

©(2018) Society of Photo-Optical Instrumentation Engineers (SPIE). One print or electronic copy may be made for personal use only. Systematic reproduction and distribution, duplication of any material in this paper for a fee or for commercial purposes, or modification of the content of the paper are prohibited.

Citation:

Matthew Pelton, "Single quantum dots strongly coupled to plasmons: from Fano interference to Rabi splitting (Conference Presentation)," Proc. SPIE 10721, Active Photonic Platforms X, 1072126 (17 September 2018); <https://doi.org/10.1117/12.2319994>

DOI:

<https://doi.org/10.1117/12.2319994>

Access to this work was provided by the University of Maryland, Baltimore County (UMBC) ScholarWorks@UMBC digital repository on the Maryland Shared Open Access (MD-SOAR) platform.

Please provide feedback

Please support the ScholarWorks@UMBC repository by emailing scholarworks-group@umbc.edu and telling us what having access to this work means to you and why it's important to you. Thank you.

Great. So like I said, thank you, those of you who are stick around to the very end of this thing. I hope that it's worth it.

I'm going to be talking, as he said, about strong coupling between single quantum dots and plasmon resonances. And I guess the original motivation for this, how we got into it, was looking at plasmon-enhanced processes, plasmon-enhanced emission.

And plasmon enhancement is probably how the field of plasmonics got started back in the '70s when people first saw surface-enhanced Raman scattering and understood that it was because of the localization of electromagnetic fields next to plasmonic metal nanostructures. And that sort of localization can enhance any kind of optical process of not just Raman scattering, but harmonic generation, or wave mixing, or even linear processes like absorption or emission.

And so since then, there's been a lot of work on looking at using plasmons to increase spontaneous emission rates to enhance emissions from all sorts of different emitters. [INAUDIBLE] a few different communities have been spending the last 20 years or so doing this and have developed different ways of looking at it, which of course are ultimately equivalent, but give a different insight into what's going on.

So in a place like this, if you're an electrical engineer, you might like to think of a plasmonic metal nanoparticle as an antenna. You put an emitter next to it, and the optical antenna broadcasts radiation from that emitter to the far field by overcoming the impedance mismatch between that microscopic emitter and free space the same way that the antenna inside your cell phone broadcasts radiation from a electronic oscillator to the far-field.

If you're a chemist, you might like to think of it in energy transfer picture where the emitter next to a nanostructure, and undergoes an efficient near-field energy transfer similar to a Forster energy transfer from its dipole to the

the plasmon inside the particle. And then that plasmon can either itself radiate to the far field or decay non-radiatively and produce heat in the metal nanoparticle. You can look at the balance of the different rates and use that to characterize the so-called Purcell factor, how much faster the emitter is coupling to the metal nanostructure than it would couple to the free space in the absence of that metal nanoparticle.

[INAUDIBLE] if you have a physics background like me, you might see this in terms of a cavity QED kind of picture. What the plasmonic metal nanoparticle is doing in this case is very similar to what people have seen since the '80s when they've taken atoms or molecules, or solid-state emitters like quantum dots and put them inside microscopic cavities. The spontaneous emission is not just a property of the emitter itself, but it's an interaction between the emitter and the electromagnetic state surrounding it. So by changing the density of those available [INAUDIBLE] states to emit into, you change the rate at which the emitter spontaneously emits radiation.

And so this Purcell factor in this picture is very roughly proportional to the quality factor of the cavity, how long electromagnetic energy is stored inside the resonator, and is inversely proportional to the mode volume. How tightly confined the electromagnetic fields are inside that resonator.

So dielectric cavities can have very high quality factors, with crystal cavities can have 10,000-- there have been 100,000 Qs. But their mode volumes are fundamentally limited by the diffraction limit. So you can't get much smaller than about half a wavelength on each side here.

On the other hand, for plasmonics, Qs are low. Plasmons dephase [INAUDIBLE] femtoseconds. Tens of femtosecond time scales. So Qs are in the order of 10. But mode volumes are no longer restricted by the diffraction limit. So you can go down, really, to the single nanometer dimensions and more than

can make up for the loss of Q and get the largest [INAUDIBLE] factors that have been seen have been seen in plasmonic structures.

In the context of cavity QED, this is what's known as the weak coupling regime, where you're enhancing a process that's already happening. But this leaves the question, if we can do so well in the weak coupling regime, can we get into the strong coupling regime of cavity QED using plasmonics instead of optical cavities?

So strong coupling means, instead of having this irreversible emission of radiation, now there's a coherent oscillation of energy from the emitter to the plasmonic back to the emitter and back and forth. And this coherent exchange leads to a splitting of the normal modes of the system, the formation of hybrid plasmon, exciton, or plexciton states that are now the new fundamental excitations of the system.

So in a very simple picture, you can just think of the plasmons, or the quantum dots, or the photon field in the cavity. As a boson, you have your emitter which you can approximate as a two-level system. You put in an interaction between the two of them, and where the plas--

absorbs a photon from the emitter, or the plasmon excites the emitter, diagonalize the Hamiltonian, then we get the splitting of these new normal modes, which on resonance is just given by this interaction strength. And it's called the so-called vacuum Rabi frequency.

So if you liked--

of course, this is all considering a completely closed system without any damping going on. In reality, if you actually want to talk about modes, if you're going to claim to be forming new excitations in the system, you're going to have to have this mode splitting to be larger than the rate at which energy is dissipated from the system. So that means this coupling strength is going to have to be larger than whatever the fastest damping rate in the system is. And this coupling rate goes with the dipole moment of the emitter, but it's inversely proportional to the square root of the mode volume here.

So this isn't quite as good for plasmonics as the coupling is, because it's only with a square root instead of the inverse of the mode volume. Which is one of

the reason why it's been a lot easier to get weak coupling than strong coupling in plasmonics, because the plasmon damping rate is so large. And the solution a lot of people have been doing to

this is to use a whole bunch of emitters and coherently couple them to one plasmon mode. So if you do that, you get another factor here, which goes as the square root of the number of emitters that are effectively coupled simultaneously to the plasmon. If you make that the macroscopic number of emitters, then you can get large coupling strengths and get strong coupling between plasmons and emitters.

But we're really interested in trying to do this down at the level of a single molecule, or a single quantum dot. Because that opens up possibilities that people have been exploiting in the cavity QED community for quantum information processing. You can do things like entanglement generation, or [INAUDIBLE] exchange, quantum state exchange between the emitter and the photon.

And all those basically are enabled by the fact that you have strong single-photon nonlinearities here. Basically mediated by this extremely strong coherent light-matter interaction inside your cavity QED system, you can get single photons to interact strongly with one another, and thereby get quantum information control.

So single-emitter strong coupling is an enabling technology for quantum information. But so far, this has all been done in cavity QED at low temperatures. And again, that comes down to the fact that the mode volumes in dielectric cavities are ultimately limited to half a wavelength cubed, which means that you have to make the damping rates in your system small enough that they're less than this coupling strength that you can get in that case.

So again, the high Qs help. You can make very high Q dielectric cavities. And so their cavity damping rate isn't that large, but of course, you still have your emitter damping rate, which at room temperature, the emitter dephasing is still going to be larger than the coupling strength you can get in dielectric cavities. And you have to cool your emitters down. Really solid state emitters have to be cooled down to low temperatures in order to get strong coupling.

In plasmonics, in principle, this can become much smaller. And if we can get down to very large mode volumes, then we can get coupling strengths that would be large enough to overcome the damping at room temperature and enable the same sorts of things at the single-emitter level at room temperature.

So I've been talking about this stuff for a while, but now I do have data. But before I show that, the question is, of course, how do you know when you have strong coupling? What's your experimental signature?

There have been a couple of reports out there that have been claiming strong coupling at the single-emitter level in plasmonics. And so far, almost all of those have involved making scattering measurements. You look at the spectrum of light scattered from the plasmonic structure. You see the two peaks in that scattering spectrum.

You say, OK, now, two peaks, we have two normal modes. The separation between those two peaks is our mode splitting frequency. Look, I have strong coupling. But we'd like to be a little more rigorous about that.

So we can do it quantum mechanically. And if I get there, I'll show you what it looks like when we actually deal with the quantum mechanical equations. You have to add damping into it, so now you end up having to solve the quantum master equation. You put in the [INAUDIBLE], and you can do all of this. But if you don't have to, it's nice not to have to do that.

And if you're only interested in the linear response of this system, meaning the response to low light levels so that your emitter stays almost entirely in the ground state, then the quantum equations of motion reduce just to classical equations of motion. And those classical equations of motion are two coupled oscillators, one referencing the dipole of the plasmon which is driven by this external field, one of them representing the dipole moment of your emitter, which is much smaller, so it basically doesn't couple the external field, and it's only driven by the coupling to the plasmon through the near field.

So in sophomore physics, we can solve these equations of motion. We can infer extinction cross-sections or scattering cross-sections from those solutions and figure out, really, what the spectral response of these systems should be for different parameters that go in there.

So let's say we put in realistic parameters for a plasmon quantum dot coupled system and just tune up the interaction strength between those two and see what happens. So again, in the weak coupling regime, where we see the Purcell effect, the scattering spectrum is essentially unchanged. It looks just like the plasmonic nanostructure would look by itself.

If we go into the strong-coupling regime, you see here this clear mode splitting, two well-resolved peaks. But in between, if we put in the coupling strength that we know is not large enough to reach strong coupling, we still see two maxima in the scattering spectrum.

But if we look at this more carefully, this is not actually two peaks. What is really going on here is we have one peak, and then we have a dip in the middle of that central peak. So rather than Rabi splitting going on, we have what I would call induced transparency, or Fano interference.

So at the frequency of the quantum dot, we have destructive interference between the quantum

dot dipole and the plasmon dipole induced by this coupling between the two of them. That cancels out the scattering of the plasmonic nanoparticle on resonance. So seeing two peaks in your spectra is certainly not sufficient to know that you're in the strong coupling regime. And the separation between these two peaks in this intermediate coupling, or in the strong coupling regime, is not your vacuum Rabi frequency.

We can do a detailed line shape analysis, and we will do that to figure out what the actual line shapes are. But it would be nice to have a quality measure that lets us know, are we in the strong coupling regime or not? And so the reason this arises fundamentally is because scattering is a coherent process. And this is coming from an interference effect that's driven by this coherent incident field.

But we could look at photoluminescence instead. Photoluminescence, where we excite the emitter into an excited state, decays down into its ground state and then emits photons as an incoherent process. There's no phase relationship between the light that's coming out and the light that's going in, so it's not subject to this Fano interference.

So we can do a semi-classical model for what you expect, again, based on what's well known in the cavity QED community, and calculate what you expect to see in the photoluminescence spectra for the weak, intermediate, and strong coupling regime. So blue is scattering, green is photoluminescence. So again, in weak coupling, nothing changes from the uncoupled system spectrally.

Strong coupling, we have well-resolved peaks in the photoluminescence. And then weak coupling, basically, we still have just the single peak in the photoluminescence spectra. So here, we can really qualitatively distinguish weak and strong coupling by looking at both scattering and photoluminescence instead of just scattering by itself.

So this is what we'll do. We need a system that gives us as small of a mode volume as possible. It's been known for quite a while that the best way to localize fields in plasmonic structures is to use so-called gap plasmons. So either two particles next to each other, or a metal particle next to a metal film can lead to much stronger light confinement in the gap between the two of them-- especially if this is a nanometer-scale gap-- than you would get from any kind of single nanoparticle structure.

So we'd like to make this gap as small as possible. We've put quantum dots inside the gap instead of the molecules that have been used in some previous structures, because they're much brighter and more stable emitters and we want to look at photoluminescence, and also because they have very large transition dipole moments. So that also helps in terms of getting a large coupling strength, in this case.

So this is our target structure. It's easy to draw. It's not very easy to make. For our first demonstrations, we're just going to rely on luck, make a whole bunch of things, and find a few that give us the structure that we want.

So in collaboration with the group of Mary Christine Danielle at UNBC, we synthesize gold-- metal nanoparticles, we synthesize colloidal cadmium selenide and cadmium sulfide quantum dots. We functionalize their surfaces with complementary ligands and then link the two of them together at low concentrations so there's just a few quantum dots on the surface of each metal nanoparticle and they're well separated from each other. So you can see here, you have a couple of quantum dots on the surface of one metal nanoparticle.

Then if we take these coupled particles, put them down onto a silver film, eventually now and then, we're going to end up with one quantum dot where it's supposed to be in the gap between the two of them. Most of the time, it won't happen. But if it does happen, we know we've only got one dot there, because they're well isolated from each other, and because the mode volume is so small, and so the localized field only extends over an area that's just a little bit bigger than the quantum dots themselves here.

And then we have a very patient graduate student who looks at a few hundred of these things and 99% of the time sees weak coupling. So the scattering spectrum, the photoluminescence spectrum look the same as the metal nanoparticle by themselves, the quantum dot by themselves, apart from some background that comes from luminescence from the metal itself.

We still have very strong Purcell effect. So the decay rate from the quantum dot compared to the decay rate of an isolated quantum dot is increased on the order of 100 times. Still, on all of these structures, despite the large Purcell effect, we're not seeing any evidence of strong coupling.

Every now and then, we find a structure that does show two peaks in the scattering spectrum. And in this case, we still see one peak in the photoluminescence spectrum. There's a shoulder here, but this is coming from charged states, or by exciton states inside quantum dot. This is the main luminescence peak, and you can see it's pretty much lined up with the minima in the scattering spectrum, like we saw in our theoretical predictions.

And just to make it clear that this is really because of quantum dot plasmon coupling and not some strange thing in the metal nanostructure, and also that it's due to one quantum dot, what we do is we take the structure, we turn off the laser power, illuminate it with really bright light for a little while-- couple of minutes-- turn the laser power back down, and then measure the same thing again. And so the quantum dot blue shifts because it goes through photo-oxidation and eventually becomes a little bit smaller. The plasmon shifts a little bit, probably because of migration of metal atoms on the surface of the particle.

But even after everything moves around, we still have a dip. We still have a luminescence peak right where that dip is. Everything moves together and ends up in the right place. And if we had more than one quantum dot, we wouldn't expect them to all move exactly the same way as one another.

And we can take that theory that I showed before, a very simple analytical theory, and compare it to what we see experimentally. So what we've done here is we've taken the scattering spectrum, fit it to our formula, and then used the parameters that we got from that fit to predict the photoluminescence spectrum. And we can see we have really good agreement, in this case, between what we see experimentally and what we predict.

And we get explicitly a coupling strength that is not in the strong coupling regime. That's in this intermediate Fano interference regime. And I guess I shouldn't talk much about other people, but I should say that this gives you more-- a hint that maybe some previous reports were more in this regime than in the true strong coupling regime.

And with some patience, we do see some stuff that we think really is in the strong coupling regime. So we see not only two very well-resolved peaks in the scattering spectrum, but we see photoluminescence peaks that are well-resolved and that line up with those scattering peaks.

And again, we can compare it to our theory. Here, the agreement is a little less good. That's probably because we're using a very simple theory. The theory corresponds to light that comes from the plasmonic metal nanoparticles.

So radiation that's coupled from the quantum dot, and then to the plasmon, and then from the plasmon to free space. But in the strong coupling regime, there should be a significant amount of light that comes from the quantum dot directly to free space that will have a slightly different spectrum that may account for the differences that we're seeing.

But in some sense, it doesn't really matter, because qualitatively, we have two peaks in photoluminescence. We have two peaks in scattering. This is a definitive measurement of strong coupling at the single-emitter level at room temperature in a plasmonic system.

So one natural question is, why do we see strong coupling sometimes? Intermediate coupling sometimes? What's the difference among the systems?

For this, we turn to calculations. We're not quantitatively reproducing it, but qualitatively, this gives us some idea of what's going on. And so if we actually do these finite element calculations, even for our ideal case where the quantum dot is right where it's supposed to be between a spherical metal particle and a metal surface, we actually don't see strong coupling. We barely even see any Fano interference.

In previous calculations, people have used unrealistic dipole moments for their emitters to try to make up for this. But what we're thinking is, this is actually not what our metal nanoparticles look like.

Anybody that's used gold particles knows that their quasi-spherical particles are not spherical at all. They're faceted. It's a cubic crystal structure. And so they have edges, they have vertices.

If we put edges or vertices into our calculations, we get stronger Fano interference, and we get into the strong coupling regime. So we really think it's not just the overall structure, but the details of the nanometer scale structure of our metal nanoparticles that's determining whether we have strong coupling in the coupling strength between the plasmon and the emitter.

So this really shows that if we want to do this more than a couple percent of the time, if we want to have a well-controlled system where we can do this with high yield, we really need detailed control over this nanometer scale structure. That's the big challenge.

There's one way to get slightly better control. We worked in collaboration with a group of Markus Raschke at the University of Colorado. And essentially the same system, but now instead of a metal nanoparticle on top of the quantum dot, their group uses a scanning probe tip. So a sharp gold tip that's scanned over this substrate with a gold film underneath, quantum dots on

top. And we can place the gold tip exactly over individual quantum dots and look at the photoluminescence spectra that they emit.

So this is one example. Blue is the quantum dot by itself. If the tip comes close but is not quite above the quantum dot, we seem more like coupled [? out, ?] but the spectrum is basically unchanged. But when the tip is exactly the right place right above the quantum dot, you have this clear splitting of the photoluminescence peak into these two peaks as a real signature of strong coupling.

And here, the nice thing is we can control this thing very precisely by moving the tip around. So we can turn this strong coupling on and off by taking the tip that's far away from the quantum dot and moving it closer and closer to the quantum dot. We go from a single peak into two peaks that split further and further apart until we have this clear strong coupling where the tip is right above the quantum dot.

And then once it's above it, you can pull the tip away vertically. The fields drop off very quickly so luminescence gets weaker and goes down into a single peak with only a couple of nanometer distance. So showing that there's a very tightly-localized field that's controlling the strong coupling between the quantum dot and the tip.

This also gives us the opportunity to control this dynamically, in a way. First by moving the tip by a couple nanometers, we could turn the strong coupling on and off and achieve things in quantum information that are based on dynamical control over strong coupling. Do I have two more minutes? OK.

So very quickly, having done this, one question is, what is it good for? And I think I motivated the idea of using strong coupling for quantum information processing. I don't actually mean to be too negative about this intermediate coupling or Fano interference regime, because I think there are actually a lot of interesting things that can be done there as well.

So here's where I'll turn back to those quantum mechanical calculations, because we want to look at what happens not just in the linear regime, but in the non-

linear regime. As we increase the intensity of the light on the system, how does this Fano interference change?

And we see we start with a cleared induced transparency dip. At low light, it has [INAUDIBLE]. As we turn off our intensity to relatively modest values, this dip goes away.

And it's not hard to understand qualitatively. The quantum dot, once it absorbs a single photon, goes from being absorptive now to being transparent. So as far as this coupled systems goes, optically, it's as if the quantum dot isn't there. This interference disappears as soon as you absorb a single photon inside that quantum dot.

So this very large change in the extinction cross-section, absorption or scattering cross-section of this coupled plasmon-exciton system has the potential to serve as a low-power optical modulator. It should also work at very high speeds because the coherence times of these things are on the femtosecond time scale. It's very small.

So these are the ingredients that are needed for next-generation optical information processing. We want [INAUDIBLE] joules or sub-

picosecond nanoscale optical components for integrated photonics, and this is a potential way of achieving that in a coupled system. And since you're standing up, I'll leave it at that, and I'll be happy to take questions.

[APPLAUSE]