid as a whole. To identify the nanobeam resonances we use a combination of techniques. First, we study the nanobeam cavities by using photoluminescence of the Si_3N_4. When we shine a green laser on the center of the cavity, the Si_3N_4 will emit light in the yellow and red, which have lower energy than green light. At the resonance frequency of the cavity, the light experiences interaction with the cavity, and slightly more light is emitted, allowing us

to determine the resonance of the cavity. We find that these resonances lie between 777 nm and 802 nm, where we can control the resonance by scaling the size of the beam during fabrication. A more precise measure of the cavity resonances and quality factors is obtained in a next step, where we use a narrowband tunable laser to measure resonant transmission through the nanobeam cavities. To study the antennas, we use a so-called dark field microscopy setup, where light is sent onto the antennas at a glancing angle, and only light that is closer to straight up is collected. This way, we look only at light that interacts with the antennas. By looking at this light as a function of its wavelength, we can determine the resonances of the antennas. We find that the antennas have resonances between 600 nm and 680 nm, where the resonance wavelength again scales with antenna size.

We study the behavior of the hybrids systems using optical gratings, fabricated on either end of the beams. This allows us to send light into the beam on one end and collect it from the other. We find that the presence of an antenna shifts the resonance of the hybrids with respect to that of the cavity by an average of 1.2 nm, with stronger shifts observed for larger antennas, and vice versa. We also find that the presence of an antenna broadens the resonance of the hybrids with respect to that of the bare antennas, shifting from a width of about 1.4 nm to 1.7 nm to 3 nm. This is in line with our expectations, where we had predicted a maximum shift and broadening of around 2 nm. Finally, we have also explored LDOS effects of hybrids, using the vibrational Raman scattering signature of reporter organic molecules deposited on the antennas. The gold antennas were functionalized with BPT (biphenyl-4-thiol) molecules which show a distinct SERS (surface enhanced Raman) signal. The SERS signal is proportional to the product of pump field enhancement, and the LDOS enhancement at the frequency of the interrogated vibrational sideband. By using our tuneable diode laser as pump, we can arrange either the pump wavelength, or the frequency of an interrogated vibration, to be resonant with a hybrid mode. We indeed find that if the pump wavelength is resonant with a hybrid mode, all vibrational lines in Raman spectra are equally affected. Conversely, if the pump is not resonant but just one vibrational line is, only that vibrational line is modified by the hybrid resonance. The enhancements generally show Fano lineshapes consistent with the expected behavior for hybrids. We can obtain SERS signal both by purely free-space addressing, and through using the waveguide as either the pump, or the detection channel.

The second building block of this thesis concerns single photon emitters. Many applications in, for instance, quantum communication and information processing require the combination of single quantum emitters that emit single photons at a time with resonators such as microcavities, antennas, or the hybrids that we realized. A significant challenge for building such applications on the basis of single emitters is presented by the fact that many emitters show blinking, also known as intermittency. The signature of intermittency is that emitters seemingly randomly lose and regain their ability to emit light. That is, an emitter that is continuously pumped with light will, seemingly randomly, switch from emitting many single photons, to emitting only a few, and back. A main challenge in the field is to identify novel promising single photon emitters, and to assess if they suffer from intermittency, and if so, why they suffer from intermittency.

In Chapter 5 we discuss a Python-based toolbox for statistical analysis of fluorescence intermittency properties of single emitters. When analyzing data collected from

blinking single emitters, many methods introduce some form of bias in determining the characteristics of intermittency, such as in the identification of intensity levels and in the identification of the switching timescales. This results in artifacts that hinder accurate interpretation of the intermittency. The method used in this toolbox maximally eliminates bias, allowing one to retrieve maximum information from the data at hand, with minimal artifacts.

The basis of the Python toolbox lies in Bayesian inference, which objectively identifies the most likely hypothesis for the process that generates a certain observed behavior, given the available dataset. This paradigm should be compared to usual probability theory, which given a hypothesis describes the likelihood of drawing a certain dataset. For intermittency, the Bayesian viewpoint leads to socalled Bayesian changepoint analysis (CPA). Applied to the data from single emitters, CPA provides the most likely segmentation into time segments, where each segment has a constant intensity, and between which an intensity jump occurs. After the segmentation, one can again use tools from Bayesian statistics to find the most likely description of the data in a small number of intensity levels, where the challenge is both to find the number of states, and their brightnesses. The toolbox can be applied to both experimental data to quantify its behavior, and to simulated data, to investigate the implications of intermittency models. We provide a detailed analysis of simulated quantum dot behavior to illustrate these tools. In the spirit of Open Science, this toolbox is posted on GitHub and Zenodo. We benchmark the toolkit with a large amount of simulated data. We discuss that there is an upper limit to the amount of information that can be extracted from a given data set and that this upper limit is set by the total number of measured photons.

In Chapter 6 we report measurements on a particular type of single emitter: a nanocrystal with chemical formula $CsPbBr_3$. This is one example of the emerging class of perovskite materials, which are solution-processable semiconducting materials with very advantageous properties for light emission and absorption. Sufficiently small nanocrystals show single photon emission behavior, indicative of quantum confinement, much like traditional colloidal quantum dots. We measure a large number of these nanocrystals individually, pumping them with blue light, The nanocrystal under study will hold on to the energy from an absorbed photon briefly, and then release it with a slightly lower energy. We measure the times at with these photons are emitted, and apply the toolbox described in Chapter 5 to the resulting data.

We find that perovskite quantum dots display a large number of gray states. These are states that are not quite bright, but not quite dark either. We see that brightness broadly speaking correlates with the time between absorption of a photon, and emission of another. That is, an emitter in a bright state will hold on to the energy for a longer time before re-emitting it, whereas the same emitter in a dark state will release the energy more quickly. This behavior is markedly different from that of usual CdSe colloidal quantum dots. At the same time, the behavior is consistent with the so-called multiple recombination centers model, which hypothesizes that there is a multitude of defect centers that can cause nonradiative decay. The random switching creates a large number of intensity levels. We use our toolbox to examine the effects of aging on the quantum dots, as well as to investigate if the behavior has some memory. We find that dots tend to always return to the bright state before jumping to another dim state, and that when choosing a dim state they tend to explore the entire set of states available.

SAMENVATTING

" *'Hebben jullie niets beters te doen dan raadsels opgeven waar geen antwoord op bestaat?*
Daar ga ik mijn tijd niet aan verspillen, hoor.' "
— Lewis Carroll, *De Avonturen van Alice in Wonderland*, 1865

Lıcht speelt een belangrijke rol in het menselijk bestaan. Onze ogen zijn het dominante zintuig waarmee we de wereld waarnemen, en met de komst van de radio, glasvezelkabels, en van nanofabricage op basis van optische lithografie, is de veelheid aan toepassingen waarvoor we licht gebruiken sinds de 20e eeuw sterk toegenomen. Optische fibers wikkelen zich om de aarde en verbinden ons met het Wereldwijd Web. Verbeteringen in nanofabricage stellen ons in staat om mobiele telefoons bij ons te dragen met meer rekenkracht dan de computers die de mensheid naar de maan bracht. De lengteschalen waarop licht toepassingen heeft gaat van het in beeld brengen van enkele atomen en moleculen aan het fundament van het bestaan, tot het in kaart brengen van ons heelal.

In dit proefschrift bestuderen wij de interactie van licht en materie op de nanoschaal, ter grootte van de golflengte van het licht zelf. Daarbij willen we een aantal verschillende bouwblokken bij elkaar brengen. Het eerste bouwblok betreft het vangen en opsluiten van licht in zogeheten optische hybride resonatoren. Deze combineren verschillende methodes om licht te vangen om licht zo lang mogelijk vast te houden en tegelijk in een zo klein mogelijk volume te persen. Het tweede bouwblok gaat over enkele fluoroforen. Dit zijn kleine stukjes materiaal, soms zelfs een enkel molecuul, die slechts een enkel foton per keer kunnen absorberen, en dan ook maar een enkel foton per keer uitstralen. Dit gedrag maakt ze onmisbaar voor informatieverwerking op de kleinste schaal. Het tweede deel van dit proefschift betreft het gedrag van een nieuw soort fluoroforen, en statistische methoden om het beste naar hun gedrag te kijken.

Als het gaat over het vangen en vasthouden van licht, zijn er grofweg twee aanpakken te onderscheiden. De eerste aanpak is die van de optische trilholtes. Zij vangen licht en houden het gedurende lange tijd vast, terwijl het licht honderden of duizenden keren rondzingt in een relatief groot volume. De andere aanpak is die van optische antennes. Zij houden licht maar kort vast, maar persen het in volumes vele malen kleiner dan de golflengte.

Wanneer je deze twee methodes van licht vangen combineert, geeft dat een zogeheten optische hybride resonator, die onder bepaalde voorwaarden de positieve eigenschappen van beide componenten kan combineren. In Hoofdstuk 2 bestuderen wij op basis van numerieke en analytische theorie hoe klein we het licht kunnen samenpersen en hoe lang we het vast kunnen houden in deze hybride systemen. Deze karakteristieken samen definiëren een belangrijke maat voor licht-materie interactie, vaak gevangen in een getal dat "Purcell factor" heet. Op basis van Purcell factoren die

bereikt kunnen worden beoordelen we de geschiktheid van hybrides voor verschillende toepassingen, waarbij we ons met name afvragen of men door het gebruik van hybrides niet hoeft te koelen tot nagenoeg het absolute nulpunt van $-273.15\,°C$. We leiden algemene regels af voor het ontwerp van een goede hybride en verifiëren deze regels door te vergelijken met numerieke simulaties van ons eigen ontwerp voor een hybride, van silicium nitride (Si_3N_4) en goud. Si_3N_4 is, net als glas, grotendeels doorzichtig voor zichtbaar licht, maar staat ons toe om er licht in te vangen.

Onze hybride lijkt een beetje op een stukje glasvezel. Het is een langgerekte balk Si_3N_4 waar op regelmatige afstanden gaten in zijn geboord. Deze gaten weerspiegelen het licht en vormen een soort halfdoorlatende spiegel. In het midden is een set gaten met afwijkende grootte en onderlinge afstand gemaakt, waar het licht tussen twee van de halfdoorlatende spiegels kan worden opgesloten. Precies in het midden, bovenop de bundel, plaatsen wij een gouden antenne. Deze heeft de vorm van twee ellipsen die met de punten naar elkaar wijzen, met slechts $20\,nm$ er tussen. Samen vormt dit onze optische hybride. Wij laten in dit hoofdstuk zien dat hybrides hun componenten kunnen overtreffen, omdat ze de beste eigenschappen uit beide kunnen halen. Daarbij laten wij zien dat, door de resonantie frequentie van de antenne en de trilholte ten opzichte van elkaar te verstemmen, we het volume van de opsluiting van het licht kunnen uitruilen tegen de duur ervan, terwijl de verhouding hiertussen ongeveer gelijk blijft. Door deze verstemming laten wij zien dat het mogelijk is alle parameterwaarden te bereiken tussen die van de onafhankelijke componenten.

In Hoofdstuk 3 en 4 bespreken we onze experimenten aan onze optische hybrides. We passen het ontwerp enigszins aan, zodat de trilholte op glas ligt in plaats van in de lucht te hangen. In Hoofdstuk 3 beschrijven we het complexe proces dat we doorlopen om deze hybrides te realiseren door middel van nanofabricage. We beginnen met een plakje silicium wat als een mechanisch stabiele onderlaag dient. Daarbovenop ligt een veel dunner plakje glas, met daar weer bovenop een aller dunst plakje silicium nitride (Si_3N_4) van $200\,nm$ dik. Dit is meer dan $100\,$maal dunner dan een menselijke hoofdhaar. Omdat Si_3N_4 een hogere brekingsindex heeft dan glas kan het licht geleiden zonder dat het weglekt in het glas. Met een gefocuseerde elektronenbundel schrijven wij het patroon van onze nanobundel trilholtes in een polymeer dat we voor dit doel neerleggen op het Si_3N_4, waarna we dit patroon in het Si_3N_4 overbrengen door middel van een hoog-energetisch plasma dat het Si_3N_4 etst. We gebruiken de elektronenbundel een tweede keer, nu met een ander polymeer, om een gat in het polymeer, in de vorm van de antennes te schrijven. Wanneer we goud verdampen en vervolgens laten neerslaan op het monster, daalt het neer in de gaten, in de vorm van de antennes. De antenne en de trilholte samen vormen onze optische hybride.

In Hoofdstuk 4 bespreken we optische experimenten aan de hybrides. We bestuderen de eigenschappen van de hybrides eerst onafhankelijk van elkaar, voordat we de hybride in zijn geheel bestuderen. We bekijken het gedrag van de trilholtes door middel van fotoluminescentie van het Si_3N_4. Wanneer we een groene laser op de Si_3N_4 trilholte schijnen, wordt het deels geabsorbeerd, en weer uitgestraald met een lagere energie, in het geel en rood. Bij de golflengte waar de trilholte resonant is, wordt er een klein beetje extra licht gestraald, waardoor we deze resonantie kunnen aflezen. We vinden dat onze resonanties tussen de $777\,nm$ en de $802\,nm$ liggen, afhankelijk van een schaalfactor waarmee we de trilholtes groter maken en daarmee verstemmen. We bestuderen ook het gedrag van de antennes, waarbij we een zogeheten donker-veld microscoop

gebruiken. Hierbij wordt licht onder een scherpe hoek op de antennes gestuurd, en het licht dat min of meer loodrecht op het monster wordt uitverstrooid, wordt opgevangen. Op deze manier zorgen we er voor dat alleen licht dat verstrooit aan de antennes, wordt bekeken. Hiermee vinden wij dat de golflengte van de antennes tussen de 600 nm en 680 nm ligt, waar de resonantie schaalt met de grootte van de antenne, en grotere antennes resoneren bij langere golflengtes. We bestuderen vervolgens het gedrag van de hybrides door middel van een tweetal optische tralies, één aan elk eind van de langwerpige trilholte. Deze tralies zorgen dat wij licht aan een kant van de trilholte kunnen inkoppelen, en aan de andere kant kunnen uitlezen. Wij vinden dat de antenne de resonantie van de trilholte gemiddeld verschuift met 1.2 nm, en verbreedt met 1.7 nm tot 3 nm, wat in lijn ligt met onze verwachtingen.

Tot slot hebben we ook de licht-materie interactiesterkte bestudeerd door gebruik te maken van het Raman-signaal van op de antennes aangebrachte organische moleculen. Organische moleculen op goud antennes geven een duidelijk SERS (oppervlakte-versterkte Raman verstrooiing) signaal, waarbij Raman verstrooiing door moleculaire trillingen vele ordes van grootte wordt versterkt. Dit signaal is evenredig met het product van veld versterking van het aandrijvende pomp licht enerzijds, en anderzijds de verhoging van de optische toestandsdichtheid bij de Raman-verschoven frequentie. Door middel van onze verstembare diode laser kunnen we zorgen dat ofwel de pomp laser, ofwel één van de trillingen in het Raman spectrum precies afgestemd is op de hybride resonantie. We vinden dat als de pomp laser precies op de hybride resonatie past alle lijnen in het Raman spectrum op dezelfde manier door de hybride mode worden gewijzigd. Anderzijds, als niet de pomp maar een Raman lijn resonant is met de hybride wordt alleen deze lijn door de hybride mode extra versterkt. De versterkingen laten over het algemeen Fano-lijnvormen zien zoals verwacht voor hybrides. We kunnen de Raman signalen bovendien aanspreken dooor middel van de golfgeleiders.

Het volgende bouwblok van dit proefschrift gaat in op enkele fluoroforen. Veel toepassingen in quantum communicatie en quantum informatie technologie vereisen het samenbrengen van fluoroforen die enkele fotonen per keer kunnen uitzenden met een optische trilholte, antenne, of hybride. Helaas hebben veel van zulke fluoroforen, vaak individuele moleculen of enkele quantum dots de eigenschap dat ze knipperen. Dat wil zeggen dat ze snel en schijnbaar willekeurig tussen helderheden wisselen. Ook als ze steeds even sterk door pomplicht worde aangedreven stralen ze soms veel fotonen uit, en andere periodes juist minder. Een belangrijke uitdaging is om betere fluoroforen te vinden als geschikte kandidaat voor bronnen van enkele fotonen. Bij deze zoektocht is het belangrijk om vast te stellen of en waarom knipperen optreedt.

In Hoofdstuk 5 bespreken wij een statistische analyse gereedschapskist die ge-realiseerd is in de programmeertaal Python. Wanneer metingen van knipperende fluoroforen worden geanalyseerd, liggen tal van artefacten op de loer, die worden veroorzaakt doordat veel voor de hand liggende analyse methoden helaas ongemerkte vooroordelen introduceren. De methode die hier wordt gebruikt, is maximaal vrij is van vooroordeel. Dit zorgt er voor dat de maximale hoeveelheid informatie uit de data gehaald kan worden met minimale artefacten. De basis hiervoor ligt in Baysesiaanse statistische methoden die objectief op basis van data een uitspraak kunnen doen over hoe waarschijnlijk een hypothese is over het proces waardoor de data is gegenereerd. Dit is juist omgekeerd ten opzichte van het normale perspectief van waarschijnlijkheids theorie, die beantwoordt hoe waarschijnlijk data is gegeven een hypothese, in plaats van hoe waarschijnlijk een hypothese is gegeven verzamelde data.

Toegespits op metingen aan knipperen van fluorforen heet de methode "Bayesiaanse changepoint analyse (CPA)". Deze methode geeft de meest objectieve opslitting van data in tijdsegmenten, waarbij in elk blokje tijd de helderheid van de fluorfoor even constant is, en waartussen de fluorofoor verspringt. Na het ophakken in blokjes, gebruiken we Bayesiaanse statistiek om te bepalen hoeveel onderliggende fysische intensiteitsniveaus metingen het best verklaren, en wat de helderheden zijn van deze toestanden. De toolkit kan worden toegepast op experimentele data om het gedrag van gemeten fluorforen in kaart te brengen, en op gesimuleerde data, om modellen voor het gedrag van fluorforen te bestuderen. We ijken het gedrag van onze toolkit aan de hand van een grote hoeveelheid gemeten data. We bespreken dat er een bovengrens is aan de hoeveelheid informatie die uit een gegeven meting gehaald kan worden, en dat deze bovengrens afhankelijk is van het gemeten aantal fotonen.

In Hoofdstuk 6 besprekingen we metingen aan een nieuw type recent ontdekte fluoroforen, namelijk nanokristallen met chemische formule $CsPbBr_3$. Deze materialen heten perovskieten, naar hun kristalstructuur, en zijn uniek omdat ze eenvoudig in oplossing gemaakt kunnen worden en toch uitmuntende eigenschappen hebben als halfgeleider. Perovskieten zijn zeer veelbelovend voor zonnecellen en LEDs. Nanokristallen van dit materiaal die klein genoeg zijn kunnen, door quantummechanische effecten die optreden door de opsluiting van ladingsdragers in het materiaal, zelfs enkele fotonbronnen zijn. Wij bestuderen in dit hoofdstuk vele individuele nanokristallen. Wij pompen om beurten steeds een ander, *enkel* nanokristal, die elk een enkel fluorofoor zijn, door aan te stralen met blauw licht, dat een relatief hoge fotonenergie heeft. Een fluorofoor die een pompfoton absorbeert houdt het licht eventjes vast, en straalt dan een enkel foton in het groen, wat een lagere foton energie heeft. Deze cyclus kan zich steeds herhalen: na het uitstralen van een fonton kan er weer een foton geabsorbeerd kan worden enzovoorts. Wij meten heel precies de tijden waarop fotonen worden uitgezonden, en gebruiken vervolgens de toolkit die in Hoofdstuk 5 wordt besproken om deze data te analyseren.

Wij vinden dat deze fluorforen een zeer groot aantal helderheden hebben waartussen zij wisselen, veelal groter dan 20. Dit gedrag verschilt sterk van dat van de meeste bekende colloidale quantum dot stralers. Wij vinden ook dat de helderheid correleert met de tijd gedurende welke energie wordt vastgehouden voordat het weer wordt uitgezonden. Dat wil zeggen: een fluorfoor in een heldere toestand houdt het licht langer vast dan deze zelfde fluorfoor in een donkere toestand. Hiermee bevestigen we het zogeheten "multiple recombination model", dat stelt dat er in een enkel fluorfoor rond een tiental centra zijn die elk afhankelijk van hun toestand niet stralend verval kunnen veroorzaken. Dit kunnen bijvoorbeeld defecten zijn op het oppervlak van de nanokristallen, of het gevolg zijn van uitwisseling van enkele ladingsdragers met de omgeving. Door het gezamenlijke effect van vele centers kan effectief een groot aantal helderheden ontstaan, en een brede verdeling aan tijdsduren tussen knippermomenten. Wij gebruiken onze toolkit om de effecten van blootstelling aan licht te bekijken, en om te onderzoeken of het verspringen tussen helderheden ongecorreleerd is in de tijd, of een geheugen kent in haar patroon. Wij zien dat de fluoroforen de neiging hebben om vanuit een donkere toestand het liefst steeds naar de heldere toestand te verspringen. Als ze daarentegen in de heldere toestand zijn, kiezen ze steeds willekeurig uit een veel breder scala aan minder heldere toestanden.

List of publications

Related to this thesis

1. I.M. Palstra, H.M. Doeleman, and A.F. Koenderink, *Hybrid cavity-antenna systems for quantum optics outside the cryostat?*, Nanophotonics **8**, 1513—1531 (2019) (**Chapter** 2).
   All authors conceived of the project idea jointly. I.M.P. implemented the full-wave simulations and designed the nanobeam-gap antenna hybrid. H. M. D. implemented the coupled oscillator model. I.M.P. and H. M.D jointly performed the analysis. All authors contributed to writing the manuscript. A. F. K. supervised the project.

2. I.M. Palstra, I. Shlesinger, and A.F. Koenderink, *Side-band resolved molecular optomechanics with hybrid cavity-antenna systems*. To be submitted 2021 (**Chapter** 3 **and** 4).
   I.M.P. developed the full fabrication process and realized the hybrid nanobeam-gap antenna resonators. I. M. P. and I. S. developed the tuneable laser mode spectroscopy method, and I. M. P. collected and analyzed all mode spectroscopy data. I.S. collected SERS data. I. M. P. wrote the text. A. F. K. supervised the project.

3. I.M. Palstra and A.F. Koenderink, *A Python Toolbox for Unbiased Statistical Analysis of Fluorescence Intermittency of Multilevel Emitters*, J. Phys. Chem. C **125**, 12050—12060 (2019) (**Chapter** 5).
   The authors jointly conceived the project idea and jointly implemented individual algorithms that form part of the code. I. M. P. combined all algorithms in a full workflow, optimized the Python code, and defined and performed the code benchmarks reported on in the paper and wrote the first version of the manuscript. A.F.K. supervised the project.

4. I.M. Palstra, and A.F. Koenderink, *A Python Toolbox for Unbiased Statistical Analysis of Fluorescence Intermittency of Multilevel Emitters. GitHub repository*, first commit archived at ZENODO DOI: 10.5281/zenodo.4557226 (2021) (**Chapter** 5).
   Same as above. I.M.P. documented the code and packaged the code with example data files.

5. I.M. Palstra and I. Maillette de Buy Wenniger and B.K. Patra and E. C. Garnett and A. F. Koenderink, *Intermittency of CsPbBr3 Perovskite Quantum Dots Analyzed by an Unbiased Statistical Analysis*, J. Phys. Chem. C **125**, 12061–12072 (2021) (**Chapter 6**).

The project idea was conceived by I.M.P. and A.F.K. Perovskite samples were provided by B.K.P. and photoluminescence measurements were done by I.M.B.W., with I.M.P. providing supervision. I.M.P. implemented the data analysis. A.F.K. and E.C.G. supervised the project. The manuscript was written by I.M.P. and A.F.K.; it was reviewed by all authors.

Other

1. A.I. Barreda, M. Zapata-Herrera, I.M. Palstra, L. Mercadé, J. Aizpurua, A.F. Koenderink, and A. Martínez, *Hybrid Photonic-Plasmonic Cavities based on the Nanoparticle-on-a-Mirror Configuration*, submitted. Preprint arXiv:2106.01931.

BIBLIOGRAPHY

1. M. G. Minnaert, *De natuurkunde van 't vrije veld* (Koningklijke drukkerij G.J. Thieme, Zutphen, 1937-1942).

2. P. Vukusiç and J. R. Sambles, *Photonic structures in biology,* Nature **424**, 852 (2003).

3. P. Ball, *Nature's color tricks,* Scientific American **306**, 74 (2012).

4. A. R. Parker and H. E. Townley, *Biomimetics of photonic nanostructures,* Nat. Nanotechnol. **2**, 347 (2007).

5. C. M. Eliason, R. Maia, J. L. Parra, and M. D. Shawkey, *Signal evolution and morphological complexity in hummingbirds (aves: Trochilidae),* Evolution **74**, 447 (2020).

6. J. W. Galusha, L. R. Richey, J. S. Gardner, J. N. Cha, and M. H. Bartl, *Discovery of a diamond-based photonic crystal structure in beetle scales,* Phys. Rev. E **77** (2008).

7. H. Ghiradella, *Light and color on the wing: structural colors in butterflies and moths,* Applied Optics **30**, 3492 (1991).

8. S. Vignolini, P. J. Rudall, A. V. Rowland, A. Reed, E. Moyroud, R. B. Faden, J. J. Baumberg, B. J. Glover, and U. Steiner, *Pointillist structural color in pollia fruit,* Proc. of the Natl. Acad. Sci. **109**, 15712 (2012).

9. S. J. Stachel, S. A. Stockwell, and D. L. V. Vranken, *The fluorescence of scorpions and cataractogenesis,* Chem. Biol. **6**, 531 (1999).

10. G. Marconi and K. Braun, *Nobel physics prize,* https://www.nobelprize.org/prizes/physics/1909/summary/ (1909).

11. C. Kao, *Nobel physics prize,* https://www.nobelprize.org/prizes/physics/2009/summary/ (2009).

12. M. B. Jr., *Semiconductor nanocrystals as fluorescent biological labels,* Science **281**, 2013 (1998).

13. W. C. Chan, D. J. Maxwell, X. Gao, R. E. Bailey, M. Han, and S. Nie, *Luminescent quantum dots for multiplexed biological detection and imaging,* Curr. Opin. in Biotech. **13**, 40 (2002).

14. J. M. Costa-Fernández, R. Pereiro, and A. Sanz-Medel, *The use of luminescent quantum dots for optical sensing,* TrAC Trends in Analytical Chemistry **25**, 207 (2006).

15. C. Monroe, *Quantum information processing with atoms and photons,* Nature **416**, 238 (2002).

16. G. Wendin, *Quantum information processing with superconducting circuits: a review,* Rep. Prog. Phys. **80**, 106001 (2017).

17. F. Flamini, N. Spagnolo, and F. Sciarrino, *Photonic quantum information processing: a review,* Rep. Prog. Phys. **82**, 016001 (2018).

18. T. Ha and P. Tinnefeld, *Photophysics of fluorescent probes for single-molecule biophysics and super-resolution imaging*, Annu. Rev. Phys. Chem. **63**, 595 (2012).

19. H. Miller, Z. Zhou, J. Shepherd, A. J. M. Wollman, and M. C. Leake, *Single-molecule techniques in biophysics: a review of the progress in methods and applications*, Rep. Prog. Phys. **81**, 024601 (2017).

20. C. H. Bennett and S. J. Wiesner, *Communication via one- and two-particle operators on einstein-podolsky-rosen states*, Phys. Rev. Lett. **69**, 2881 (1992).

21. R. Raussendorf and H. J. Briegel, *A one-way quantum computer*, Phys. Rev. Lett. **86**, 5188 (2001).

22. R. P. Feynman, *Simulating physics with computers*, Int. J. Theor. Phys. **21**, 467 (1982).

23. C. H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W. K. Wootters, *Teleporting an unknown quantum state via dual classical and einstein-podolsky-rosen channels*, Phys. Rev. Lett. **70**, 1895 (1993).

24. N. A. Gershenfeld and I. L. Chuang, *Bulk spin-resonance quantum computation*, Science **275**, 350 (1997).

25. S. Muralidharan and P. K. Panigrahi, *Perfect teleportation, quantum-state sharing, and superdense coding through a genuinely entangled five-qubit state*, Phys. Rev. A **77**, 032321 (2008).

26. H. J. Kimble, *The quantum internet*, Nature **453**, 1023 (2008).

27. B. Lounis and M. Orrit, *Single-photon sources*, Rep. Prog. Phys. **68**, 1129 (2005).

28. P. Lodahl, S. Mahmoodian, and S. Stobbe, *Interfacing single photons and single quantum dots with photonic nanostructures*, Rev. Mod. Phys. **87**, 347 (2015).

29. J. R. Mejía-Salazar and O. N. Oliveira, *Plasmonic biosensing*, Chem. Rev. **118**, 10617 (2018).

30. N. Toropov, G. Cabello, M. P. Serrano, R. R. Gutha, M. Rafti, and F. Vollmer, *Review of biosensing with whispering-gallery mode lasers*, Light-Sci. Appl. **10** (2021).

31. R. Pandey, S. K. Paidi, T. A. Valdez, C. Zhang, N. Spegazzini, R. R. Dasari, and I. Barman, *Noninvasive monitoring of blood glucose with raman spectroscopy*, Acc. Chem. Res. **50**, 264 (2017).

32. B. Sharma, R. R. Frontiera, A.-I. Henry, E. Ringe, and R. P. V. Duyne, *SERS: Materials, applications, and the future*, Mater. Today **15**, 16 (2012).

33. T. Itoh, Y. S. Yamamoto, and Y. Ozaki, *Plasmon-enhanced spectroscopy of absorption and spontaneous emissions explained using cavity quantum optics*, Chem. Soc. Rev. **46**, 3904 (2017).

34. E. C. Garnett, B. Ehrler, A. Polman, and E. Alarcon-Llado, *Photonics for photovoltaics: Advances and opportunities*, ACS Photonics **8**, 61 (2020).

35. G. J. Supran, Y. Shirasaki, K. W. Song, J.-M. Caruge, P. T. Kazlas, S. Coe-Sullivan, T. L. Andrew, M. G. Bawendi, and V. Bulović, *QLEDs for displays and solid-state lighting*, MRS Bulletin **38**, 703 (2013).

36. G. Lozano, D. J. Louwers, S. R. Rodríguez, S. Murai, O. T. Jansen, M. A. Verschuuren, and J. G. Rivas, *Plasmonics for solid-state lighting: enhanced excitation and directional emission of highly efficient light sources*, Light-Sci. Appl. **2**, e66 (2013).

37. R. Loudon, *The Quantum Theory of Light* (OUP Oxford, 2000).

38. S. Haroche and J.-M. Raimond, *Exploring the Quantum* (Oxford University Press, 2006).

39. C. Santori, D. Fattal, and Y. Yamamoto, *Single-photon devices and applications* (John Wiley & Sons, Weinheim, 2010).

40. C. W. Hsu, B. Zhen, A. Stone, J. Joannopoulos, and M. Soljačić, *Bound states in the continuum,* Nat. Rev. Mat. **1**, 16048 (2016).

41. L. Novotny and B. Hecht, *Principles of nano-optics*, 2nd ed. (Cambridge University Press, New York, 2012).

42. R. Sprik, B. A. v. Tiggelen, and A. Lagendijk, *Optical emission in periodic dielectrics,* Europhys. Lett. **35**, 265 (1996).

43. K. H. Drexhage, H. Kuhn, and F. P. Schäfer, *Variation of the fluorescence decay time of a molecule in front of a mirror,* Ber. Bunsen. Phys. Chem. **72** (1968).

44. K. Drexhage, *Influence of a dielectric interface on fluorescence decay time,* J. Lumin. **1-2** (1970).

45. E. M. Purcell, *Spontaneous emission probabilities at radio frequencies,* Phys. Rev. **69**, 681 (1946).

46. Q. Wang, S. Stobbe, and P. Lodahl, *Mapping the local density of optical states of a photonic crystal with single quantum dots,* Phys. Rev. Lett. **107** (2011).

47. J. Vučković, D. Fattal, C. Santori, G. S. Solomon, and Y. Yamamoto, *Enhanced single-photon emission from a quantum dot in a micropost microcavity,* Appl. Phys. Lett. **82**, 3596 (2003).

48. D. Englund, D. Fattal, E. Waks, G. Solomon, B. Zhang, T. Nakaoka, Y. Arakawa, Y. Yamamoto, and J. Vučković, *Controlling the spontaneous emission rate of single quantum dots in a two-dimensional photonic crystal,* Phys. Rev. Lett. **95** (2005), 10.1103/physrevlett.95.013904.

49. A. Kress, F. Hofbauer, N. Reinelt, M. Kaniber, H. J. Krenner, R. Meyer, G. Böhm, and J. J. Finley, *Manipulation of the spontaneous emission dynamics of quantum dots in two-dimensional photonic crystals,* Phys. Rev. B **71** (2005).

50. D. Najer, I. Söllner, P. Sekatski, V. Dolique, M. C. Löbl, D. Riedel, R. Schott, S. Starosielec, S. R. Valentin, A. D. Wieck, N. Sangouard, A. Ludwig, and R. J. Warburton, *A gated quantum dot strongly coupled to an optical microcavity,* Nature **575**, 622 (2019).

51. R. Albrecht, A. Bommer, C. Deutsch, J. Reichel, and C. Becher, *Coupling of a single nitrogen-vacancy center in diamond to a fiber-based microcavity,* Phys. Rev. Lett. **110** (2013).

52. A. Faraon, P. E. Barclay, C. Santori, K.-M. C. Fu, and R. G. Beausoleil, *Resonant enhancement of the zero-phonon emission from a colour centre in a diamond cavity,* Nat. Photonics **5**, 301 (2011).

53. T. B. Hoang, G. M. Akselrod, and M. H. Mikkelsen, *Ultrafast room-temperature single photon emission from quantum dots coupled to plasmonic nanocavities,* Nano Lett. **16**, 270 (2015).

54. F. Marquier, C. Sauvan, and J.-J. Greffet, *Revisiting quantum optics with surface plasmons and plasmonic resonators,* ACS Photonics **4** (2017).

55. A. Auffèves, D. Gerace, J.-M. Gérard, M. F. Santos, L. C. Andreani, and J.-P. Poizat, *Controlling the dynamics of a coupled atom-cavity system by pure dephasing,* Phys. Rev. B **81**, 245419 (2010).

56. J. M. Raimond, M. Brune, and S. Haroche, *Manipulating quantum entanglement with atoms and photons in a cavity,* Rev. Mod. Phys. **73**, 565 (2001).

57. G. Rempe, H. Walther, and N. Klein, *Observation of quantum collapse and revival in a one-atom maser,* Phys. Rev. Lett. **58**, 353 (1987).

58. R. J. Thompson, G. Rempe, and H. J. Kimble, *Observation of normal-mode splitting for an atom in an optical cavity,* Phys. Rev. Lett. **68**, 1132 (1992).

59. S. Haroche and D. Kleppner, *Cavity quantum electrodynamics,* Phys. Today **42**, 24 (1989).

60. R. Chikkaraddy, B. de Nijs, F. Benz, S. J. Barrow, O. A. Scherman, E. Rosta, A. Demetriadou, P. Fox, O. Hess, and J. J. Baumberg, *Single-molecule strong coupling at room temperature in plasmonic nanocavities,* Nature **535**, 127 (2016).

61. E. Peter, P. Senellart, D. Martrou, A. Lemaître, J. Hours, J.-M. Gérard, and J. Bloch, *Exciton-photon strong-coupling regime for a single quantum dot embedded in a microcavity,* Phys. Rev. Lett. **95**, 067401 (2005).

62. K. Srinivasan and O. Painter, *Linear and nonlinear optical spectroscopy of a strongly coupled microdisk-quantum dot system,* Nature **450**, 862 (2007).

63. T. Yoshie, A. Scherer, J. Hendrickson, G. Khitrova, H. M. Gibbs, G. Rupper, C. Ell, O. B. Shchekin, and D. G. Deppe, *Vacuum Rabi splitting with a single quantum dot in a photonic crystal nanocavity,* Nature **432**, 200 (2004).

64. A. Reinhard, T. Volz, M. Winger, A. Badolato, K. J. Hennessy, E. L. Hu, and A. Imamoğlu, *Strongly correlated photons on a chip,* Nat. Photonics **6**, 93 (2012).

65. J. P. Reithmaier, G. Sęk, A. Löffler, C. Hofmann, S. Kuhn, S. Reitzenstein, L. V. Keldysh, V. D. Kulakovskii, T. L. Reinecke, and A. Forchel, *Strong coupling in a single quantum dot-semiconductor microcavity system,* Nature **432**, 197 (2004).

66. K. J. Vahala, *Optical microcavities,* Nature **424**, 839 (2003).

67. T. Asano, Y. Ochi, Y. Takahashi, K. Kishimoto, and S. Noda, *Photonic crystal nanocavity with a q factor exceeding eleven million,* Opt. Express **25**, 1769 (2017).

68. M. L. Gorodetsky, A. A. Savchenkov, and V. S. Ilchenko, *Ultimate Q of optical microsphere resonators,* Opt. Lett. **21**, 453 (1996).

69. F. Vollmer and S. Arnold, *Whispering-gallery-mode biosensing: Label-free detection down to single molecules,* Nat. Methods **5**, 591 (2008).

70. J. B. Khurgin, *How to deal with the loss in plasmonics and metamaterials,* Nat. Nanotechnol. **10**, 2 (2015).

71. S. Hu, M. Khater, R. Salas-Montiel, E. Kratschmer, S. Engelmann, W. M. J. Green, and S. M. Weiss, *Experimental realization of deep-subwavelength confinement in dielectric optical resonators,* Science Adv. **4**, eaat2355 (2018).

72. J. D. Ryckman and S. M. Weiss, *Low mode volume slotted photonic crystal single nanobeam cavity,* Appl. Phys. Lett. **101**, 071104 (2012).

73. P. Seidler, K. Lister, U. Drechsler, J. Hofrichter, and T. Stöferle, *Slotted photonic crystal nanobeam cavity with an ultrahigh quality factor-to-mode volume ratio,* Opt. Express **21**, 32468 (2013).

74. S. Reitzenstein and A. Forchel, *Quantum dot micropillars,* J. Phys. D - Appl. Phys. **43**, 033001 (2010).

75. Z.-S. Chen, B. Ma, X.-J. Shang, H.-Q. Ni, J.-L. Wang, and Z.-C. Niu, *Bright single-photon source at 1.3 μm based on InAs bilayer quantum dot in micropillar,* Nanoscale Res. Lett. **12** (2017).

76. H. P. Springbett, J. Jarman, T. Zhu, M. Holmes, Y. Arakawa, and R. A. Oliver, *Improvement of single photon emission from InGaN QDs embedded in porous micropillars,* Appl. Phys. Lett. **113**, 101107 (2018).

77. S. Noda, M. Fujita, and T. Asano, *Spontaneous-emission control by photonic crystals and nanocavities,* Nat. Photonics **1**, 449 (2007).

78. K. Hennessy, A. Badolato, M. Winger, D. Gerace, M. Atatüre, S. Gulde, S. Fält, E. L. Hu, and A. Imamoğlu, *Quantum nature of a strongly coupled single quantum dot–cavity system,* Nature **445**, 896 (2007).

79. A. Faraon, I. Fushman, D. Englund, N. Stoltz, P. Petroff, and J. Vučković, *Coherent generation of non-classical light on a chip via photon-induced tunnelling and blockade,* Nat. Phys. **4**, 859 (2008).

80. A. Sipahigil, R. E. Evans, D. D. Sukachev, M. J. Burek, J. Borregaard, M. K. Bhaskar, C. T. Nguyen, J. L. Pacheco, H. A. Atikian, C. Meuwly, R. M. Camacho, F. Jelezko, E. Bielejec, H. Park, M. Lončar, and M. D. Lukin, *An integrated diamond nanophotonics platform for quantum-optical networks,* Science **354**, 847 (2016).

81. K. Kuruma, Y. Ota, M. Kakuda, S. Iwamoto, and Y. Arakawa, *Surface-passivated high-q GaAs photonic crystal nanocavity with quantum dots,* APL Photonics **5**, 046106 (2020).

82. M. Barth, S. Schietinger, S. Fischer, J. Becker, N. Nüsse, T. Aichele, B. Löchel, C. Sönnichsen, and O. Benson, *Nanoassembled plasmonic-photonic hybrid cavity for tailored light-matter coupling,* Nano Lett. **10**, 891 (2010).

83. A. G. Curto, G. Volpe, T. H. Taminiau, M. P. Kreuzer, R. Quidant, and N. F. van Hulst, *Unidirectional emission of a quantum dot coupled to a nanoantenna,* Science **329**, 930 (2010).

84. E. B. Ureña, M. P. Kreuzer, S. Itzhakov, H. Rigneault, R. Quidant, D. Oron, and J. Wenger, *Excitation enhancement of a quantum dot coupled to a plasmonic antenna,* Adv. Mater. **24**, OP314 (2012).

85. H. M. Doeleman, C. D. Dieleman, C. Mennes, B. Ehrler, and A. F. Koenderink, *Observation of cooperative purcell enhancements in antenna–cavity hybrids,* ACS Nano **14**, 12027 (2020).

86. C. A. Balanis, *Antenna Theory - Analysis and Design*, 3rd ed. (John Wiley & Sons, Inc., Hoboken, New Jersey., 2005).

87. L. Novotny and N. van Hulst, *Antennas for light*, Nat. Photonics **5**, 83 (2011).

88. F. Benz, M. K. Schmidt, A. Dreismann, R. Chikkaraddy, Y. Zhang, A. Demetriadou, C. Carnegie, H. Ohadi, B. de Nijs, R. Esteban, J. Aizpurua, and J. J. Baumberg, *Single-molecule optomechanics in "picocavities"*, Science **354**, 726 (2016).

89. H. Groß, J. M. Hamm, T. Tufarelli, O. Hess, and B. Hecht, *Near-field strong coupling of single quantum dots*, Sci. Adv. **4**, 4906 (2018).

90. A. F. Koenderink, *On the use of Purcell factors for plasmon antennas.* Opt. Lett. **35**, 4208 (2010).

91. P. T. Kristensen, C. Van Vlack, and S. Hughes, *Generalized effective mode volume for leaky optical cavities.* Opt. Lett. **37**, 1649 (2012).

92. C. Sauvan, J. P. Hugonin, I. S. Maksymov, and P. Lalanne, *Theory of the spontaneous optical emission of nanosize photonic and plasmon resonators*, Phys. Rev. Lett. **110**, 237401 (2013).

93. J. Yang, C. Sauvan, A. Jouanin, S. Collin, J.-L. Pelouard, and P. Lalanne, *Ultrasmall metal-insulator-metal nanoresonators: impact of slow-wave effects on the quality factor*, Opt. Express **20**, 16880 (2012).

94. T. Coenen, E. J. R. Vesseur, A. Polman, and A. F. Koenderink, *Directional emission from plasmonic Yagi–Uda antennas probed by angle-resolved cathodoluminescence spectroscopy*, Nano Lett. **11**, 3779 (2011).

95. L. Novotny, *Effective wavelength scaling for optical antennas*, Phys. Rev. Lett. **98** (2007).

96. G. M. Akselrod, C. Argyropoulos, T. B. Hoang, C. Ciracì, C. Fang, J. Huang, D. R. Smith, and M. H. Mikkelsen, *Probing the mechanisms of large Purcell enhancement in plasmonic nanoantennas*, Nat. Photonics **8**, 835 (2014).

97. C. Belacel, B. Habert, F. Bigourdan, F. Marquier, J.-P. Hugonin, S. M. de Vasconcellos, X. Lafosse, L. Coolen, C. Schwob, C. Javaux, B. Dubertret, J.-J. Greffet, P. Senellart, and A. Maitre, *Controlling spontaneous emission with plasmonic optical patch antennas*, Nano Lett. **13**, 1516 (2013).

98. W. Kim, V. P. Safonov, V. M. Shalaev, and R. L. Armstrong, *Fractals in microcavities: Giant coupled, multiplicative enhancement of optical responses*, Phys. Rev. Lett. **82** (1999).

99. M. R. Gartia, S. Seo, J. Kim, T.-W. Chang, G. Bahl, M. Lu, G. L. Liu, and J. G. Eden, *Injection- seeded optoplasmonic amplifier in the visible*, Sci. Rep. **4** (2014).

100. J.-N. Liu, Q. Huang, K.-K. Liu, S. Singamaneni, and B. T. Cunningham, *Nanoantenna-microcavity hybrids with highly cooperative plasmonic-photonic coupling*, Nano Lett. **17**, 7569 (2017).

101. J. Losada, A. Raza, S. Clemmen, A. Serrano, A. Griol, R. Baets, and A. Martinez, *SERS detection via individual bowtie nanoantennas integrated in Si3N4 waveguides*, IEEE J. Sel. Top. Quant. **25** (2019).

102. S. I. Shopova, C. W. Blackledge, and A. T. Rosenberger, *Enhanced evanescent coupling to whispering-gallery modes due to gold nanorods grown on the microresonator surface*, Appl. Phys. B **93**, 183 (2008).

103. M. A. Santiago-Cordoba, M. Cetinkaya, S. V. Boriskina, F. Vollmer, and M. C. Demirel, *Ultrasensitive detection of a protein by optical trapping in a photonic-plasmonic microcavity,* J. Biophotonics **5** (2012).

104. V. R. Dantham, S. Holler, C. Barbre, D. Keng, V. Kolchenko, and S. Arnold, *Label-free detection of single protein using a nanoplasmonic-photonic hybrid microcavity,* Nano Lett. **13** (2013).

105. A. Bozzola, S. Perotto, and F. D. Angelis, *Hybrid plasmonic–photonic whispering gallery mode resonators for sensing: a critical review,* The Analyst **142** (2017).

106. M. D. Baaske and F. Vollmer, *Optical observation of single atomic ions interacting with plasmonic nanorods in aqueous solution,* Nat. Photonics **10**, 733 (2016).

107. F. Pan, K. C. Smith, H. L. Nguyen, K. A. Knapper, D. J. Masiello, and R. H. Goldsmith, *Elucidating energy pathways through simultaneous measurement of absorption and transmission in a coupled plasmonic–photonic cavity,* Nano Lett. **20**, 50 (2019).

108. N. Thakkar, M. T. Rea, K. C. Smith, K. D. Heylman, S. C. Quillin, K. A. Knapper, E. H. Horak, D. J. Masiello, and R. H. Goldsmith, *Sculpting fano resonances to control photonic-plasmonic hybridization,* Nano Lett. **17**, 6927 (2017).

109. F. De Angelis, M. Patrini, G. Das, I. Maksymov, M. Galli, L. Businaro, L. C. Andreani, and E. Di Fabrizio, *A hybrid plasmonic-photonic nanodevice for label-free detection of a few molecules,* Nano Lett. **8**, 2321 (2008).

110. F. Liang, Y. Guo, S. Hou, and Q. Quan, *Photonic-plasmonic hybrid single-molecule nanosensor measures the effect of fluorescent labels on DNA-protein dynamics,* Science Adv. **3** (2017).

111. Y.-W. Hu, B.-B. Li, Y.-X. Liu, Y.-F. Xiao, and Q. Gong, *Hybrid photonic–plasmonic mode for refractometer and nanoparticle trapping,* Opt. Commun. **291** (2013).

112. P.-T. Lin, H.-Y. Chu, T.-W. Lu, and P.-T. Lee, *Correction: Trapping particles using waveguide-coupled gold bowtie plasmonic tweezers,* Lab on a Chip **15** (2015).

113. D. Conteduca, F. Dell'Olio, F. Innone, C. Ciminelli, and M. Armenise, *Rigorous design of an ultra-high Q/V photonic/plasmonic cavity to be used in biosensing applications,* Opt. Laser Technol. **77** (2016).

114. D. Conteduca, C. Reardon, M. G. Scullion, F. Dell'Olio, M. N. Armenise, T. F. Krauss, and C. Ciminelli, *Ultra-high q/v hybrid cavity for strong light-matter interaction,* APL Photonics **2** (2017).

115. T. Zhang, S. Callard, C. Jamois, C. Chevalier, D. Feng, and A. Belarouci, *Plasmonic-photonic crystal coupled nanolaser,* Nanotechnol. **25**, 315201 (2014).

116. M. Mivelle, P. Viktorovitch, F. I. Baida, A. E. Eter, Z. Xie, T.-P. Vo, E. Atie, G. W. Burr, D. Nedeljkovic, J.-Y. Rauch, S. Callard, and T. Grosjean, *Light funneling from a photonic crystal laser cavity to a nano-antenna: overcoming the diffraction limit in optical energy transfer down to the nanoscale,* Opt. Express **22** (2014).

117. M. Chamanzar and A. Adibi, *Hybrid nanoplasmonic-photonic resonators for efficient coupling of light to single plasmonic nanoresonators,* Opt. Express **19** (2011).

118. W. Ahn, S. V. Boriskina, Y. Hong, and B. M. Reinhard, *Photonic–plasmonic mode coupling in on-chip integrated optoplasmonic molecules,* ACS Nano **6** (2011).

119. F. Ren, H. Takashima, Y. Tanaka, H. Fujiwara, and K. Sasaki, *Two-photon excited fluorescence from a pseudoisocyanine-attached gold tip via a plasmonic-photonic hybrid system,* Opt. Express **23**, 21730 (2015).

120. W. Ahn, X. Zhao, Y. Hong, and B. M. Reinhard, *Low-power light guiding and localization in optoplasmonic chains obtained by directed self-assembly,* Sci. Rep. **6** (2016).

121. B. Chen, R. Bruck, D. Traviss, A. Z. Khokhar, S. Reynolds, D. J. Thomson, G. Z. Mashanovich, G. T. Reed, and O. L. Muskens, *Hybrid photon–plasmon coupling and ultrafast control of nanoantennas on a silicon photonic chip,* Nano Lett. **18** (2017).

122. H. M. Doeleman, E. Verhagen, and A. F. Koenderink, *Antenna-Cavity Hybrids: Matching Polar Opposites for Purcell Enhancements at Any Linewidth,* ACS Photonics **3**, 1943 (2016).

123. S. Cui, X. Zhang, T. li Liu, J. Lee, D. Bracher, K. Ohno, D. Awschalom, and E. L. Hu, *Hybrid plasmonic photonic crystal cavity for enhancing emission from near-surface nitrogen vacancy centers in diamond,* ACS Photonics **2**, 465 (2015).

124. A. Jablonski, *Efficiency of anti-stokes fluorescence in dyes,* Nature **131**, 839–840 (1933).

125. J. R. Lakowicz, *Principles of Fluorescence Spectroscopy* (Springer, 2011).

126. E. Nir, M. Jäger, and S. Weiss, *Chapter 2 - alternating laser excitation spectroscopy of freely diffusing single molecule: Applications to bimolecular structure, dynamics and interactions,* in *Nano Biophotonics,* Handai Nanophotonics, Vol. 3, edited by H. Masuhara, S. Kawata, and F. Tokunaga (Elsevier, 2007) pp. 23–42.

127. C. Toninelli, I. Gerhardt, A. Clark, A. Reserbat-Plantey, S. Götzinger, Z. Ristanovic, M. Colautti, P. Lombardi, K. Major, I. Deperasińska, W. Pernice, F. Koppens, B. Kozankiewicz, A. Gourdon, V. Sandoghdar, and M. Orrit, *Single organic molecules for photonic quantum technologies,* preprint , arXiv:2011.05059 (2020).

128. L. Allen and J. Eberly, *Optical Resonance and Two-Level Atoms,* Dover Books on Physics (Dover Publications, 2012).

129. B. Korzh, Q.-Y. Zhao, J. P. Allmaras, S. Frasca, T. M. Autry, E. A. Bersin, A. D. Beyer, R. M. Briggs, B. Bumble, M. Colangelo, G. M. Crouch, A. E. Dane, T. Gerrits, A. E. Lita, F. Marsili, G. Moody, C. Peña, E. Ramirez, J. D. Rezac, N. Sinclair, M. J. Stevens, A. E. Velasco, V. B. Verma, E. E. Wollman, S. Xie, D. Zhu, P. D. Hale, M. Spiropulu, K. L. Silverman, R. P. Mirin, S. W. Nam, A. G. Kozorezov, M. D. Shaw, and K. K. Berggren, *Demonstration of sub-3 ps temporal resolution with a superconducting nanowire single-photon detector,* Nat. Photonics **14**, 250 (2020).

130. H. Bernien, B. Hensen, W. Pfaff, G. Koolstra, M. S. Blok, L. Robledo, T. H. Taminiau, M. Markham, D. J. Twitchen, L. Childress, and R. Hanson, *Heralded entanglement between solid-state qubits separated by three metres,* Nature **497**, 86 (2013).

131. I. Aharonovich, A. D. Greentree, and S. Prawer, *Diamond photonics,* Nat. Photonics **5**, 397 (2011).

132. A. Gruber, A. Dräbenstedt, C. Tietz, L. Fleury, J. Wrachtrup, and C. von

Borczyskowski, *Scanning confocal optical microscopy and magnetic resonance on single defect centers,* Science **276**, 2012 (1997).

133. C. Kurtsiefer, S. Mayer, P. Zarda, and H. Weinfurter, *Stable solid-state source of single photons,* Phys. Rev. Lett. **85**, 290 (2000).

134. R. Brouri, A. Beveratos, J.-P. Poizat, and P. Grangier, *Photon antibunching in the fluorescence of individual color centers in diamond,* Opt. Express **25**, 1294 (2000).

135. J. Wolters, A. W. Schell, G. Kewes, N. Nüsse, M. Schoengen, H. Döscher, T. Hannappel, B. Löchel, M. Barth, and O. Benson, *Enhancement of the zero phonon line emission from a single nitrogen vacancy center in a nanodiamond via coupling to a photonic crystal cavity,* Appl. Phys. Lett. **97**, 141108 (2010).

136. D. Englund, B. Shields, K. Rivoire, F. Hatami, J. Vučković, H. Park, and M. D. Lukin, *Deterministic coupling of a single nitrogen vacancy center to a photonic crystal cavity,* Nano Lett. **10**, 3922 (2010).

137. B. J. M. Hausmann, B. J. Shields, Q. Quan, Y. Chu, N. P. de Leon, R. Evans, M. J. Burek, A. S. Zibrov, M. Markham, D. J. Twitchen, H. Park, M. D. Lukin, and M. Lončar, *Coupling of NV centers to photonic crystal nanobeams in diamond,* Nano Lett. **13**, 5791 (2013).

138. A. Faraon, C. Santori, Z. Huang, V. M. Acosta, and R. G. Beausoleil, *Coupling of nitrogen-vacancy centers to photonic crystal cavities in monocrystalline diamond,* Phys. Rev. Lett. **109** (2012).

139. P. E. Barclay, C. Santori, K.-M. Fu, R. G. Beausoleil, and O. Painter, *Coherent interference effects in a nano-assembled diamond NV center cavity-QED system,* Opt. Express **17**, 8081 (2009).

140. A. Mohtashami and A. F. Koenderink, *Suitability of nanodiamond nitrogen–vacancy centers for spontaneous emission control experiments,* New J. Phys. **15**, 043017 (2013).

141. P. Appel, E. Neu, M. Ganzhorn, A. Barfuss, M. Batzer, M. Gratz, A. Tschöpe, and P. Maletinsky, *Fabrication of all diamond scanning probes for nanoscale magnetometry,* Rev. of Sci. Instrum. **87**, 063703 (2016).

142. L. Fleury, B. Sick, G. Zumoffen, B. Hecht, and U. P. Wild, *High photo-stability of single molecules in an organic crystal at room temperature observed by scanning confocal optical microscopy,* Mol. Phys. **95**, 1333 (1998).

143. B. Lounis and W. E. Moerner, *Single photons on demand from a single molecule at room temperature,* Nature **407**, 491 (2000).

144. C. Toninelli, K. Early, J. Bremi, A. Renn, S. Götzinger, and V. Sandoghdar, *Near-infrared single-photons from aligned molecules in ultrathin crystalline films at room temperature,* Opt. Express **18**, 6577 (2010).

145. F. Jelezko, P. Tamarat, B. Lounis, and M. Orrit, *Dibenzoterrylene in naphthalene: a new crystalline system for single molecule spectroscopy in the near infrared,* J. Phys. Chem. **100**, 13892 (1996).

146. C. Hofmann, A. Nicolet, M. A. Kol'chenko, and M. Orrit, *Towards nanoprobes for conduction in molecular crystals: Dibenzoterrylene in anthracene crystals,* Chem. Phys. **318**, 1 (2005).

147. S. Pazzagli, P. Lombardi, D. Martella, M. Colautti, B. Tiribilli, F. S. Cataliotti, and C. Toninelli, *Self-assembled nanocrystals of polycyclic aromatic hydrocarbons show photostable single-photon emission,* ACS Nano **12**, 4295 (2018).

148. R. Osovsky, D. Cheskis, V. Kloper, A. Sashchiuk, M. Kroner, and E. Lifshitz, *Continuous-wave pumping of multiexciton bands in the photoluminescence spectrum of a single CdTe-CdSe core-shell colloidal quantum dot,* Phys. Rev. Lett. **102** (2009).

149. M. Kuno, D. P. Fromm, H. F. Hamann, A. Gallagher, and D. J. Nesbitt, *Nonexponential "blinking" kinetics of single CdSe quantum dots: A universal power law behavior,* J. Chem. Phys **112**, 3117 (2000).

150. K. Zhang, H. Chang, A. Fu, A. P. Alivisatos, and H. Yang, *Continuous Distribution of Emission States from Single CdSe/ZnS Quantum Dots,* Nano Lett. **6**, 843 (2006).

151. M. Y. Odoi, N. I. Hammer, K. T. Early, K. D. McCarthy, R. Tangirala, T. Emrick, and M. D. Barnes, *Fluorescence Lifetimes and Correlated Photon Statistics from Single CdSe/Oligo(phenylene vinylene) Composite Nanostructures,* Nano Lett. **7**, 2769 (2007).

152. R. Schmidt, C. Krasselt, and C. Von Borczyskowski, *Change point analysis of matrix dependent photoluminescence intermittency of single CdSe/ZnS quantum dots with intermediate intensity levels,* Chem. Phys. **406**, 9 (2012).

153. F. T. Rabouw, F. V. Antolinez, R. Brechbühler, and D. J. Norris, *Microsecond blinking events in the fluorescence of colloidal quantum dots revealed by correlation analysis on preselected photons,* J. Phys. Chem. Lett. **10**, 3732 (2019).

154. R. W. Schoenlein, D. M. Mittleman, J. J. Shiang, A. P. Alivisatos, and C. V. Shank, *Investigation of femtosecond electronic dephasing in CdSe nanocrystals using quantum-beat-suppressed photon echoes,* Phys. Rev. Lett. **70**, 1014 (1993).

155. F. Masia, N. Accanto, W. Langbein, and P. Borri, *Spin-flip limited exciton dephasing inCdSe/ZnSColloidal quantum dots,* Phys. Rev. Lett. **108** (2012).

156. L. Protesescu, S. Yakunin, M. I. Bodnarchuk, F. Krieg, R. Caputo, C. H. Hendon, R. X. Yang, A. Walsh, and M. V. Kovalenko, *Nanocrystals of Cesium Lead Halide Perovskites (CsPbX 3 , X = Cl, Br, and I): Novel Optoelectronic Materials Showing Bright Emission with Wide Color Gamut,* Nano Lett. **15**, 3692 (2015).

157. A. Swarnkar, R. Chulliyil, V. Ravi, M. Irfanullah, A. Chowdhury, and A. Nag, *Colloidal CsPbBr$_3$ perovskite nanocrystals: Luminescence beyond traditional quantum dots,* Angew. Chem. Int. Ed. (2015), 10.1002/anie.201508276.

158. H. Huang, M. I. Bodnarchuk, S. V. Kershaw, M. V. Kovalenko, and A. L. Rogach, *Lead halide perovskite nanocrystals in the research spotlight: Stability and defect tolerance,* ACS Energy Lett. **2**, 2071 (2017).

159. R. Chatterjee, I. M. Pavlovetc, K. Aleshire, and M. Kuno, *Single semiconductor nanostructure extinction spectroscopy,* J. Phys. Chem. C **122**, 16443 (2018).

160. F. Vollmer and L. Yang, *Label-free detection with high-Q microcavities: a review of biosensing mechanisms for integrated devices,* Nanophotonics **1**, 267 (2012).

161. M. Aspelmeyer, T. J. Kippenberg, and F. Marquardt, *Cavity optomechanics,* Rev. Mod. Phys. **86**, 1391 (2014).

162. I. Aharonovich, D. Englund, and M. Toth, *Solid-state single-photon emitters,* Nat. Photonics **10**, 631 (2016).

163. P. Senellart, G. Solomon, and A. White, *High-performance semiconductor quantum-dot single-photon sources,* Nat. Nanotechnol. **12**, 1026 (2017).

164. M. T. Hill and M. C. Gather, *Advances in small lasers,* Nat. Photonics **8**, 908 (2014).

165. D. E. Chang, V. Vuletić, and M. D. Lukin, *Quantum nonlinear optics - photon by photon,* Nat. Photonics **8**, 685 (2014).

166. K. J. Russell, T.-L. Liu, S. Cui, and E. L. Hu, *Large spontaneous emission enhancement in plasmonic nanocavities,* Nat. Photonics **6**, 459 (2012).

167. S. Bidault, A. Devilez, V. Maillard, L. Lermusiaux, J.-M. Guigner, N. Bonod, and J. Wenger, *Picosecond lifetimes with high quantum yields from single-photon-emitting colloidal nanostructures at room temperature,* ACS Nano **10**, 4806 (2016).

168. N. Somaschi, V. Giesz, L. De Santis, J. C. Loredo, M. P. Almeida, G. Hornecker, S. L. Portalupi, T. Grange, C. Antón, J. Demory, C. Gómez, I. Sagnes, N. D. Lanzillotti-Kimura, A. Lemaítre, A. Auffeves, A. G. White, L. Lanco, and P. Senellart, *Near-optimal single-photon sources in the solid state,* Nat. Photonics **10**, 340 (2016).

169. I. M. Palstra, H. M. Doeleman, and A. F. Koenderink, *Hybrid cavity-antenna systems for quantum optics outside the cryostat?* Nanophotonics **8**, 1513 (2019).

170. M. A. Santiago-Cordoba, M. Cetinkaya, S. V. Boriskina, F. Vollmer, and M. C. Demirel, *Ultrasensitive detection of a protein by optical trapping in a photonic-plasmonic microcavity,* J. Biophotonics **5**, 629 (2012).

171. M. Frimmer and A. F. Koenderink, *Superemitters in hybrid photonic systems: A simple lumping rule for the local density of optical states and its breakdown at the unitary limit,* Phys. Rev. B **86**, 235428 (2012).

172. Y.-F. Xiao, Y.-C. Liu, B.-B. Li, Y.-L. Chen, Y. Li, and Q. Gong, *Strongly enhanced light-matter interaction in a hybrid photonic-plasmonic resonator,* Phys. Rev. A **85**, 031805 (2012).

173. T. Zhang, A. Belarouci, S. Callard, C. Jamois, X. Letartre, C. Chevalier, P. Rojo-Romeo, B. Devif, and P. Viktorovitch, *Far-field and near-field investigation of plasmonic-photonic hybrid laser mode,* arXiv , 1411.7201:1 (2014).

174. Y. Hong, W. Ahn, S. V. Boriskina, X. Zhao, and B. M. Reinhard, *Directed assembly of optoplasmonic hybrid materials with tunable photonic-plasmonic properties,* J. Phys. Chem. Lett. **6**, 2056 (2015).

175. M. K. Dezfouli, R. Gordon, and S. Hughes, *Modal theory of modified spontaneous emission of a quantum emitter in a hybrid plasmonic photonic-crystal cavity system,* Phys. Rev. A **95**, 013846 (2017).

176. B. Gurlek, V. Sandoghdar, and D. Martín-Cano, *Manipulation of quenching in nanoantenna-emitter systems enabled by external detuned cavities: A path to enhance strong-coupling,* ACS Photonics **5**, 456 (2018).

177. P. B. Deotare, M. W. McCutcheon, I. W. Frank, M. Khan, and M. Lončar, *High quality factor photonic crystal nanobeam cavities,* Appl. Phys. Lett. **94**, 121106 (2009).

178. M. K. Dezfouli, R. Gordon, and S. Hughes, *Molecular optomechanics in the anharmonic cavity-QED regime using hybrid metal–dielectric cavity modes*, ACS Photonics **6**, 1400 (2019).

179. M. Grande, G. Calo, V. Petruzzelli, and A. D'Orazio, *High-Q photonic crystal nanobeam cavity based on a silicon nitride membrane incorporating fabrication imperfections and a low-index material layer*, Prog. Electromagn. Res. B **37**, 191 (2012).

180. J. D. Ryckman and S. M. Weiss, *Low mode volume slotted photonic crystal single nanobeam cavity in silicon*, in *The 9th International Conference on Group IV Photonics (GFP)*, Vol. 071104 (IEEE, 2012) pp. 24–26.

181. E. N. Economou, *Surface Plasmons in Thin Films*, Phys. Rev. **182**, 539 (1969).

182. S. I. Bozhevolnyi and T. Søndergaard, *General properties of slow-plasmon resonant nanostructures: nano-antennas and resonators*, Opt. Express **15**, 10869 (2007).

183. R. Faggiani, J. Yang, and P. Lalanne, *Quenching, plasmonic, and radiative decays in nanogap emitting devices*, ACS Photonics **2**, 1739 (2015).

184. U. Fano, *Effects of configuration interaction on intensities and phase shifts*, Phys. Rev. **124**, 1866 (1961).

185. H. Doeleman, *Hybrid resonators for light trapping and emission control*, Ph.D. thesis, Universiteit van Amsterdam, Amsterdam, The Netherlands (2019).

186. F. Ruesink, H. M. Doeleman, R. Hendrikx, A. F. Koenderink, and E. Verhagen, *Perturbing open cavities: Anomalous resonance frequency shifts in a hybrid cavity-nanoantenna system*, Phys. Rev. Lett. **115**, 203904 (2015).

187. P. De Vries, D. van Coevorden, and L. A., *Point scatterers for classical waves*, Rev. Mod. Phys. **70** (1998).

188. C. F. Bohren and D. R. Huffman, *Absorption and scattering of light by small particles* (John Wiley & Sons, New York, 1983).

189. R. Waldron, *Perturbation theory of resonant cavities*, Proc. IEE Part C: Monographs **107**, 272 (1960).

190. H. Bethe and J. Schwinger, *Perturbation theory for cavities* (Massachusetts Institute of Technology, Radiation Laboratory, Cambridge, MA, 1943).

191. S. Haroche, *Nobel lecture: Controlling photons in a box and exploring the quantum to classical boundary*, Rev. Mod. Phys. **85**, 1083 (2013).

192. A. Kwadrin and A. F. Koenderink, *Gray-tone lithography implementation of Drexhage's method for calibrating radiative and nonradiative decay constants of fluorophores*, J. Phys. Chem. C **116**, 16666 (2012).

193. A. P. Green and A. R. Buckley, *Solid state concentration quenching of organic fluorophores in PMMA*, Phys. Chem. Chem. Phys. **17**, 1435 (2015).

194. K. Guo, G. Lozano, M. A. Verschuuren, and J. Gómez Rivas, *Control of the external photoluminescent quantum yield of emitters coupled to nanoantenna phased arrays*, J. Appl. Phys. **118**, 073103 (2015).

195. J. Olmsted, *Calorimetric determinations of absolute fluorescence quantum yields*, J. Phys. Chem. **83**, 2581 (1979).

196. M. D. Leistikow, J. Johansen, A. J. Kettelarij, P. Lodahl, and W. L. Vos, *Size-dependent oscillator strength and quantum efficiency of CdSe quantum dots controlled via the local density of states*, Phys. Rev. B **79**, 045301 (2009).

197. P. Lunnemann, F. T. Rabouw, R. J. A. van Dijk-Moes, F. Pietra, D. Vanmaekelbergh, and A. F. Koenderink, *Calibrating and Controlling the Quantum Efficiency Distribution of Inhomogeneously Broadened Quantum Rods by Using a Mirror Ball*, ACS Nano **7**, 5984 (2013).

198. K. Sebald, P. Michler, T. Passow, D. Hommel, G. Bacher, and A. Forchel, *Single-photon emission of CdSe quantum dots at temperatures up to 200 K*, Appl. Phys. Lett. **81**, 2920 (2002).

199. K. Choi, S. Kako, M. J. Holmes, M. Arita, and Y. Arakawa, *Strong exciton confinement in site-controlled GaN quantum dots embedded in nanowires*, Appl. Phys. Lett. **103**, 171907 (2013).

200. W. Yang, J. Li, Y. Zhang, P.-K. Huang, T.-C. Lu, H.-C. Kuo, S. Li, X. Yang, H. Chen, D. Liu, and J. Kang, *High density GaN/AlN quantum dots for deep UV LED with high quantum efficiency and temperature stability*, Sci. Rep. **4**, 5166 (2015).

201. M. Bayer and A. Forchel, *Temperature dependence of the exciton homogeneous linewidth in InGaAs/GaAs self-assembled quantum dots*, Phys. Rev. B **65**, 041308 (2002).

202. E. Neu, C. Hepp, M. Hauschild, S. Gsell, M. Fischer, H. Sternschulte, D. Steinmüller-Nethl, M. Schreck, and C. Becher, *Low-temperature investigations of single silicon vacancy colour centres in diamond*, New J. Phys. **15**, 043005 (2013).

203. F. B. Zhelezko, I. M. Gulis, B. Lounis, and M. Orrit, *Spectroscopic characteristics of single dibenzanthanthrene molecules isolated in a low-temperature naphthalene matrix*, J. Appl. Spectrosc. **66**, 344 (1999).

204. A. Boiron, B. Lounis, and M. Orrit, *Single molecules of dibenzanthanthrene in n-hexadecane*, J. Chem. Phys. **105**, 3969 (1996).

205. T. W. Ebbesen, *Hybrid Light–Matter States in a Molecular and Material Science Perspective*, Acc. Chem. Res. **49**, 2403 (2016).

206. J. Kasprzak, M. Richard, S. Kundermann, A. Baas, P. Jeambrun, J. M. J. Keeling, F. M. Marchetti, M. H. Szymańska, R. André, J. L. Staehli, V. Savona, P. B. Littlewood, B. Deveaud, and L. S. Dang, *Bose–einstein condensation of exciton polaritons*, Nature **443**, 409 (2006).

207. M. Ramezani, Q. Le-Van, A. Halpin, and J. Gómez Rivas, *Nonlinear emission of molecular ensembles strongly coupled to plasmonic lattices with structural imperfections*, Phys. Rev. Lett. **121**, 243904 (2018).

208. A. Shalabney, J. George, J. Hutchison, G. Pupillo, C. Genet, and T. W. Ebbesen, *Coherent coupling of molecular resonators with a microcavity mode*, Nat. Commun. **6**, 5981 (2015).

209. R.-Q. Li, D. Hernángomez-Pérez, F. J. García-Vidal, and A. I. Fernández-Domínguez, *Transformation Optics Approach to Plasmon-Exciton Strong Coupling in Nanocavities*, Phys. Rev. Lett. **117**, 107401 (2016).

210. J. del Pino, F. J. Garcia-Vidal, and J. Feist, *Exploiting vibrational strong coupling to make an optical parametric oscillator out of a raman laser,* Phys. Rev. Lett. **117**, 277401 (2016).

211. D. Sanvitto and S. Kéna-Cohen, *The road towards polaritonic devices,* Nat. Mater. **15**, 1061 (2016).

212. P. Törmä and W. L. Barnes, *Strong coupling between surface plasmon polaritons and emitters: a review,* Rep. Prog. Phys. **78**, 013901 (2015).

213. T. Byrnes, N. Y. Kim, and Y. Yamamoto, *Exciton–polariton condensates,* Nat. Phys. **10**, 803 (2014).

214. G. Lerario, A. Fieramosca, F. Barachati, D. Ballarini, K. S. Daskalakis, L. Dominici, M. De Giorgi, S. A. Maier, G. Gigli, S. Kéna-Cohen, and D. Sanvitto, *Room-temperature superfluidity in a polariton condensate,* Nat. Phys. **13**, 837 (2017).

215. A. Amo, J. Lefrère, S. Pigeon, C. Adrados, C. Ciuti, I. Carusotto, R. Houdré, E. Giacobino, and A. Bramati, *Superfluidity of polaritons in semiconductor microcavities,* Nat. Phys. **5**, 805 (2009).

216. J. Fischer, I. G. Savenko, M. D. Fraser, S. Holzinger, S. Brodbeck, M. Kamp, I. A. Shelykh, C. Schneider, and S. Höfling, *Spatial Coherence Properties of One Dimensional Exciton-Polariton Condensates,* Phys. Rev. Lett. **113**, 203902 (2014).

217. R. Su, J. Wang, J. Zhao, J. Xing, W. Zhao, C. Diederichs, T. C. H. Liew, and Q. Xiong, *Room temperature long-range coherent exciton polariton condensate flow in lead halide perovskites,* Sci. Adv. **4**, eaau0244 (2018).

218. D. R. Gulevich, D. Yudin, D. V. Skryabin, I. V. Iorsh, and I. A. Shelykh, *Exploring nonlinear topological states of matter with exciton-polaritons: Edge solitons in kagome lattice,* Sci. Rep. **7**, 1780 (2017).

219. S. Lloyd, *Universal quantum simulators,* Science **273**, 1073 (1996).

220. W. Du, S. Zhang, Q. Zhang, and X. Liu, *Recent Progress of Strong Exciton-Photon Coupling in Lead Halide Perovskites,* Adv. Mater. , 1804894 (2018).

221. J. Wang, R. Su, J. Xing, D. Bao, C. Diederichs, S. Liu, T. C. Liew, Z. Chen, and Q. Xiong, *Room temperature coherently coupled exciton–polaritons in two-dimensional organic–inorganic perovskite,* ACS Nano **12**, 8382 (2018).

222. M.-E. Kleemann, R. Chikkaraddy, E. M. Alexeev, D. Kos, C. Carnegie, W. Deacon, A. C. de Pury, C. Große, B. de Nijs, J. Mertens, A. I. Tartakovskii, and J. J. Baumberg, *Strong-coupling of WSe2 in ultra-compact plasmonic nanocavities at room temperature,* Nat. Commun. **8**, 1296 (2017).

223. B. Munkhbat, D. G. Baranov, M. Stührenberg, M. Wersäll, A. Bisht, and T. Shegai, *Self-Hybridized Exciton-Polaritons in Multilayers of Transition Metal Dichalcogenides for Efficient Light Absorption,* ACS Photonics **6**, 139 (2019).

224. J. A. Hutchison, T. Schwartz, C. Genet, E. Devaux, and T. W. Ebbesen, *Modifying Chemical Landscapes by Coupling to Vacuum Fields,* Angew. Chem. Int. Ed. **51**, 1592 (2012).

225. J. A. Hutchison, A. Liscio, T. Schwartz, A. Canaguier-Durand, C. Genet, V. Palermo,

P. Samorì, and T. W. Ebbesen, *Tuning the Work-Function Via Strong Coupling*, Adv. Mater. **25**, 2481 (2013).

226. J. d. Pino, J. Feist, and F. J. Garcia-Vidal, *Quantum theory of collective strong coupling of molecular vibrations with a microcavity mode*, New J. Phys. **17**, 053040 (2015).

227. T. Schwartz, J. A. Hutchison, C. Genet, and T. W. Ebbesen, *Reversible Switching of Ultrastrong Light-Molecule Coupling*, Phys. Rev. Lett. **106**, 196405 (2011).

228. E. Orgiu, J. George, J. A. Hutchison, E. Devaux, J. F. Dayen, B. Doudin, F. Stellacci, C. Genet, J. Schachenmayer, C. Genes, G. Pupillo, P. Samorì, and T. W. Ebbesen, *Conductivity in organic semiconductors hybridized with the vacuum field*, Nat. Mater. **14**, 1123 (2015).

229. P. Roelli, C. Galland, N. Piro, and T. J. Kippenberg, *Molecular cavity optomechanics as a theory of plasmon-enhanced Raman scattering*, Nat. Nanotechnol. **11**, 164 (2016).

230. J. Flick, N. Rivera, and P. Narang, *Strong light-matter coupling in quantum chemistry and quantum photonics*, Nanophotonics **7**, 1479 (2018).

231. D. Lidzey, D. Bradley, M. Skolnick, T. Virgili, S. Walker, and D. Whittaker, *Strong exciton-photon coupling in an organic semiconductor microcavity*, Nature **395**, 53 (1998).

232. J. L. O'Brien, A. Furusawa, and J. Vučković, *Photonic quantum technologies*, Nat. Photonics **3**, 687 (2009).

233. N. Gisin, G. Ribordy, W. Tittel, and H. Zbinden, *Quantum cryptography*, Rev. Mod. Phys. **74**, 145 (2002).

234. G. Zhu and Q. Liao, *Highly efficient collection for photon emission enhanced by the hybrid photonic-plasmonic cavity*, Opt. Express **26**, 31391 (2018).

235. E. Knill, R. Laflamme, and G. J. Milburn, *A scheme for efficient quantum computation with linear optics*, Nature **409**, 46 (2001).

236. L.-M. Duan, M. D. Lukin, J. I. Cirac, and P. Zoller, *Long-distance quantum communication with atomic ensembles and linear optics*, Nature **414**, 413 (2001).

237. M. A. Broome, A. Fedrizzi, S. Rahimi-Keshari, J. Dove, S. Aaronson, T. C. Ralph, and A. G. White, *Photonic Boson Sampling in a Tunable Circuit*, Science **339**, 794 (2013).

238. N. Spagnolo, C. Vitelli, M. Bentivegna, D. J. Brod, A. Crespi, F. Flamini, S. Giacomini, G. Milani, R. Ramponi, P. Mataloni, R. Osellame, E. F. Galvão, and F. Sciarrino, *Experimental validation of photonic boson sampling*, Nat. Photonics **8**, 615 (2014).

239. C. Santori, D. Fattal, J. Vučković, G. S. Solomon, and Y. Yamamoto, *Indistinguishable photons from a single-photon device*, Nature **419**, 594 (2002).

240. A. Pitanti, M. Ghulinyan, D. Navarro-Urrios, G. Pucker, and L. Pavesi, *Probing the spontaneous emission dynamics in Si-nanocrystals-based microdisk resonators*, Phys. Rev. Lett. **104**, 103901 (2010).

241. D. Miller, *Device requirements for optical interconnects to silicon chips*, Proc. IEEE **97**, 1166 (2009).

242. K. L. Tsakmakidis, R. W. Boyd, E. Yablonovitch, and X. Zhang, *Large spontaneous-emission enhancements in metallic nanostructures: Towards LEDs faster than lasers*, Opt. Express **24**, 17916 (2016).

243. K. Liu, S. Sun, A. Majumdar, and V. J. Sorger, *Fundamental Scaling Laws in Nanophotonics*, Sci. Rep. **6**, 37419 (2016).

244. D. A. Long, *The Raman effect* (John Wiley & Sons Ltd., Chichester, United Kingdom, 2002).

245. K. Kneipp, Y. Wang, H. Kneipp, L. T. Perelman, I. Itzkan, R. R. Dasari, and M. S. Feld, *Single Molecule Detection Using Surface-Enhanced Raman Scattering (SERS)*, Phys. Rev. Lett. **78**, 1667 (1997).

246. S. Nie, *Probing Single Molecules and Single Nanoparticles by Surface-Enhanced Raman Scattering*, Science **275**, 1102 (1997).

247. G. C. Schatz, M. A. Young, and R. P. Van Duyne, *Electromagnetic Mechanism of SERS*, in *Surface-Enhanced Raman Scattering* (Springer-Verlag, Berlin/Heidelberg, 2006) pp. 19–45.

248. S. H. Huang, X. Jiang, B. Peng, C. Janisch, A. Cocking, S. K. Özdemir, Z. Liu, and L. Yang, *Surface-enhanced Raman scattering on dielectric microspheres with whispering gallery mode resonance*, Photonics Res. **6**, 346 (2018).

249. T. Hümmer, J. Noe, M. S. Hofmann, T. W. Hänsch, A. Högele, and D. Hunger, *Cavity-enhanced Raman microscopy of individual carbon nanotubes*, Nat. Commun. **7**, 12155 (2016).

250. M. K. Schmidt, R. Esteban, F. Benz, J. J. Baumberg, and J. Aizpurua, *Linking classical and molecular optomechanics descriptions of SERS*, Faraday Discuss. **205**, 31 (2017).

251. S. V. Boriskina and B. M. Reinhard, *Spectrally and spatially configurable superlenses for optoplasmonic nanocircuits*, Proc. Natl. Acad. Sci. U. S. A. **108**, 3147 (2011).

252. K. Santhosh, O. Bitton, L. Chuntonov, and H. G., *Vacuum Rabi splitting in a plasmonic cavity at the single quantum emitter limit*, Nat. Commun. **7**, 11823 (2016).

253. D. K. Gramotnev and S. I. Bozhevolnyi, *Plasmonics beyond the diffraction limit*, Nat. Photonics **4**, 83 (2010).

254. J. A. Schuller, E. S. Barnard, W. Cai, Y. C. Jun, J. S. White, and M. L. Brongersma, *Plasmonics for extreme light concentration and manipulation*, Nat. Mater. **9**, 193 (2010).

255. C. Vieu, F. Carcenac, A. Pépin, Y. Chen, M. Mejias, A. Lebib, L. Manin-Ferlazzo, L. Couraud, and H. Launois, *Electron beam lithography: resolution limits and applications*, Appl. Phys. Lett. **164**, 111 (2000).

256. A. Kinkhabwala, Z. Yu, S. Fan, Y. Avlasevich, K. Müllen, and W. E. Moerner, *Large single-molecule fluorescence enhancements produced by a bowtie nanoantenna*, Nat. Photonics **3**, 654 (2009).

257. M. R. Foreman and F. Vollmer, *Level repulsion in hybrid photonic-plasmonic microresonators for enhanced biodetection*, Phys. Rev. A **88**, 023831 (2013).

258. F. Ruesink, H. M. Doeleman, E. Verhagen, and A. F. Koenderink, *Controlling*

nanoantenna polarizability through backaction via a single cavity mode, Phys. Rev. Lett. **120**, 206101 (2018).

259. I. Mukherjee, G. Hajisalem, and R. Gordon, *One-step integration of metal nanoparticle in photonic crystal nanobeam cavity,* Opt. Express **19**, 22462 (2011).

260. A. F. Koenderink, M. Kafesaki, B. C. Buchler, and V. Sandoghdar, *Controlling the resonance of a photonic crystal microcavity by a near-field probe,* Phys. Rev. Lett. **95** (2005).

261. M. Burresi, T. Kampfrath, D. van Oosten, J. C. Prangsma, B. Song, S. Noda, and L. Kuipers, *Magnetic light-matter interactions in a photonic crystal nanocavity,* Phys. Rev. Lett. **105** (2010).

262. K. G. Cognée, W. Yan, F. L. China, D. Balestri, F. Intonti, M. Gurioli, A. F. Koenderink, and P. Lalanne, *Mapping complex mode volumes with cavity perturbation theory,* Optica **6**, 269 (2019).

263. K. D. Heylman, N. Thakkar, E. H. Horak, S. C. Quillin, C. Cherqui, K. A. Knapper, D. J. Masiello, and R. H. Goldsmith, *Optical microresonators as single-particle absorption spectrometers,* Nat. Photonics **10**, 788 (2016).

264. A. F. Koenderink, *Single-photon nanoantennas,* ACS Photonics **4**, 710 (2017).

265. N. Sadzak, M. Héritier, and O. Benson, *Coupling a single nitrogen-vacancy center in nanodiamond to superparamagnetic nanoparticles,* Sci. Rep. **8** (2018).

266. F. T. Rabouw, P. Lunnemann, R. J. A. van Dijk-Moes, M. Frimmer, F. Pietra, A. F. Koenderink, and D. Vanmaekelbergh, *Reduced Auger Recombination in Single CdSe/CdS Nanorods by One-Dimensional Electron Delocalization,* Nano Lett. **13**, 4884 (2013).

267. C. B. Murray, D. J. Norris, and M. G. Bawendi, *Synthesis and characterization of nearly monodisperse CdE (E = sulfur, selenium, tellurium) semiconductor nanocrystallites,* J. Am. Chem. Soc. **115**, 8706 (1993).

268. M. Fleischmann, P. Hendra, and A. McQuillan, *Raman spectra of pyridine adsorbed at a silver electrode,* Chem. Phys. Lett. **26**, 163 (1974).

269. D. L. Jeanmaire and R. P. V. Duyne, *Surface raman spectroelectrochemistry,* J. Electroanal. Chem. and Interfacial Electrochem. **84**, 1 (1977).

270. E. C. Le Ru and P. G. Etchegoin, *Principles of Surface-Enhanced Raman Spectroscopy,* edited by E. C. Le Ru and P. G. Etchegoin (Elsevier, 2009).

271. C. Vericat, M. E. Vela, G. Benitez, P. Carro, and R. C. Salvarezza, *Self-assembled monolayers of thiols and dithiols on gold: new challenges for a well-known system,* Chem. Soc. Rev. **39**, 1805 (2010).

272. I. Shlesinger, K. G. Cognée, E. Verhagen, and A. F. Koenderink, *Integrated molecular optomechanics with hybrid dielectric-metallic resonators,* Arxiv , 2105.07887 (2021).

273. S. Fan, W. Suh, and J. D. Joannopoulos, *Temporal coupled-mode theory for the fano resonance in optical resonators,* J. Opt. Soc. Am. A **20**, 569 (2003).

274. M. Orrit and J. Bernard, *Single pentacene molecules detected by fluorescence excitation in a p-terphenyl crystal,* Phys. Rev. Lett. **65**, 2716 (1990).

275. M. W. Doherty, N. B. Manson, P. Delaney, F. Jelezko, J. Wrachtrup, and C. L. Hollenberg, *The nitrogen-vacancy colour centre in diamond,* Phys. Rep. **528**, 1 (2013).

276. S. Castelletto and A. Boretti, *Silicon carbide color centers for quantum applications,* J. Phys: Photonics **2**, 022001 (2020).

277. M. Toth and I. Aharonovich, *Single photon sources in atomically thin materials,* Annu. Rev. Phys. Chem. **70**, 123 (2019).

278. D. V. Talapin, J.-S. Lee, M. V. Kovalenko, and E. V. Shevchenko, *Prospects of colloidal nanocrystals for electronic and optoelectronic applications,* Chem. Rev. **110**, 389 (2010).

279. Y. Shirasaki, G. Supran, M. Bawendi, and V. Bulovic, *Emergence of colloidal quantum-dot light-emitting technologies,* Nat. Photonics **7**, 13 (2013).

280. Y.-S. Park, S. Guo, N. S. Makarov, and V. I. Klimov, *Room temperature single-photon emission from individual perovskite quantum dots,* ACS Nano **9**, 10386 (2015).

281. B. Li, H. Huang, G. Zhang, C. Yang, W. Guo, R. Chen, C. Qin, Y. Gao, V. P. Biju, A. L. Rogach, L. Xiao, and S. Jia, *Excitons and biexciton dynamics in single $CsPbBr_3$ perovskite quantum dots,* J. Phys. Chem. Lett. **9**, 6934 (2018).

282. N. A. Gibson, B. A. Koscher, A. P. Alivisatos, and S. R. Leone, *Excitation Intensity Dependence of Photoluminescence Blinking in $CsPbBr_3$ Perovskite Nanocrystals,* J. Phys. Chem. C **122**, 12106 (2018).

283. S. Seth, N. Mondal, S. Patra, and A. Samanta, *Fluorescence blinking and photoactivation of all-inorganic perovskite nanocrystals $CsPbBr_3$ and $CsPbBr_2I$,* J. Phys. Chem. Lett. **7**, 266 (2016).

284. G. Yuan, C. Ritchie, M. Ritter, S. Murphy, D. E. Gómez, and P. Mulvaney, *The degradation and blinking of single $CsPbI_3$ perovskite quantum dots,* J. Phys. Chem. C **122**, 13407 (2018).

285. L. Hou, C. Zhao, X. Yuan, J. Zhao, F. Krieg, P. Tamarat, M. V. Kovalenko, C. Guo, and B. Lounis, *Memories in the photoluminescence intermittency of single cesium lead bromide nanocrystals,* Nanoscale **12**, 6795 (2020).

286. P. Frantsuzov, M. Kuno, B. Jankó, and R. Marcus, *Universal emission intermittency in quantum dots, nanorods, and nanowires,* Nat. Phys. **4**, 519 (2008).

287. T. D. Krauss and J. J. Peterson, *Bright future for fluorescence blinking in semiconductor nanocrystals,* J. Phys. Chem. Lett. **1**, 1377 (2010).

288. A. Efros and D. Nesbitt, *Origin and control of blinking in quantum dots,* Nat. Nanotechn. **11**, 661 (2016).

289. C. Galland, Y. Ghosh, A. Steinbrück, M. Sykora, J. A. Hollingsworth, V. I. Klimov, and H. Htoon, *Two types of luminescence blinking revealed by spectroelectrochemistry of single quantum dots,* Nature **479**, 203 (2011).

290. L. P. Watkins and H. Yang, *Detection of Intensity Change Points in Time-Resolved Single-Molecule Measurements,* J. Phys. Chem. B **109**, 617 (2005).

291. J. P. Hoogenboom, W. K. den Otter, and H. L. Offerhaus, *Accurate and unbiased estimation of power-law exponents from single-emitter blinking data,* J. Chem. Phys **125**, 204713 (2006).

292. D. L. Ensign and V. S. Pande, *Bayesian single-exponential kinetics in single-molecule experiments and simulations,* J. Phys. Chem. B **113**, 12410 (2009).

293. D. L. Ensign and V. S. Pande, *Bayesian detection of intensity changes in single molecule and molecular dynamics trajectories,* J. Phys. Chem. B **114**, 280 (2010).

294. C. H. Crouch, O. Sauter, X. Wu, R. Purcell, C. Querner, M. Drndic, and M. Pelton, *Facts and artifacts in the blinking statistics of semiconductor nanocrystals,* Nano Lett. **10**, 1692 (2010).

295. J. Houel, Q. T. Doan, T. Cajgfinger, G. Ledoux, D. Amans, A. Aubret, A. Dominjon, S. Ferriol, R. Barbier, M. Nasilowski, E. Lhuillier, B. Dubertret, C. Dujardin, and F. Kulzer, *Autocorrelation analysis for the unbiased determination of power-law exponents in single-quantum-dot blinking,* ACS Nano **9**, 886 (2015).

296. Y. J. Bae, N. A. Gibson, T. X. Ding, A. P. Alivisatos, and S. R. Leone, *Understanding the bias introduced in quantum dot blinking using change point analysis,* J. Phys. Chem. C **120**, 29484 (2016).

297. F. R. Hill, E. Monachino, and A. M. van Oijen, *The more the merrier: high-throughput single-molecule techniques,* Biochem. Soc. Trans. **45**, 759 (2017).

298. M. C. A. S. Hadzic, R. Börner, S. L. B. König, D. Kowerko, and R. K. O. Sigel, *Reliable state identification and state transition detection in fluorescence intensity-based single-molecule förster resonance energy-transfer data,* J. Phys. Chem. B **122**, 6134 (2018).

299. S. Schmid, M. Götz, and T. Hugel, *Single-molecule analysis beyond dwell times: Demonstration and assessment in and out of equilibrium,* Biophys. J. **111**, 1375 (2016).

300. S. Pressé, J. Peterson, J. Lee, P. Elms, J. L. MacCallum, S. Marqusee, C. Bustamante, and K. Dill, *Single molecule conformational memory extraction: P5ab RNA hairpin,* J. Phys. Chem. B **118**, 6597 (2014).

301. D. S. White, M. P. Goldschen-Ohm, R. H. Goldsmith, and B. Chanda, *Top-down machine learning approach for high-throughput single-molecule analysis,* eLife **9**, e53357 (2020).

302. B. Shuang, D. Cooper, J. N. Taylor, L. Kisley, J. Chen, W. Wang, C. B. Li, T. Komatsuzaki, and C. F. Landes, *Fast step transition and state identification (stasi) for discrete single-molecule data analysis,* J. Phys. Chem. Lett. **5**, 3157 (2014).

303. H. Li and H. Yang, *Statistical learning of discrete states in time series,* J. Phys. Chem. B **123**, 689 (2019).

304. D. L. Ensign, *Bayesian statistics and single-molecule trajectories*, Ph.D. thesis, Stanford University (2010).

305. I. M. Palstra and A. F. Koenderink, *A python toolbox for unbiased statistical analysis of fluorescence intermittency of multi-level emitters. github repository* `https://github.com/AMOLFResonantNanophotonics/CPA/`. *first commit archived at* `DOI:10.5281/zenodo.4557226`, (2021).

306. I. M. Palstra and A. F. Koenderink, *A python toolbox for unbiased statistical analysis of fluorescence intermittency of multilevel emitters,* J. Phys. Chem. C (2021).

307. D. E. Gómez, J. van Embden, P. Mulvaney, M. J. Fernée, and H. Rubinsztein-

Dunlop, *Exciton-trion transitions in single cdse-cds core–shell nanocrystals*, ACS Nano **3**, 2281 (2009).

308. A. A. Cordones, T. J. Bixby, and S. R. Leone, *Direct measurement of off-state trapping rate fluctuations in single quantum dot fluorescence*, Nano Lett. **11**, 3366 (2011).

309. R. Schmidt, C. Krasselt, C. Göhler, and C. von Borczyskowski, *The fluorescence intermittency for quantum dots is not power-law distributed: A luminescence intensity resolved approach*, ACS Nano **8**, 3506 (2014).

310. M. Wahl, I. Gregor, M. Patting, and J. Enderlein, *Fast calculation of fluorescence correlation data with asynchronous time-correlated single-photon counting*, Opt. Express **11**, 3583 (2003).

311. Ž. Bajzer, T. M. Therneau, J. C. Sharp, and F. G. Prendergast, *Maximum likelihood method for the analysis of time-resolved fluorescence decay curves*, Eur. Biophys. J. **20**, 247 (1991).

312. M. Köllner and J. Wolfrum, *How many photons are necessary for fluorescence-lifetime measurements ?* Chem. Phys. Lett. **200**, 199 (1992).

313. J. Kang and L.-W. Wang, *High defect tolerance in lead halide perovskite CsPbBr$_3$*, J. Phys. Chem. Lett. **8**, 489 (2017).

314. X. Ling, S. Zhou, J. Yuan, J. Shi, Y. Qian, B. W. Larson, Q. Zhao, C. Qin, F. Li, G. Shi, C. Stewart, J. Hu, X. Zhang, J. M. Luther, S. Duhm, and W. Ma, *14.1% CsPbI$_3$ perovskite quantum dot solar cells via cesium cation passivation*, Adv. Energy Mater. **9**, 1900721 (2019).

315. Q. Van Le, H. W. Jang, and S. Y. Kim, *Recent advances toward high-efficiency halide perovskite light-emitting diodes: Review and perspective*, Small Methods **2**, 1700419 (2018).

316. C. Zhang, L. Turyanska, H. Cao, L. Zhao, M. W. Fay, R. Temperton, J. O'Shea, N. R. Thomas, K. Wang, W. Luan, and A. Patanè, *Hybrid light emitting diodes based on stable, high brightness all-inorganic CsPbI$_3$ perovskite nanocrystals and ingan*, Nanoscale **11**, 13450 (2019).

317. F. Cichos, C. Borczyskowski, and M. Orrit, *Power-law intermittency of single emitters*, Curr. Opin. Colloid Interface Sci. **12**, 272 (2007).

318. A. A. Cordones and S. R. Leone, *Mechanisms for charge trapping in single semiconductor nanocrystals probed by fluorescence blinking*, Chem. Soc. Rev. **42**, 3209 (2013).

319. J. Cui, A. P. Beyler, T. S. Bischof, M. W. B. Wilson, and M. G. Bawendi, *Deconstructing the photon stream from single nanocrystals: from binning to correlation*, Chem. Soc. Rev. **43**, 1287 (2014).

320. P. A. Frantsuzov, S. Volkán-Kacsó, and B. Jankó, *Model of fluorescence intermittency of single colloidal semiconductor quantum dots using multiple recombination centers*, Phys. Rev. Lett. **103**, 207402 (2009).

321. S. Volkán-Kacsó, P. A. Frantsuzov, and B. Jankó, *Correlations between subsequent blinking events in single quantum dots*, Nano Lett. **10**, 2761 (2010).

322. P. A. Frantsuzov, S. Volkán-Kacsó, and B. Jankó, *Universality of the Fluorescence Intermittency in Nanoscale Systems: Experiment and Theory*, Nano Lett. **13**, 402 (2013).

323. A. Merdasa, Y. Tian, R. Camacho, A. Dobrovolsky, E. Debroye, E. L. Unger, J. Hofkens, V. Sundström, and I. G. Scheblykin, *"Supertrap" at Work: Extremely Efficient Nonradiative Recombination Channels in MAPbI3 Perovskites Revealed by Luminescence Super-Resolution Imaging and Spectroscopy*, ACS Nano **11**, 5391 (2017).

324. F. Stefani, X. Zhong, W. Knoll, M.-Y. Han, and M. Kreiter, *Memory in quantum-dot photoluminescence blinking*, New J. Phys. **7**, 197 (2005).

325. B. Patra, H. Agrawal, J.-Y. Zheng, X. Zha, A. Travesset, and E. Garnett, *Close-packed ultrasmooth self-assembled monolayer of CsPbBr3 perovskite nanocubes*, ACS Appl. Mater. Interfaces **12**, 31764 (2020).

326. G. Baffou and R. Quidant, *Thermo-plasmonics: using metallic nanostructures as nano-sources of heat*, Laser Photon. Rev. **7**, 171 (2013).

327. S. Seth, T. Ahmed, and A. Samanta, *Photoluminescence flickering and blinking of single CsPbBr3 perovskite nanocrystals: Revealing explicit carrier recombination dynamics*, J. Phys, Chem. Lett. **9**, 7007 (2018).

328. T. Ahmed, S. Seth, and A. Samanta, *Mechanistic investigation of the defect activity contributing to the photoluminescence blinking of CsPbBr3 perovskite nanocrystals*, ACS Nano **13**, 13537 (2019).

329. V. S. Chirvony, S. González-Carrero, I. Suárez, R. E. Galian, M. Sessolo, H. J. Bolink, J. P. Martínez-Pastor, and J. Pérez-Prieto, *Delayed luminescence in lead halide perovskite nanocrystals*, J. Phys. Chem C **121**, 13381 (2017).

330. Y. Wang, M. Zhi, and Y. Chan, *Delayed exciton formation involving energetically shallow trap states in colloidal CsPbBr3 quantum dots*, J. Phys. Chem. C **121**, 28498 (2017).

331. S. J. W. Vonk, M. B. Fridriksson, S. O. M. Hinterding, M. J. J. Mangnus, T. P. van Swieten, F. C. Grozema, F. T. Rabouw, and W. van der Stam, *Trapping and detrapping in colloidal perovskite nanoplatelets: Elucidation and prevention of nonradiative processes through chemical treatment*, J. Phys. Chem. C **124**, 8047 (2020).

332. H.-H. Fang, S. Adjokatse, H. Wei, J. Yang, G. R. Blake, J. Huang, J. Even, and M. A. Loi, *Ultrahigh sensitivity of methylammonium lead tribromide perovskite single crystals to environmental gases*, Science Adv. **2**, e1600534 (2016).

333. K. Miyata, T. L. Atallah, and X.-Y. Zhu, *Lead halide perovskites: Crystal-liquid duality, phonon glass electron crystals, and large polaron formation*, Science Adv. **3**, e1701469 (2017).

Acknowledgments

This has been a wild ride. There have been disappointing setbacks and exhilarating successes. Sometimes the route to our goal was a thrilling chase, other days the only option was a slow, grinding trudge. By myself, I could have never made it through. Without support from friends, the difficult moments would have been that much heavier. Without the ideas and input from my colleagues, I could have never figured all this out. In my moments of success, I could never have felt as triumphant if I could not celebrate with others. So this section is where I will take a moment to thank you all. All you sweet, wonderful, mad, beautiful people that I have had the privilege and the honor to work with, to talk with, to share food and fun and laughter. It has been an absolutely electric experience. Thank you for helping me when I was stuck, for believing in me when I didn't think I could do it, and for cheering with me when I found out that I could.

To Femius, I want to extend my deepest thanks for the time I spent here. You have taught me so much over the years, with one of the first lessons on the mitigation of risk when embarking on a chancy project. Our journey to this thesis could hardly be called smooth. I remember coming up to you at the end of my second year, fretful that I didn't have any almost any results to share in my colloquium. Yet, with your help, and another switch in projects, somehow we pulled through. We made it happen. Without your skills, knowledge, and your dedication to the group, I don't think it would have been possible. Katerina, thank you for helping me on the last leg of this work. Stepping into the role of my co-promotor relatively recently, you managed to give invaluable feedback on my manuscript, even though a deal of the content was new to you. Tom, I also want to thank you for being my initial co-promotor. I enjoyed the discussions we had, and I am sorry that you could not see the culmination of my work.

When I joined the group, it had an almost entirely different makeup of people. Marko, you have been a beacon of constancy in an ever-changing patchwork of people. Your wide array of skills is remarkable, and your dry sense of humor during coffee breaks was always greatly appreciated. Hugo, you were my office mate from day one. Your in-depth knowledge of such a broad range of topics is truly impressive. I deeply enjoyed working with you, though I do hope you have finally stopped using Python 2. Ruslan, we have had a long run together, with you starting only a few months before I did. This gave you a small head start with which I feel I have never truly been able to catch up. I enjoyed our dinners and the talks we had. You are good company and a good friend. Thank you for always making sure I was invited to your kick-ass house parties, I hope I will be invited again in the future. Let's bring the house down. Annemarie, you have a particular brand of wisdom that I have always found inspiring, though I find it difficult to put into words. Where others might fight to effect change at

the cost of their mental well-being, you seem to have a healthier balance. Of the many people I know, you are one of the few who seems to live with an innate understanding of human capabilities and limitations. It gives you a remarkable strength, and I hope you can continue to draw from that. Kévin, you are one clever dude, to the point where your work has always seemed part witchcraft to me. The months you were not in Bordeaux brightened my days, I just wish there were more of them. You seem to have tremendous energy to learn and to explore new topics. You always were able to set aside time to help me out, and were at my side when I needed support. You have always been an excellent conversational partner, with fresh and interesting insights. We might not always agree, but I do always have fun arguing with you, and always learn something new. Christiaan, thank you for all the fun conversations we shared. I hope we can share drinks again, and I really liked your idea of a cross-USA road trip with the members of the group. Tom W., your presence in the office brought cheer and laughter. Your boundless positivity is infectious, and it went a long way to balance out my own nervous nature. Whatever the issue, you will find a way to argue that it will all be all right in the end. You are kind and helpful, have a good sense of fun and a good sense of when to celebrate success (which is always). Sharing an office with you made me leave for work with a smile. It has been fantastic working with you.

Ilse, you were the first student that I supervised. Not just through one, but through two projects. I had great fun supervising you, and I count your project as one of the main successes of my second year. In fact, the measurements you did were of such high quality that they kept me entertained well into my extension. I am confident you will do well in your endeavors, if only you do not underestimate yourself. You got this. Beniamino, your project was one of twists and turns, in part because we bit off somewhat more than we could chew with the initial project plan. I have learned a lot from you, and have had fun in our time working together. Radek, scientific discussions with you frequently took a deeper turn than I was initially expecting them to, which just made it more fun. Ilan, you have helped me take the work on hybrids to the next level. With your energy and passion, we pulled together the measurements and analysis we had been working towards in the 11th hour. It was scary, it was stressful, but you helped me make it happen. It's been great working with you. Your skills and good sense of fun were a joy to be around. Deba, it feels like you only just joined our group, but at the time of writing, you've been with us well over a year. It's always great hanging out and discussing with you. With your inquisitive nature and unscrupulous approach to asking questions, I'm sure you'll do great. Nelson, as the most junior member of a group that has seen so many fantastic people come and go, you have some big shoes to fill. With your friendly nature and always-appropriate demeanor, I'm sure they'll fit you perfectly.

Of course, my interactions and collaborations extended beyond the members of the Resonant Nanophotonics group. Much of my work could not have been completed without the help from the amazing nanolab staff. Hans, your position as the most senior staff member at AMOLF gives you a unique perspective on the history of the institute. You know the ins and outs of every device in the nanolab, making it a joy to

have their function explained by you. Bob, Dimitry, and Igor, you share that endless willingness to assist the users of the lab. Between the four of you, you make beautiful things come together. On top of that, you are great fun over drinks.

The great collaborative environment at AMOLF has an opportunity to shine at our recurring colloquia and poster sessions, where we can get together to spread ideas. Albert, Femius, Ewold, Esther, Bruno, Said, and Erik, thank you for organizing these for us. I want to thank Erik and Biplab for your help with the perovskites project. Biplab, for helping Ilse make the perovskites needed, and Erik for co-supervising this project. Your ideas and input were greatly appreciated, and your enthusiasm for the project never seemed to reduce, even as it steadily grew far beyond our initial plans. Marco Konijnenburg, I want to thank you for your help with the initial steps into the analysis of quantum dot data. Paula, Esther, Henriette, and Sylvianne, it was a great experience to organize the first AMOLF and ARCNL Women in Science day with you. I learned so much, had a lot of fun, and hope that it helped to build towards a fairer, more inclusive working environment. A skill set where I have always felt comfortable are those involved in throwing a good party. Joining the PV for a year came with ample opportunity to scatter around my wild ideas. I loved playing Which One Is Rum, building a spaceship to kidnap Paula's goat, and taking the institute off adventuring at the fortress. Wessel, Jente, Yvonne, Alexander, Christian, Jeffrey, and all the rest, it's been a blast working on all these silly things with you. Stay crazy, my friends.

There are so many other people that I was not directly associated with through any collaboration but have nevertheless filled AMOLF life with color. Arno, you seem to possess an otherworldly sense of selflessness. It suits you beautifully. I'll never forget how you came to rescue me with your lock-picking skills when my bike keys were stolen, and it remains one of my favorite bar stories. Harshal, your joyful approach to work and life is an inspiration, and I have loved ripping up the dance floor with you. Benjamin B., you are such an energetic and sociable person, I don't know how you can keep up with yourself. Hans H., I hope you keep that ocarina close to your heart, so it can stay filled with music. Lukas, your work is some of the most visually appealing of any I have seen at AMOLF, and it could not have been in better hands than yours. Luuk, just one door down the hall, I have always enjoyed our coffee-side chats and your unique take on the topics we discussed. Wim, we didn't get the opportunity to talk much in recent times, but your collaborative spirit and willingness to chase the white rabbits in your life is an example to live by. Robin, conversations with you were ever a breath of air. Your boundless curiosity for topics both in and out of physics will lead you to the most interesting places. I'll meet you there for drinks. Jesse, how fun to run into each other at AMOLF again. Nasim, Jenny, Lucie, Marloes, Susan, Mareike, Magda, Verena, and Julia, your company and conversation were always great. Keep smashing that patriarchy. Nick, Roel, Jente, Sander, Moritz, Benjamin D., Nikhil, Giorgio, Agustin, Niels, Iarik, Andrea, Tom V., and René, you are a collection of great people, near and dear to my heart. To all of you, stay the awesome, lovely people you are.

Liefste Sarah, wat hebben we toch al een hoop meegemaakt samen. Wie had dat

gedacht, een vriendschap van al bijna 20 jaar. Ik kan geen genoeg krijgen van samen thee doen, de vragen van het leven overpeinzen, en lekker Boonanza spelen. Ik hoop dat ik snel kan komen vieren in jullie nieuwe huis. Beste Aleid, wij kennen elkaar ook al zo lang, sinds die vroege jaren op school. Wat voel ik me toch een geluksvogel dat ik jou heb leren kennen. Er breken nieuwe en spannende tijden aan, ook met de komst van de kleine Meijke en Tjitske, en ik kan niet wachten om ze met jullie te delen. Wladick, with a friendship starting in a Sicily restaurant with a conversation about cameras, I have had the best of times with you. Tobias, with some distance between us, I hope that we can meet up in either of our hometowns again soon, and sing some silly songs together. Nicolaas, ook al woon je nu op een afstand, geniet ik altijd van jouw gezelschap. Bij onze telefoontjes is het altijd gezellig, en ik wil je graag binnenkort opzoeken. Richard, Chayenne, Mardie, Nikki, Mike, Michelle, Tim, Michiel, Thalia, en Julia, wat een leuke groep heb ik overgehouden aan mijn tijd bij U Dance. Laten we een groot feest houden, samen dansen, en genieten. Gordian, met zo'n lange tijd die we samen hadden, is het geen wonder dat we beiden veel zijn gegroeid. Ik heb zo veel geleerd in mijn tijd met jou. Dank je wel voor jouw partnerschap toen, en dank je wel voor je vriendschap nu.

Romeo and Thor, thank you for helping me keep stress levels under control in this time. There are few things quite a soothing as having you purr away happily on my lap. Margot, I hope that Mel will always run the tap for you. Robin-Yann, wie had dit nu gedacht. Omdanks de afstand die ons het leven niet altijd even makkelijk maakt, kan ik altijd mijn wel en wee met jou delen. Je hebt mij zo geholpen in deze tijd, en ik kan niet wachten je weer te zien. Bedankt voor jouw partnerschap, jouw eerlijkheid en jouw openheid naar mij en anderen. Daniel, you came into my life as a surprise, as most people do. It was always helpful to have you remind me that things will work out, and to get the perspective of someone who has the benefit of having already completed their PhD. You are one smart cookie. Let's keep having all the strange, mad, and sometimes surprisingly deep conversations that we love. Thank you for taking care of me through these tough times, and thank you for letting me take care of you. I think we make a good team.

Mijn lieve ouders, mama, papa, ik ben jullie zo dankbaar voor alle ondersteuning die jullie mij gegeven hebben. De herinneringen dat ik niet altijd hoef te stressen, waren zeer gewaardeerd. Pap, jouw adviezen waren mij zeer waardevol. Mam, jouw perspectief als natuurkunde outsider hielp mij herinneren dat er meer is in het leven. Dank jullie wel dat ik altijd mijn hart kon komen luchten, en lekker kon uitblazen in Enschede. Mijn lieve zusjes, Marieke en Eline. Net als tijdens mijn onderzoek, staan jullie bij mijn verdediging naast me. Ik had me geen betere paranymphen kunnen voorstellen. Jullie zijn in de tijd van mijn PhD mijn steun en toeverlaat geweest. Bij jullie kon ik, ongeacht het moment, komen uithuilen of feest vieren. Men zegt dat de vrienden waarmee je jezelf omgeeft, je gekozen familie vormen. Als dat zo is, zijn we dubbel familie. Drie paar handen op één hart.

About the Author

Isabelle was born in Summit, NJ, in 1992. At age 4 she moved to the Netherlands, the home country of her family. She received her high school diploma from Praedinius Gymnasium in Groningen at age 18, after which she moved to Utrecht for her physics Bachelor at Utrecht University. After her Bachelor, Isabelle took a break from studying to work for a humanitarian organization in Sri Lanka, and did an internship at TNO Delft, studying the Raman spectra of asbestos. Returning to Utrecht, she completed the first year of her Masters program Nanomaterials: Chemistry and Physics. For her Masters research project, she applied for a position at AMOLF, in the group of prof. dr. Kobus Kuipers, where she worked on the design of photonic crystal waveguides to create superchiral light. Having greatly enjoyed her time at AMOLF, Isabelle started her PhD project in the same institute, now in the group of prof. dr. Femius Koenderink. The results of her PhD research are published in this thesis. When she's not doing science, Isabelle likes to spend her time creating. She enjoys sewing, building, gardening, and historical crafts. She is passionate about dancing, and practices a variety of styles.