

Single Nanoparticle Absorption Spectroscopy: Chemical Dynamics using Optical Microresonators

By
Levi T. Hogan

A Dissertation submitted in partial fulfillment
of the requirements for the degree of

Doctor of Philosophy
(Chemistry)

At the
UNIVERSITY OF WISCONSIN-MADISON
2022

Date of final oral examination: September 22nd, 2022

The dissertation is approved by the following members of the Final Oral Committee:

Randall Goldsmith, Professor, Chemistry, Physical

Silvia Cavagnero, Professor, Chemistry, Analytical

Thomas Brunold, Professor, Chemistry, Chemical Biology

Judith Burstyn, Professor, Chemistry, Inorganic

For my Mom, Grandma, and Grandpa.

TABLE OF CONTENTS

TABLE OF CONTENTS	iii
TABLE OF FIGURES.....	viii
Abstract.....	xii
Acknowledgements.....	xiii
Thesis Roadmap	xxi
1. Can you see a single molecule?	1
2. Scientific background for this thesis	8
2.1. Introduction	8
2.1.1. Recommended reading.....	8
2.1.2. Overview	9
2.2. WGM Microresonators	10
2.3. Our Photothermal Measurement	12
2.4. Quality Factor (Q).....	13
2.5. Mode Volume (V).....	14
2.6. Thermo-Optic Coefficient (dn/dT)	15
2.7. Sensitivity, Responsivity, and Limit-of-Detection (LOD)	16
2.8. Energy Matching and Phase Matching/Conservation of Momentum	18
2.9. Polarization	19
2.9.1. Why write (or read) a long section on polarization?	19
2.9.2. Fundamentals of polarization	21
2.9.3. Linear polarization designations: S & P, TE & TM	23
2.9.4. The Poincaré sphere	25
2.9.5. Jones matrices.....	28
2.9.6. Measuring Polarization	31
3. Building, Calibrating, and Operating a Microbubble Photothermal Spectrometer	32
3.1. Introduction	32
3.2. Overview: The anatomy of our photothermal microscope	33
3.3. Setting up the Probe Beam	35
3.3.1. Laser choice, optical path and components	35
3.3.2. Polarization stability.....	42
3.3.3. Pound-Drever-Hall locking & laser feedback	43

3.3.4.	PDH components (and how they connect to other electronics)	46
3.4.	The Pump Beam Path	48
3.4.1.	Overview	48
3.4.2.	Alignment safety, tips, and tricks	50
3.4.3.	The beam-scanning relay system	54
3.4.3.2.	Aligning the pump relay	59
3.4.4.	The Liquid Crystal Polarization Rotator (LCPR)	62
3.5.	Amplitude Modulation & Lock-in Detection	70
3.5.1.	Method of modulation	71
3.5.2.	The math of LIA	71
3.5.3.	Modulation speed	72
3.5.4.	Time constant	73
3.5.5.	Sensitivity	74
3.5.6.	Dynamic reserve	75
3.5.7.	Other lock-in settings	75
3.6.	The M Squared Laser System	75
3.7.	The Control Computer	78
3.8.	Machined Pieces	79
3.9.	Taper Pulling: Protocol & Considerations	79
3.9.1.	Overview	79
3.9.2.	What makes a good taper?	80
3.9.3.	Fabricating a taper	82
3.9.4.	Future improvements for tapering	88
3.9.5.	Ultrathin tapers	89
3.10.	Calibration Procedures	89
3.10.1.	Probe beam piezo	89
3.10.2.	Galvo mirror calibration	91
3.10.3.	Pump beam focusing	92
3.10.4.	Pump beam power calibrations	93
3.10.5.	FOV calibrations	95
3.10.6.	Extinction map calibrations	96
3.10.7.	Pump beam polarization	99
3.11.	Controlling the Microscope (Code and Software)	106

3.11.1.	Measuring error slopes	107
3.11.2.	Measuring quality factors	108
3.11.3.	Photothermal mapping	109
3.11.4.	Polarization rotation	112
3.11.5.	Under construction: Spectra acquisitions	113
3.11.6.	Interactive locations tab	117
3.11.7.	Optimization tab	118
3.11.8.	Time series collection tab	120
3.11.9.	Pros/cons of transitioning to other control software	121
3.11.10.	Quirks and tips for various drivers, GUIs, etc.....	121
3.11.11.	Future code improvements.....	123
3.12.	Taking Photothermal Data	124
3.12.1.	Aligning the microresonator, taper, camera, and pump beam	124
3.12.2.	Finding and coupling to resonances	127
3.12.3.	Error slope measurement	128
3.12.4.	Having a robust and trustworthy resonance lock.....	128
3.12.5.	Locking procedure.....	130
3.12.6.	Initial mapping	134
3.12.7.	Optimizing signal.....	135
3.12.8.	Taking “real” data	135
3.13.	Understanding photothermal data.....	135
3.13.1.	Calculating resonance shift.....	135
3.13.2.	Polarization results.....	136
3.13.3.	Post-acquisition application of correction factors.....	137
3.13.4.	COMSOL modeling to obtain absorption cross-section.....	137
3.14.	Photographs of the Photothermal Setup.....	138
4.	Toward Real-Time Monitoring and Control of Single Nanoparticle Properties with a Microbubble Resonator Spectrometer.....	140
4.1.	Abstract.....	140
4.2.	Main Text: Introduction	141
4.3.	Main Text: Experimental Design	144
4.4.	Main Text: Operation of Single-Particle Microresonator Spectrometers.....	146
4.5.	Main Text: Probing Photophysical Features of Single AuNRs	149

4.6.	Main Text: Selecting an etchant	151
4.7.	Main Text: Single AuNR Reactions	152
4.8.	Main Text: Mechanistic Discussion	157
4.9.	Main Text: Single AuNR Rotations	160
4.10.	Main Text: Conclusions	163
4.11.	Main Text: Methods/Experimental	164
4.11.1.	Microbubble fabrication	164
4.11.2.	Tapered optical fiber fabrication	165
4.11.3.	Instrumentation for photothermal spectroscopy	165
4.11.4.	Polarization plots	166
4.11.5.	Bulk/UV-Vis studies	166
4.11.6.	Nanorod deposition in microbubbles	167
4.11.7.	Reactions in microbubbles	167
4.12.	Supporting Information: Geometric Parameters of Microbubbles	168
4.13.	Supporting Information: COMSOL Simulations	169
4.14.	Supporting Information: Comparison with Microtoroids	173
4.15.	Supporting Information: Static Offset vs. Modulated Signal	175
4.16.	Supporting Information: Mode Shifting for Different dn/dT Values (Experimental)	176
4.17.	Supporting Information: Modeling of LPB and TPB During Etching	178
4.18.	Supporting Information: Thermal Expansion	184
4.19.	Supporting Information: Diagram for Mode Indices	184
4.20.	Supporting Information: Additional Single AuNR Etching Data	185
4.21.	Supporting Information: Additional Rotation Data	187
4.22.	Supporting Information: Bulk Reaction Results	188
4.23.	Supporting Information: Brief Discussion on Extracting Reaction Kinetics	190
4.24.	Supporting Information: Background on AuNR etchants	191
4.25.	Supporting Information: Impervious nanorods	191
4.26.	Supporting Information: Discussion of Role of CTAB in Etching Mechanism	192
4.27.	Supporting Information: Estimation of CTAB Remaining on AuNRs after Deposition	194
4.28.	Supporting Information: Estimation of Nanorod Temperature Increase	199
4.29.	Supporting Information: Theory of Nanorod Rotation	199
5.	Towards Single Molecules: Reducing Photothermal Background	202
5.1.	Abstract	202

5.2.	About This Chapter	203
5.3.	Background on Photothermal Background.....	203
5.4.	Origins of Photothermal Background	205
5.5.	Two Types of Background.....	206
5.6.	How Low Does Background Need to Be for Single Molecule Experiments?	207
5.6.1.	Toroid calculations.....	207
5.6.2.	Our molecule of interest.....	208
5.6.3.	Microbubbles calculations	209
5.6.4.	Impact of dn/dT	209
5.6.5.	Impact on Q-factors	210
5.6.6.	Other considerations	210
5.7.	Measuring Photothermal Background.....	211
5.7.1.	Pump off map.....	211
5.7.2.	Pump on Maps	212
5.7.3.	Mapping parameters to consider	213
5.7.4.	More instrumental parameters to consider	215
5.8.	Processing Photothermal Background Data	216
5.8.1.	Pump off maps.....	216
5.8.2.	Region maps.....	217
5.8.3.	Object maps	218
5.9.	Instrumental Background	220
5.9.1.	Reference bleed through in the lock-in amplifier.....	220
5.9.2.	Ground loops.....	223
5.9.3.	Other pieces of equipment	224
5.9.4.	Ensuring a correct lock.....	230
5.9.5.	Resonance stability	231
5.10.	Materials Background	231
5.10.1.	Background patterns: reproducibility and chromaticity.....	232
5.10.2.	Pump power dependence.....	234
5.10.3.	The carbon black hypothesis.....	235
5.10.4.	The glass-stress/in-coupling hypothesis	236
5.10.5.	The dusty bubble hypothesis	239
5.10.6.	Other mapping experiments.....	243

5.10.6.1. Mapping the tapered fiber.....	243
5.10.6.2. Mapping the tapered capillary.....	244
5.10.7. Changing the focusing objective	244
5.10.8. Future work: distance-dependence of background variation	249
5.11. So, are We Good Enough to Look at Single Molecules?	250
5.11.1. The first single molecule attempt	251
5.11.2. Additional single molecule attempts	253
5.12. Fundamental Materials Absorption	253
5.13. Getting to Single Molecules in the Future	256
5.14. Conclusion.....	257
6. Photoluminescence Studies of Single Carbon Dots	258
6.1. Motivation.....	258
6.2. Introduction to Carbon Dots	259
6.3. Background	260
6.4. Previous Single CD Studies.....	263
6.5. Results and Discussion	265
6.6. Conclusions & Future Work on Carbon Dots	275
6.7. Experimental Methods.....	276
7. Thesis Recap & Outlook	278
8. References	279

TABLE OF FIGURES

Figure 2-1. Constructive interference in microresonators	11
Figure 2-2. How a resonance looks in our experiment.....	11
Figure 2-3. Comparison of shift in a high-Q vs. a low-Q resonance.....	14
Figure 2-4. The photothermal impact of large vs. small mode volume.....	15
Figure 2-5. How S- and P-polarized light interact with a surface	24
Figure 2-6. A TE mode in a microbubble resonator.....	25
Figure 2-7. The Poincaré Sphere.....	27
Figure 3-1. A simplified diagram of the probe beam path.....	35
Figure 3-2. Probe beam losses	38
Figure 3-3. Thermal broadening of a resonance.....	38
Figure 3-4. Comparison of resonances and error signals	44
Figure 3-5. A simplified schematic of our experimental electronics	46
Figure 3-6. A simplified diagram of the probe beam path.....	50
Figure 3-7. A simplified schematic of the beam-scanning relay system.....	55
Figure 3-8. Somewhat reproducible beam-walking at resonator surface due to HWP rotation.....	63
Figure 3-9. Polarimeter positioning for LCPR alignment	68
Figure 3-10. Transmission of light through the hypothetical two-LC-LCPR.....	70
Figure 3-11. Lock-in amplification diagrams	72
Figure 3-12. A cartoon of how a lock-in amplifier rejects frequency noise.....	74
Figure 3-13. The M Squared Laser System.....	76
Figure 3-14. The double decker diode holder.....	80
Figure 3-15. Illustrative data of mode beating of a tapered optical fiber while it is being pulled	87
Figure 3-16. Probe voltage piezo calibration.	91
Figure 3-17. Calibration of the beam scanning relay system (galvo mirror parameter)	92
Figure 3-18. Diode power calibrations.....	95
Figure 3-19. Example FOV calibration.....	96
Figure 3-20. Extinction map calibration	98
Figure 3-21. Calibration of the LCPR using the proxy-dipole method	101
Figure 3-22. Polarimeter calibration data for LCPR as a function of M Squared Wavelength	104
Figure 3-23. Example Jones Matrix Simulation.....	105
Figure 3-24. The “Error Slope” tab of the master photothermal code (version 3g).....	108
Figure 3-25. The “Quality Factor” tab of the master photothermal code (version 3g).....	109
Figure 3-26. The “Mapping” tab of the master photothermal code (version 3g)	112
Figure 3-27. The “Polarization” tab of the master photothermal code (version 3g).....	113
Figure 3-28. Control of the M Squared system in the photothermal master code (version 4a)	117
Figure 3-29. The “Interactive Locations” tab of the photothermal master code (version 3g)	118
Figure 3-30. The “Optimization” tab of the photothermal master code (version 3g).....	120
Figure 3-31. Robust lock assurance.	129
Figure 3-32. Locked-resonance oscillations.....	133
Figure 3-33. Pictures of microscope setup	138
Figure 3-34. Pictures of microscope setup	139

Figure 4-1. Microbubble Absorption Spectroscopy	145
Figure 4-2. Optical resonances in microbubble resonators.....	149
Figure 4-3. Probing photophysical features of single AuNRs.....	151
Figure 4-4. Etching single AuNRs	154
Figure 4-5. Etching reactions driven at two different pump wavelengths	157
Figure 4-6. Proposed mechanistic explanation for etching rates	160
Figure 4-7. Orientation control of single AuNRs	163
Figure 4-8. Optical micrographs of the microbubbles used to collect the data in the main text.....	169
Figure 4-9. Steps in a simulating the resonance shift from a point-heat source.....	170
Figure 4-10. Pictures of microbubble models used for the COMSOL models in this report	171
Figure 4-11. Steady-state COMSOL simulations	173
Figure 4-12. Comparison of photothermal static offset to modulated signal	176
Figure 4-13. Error signal shown for multiple resonances	177
Figure 4-14. Modeled spectra for the TPB and LPB	180
Figure 4-15. For TPB and LPB driven, tip-only etching.....	181
Figure 4-16. LPB and TPB driven etching	181
Figure 4-17. The isotropic etching scenario for both TPB and LPB driven reactions.....	182
Figure 4-18. The isotropic etching scenario for both TPB and LPB driven reactions.....	183
Figure 4-19. Mode number nomenclature	185
Figure 4-20. Additional examples of LPB-driven reactions	186
Figure 4-21. TPB-driven reactions.....	187
Figure 4-22. Reproducible staircase plots.....	187
Figure 4-23. Bulk reactions with large nanorods	188
Figure 4-24. Bulk reactions with small nanorods.....	189
Figure 4-25. Langmuir isotherm data for CTAB	198
Figure 5-1. BHQ-3 structure, spectral parameters, and solubility.	209
Figure 5-2. Representative pump on map, pump off map, pump off histogram	217
Figure 5-3. Excision map.	220
Figure 5-4. Lock-in amplifier measurements, before and after the post-offset procedure	223
Figure 5-5. DAQ input and lock-in output tests	227
Figure 5-6. Wavelength-dependent background mapping patterns.	233
Figure 5-7. Materials background is linear with pump power.....	234
Figure 5-8. Higher photothermal background at bubble edges	237
Figure 5-9. Background pattern as a function of resonance choice and pump polarization.....	238
Figure 5-10. Results of plasma cleaning a dusty microbubble.....	241
Figure 5-11. Testing different pump objectives.....	246
Figure 5-12. Comparison of beam spot size for three different objectives.....	247
Figure 5-13. Galvo calibration for different objectives.....	248
Figure 5-14. FOV calibrations for the red beam, for two different objectives	249
Figure 5-15. Comparing instrumental and materials background for an air- vs. water-filled bubble.....	255
Figure 6-1. Carbon dot samples	266
Figure 6-2. Bulk data for CDs	266
Figure 6-3. Spuncoat CD Dilution Series	267
Figure 6-4. Fluorescence trajectories for single CDs.....	268

Figure 6-5. Calibration of HWP for excitation polarization studies	268
Figure 6-6. Excitation polarization results for one set of CDs using the rotating HWP method	269
Figure 6-7. The quadruple polarization experiment.....	270
Figure 6-8. Single frames from movies for Trolox experiments	272
Figure 6-9. Excitation polarization results for CDs in air, spuncoat from water.....	272
Figure 6-10. Excitation polarization results for CDs in air, spuncoat from Trolox solution	273
Figure 6-11. Excitation polarization results for CDs in water.....	273
Figure 6-12. Excitation polarization results for CDs in Trolox solution.....	274
Figure 6-13. Carbon Dot spectral acquisition	275

Abstract

Light-matter interactions underpin many of the ways in which our world operates and how we seek to understand it. This is especially true in science, where we use different types of light to learn about the chemistry and physics of various materials. Importantly, while traditional “bulk” measurements are good for many applications, they unfortunately can miss out on the rich heterogeneity present at the nanoscopic scale. Thus the past few decades have seen a revolution in chemical analysis, aimed at studying molecules and nanoparticles one-at-time, with the majority of this revolution being in the field of fluorescence microscopy, to phenomenal results. However, there are many analytes of interest that do not emit light and thus must be studied with a method other than fluorescence. To meet this pressing need, our research group in recent years leveraged advancements in nanofabrication to construct a novel photothermal absorption spectrometer using optical microresonators. In this thesis, I describe my work in advancing this microresonator spectrometer platform for aqueous single molecule and nanoparticle studies using microbubble resonators. First, I discuss in detail how to build such a spectrometer from scratch, calibrate it, and operate it. Then, I include a story wherein we used this new platform to study and control the chemical and orientational dynamics of gold nanorods. Next, I delve deep into diagnosing and remedying photothermal background, which is the last thing standing between our platform and single molecule absorption spectroscopy, and how we can further improve our experimental design to reach single molecules. Lastly, I describe preliminary results in ongoing photoluminescence studies of single Carbon Dots, with the eventual goal of studying this intriguing and quickly rising class of nanoparticles using simultaneous photophysical measurements.

Acknowledgements

*"In 900 years of time and space, I've never met anyone who wasn't important."
~The Doctor*

Similarly, in 6 years of graduate school, I have not met anyone who wasn't important. There have been so many different people who have made this long PhD journey, and the long route leading up to it, possible. I would surely fail if I tried to list them all, and I suspect that the bookbinders only have so much paper to print this thesis. However, I must still spend some time thanking quite a few folks.

Firstly, the staff here in the Department of Chemistry at the University of Wisconsin-Madison is fantastic. Thank you to Arrietta, Kristi, Beatriz, Mary, Alicia, Kayla, Marc, Francisca, Ariel, Taylor, Char, Matt, Pat, and other administrators and support staff for keeping the department running. Thank you to Kendall and Steve for all of your teachings and advice in the machine shop. Thank you to Rob, Blaise, and Bill for teaching me so much about electronics and other interesting aspects to instrumentation. Thank you to Jim Maynard for furnishing whatever is needed for a given chemical demonstration in a timely and high-quality manner. Thank you to the folks at EH&S, including Tilak, Aaron, Bindu, and Karisa for prioritizing safety, and to Ive (and randy!) for emphasizing the importance of safety in our department. Thank you to Mike and Aaron for your shipping expertise. Thank you to Chad, Emily, the rest of Purchasing for putting up with my many, many, many purchases in the past few years, whether from PunchOut vendors or random electronics vendors from basements in Oregon. Thank you to Bruce and Kate for running the print shop, and always being so helpful and welcoming to our group. Thank you to Alan, Jim, and other Comphelp folks for keeping the digital infrastructure of the department intact. Thank you to Tristan and James for running the stockroom. Thank you to Judith Burstyn, for chairing the department for what was surely more disasters than any other chair has had to deal with in department history, and with excellent results given the circumstances. Thank you to Clark for your chairing, and continuing the now-tradition of steering the department through difficult times. Thank you to Mike Schwartz for your

wonderful leadership within the CSN, as well as Miriam Krause (UMN-TC) for your wonderful leadership within the CSN. And thank you to Jeff Nielsen, for your dedication to and excellent performance of your herculean job of keeping our cluster of buildings as functioning as possible (much to my surprise, I was never trapped in an elevator!). Whatever the University pays you, I wish they would double it. The above is not a comprehensive list of folks, and I also thank all of the staff I omitted that also have kept the department running with their hard work.

Thank you to my committee members, especially Silvia and Thomas for being on my committee for years now. Thank you to Judith for agreeing to join my committee in this final stretch of my PhD; I enjoyed my time TA-ing for you years ago. Randy, I will thank you in more detail in short order.

To the members of the goldsmith group [*sic*]thank you so, so much. You are such a brilliant and creative group of human beings. And more importantly, you are a loving and caring group of human beings. I will now elaborate for multiple pages.

Erik, I couldn't have asked for a more patient, intelligent, and helpful mentor for my first few years of graduate school. While your knowledge of the scientific literature was sometimes bewildering, you never talked down to me or anyone else in the group. Your calm demeanor and endless support towards your groupmates contributed value to this group's culture, and still does. Thank you for teaching me so much!

Feng, it was a pleasure to enter the group at the same time as you. The projects you accomplished during your time here would not have been possible without your impressive persistence, especially given that you became a dad during graduate school! I particularly enjoyed getting to attend your wedding and reception. Thank you for all the conversations throughout the years, and I hope that you are having an enjoyable time at Stanford!

Ceci, we've worked together a lot the past few years both in the photonics world as well as with Science is Fun. In fact, you might be the groupmate that I've interacted the most with, save maybe Erik.

While the lab work has been intellectually meaningful, the SciFun work has been more personal, as we have facilitated workshops, farmers markets, gala events, and oh so many other events at such a wide variety of venues. You are going to make a fantastic professor, teacher, and mentor in the foreseeable future, and I look forward to seeing your successes! And your completed ring resonators. Veronica, you are certainly one of the anchors that keeps the group grounded. You're always willing to help your colleagues and always willing to keep the boss in check if needed. In addition to all that, you are also a great friend and fantastic scientist. Please drop me a SpongeBob reference every once in a while.

Brendan, thank you for the many scientific brainstorming sessions. You are always willing to talk about science, whether it be giving advice on things that you know much about, or asking insightful questions about the work of others, which often lead to interesting understandings and discoveries. Beau, your persistence with optimizing the FFPC fabrication process is so impressive, and your attention to detail is fantastic. You also clearly care about the wellbeing of your friends and coworkers, which I greatly appreciate. Katherine, you have shown a deep level of independence on your research project, and I look forward to seeing your successes with it! I never thought I'd learn about k-space imaging of metamaterials, but somehow you taught me about that. I also have greatly enjoyed the Science is Fun events that we have done together. Brandon H., you always have something intriguing to contribute to a conversation. And you always are ready to talk about a science concept with a colleague to help them out, or chat about a non-science concept that you somehow are literally an expert in, to pass the time by and make the workday a bit easier. Thank you for teaching me how to found a super-PAC as well. Ray, you are admirably persistent while doing some very difficult science. Your sarcastic attitude can't possibly obscure your caring nature, and both add greatly to the group culture. You're a great colleague and friend, and I wish that we had had more opportunities to facilitate science literacy events together, as I'm confident that this is something you will be excellent at in your future pursuits! Mackinsey, our different research pursuits have led to us interacting less than I wish scientifically! Your positive attitude has been a bright

addition to the group culture and office atmosphere. I think that you will do great things with ZMWs, and be an excellent CHO! Additionally, your ability to apparently perform for almost any role in a circus is quite impressive, and I hope that you will be ringmaster for a goldsmith group performance someday.

Brandon M., it was a pleasure to work with you while you were in the group. I admired you for your Letterkenny quoting prowess, and admired you even more for your dedication to improving the coding infrastructure and maintenance within the group. Together we also built a good deal of the new lab room's infrastructure, which while difficult, was made much better by getting to have fun with you while doing it. Alex Fairhall, you have been a delight to work with, and helped me to keep my sanity through the M Squared saga. You took being tasked with helping maintain the most expensive piece of lab equipment in the group, and constructing the second most expensive, in stride, even as a quite early graduate student! Best of luck on reforming group code management practices, and on your cryostatic adventures.

Simi, despite your disclaimers, every question you have asked and suggestion you have made during group presentations has been insightful and on-point! Additionally, your baking skills are unmatched; thank you for all of the delicious baked foods the past couple of years. Best of luck trapping the ABELs. Julia, you joined the group and were immediately handed a quite difficult project, which after a bit you took sole custody of; you've been doing an excellent job working on it! You also have been a splendid person to work with, and contribute admirably to group sass levels. Keep up the good work, and I look forward to hearing about Raman photons soon! Daniel and Sam, I have had less time to get to know you than I'd like this past year of graduate school, but I know that you are already both making great additions to the group through your resonator fabrications, and I'm excited to see your photonic journeys continue in the future. Daniel, please delegate that M Squared safety stuff to someone soon!

Save for Erik, I have thus far only mentioned students who were my contemporaries or younger in terms of graduate school years, but it is important to acknowledge my elders as well. Andrew, your

calm, kind, and helpful demeanor has been a rock for this group for many years. I know that you will graduate soon (as you should!), but I also know that the group will miss your presence. You have never even given a second thought to helping me when I've asked for advice. Also, your GeoGuessr skills are quite impressive. Morgan, the clarity with which you taught concepts to the group, and the obvious time and dedication that went into preparing your presentations were inspiring, and your wall of dachshunds was impressive. Kasie, you were a heroic fabrication guru for the group, and your personality made the goldsmith-refugee office so much fun. Alex Foote, your ability to conjure analogies to teach complicated topics was inspiring. You are a genuinely kind and loving human being, and I know that any future workplaces and communities of which you are a part will benefit immensely from your presence.

There are some other old guard group members who I had the joy of interacting with that I have yet to mention. Kevin, we only overlapped for a couple of weeks, but those weeks were enjoyable. And of course, your legacy as resonator grandpa still resonates in the group today! Thank you to Daniel Hinton, James, Katie, and Angela as well, for forming such an excellent group culture that I got to join and be a part of. And thank you to Lydia too, even though you confiscated my taper-holder/tuning-fork for several months.

To Nicolas, Windy, Evan, Jordan, Lucy, and other undergraduates I have missed in this list: you also play a vital role in the culture of the goldsmith group, and not just in terms of preparing fibers for ablation. Keep up the good work, and best of luck in your future pursuits, grad school or otherwise!

This group has also had some excellent post-docs in my time here (exclusively excellent, in fact). LM, your prowess as a fluorescence sorceress has been a godsend these past few months, and the tenacity with which you've attacked FFPCs is inspiring. You also throw some kick-ass parties, and are an excellent friend. It's been splendid to have you on this side of the pond for the past couple of years, and I know that the group will miss you sorely when you hop back to the commonwealth. Hoang, you were a kind and thoughtful group member, and your persistence in pursuing your experiments was quite impressive. I

have no doubt that you are doing an excellent job at Washburn! Sudheer, working on you on fiber ablation was quite the ordeal. While I can't say that my memories of fiber ablation itself are all fond, the memories of interacting with you certainly are. I hope that you're doing well in Singapore!

Mike Mattei, your inquisitive nature seems to know no bounds, and neither does your helpfulness towards your colleagues. I already miss the deep scientific ponderings, where a simple question somehow would turn into an hour of trying to figure out the momentum matching of three-wave mixing processes. With your dedication to explaining topics clearly and mentorship you will make an excellent PUI professor in the near future! Johannes, I only got to spend a couple of months as desk-neighbors with you, but during that time I really enjoyed all of our conversations, whether about physics, world events, local restaurants, or American colloquialisms. Not only did you bring some much-needed expertise to our group, but you connected with the group immediately on a personal level. It's clear that you are a caring human being who contributes to the communities you are a part of, and I know this will benefit you in life, as well as your pursuits in sustainable nuclear fusion. Carlos, it has been a pleasure having you as one of the photonics pholks! I enjoy the in-depth science discussion we have, and the questions you ask. Your understanding of PDH locking has also been wonderful for the group!

Ruohan, our time with you was cut short due to logistical constraints surrounding the pandemic, but I enjoyed going to dinner with you when you interviewed with the group, and wish you well in future pursuits. Tzu-Ling, you were only with the group for several months, but I enjoyed our conversations during that time. You always had excellent questions to ask, and taught me a good deal about chiral microresonators!

randy, thank you so much for being my advisor these past few years. Your rapid-paced intellect and immense creativity always lead to engaging and enthusing scientific conversations, and the sometimes out-of-the-blue brainstorming sessions on whatever new whacky idea you'd thought-up kept me on my toes. You make sure that the group is producing top-notch science in an invigoratingly

multidisciplinary and collaborative atmosphere. And as you already know, you have a truly fantastic research group that any other professor should be jealous of. Whether by your direct influence, or by your sitting back and let the group set its own culture, our group is vibrant, weird, super intelligent, and remarkably creative. And most importantly, they're a caring group of people who looks out for one another and supports each other. All of the above has made my graduate school experience so meaningful, both at a scientific and at a personal level, and I know that I am going to miss this group of people immediately upon graduating.

To the Science is Fun crew, you have been such a vital part of my graduate school experience. I have already mentioned a few of you, and must here thank a few more. Jerry, Rodney, and Ron, your many years of experience have been wonderful to learn from. Cayce, without your organization of events, among so many other roles you fill, our group would not have functioned. Jill, it was a pleasure getting to work with you for the couple of years that we overlapped. Amanda, it was fantastic getting to work farmers markets, teacher workshops, and other events with you, and I know that your students are lucky to have you as a teacher! Tom, it has been brilliant getting to publish a paper with you. Your persistence in science education is admirable, and our conversations throughout the first year or so of the pandemic were very valuable and meaningful to me. Other members of the group and affiliates that I have not mentioned: thank you as well.

Bassam, you have been a great mentor, and an even greater friend these past few years. I (usually) leave our conversations with more clarity than I went into them with, and getting to absorb your wisdom has been a remarkable experience. Thank you for everything that you do and have done for society, and thank you for everything you have done for me. The connections you strive for in both your personal and professional pursuits make relationships with you, and life overall, much more meaningful. I hope you enjoy your "retirement"! And give Oliver a treat for me, please.

To all of my past teachers, thank you. Rhonda and Amy so especially, you were early and central pillars in my training as a scientist. The freedom you afforded me to run amok within the chemistry closets everyday (literally) after school for a couple of hours was splendid. I hope that we continue to keep in touch, and I hope you know what an incalculable impact you make in the lives of your students, and those that they go to teach later on.

To my Madison friends. Adam, you were a great roommate for five years, and I am very glad that we signed that lease after knowing each other for about 6 hours. You are a truly excellent friend. Emily, I am glad that Adam met you, and that you have continued to put up with him. You are a thoughtful friend with whom I look forward to further years of gardening discussions and jointly chastising Adam. Chris, you always look out for me, and have made sure to pull me out of my graduate research hole when I really needed it. Somehow we have known each other for over 10 years now, which I think means that you're stuck being my friend forever, for which I'm very grateful. Elainie, you are always so welcoming and kind, and I appreciate having you as my friend! Your students are clearly lucky to have you.

I also have friends from previous parts of my life, only a few of which I will mention here, but all of which are so important. Niko, you have continually supported my endeavours and been an excellent friend, even as much as you try and nudge me into pursuing a career in venture capitalism. Becca, you always keep it real, and are a supportive and great friend. Ryan, while you now live in a desert that no human being should inhabit, you have continued to be a friend throughout the years, which I appreciate greatly. Alec, we had some hoopla times together those years ago which I still think of fondly, and I am glad that we have stayed connected throughout the years. Erich, it has been splendid returning to much more frequent hangouts in recent years, even if they have been digital; thank you for your commitment to organizing activities for our friend group. To the many other friends that I connect with, whether frequently or infrequently, thank you as well.

Jenna, my partner, it was serendipitous to meet you whilst working the 2020 election polls, and I think that boded well from the start. I so enjoy our adventures together, whether hiking various terrains, traveling, or even just watching movies together at home. Your positivity is a welcome tempering to my cynicism, and your persistence once you set your mind to something is impressive. Your adventurous attitude is contagious, and speaks to your bravery. Someday, you shall plan marvelous cities. I love you very much, and look forward to our many adventures together ahead!

Last, but absolutely not least, I thank my family for all of their support, throughout my entire life as well as my time in graduate school. My grandpa, a father figure to me, unfortunately did not make it to see my graduation with my PhD, as much as he wanted to. He has been a continual source of comfort for me during my PhD, both before and after his passing. My uncle Tim also passed away during my graduate school years, and is dearly missed. Grandma, you have been a constant positive influence in my life, and I love you dearly. You and grandpa showed me how to love other people unconditionally, despite their flaws. And while the Covid pandemic has been largely tragedy, I treasure the several weeks that I spent in lockdown with my you in the beginning of it all. Simone, I love you a lot, and am very proud of the hard-working and loving woman that you are. I'm confident that you can do anything you set your mind to, and look forward to watching your future pursuits! Mom, thank you so, so much for raising Simone and me. You always put us first, often to your own personal sacrifice. Your humor is contagious, and your persistence in the face of adversity is inspiring. I love you very much.

"Everything's got to end sometime. Otherwise, nothing would ever get started."

~The Doctor

Thesis Roadmap

Chapter 1 of this thesis is my “chapter for the public”, supported by the Wisconsin Initiative for Science Literacy, and written with the conviction that it is important that scientists share their work with the public, and critical that we engage in science literacy and appreciation efforts. I deliberately put this chapter as the first chapter to emphasize to my colleagues this importance. Additionally, I think that this chapter is also a good introduction for scientists, too! Jargon shouldn’t be necessary to appreciate the importance of a scientific endeavor, and I didn’t want to bust out of the gate with dense material.

Chapter 2 provides technical background that will hopefully be of assistance later in this thesis. It is far from a comprehensive treatise on light-matter interactions, or on microresonators for that matter, and I explicitly direct the reader to other documents, including previous theses from my research group, for such treatments. Instead, I introduce specialized topics that will be repeatedly used throughout this thesis, and try to distill them to lift the jargon-veil. While I do include some math, I also try to do this in an intuitive, qualitative way where possible.

Chapter 3 is arguably the center-piece of this thesis (it’s certainly the longest chapter). This chapter is my attempt to explain how to build, calibrate, and operate a microresonator spectrometer from the ground up. As I wrote this, I found that there was no end to the level of detail I could go into, and thus some topics receive moderate explanations while some take a very deep dive. I hope that I have not left out any particularly critical topics. I also hope that this document is a useful resource for saving future group members time when working with our instruments, especially if it can help them avoid some of the mistakes that I have learned from. Lastly, perhaps this chapter can be a helpful blueprint for more photothermal microresonator setups to be built around the world!

Chapter 4 is a paper that I wrote (and co-authored with Erik Horak, who worked with me in lockstep while also being a fantastic mentor), wherein we took the microresonator platform our group

had developed with microtoroids and moved it to a solution-exposed environment. Minus some formatting adjustments, the main text and supporting information appear in this thesis as they did in the published article. For the reader interested in denser and more technical introduction to microresonator photothermal spectroscopy as compared to Chapter 1 and Chapter 3, I suggest reading Chapter 4 first!

While the limits-of-detection of both of our photothermal microscopes easily surpass the necessary threshold for single molecule studies, photothermal background still plagues us. Thus, I undertook an extensive exploration into the sources of this background in microbubbles resonators. In Chapter 5 I describe my characterizations of photothermal background and its various sources, how we achieved some significant reductions in background, initial (if so far unsuccessful) attempts at single molecule studies, and possible routes for further reduction of photothermal background.

Chapter 6 does a bit of a pivot! I switch gears to describe an up-and-coming class of nanoparticles called “Carbon Dots”, and the outstanding need for single-particle studies in this field. We hope to eventually do simultaneous photothermal and photoluminescence measurements on this and other systems, but for now, I only discuss some recent and preliminary photoluminescence microscopy results on single carbon dots, particularly regarding excitation polarization dependence of emission. This story is still evolving.

Finally, Chapter 7 gives a brief recap of this thesis and a few comments on future directions for getting our photothermal technique to the single-molecule limit. Allons-y!

Can you see a single molecule?

What's the smallest object you can see?


Would you believe that you can see *a single molecule*? Not only would that be **really**

cool...studying single molecules would also be very useful for scientists, especially chemists. Traditional chemical measurements


look at many molecules, often billions of billions, at the same time. While this is good for some experiments, it can cause you to miss

out on a lot of information. This is because different copies of the same molecule can behave differently! An analogy is shown to the right. Up above, you see a confusing jumble of an image.

Below, you see six images of artwork by my favorite painter, Vincent van Gogh. They're so beautiful! The above image is actually an overlaid composite I made of many copies of the bottom images. It's hard to understand what's happening in the composite image on top! Similarly, it can be hard to know what's going on in chemistry if you look at a billion billion molecules at once, all doing their own thing at their own speed. Therein lies the reason why we want to look at single molecules.




Now that you're excited about looking at single molecules...how should we go about seeing them? Do you have any ideas? A water molecule, for example, is about 0.1 *nanometers*



across (that's one ten-billionth of a meter!). Just how small is that? Well, if we consider a drop of water, it will have about 1.6×10^{21} molecules within it (that's a billion trillion molecules!). If I had that


many US pennies, and gave you \$1 million in pennies **every second**, it would take well over **500,000 years** to give you all of my money! It turns out that molecules are just too small for your eyes to pick up on. So how in the world are we going to see single molecules? We'll explore two different ways to do just that in this chapter, both of which I've employed during my thesis work.

The first way to see single molecules is using a process called *fluorescence*, where we make molecules emit **light** which we then measure. While your eyes can't see something as small as a single molecule, they happen to be **very**



sensitive to light, and laboratory cameras can be made even more sensitive. I like to think about a fluorescence experiment like standing in a dark American football stadium, perhaps during halftime show at a Super Bowl game (where of course Kansas City is winning the game). Perhaps I'm at the game with my grandma (an aggressive Kansas City football fan), and am trying to see where she is in the crowd in the stands. This is impossible in the dark! However, if she takes out her cell phone, and takes a picture with the flash on, I'll easily be able to see the light from that camera flash, telling me where she is located. Similarly, if molecules emit light, we can use a microscope to easily see where those molecules are, and study that light to learn about molecular properties. In chapter 6 of my thesis, I describe my fluorescence experiments looking at a special

type of *nanoparticle* called “Carbon Dots”. I studied these dots to better learn how their chemical components give rise to their light-emitting properties, which have a variety of applications in medicine and other technologies such as television displays.



Fluorescence sounds great! Can we

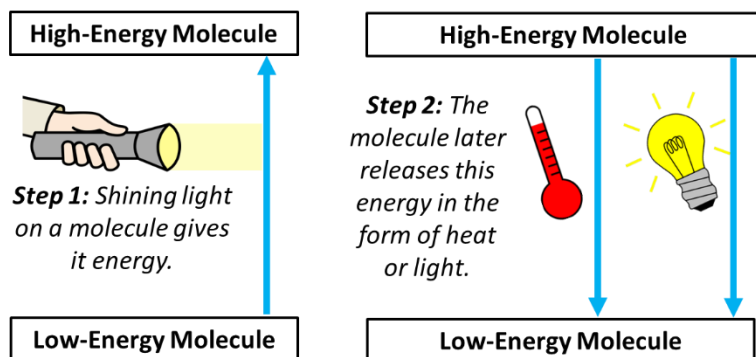
just use it for every single-molecule study that we want to do? Unfortunately not,

because not everything fluoresces. So what do we do in the scenario where our molecule won't emit light for us? What if instead, we shine a light on a molecule, and look for the shadow that the molecule casts? This is actually a common technique when looking at a collection of many molecules, but **is it possible to see the shadow of a single molecule?** Let's return to the stadium analogy...

Imagine you're back in the people-filled stadium, but this time, the stadium is lit brightly. Your job is to count the number of the people in the lit-up stadium over and over again, so that you know when a single person has left. If the stadium starts with 80,000



people, that means you have to be able to quickly tell the difference between 80,000 people and 79,999 people. That sounds pretty hard! This example is like doing the single-molecule shadow experiment described above...shining light at a molecule and trying to measure the teensy, tiny part of the light missing on the other side.




So, measuring the missing light from a single molecule's absorption might be harder than we want to deal with! What else could we measure? If we shine light on a molecule, let's think for a bit about

where the absorbed energy goes. If it doesn't get released as light (fluorescence), then the only other way for it to be released is as **heat!** And when heat is released from an object, that changes the temperature of the object's surroundings. Can we measure that temperature change using a thermometer? The heat output from a single molecule is going to be quite small...so we'll need a **very** sensitive thermometer to detect the temperature change from a single molecule after it absorbs light! Experiments that measure temperature changes from the absorption of light are called *photothermal* experiments.

In many of my experiments, I use a tool called an *optical microresonator* to make such photothermal measurements! "Optical" means light, and "micro" means small (a micrometer is 1 millionth of a meter, or about 1/100th the width of a human hair). "Resonator" means that



certain wavelengths, or **colors** of light "resonate" within the structure. You're probably already familiar with a similar kind of resonance, an *acoustic* resonance. Have you ever heard a trombone played? When a trombonist plays, the instrument emits a specific note based on the length of the instrument. By changing the trombone slide length, the note that resonates also changes! This is because for a specific trombone length, only certain *wavelengths* (notes) of sound will build up



in volume (get louder). For our photothermal experiments, we study *optical* resonances, meaning that rather than the volume of sound building up, the brightness of light builds up! My resonator is a hollow


glass bubble around 50 micrometers across. I trap light in the resonator's walls, similar to how light travels through a fiber optic cable. Light can then circulate around the resonator many times, but only very specific wavelengths (colors) of light will resonate.

Now we know how resonators work...but how do we use them to measure temperature changes? It turns out when a material changes temperature, that also changes how fast light travels through it! You're likely already familiar with this fact. Have you ever



noticed that the asphalt of a road, and the air above it, looks strange on a hot day? That's because the heat from the asphalt is changing the air's density above it, and thus the speed of light through the air. That hot air then bends sunlight differently, making pavement look shiny, and the air above it blurry. In the case of my microresonators, temperature increase changes the resonance wavelength! Because the light travels around our resonator many, many times, we can measure the change in the light's properties due to even a very tiny temperature change in the resonator.

In my experiments, I measure an optical resonance using a laser going through a fiber optic (the blue line crossing the picture below to the left). I also shine a second laser onto the resonator (the vertical green part of the image). When the second laser hits a molecule (shown as the gold oval), that molecule releases heat (shown by the red plume). By monitoring the resonance while heating up the molecule, I learn information about how the molecule absorbs light! For



example, I have employed this technology to study the chemical reactions of single gold nanorods (see Chapter 4 of this thesis for more!). This technique has also been used by my research group to study the physics of electrically conductive plastics, with applications in solar cell technology.

In conclusion, single molecule studies can be quite challenging, but they are worth all the effort! In my research, I have studied a variety of interesting nanoparticles and molecules by examining the energy that they release, either by looking for light they emit or by making very sensitive temperature measurements. **Thank you** for joining me in this fascinating journey to see single molecules, and please feel free to peruse the other chapters of my thesis if you wish!

The reason that I wrote this chapter is twofold:.

- (1) I believe that science literacy and appreciation is crucial in our society, and wanted to communicate my research (largely paid for by taxpayer dollars!) to the public at large. More specifically, I think that optics and nanophotonics have the potential for revolutionary impacts in computers, healthcare, climate monitoring, and other critical research areas in the coming years. Plus, lasers are cool. Who doesn't like lasers?!?*
- (2) I believe that it is important for scientists to practice communicating their work to the public, and to keep in mind how their research will impact society at large. Thus, this chapter was both a challenge to myself, and a suggestion to my colleagues to pursue similar challenges.*

I thank Elizabeth Reynolds for her editing and suggestions for this chapter. I thank Cayce Osborne for all of her support in my public engagement pursuits these past few years. And I deeply thank my friend and mentor Bassam Shakhshiri for instilling in me a deep sense of responsibility as a scientist-citizen.

Image sources: All graphics and components of graphics in this chapter not created by me are licensed under Creative Commons or are public domain, and were accessed from the following sources.

van Gogh, The Bedroom: <https://www.rawpixel.com/image/537441/the-bedroom-van-gogh>. Accessed Sept. 6th, 2022.

van Gogh, Starry Night: <https://www.rawpixel.com/image/537438/the-starry-night-van-gogh>. Accessed Sept. 6th, 2022.

van Gogh, The Potato Eaters: <https://www.rawpixel.com/image/3865203/illustration-image-art-vincent-van-gogh-people>. Accessed Sept. 6th, 2022.

van Gogh, Vase with Twelve Sunflowers: <https://www.rawpixel.com/image/3865867/illustration-image-flower-art-vincent-van-gogh>. Accessed Sept. 6th, 2022.

van Gogh, The Church at Avers: <https://www.rawpixel.com/image/3864730/illustration-image-art-vincent-van-gogh-person>. Accessed Sept. 6th, 2022.

van Gogh, Café Terrace at Night: <https://www.rawpixel.com/image/3864780/illustration-image-art-vincent-van-gogh>. Accessed Sept. 6th, 2022.

Magnifying glass: <http://www.clker.com/clipart-ssss-4>. Accessed Sept. 6th, 2022.

Water molecule: https://commons.wikimedia.org/wiki/File:Water_molecule_2.svg. Accessed Sept. 6th, 2022.

Water droplet: <http://www.clker.com/clipart-water-droplet-icon.html>. Accessed Sept. 6th, 2022.

Trombone: https://commons.wikimedia.org/wiki/File:Trombone_CG_Bach42AG.jpg. Accessed Sept. 6th, 2022.

Flashlight: <https://openclipart.org/detail/326463/flashlight-2>. Accessed Sept. 6th, 2022.

Thermometer: <https://openclipart.org/detail/218124/thermometer>. Accessed Sept. 6th, 2022.

Light bulb: <https://openclipart.org/detail/48397/light-bulb-on>. Accessed Sept. 6th, 2022.

Mirage: https://commons.wikimedia.org/wiki/File:Hot_road_mirage.jpg. Accessed Sept. 6th, 2022.

Light stadium: [https://commons.wikimedia.org/wiki/File:Damien_William_Run_\(49607462502\).jpg](https://commons.wikimedia.org/wiki/File:Damien_William_Run_(49607462502).jpg). Accessed Sept. 6th, 2022.

Dark stadium:

https://commons.wikimedia.org/wiki/File:BTS_in_concert_at_Wembley_Stadium,_2_June_2019_01.jpg. Accessed Sept. 6th, 2022.

2. Scientific background for this thesis

2.1. Introduction

2.1.1. Recommended reading

While any doctoral thesis can be expected to be filled with a significant quantity of technical language, the world of nanophotonics,^I while burgeoning, is still niche enough that there is an extra thick layer of technical understanding necessary to discuss the minutiae of the experiments within this thesis. I will not attempt to give an exhaustive technical background in this Chapter, but rather provide useful information that will assist in understanding the most important technical areas of this thesis. And while throughout the thesis I cite many individual papers and a number of reviews for many specific topics, there are also a few textbooks out there that provide invaluable background for the content of this thesis.^{II}

The first is *Introduction to Modern Optics* by Fowles,¹ which provides a superb introduction on many optics topics. While most of the content is from 1968, and thus huge swaths of now-modern optics are not included, the book is still an excellent resource.^{III} The second is *Optics* by Hecht, which is another staple for learning and relearning optics topics, obtaining useful equations and relationships, etc.² The third textbook I will suggest is *Introduction to Nanophotonics* by Gaponenko.³ While this book will be less comprehensive than the first two on fundamentals of optics, and the breadth of topics covered, it will provide meaningful insights into how light interacts with nanoscale structures, and draw important connections to the quantum realm.^{IV} The fourth textbook I suggest having a copy accessible of is *Principles of Fluorescence Spectroscopy* by Lakowicz.⁴ This is only directly relevant to Chapter 6 of this thesis, though this treatise on fluorescence experiments is also a great resource on understanding

^I Probably the most relevant technical field for this thesis, though single molecule spectroscopy and a variety of other communities also are quite relevant.

^{II} Not that I am claiming to have read them all cover-to-cover!

^{III} It at least came out after lasers were invented, which means that the foundation of modern optics (or so I would argue) is well-included.

^{IV} Though pretty much all of the nanophotonics in this thesis can be understood from a solely-classical perspective.

fundamental light-matter interactions. There are also very useful online resources, such as the introduction to polarization from Edmund Optics.⁵ Other theses from our group are also excellent resources; for learning more on the fundamentals of light and light-matter interactions, I particularly suggest the theses of Dr. Morgan Rea and Dr. Erik Horak.⁶⁻⁷

I will lastly note on the topic of textbooks above that I have stuck to optics resources. If you, as Chapter 3 outlines, want to build your own microresonator photothermal spectrometer, then additional background in other areas such as electronic circuitry and computer programming may also be necessary.

2.1.2. Overview

This chapter will **not** delve deeply into physics derivations, or cover comprehensive mathematical treatments. Rather, I will point to suggested tools and frameworks for understanding, and hopefully intuiting, some of the complex topics involved in this thesis work, favoring qualitative understanding as much as quantitative understanding. This chapter also is written with the assumption of introductory optics and physics knowledge by the reader, and is largely provided to give some relevant technical background for Chapter 3. Thus, the reader may wish to instead move ahead to that chapter, which describes how to build a microresonator spectrometer from the ground up, and then reference back to this chapter when relevant for technical information.

With that in mind, there are three particularly important resonator parameters to know about in our photothermal experiments : Quality Factor, Mode Volume, and Thermo-optic Coefficient, which are discussed below, after we first explore briefly what an optical microresonator is in the first place. Then, we will discuss the sensitivity of our experiments, and how to define this sometimes-ambiguous parameter. Next, a short section on conservation of energy and momentum will fortify how we couple light into and out of our microresonators. Lastly, a substantial section will be devoted to the polarization of light, and useful tools on understanding in it the context of our experiments.

2.2. WGM Microresonators

Here I will give a very brief introduction to optical microresonators, specifically of the Whispering Gallery Mode (WGM) flavor.¹ If the reader wishes a more thorough primer on microresonators, previous theses from our research group are a great resource, as is this review,⁸ amongst others.

As described in in Chapter 1, optical microresonators have “resonances”, where certain wavelengths of light meet the condition for resonance. This condition is mathematically described by **Equation 2-1**, where λ_{res} is the resonance wavelength, R is the radius of the (here presumed circular) resonator, and n is the refractive index of the resonator. The parameter m can take the form of any positive integer.

Equation 2-1

$$\lambda_{\text{res}} = \frac{2\pi Rn}{m}$$

At these resonant frequencies, the **optical path length** (circumference multiplied by the refractive index¹¹) is equal to an **integer number of wavelengths**, which results in the build-up of power through **constructive interference** (**Figure 2-1**). At all other wavelengths, after enough trips around the resonator, light destructively interferes with itself, leading to zero propagating power in the resonator.

¹ I do not condone tasting of optical microresonators.

¹¹ I think of the optical path length as “the distance that light sees”.




Figure 2-1. Constructive interference in microresonators. As light travels around a microresonator multiple times, it will constructively interfere and build up power if it is at a resonant wavelength, but will destructively interfere and lose power if not at a resonant wavelength.

In our experiment, where we measure the light transmitted through an optical fiber which interacts with the resonator, we see our resonances as a **dip in transmission** at the resonant wavelength (Figure 2-2).¹




Figure 2-2. How a resonance looks in our experiment. (A) The probe laser travels through an optical fiber, which tapers as it goes near the WGM resonator. The probe then continues to an avalanche photodiode (APD), which measures the power. (B) As the probe wavelength is scanned, the detected light at the APD (and thus its output voltage) suddenly drops when the wavelength is resonant, as light now circulates around the resonator instead of propagating through the fiber. The resonance can be mathematically represented by a Lorentzian function.

It is important to note that reality tends to be more complicated than the above equation, which it turns out is rather naïve (though still often useful, and a great pedagogical tool). The parameter m ,

¹ One might wonder “but shouldn’t light couple back out of the resonator into the fiber?” The answer is yes. But due to the phase shift of coupling into and out of the resonator, this light destructively interferes with any light still propagating through the fiber. When coupling is optimized, a condition called “critical coupling” can be reached, where the dip depth is so deep that measured light is essentially zero at the APD. Previous theses in the group and other technical sources have overviewed this in-depth mathematically, so I shan’t recapitulate it here.

often called the **azimuthal mode number**, is only one of three such mode numbers. The other mode numbers¹ are assigned to higher-order modes that exist, some with electromagnetic energy going into the resonator (often called **radial** modes), and/or some with electromagnetic energy going above and below the plane in which the fundamental mode (where all mode numbers equal zero, i.e. a TE₀₀ or TM₀₀ mode). These higher order modes (see **Figure 4-19** for more intuition) will have nodes and antinodes in their electric field, whereas the fundamental mode will have only one local maximum of energy density. These higher-order modes also get even more complicated in resonators composed of multiple materials, such as water-filled microbubbles (**Figure 4-2**). The various modes a resonator will support can be calculated using COMSOL, a finite element method (FEM) software package, employing code borrowed from Oxborrow.⁹ These codes, explained elsewhere in this thesis and more extensively in other theses from our group, essentially solve Maxwell's equations for the propagation of light through a resonator of given geometry and material, resulting in eigenvalues (resonant frequencies) and eigenmodes (the spatial arrangement of the electric field at those frequencies).

2.3. Our Photothermal Measurement

As detailed elsewhere in this thesis, our photothermal measurement is in essence a **refractive index measurement**. This is because the refractive index, n , is in-fact temperature dependent. Thus, as we pump energy into molecules or nanoparticles, and they transduce that absorbed energy into heat, the refractive index of the microresonator changes, and thus so does the resonance frequency condition. In this section, I won't recapitulate the equipment involved in our photothermal measurement, as this is discussed in great detail in Chapter 3 and Chapter 4.

¹ Which feels somewhat analogous to quantum numbers, but notably, we are in a strictly classical regime in this current discussion. Dr. Erik Horak's thesis has some enlightening discussion comparing the mathematical forms of classical electromagnetic equations and quantum wave equations, at which I suggest the reader take a gander.

2.4. Quality Factor (Q)

The quality factor of a resonator is intuitively simple: a resonator of higher quality resonates better! If we go back to the trombone analogy from Chapter 1, a higher quality trombone could perhaps be expected to produce purer and louder notes. Similarly, an optical resonator with higher Q has purer resonances, i.e. resonances with a more specific wavelength (narrower linewidth), that also build up higher optical power. This leads to a very important series of three relationships: **High Q → Low loss → Narrow resonance linewidth.**

There are three components that are important to consider regarding quality factor: (1) **smoothness** of the resonator (does light scatter out due to **surface roughness**?), (2) **transparency** of the resonator (is light **absorbed** by the resonator **material**?), and (3) **size** of the resonator (are there **bending losses** due to forcing light to take to change direction too quickly?). These components are added in a reciprocal fashion to achieve the total quality factor. We will not explore these different components quantitatively here, but they are shown in **Equation 2-2**, where Q_{measured} is the measured Q-factor, $Q_{\text{materials}}$ represents the components due to materials absorption, Q_{surface} represents the components due to surface roughness, and $Q_{\text{curvature}}$ represents the components due to bending losses.

Equation 2-2

$$\frac{1}{Q_{\text{measured}}} = \frac{1}{Q_{\text{materials}}} + \frac{1}{Q_{\text{surface}}} + \frac{1}{Q_{\text{curvature}}}$$

As described above, a higher quality factor means a narrower linewidth (sharper resonance). In **Figure 2-3** are shown two cartoon resonances, one broader and one narrower, and hopefully an intuitive depiction of why resolving small shifts in the narrow resonance should be easier than resolving small shifts in the broader resonance.




Figure 2-3. Comparison of shift in a high-Q vs. a low-Q resonance. On the left, a high-Q resonance shifts (from black trace to grey trace). The same vertical point is shown by the two golden stars, one on the black trace and the other on the gray trace, indicating the large voltage shift observed when monitoring the APD at the starting resonance wavelength, after the resonance shifts. On the right, the same analysis for a low-Q resonance, where it is evident that the voltage change is much smaller than for the high-Q case, indicating that small resonance shifts are more difficult to detect with the low-Q resonance.

As discussed in multiple sections of Chapter 3, what we actually track in our experiments is the **Error Slope**, a phase-sensitive measurement that provides information on the **direction** of resonance shift as well as its magnitude. We won't delve into this further here, except to say that a **higher quality factor = steeper error slope**, and a steeper error slope gives us more sensitivity (more on this term below) for resolving resonance shifts.

2.5. Mode Volume (V)

The “mode volume” of a resonator is perhaps even more intuitive than the quality factor, as V simply describes **the volume that light occupies in a resonator**. For our experiments, we want to **minimize** mode volume. This can be understood straightforwardly as wanting to **maximize the overlap** between the thermal plume (from an analyte absorbing pump beam light and emitting heat) and the optical mode, as shown in cartoon form in **Figure 2-4**. Because the resonance shift condition is based on the **average refractive index** of the entire resonator, shrinking the resonator means that less of the optical mode is outside of the thermal plume, and thus the same heat leads to more resonance shift.




Figure 2-4. The photothermal impact of large vs. small mode volume. Two mode volumes (blue toruses) are shown. The purple sphere represents a single molecule or nanoparticle, and the red fuzzy region the heat plume output by the analyte when pumped with light. It is evident that more of the mode volume in the “Small V” case is overlapping with the thermal plume, meaning that the average refractive index will change more, meaning that the resonance shift will be more prominent.

There are actually a variety of different equations in the literature for calculating the mode volume of optical resonances, and in practice, you will have to choose a given formula to calculate the mode volume for a certain mode (ex. in a COMSOL simulation). I will not detail these various calculations here, but wanted to make the fact that multiple options exist to the reader. For a reasonable approximation for a **fundamental** mode, one might measure the full-width half-max (FWHM) of an optical mode generated in COMSOL and calculate the area of the circle (or FWHM on major and minor axes for an ellipse area calculation), and then multiply that by the circumference of the resonator (if it is a WGM resonator) to obtain the volume of the torus that the mode fills. For higher-order modes, it is prudent to have COMSOL integrate the over the mode to calculate mode volume for you.

2.6. Thermo-Optic Coefficient (dn/dT)

The thermo-optic coefficient, colloquially referred to as “ $d n d t$ ”, and mathematically shown as dn/dT , quantifies how the refractive index of a given material changes as a function of the material’s temperature. This parameter may itself be a function of temperature, however given the very small temperature changes we measure in our experiments we can readily approximate dn/dT as constant for a given material. Because our method is fundamentally a refractive index change measurement, the dn/dT is a crucial parameter for forming a quantitative relationship between resonance shift and heat

output (we do this via COMSOL simulations, see *Section 4.13*). Specifically for microbubble resonators (which are a central component of Chapters 3, 4, and 5), the fact that optical modes can reside both in the glass of the resonator and the solvent within means that dn/dT is even more important, particularly on a per-pixel value in our thermo-optic simulations.

2.7. Sensitivity, Responsivity, and Limit-of-Detection (LOD)

The **sensitivity** of our photothermal microresonator measurements increases with higher Q-factor, and increases with smaller mode volume. Given the above discussions, this should make good sense, as higher Q means that light is confined better temporally, and higher V means light is confined better spatially, both resulting in enhanced light-matter interactions. **But what does “sensitivity” mean?** To explain, I want to also introduce two other terms: **responsivity** and **limit-of-detection (LOD)**. I define these three concepts in my own words as follows, with some important descriptors in [brackets]:

- (1) [Voltage] Sensitivity:** The word sensitivity for our experiment is best thought of as the how **our readout scales with resonance shift**.¹ In our case, this means the DC output voltage from the lock-in amplifier. When our technique has been made more sensitive, this has been because our ability to resolve resonance shifts has improved, typically by being able to measure a **smaller voltage change**. For example, when PDH was introduced, this permitted Lock-in amplification,¹⁰ which through noise rejection improved the limit-of-detection of our experiments by multiple orders of magnitude compared to the previous method of resonance scanning to monitor shifts. Typically, we say that our **sensitivity scales as Quality factor divided by mode volume (Equation 2-3)**.

¹ In my view, at least. The reader should certainly think for themselves if this concept, and the following concepts, can be better understood in another way!

Equation 2-3

$$\text{Sensitivity} \propto \frac{Q}{V}$$

(2) [Thermal] Responsivity: Another useful term is responsivity. For our experiments, I define responsivity as thermal responsivity, which **indicates how much a resonance shifts for a given thermal input**. For example, the water modes of microbubbles are more responsive than the glass modes of microbubbles, due to the higher dn/dT , which results in more resonance shift for the same heat output. The fact that dn/dT is a function of temperature and not heat directly means that properties such as **heat capacity** (which determines the temperature increase for a given thermal energy addition), and **heat conductivity** (which determines how fast heat will be wicked towards or away from the mode-containing portion of the resonator) play important roles in thermal responsivity. What I've described for sensitivity above may also seem like a responsivity of sorts. I'm not sure the two terms can be defined in a way that fully sets them apart in every context, but from my perspective, the divorce of terms becomes easier when considering responsivity to *thermal* effects, specifically.

(3) Limit-of-Detection [of Resonance Shift, or of Absorption Cross-Section?]: This parameter is in fact what we've typically meant when we talk about the "exquisite sensitivity" of our microresonator spectrometers, as this parameter is usually given as the **smallest resonance shift** we can detect. However, for the purpose of our experiments, what we care the *most* about is the **smallest object** we can detect (ex. a single dye molecules). This conversion requires other parameters discussed elsewhere in this thesis, but it is worth noting that if smallest object is being examined as the limit of detection, then thermal responsivity also matters, whereas if it is only a change in voltage being considered, the sensitivity should be sufficient.

For photothermal microresonator experiments, I think it is important to move towards a common language with how “sensitivity” is defined. In particular, the one communicating such work should consider the audience, as the connotations of these various words may vary in specificity based on the listener’s scientific background.

2.8. Energy Matching and Phase Matching/Conservation of Momentum

A critical aspect to success with our current microresonator platforms (microtoroids and microbubbles) is the coupling of light to and from the resonator via a **tapered optical fiber**, the desirable characteristics and fabrication of which are discussed in *Section 3.9*. Here, I instead focus on two conditions that govern the efficiency of this coupling, namely the matching of **energy** and of **phase**. The polarization state of the light is also important, as will be discussed later in this chapter, and will be skipped over here.

Due to the constraint of **conservation of energy**, the energy of the light to be coupled into a resonator must match that of the optical mode being coupled to. Equivalently, this means that the **frequency** of light must be matched to the **eigenvalue** of a particular optical mode! This is in practice done using a **tunable laser**, where wavelength of the laser is tuned until a resonance is observed. Thus, it is easiest to actually think of this matching condition as a wavelength condition, but it is good to recall that the frequency of light is what stays constant, while the wavelength will change as the material light is propagating within changes. However, usually the free-space wavelength (λ_0) is the one used in relevant-to-us equations (whether explicitly stated or not. Otherwise, refractive index gets brushed under the rug within the new wavelength variable in some situations). The tunable laser method is particularly salient for WGM resonators. In the case of a resonator which has easily tunable dimensions, such as a fiber Fabry-Pérot cavity (FFPC), a single-wavelength diode can be used (though in either case, the laser linewidth must be narrower than the cavity linewidth, or the resonance linewidth will be compromised!).

Energy matching is not the only requirement for successful mode coupling, as the **phase** of the light must also be matched to the desired mode. This has been discussed previously in the literature specifically for the case of a fiber taper and a microsphere,¹¹ thus we won't recapitulate the math here. What is worth knowing, and what will probably make phase matching more intuitive, is that arises **due to the conservation of momentum**. Mathematically, this condition is met when the **propagation constant** (often known alternatively as the **wavenumber, k**) of the light in the tapered fiber matches the propagation constant of the optical mode in the resonator. In practice, phase matching is achieved by either **changing the taper thickness** (for microbubbles, translating a taper along the bubble to which it is stuck will do), or **changing the taper-resonator distance** (as we do with microtoroids, where we do not stick the taper to the resonator).

The reader may also notice the term "phase matching" occurring in a variety of other places in optics and nanophotonics literature. In such cases, the term is also referring to a momentum conservation condition, often very literally corresponding to the **angle of light** (this is actually quite common in materials science as well, to interchange angle and momentum!). These uses of phase matching often are brought up in non-linear optical contexts such as for sum frequency generation (SFG), difference frequency generation (DFG), etc.

2.9. Polarization

2.9.1. Why write (or read) a long section on polarization?

Control and/or measurement of the polarization state of light is a critical component of many optical measurements, and our photothermal microresonator experiments are no exception. In fact, we must exert precise control over the polarization of both the probe beam and the pump beam, each of which is discussed in great detail in Chapter 3. Despite the importance of polarization, the topic has been a challenging one that I have continually striven to improve my understanding of during the past few years. There are so many ways that polarization is described! S-polarization vs. P-polarization. Transverse

Electric vs. Transverse Magnetic polarization. Linear vs. circular vs. elliptical polarization. How does one fit all of these concepts into a single framework, either intuitively or mathematically?

Even if they are not familiar with optics, the reader may be somewhat familiar with polarization already, for example in the case of Circular Dichroism, which is used in a popular family of methods for studying the chirality of chemical and biological samples. Someone with a moderate amount of optics experience likely already knows that a half-wave plate (HWP) will rotate the angle of linearly polarized light (LPL), and that a quarter-wave plate (QWP) can be used to transform LPL to circularly polarized light (CPL) and vice-versa. The reader may also already understand that these optics work through **birefringence**, where along one axis, the electromagnetic wave (light) propagation gets retarded by a certain amount (half a wave or quarter a wave in the above examples) relative to the electromagnetic wave component along the orthogonal axis of the optic.

However, life often isn't that simple! For example, during my PhD, all of the following questions have arisen, often in multiple contexts:

- What happens when the polarization input isn't as friendly (LPL or CPL) as above?
- How does the rotational alignment of our optics impact the transformation?
- What if the optics aren't quite perfect, ex. not exactly a true QWP?
- What if we want to employ birefringence at other retardance values, not quarter- or half-wave?
- What happens when we start putting multiple polarization optics in the path of our laser beam?

While the optics books I suggested towards the beginning of this chapter for further technical reading are great, not one of those has given me a comprehensive understanding of even just the above polarization terms all together yet.¹ In fact, the Wikipedia page¹² on Polarization is probably the most

¹ Perhaps the challenge in holistically understanding polarization (at least for me) comes from the way different polarization concepts are taught in different contexts, ex. RHCPL and LHCPL, ex. transforming LPL to CPL, etc.

comprehensive source I have seen yet (albeit without the practice problems that a textbook includes to help strengthen understanding),ⁱ perhaps because a computer-displayed format is very helpful for conveying intuitive information, ex. through GIFs explaining CPL.

Given the above, having some additional tools in our kit would be quite helpful in understanding the polarization of light, thereby allowing us to improve existing experiments and perhaps even conceive of new ones. Thus is the motivation for this section of my thesis. Here, I try not to give a comprehensive explanation of polarization,ⁱⁱ but will share some insights and specific tools that I have found useful for understanding polarization in the contexts most relevant in my experiments, particularly photothermal microresonator spectroscopy. First, we will discuss some useful terminology regarding the polarization of light. Then, we will learn about two tools that are particularly useful: (1) the Poincaré Sphere, which I suggest as a helpful visual aid (though it can also be a mathematical one), and (2) Jones Matrices, which I suggest as a go-to mathematical tool for modeling the impact of polarization optics on experiments.

2.9.2. Fundamentals of polarization

First, the most fundamental question: what *is* polarization? Polarization describes the orientation of the **oscillations** in a **transverse wave**. For example, when a string on a ukulele is plucked, that string will vibrateⁱⁱⁱ along a certain axis, the polarization axis, orthogonal to the length of the string. Returning to optics, light is an electromagnetic wave, composed of an electric field that is perpendicular to a magnetic field. Each field oscillates in a direction orthogonal to the propagation of light, and by convention, the direction of **electric** field oscillation is defined as the direction of polarization.

Importantly in this section we are talking entirely about **polarized light**; light can also come in an unpolarized state (ex. sunlight) where the polarization wanders randomly over time. However, given that

ⁱ And I suggest that the reader give this Wikipedia page a readthrough, perhaps once a year, to gain new insights and fortify previous ones.

ⁱⁱ As valuable as this would be, it is outside the scope of this thesis.

ⁱⁱⁱ And that vibration will result in air-vibrations that make sounds. Because sound waves are longitudinal waves (oscillate along the direction of propagation), they do not have polarization (sad).

all experiments in this thesis employ lasers (as do most optical experiments nowadays), we will keep this discussion almost entirely to polarized light from coherent sources. First, let us consider five situations qualitatively:

1. In the case that the direction of electric field oscillation **stays constant** as light propagates, the light is said to be **linearly polarized**.
2. In the case where the above is true, but the oscillation is in an orthogonal direction to the original case, the light is also said to be linearly polarized.
3. In the case there are electric field components at two orthogonal directions (i.e. a summation of situations #1 and #2), and the waves are perfectly **in phase**, this still describes linearly polarized light, but at an angle in-between the orthogonal directions we started with.
4. If there are two such components as in case #3, but the components are **not in phase** with each other, this describes **elliptically polarized** light (EPL).
5. Lastly, in the special case of elliptical polarization where the phase difference between orthogonal electric field components is exactly 90-degrees (one quarter wavelength), we call this **circularly polarized** light, which, depending on the sign of the phase difference, will be either right-handed (RHCPL) or left-handed (LHCPL).

Fowels¹ goes through a useful derivation for arriving at mathematical descriptions of light with different polarizations. While we will not explore this in full here, we will take a look at the resulting equations, where a **complex vector amplitude** \mathbf{E}_0 is defined by **Equation 2-4**, where E_0 and E_0' are the real amplitudes of electric field in two orthogonal directions.

Equation 2-4

$$\mathbf{E}_0 = \hat{i}E_0 + i\hat{j}E_0'$$

We then employ the above in the **wave equation** shown in **Equation 2-5**, where **t** is time, **k** is the wavenumber, **ω** is the frequency of light, and **z** is distance.

Equation 2-5

$$\mathbf{E} = \mathbf{E}_0 e^{i(kz - \omega t)}$$

This compact equation is quite useful, as it can represent any type of polarization, for example in the three following cases:

- (1) *LPL*: \mathbf{E}_0 is **real**.
- (2) *EPL*: \mathbf{E}_0 is **complex**.
- (3) *CPL*: \mathbf{E}_0 is **complex**, and the real and imaginary parts are **equal**. RHCPL and LHCPL are distinguished based on the sign in the equation for the complex value for \mathbf{E}_0 .

2.9.3. Linear polarization designations: S & P, TE & TM

When dealing with linear polarization, there are two common ways¹ to see it described, which are particularly useful in the context of light matter interactions where boundary conditions arise, such as light interacting with a surface (ex. bouncing off of a mirror), or light propagating through a waveguide (ex. traveling around inside a microresonator).

The first description uses two terms: **S-polarization** and **P-polarization**. These terms, as shown in **Figure 2-5**, arise when considering the **plane of incidence** of light on a surface, which is defined by the vector along which the light is traveling, and the vector perpendicular to the surface. Electric field components *parallel* to the plane are dubbed P-polarized, while components *perpendicular* to the plane are termed S-polarized (S from Senkrecht¹¹). S and P descriptors are particularly useful in circumstances that determine how much light will be reflected vs. transmitted at a surface (using Fresnel equations),

¹ Actually, there are other sets of terms as well, but these two sets will be most if not all contexts polarization is described in for experiments such as ours.

¹¹ Which is, of course, German for perpendicular. To be fair, both parallel and perpendicular start with the same letter, so something had to be done.

such as that of a wire grid polarizer. The terminology is also useful when considering how polarization is **scrambled** as it interacts with optics, especially mirrors, at certain angles.




Figure 2-5. How S- and P-polarized light interact with a surface. Incoming mixed-polarized light meets a surface. S-polarized (TE) light is reflected, and P-polarized (TM) light is transmitted. This figure is inspired by a figure on the Edmund Optics website.⁵

The second common description for linearly polarized light also uses two terms: **transverse-electric (TE)** and **transverse-magnetic (TM)**, which are terms quite commonly used in the nanophotonics community! Luckily, we do not have to learn a new model, as **TM is the same as P-polarized**, and **TE is the same as S-polarized** (also included in **Figure 2-5**) as these new terms simply described which component of the electromagnetic field is oscillating transverse to a surface.

In the case of light propagating in free-space, light is described by **transverse electric magnetic (TEM)** modes, as both the electric and magnetic fields are orthogonal to the direction of propagation, and there is no surface to reference, as in our previous discussion. However, things get a bit more complicated in the context of, for example, a microresonator. Now, there is a surface interface to consider, the same very one that contains the light via total internal reflection! When Maxwell's equations are solved for the propagation of light through the resonator (in our case, using code from Oxborrow⁹ in COMSOL), eigenvalues are found for both TE and TM modes, where for the TE modes, the electric field vector points parallel with the surface of the resonator, as shown in **Figure 2-6**. In our

experiments, we use polarization paddles to alter the **probe beam polarization**, which is an important factor when maximizing coupling efficiency to a given resonance.




Figure 2-6. A TE mode in a microbubble resonator. More specifically, the 2nd order radial and 0th order polar mode. The black arrow indicates the direction of electric field oscillation.

2.9.4. The Poincaré sphere

The Poincaré sphere is a helpful visual aid for understanding the relationship between different polarization states. Unfortunately, I have found this aid to be absent in many major optics textbooks such as the ones I suggested at the beginning of this chapter.¹ Thus, I want to explicitly draw the reader's attention to the Poincaré sphere, and discuss some potentially uses for it here.

First, we need to know about something called **Stokes parameters**, which form a **basis** (the minimum number of vectors that together define a parameter space, here any polarization). It is important to know that the Stokes parameters work for *any* light; the light need not be fully polarized. Stokes parameters were originally defined in a way (by Sir George Gabriel Stokes) that described how light would be impacted by four different optical filters. We will not discuss that definition of the parameters here, but rather jump right to a definition that is much more useful for our discussion

¹ And I'm not sure why. Perhaps folks writing and reading those books benefit more from the equations-approach than a graphic illustration, whereas I best learn when provided with both.

regarding the Poincaré sphere.¹ We define a **Stokes Vector**, comprising the four Stokes parameters, in **Equation 2-6**, where I is the intensity of light, P is the degree of polarization, and ψ and χ are geometric parameters.

Equation 2-6

$$\mathbf{S} = \begin{bmatrix} S_0 \\ S_1 \\ S_2 \\ S_3 \end{bmatrix} = \begin{bmatrix} I \\ IP \cos(2\psi) \cos(2\chi) \\ IP \sin(2\psi) \cos(2\chi) \\ IP \sin(2\chi) \end{bmatrix}$$

For our treatments, we will consider that $P = 1$ (i.e. the light is fully polarized), and we will also divide the Stokes vector \mathbf{S} by S_0 (intensity) to yield a normalized stokes vector, \mathbf{S}' . The geometric parameters ψ and χ can be understood through the **polarization ellipse**, which is the ellipse that light would trace out over one optical cycle. This is not shown here, but it is worth knowing that ψ describes the orientation of the ellipse, and χ is the ellipticity angle.

So what do all these parameters have to do with a sphere? First, we will take the three Stokes parameters S_1 , S_2 , and S_3 , and set them as the **cartesian coordinates** (axes) in a 3-dimensional space (**Figure 2-7A**). If you inspect our earlier equations for the Stokes parameters, you may realize that we have also already provided ourselves **spherical coordinates**, where the angles are defined by the variables 2ψ and 2χ , and the radius of the sphere equals I for a non-normalized system, or just equals unity for our normalized, fully-polarized system (i.e. our sphere is a unit sphere for this special case). Now is where things get really useful! Shown in **Figure 2-7B** are some select polarization states for various coordinates on the sphere's surface. Look! The north pole of the sphere represents RHCPL, and the south pole represents LHCPL! And everything along the equator is linearly polarized! This is why I love the Poincaré sphere so much: it shows all possible polarization states in one figure, and with some math, relates them to intuitive geometric parameters of the polarization ellipse. That's pretty neat!

¹ The same definitions used in the Wikipedia article on polarization, in fact.




Figure 2-7. The Poincaré Sphere. To the left is shown the 3D cartesian space generated by the Stokes parameters. To the right is shown a simplistic version of the Poincaré sphere, with a few different polarization states shown by arrows.

A particularly useful fact about the sphere is that any two points opposite on the sphere (antipodal) are *orthogonal* polarization states (mathematically meaning that their inner product is equal to zero). This is conceptually clear by considering two opposing points on the equator the sphere, as the two linear polarizations found will in fact be perpendicular lines. Another pair that may be intuitive is considering the north pole of the sphere and the south pole, which are CPL of opposite handedness. In a variety of contexts in chemistry, anatomy, etc. we may consider “right hand” and “left hand” as opposites, therefore it makes sense that the same holds true here as well. This truth of orthogonal states holds for **any** two antipodal points on the sphere!

Lastly,¹ the sphere can be used as a visual aid to understand some polarization optics. For example, imagine that linearly polarized light is incident upon an HWP, and that HWP is rotated. As the HWP is rotated, the polarization state will circumnavigate the equator of the sphere! Unfortunately, other optics such as QWPs¹³ and Liquid Crystals¹⁴ are less intuitive and require closer inspection. Still,

¹ As one final note, if the reader has found themselves captivated by The Sphere, consider searching online for 5-dimensional variants of the sphere, which have been applied in contexts such as optical vortex beams and topological photonics.

mapping out the journey of polarization state on a sphere as a given optic is rotated (or voltage-adjusted) can be worthwhile.

2.9.5. Jones matrices

Another useful mathematical tool for understanding the polarization of light, and for tracking the polarization state of a laser beam as it passes through optics, is the “Jones calculus”. I personally don’t like this term, as what we’ll discuss below doesn’t require the use of any derivatives or integrals. Since we *will* be heavily employing matrices, I will primarily use the term “Jones matrices” in this thesis (which can be found in use elsewhere in the literature as well). It is important to know that Jones matrices are only applicably to **fully polarized** light (in contrast to the Stokes formalism we explored above). If one wants to apply the following treatment to partially polarized light, they should use Mueller matrices (which employ 4 dimensional vectors and 4x4 matrices); these not be discussed further here.

To start understanding Jones matrices, we must first consider a beam of polarized light, where the polarization state is entirely described by a Jones vector, defined in **Equation 2-7**, where E_{0x} describes the amplitude of the x-component of the electric field (and similar for the y-component), and the two ϕ variables describe the relative phase of the two fields.

$$\text{Equation 2-7} \quad \mathbf{E}_t = \begin{bmatrix} E_{0x}e^{i\phi_x} \\ E_{0y}e^{i\phi_y} \end{bmatrix}$$

It is common practice, and quite convenient, to normalize the amplitude coefficient. After doing this, then a simple consideration of the relative amplitudes of the components, as well as their phases, should lead to some fairly intuitive Jones vectors. In **Table 2-1** are shown the vectors for horizontal LPL, LPL at 45-degrees, and RHCPL. **What would the vector for vertically polarized light look like? For LHCPL?** I will leave these and other Jones vectors for the reader to figure out, should they wish.

Table 2-1. The Jones vectors associated with example polarization states of light.

Polarization State	Jones Vector
Horizontal LPL	$\begin{bmatrix} 1 \\ 0 \end{bmatrix}$
45-Degree LPL	$\frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \end{bmatrix}$
RHCPL	$\frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ -i \end{bmatrix}$

It seems that we are already off to a promising start, being able to describe any state of polarization with a fairly simplistic vector! Now, we can use 2x2 **matrices**, which **represent the impact of polarization optics** on a given incoming polarization state! A select few such matrices are shown in **Table 2-2** (longer lists can easily be found elsewhere). First is shown the very simple matrix representing a linear polarizer with its transmission axis set to horizontal. Second is shown the generic Jones matrix for a linear polarizer, where θ is the angle of the transmission axis relative to the horizontal. A good check for the reader here would be to prove that the latter matrix gives rise to the former, given a horizontally aligned transmission axis!

Polarizers are all good and well, but we would really benefit from modeling birefringent optics as well, such as QWPs. And in fact, we can! The third entry for the table is for a QWP with its fast axis oriented vertically. Notice the introduction of the imaginary number i in both the matrix and in the coefficient out front. The last entry in the table is for an **arbitrary** birefringent material (i.e. an elliptical phase retarder). This matrix can be used to represent **any polarization transformation** (i.e. it can be reduced to all other matrices with the proper parameters). For this matrix, η is the relative phase retardation between the fast and slow axis (see that in the case of a QWP, this number was pi divided by 2, i.e. a quarter wave, as it should be!). θ is the angle of the fast axis with respect to the horizontal. Lastly,

ϕ is the **circularity** of the retarder. For linear retarders (QWPs, HWPs, variable liquid crystals), this parameter is equal to zero. I still include here this most general form, as it is then useful for circular retarders and elliptical retarders if needed.

Table 2-2. The Jones matrices for a few example polarization optics.

Optic	Jones Matrix
Linear Polarizer set to Horizontal	$\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$
Linear Polarizer at any angle	$\begin{bmatrix} \cos^2(\theta) & \cos(\theta)\sin(\theta) \\ \cos(\theta)\sin(\theta) & \sin^2(\theta) \end{bmatrix}$
QWP set to vertical	$e^{\frac{i\pi}{4}} \begin{bmatrix} 1 & 0 \\ 0 & -i \end{bmatrix}$
Arbitrary phase retarder	$e^{-\frac{i\eta}{2}} \begin{bmatrix} \cos^2(\theta) + e^{i\eta} \sin^2(\theta) & (1 - e^{i\eta})e^{-i\phi} \cos(\theta)\sin(\theta) \\ (1 - e^{i\eta})e^{i\phi} \cos(\theta)\sin(\theta) & \sin^2(\theta) + e^{i\eta} \cos^2(\theta) \end{bmatrix}$

While some of these Jones matrices may look a bit messy, there are quite convenient, especially for modeling the impact on polarization. And by using traditional matrix multiplication, multiple matrices can be strung together to represent stacks of polarization optics. Note that the matrices **do not commute**. Just as with standard matrix multiplication, you must multiply the matrices properly in order by the starting Jones vector to obtain a correct output. This may seem a bit familiar, as in an intro optics course, one will have used **ray transfer matrices** (also known as ABCD matrices) to calculate how a beam will be shaped by optics such as lenses. Indeed, I like to think of Jones matrices as the polarization-

equivalent to the ray transfer approach. A final note I'll make on Jones matrices is that **nowhere do you see wavelength** in the above matrices. However, there is a variable for **retardation**, which for most birefringent optics is a function of wavelength.

A couple of examples for how I have used Jones matrices, personally:

1. Determining alignment tolerances for the variable liquid crystal rotator (which contains a linear polarizer, a variable LC, and a QWP) in our pump beam path, by diagnosing how the extinction ratio of the output light varies for misalignments and an imperfect achromatic QWP.
2. *In progress*: modeling the impact of polarization paddles in our probe beam path, with the goal of gaining rational insights that can help with writing feedback algorithms for motorized paddles, in order to provide a more stable probe polarization.

The various analysis and codes for the above can be found on group's internal data server.

Beyond the application of Jones matrices (and furthermore Mueller matrices, as briefly mentioned earlier) are even more advanced treatments such as the **coherence matrix**. This will not be discussed further here, but the reader can find more advanced discussions of matrix mathematics for polarization elsewhere.¹⁵

2.9.6. Measuring Polarization

If the reader wants to learn more about *measuring* polarization state, I point them to the sections of Chapter 3 on the liquid crystal polarization rotator we employ, which also discuss the use of an automated polarimeter, a very useful tool.

3. Building, Calibrating, and Operating a Microbubble Photothermal Spectrometer

3.1. Introduction

The microresonator photothermal absorption technique that I detail throughout this thesis is very powerful, and holds much promise for elucidating the chemical and photophysical dynamics of single molecules and nanoparticles. To my knowledge, however, there are only two microscopes in the world for using microresonators to do single-particle/molecule photothermal absorption spectroscopy. Both microscopes are in the Goldsmith group at UW-Madison. One microscope is optimized for NIR spectroscopy studies, and has mostly been used with toroidal resonators (though some recent work using microbubbles has also been done). The other microscope is optimized for studies with visible pump light, and has thus far been used exclusively with microbubbles, though it will expand to other resonator types in the future. I have two main goals in writing this chapter. The first is to provide an easily-accessible resource for my own group understanding, constructing, and using a microresonator photothermal spectrometer for chemically dynamic studies. The second goal is to lay a blueprint so that others outside of our group might build their own photothermal microresonator instruments, so-as to disseminate this technology and allow our technique to flourish in a variety of other studies.

Previous thesis in our group have done well in explaining various aspects of microresonator photothermal spectrometers,^{6-7, 16-18} and I will cite specific theses where I suggest further reading. Here, I provide a detailed roadmap on how to construct such a microscope's many parts, calibrate critical components, pilot the microscope using mostly custom software, and meaningfully interpret the data produced by the microscope. In particular, I will focus here on the construction of a visible-wavelength microresonator absorption spectrometer for use in solution-exposed experiments. A briefer description of the microbubble spectrometer can be found in Chapter 4.

While I provide tips and tricks in a variety of places throughout this chapter, the fact is that a variety of skillsets that are necessary for constructing a microresonator spectrometer. You and/or your colleagues must be able to align lasers, understand polarization control, machine custom components out of various materials, work with custom electronics, and write custom code. Such a custom instrument typically requires cobbling together a dozen or more pieces of equipment from perhaps almost as many manufacturers, and perhaps spanning multiple decades in manufacture date. In particular, getting custom code to employ drivers for all these pieces of equipment for simultaneous communication can be quite challenging. With all of that in mind, if you are up to the challenge¹, let us begin! We will first explore the anatomy of our experiment from a birds-eye view, and later sections of this chapter will explore the spectrometer components in much greater detail. **At the end of this chapter (Section 3.14) are several pictures of the photothermal microscope**, some that may have utility, and some that are just pretty.

3.2. Overview: The anatomy of our photothermal microscope

Ours is a **two-beam experiment**. The first beam we refer to as the **probe beam**, as we use it to interrogate the microresonator on its local environment via probe-resonator interactions. The probe beam is, at least in our current designs, completely contained within optical fiber components. The second beam we refer to as the **pump beam**, which is used to input energy into our system of interest (ex. a single gold nanorod). Some or all of that energy is then dissipated as heat. This raises the temperature of the microresonator, which impacts its optical properties. That change in optical properties is then encoded onto our probe beam, which is analyzed to understand the light-matter interactions of the pump beam and analyte. The pump beam, at least in our current designs, is entirely or almost-entirely free-space. Further below, we delve into detail about how both of these beams, probe and pump, accomplish their assigned tasks.

¹ I believe in you!

If you read a paper from our group, you may see statements referencing to the “exquisite sensitivity” of microresonator spectrometers, often due to what we call a “double-modulation scheme”. This double-modulation refers to the two kinds of modulation we employ in our experiments (though early versions of our microresonator spectrometer did not use these modulations), namely **amplitude modulation** and **phase modulation**. The amplitude modulation is applied to our pump beam using an optical chopper, which effectively turns on and off our pump beam at a pre-determined rate, at which we’re able to later isolate our desired photothermal signal with powerful sensitivity. The phase modulation is applied to our probe beam, and allows for staying on-resonance constantly by providing real-time feedback to the probe laser. Both types of modulation will be discussed at length later in this chapter.

There are essentially three tools we employ, all related to the pump beam, to find single nanoparticles or molecules, confirm their identity, and study their dynamics. These diagnostics are (1) spatial control of the pump beam, (2) control of its polarization state, and (3) control of its wavelength. I won’t include a repeat figure here (see **Figure 4-2** for a good diagram), but will still detail each of these controls a bit further below.

- (1) **Spatial:** By scanning the position of the focused pump beam across the resonator surface, the location of nanoscopic objects is confirmed. The result of this is called a “photothermal map”, where the hot spots indicate single objects of initial interest.
- (2) **Polarization:** Next, we decide which objects are worth studying further, and which objects are dust of unknown origin to be subsequently ignored. To do this, we rotate linearly polarized light. When the linear polarization is parallel to the excitation transition dipole (Ex-TD), maximum light will be absorbed, leading to maximum heat output, and thus maximum resonance shift. In contrast, when the polarization is perpendicular to the Ex-TD, no light will be absorbed, and thus no resonance shift will be observed. Thus, resonance shift as a function of pump polarization can

diagnose single molecules or single nanoparticles, provided they have an Ex-TD, which is true for the spectral absorption features of many interesting targets, such as anisotropic nanoparticles (ex. gold nanorods) and single dye molecules.

- (3) **Spectral:** The third useful tool for object identification (and tracking dynamics) is the spectral response an analyte. If one sweeps the pump and monitors the resonance shift, they will arrive at a **single-particle/molecule absorption spectrum!** While we've previously acquired such absorption spectra in the NIR, ex. for gold nanorods and for conductive polymers,^{10, 19} our work published thus-far in the visible has only employed discrete diode wavelengths, albeit sometimes multiple wavelengths for a single nanoparticle. As described in *Section 3.6*, we now have the ability to pump with variable-wavelength light, which will enable many future scientific pursuits.

3.3. Setting up the Probe Beam




Figure 3-1. A simplified diagram of the probe beam path. WGM = Whispering Gallery Mode, APD = Avalanche Photodiode, Mod. = Modulator, Pol. = Polarization, Atten. = Attenuation.

3.3.1. Laser choice, optical path and components

As mentioned above, the probe beam is entirely within fiber optical components. **Figure 3-1** shows a simplified diagram containing the essential components of the probe path. Working with fiber optics requires certain care be taken to avoid damage, for example to not scratch the connectors of patchcords. Extra care must also be taken for using the correct types of patchcords, mating sleeves, etc., for example to maintain consistent beam polarization and power through the probe path. The reader

should familiarize themselves with proper techniques on working with fiber optics prior to setting up the probe path!

In all iterations of our photothermal experiments thus far, we have used lasers from the product line Velocity™ TLB-6700 (sold by Newport) as our probe laser. These external cavity diode lasers (ECDLs) provide an ultranarrow linewidth output (<200 kHz for a 50 ms integration time), substantial mode-hop-free tunability, excellent wavelength and power stability, and an optional and convenient fiber-coupled output. While we have had some issues in maintaining good polarization control with these lasers in the visible regime (see *Section 3.3.2*), overall this line of lasers has proved very useful for working with ultra-high-Q microresonators.

All of our toroid work and some of our microbubble work has used a probe beam that tunes from 1520-1570 nm. Almost exclusively, the wavelength band around 1550 nm is the choice for fiber optic transmission of telecommunications and internet across the globe, due to the fact that silica glass absorbs the least amount of light at this wavelength. This minimized absorption ensures maximum transmission distance before needing to boost a signal in fiber optics, and concomitantly, using 1550 nm light ensures minimum absorption of probe beam by the resonator, helping attain ultra-high Q-factors. Another very convenient aspect of working at telecom. wavelengths is the economy-of-scale for optical components, driving prices often one to two orders of magnitude cheaper than for visible-wavelength analogs of such fiber-based devices (ex. splitters, isolators, circulators, attenuators)¹.

However, as soon as one wants to start doing aqueous experiments, working in the NIR becomes a big problem, as any overlap between water and the probe-beam will result in absorption of probe of light and deterioration of Q-factor. Thus, for our water-exposed microbubble work, we transitioned to using a probe beam tunable from 765-781 nm.²⁰ More recently, I have transitioned the microbubble

¹ Fiber Store (fs.com) is a website from which very cheap telecom wavelength fiber optic components can be purchased.

platform to a probe beam tunable from 635-638 nm, driven by the knowledge that water drops yet another order of magnitude in extinction coefficient, thus reducing any possible water-limitations on Q-factor. See *Chapter 5* on photothermal background studies for data using the 635 nm probe beam.

Future improvement: It is possible that going even bluer in probe wavelength could be fruitful, but in this context, only if water-absorption of probe light is a limiting factor. Notably, while there is a 407-412 nm tunable option in the same laser product line, this would make fabrication of single-mode tapered-fibers exceedingly difficult. Also, the increase of light by glass absorption might negate gains made from decreased water-absorption of light.

Once our photons are out of our probe laser itself, where do they go next? First, they may encounter one or more *fiber-based power splitters*. As various experiments in our lab use the Velocity lasers (and they're quite expensive), this is a convenient way to pipe light around the lab from one laser to multiple experiments. While only one experiment can actively tune the laser's wavelength, the splitters are rarely an issue as far as power is concerned. For photothermal experiments, we tend to use resonator input powers on the order of 1-10 uW so-as to avoid thermal nonlinearities. A good back-of-the-envelope calculation for our experiments is shown in Equation 3-1.

Equation 3-1

$$\text{Number of Round Trips} \approx \frac{\text{Quality Factor}}{100}$$

Thus, for a 10^7 resonances, the input power gets amplified by a factor of 100,000, so 10 uW becomes 1 W of circulating power. Going to higher powers can be useful for techniques such as thermal locking to resonances,²¹ and for deliberately achieving non-linear resonator responses such as lasing.²²⁻
²⁴ However, a couple of uW *input at the resonator* is usually sufficient for robust photothermal measurements where we want to avoid any nonlinear optical effects. Higher optical powers may be needed up-stream to deliver the requisite power to the microresonator, as the probe path has many stages for losses to build up. **Figure 3-2** shows one example of how these losses can compound significantly.

This particular example is worse than usual for our experiments, likely due to a mediocre tapered fiber, and perhaps some bad couplings or fiber damage elsewhere in the probe path.




Figure 3-2. Probe beam losses. A particularly nasty example of losses in the probe beam path. To the left is shown the net % transmission after each optical element (log scale). To the right is shown the % loss at each element.

If there is too much power in the probe beam at the resonator, this often manifests as thermal broadening of the resonance (**Figure 3-3**),²¹ where the resonance starts to look asymmetric on one side of the probe sweep. The probe laser has a lasing limit in power that may still cause such thermal issues, and thus optical attenuation may be needed, even considering the many losses that can occur. While under-coupling the bullet adaptor that puts light into the tapered fiber can be one way to reduce input power, a more elegant way is to employ an in-line fiber attenuator (Thorlabs VOA630-APC) that can be used adjust the probe power until resonances look qualitatively satisfactory. An even more advanced option is a voltage-tunable attenuator, though we have yet to employ these in our experiments.




Figure 3-3. Thermal broadening of a resonance. The resonance shows up twice due to the symmetric nature of the probe piezo voltage sweep. Notably, the resonance is broadened for one sweep direction, and not for the other

sweep direction. This is a typical indication of thermal broadening, which means that input probe power into the microresonator should be lowered.

Next on their journey, the probe photons encounter the first of two sets of *polarization paddles* (Thorlabs FPC560), each containing a single-mode, 900 um clad patchcord of appropriate wavelength specifications (ex. Thorlabs P1-630Y-FC-2). The thin-cladding is necessary to insert the patchcord loops into the paddles, and the patchcord *cannot* be polarization-maintaining fiber, as this would defeat the purpose of the paddles. These paddles provide arbitrary control of the probe beam's polarization from any polarization state into any other state. The fiber patchcord is looped around inside of each paddle a specific number of times to achieve a desired retardation.

The paddles work through stress-induced birefringence, where each paddle acts as its own waveplate. The number of times the fiber loops around a paddle determines the retardation of that paddle. ex. if is a half-wave plate (HWP), quarter-wave plate (QWP), or a plate with another value of retardation. The angle of the paddle can then be thought of as the equivalent of rotating such an optic in free-space, which greatly impacts how the input light is transformed (see *Section 2.9* for an in-depth discussion of polarization). Also, the Thorlabs manual for these paddles gives a helpful, quantitative descriptions of how the paddles work as waveplates. For a three-paddle system, the goal should be to make the first and third paddles approximately half-wave plates, and the middle paddle approximately a quarter-wave plate, to provide *somewhat* intuitive control over polarization¹. A two-paddle system can instead be used for arbitrary polarization transformation, with each paddle approximating a quarter-waveplate, however this system will likely be less intuitive for the user. One-rotator options also exist, and can in theory also arbitrarily transform polarization states, as they have both a pressure knob for adjusting waveplate retardation and a rotational motion for adjusting waveplate angle. We have yet to

¹ See the paddle manual to help ascertain how many loops this means for a given type of optical fiber.

try these out in photothermal experiments to my knowledge. I suspect they will be more difficult to optimize resonances with compared to a polarization paddle set, but they are probably worth trying out!

In the paddle-based systems at least, it's important to note that none of the paddles will be the ideal desired retardation, and thus the system will likely never be fully intuitive or predictable. In practice, monitoring a resonance on an oscilloscope while adjusting the paddles is the best way to empirically optimize polarization of the probe beam, though recent work towards using automated paddles (Thorlabs MPC320), discussed briefly in *Section 3.3.2*, may improve rationality and/or automation of this process in the future. Free-space fiber launch benches are also available to take fiber-based light and send it free-space, manipulate the light, then relaunch it into fiber, however this would be an expensive, arduous, and power-limiting option.

Above, I specified *two* sets of fiber paddles. The first set of paddles is to prepare the polarization to be optimized for a fiber-based phase modulator (see next paragraph). The second set of paddles is to prepare the phase-modulated light to couple maximally into a resonance of choice (See *Section 2.8* for description of phase/momentum matching). Typically, we've found that the first set of paddles only needs to be adjusted when a new tapered optical fiber is installed, whereas the second set needs frequent resetting, even sometimes mid-experiment with the same resonance if polarization stability is poor. In practice, I've actually found that adjusting the first paddle set can be useful for getting good resonances, while disregarding optimizing polarization state for the fiber phase-modulator. Results may vary here, but I suggest having both sets of paddles to provide maximum control.

In-between the aforementioned polarization paddles lies the *fiber phase modulator*. This small and expensive optic permits locking the probe laser to resonances, and *also* permits the amplitude modulation described further below! The fiber-based phase modulator comprises a piece of lithium niobate, to which an alternating voltage is applied. This alternating voltage, through a process called Pockel's effect, induces birefringence in the lithium niobate crystal within the modulator. This process

causes an alternating change in the refractive index of the lithium niobate, which alternatively compresses and expands the wavelength of the light traveling through it. In frequency-space, this corresponds to two new frequencies of light emerging! Thus, for the purposes of our system, the phase-modulation can be equivalently thought of as frequency modulation. We have now taken our single input wavelength, and turned it into three output wavelengths. The new wavelengths are termed “sidebands”, and how we use them optically and electronically will be further discussed in *Section 3.3.3* on PDH locking. We phase-modulate at somewhere between 200 and 300 MHz for our photothermal experiments, though other resonator experiments in our lab use higher frequencies for their own purposes.

Once the probe photons are beyond the phase-modulator and the second set of polarization paddles, they then transit to the *tapered optical fiber*, mostly referred to henceforth as the “taper”. We have to get the probe beam into the microresonator somehow, so we can monitor resonances. However, we deliberately fabricate our resonators to be very low loss, so that they confine light for many round trips and possess high Q-factors (See *Section 2.4* for a discussion on Q). This constraint for light to leak out, vis-à-vis time-reversal symmetry, also means that it’s hard to get light **into** the resonator, at least from free-space! While it would be nice to somehow free-space couple light into our resonator, and while there are some such examples, ex. by using resonators with deformations in them,²⁵ the best route for us is to provide a route to “leak” light into our high-Q microresonator.

One way of doing this is by using prisms, as has been shown with microsphere resonators.²⁶ We have not used prisms in our toroidal work since there is not enough space above the toroid substrate to put in a crystal. While this constraint does not exist for microbubbles, another aspect to using prisms is that their multimodal nature is non-ideal for robustly coupling into single modes of the microresonator.²⁷ Thus we instead use a different platform: the taper. The properties and fabrication of this custom-made

nanophotonic device are detailed in *Section 3.9*. Further discussion on the merits of prisms vs. tapers, and other coupling methods, can be found in Dr. Kevin Heylman's thesis.¹⁰

The last stop in the *optical* path for our probe beam is at an avalanche photodiode (APD, Thorlabs APD430A). Importantly, this APD must have sufficient *bandwidth* to account for our phase modulation, as we need that optical information to be transduced into an electrical signal that will undergo further processing in the PDH electronics (*Section 3.3.3*). Using a high-bandwidth APD typically comes with a tradeoff in sensitivity, however the one we use easily provides the necessary sensitivity for our uW-level probe beam.¹

3.3.2. Polarization stability

Special care is needed when working with fiber optics to maintain desired and stable polarization of light. If the polarization is unstable, this can lead to substantial experimental issues. For example, WGM resonances are polarization-dependent, and thus wandering of the probe beam polarization will result in a change in the coupling efficiency to a given resonance. If this occurs during photothermal experiments, the result is often either a breaking of the probe beam lock, or at the very least, incorrect calculations of resonance shift due to a changing error slope.

Our group has had significant difficulties over the past couple of years with polarization stability of our Velocity (Newport) lasers. It seems that while some of these issues may be coming from the lasers themselves, the issues have sometimes instead caused by (or compounded by) the improper use of polarization maintaining fibers. I will not delve deeper into these issues here, but encourage the reader to thoroughly read any supporting literature for polarization maintaining fibers before using them. If polarization drift of the probe beam is unavoidable, an active feedback mechanism for adjusting the polarization is likely needed. Efforts are underway in our group to employ motorized polarization paddles

¹ For our power levels, the gain knob on this APD is usually set somewhere between 50% and 75% of the way to maximum on the gain knob rotation.

(Thorlabs MPC220) and a polarimeter (Thorlabs PAX1000IR1) readout to provide active feedback that stabilizes polarization. I have also written some preliminary code using Jones matrices (see *Section 2.9.5* for more about these) to model the effect of polarization paddles, and I suggest further refinement of this code in the future for understanding and controlling polarization in the probe path.

3.3.3. Pound-Drever-Hall locking & laser feedback

A critical component to the limit-of-detection of our photothermal experiments is a technique called Pound-Drever-Hall (PDH) locking.¹ This technique has been used in many laser cavity experiments, including the LIGO gravitational observatory, as well as multiple instances involving microresonators.²⁸⁻
³¹ The implementation of PDH in our experiments, in-tandem with lock-in amplification (discussed in *Section 3.5*), lowered the limit-of-detection of our experiments by multiple orders of magnitude.¹⁰

To motivate why we would want to implement PHD, let's first examine the resonance shift we need to measure in our photothermal experiments. Consider a whispering gallery mode resonance, typically measured as a dip in transmission through a tapered optical fiber (*Section 2.2*). In older renditions of our group's photothermal experiments, the wavelength of the probe beam was repetitively scanned over the resonance, and if the resonance blue-shifted or red-shifted, that shift could be measured by a change in the wavelength-position of the resonance. While useful for some experiments, this repetitive resonances scanning is both slow and limited in its ability to measure very small resonance shifts. What if we were instead to somehow "lock" onto the bottom of the resonance? In such an experiment, we would only read out a voltage output while the probe beam was parked on a specific wavelength. If the resonance shifts to the left, that voltage will increase. If the resonance shifts to the right, that voltage will identically increase (**Figure 2-3** is helpful for visualizing this). Thus, when the signal when locked to the bottom of the resonance will be agnostic to the direction of resonance shift! Since this method fails to diagnose which direction the resonance shifted, it also cannot provide a feedback

¹ Dr. Kevin Heylman's thesis is a good resource for more on the history and namesakes of this technique.

mechanism with which the probe laser could be corrected to remain on-resonance. One way of getting around this issue is to instead lock to the *side* of the resonance, however locking to the bottom of the resonance is necessary to minimize probe laser amplitude noise.¹⁶ Therefore, a way of locking to the bottom of the resonance (the “dip depth”) is necessary.

Enter the phase-sensitive measurement, PDH. Instead of the Lorentzian resonance shape, PDH furnishes a new function, called the *error signal*.¹ In **Figure 3-4A** several resonances are shown. In **Figure 3-4B**, a high-Q (side-band resolved) resonance is shown on top, along with the associated error signal below. In **Figure 3-4C**, a much lower-Q resonance is shown, along with its correspondingly shallower error slope.




Figure 3-4. Comparison of resonances and error signals. (A) A field of resonances is shown. The yellow line comes from a user-selection interface for measuring Q-factors. PDH electronics are turned off for this data, thus no sidebands are generated. Each resonance shows up twice, due to the symmetric nature of the probe piezo voltage sweep. (B) A resonance with sidebands on top, and the corresponding error signal below. This is illustrative of a high-Q resonance in the side-band resolved regime. (C) A lower-Q resonance than the previous example, with the corresponding error slope below. This is certainly a lockable resonance, but one we would say is not in the sideband-resolved regime.

If we lock to zero-crossing point of the error slope (center point), this will correspond to the dip-depth described above. **Now**, if the resonance shifts to the left or right, what happens to the voltage readout at the designated wavelength? The sign of voltage change will be different! This gives us what is known as a **phase sensitive** measurement. The error signal is then sent to a servo-controller (“lockbox”).

¹ The measure of difference between a laser source and a cavity resonance is also termed the *detuning* across a wide variety of literature, so this is a good term to know

Now, if the resonance and/or laser shift relative to one-another in wavelength, the error signal shifts, and the lockbox applies a correction feedback to the probe laser controller, which shifts the wavelength to keep it on resonance. Tips on locking to resonances can be found in *Section 3.12.4* and *Section 3.12.5*.

The apt reader might be now saying “but that error signal just looks like the first derivative of the resonance trace!” And this is indeed true! Importantly, though, we do not have to sweep the probe wavelength here in time, but instead extract the phase-sensitive curve due to the interaction of both the carrier frequency and sideband frequencies with the resonator. In fact, the heart of PDH is that it extracts the out-of-phase component (the mathematically imaginary component) in the beating between the carrier and sidebands.¹ Since this process has been detailed well mathematically elsewhere,^{7, 29} I won’t replicate that derivation here, but will include the resulting equation for the error signal. This is shown in **Equation 3-2** where J_0 and J_1 are zeroth and first Bessel functions of the first kind, which contain information about the carrier frequency and sideband frequency respectively, F is a function describing transmission of light through the resonator^{II} (and F^* is its complex conjugate), Ω is the frequency of phase modulation, and ω is the frequency of the resonance (carrier band).

Equation 3-2

$$\text{Error Signal} \propto 2J_0(\beta)J_1(\beta)\text{Im}(F(\omega)F^*(\omega + \Omega) - F^*(\omega)(\omega - \Omega))$$

Notably, if this equation is plotted for different resonance linewidths (Q-factors), which are included via loss parameters in the equation for F (not shown), vs. the modulation frequency, the visual form of the error signal will change significantly,^{III} as observed for the two error slope examples in **Figure 3-4**.

Besides being a crucial component for PDH locking, the error signal is also very useful for measuring resonance shift, as the **error slope is linear** when close to its zero point. Therefore, the error

^I In jargon, a heterodyne process since it involves interference of two different frequencies.

^{II} The “ F ” comes from an original derivation using a Fabry-Perót cavity.

^{III} I once again direct the reader to Erik Horak’s thesis for excellent visual aids on this.

slope is the signal that we feed to the lock-in amplifier. See *Section 3.5* for more on lock-in amplification, and *Section 3.13.1* for details on how to convert the output signal into resonance shift. Now that we know how the error slope should look, and why it is useful, we will next discuss how to actually generate the error signal in the first place.

3.3.4. PDH components (and how they connect to other electronics)

The anatomy of PHD locking in our photothermal experiments is discussed in detail in this section. I will not list most electronic part numbers, as they are listed in the Methods section of *Chapter 4*. **Figure 3-5** shows a simplified schematic of the electronics of our experiment, including those involved in PDH. Various static attenuators and active (powered) amplifiers are left out; previous theses from our group detail these further,^{6, 16} though when a new PHD system is designed, different components may be optimal.

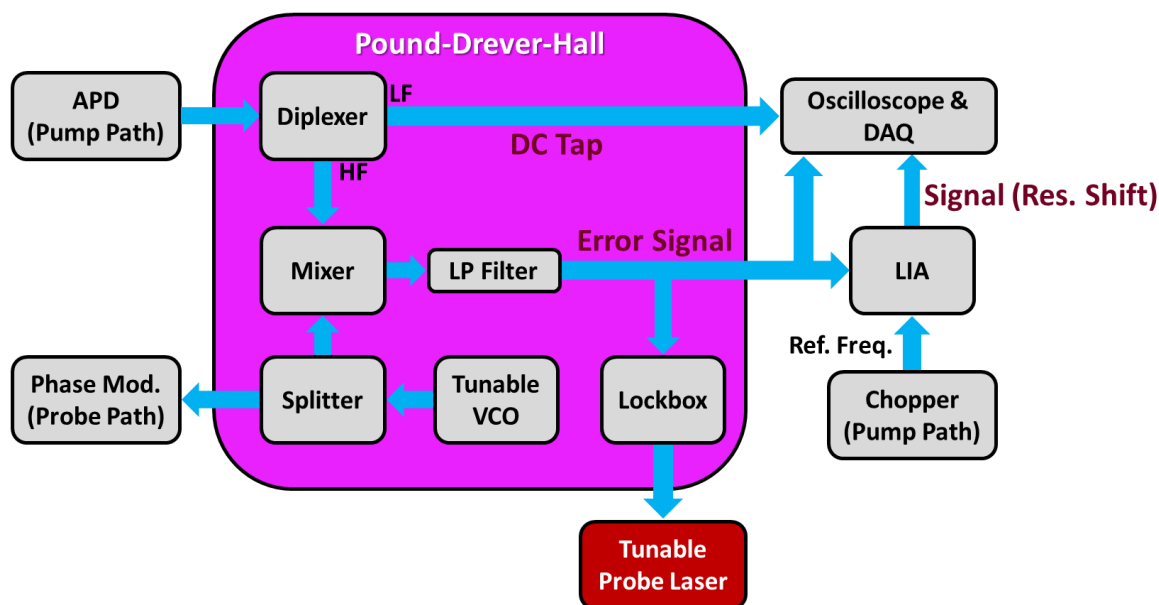


Figure 3-5. A simplified schematic of our experimental electronics. VCO = voltage-controlled oscillator. LP = low-pass. LF = low frequency. HF = high-frequency. APD = Avalanche photodiode. Mod = modulation. LIA = lock-in amplifier. DAQ = data acquisition card. Res. = Resonance.

To generate the error signal, we first phase-modulate the probe-beam to generate resonance sidebands, as discussed in *Section 3.3.3*, by sending signal from a VCO (voltage controlled oscillator) to the phase modulator (EOSpace, for 785nm and 635nm probe beams). The VCO also has some of its signal

(separated by an electronic splitter) sent to a mixer. The mixer multiplies the oscillator signal with the resonance detection signal coming from the APD, which is also oscillating at the phase-modulation frequency due to beating between the carrier and sideband frequencies. This multiplication produces signal at $2f$ (twice the phase-modulation frequency), and at DC. After passing through an electronic low-pass filter, only the latter remains. If you understand the basics of lock-in amplification (or have read ahead to *Section 3.5*), this technique of multiplying a signal by a reference, and then filtering it through a low-pass filter, will seem quite familiar! In both instances, an AC signal is down-converted to DC for easier measurement.ⁱ

If one works through the trigonometry of the electronic mixing described above, they shall notice that the product of two oscillating waves of the same frequency is also dependent on the phase-difference between them. Unlike in the case of lock-in amplification, where we operate in quadrature mode to get around this fact (*Section 3.5.2*), we do something different in PHD: we employ a *tunable* local oscillator. By tuning the frequency of the oscillating signal driving the phase-modulation, we account for distance-mismatch in the different arms of the electrical path that end up being multiplied, eliminating the phase-difference of the two signals to zero.

As described above, after the Error Signal is generated, it is then sent to both the lockbox and the lock-in amplifier. The lockbox we use is one of the components to our PDH set-up that has **changed** in the last few years. We now use a lockbox from Vescent (part D2-125), rather than the previously used box from Newport (part LB1005-S)ⁱⁱ, as we have found the former to provide superior locking capabilities.ⁱⁱⁱ

ⁱ And here, since the phase modulation frequency and the beat notes of carrier and sideband are at the same frequency, we call this a homodyne measurement.

ⁱⁱ Even though this lockbox is sold by Newport essentially to run this exact line of probe lasers that we have, which are from that same company.

ⁱⁱⁱ And break less than the Newport lockbox.

Future improvements for PDH: There are a variety of future improvements that could be made to our PDH electronics system. Firstly, attenuations within the PHD system, especially between the mixer and the lockbox, and between the lockbox and laser, could be optimized. Secondly, the Vescent lockbox that we've transitioned to for this experiment (and on other ongoing resonator experiments in the lab) has substantially more control options than the Newport lockbox formerly used. This is because while the Newport box is only a PI controller (proportional integral), the Vescent lockbox has PI²D capabilities (Proportional, Integral, Integral, Derivative). Thus, the Vescent system has a more extensive set of parameters that can be employed for optimizing locking. However, the control theory needed to optimize this is probably quite complicated, and no-one in the group has yet figured out how to use all the knobs we have to maximize PHD quality for photothermal measurements. Thus, for the time being, we only use the PI component of the Vescent, albeit providing a higher quality lock than with the Newport box. Another possibility, as pointed out by Newport on their website, is to stack multiple servo controllers in-sequence to provide tighter locking control; we have not tried this.

3.4. The Pump Beam Path

3.4.1. Overview

Thus far, we've focused on the probe beam. A probe-beam only experiment would hardly be different from many other microresonator reports out there, such as the impressively sensing reports by Vollmer *et al.*,^{26, 32-35}, for one set of examples. Even PDH-locking, described in *Section 3.3.3*, has been employed in microresonator experiments. What sets our experiments apart from the rest of the literature is the second laser beam, the **pump beam**, which gives us a flexible and very powerful **spectroscopic** handle to leverage in our experiments. This leverage widens the applicability of microresonator technology tremendously; in the case of the pursuits of this thesis, for studying chemical dynamics of single nano-objects.

Many different lasers could be chosen as pump beam, depending on the desired experiments. For example, our toroidal resonators have mostly employed a near-IR (NIR) pump beam (Thorlabs TLK-L1300R). In a different vein, our initial microbubble work with gold nanorod etching utilized multiple single-wavelength diode lasers. In future work, the new tunable laser system from M Squared (*Section 3.6*) will prove invaluable. Regardless of the above, having a few lasers co-aligned has a lot of utility in practice.^I With lasers of different wavelengths, this can be done with suitable dichroic mirrors. I strongly caution against cheap dichroic mirrors for beam-combining, as the wavefront aberration introduced can wreak havoc on the laser PSF.^{II} If one wants to combine wavelengths that are the same (ex. discrete diodes that are at wavelengths included in a tunable laser's wavelength, which may be desirable if the tunable laser is not always available for use), multiple HWPs and a polarizing beam splitter (PBS) may be employed.

There are a variety of important components in the path of the pump beam, including a beam relay system using voltage-controlled galvanometer mirrors (galvos), a chopper for amplitude modulation, various optics for polarization control, a specially chosen focusing objective, and a witness camera. Some of these important components are shown in the simplified pump beam path diagram, **Figure 3-6**. These various components will be described in more detail in the following sections, after a brief section on alignment safety and tips.

^I Of course, different wavelength ranges will necessitate lenses, mirrors, and polarization components with different anti-reflection coatings.

^{II} Often, this means one must go for a super-resolution, TIRF quality dichroic (often \$600+) for combining beams (and a thicker dichroic, while even more expensive is less likely to bend in a mount, and thus worth the extra money)

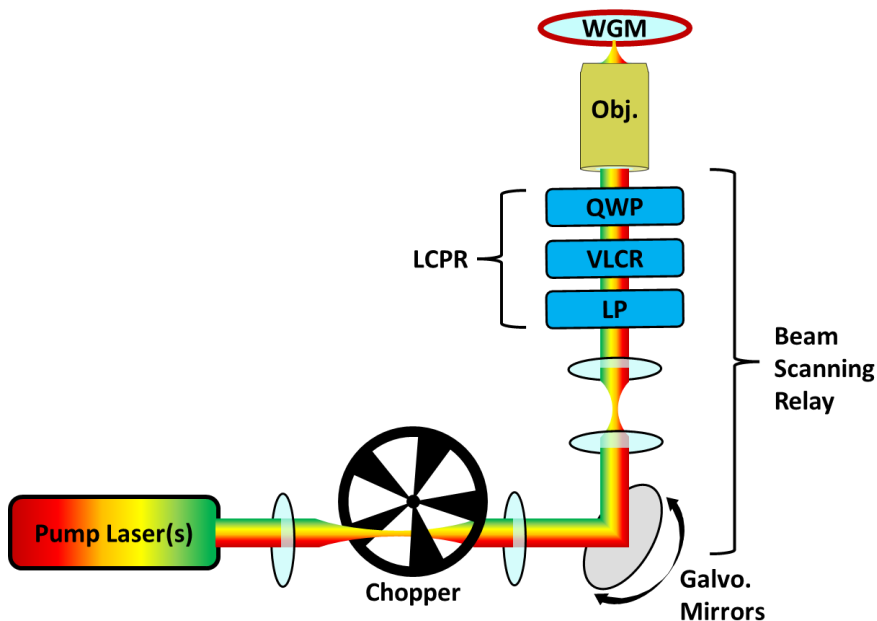


Figure 3-6. A simplified diagram of the probe beam path. Galvo = galvanometer. LP = linear polarizer. VLCR = variable liquid crystal retarder. QWP = quarter wave plate. LCPR = liquid crystal polarization rotator. Obj. = objective, WGM = whispering gallery mode. “WGM” indicates the location of the microresonator, **not** that the pump beam is interacting directly with a WGM (which it is not).

3.4.2. Alignment safety, tips, and tricks

Here, I would like to pass along some tips and tricks for aligning the pump beam, passed down to me or discovered by myself. Firstly, **this section does NOT substitute for proper laser SAFETY training**, nor for authorization to align lasers in lab. All safety trainings and authorizations must be met *before* starting laser alignment. That said, here are some important safety tips to remember when aligning lasers (this is **not** a comprehensive laser safety list!):

- (1) **Power of beam:** The laser beam should be set to a low power. For visible wavelengths, most alignment steps should be doable with <math><1\text{mW}</math> of power, a regime where hazards to the eye are mostly eliminated. If the laser cannot be set to such a low power while still lasing, Neutral Density (ND) filters can be employed to lower the laser power immediately after the laser head/collimator.

- (2) **Invisible beams:** The alignment of an invisible laser beam is trickier than for a visible beam. If possible, a visible laser beam should be co-aligned to the invisible beam with high precision (using a power meter and irises) early in the beam-path, such that the visible laser can be used for subsequent alignment, which the invisible beam will then follow the path of. If such coalignment is not possible, other sources than this thesis should be sought for safety advice prior to attempting alignment!
- (3) **Reflective objects and accessories:** Care should be taken to minimize reflective objects in the vicinity of the beam being aligned, for example any mirrors not bolted to the table in already-aligned positions should be removed. Additionally, accessories such as watches and rings can be quite hazardous, especially given how much their angles change during your movements. Please take off such accessories before doing any kind of beam alignment!
- (4) **Eye/Beam height:** It is best practice to keep your eyes away from the plane(s) in which free-space laser beams propagate on the optical table. Sometimes this may not be possible to avoid, but whenever possible, obey this safety rule. While you hopefully track down and block any stray collimated beams on the table, it would be unfortunate to find a stray beam with your eyeballs.
- (5) **Translating Optics:** For most optics, even if they have angular degrees of freedom, the position of the optomechanics base will need to be translated along the optics table during alignment. This comes with some hazards, as if the optic tips at all, this can send a laser beam upward into an eye. Thus, extra safety should be employed when translating optics, such as turning off the beam, or using a sufficiently low laser power.
- (6) **Periscopes:** Periscopes can be used to compactly change the height of a beam path, in comparison to only traditional kinematic mirror mounts, which typically have around 4 degrees of angular variability. However, periscopes also can be a safety hazard since they send the laser beam vertical. Thus, using a robust periscope (preferably a periscope set bought from a

manufacturer, rather than one cobbled together) is important, as is ensuring all screws are sufficiently tight to keep the periscope exactly in-place once aligned.

(7) **Beam blocks:** Having multiple beam blocks in hand (at least notecards, preferably metal, magnetic beam blocks) is very useful when aligning lasers, and for blocking stray beams after alignment (Ex. back-reflected beams, and those transmitted through mostly-reflective optics).

(8) **Lights, signs, curtains, glasses etc.:** Lastly, but certainly not least, are the various safety features that keep you and others safe if all of the above precautions fail. Such things as laser warning lights, laser warning signs on doors, laser curtains, laser glasses, etc. are very important, and should be employed as needed in specific situations.

Now that safety has been discussed, I will move on to some suggestions for aligning various optical components. Much of this can be determined by trial and error, and honestly it will take experience to have much of this become intuitive. However, I include some suggestions here that will hopefully be helpful in setting up your own microresonator spectrometer pump path! I personally prefer to have as many of much of my laser path as possible following rows of holes on the optics table, as straight and flat as possible, since while this isn't strictly necessary, I find that it makes alignment easier and more reliable. These methods are useful for aligning optics such as lenses and mirrors.

The first thing to know is that for aligning pretty much any optic, you will want a **near alignment point** and a **far alignment point**, as two points make a straight line. These will each be used for respective degrees of freedom, such as mirror translation on the table for the former degrees, and angular changes for the latter degrees, or such as using two mirrors to optimize to the two alignment points for walking a beam to a new height and path. The exact degrees of freedom being optimized to each point will vary depending on the application, and sometimes trial and error is the easiest way to determine the best alignment method. Some additional tips:

- (1) Power meter:** For quick alignment, a power meter often isn't needed, as by eye you can get a beam reasonably optimized through a constricted iris. Nonetheless, for more precision alignment, ex. co-alignment of laser beams, using a power meter will give significantly tighter alignment.
- (2) Two irises:** Irises are great tools for precision alignment. I will not elaborate too much here, but will suggest that having two irises is better than having one, as this allows for easier swapping between the near-degree of freedom and far-degree of freedom when optimizing alignment. If doing this, and presuming that the beam is traveling flat (otherwise, the two-iris method is going to be quite hard to employ), the irises must be **set to the exact same height**. This can be done by setting the heights directly before the most recently aligned optic (ex. laser collimator, ex. most recently set mirror), and using collars to ensure that the iris heights don't change.
- (3) Alignment railings:** For mirrors (which are not usually post-centered) and lenses (even though they should in-theory be post-centered) I like to use what I call "alignment rails" (not to be confused with the optical rails that come in 30 mm, 16 mm, etc. varieties, that are also nice tools for keeping alignment clean). I first set an "H base" (Thorlabs BA1), screwed in such that it is pushed against a set of holes on the desired direction. I then use that H base to set up a "full base" (Thorlabs BA2), which uses the H base as an axis to slide along. Then, the optic (usually a lens or mirror) can, in conjunction with the full base, be translated in a nicely controlled manner along the table in X and Y. Ultimately a tie-down is used to hold down the optic being aligned, as not being perfectly along a row of holes will prevent using only a screw and a washer to buckle it down while keeping the optic base flush with the full base. I find these alignment railing to be both more deterministic and to provide a better alignment than just translating bases with my hands. Notably, this works for **optics mounted on bases (ex. "U" bases or "H" bases)**, but does **not** work easily for the 1" posts that use clamping forks to be held down. In some circumstances,

it's actually useful to employ **two sets** of optical rails for alignment, particularly for aligning large items such as the pillar support that we use to mount our focusing objective, focusing piezo, etc., as this allows for iterations of the X and Y translational degrees of freedom more easily. This comes at significant cost to alignment time.

(4) Minimizing optics: I also try to minimize the number of optics in my path, both for simplicity, and to reduce polarization scrambling, as will happen a little bit with each mirror a laser beam encounters. While this scrambling can be either pre-corrected for, post-corrected for, or just eliminated by a suitable polarizer, this adds optics and reduces transmission efficiency.

(5) Other skills: Other skills necessary for operating the photothermal microscope may include fiber launching free-space light into optical fibers, co-aligning lasers (sometimes including invisible lasers), and beam-shaping, which are skills that will not be discussed further here.

3.4.3. The beam-scanning relay system

3.4.3.1. Description of relay system

A crucial, and important to carefully align, component of the pump beam path is the beam-scanning relay system. We often refer to this beam relay as the “4f system”, where “f” stands for “focus”. This relay allows us to scan the focused beam¹ on the resonator surface to construct a “photothermal map”, where the scanning is done by a voltage-controlled pair of galvanometer mirrors (“galvos”) from Thorlabs (GVS012). In essence, this relay takes the angular displacement of the pump beam imposed by the galvo mirrors and converts that into a lateral displacement in the focal plane of the pump beam focusing objective (**Figure 3-7**).

¹ Not a truly confocal measurement, as no pinhole is involved. I call this a “confauxcal” measurement.




Figure 3-7. A simplified schematic of the beam-scanning relay system. The total of four focal lengths between the galvo mirrors and the objective back-aperture gives rise to the nomenclature “4f”. The colors in the diagram do **not represent different wavelengths**, but instead **different angles** of the galvo. mirrors. Only one of two mirrors is shown for simplicity. Some other optics are also omitted, such as the final mirror and polarization optics prior to the objective.

Pro tip: the constraints of the 4f -system should be considered **before beginning alignment** of the pump path, as an incorrectly chosen initial height for the lasers will likely result in doing the whole process over again.¹ There are two main considerations here:

- (1) The pump beam eventually gets transferred from horizontally to vertically propagating (at least in the current rendition of the set-up) by a 45-degree mirror. The vertical-stack of optics leading to the resonator is tightly constrained in geometry, essentially setting that final mirror’s height. This geometry could be reinvented, either in a vertical or horizontal arrangement, depending on the desired experiments, though this will require custom-machining many new pieces.² The constrained beam height for that final mirror subsequently tightly constrains the height at which the galvanometer mirrors must be set! And because the two galvanometer mirrors are factory-set to a specified distance apart, this sets the incoming-to-galvo height of the beam. Although beam heights can certainly be changed if needed, to simplify the number of optics included, and footprint of the setup, the best course of action is thus to set any pump laser beam sources (i.e.

¹ Trust me, you want to save yourself the time here.

² Erik Horak was a big proponent of reengineering the pump path to be completely horizontal for microbubbles. This would indeed alleviate a number of headaches once done. However, given the movement towards other resonator geometries in the group, I suggest maintaining the inverted vertical geometry for now.

collimators or free-space laser heads) close to the **height of the lower galvo mirror**, once the galvo has been set to its required height per the final mirror, per the necessary focusing objective location.

- (2) The **two lenses** chosen may be of the same focal length, or different focal lengths, such that they also act as a beam expander. The necessary beam size at the focusing objective should be considered here, as well as the beam-size entering the relay system. The incoming beam(s) can also be adjusted as necessary in size with additional lenses, but reducing the number of optics is wise where possible.

SAFETY: The beam scanning system comes with safety hazards, most notably the fact that the angle of a laser beam is being changed. Coupled with back-reflections from downstream optics, this can lead to a bit of a disco in the room! Except the disco is sometimes composed of collimated lasers.¹ To remedy this hazard, firstly, any software that runs the galvanometer mirrors should have encoded voltage limits for the galvos, since while they can go -10V to 10V, this range is unlikely to ever be needed. Secondly, for the needed voltage range, a low-power beam should be scanned to all voltage extremes, and beam blocks should be installed in the pump path to block back-reflections. Thirdly, if particularly high-powers or invisible beams are being used, it is recommended that the user enclose the entire pump beam path, and/or wear appropriate laser safety glasses. Given the open galvo electronics, which we sometimes modulate further with our own custom electronics,² there are also potential electrical hazards associated with the galvanometers, at least anywhere that AC comes into play.

We use something called a “raster pattern” for our beam scanning, which moves the beam along points on one axis of the resonator (thus filling in one column of the photothermal map), and then rapidly backtracks to start again at the next column, once again at the bottom . This avoids hysteresis in the directionality of the galvos, though it can introduce other artefacts depending on the scanning speed and the acquisition time.

¹ Which is, in fact, not a great thing for eye safety.

² Goldsmith group members can reference my files on spatial modulation spectroscopy (SMS).

In future versions of our experiment, other methods may be used for photothermal mapping. While the galvanometer scanning is both fast and quite precise due to the reproducibility of voltage outputs, it also comes with drawbacks. For example, there is a “pillowing effect” due to the fact that the two galvo mirrors are in different locations along the pump beam’s optical axis. Luckily, this is most prominent in the extremes of the galvo ranges, which aren’t typically used, and it could probably be pre-accounted for if needed in the operating code that selects voltages for positions on the resonator.

Another drawback of the beam-scanning system is that the optics in the beam path themselves have limited size. Even using 2” mirrors and lenses, there is only so much scanning range that can be achieved before clipping the beam. Also, the selection of lens focal lengths will impact the relay pathlength, thereby impacted the range of distance the beam uses on lenses. This changing of focal lengths may however have upstream impacts, for example, requiring careful expansion or shrinking of the beam size prior to the galvos, which themselves have a limited span in size. If one truly wanted to optimize this system, I suggest using a software package such as Zemax.¹

Another potential pitfall of the beam scanning must be considered when multiple pump wavelengths are to be used. In particular, it is worth very carefully **centering the beam(s) through the relay lenses**. Otherwise, dispersion will cause a slight angular difference in the beam path through the lenses as a function of wavelength, which will in turn yield a **wavelength-dependent galvo. calibration** (see *Section 3.10.2* for information on this calibration).¹¹

Perhaps the most significant drawbacks of the beam-relay system surround how the quality of the pump laser itself is impacted. For one, the beam’s power transmission through the system will vary by galvo position, largely due to clipping of optics. While a more perfect alignment largely negates this,

¹ Or an open-source equivalent if you don’t want to pay \$15,000 for a license.

¹¹ Though this is not something I have measured yet, since I have not yet taken photothermal data with the M Squared laser system, I think this concern merits careful alignment and diagnosis in the future!

as the beam should always hit the same spot on the back of the focusing objective, in reality, the beam power will change across the FOV used to map the resonator.

Additionally, the fact that the beam will encounter different portions of lenses and mirrors in the path means that it will accrue different alterations to its polarization as a function of galvo voltages. Once the beam hits the final pump polarization optics, the power transmitted will be further impacted, and the continued angular displacement of the beam through these polarization optics will further alter the polarization state of the light even once it hits the resonator, which greatly impacts the depth-of-modulation achievable. Thus, while characterization of the extinction capabilities of the beam may initially be done for a well-centered beam, **characterizing extinction across the needed Field-of-view (FOV)** is also an important, time-consuming, and sometimes geometrically challenging process.

Future Improvements: To avoid the pitfalls of beam-scanning,¹ perhaps the most obvious alternative is to instead raster scan the *resonator*. This means moving the resonator, using a piezo stage, through a statically-located pump beam. Unfortunately, this is difficult in our case, given the necessity of a taper for coupling light into and out of the resonator. The introduction of ring resonators, or similar platforms with an integrated light-coupling mechanism, may permit a transition to a resonator-scanning geometry, though this will come with new challenges of stability in the light in-coupling and out-coupling devices. If implemented, however, resonator scanning will not only make the quality of pump beam far more homogenous across our FOV, but it will also allow a significant reduction in the footprint of our instrument! This is because with sample scanning, the necessity of the 4f relay is obviated, essentially rendering our measurement compatible with a more traditional inverted microscope design. An alternative, suggested by Erik Horak, is to instead use an objective-scanning system. While this might still result in a change in beam intensity at the focal plane (depending on alignment constraints), it would

¹ And to avoid the obnoxious ringing noises that the Galvos make.

likely give a more homogenous quality of beam, especially regarding polarization, than the beam-scanning system. However, to date we have not found the parts to design and build such a system.

3.4.3.2. Aligning the pump relay

Aligning the galvo mirrors is an important step in setting a good pump beam path. The galvos have a few degrees of freedom: the X-Y translation of the mirror system on the table, the rotation of the mirror system on the post, and the voltages applied to each mirror. The incoming beam may also be adjusted if walking mirrors are in place, though this adds to the number of necessary optics. In some permutations, the output beam from the galvos can also be adjusted with further mirrors, though I strongly suggest having the galvo, two lenses, and 45-degree mirror all within **a single vector** if possible, as this greatly reduces the hassle of positioning the beam relay optics. I suggest the following steps, in order, when planning out the beam relay (which should absolutely be done before beginning alignment!):

- (1) See the “Pro Tip” in the previous section for constraints regarding the final beam path height!
- (2) The **planned location of the focus objective should be ascertained.**
- (3) The **two lenses to make the beam relay should be chosen** (which may include factoring in a beam expansion, for example).
- (4) The **position at four foci** from the objective’s back aperture ($f_2+f_2+f_1+f_1$) should be marked, as this is where the galvo system will be placed. If you’ve planned ahead well, this position is hopefully already in the path of the planned path of pump beam; if not, a couple of mirrors should remedy the situation. The downstream microscope components (objective, taper stages, resonator stages, above witness camera, etc.) can also be adjusted to the galvo position, but in

¹ Well, if you want to only do this alignment one time. At least in my experience, lack of planning here before beginning alignment almost always leads to realignment.

practice these later components may be tightly constrained on the table, or at least will take significant effort to move.

With the component locations planned for, it is now time to get the pump beam going flat and straight relative to the optical table after it exits the Galvos. After getting the angle of the galvo system aligned flush with the table, the alignment can be adjusted by translating the mounted galvos, the same as one would adjust the position of a mirror. The other two parameters at our disposal are mirror angles, which here will be controlled by voltages applied to the Galvo control boards. I suggest setting up two pinholes along the line of holes to align the beam to, and using the translational degrees of freedom for optimizing to the first iris, and the applied voltages for optimizing to the second iris. The “center voltages” for the galvos are **unlikely to be exactly 0V**, but if you find centers between -1 and 1V, you’ve done great! Even within -2V and 2V will likely be sufficient, but beyond this, you’ll want to troubleshoot alignment.

Now that the beam is going straight and flat, the two lenses of the relay can be inserted. It is likely best to do the first lens first, followed by the second lens, though the divergence of the beam coming from the first lens makes fine alignment difficult. Once both lenses are placed laterally with the table, and in height and angle such that the beam is going as close through the center of the lenses as possible, and the beam is well-collimated coming from the second lens and going down the same path as initially aligned with the galvos,¹ we can move on to the next component.

Next, the 45-degree mirror should be installed to send the pump beam from horizontally traveling to vertically traveling.

SAFETY: Besides the hazards of the vertical path itself, this leaves room during alignment to send the beam at other angles, which could be towards your face. Exercise extra caution in aligning this optic.

¹ A lot of considerations!

For this step, some sort of vertical-alignment-contraption will be helpful, perhaps employing a cage system or lens tubes. As with horizontal alignment, having two apertures will be ideal for optimizing the degrees of freedom, with the aperture closer to the mirror setting the X-Y translation of the mirror on the table, and the second aperture setting the tip/tilt of the mirror, with a power meter above the second aperture.¹ It is likely that this step can be done prior to setting up the focusing objective, though the expected position of the objective should be considered, given that the many components in this part of the microscope leaves little wiggle room for the position of components.

With the mirror now all aligned, the alignment contraption can be removed, and the focusing objective installed. A power meter should then be mounted above the objective, and the FOV calibration (see *Section 3.10.5*) performed. The FOV will give new “center” coordinates for the galvos that will likely deviate a bit from the starting coordinates. If this value is too different than before, this may cause aberrations, as the FOV will not be centered in the relay lenses. Adjustments to the final mirror may be used to adjust the FOV center to the original galvo center coordinates, though in practice some deviation is fine. Next, polarization components will be inserted into the beam path, and while these ideally will not impact the beam path,² a final FOV calibration is worthwhile once *all* components have been installed.

There are certainly improvements that could be made in this system. For example, optical rails (note that this is not the same as the “alignment rails” I mentioned earlier) could be employed, as they have been in other set-ups in our group. The last time I looked into this, our 2-mirror galvo system was not fully compatible with a rail system (though it might be on one side), but I’m sure that either new pieces, or custom-machined pieces could get around this issue.

¹ In practice, this step can be done a little less rigorously, though this may come with sacrifices in the FOV, polarization control, etc.

² They might clip the beam, though!

3.4.4. The Liquid Crystal Polarization Rotator (LCPR)

3.4.4.1. Overview of the LCPR

An important and somewhat complex component in the pump path optics is the linear polarization rotator. As mentioned near the beginning of this chapter, the rotation of linear polarization provides a readout by which we can confirm nano-objects to have single absorption transition dipoles. Importantly, with this and other polarization methods, it is prudent to put the polarization-rotating optics immediately prior to the focusing objective in the pump path. This helps to ensure that the polarization state of the pump beam at the resonator surface is as pure as possible, whereas if the polarization rotation was done further upstream in the pump path, subsequent mirrors would scramble the polarization somewhat, impacting the extinction ratio, and thus the confidence with which single transition dipoles could be confirmed.

In our previous photothermal microresonator experiments, we used a rotating halfwave plate (HWP) method. As long as the pump light is linearly polarized before it hits the halfwave plate, the rotation of the halfwave plate will then efficiently and proportionally rotate the azimuthal angle of the polarization vs. the HWP rotation angle. Notably, the control of this does necessitate a HWP that is tailored to the pump wavelength being used, or a reasonably achromatic waveplate that can handle multiple wavelengths if the experiment demands this flexibility. When a suitably achromatic waveplate is used, this system is wonderfully robust when operating in the NIR.

However, the second-generation photothermal microscope in our group moved our pump excitation from the NIR (~ 1300nm) to the visible regime (~532-780nm) for doing aqueous studies. What we found was that the rotating HWP method now introduced an artefact of beam-walking (

Figure 3-8), where the focus position moved as a function of waveplate angle. Even moving to a more precise rotation stage (Thorlabs PRM1Z8 vs. K10CR1) we found that the beam walking persisted. We considered trying to measure and compensate for this walking by adjusting the galvo positions while

rotating the HWP, however this would have to be very precise and consistent for us to be confident that beam displacement would not convolute the photothermal measurement. Thus, we sought out another method for polarization rotation.




Figure 3-8. Somewhat reproducible beam-walking at resonator surface due to HWP rotation. Diagnosed by iterative maps at each polarization, for each of which a 2D-Gaussian fit determined the galvo coordinates corresponding to the center of a piece of dust.

What we found was the Liquid Crystal Polarization Rotator (LCPR).¹ The components of an LCPR are not so complicated, though the alignment and calibration of one is less straightforward, and thus described in further detail in *Section 3.10.7* (calibration). The impact of the LCPR is fantastic, as it (1) does not require the physical rotation of any optics once set up, and (2) is significantly faster than a rotating HWP (tens of milliseconds to change polarization angle, as opposed to several seconds).

Our LCPR comprises three optics: (1) a linear polarizer (LP), (2) a variable liquid crystal retarder (LC = liquid crystal), and (3) a quarter waveplate (QWP). For the pre-made LCPRs, at least from Thorlabs, the QWPs are unfortunately wavelength-specific, as for our photothermal spectroscopy experiments, we want something that is much more achromatic in nature. Thus, we needed to design a custom LCPR that would be usable over a wider wavelength range. The optics of the LCPR must be aligned in a specific

¹ This design was envisioned, planned, and initially executed with the significant contribution of Erik Horak.

configuration of rotation angles as described below. Here I further detail each of the three optics in the LCPR:

- (1) **LP:** We chose an economy thin-film polarizer (Thorlabs LPVISE100-A) for compactness, and because the higher-extinction ratio afforded by a nicer polarizer, such as a Glan-Thompson polarizer, is not necessary for our experiments.
- (2) **LC :** We chose a full-wave rotator (Thorlabs LCC1223-B), which permits is a full 180-degree rotation of polarization from the LCPR system. We intentionally selected a tunable LC with *compensation*, which gives the option to apply zero-retardance to the light passing through the liquid crystal. We also intentionally selected a 1" VLCR, instead of a smaller one, to allow for a large enough beam to be transmitted as to overfill the focusing objective, thereby minimizing the spot-size of the laser at the microresonator (sample plane).¹
- (3) **QWP:** We chose a fairly achromatic waveplate (Thorlabs AQWP05M-600), though it leaves a bit to be desired for very precise polarization control (see *Section 3.4.4.4*) for suggested future improvements).

As will be made apparent later in this chapter, while the QWP was selected to be achromatic, the LCPR overall is very much chromatically-dependent. For any wavelength that the LC give a full wave of retardation for, the light will be able to be rotated 180 degrees. However, the **voltage-span** necessary to do so will vary depending on wavelength, and this must be rigorously calibrated (*Section 3.10.7*). This can be rationalized by considering the thickness of the LC in relation to the wavelength of light. For bluer wavelength, the light experiences the LC for more wavelengths, and thus is impacted more by the optic than for redder wavelengths, meaning that a smaller voltage range needs to be applied to get the same rotation the bluer in wavelength one goes.

¹ Though in the current version of the LCPR, the ½" QWP is the limiting optic for transmitted beam diameter.

There are specific constraints to our LC system which govern the order and way in which the optics will be aligned. A different alignment scheme may make sense for future iterations of this system. Our system is mounted partially on a set of optical rails, which permits removing and replacing optics as necessary while maintaining their rotational positions, which is important for the alignment steps outlined below. The alignments below will be easiest to do with the **focusing objective temporarily removed**.

In our LCPR, the third optic (the QWP) is actually the first one installed, as it rests in the objective-holder, and we do not have rotational control over it. This geometry was employed solely to save on space in the vertical beam path, so we didn't have to re-machine many microscope pieces when we switched over from a rotating HWP method. Next, the LP and the LC must be aligned properly and calibrated. There are **two methods** for this process, both of which I will outline below. The first method is more manual and time-intensive, and employs what I call a "calibration polarizer", or "proxy dipole". The second method, which requires an automated polarimeter, is much faster.

How this system works: The LC takes the nicely prepared linear polarization from the LP, and converts it to some elliptical polarization with its axes at 45 degrees relative to those of the LC, where the **degree of ellipticity depends on the retardation** (which depends on the applied voltage). The QWP then takes this elliptically polarized light and converts it back to pure linearly polarized light, where the **degree of ellipticity decides the angle** of the resulting polarization.

3.4.4.2. Aligning the LCRR with a calibration polarizer

The steps that I take, **after the QWP is installed** in the objective holder:

- (1) **Align LP:** Install rotation-mounted LP between the 45-degree mirror and the QWP. Temporarily install a mirror facing downwards after the QWP. Send the pump beam through the LP/QWP/mirror system, and monitor a back-reflection of the beam somewhere upstream.

Rotate the LP until the power of that back-reflection is *minimized*. This works because the LP sets the incoming light to a specific linear polarization (by filtering out all light at other polarizations). That linearly polarized light passes through the QWP, turning into CPL. Upon reflection on the temporary mirror, the CPL flips the light's handedness. As the flipped CPL traverses the QWP, it is transformed into linearly polarized light orthogonal to the original post-polarizer light. Thus, the light will be completely extinguished when it passes through the LP again. However, this orientation is for when the slow axis (and fast axis) of the QWP is 45 degrees rotated relative to the polarizer angle, and we actually want the QWP slow axis (or fast axis) **aligned with** the polarizer angle. Thus, after doing the above steps, the LP should be **rotated 45 degrees**. While this final position could be optimized to begin with, I suggest the procedure outlined above, as optimizing to a minimum will be more precise than optimizing to a maximum. The **temporary mirror** above the QWP can now be removed.

- (2) **Align LC:** First, **remove** the QWP from the system (by removing the objective holder...keep the QWP retained within the holder, as its angle should not be changed!). Next, put in a **second polarizer** (a temporary one) above the LC, followed by a power meter, and *cross the polarizers* such that power through the system is *minimized*. Next, put in the LC (on a 2" rotation stage) between the two polarizers, such that the optics order is LP-LC-temporary polarizer-power meter. The LC controller can be left off, or if on, should be applying 0V to the LC. The LC should then be rotated so that power is once again minimized passing through the optics train. When this happens, this means that the slow axis of the LC (the axis of the molecules at 0 V) is 45 degrees relative to the polarizer angle, and the LC will act as a **phase retarder** for the incident linearly polarized light, variable with applied voltage.
- (3) Reinstall the objective holder (with QWP inside). The system should be ready for testing and calibration!

3.4.4.3. Aligning the LCPR with a polarimeter

An automated polarimeter is a wonderful tool which, for its specified wavelength range, will measure the polarization state of an incoming light source and report to the user its azimuthal angle, degree of ellipticity, beam power, and degree of polarization. Alternatively, the polarimeter can give other parameters carrying the same information, for example the Stokes parameters, which can then be easily used for understanding the polarization state in terms of the Poincaré Sphere (see *Section 2.9.4*).

In brief, the polarimeter we use (Thorlabs PAX1000IR1) works using something called the “rotating quarter waveplate method”, where the incident, collimated beam being studied first hits a rapidly rotating QWP, then a linear polarizer, and finally a photodiode (PD). A fast Fourier transform (FFT) is then applied to the PD signal. For linearly polarized light, the PD will measure oscillating power at **four times** the QWP rotation frequency, for CPL **two times** the frequency, and for unpolarized light, the signal will be DC. The polarimeter then converts the ratio of power at these frequencies to determine the Stokes parameters of the incoming beam, as well as the **azimuthal angle** and **ellipticity angle** of the laser beam (as well as degree of polarization and total power). The manual from Thorlabs for the polarimeter is a great resource to learn more about its operation.

For the purpose of aligning and calibrating the LCPR, having an automated polarimeter is a huge time saver, since a direct readout of polarization state is provided! While it might be possible to align all optics cleverly without having to take any out and put them back in, I suggest the following order for alignment, given the current arrangement of the LC system:

- (1) Install QWP in objective holder.
- (2) Place polarimeter after the QWP in the beam path. The taps in the polarimeter for rail set screws can be helpful for this (**Figure 3-7**).

- (3) Place LP on rotation stage in its intended location. **Rotate the LP** to make the measured light **linear** as possible, which will occur when the slow or fast axis of the QWP is at the same angle as the linear polarization.
- (4) Place the LC on its rotation stage in its intended location, with an applied voltage of **zero volts**. Next, **rotate the LC** until the polarization read out is once again **linear** as possible, at an angle **orthogonal to the linear angle in step #3**.



Figure 3-9 Polarimeter positioning for LCPR alignment.

In practice, the above sequence may or may not give the desired degree of polarization control, and some trial and error may be needed to obtain the desired polarization control. Luckily, the calibration procedure enabled by the polarimeter is much smoother and faster than the calibration-polarizer method (see *Section 3.10.7.2*), readily allowing for iterative alignment tests.

I actually suggest that the reader try setting the different LCPR polarization optics at various angles of that maximize or reduce some parameter, for example the ellipticity of the beam. Ex. One might rotate the LP until ellipticity is maximized, and then rotate the LCPR until the beam is linear again, or vice versa, and see what the calibration result is. Another parameter than can be adjusted is the voltage

applied to the LC. This iterative exploration will hopefully lead to some intuition on how the optics in the LCPR work in sequence, and may also reveal a better way to align the LCPR!¹

3.4.4.4. Future improvements for the LCPR

One of the biggest improvements to be made in quality of linear polarization control, namely getting better extinction ratios, is to have a QWP with true 0.25λ retardation at all pump wavelengths. There are two ways I've thought of to achieve this.

The first is to use a "superachromatic" waveplate (Thorlabs SAQWP05M-700). These, while expensive, provide much closer to quarter-waveplate control over a wider wavelength range, though they're only available in small sizes that may be limiting if one wants to overfill an objective with a large back-aperture to minimize the focused beam size.

The second method is to swap the QWP for yet **another** voltage-tunable liquid crystal variable retarder, which for any wavelength within its operational retardation range can be tuned to be exactly a QWP! This will come with the overhead of having to **tune two LCs simultaneously**, which could be quite challenging to calibrate (though an automated polarimeter will help greatly here) and to code. However, I think the effort here is worth it for the results.

There is another benefit to the two-LC system! If a compensated LC is utilized for the second LC, the voltage can be set to produce a "zero waveplate" (my words, not to be confused with zero-order, but instead meaning zero retardation). This means that the second LC can be thought of as effectively non-existent, meaning that whatever polarization state exits the first LC **passes on through** to the sample. Because the first LC can produce RHCPL and LHCPL, the two-LC system can produce **circularly polarized light of both handedness**, as well as the aforementioned achromatic linear polarization rotation (**Figure 3-10**). This ability to generate CPL could be very useful for studying dipoles that might have orientational

¹ If you do find a successful procedure, write it down! This documentation will help future you, and others, greatly.

dynamics on a speed that could convolute with dynamic chemical or conformational information. Additionally, the ability to rapidly swap between handedness of light could perhaps lead the way to **microresonator-based circular dichroism studies.**

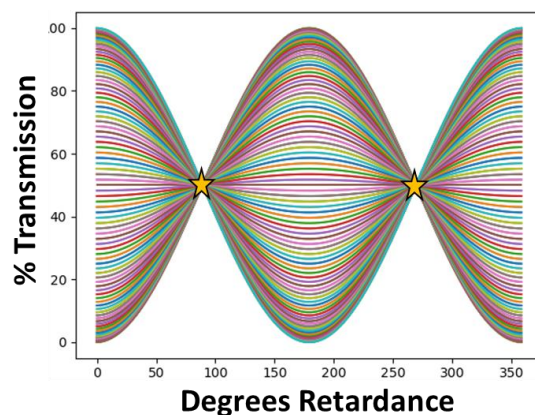


Figure 3-10. Transmission of light through the hypothetical two-LC-LCPR, and a calibration polarizer. The x-axis represents the retardance angle of the second LC. Each trace is for a different retardance angle of the first LC. The two yellow stars denote the points at which all incoming polarization states result in the same transmission power through a calibration polarizer, indicating that the light has been transformed by the first LC into CPL, and passed through unperturbed by the second LC.

3.5. Amplitude Modulation & Lock-in Detection

If PDH is the first half of instrumentation that enables the exquisite sensitivity of our experiments, then the second half is lock-in amplification (LIA).¹ At its heart, LIA is actually quite simple, at least from a trigonometric perspective. We first amplitude-modulate the pump beam, which concomitantly encodes the photothermal signal at the amplitude modulation frequency, thus encoding the signal in the probe beam. Then, with a lock-in amplifier, we look for signal only at the amplitude modulation frequency, and throw all other signal away, thus drastically lowering the noise floor of the measurement.

¹ Put more formally: PDH and LIA are the pieces of bread in our exquisitely sensitive sandwich.

3.5.1. Method of modulation

In the microbubble experiments, I used an optical chopper (Thorlabs MC2000B) to do this amplitude modulation. Other options include employing an acousto-optic modulator (AOM), or for some lasers applying an alternating current to the laser driver. Alternatively, the overlap of the beam and sample can be dithered with spatial modulation spectroscopy, or for anisotropic absorbers, the polarization of the pump light can be modulated. However, at our modulation frequencies, the chopper works great, and is relatively cheap and straightforward! Notably, we approximate the chopper frequency as a sine wave rather than a square wave, and during signal processing, the root mean squared (RMS) of the sine wave encoded in the signal fed to the lock-in amplifier must be considered to convert to peak-to-peak oscillation voltage of the error signal.

3.5.2. The math of LIA

In our experiments, the *error signal* is fed to the input of the LIA (Signal Recovery 7265). Additionally, the modulation signal is fed from the chopper control to the *reference* channel of the LIA. The LIA then multiplies the reference by the servo output, resulting in signal at twice the modulation frequency as well as DC (**Figure 3-11A**). The error signal is then passed through an internal low-pass filter to remove the $2f$ components, leaving only the DC component of the signal. Thus, the signal has been down-converted from AC to DC, with the magnitude of the DC current proportional to the amplitude of the input signal, making for a very convenient readout! Importantly, anything in the signal not at the modulation frequency is not passed to the DC output.

The reader may have already noted that the output of the above transformation is *phase-sensitive*, where the relative phase of the signal and reference channels will also contribute to the output DC signal (this was also discussed in *Section 3.3.3*) on PHD locking, though with a different solution employed). To get around this drawback, we operate our LIA in **quadrature**, wherein the signal channel is multiplied by both the reference channel, and the reference channel phase-shifted by 90 degrees. This

gives cosine and sine terms, which can then together be used to calculate a phase-insensitive readout of photothermal signal (**Figure 3-11B**). These two calculations can also be used to extract the phase difference between signal and reference, which our group has found to sometimes yield useful information.³⁶

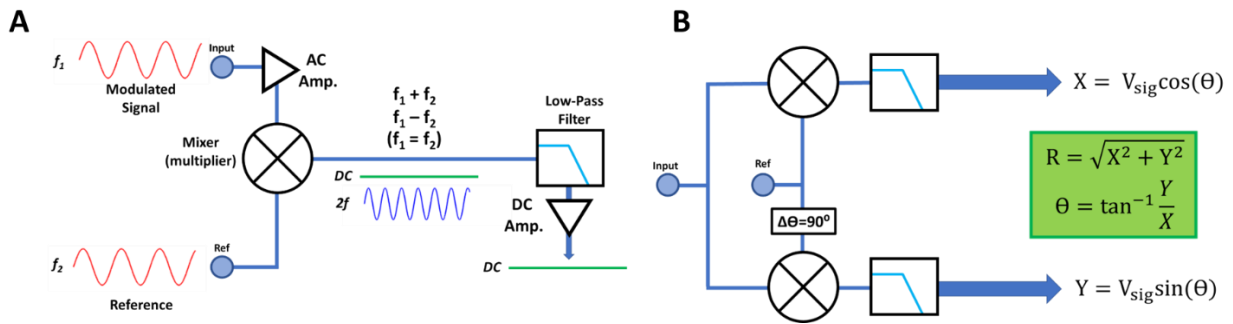


Figure 3-11. Lock-in amplification diagrams. (A) The modulated signal is multiplied by the reference signal, resulting in signal components at $2f$ and DC. A low-pass filter passes through only the DC component. (B) In quadrature, the signal is independently multiplied by the reference and the reference phase-shifted by 90 degrees, providing two orthogonal channels that are used to generate a phase-independent output signal.

3.5.3. Modulation speed

If we want to reduce noise, why not amplitude modulate as fast as possible to reduce $1/\text{frequency}$ noise?

After all, optical modulators readily permit modulations at many megahertz or even gigahertz. So why not modulate at those speeds? Unfortunately, our photothermal experiments have a speed limit, which is governed by the heat transport properties of our resonators. In short, heat can only travel so fast through a microresonator. If the amplitude modulation speed surpasses this heat transport rate by too much, then the modulation gets averaged out, and the LIA doesn't have a modulated signal from which to extract resonance shift information.

In practice, we have found that amplitude modulating at few hundred Hz is a good range for microbubble experiments.¹ We previously demonstrated that 433 Hz modulation frequency results in the resonance shift **modulation** at about 10% of an **offset** resonance shift (See Figure 4-2C),²⁰ both

¹ I like to pick prime numbers to avoid higher harmonics of any integer-frequency noise, though of course noise can happen at non-integer frequencies.

experimentally and in numerical simulations. We found that this value was the best balance between letting through low-frequency noise, and losing photothermal signal to thermal transport limitations. For microtoroids, we typically use something like 2.1 kHz,¹⁰ as the proximal heat sink of the toroid pillar wicks away heat so fast that this higher modulation speed is permitted.^I Further work in considering the thermal parameters of our microresonator systems (thermal transport speed, heat capacity, heat sinks, etc.) would help in the rational design of future photothermal microresonator platforms. It is next worth walking through the most important settings of the lock-in amplifier for our photothermal experiments.

3.5.4. Time constant

One of the most important parameters on a lock-in amplifier is the **time constant (TC)**. In the time-domain, the TC determines how long the signal is averaged for. The longer the signal is averaged, the better the signal-to-noise, provided that one doesn't average so long that another noise source starts to creep-in significantly.^{II} In the frequency-domain, this increase in time averaging corresponds to a reduction in **bandwidth**, which in this situation^{III} means that a smaller frequency range of noise is integrated (converted from AC to DC) along with the desired signal. In which case, since less noise is converted to DC, signal-to-background improves, presuming that the bandwidth of the measurement is still wider than the bandwidth of the signal. Of course, this comes at the cost of a slower experiment! This TC-bandwidth relationship can also be understood in terms of changing the terms in the equation for a low-pass RC (resistor capacitor) filter, though we will not discuss that further here.^{IV}

In practice for microbubble resonators, where we amplitude-modulate at a few hundred Hz, 50 ms is the fastest TC that gives high-quality data. A 20 ms TC can be employed if needed, however artefacts in the photothermal signal start to arise. Thus, for quick searches for big (high cross-section) objects, a

^I I suspect that chip-based platforms such as microrings will allow for even faster modulation speed.

^{II} An Allen-Deviation plot is a typical way to show where the "corner" (optimal frequency) lies.

^{III} Bandwidth can mean many things, and often its use as a term isn't immediately clear.

^{IV} This is particularly useful from an electronics perspective.

20 ms TC is fine, but for publication quality data, and for looking for smaller objects, higher TCs are advisable. *Chapter 5* explores in-depth how various lock-in parameters impact background in our photothermal measurements.

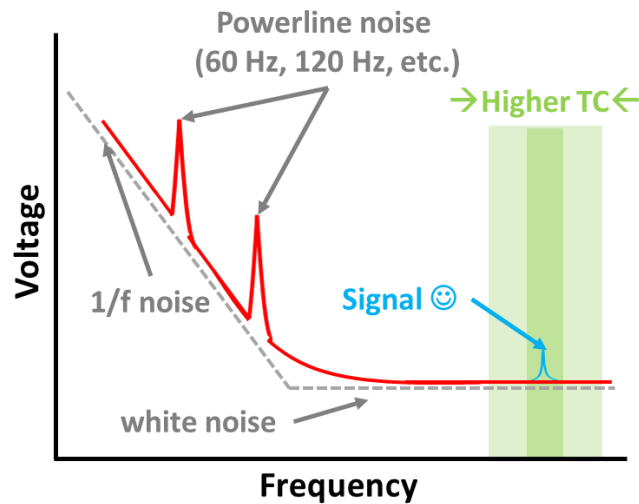


Figure 3-12. A cartoon of how a lock-in amplifier rejects frequency noise. TC = time constant, $1/f = 1/\text{frequency}$. Shown in the cartoon (i.e. not experimental or simulated data plotted) are $1/f$ noise at low frequencies, and additional noise peaks (powerline noise and a higher harmonic shown as examples). At higher frequencies, $1/f$ noise is no longer dominant, and sources of white noise (frequency independent noise) take over as the dominate noise source. The signal has been encoded at a high frequency (blue peak), where it can be detected over the noise floor. The green-shadowed regions represent a shorter TC (wider region) and longer TC (narrower region), to indicate that by increasing the TC, less of the frequency floor is integrated into the measurement (bandwidth is reduced).

3.5.5. Sensitivity

Another important parameter is the Sensitivity (“Sens.”). The lock-in amplifier outputs a DC voltage from 0-10 V. While this output range is constant, it can be proportionally applied to a variety of input voltage ranges. For example, if Sens. is set to 1 V, then an input of 1 V_{AC} oscillating signal will generate a 10 V_{DC} output, whereas a 100 mV input would convert to 1 V output. Smaller values for “Sensitivity”¹ give higher output for a given input, meaning that smaller signals are measurable. However, this comes at the risk of maxing out the dynamic range: any signal above the Sens. will have the same lock-in output of 10 V.

¹ It is easy to get turned-around by this parameter, as saying that something has “higher sensitivity” lends itself to thinking of measuring smaller signals. However, here the term means the opposite.

3.5.6. Dynamic reserve

The last lock-in parameter that we will discuss is **Dynamic Reserve** (DR), which quantifies the ability of the lock-in to reject noise, and is usually expressed in terms of decibels (dB). For example, with 120 dB of DR, the lock-in can fish out a 1 μ V signal with nearby noise of 1 V! While DR will not be considered further here, as it is not a parameter that we directly change in our measurements, it is a worthwhile parameter to be aware of.

3.5.7. Other lock-in settings

The lock-in amplifier currently used for all of our photothermal experiments has many other features, most of which we don't make use of. There are a few settings that are quite useful, however, for improving sensitivity limits. For example, the auto-offset and auto-phase protocols are very useful. These protocols, and their benefits, are outlined further in *Section **Error! Reference source not found.***, with the aim of reducing instrumental background signal in our measurements.

3.6. The M Squared Laser System

Tired of the limitations of visible single-wavelength diodes as pump beam, long did our group search for a visible, high-power, and tunable-without-mode-hopping laser. While we had some false starts, we eventually acquired our current and impressive system from the company M Squared. Despite a variety of significant delays and setbacks, this laser system is nothing short of miraculously cool! As far as the photothermal microscope described in this chapter is concerned, this laser system will be used as the **pump** beam, though some other resonator experiments in the group, or future photothermal experiments, could certainly employ it as a probe beam.

Safety note: This system has many hazards, including **five** Class IV laser beams (though not all are exposed in free-space)! I will not discuss safety concerns in detail here, but any user should be **thoroughly** trained regarding this system before attempting to use it.

The entire laser system comprises three ports that together output light covering much of the visible, NIR, and IR spectrum. **Figure 3-13** gives an overall schematic of the system's main components, as well as some pretty pictures of the output light from the visible mixing module.

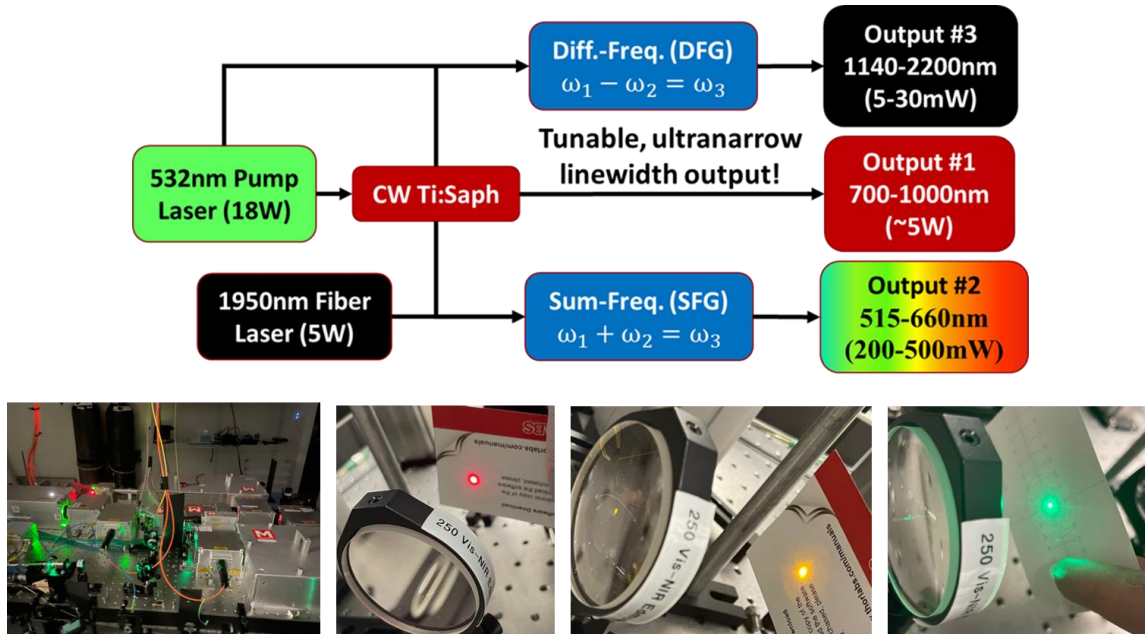


Figure 3-13. The M Squared Laser System. Specifically, our group's system, as this company offers a variety of modules that can be mixed and matched. On the top is shown a diagram of the various modules and outputs. On the bottom are shown a picture of the full system, and a few pictures of pretty wavelengths of light that the system can produce.

The heart of the M Squared system is a continuous wave (CW) Ti:Saph laser, dubbed the “SolsTis”, which can tune from 700-1000nm with roughly 5 W of output, and has an impressive linewidth of <100 kHz over <100 us!¹ The SolsTis is driven by an 18 W, 532 nm CW laser dubbed the “Equinox”. The SolsTis can be output itself for use in various experiments, or instead be steered using polarization optics that vary power between beam paths into two different mixing modules. While power could be coming out of more than one port at a given time, it is important to note that since all ports rely on the tuning of the Ti:Saph, only one output can be tuned as desired at a time, and the other outputs will be tuned along with it.

¹ We will not go through the physics of Ti:Saph lasers here, but I encourage the reader to learn about this!

The first mixing module is the “EMM1950” (EMM = External Mixing Module), often referred to for short as the “EMM” by M Squared or as the “Visible module” by us. The EMM1950 employs sum-frequency generation (SFG) to create visible wavelength light by adding together the energy of two photons to produce a third, higher energy (shorter wavelength) photon. More specifically, the Solstis light is mixed with a 5 W 1950 nm beam (which comes from a seed laser that is fibre-amplified) using a Periodically Poled Lithium Niobate (PPLN) crystal. As the Solstis is tuned from 700-1000nm, the visible output of the EMM1950 is tuned from 515-661nm, with a maximum power of around 600 mW. We have not measured the linewidth of this output, though it is likely limited by the linewidth of the 1950 nm laser.

The second mixing module is the EMM532, often referred to simply the “DFG” module by M Squared, or as the “IR” module by our group. This module also employs a PPLN crystal, but this time for Difference Frequency Generation (DFG), which takes two photons and subtracts one’s energy from the other to produce a third photon of lower energy than either incident photon. The second beam used (other than the Solstis) is in fact a couple of watts of the Equinox that is picked off prior to the Solstis. The tuning range of this module is 1140-2200 nm, with a max power of around 30 mW.¹ Because the two beams used are the Equinox and the Solstis, the latter of which has a linewidth inherited from the former, the DFG should inherit a similar linewidth.

While traditional PPLN crystals are often temperature-tuned to get a small amount of wavelength-mixing range, these PPLN crystals are kept at a very stable and constant temperature. To gain the impressive tunability that this system affords, the PPLN crystals have a fan-grating build into

¹ The maximum power of this module is limited by the power-tolerance of the PPLN to the green beam, as more than around 2 W of power of the green beam will likely destroy the crystal.

them.¹ The DFG module only requires one PPLN crystal to span its tuning range, while the visible module must be switched between two PPLN crystals at around 560 nm wavelength.

Within our group we have lengthy documented protocols regarding many hardware and software aspects of this laser system. For example, the various degrees of tuning afforded, how different control loops operate, how to tune the 1950 nm seed laser using the reference cavity included in the laser system, and many other aspects of safety and operations. Due to intellectual property concerns and the fact that many of these aspects do not directly pertain to this chapter, I will not include these many protocols here, but suffice it to say that this system is a lot of effort to learn and operate, but that the effort is well worth it!

3.7. The Control Computer

A central computer is required to orchestrate the many components of the photothermal spectrometer. A couple of suggestions to keep in-mind when selecting a computer, beyond the typical considerations of processing speed, RAM, etc.:

- For our experiments, we use a Data Acquisition (DAQ) card^{II} (National Instruments part numbers PCIe-6353, NI BNC-2120, and SHC68-68-EPM), which works via PCI-e (Peripheral Component Interconnect Express). This actually ends up being quite a bit more expensive (>\$2k) than most computers. The DAQ plugs directly into the computer's motherboard, allowing for very rapid control of voltage inputs and outputs. Thus, the **computer tower must be large enough** to hold the DAQ card, which is not true of all modern computer towers. For our current experiments, the impressive clock speed of our DAQ is not usually necessary, so cheaper alternatives, such as a USB DAQ, could likely be employed to save on expense.

¹ To be honest, I'm not sure what this means, but I suspect it means that the periodicity of the periodic poling varies across the crystal sideways (orthogonal to the optical axis). This idea is supported by the necessity of lateral crystal translation during wavelength tuning (which is automated).

^{II} Sometimes also abbreviated as "DAC", though that can also mean digital-to-analog converter.

- Many different components connect via USB, for example the chopper, liquid crystal, focusing piezo, Elliptic^l beam shutters, witness camera(s), to name just a few beyond the usual mouse, keyboard, etc. Thus, maximizing the number of USB ports on the computer tower chosen, having these USB ports able to output extra power, and having as many ports as possible be USB3 (for communication speeds) will all be helpful. One or more USB hubs will still likely be needed, and I suggest using hubs that have independent power sources to reduce the necessary power drawn from the computer tower, which already is handling a lot!

3.8. Machined Pieces

Many custom pieces required machining for our microscope, such as adaptor plates to connect metric to imperial components, custom-height posts, and holders for microresonators and tapered fibers, to name just a few. Most such pieces have been made in a machine shop from aluminum, though with the rapid proliferation of accessible 3D-printing, some pieces can now be printed in plastic. Various schematics for such pieces can be found on our group's data backup server.

3.9. Taper Pulling: Protocol & Considerations

3.9.1. Overview

I hope that microresonator geometries other than toroids and bubbles will render tapered optical fibers (tapers) into Goldsmith Group history only, as tapers are pain to fabricate and a pain to work with. But alas, in case tapers are needed in the future, here I will detail the art^{ll} of taper fabrication ("pulling tapers"). A great deal on tapered optical fiber fabrication can also be found in the thesis of Dr. Kevin Heylman,¹⁶ particularly on the many ways in which tapers can break, and how to avoid this sad occurrence.

^l A software name.

^{ll} A somewhat arcane art.

We use a fairly simplistic method for pulling our tapered optical fibers, namely two motors (Thorlabs PT1-Z8), and a hydrogen torch on a manual three-axis translation stage for the tapering, along with a laser source and power meter for diagnostics. There are fancier taper pulling schemes out there, however we've found the hydrogen torch to be sufficient in our needs. Nonetheless, investing in a more precise method may be worthwhile in the future if tapers continue to be needed, *especially* for tapers operating at visible wavelengths, which requires narrower tapers as described shortly below.

While the overall taper pulling set-up was already designed before I started my thesis work (and thus I will not include full diagrams here), I did spend significant time optimizing parameters when we moved from using 1550 nm wavelength for the probe beam to 785 nm, and again when we moved to 635 nm. Regarding the two newer wavelengths, I also designed and machined an apparatus, the **double decker diode holder**, which is shown in **Figure 3-14**.




Figure 3-14. The double decker diode holder. Two relatively cheap diode laser from Thorlabs, with driver boards, mounted on a custom piece that saves on space, and includes heat sinks for each diode.

3.9.2. What makes a good taper?

Firstly, the correct type of optical fiber should be selected for the desired wavelength range of taper operation. Importantly, the core of the fiber needs to be narrow enough to ensure that the probe beam stays single-mode while traversing the non-tapered portions of the fiber. However, having a fiber core that is unnecessarily narrow will introduce losses in the probe beam path.

There are at least four additional parameters for a successful taper. Two of these parameters are optical in nature, and two are mechanical. The first optical characteristic is that the taper should be single-mode in its thinnest region. While not always strictly necessary, having a taper be single mode helps make the process of robustly coupling into single modes of a microresonator much more reliable. An optical fiber is single-mode when it is thin enough such that only the TEM₀₀ Gaussian mode is supported by the fiber. If one does a calculation of glass in air (the refractive index contrast between the fiber and the surrounding media matters), they will find that a fiber optic needs to be subwavelength in diameter to support only one mode.¹ That is quite thin, especially for visible wavelengths! Thus, the way fiber optics actually achieve single mode operation is by having a gradient refractive index, where the core of the fiber is at a slightly different refractive index than the glass cladding it, allowing for fibers >100 μm in diameter to still be single mode. However, we need a “leaky” taper, where the evanescent field of the fiber mode contains electromagnetic energy *outside* of the taper itself, and this is not possible with standard optical fiber thickness. Therein lies the impetus for tapering fibers in the first place. And unfortunately, almost immediately upon softening the fiber while pulling with the motors, the refractive index contrast that maintains the single-mode core is lost, and the fiber optic becomes multimode. The process of returning to single-mode operation (and diagnosing that this is the case) is described below in the fabrication protocol.

The second optical parameter is that the taper should be adiabatic. This means that the tapering of the fiber (ex. from 125 μm to <1 μm) should be gradual enough such that the modes traversing the fiber can transition with the changing fiber diameter. If the fiber tapers too quickly, the non-adiabaticity will result in transmission losses. These losses aren't as big of a deal in the tapered fiber *before* the resonator (provided there is enough input power to deliver a few μW to the resonator), however losses

¹ I direct the reader to Kevin Heylman's thesis for further information on the necessary thinness of fibers to achieve single-mode operation

after the resonator throw away valuable photons. In practice, making an adiabatic taper is much easier for telecom fibers (at around 1550 nm) than for visible fibers at 780nm or 635nm. This makes sense, as the visible fibers must be much thinner to achieve single mode operation, making it significantly more difficult to achieve adiabatic tapering. Thus, while a transmission % (as defined by power transmitting through the fiber after tapering vs. before tapering) for a telecom fiber can reach 90-95%, we in practice consider a visible taper transmission of 60% to be quite good, and we often accept significantly less.

The first mechanical property a good taper will possess is that it will have the correct length. If the taper is too short, it is unlikely to be adiabatic, and will offer coarser tunability for phase-matching to a resonator, especially for microbubbles. However, if the taper is too long, it becomes quite fragile! In particular, if the fiber is tapered beyond where it is glued to its mount, this makes for easily-breakable regions of the fiber beyond the gluing epoxy. Additionally, a longer fiber means an extended region that is thin enough to be susceptible to large losses due to dust binding to the resonator. While having a clean experimental environment (ex. positive pressure of HEPA filtered air) reduces this risk, a too-long taper should still be avoided.

The second mechanical parameter is the tautness of the fiber. Once a taper has been pulled, then further tensioned as needed so that there are no bends remaining in it, it should be pulled a bit tauter, as this will provide a more mechanically stable taper. However, an over-tensioned taper is prone to rapid deterioration of transmission efficiency and sometimes simply breaking.

3.9.3. Fabricating a taper

SAFETY: This procedure involves the use of a hydrogen torch. Specific training and authorization on the hydrogen torch are necessary before pulling tapers! Other hazards include the use of small volumes of flammable solvents, and harps hazards.

Here I describe a step-by-step protocol for taper fabrication. If this is the first taper of the day (this may take several attempts), the fabricator should anticipate running a “test fiber”. The test fiber does not

necessarily need light going through it, but is simply used to ensure that the torch is set correctly, the motors are functioning properly, etc. The steps for pulling a taper are thus:

- (1) **Turn on motors and initialize software.**
- (2) **Open hydrogen tank and confirm torch position.**
- (3) **Install the wind-guard around the torch.** This step is not always necessary, but if the room is drafty, having the plexiglass box in place (with open front and opening on top for heat/hydrogen to escape from towards safety the snorkel) may improve reproducibility of tapering.
- (4) **Select stretch of fiber for tapering.** The stretch should be at least 1 meter in length, or however long is necessary to conveniently couple light into and out of the fiber once the taper is installed on the microresonator setup. It is best to overestimate!¹
 - a. This is most efficiently done if the fiber is unwound from a spool, where the fiber all the way at the other end of the spool (perhaps after a kilometer of fiber) is the laser input side. Thus when the tapering inevitably takes several attempts, each attempt then only requires preparing one fiber end, which saves on time.¹¹
- (5) **Prepare the fiber ends.** Laser light must be coupled into the fiber, usually from a laser patch cable. Light must also be coupled out of the fiber to be measured on a photodetector. The in-coupling side requires a particularly good preparation to achieve good coupling efficiency. To prepare the fiber ends:
 - a. Use fiber strippers to remove the polymer sheath off of the fiber.
 - b. Clean the stripped region with a Kimwipe wetted with isopropanol.

¹ While this will lead to higher fiber costs (which are not insignificant for visible-wavelength fiber), it will save on heartbreak.

¹¹ And sanity.

- c. Cleave the fiber end. For tapering, this can be a simple flat cleave done with a handheld cleaver. For use in microresonator experiments, a nicer cleave may be desired. In fact, to avoid etalon effects in the fiber, an angle-cleave is advisable. As long as you start with some extra fiber on each side of the taper, you can go back and re-cleave later as necessary.

(6) Get laser light going through the fiber.

- a. Couple light into the taper, usually using a bullet adaptor.
- b. Couple light out of the taper, usually using a bullet adaptor and a fiber plate, butted up against an optical power meter.
- c. Turn the laser on and ensure sufficient light is going through the taper, where “sufficient” means that the deviation in laser power is small enough that the readout is reasonably constant.
- d. Turn on the power meter monitoring software. As discussed in *Section 3.11.10*, there can be quirks with Thorlabs power meter software. You may need to change the drivers being used for the power meter box, or try an older or newer version of the GUI, if the power meter box isn’t communicating with the computer.

(7) Prepare the tapering region. Select where you wish to taper the fiber, and prepare a region 1-2” in length. If the region is too short, the polymer sheath will catch fiber. If the region is too long, the magnets on the pulling stages will not be able to grip the fiber well enough.

- a. Remove the polymer sheath from the region using fiber strippers.
- b. Clean the stripped region with a Kimwipe wetted with isopropanol, to remove any sheath fragments remaining. The stripped region will be quite fragile!
- c. Clean the stripped region with a Kimwipe wetted with acetone, to remove any polymer residue on the glass fiber.

- (8) **Place the fiber on the pulling apparatus.** The fiber should be placed taught on the tapering stages, with the stripped portion of the fiber centered over the hydrogen torch. This process is typically iterative, involving placing the clamping magnets on one side of the torch, and then the other side, and repeatedly adjusting each side until satisfied. A final tug on both ends of the fiber from outside the magnets will help make sure that it is taught between the magnets.
- (9) **Heat the tapering region.** Heat the region to be tapered for about 10 seconds using a hairdryer. Our group has various hypotheses on why this might help in the process, from evaporating residual solvent from the fiber cleaning step, to loosening asymmetric tensions within the fiber that would reduce adiabaticity of the taper. The latter is discussed further in the thesis of Connolly.³⁷
- (10) **Start measuring transmission.** For this, use an interface that shows the measured laser power going through the taper over time, with a measurement acquisition time of ≤ 100 ms.
- SAFETY:** At this point, make sure that flammable materials, especially the cleaning solvents, are not near the torch!
- (11) **Turn on the hydrogen torch.** If you have not opened the hydrogen tank itself yet, do so. The regulator is hopefully already set at the necessary pressure. Open the hydrogen valve and cautiously light the torch with a long-nose lighter at an angle from below. You will hear a “pop” as the hydrogen ignites.
- (12) **Start the motors.** Do this promptly after turning on the hydrogen torch, perhaps with a ~ 5 second delay to soften the taper before pulling begins.
- (13) **Check the flame and hot, glowing fiber region.** Safely turn off the room lights, and ensure that there is a *single* glowing region of the taper due to torch heating. If the torch is too low there won't be a glowing region at all, and if the torch is too high (close to the fiber), there may be *two*

glowing regions. This is because the hydrogen comes from multiple holes within the torch, and if the fiber is not high enough in the flame, the heating is not spatially homogenous.

(14) **Monitor the laser transmission.** This is the fun part! At the beginning of taper pulling, the transmission should be quite stable, as shown in **Figure 3-15** (well, in this case, it was slowly decreasing in efficiency, but the *amplitude* of power variation is quite low for the first 3 minutes of pulling). Then, as the taper continues to thin, **mode beating** arises! This beating occurs as the now-multimode taper has possible modes of transmission eliminated as the taper gets too thin to support those modes.¹ The amplitude and length of mode beating can vary quite significantly, as taper pulling with the current set-up, especially for visible wavelength tapers, is not particularly reproducible from pull to pull. Once the **mode beating ceases** (and the amplitude variance is roughly the same as at the beginning of the pulling), the taper has gone single mode! Unfortunately, the transition to single mode is typically not nearly as clear for visible tapers as for telecom tapers, and sometimes there is a last extra mode or two that take a bit to be removed even once amplitude of beating is quite low. Additionally, for visible wavelength tapers, the fall-off in transmission (some of which is observed below) can be *quite* sharp upon going single mode (and if pulling parameters aren't optimized, before going single mode).

¹ I think of this like the acoustic beating that can be heard while two instruments are tuned to each other! The analogy isn't perfect, as acoustic beating will slow down as the difference between two sound frequencies is reduced, whereas in the taper pulling case, the beating speeds up somewhat, and looks more complicated.

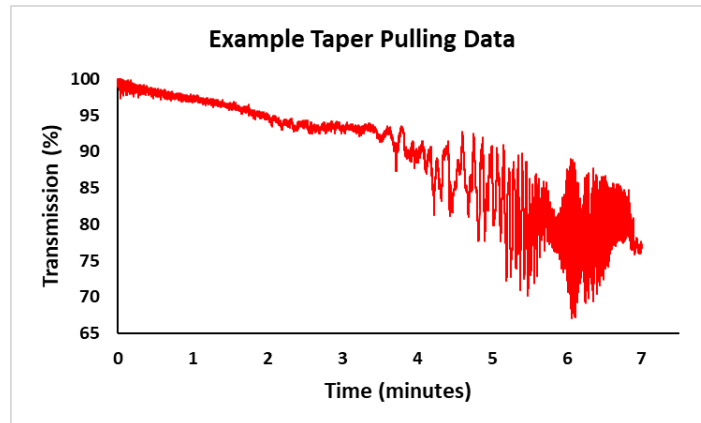


Figure 3-15. Illustrative data of mode beating of a tapered optical fiber while it is being pulled.

(15) **Stop the motors.**

(16) **Turn off the hydrogen torch.**

(17) **Tension the taper.** In practice, something like 6 jogs of 0.01 mm on each motor will do the trick.

Unfortunately, while for telecom. tapers a slight increase in transmission designates a good tensioning, for visible tapers this is often not the case, and results seem to vary wildly for optimal tensioning jog numbers, especially for different parameter regimes of taper pulling.

(18) **Mount the taper.**

- a. Bring in mount after installing the translation stage contraption.
- b. Raise up mount (the aluminum taper “U” mount), adjusting the lateral position of the mount and its rotation as necessary, until it just touches the taper. Raise the mount a little more to provide some extra tension, such that you visibly see the taper rise.
- c. Apply epoxy to the top of each side of the mount, and cure with an ultraviolet flashlight.
- d. Tape the fiber back on the sides of the mount, so that should the fiber get tugged on accidentally, the tape provides a point of resistance, rather than the fragile epoxy points.

(19) **Remove the taper stage, and carefully install it on your set-up!**

Troubleshooting: Often, your first taper will not go how you want it to! A variety of things can go wrong, and there are a variety of ways to troubleshoot. The first thing to do when a taper doesn't turn out as desired is to **try pulling another taper with the same parameters** (motor speed, torch position, torch pressure) and see if the results of the first pull are reproduced. Sometimes, the exact same parameters will give different results due to variation in how the taper pulls near the end due to air currents, the stochastic introduction of dust as a loss mechanism within the taper, etc. If the results are similar in the second pull, here are a few things you might try for different issues.

- *Taper is pulling too long before going single mode:* The fiber is probably melting too much compared to the rate of motor pulling. Suggested remedies are increasing the speed of the motors, lowering the torch position, or lowering hydrogen pressure.
- *Taper is too lossy:* The fiber may not be melting enough to keep up with the pulling. Suggested remedies are reducing the pulling speed of the motors, raising the torch to heat more of the taper, or increasing the hydrogen pressure.
- *Taper is breaking while pulling:* Usually, this occurs for similar reasons as the taper being too lossy, and similar adjustments need to be made as for that problem. However, sometimes the taper will break later in the tapering stage, which may indicate a need to speed up pulling later in the pull to avoid the taper softening so much that it breaks due to pressure from the torch.

If the above aren't working, you may want to try disassembling the tapering apparatus, cleaning surfaces, greasing the rails, and reassembling the apparatus.

3.9.4. Future improvements for tapering

Besides other tapering options such as using a dithered flame, an electrical arc, etc., there is substantial room for improvement on the tapering protocol described above. This is particularly true for visible-

wavelength tapers, as these fibers often bow upward quite severely during the tapering process, which we expect is a major factor in the low-reproducibility of taper pulling. Luckily, Steven Carpenter in the lab of Deniz Yavuz, in the UW-Madison Physics Department, has made good progress in improving this situation! In particular, he has worked to optimize a pulling sequence that changes the pulling speed (generally to faster) as the taper progresses to ensure that bowing doesn't occur. Their tapers are for around 900 nm wavelength, and they prioritize >2 μm taper waists over going completely single mode, as they use a very high input power of >1 W in their resonator experiments.²⁴ I have so far had some promising results with modifying this sequence for visible-wavelength tapers, though further work is needed to make this process robust and reproducible.

3.9.5. Ultrathin tapers

If the taper is particularly thin in the center, WGM resonances can actually manifest in peaks instead of dips! This phenomenon, where the resonator is effectively acting as a bandpass filter that mediates energy transfer of resonance wavelengths through an otherwise lossy taper, has been reported in the literature for microbubbles,³⁸ and we have seen similar behavior with microbubbles and very thin tapers in our lab.

3.10. Calibration Procedures

3.10.1. Probe beam piezo

One important calibration to perform is the voltage-to-wavelength parameter for the Velocity laser's (probe beam's) cavity piezo.ⁱ This parameter is vital for backtracking a resonance shift from the lock-in amplifier output (which is the oscillating change in error signal, which is proportional, as described earlier, to the amplitude of resonance shift oscillation due to amplitude modulation of the pump beam).ⁱⁱ

ⁱ If interested in how this works, brochures from Newport describe how their tunable cavity lasers work, in general.

ⁱⁱ A possible alternative to needing this calibration, or accounting for a non-linear feedback response in the probe beam's wavelength (should this ever be the situation), would be to do a live-reading of the probe wavelength using a wavemeter or OSA while doing photothermal experiments.

This calibration can be done with an Optical Spectrum Analyzer (OSA), which has the sensitivity necessarily to precisely pick out the laser's peak wavelength and also its linewidth.

Given that new OSAs are very expensive, and older OSAs are increasingly challenging to use,ⁱ I have more recently employed a wavemeter (WS5, made by Ångström/High-Finesse.) The wavemeter only gives a single wavelength, and thus doesn't provide a measure of linewidth.ⁱⁱ Nonetheless, the more modern wavemeters on-hand in our lab make this an easily calibration to set-up programmatically, by sweeping the piezo voltage (using, for example, a DAQ output) and measuring the wavelength. The current version of the program I wrote controls the wavemeter through LabVIEW. There are some quirks, such that the GUI must *also* be running, and the associated dynamic library *must* be in the System 32 directory of the computer.

Figure 3-16 shows calibration data using this code and method for a particular velocity laser at a variety of "piezo" settings, for a certain center wavelength, for one of the two 635 nm Velocity lasers in our lab ("older" indicating the first of the two lasers). In practice, having the piezo setting of the Velocity at "50%" seems to give the best tuning range, I suspect because this is a center point setting for the piezo.ⁱⁱⁱ As shown in the figure, the 635 nm Velocity piezo has a tuning parameter of about **20 pm/V**. In comparison the 780 nm Velocity in our lab has a calibration parameter of 60 pm/V. Thus, these lasers can **vary substantially in piezo calibration**, and each laser should be calibrated (perhaps with a set schedule, though I have not tested if calibration changes at all over time, and do not have any reason to suspect this to be the case). In fact, in *Section 3.13.3* I outline how to alter various important measurements in post-processing, should an incorrect piezo calibration parameter have been used in initial calculations.

ⁱ Unless one likes using floppy disks.

ⁱⁱ While linewidth isn't necessary for the calibration, the loss of information does require trusting the wavemeter a bit more than is necessary for an OSA, where power at all wavelengths is observable.

ⁱⁱⁱ Even though I find this not obvious. I would think that some sort of scaling factor for the piezo would be what a % parameter would mean, but based on the above data, this does not seem to be the case.

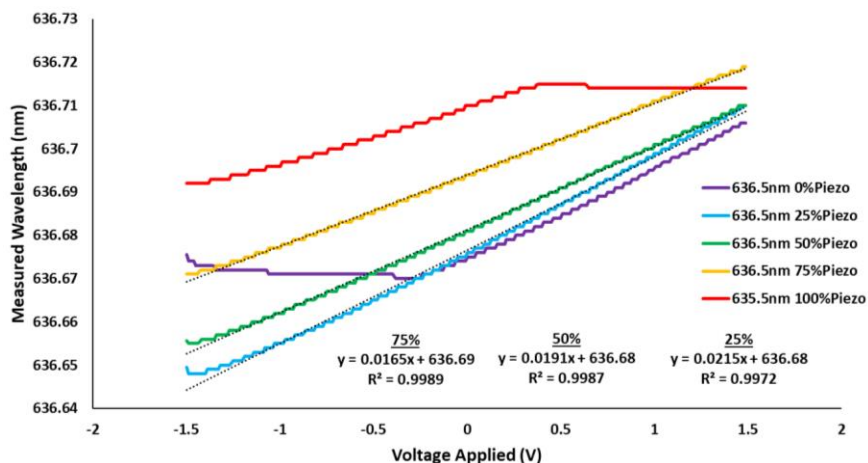


Figure 3-16. Probe voltage piezo calibration. Illustrative data of calibrating the probe piezo with an in-house LabVIEW code, for one of our 635-658 nm tunable Velocity lasers, at a central wavelength of 636.5 nm, for different “piezo” settings, across a range of applied voltage, as measured by a WS-5 Wavemeter (Ångstrom/High Finesse).

3.10.2. Galvo mirror calibration

An essential parameter, particularly for photothermal mapping, is the “galvo calibration parameter”, which describes how much the **spatial location of the pump beam** at its focal plane (on the microresonator) varies as a function of the **applied voltages to the galvo mirrors**. To diagnose this parameter, we use a calibration slide that we put between the focusing objective and the witness objective. We then can use the calibration slide to denote what galvo voltages reach known distance thresholds, and these values can be plotted to a linear fit, from which the **slope** is the parameter we need! In the past I have also done this calibration for each laser beam that is co-aligned in the pump path, though as might be expected, the calibration parameter turns out to be rather wavelength-independent. **Figure 3-17** shows the experimental set-up an example calibration for the galvo parameter. Also shown in this figure is the dependence of the galvo-calibration parameter on the magnification of the focusing objective! This is discussed further in *Section 5.10.7*, which explores the impact of different objectives on photothermal background in our measurements.

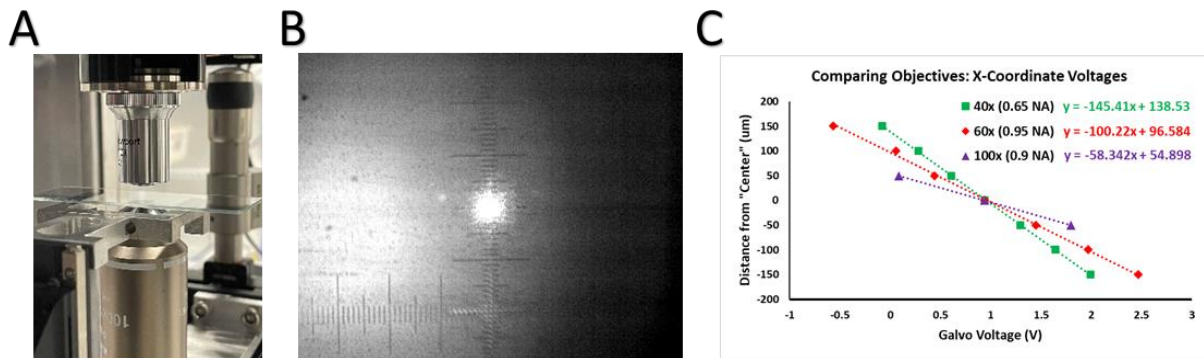


Figure 3-17. Calibration of the beam scanning relay system (galvo mirror parameter). (A) The arrangement of focusing objective (bottom), microscope slide on the attocube stage (middle), and witness objective (top). (B) Example image for calibration data, where pump beam is focused and the galvo. mirrors have been adjusted to make the beam hit a certain known distance point on the calibration slide. (C) Calibration data for three different magnification objectives (for one of the galvo. mirrors), for the same 4f-relay alignment.

3.10.3. Pump beam focusing

For our experiment to be truly a *spectroscopy* experiment, we want to study our single nano-analytes while they are pumped with different wavelengths of light. While we have used tunable pump sources in the NIR, and multiple pump wavelengths in the visible, we have yet to do tunable absorption spectroscopy in the visible. However, as outlined in *Section 3.6*, we now have a tunable, CW, mode-hop free, high-power pump laser in the visible wavelength regime! Thus, **new considerations are in order** for how to properly do photothermal absorption spectroscopy experiments using a tunable pump beam.

Multiple such considerations come into play when we consider how our pump beam will be focused onto the resonator by our microscope objective. We also recently changed to a 100x objective, chosen for its focused-beam spot size, which was determined empirically and discussed in *Section 5.10.7*. However, this objective isn't of a particularly achromatic design (based on the specific product line), and thus different wavelengths of light may interact differently with the objective in a significant way.

For example, the focal plane of the objective will be wavelength dependent to some extent. Thus, as we tune the pump laser, for example to take an absorption spectrum of an analyte, **the focal distance will also change**. Luckily, we already have our focusing objective on a piezo controller (Thorlabs DRV517), so the voltage of this controller may be tuned simultaneously (or in-step) with the wavelength (though

this will require a yet-to-be-measured calibration curve). To additionally complicate the situation, the diffraction limit itself varies as a function of wavelength, meaning that for different colors of light, even once the focal plane is corrected for, the **spot size will vary**. While this could be theoretically determined, it is worth taking the time to calibrate this out robustly at least once.

It's not immediately apparent how much either of these calibration curves/libraries will vary across the scanning FOV. That is, different positions on the resonator, which can be reached by the pump beam by scanning the galvanometer mirrors in the 4-focus lens relay, may have aberrations that alter the center-point or even the scaling of these spot-size-distance-color trends.¹ For very rigorous experiments, this may also need to be calibrated out.¹¹ In the meantime, if one has the time and patience to take photothermal maps of an object at every wavelength, then at least the spot-size will always be known, and the light intensity can be backed out!

3.10.4. Pump beam power calibrations

Some locations in the pump path, such as the fiber-launching of the M Squared visible module output, could vary significantly in efficiency if optimization is not performed for every experiment. Other steps, such as passing through polarizing optics, could be impacted in efficiency by polarization wandering of the pump beam source. Additionally, the tuning of the M Squared laser system may, depending on the settings, not always land on the exact same optimized power (a control loop tunes both the Ti:Saph and the mixing module's PPLN crystal position). Thus, it may be wise to employ a permanently-placed power-meter as **power pick-off** as a future improvement to the set-up, for example, immediately before the 4f lens relay. If the post-pick-off transmission efficiency is either consistent, or at least robustly calibrated as a function of wavelength, the pick-off can then be used to determine the pump beam power at the

¹ It's already well-documented by us, and unsurprising, that polarization control varies across the FOV, manifesting as a highly variable maximum extinction ratio across the FOV.

¹¹ In the future, hopefully changes in our microresonator platform obviate the need for calibrations across the FOV, as described earlier in this chapter.

resonator surface as a function of pixel/polarization/wavelength. Alternatively, rigorous tests to show the reproducibility of pump beam power and transmission % through the pump path could obviate the need for a constant experimental read-out. As of now, none of the visible photothermal set-up scripts include the reading of a pick-off power meter. Other scripts in the group already have some of this infrastructure designed, so it likely wouldn't be too difficult to implement.

For non-tunable pump sources, for example, the BlueSky FiberTec lasers we've traditionally used as visible pump beams, the calibration is much more straightforward. The % transmission through the entire system needs to be calibrated, which is accomplished by setting the laser to a known power, and measuring the power at the end of the pump beam path (post-objective). Measuring power at several in-between locations (ideally between every optical element) will also provide a by-optic loss diagnostic that can be used to help improve transmission efficiency if needed. This assumes that every element has a loss % that is proportional to pump power.

Another important calibration power is the **power vs. voltage** curve for some diode lasers, such as the aforementioned BlueSky diodes. These datasets can then be integrated into the photothermal master code, such that a desired laser power can be input, and the code will change a DAQ output voltage to change that laser's power accordingly. There is a LabVIEW script already written for varying DAQ voltage and reading from a Thorlabs power meter to make this calibration quite easy. Sample calibration data for the red and green BlueSky FiberTec lasers in our lab is shown in **Figure 3-18**.

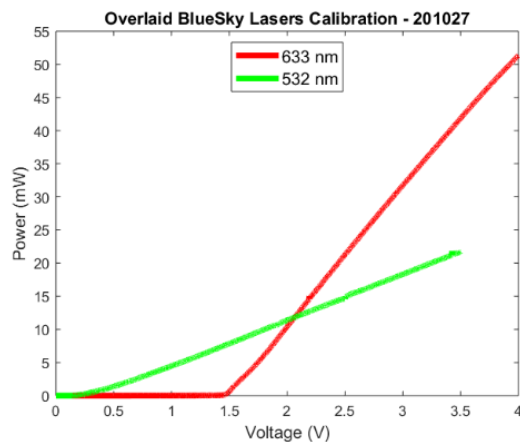


Figure 3-18. Diode power calibrations. Power vs. voltage calibrations for two BlueSky FiberTec lasers of different wavelengths.

3.10.5. FOV calibrations

It is important that the field-of-view (FOV) is calibrated. This is largely because, as shown in the work in Chapter 4 of this thesis, sometimes knowing the pump power¹ is crucial for normalizing photothermal signal in experiments. To get this, a “FOV Map” is needed, which gives the transmission % of the pump beam as a function of spatial coordinate (galvo voltages). In practice, this should be done for **each pump laser beam** (if multiple are co-aligned), as disparities in beam diameter and imperfections in co-alignment will impact clipping losses. Additionally, as discussed in Chapter 5, changing the pump objective can also change the FOV calibration, both because of clipping losses and because different objectives have different intrinsic losses. **Figure 3-19** shows a sample FOV calibration map, where galvo voltages for X and Y have been converted to distances using the afore-described galvo calibration parameter.

¹ In fact, sometimes the pump intensity, as in Chapter 4, and discussed in detail in Chapter 5.




Figure 3-19. Example FOV calibration.

3.10.6. Extinction map calibrations

Another parameter that is important to measure for beam scanning is the **extinction ratio** across the FOV, as this will determine if a single excitation transition dipole can be confirmed (and thus a single molecule identified) at each coordinate. **First the user should do the LCPR calibrations.** This “**extinction map**” will then require a somewhat complicated and perhaps unintuitive calibration, and it should be done for each pump beam that is co-aligned (or even a few different wavelengths if a tunable laser is being used).

It will be much easier to do this **without** an objective (largely just due to spatial constraints in the current microscope design). However, this will only give the extinction variation due to the beam relay system and polarization optics, and it might be that the objective itself contributes to extinction as a function of beam position. It is up to the user if they want to try and include the objective.

Either way, after the QWP of the LCPR, or after the objective if it is included, a calibration polarizer (“proxy dipole”) is placed in the beam path. After the polarizer is placed a power meter. With everything in place, the following steps should be taken:

1. Pick a given **extrema in the extinction curve (Figure 3-20A)**. It is useful to do this procedure at all such extrema, as it will give an idea of the range of best-case and worst-case extinction maps, depending on the angle of an analyte’s excitation transition dipole.

2. Set Polarizer #2 (proxy dipole) to the **relevant angle** for that extrema.
3. Sweep the LC voltage and measure the power meter reading (in a synchronous, automated fashion).
4. Set LC to the voltage which gave the **minimum** transmission **Figure 3-20B**.
5. Optimize polarizer to further minimize transmission (this should be a minor adjustment).
6. Take a FOV map (i.e. transmission % vs. galvo coordinate)! This will be called the “minimum map” (Example in **Figure 3-20C**).
7. Adjust polarizer (by around 45 deg) to **maximize** transmission.
8. Take map! This will be called the “maximum map.” (Example in **Figure 3-20D**).
9. Divide the maximum map by the minimum map. This will give the extinction at each galvo coordinate, **for the angle corresponding to the picked extinction extrema** in Step #1. (Sample calibration shown in **Figure 3-20E**).
10. Repeat for other extrema.
11. Repeat for other pump lasers/wavelengths.

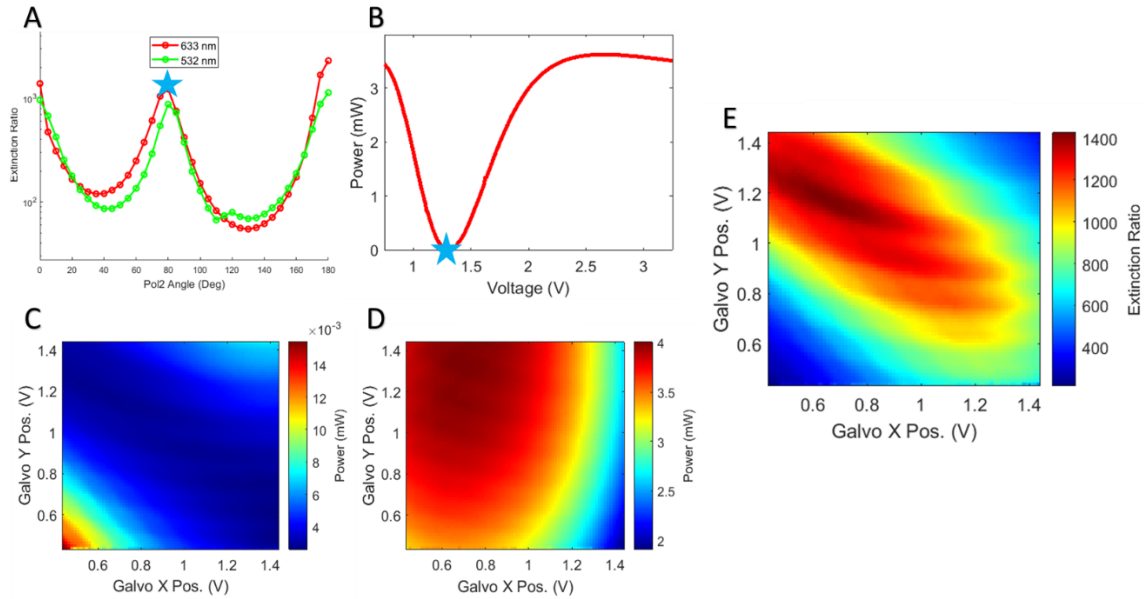


Figure 3-20. Extinction map calibration. (A) Extinction data as a function of proxy dipole angle, where the star represents the chosen extrema for this extinction map calibration. (B) An LC voltage sweep with the proxy dipole set at the angle of the star in (A). (C) A transmission map through the calibration polarizer at the voltage from the star in (B). (D) A transmission map through the calibration polarizer at the voltage from the star in (B), with the polarizer turned 45 degrees relative to (C). (E) The resulting map from dividing (C) by (D), which gives the per-pixel best-case extinction ratio.

The extinction map patterns I have seen over different calibrations in different iterations of the microscope have varied wildly, especially as a function of wavelength. Ideally, the center of the map has the best extinction (since this is where the LCPR was optimized for), though this is not always the case, as shown above. I do not have a rational explanation for the variance and odd shapes in extinction maps, except that several optics are in the pump beam relay, and they likely contribute together to polarization stage in a complicated way that is sensitive to small imperfections in alignment.

The protocol above is similar to that used for the **old** method of LCPR calibration (“proxy dipole” method), and thus a newer method might be in order now that we have an automated polarimeter for polarization alignments. Unfortunately, it is important for the polarimeter that light be coming in **collimated and straight**, thus using the polarimeter itself for Extinction Map calibration as above is likely out of the question. Still, this current method certainly has its drawbacks even beyond its time

consumption, such as the spatial constraints listed above, and the fact that changing the galvo angle will impact the angle at which light enters and exits the proxy dipole. Thus, **future work to improve this calibration** would be worthwhile!

Alternatively, an **internal standard** could perhaps be used for getting an idea about extinction across the FOV (albeit a rather spatially scattered one). For example, we know reasonably well how to deposit gold nanorods (AuNRs) into microbubbles. Thus, if one deposits AuNRs in a microbubble and measures high extinction ratios in a certain mapping region, that should imply that other dipole-containing objects should also be able to be resolved via polarization sweeps at those voltage coordinates.

In the end, whether or not extinction maps are acquired is up to the user. Just remember that the polarization calibration (described next) is unlikely to apply to the entire FOV, which will have experimental consequences, such as constraining how much of the resonator is usable for single molecule experiments.

3.10.7. Pump beam polarization

In *Section 3.4.4*, we explored alignment of the LCPR, a crucial component for many of our photothermal spectroscopy studies. Here, we discuss calibration of the LCPR, and how to predict its behavior for a given alignment using polarization theory. Before diving in, I suggest that the reader take the time to read *Section 2.9*, as concepts discussed therein will be applied in this this section.

There are actually two purposes of this calibration: (1) to determine if the maximum extinction ratio achievable with the system meets experimental requirements for depth-of-modulation diagnosis, and (2) to obtain a **calibration curve** for the liquid crystal as a function of applied voltage. Notably, the latter is *quite wavelength-dependent* as well!

As with the LCPR alignment, there are two routes to take, once of which involving a proxy dipole polarizer, and one involving an automated polarimeter. I will go over both methods here. Before moving

past this point, **make sure that the galvos are set to their FOV center voltages!** This will ensure that you are calibrating (and if adjusting optics, optimizing to) where your resonator will typically be centered.

3.10.7.1. The calibration polarizer method

The goal here is to determine what the **extinction ratio will be as a function of absorption transition dipole angle** (assuming a perfect dipole). The calibration polarizer fills the role of this dipole (thus, the “proxy dipole”). This measurement will be easier if the focusing objective is removed (with the beam focusing through the polarizer, it is hard to collect the light on the other side of it), but this may come at some cost, since the impact (if any) of the focusing objective won’t be accounted for. After the calibration polarizer, a power meter should be placed. The data collection and processing then proceed as such¹:

- (1) The calibration polarizer needs to be **iteratively set to different angles**, and at **each angle an LC sweep should be done**, with the transmitted power measured.
- (2) These traces can be normalized for easier viewing (Sample data is shown in **Figure 3-21A** and **Figure 3-21B** for two different wavelengths, with many traces from different proxy dipole angles overlaid onto the same plots).
- (3) From the above, the **Extinction ratio** as a function of pump polarization angle can be found, as shown in **Figure 3-21C**.
- (4) The *minimum* point for each trace is extracted, along with its **corresponding LC voltage**.
- (5) A plot is generated that shows **polarizer angle vs. the minimum voltages corresponding (Figure 3-21D)**. This would be enough information if we could tolerate the poor angular resolution set by the number of calibration polarizer steps, but luckily our data contains much more information that we have yet to process (after all, we in step #3 took only one data point from each LC sweep!).

¹ This process is not trivial, and the data analysis can get a bit messy. Group members can pull useful scripts, at least as starting points for future analyses, from my MATLAB codes.

(6) This is the tricky bit. To extract a lot more information, we will now take **all data points from step #4**, and in **each LC sweep**, find the point with **nearest matching normalized transmission %**. The **corresponding voltage** for each of those points is then saved. Thus, we effectively are “projecting” each LC sweep onto the minimization curve found in Step #4. **Figure 3-21E** shows all such curves overlaid for one of the pump beams.

(7) Lastly, all of the curves are **averaged together** for each pump color calibrated, to give a final **angle vs. voltage** curve (**Figure 3-21F**). The data making these curves can then be imported into codes such as the master photothermal code, so that the user may specify a wavelength and desired polarization angle, and the code interpolates to find, then applies, the necessary LC voltage to convert the pump beam to the desired polarization angle.

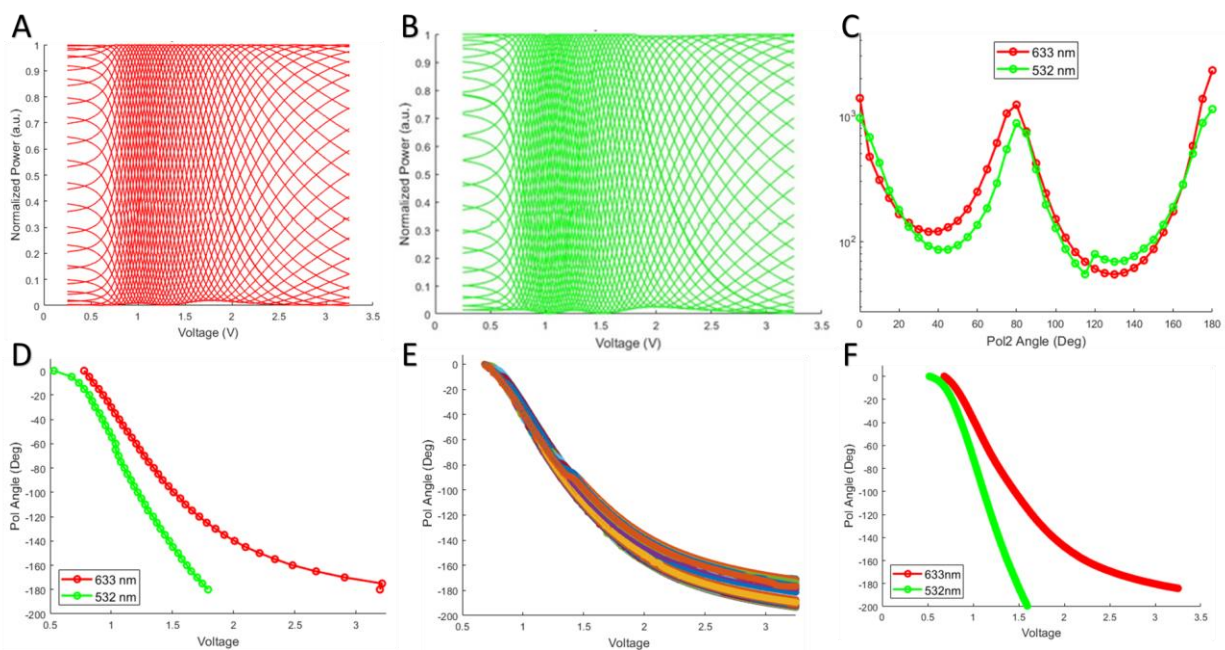


Figure 3-21. Calibration of the LCPR using the proxy-dipole method. (A) Normalized traces of power transmitted (through proxy dipole) vs. LC voltage, with each trace being a different calibration polarizer angle, in five-degree steps, for the 633 nm BlueSky diode. (B) The same as for (A), but for the 532 nm BlueSky Diode. (C) The resulting extinction curves, as a function of calibration polarizer angle, for the 633 nm and 532 nm diode data in (A) and (B), where each data point is result of dividing the maximum transmitted power value by the minimum for a given trace. (D) Each point is from one trace, designating the LC voltage at which transmission was minimum for a giving

polarizer angle. (E) For just the 633 nm beam, the given interpolated traces. (F) For both beams, the results of averaging all interpolated traces, given a much better calibration curve that was afforded solely by the data in (D).

3.10.7.2. The automated polarimeter method

As may have been evident in the previous section, calibrating with the proxy polarizer method is quite an ordeal! Additionally, once accomplished, if the calibration indicates that alignment is not good enough, an adjustment must be made, and then the entire process done again (though the next section outlines some ideas on how this might be streamlined)! Luckily, an automated polarimeter (see *Section 3.4.4.1* for further details) can speed up this process considerably.

The set-up for polarimeter-based LCPR calibration is the same as it was for the LCPR alignment (**Figure 3-9**). The polarimeter conveniently has threads for optical rails, and thus it can be inserted into the objective holder quite easily. This means that the objective cannot be a part of the alignment, but given that the polarimeter requires incoming light to be collimated, this is not a prohibitive issue anyway. The weight of the polarimeter on the objective diving board also skews alignment a little bit, but this should have minor impact for this purpose.¹

This set up is much faster to calibrate with than the calibration polarizer method, as we don't need all those steps to diagnose polarization state. The polarimeter can in fact read out azimuthal and elliptical information for us in real-time! Thus, we can simply **sweep the LC voltage**, and take the outputs, process them a bit, and determine what the **extinction vs. a function of angle will look like**. Another very convenient aspect to the polarimeter is that optics can be rotationally adjusted (or voltage adjusted) on-the-fly, with the output immediately known.

Admittedly, the necessary conversion of polarimeter outputs into the parameter we care about (Extinction ratio, or in our experiments, depth-of-modulation) is not quite straightforward. After much

¹ The polarimeter GUI has a convenient built-in tool for alignment diagnosis (which works by detecting odd harmonics of power modulation that should not be present, as a function of QWP rotation speed within the device), which can be used to check if things are set up well enough!

time staring at polarization equations and the Poincaré sphere, I came up with the following two equations, which conveniently relate both depth-of-modulation (M; **Equation 3-3**) and extinction ratio (ER; **Equation 3-4**) simply to the third Stokes vector (which describes the deviation of polarization from the equator of the Poincaré sphere, with the north and south poles being CPL).ⁱ

Equation 3-3

$$\mathbf{M} = \mathbf{1} - \mathbf{S}_3^2$$

Equation 3-4

$$\mathbf{ER} = \frac{\mathbf{1}}{\mathbf{S}_3^2}$$

The automated polarimeter is extra useful in consideration of the new M Squared laser system (*Section 3.6*). This is because as opposed to the HWP rotating method of polarization control, this LCPR system is quite wavelength dependent! And unfortunately in a quite nonlinear way, which Thorlabs does not provide data for.ⁱⁱ Thus, an automated program can be run to sweep the LC for a given pump wavelength, and then the wavelength can be swapped, and the program run again.ⁱⁱⁱ This procedure was in fact done for the EMM1950 (visible mixing module) of the system, and the results are shown in **Figure 3-22**. Shown are the azimuth as a function of LC voltage, for many wavelengths spanning 515-660 nm, the ellipticity, and the max M as calculated from the equation above. Next shown are the full range of azimuth measured, the range of ellipticity over the sweep, and the corresponding worst-case M (for the highest ellipticity) for each wavelength.

ⁱ I am not positive about these equations as I do not have an analytical derivation for them, but they follow forms I would expect, and produce results that make sense. I encourage the reader to check these for themselves, or try to come up with them (or better equations) independently!

ⁱⁱ Even if they did, it would be prudent to calibrate it ourselves anyway.

ⁱⁱⁱ Or if someone wants to invest the time, the whole process could be automated! That would be super nice, and then could include all the post-processing of data within the calibration code.




Figure 3-22. Polarimeter calibration data for LCPR as a function of M Squared Wavelength. For the left three plots, from left to right are the azimuth vs. LC voltage (with each sweep being a different wavelength), the measured ellipticity vs. azimuth, and the maximum calculated DOM vs. azimuth. For the right three plots, from left to right, are the azimuthal range vs. pump wavelength (which I think trustworthy), the range of ellipticity vs. pump wavelength (where only outside of the yellow region I trust the absolute value), and the worst-case scenario for DOM vs. wavelength (where only outside of the yellow region I trust the absolute value).

Notably, this polarimeter is only specified for wavelengths of 600-1080 nm. Thus, the measurements for the M Squared system below 600 nm are suspect. Based on data adjusting the set-point for in-range wavelengths (additional not shown), I have come to strongly suspect that the **azimuthal** measurements are still rather reliable outside of the specified range, but the **ellipticity** measurements are wrong. The latter is supported by the above data, in the yellow-shaded regions, where the trends for measured ellipticity, and subsequently worst-case M, reverse as the polarimeter continues to be set to 600 nm while the true wavelength trends bluer. This makes some intuitive sense, as it is to my understanding the QWP in the polarimeter is the wavelength limiting factor. It is a true zero-order QWP that is already quite chromatically dependent, which the polarimeter already has built-in correction factors for in the specified wavelength range. It might be that we can use data collection such as the above to further extend this correction factor, but I am unsure of how accurate and precise this would be.

3.10.7.3. Using Jones matrices for LCPR alignment

Jones matrices can be quite useful in understanding the LCPR alignment, and could perhaps even be used to predict necessary alignment adjustments. This mathematical tool is discussed in much more detail in Chapter 2. Most codes employing these matrices I have written so far are comprised of many-layered nested For Loops, which step through all possible parameter spans of interest. Just one output from such a code is shown in **Figure 3.21**, for a case where the LCPR's QWP is actually a "0.2725 λ plate". Shown are four different traces, each with the QWP and Polarizer correctly aligned, but with the LC being stepped iteratively in one-degree steps away from perfect alignment. I will not delve into all the features here (such as peaks and troughs, distance between peaks, depth of troughs, etc.), but simply draw attention to the fact that for a given set of optical parameters and alignments, extinction across the full range of possible angles can be predicted! Thus, **alignment tolerances** can be assigned, ex. "all optics must be within 1 degree of specified angle to achieve the desired level of polarization control" might be a conclusion from such an analysis.




Figure 3-23. Example Jones Matrix Simulation. Here, for an imperfect QWP with 0.2725 wave retardation, perfect QWP and LP alignment, with varied LC alignment.

In an ideal world, we would have a more advanced Jones Matrix code that would take LCPR polarization data (either from the old method or the new method of calibration), and **tell the user how**

to adjust the optics (ex. “turn the polarizer 1 degree counterclockwise as seen from above”) to improve polarization. It is not yet fully clear to me the best approach for doing this, though I have made some coding attempts. There are two routes that come to mind. One is a brute force approach that searches all parameter space with nested For Loops until one or more outputs is found that matches the calibration data, and then adjustments are recommended based on that. The other method would perhaps be numerical (I don’t think this will ever be purely analytical), but instead a code would solve for parameters based on data. This latter option is far more elegant, but also would require writing symbolic fitting code, I think. While my initial Jones matrix codes were MATLAB based, the newer codes I have worked on (for the LCPR as well as polarization paddle predictions) have been in Python. And while Python does have linear algebra packages, they tend to solve for matrix elements, and here, we want to root out parameters such as optic angles and retardations that are embedded within many matrix elements in quite complicated ways. Thus, symbolic coding is probably needed, and not just a simple linear algebra package. I hope that a future group members cracks this problem, as such deterministic code would be a fantastic aid to both alignment and intuition of polarization optics!

3.11. Controlling the Microscope (Code and Software)

This section primarily describes the “Photothermal Master Code”, written in LabVIEW, that we use to control the second-generation photothermal microresonator setup. Some parts of this code are imported from work as early as Kevin Heylman’s, while some are much newer (much written by both Erik Horak and by me). This code is essentially always under-construction, as new optical components are added to the instrument, and as new functionality is desired. The front panel of the code has many suggestions for further improvements written-out in text boxes, and the block diagram has many comments explaining various components and subVIs.

3.11.1. Measuring error slopes

The Error Slope is a crucial parameter to converting the output of the lock-in amplifier to a resonance shift. Thus, this tab is used before locking to any resonance, at least if the user wants the calculated resonance shifts to be quantitatively correct. For taking an error slope, a triangle voltage wave is applied to the probe laser piezo, which finely sweeps the laser wavelength over a desired resonance to reveal an error slope.

There are a couple of different ways that Error Slope has been calculated within our photothermal codes. The current way is the most straightforward way, which uses a slope unit of “**pm/V**” (picometers per volt). Thus, a given measured voltage from the lock-in is easily converted to a resonance shift by the slope that was locked to. For this slope to be correct, the user **must** have the correct **amplitude** and **frequency** of the voltage sweep described above. This is because what is actually measured for the error slope by the DAQ (in the “Error” channel) is just a **voltage over time**, meaning that the true units of this measurements are actually **V/s**. This voltage has a different meaning than the voltage in the pm/V unit described above! Also accounting for the **piezo calibration factor** (*Section 3.10.1*), the sweep parameters allow for this V/s to be properly converted to the more straightforward pm/V units . Some legacy versions of the photothermal code instead measure Error Slopes using the V/s units. In this case, the correct sweep parameters from above are needed in later steps (such as the Mapping tab) for calculating resonance shifts, whereas the current version of the code obviates the need to include these parameters later on.

Conveniently, the user can also select “**Auto-fill**” to have the Error Slope variable above the tab-structure auto-filled in from the measurement. If the user selects for the data to be saved, a MATLAB script is executed within the block diagram, which saves a data structure and a figure as well, with the figure also indicating with a vertical line where the user selected for fitting (**Figure 3-24**).

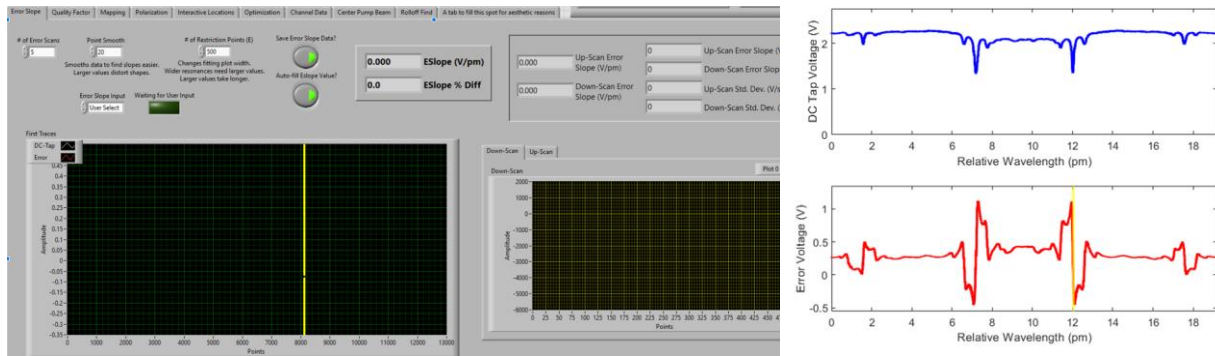


Figure 3-24. The “Error Slope” tab of the master photothermal code (version 3g). **(Left) Part of the front panel. (Right) An example figure of what the embedded MATLAB script generates, showing both the DC tap and error channels. The yellow vertical line indicates where the user placed the cursor for error slope measurement.**

3.11.2. Measuring quality factors

To measure Quality factor (Q-factor) of a resonance (See *Section 2.4* for more information), the “Quality factor” tab of the master code is used. Measuring Q-factors of resonances is actually not necessary for calculating resonance shift in our experiments. However, since Q-factors are an important figure-of-merit for microresonators in general, as well as the sensitivity of our photothermal measurement, we often measure at least a few of them for a given resonator.¹ It is also worth measuring Q-factor when taking any particularly important data (i.e. that might be published), as it is common to include Q-factors of resonances in literature.

All that is needed to calculate a Q-factor is the center wavelength and the linewidth. For the central wavelength, an estimate will do, for example when using a probe laser that tunes 635–638nm we just use the value 635 nm. The linewidth is what must be measured, and this is done while employing a voltage triangle-wave to piezo of the probe laser via the frequency modulation input on the controller. This measurement should be done with **PDH off**, as PDH will impact the apparent linewidth of the resonance.

¹ For microtoroids it was common to record all resonances, their central wavelength from the probe course-tuning, and their Q-factors, for future reference. For microbubbles, this would be a herculean task (or one needing much automation), as there are just too many resonances to keep record of.

Once the response of the “DC Tap” channel wavelength sweep is measured by the DAQ, the user is then prompted to select where in the sweep is to be fit with a Lorentzian. This user-input is not necessary for isolated resonances (where the automatic resonance-selection option can be used), but for resonators with congested mode spectra such as microbubbles, it is helpful to ensure that the correct feature is fit. The user has some options for the Lorentzian fit that the code applies to measure a linewidth for the resonance, which can be adjusted from the front panel as necessary to ensure a good fit. A “save” button allows the user to save the Q-factor data. A MATLAB script saves a data structure and a figure as well, with the figure also indicating with a vertical line where the user selected for fitting (**Figure 3-25**).

The voltage-to-wavelength calibration parameter is crucial for converting the measured DAQ times to wavelengths. Additionally, as described in the Error Slope tab description above, the frequency and amplitude of the voltage sweep must be correctly input into the code (they are not measured, though this could be a **future improvement** to the code!).

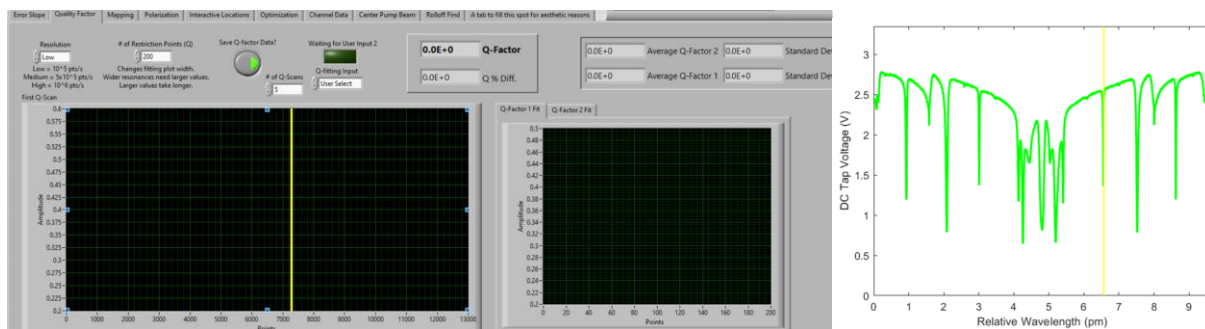


Figure 3-25. The “Quality Factor” tab of the master photothermal code (version 3g). **(Left) Part of the front panel. (Right) An example figure of what the embedded MATLAB script generates, showing the DC tap channel. The yellow vertical line indicates where the user placed the cursor for Q-factor fitting.**

3.11.3. Photothermal mapping

The “Mapping” tab (**Figure 3-26**) of the master code is an important one, and hopefully fairly intuitive, though it has features worth pointing out. The most typical map is a square-shaped map, for which the user defines center coordinates (in terms of galvo voltage, as the code currently operates), as

well as a step size (in terms of actual distance in micrometers). The mix of voltage units and distance units is a good approach for our experiments, since the resonator won't necessarily stay in exactly the same spot in space, especially day-to-day, requiring different voltages for mapping certain objects, but it is more intuitive to give a desired map size in terms of distance units. The latter requires a calibration parameter that converts the desired distance into voltage units for the galvo mirrors. With the center point and desired map size, the code then uses a nested For Loop structure to generate arrays of X and Y voltage coordinates for the map. Another way to do this would be to use logic within the mapping For Loop to calculate the position as each pixel comes up, however I prefer to deliver pre-calculated values that are indexed as needed (though we have used both paradigms in different parts of the code). The code also includes the option to set a rectangular map with different spans in X and Y, which can be useful for particularly oblong microresonators to save on map acquisition times.

A particularly handy feature is the ability to do "multi-object" maps,¹ which allow the user to take fine-resolution maps around multiple objects on the resonator in one code execution, without having to take the horrendous number of in-between, useless pixels. Using the "Interactive Locations" tab (more on that further below) is efficient way to select the desired coordinates. A **future improvement** that would be useful would be a version of this code that lets the user define regions to *excise* from a large map, as this would help avoid large objects that would ruin the resonance-lock at high pump powers.

There are additional layers to the Mapping tab that can be included. For example, the same map can be run multiple times for a series of pump polarization angles to ascertain the polarization-dependence of object excitation (this is described further below in the "Polarization Rotation" section). Another layer is number of "Runs" selected. This corresponds to the outermost For Loop in the block diagram, and lets the user execute multiple iterations of the code.

¹ This multi-object mapping was largely developed by Erik Horak.

For better or for worse, our code leans quite heavily on nested For Loops for iterating through different parameters.¹ In the current form of the code, the order of these loops is rigid. For example, the user must have any multiple maps *within* polarization sequence, and those themselves *within* the run sequence, and can't swap this order. As found elsewhere in the master code, there are a suite of "Stop" buttons available, each corresponding to a different layer in the nested loop structure, to allow the user to quickly stop the code without employing the "hard stop" LABVIEW stop button (which should be avoided! This will require restarting the LabVIEW, and sometimes even equipment!).

The Mapping tab has a MATLAB script embedded within the block diagram, which takes the measured data and processes it into figures, as well as saves a data structure of the many variables in the code. If "smart offset" is selected in the front panel, this MATLAB tab will also put the output maxima for however many objects are being mapped (one per "map region") by taking the coordinate of each map with maximum resonance shift and feeding that back to LabVIEW to be the new center coordinate for the next time that coordinate is mapped (ex. during a polarization sequence or multi-run sequence).

There is also a *second* MATLAB script within the Mapping tab, which is run if a polarization sequence was used. That script takes a 2D-Gaussian fit of the object in each map to obtain a maximum intensity, and then takes those maximum resonance shifts for each map, and the polarization angles for each map, and calculates a depth-of-modulation for the object. Exemplary data for a photothermal map and a mapping-polarization sequence are shown in **Figure 4-1** and **Figure 4-3**, respectively.

¹ I think we've reached at least 7 loops nested in one spot. Whoops.

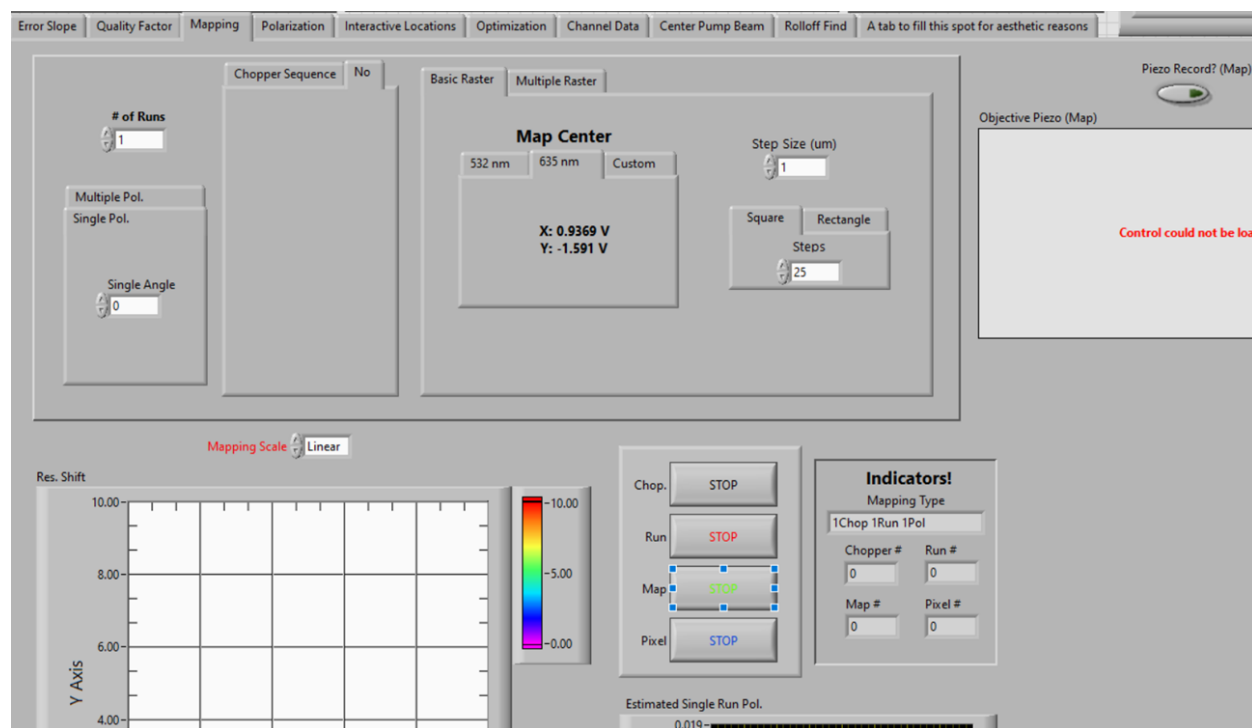


Figure 3-26. The “Mapping” tab of the master photothermal code (version 3g).

3.11.4. Polarization rotation

A nano-object can be determined to have a well-defined excitation transition dipole by studying it with rotating linearly polarized light. A faster way to acquire this data than using the pol-sequence in the mapping tab is to instead simply park the pump beam centered on an object of interest, and rotate the linear polarization. We refer to this as a **polarization sweep**. The current version of the master code does this by applying a voltage to a tunable liquid crystal, while much older versions of the master code rotate a HWP using a motorized rotation mount and APT software. **Figure 3-27** shows a snapshot of the polarization sweep code, with some of the setting options displayed.

The settings options in the code allow the user to set the breadth, resolution, and speed of polarization rotation, with a build in speed-governor of sorts that lets the user know when the requested speed has been overridden by the software (if the calculated sweep was unrealistically fast). A “dummy start” will also add extra points to the beginning of each trace to ensure that important information is not missed. An important parameter is the **wavelength** set in the code, as this will determine what

calibration vector is selected to determine the voltages that must be applied to the liquid crystal. As with mapping, this sweep can be executed for N “Runs”, with N defined by the user.

While certainly faster than iteratively mapping an object, the polarization sweep runs the risk of being convoluted with any spatial drift of the object and/or resonator relative to the pump beam. Thus, I wrote a custom algorithm for centering the beam the object. The user selects centering step size, number of centering steps, and number of refinements. The code then uses these parameters to iteratively make smaller and smaller cross-hairs, centering the beam at the point of highest signal for each sweep of each axis.

The Polarization tab has a MATLAB script within the block diagram, which takes the measured data and processes it into figures, calculates a depth-of-modulation value, and saves a data structure of the many variables set the code. Exemplary polarizations sweep data is shown in **Figure 4-3**.

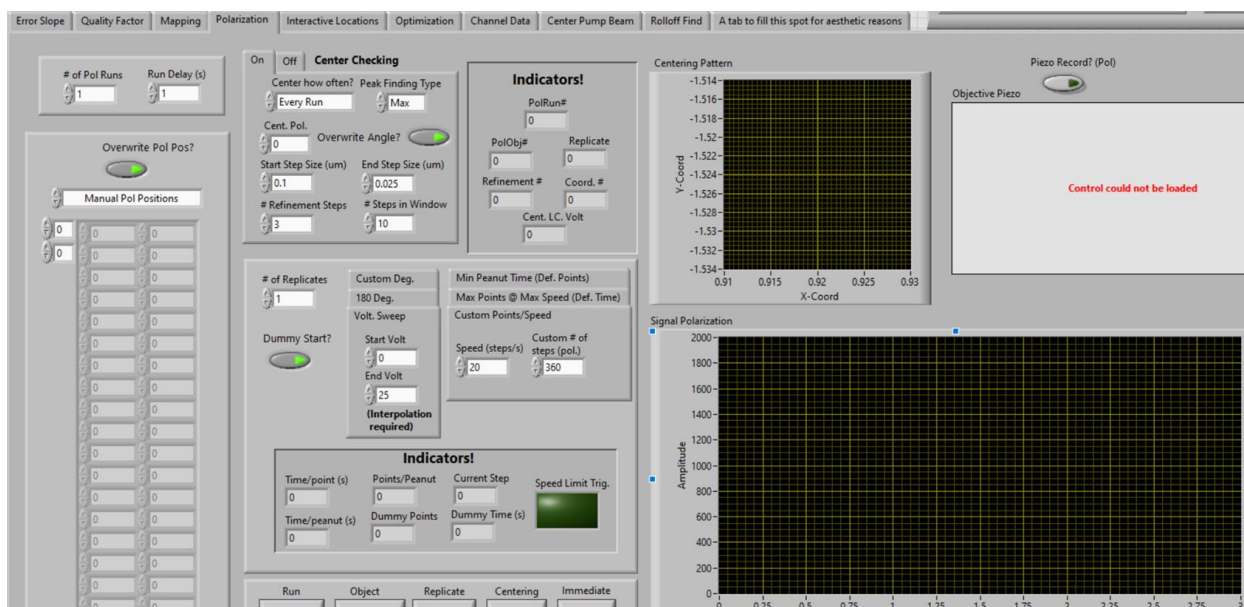


Figure 3-27. The “Polarization” tab of the master photothermal code (version 3g).

3.11.5. Under construction: spectra acquisitions

The version of the master photothermal code shown in the other screenshots in this chapter (Version 3g) is a version that does *not* include control of the M-Squared laser system, and it is currently the

functional, most-updated version of the code. With the dawn of the M-Squared era for photothermal experiments, however, control of the laser system within the master photothermal code will be crucial. Below I outline some important considerations and some of the additional components to the code that are started but not yet finished.

The most important thing to know is that the M-Squared laser system accepts TCP commands in JSON format (TCP = Transmission Control Protocol, JSON = JavaScript Object Notation). This is made easy by the fact that the entire M-Squared system is run by a set of “ICE Blocks”. There is an ICE block for the Equinox (pump laser),¹ there is another for running the SolsTis (Ti:Saph), and there is one for each mixing module. Each Block is essentially its own computer with its own IP address, which can accept relatively straightforward TCP commands and then know what settings to change, control loops to implement, variables to report on in response, etc. In our system, there is a hierarchy of Blocks. The Equinox Block can be controlled on its own. However, if the SolsTis block is sent commands, it then controls the Equinox Block as necessary, of its own accord.² Furthermore, if either the DFG or SFG Blocks are sent commands, they then send necessary commands to the SolsTis Block, which subsequently sends commands to the Equinox Block. The mixing Blocks can also control the Equinox directly.

This is all coordinated through an ethernet splitter and a central laptop. The most straightforward way to control the M-Squared system is with the manufacturer-created GUI, which runs from a web browser such as Google Chrome, on a computer connected to the ethernet splitter which connects all of the Blocks. A tab can be opened in the browser for each Block by typing in the IP (Internet Protocol) address of that Block. There are multiple levels of control for each block (ex. lab user, service engineer, factory engineer), the passwords for which will be known by those as needed already.

¹ Or two, depending on the version of the system you have

² Especially in physics and engineering context, often this is referred to as one laser being “slave” to another laser. I avoid this terminology, but I include this to let the reader know it is somewhat common in certain fields

An important note is that as long as a computer is connected to the Blocks via ethernet, it can use the IP-based GUI. However, if one wants to use a *different* software interface (ex. custom Python or LabVIEW code) to send commands to the Blocks, another change is needed. Namely, the “ports” of the Block need to include the IP address of the computer sending the command. Specifically, this IP needs to be an IP set for the **M-Squared Network** (the local network running between the Blocks and laptop), and *not* the building ethernet or Wifi network.¹ Our group has internal protocols for doing this, though the limited number of ports available given Block factory settings does limit utility. Most recently, further work by group members^{##} has made the system even more flexible by employing a router which assigns IPs to computers on the local network, though very occasionally port numbers still need to be adjusted.

Here, I discuss solely LabVIEW coding, as this is already what we have developed in great detail for photothermal code. Python development is ongoing by others in the group for the M Squared system.

With the EMM1950 (SFG) module being most relevant for visible-wavelength single molecule photothermal experiments we want to pursue, this has been my coding focus. I have written simple LabVIEW scripts that can change both the wavelength of the EMM1950 module, and the power of the system’s pump laser (though we may, for sake of stability, choose to keep Equinox power at maximum, and employ some sort of variable attenuator as needed). This has been helped greatly by the “control loop” scripts provided to us by M-Squared, as well as the documentation for the different system components, which includes for example lists of accepted TCP commands for the Blocks.

This control of wavelength will eventually hopefully come into play in the following ways in the master photothermal code:

1. Iterative maps at different wavelengths.

¹ Though the whole system could perhaps be migrated to the building network, this means that all Block IPs need to reliably exist on the network, and this also introduces security concerns of anyone with access to the building network being able to access the M-Squared system to some extent.

^{##} Shoutout to Alex Fairhall and Brandon Hacha.

2. Iterative polarization sweeps at a specific coordinates, at different wavelengths.
3. Acquisition of absorption spectra.

Rather than embed wavelength control in each LabVIEW tab,¹ I decided to place the wavelength settings outside of the tab structure (**Figure 3-28**), such that an array of wavelengths would be generated to be employed in whatever tab necessary. An area that I have not started on is including wavelength control in the “Optimization” tab. This could be useful, but will be a lot of work to code. For all the above, **further coding is needed**, including in the MATLAB scripts within the LabVIEW tabs. More details on how each tab would ideally utilize wavelength control:

(1) **The mapping tab** will be the most robust for measuring absorption spectra of analytes, as it will provide information needed to account for the change in spot-size as pump wavelength changes, as well as account for any spatial drift of the object/resonator. Mapping should also indicate if there is spectral-dependence of photothermal background. However, mapping also takes the longest, and thus may not be possible for experiments where finer temporal resolution is needed or where experiment duration is limited.

(2) **The polarization tab**: As described above, mapping does not account for change in excitation dipole orientation if the object rotates, unless the user employs the polarization nested loop, which adds significantly to the time-cost of measurements. Neither will mapping account for possible pump wavelength-dependence of optimal polarization. Thus, the polarization sweep tab should also have the ability to take polarization sweeps at iterative pump wavelengths, very similar in code structure to the wavelength-iteration in the mapping tab. This process has been started, but not completed.

(3) Lastly, it would be ideal to also have an option to simply take the **absorption spectrum** of an object without having to also iterate through spatial or polarization parameters. This could be possible

¹ As I have for example, for polarization control, which leads to an unnecessary number of local variables within the script.

by setting a map with pixel-size one, or a polarization sweep with only one polarization value, however a more customized tab is desired. The photothermal master code for the NIR setup includes spectra capabilities, which while not written for the M Squared laser, may still be helpful as a reference. A feature that may need adding is a centering algorithm, such as the one in the polarization script, which would at various points in the spectrum acquisition ensure that spatial drift isn't convoluting the results. Additional corrections and/or control components will be needed to account for varied spot size if the spectrum is wide enough in wavelength. As with most other tabs in the master code, an embedded MATLAB script for further processing, figure making, and datastructure saving would have great utility.

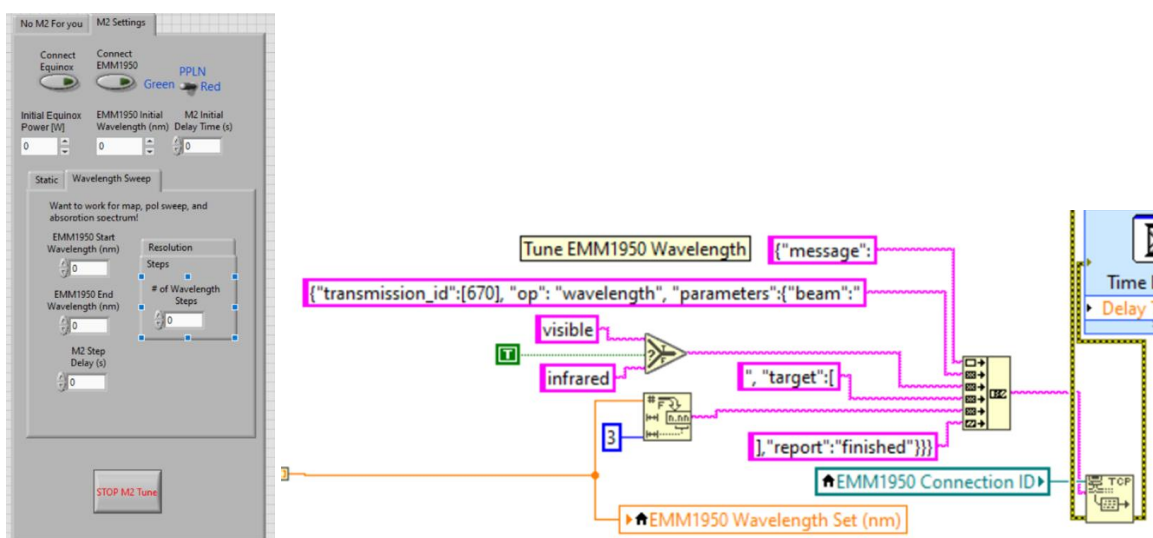


Figure 3-28. Control of the M Squared system in the photothermal master code (version 4a). *Left*: prototype control tab in the front panel. *Right*: Example of block diagram structure for sending TCP commands in JSON format. Receiving and unflattening the JSON commands is substantially more complicated than sending commands, but example code from the company is helpful, and work is underway to develop control loops within the master code (ex. for confirming that pump wavelength is set correctly via the wavemeter).

3.11.6. Interactive locations tab

The “Interactive Locations” tab (Figure 3-29),¹ is very handy when the user wants to study a variety of objects iteratively in an automated fashion. Essentially, the user may upload a photothermal map dataset

¹ Mostly coded by Erik Horak!

generated previously from the Mapping tab, and display it within LabVIEW in linear or logarithmic color scale. The user can then double-click at various positions on the map, typically where “hotspots” (objects) of interest are located. LabVIEW records the galvo voltage for each object in a pair of vectors. Positions can also be uploaded from a “positions” text file, which is saved for every photothermal map (including a file with all positions if a multi-object map is executed).

Even after the code has been stopped, these positions can then be uploaded in the “Mapping” or “Polarization” tabs for multi-object maps or sweeps, as the positions have been saved in the front panel of the interactive locations tab. In fact, vectors of positions can be selectively pulled from any of these three tabs into either of the others by varying the selection at the top of the respective array in the front panel of the relevant tab.

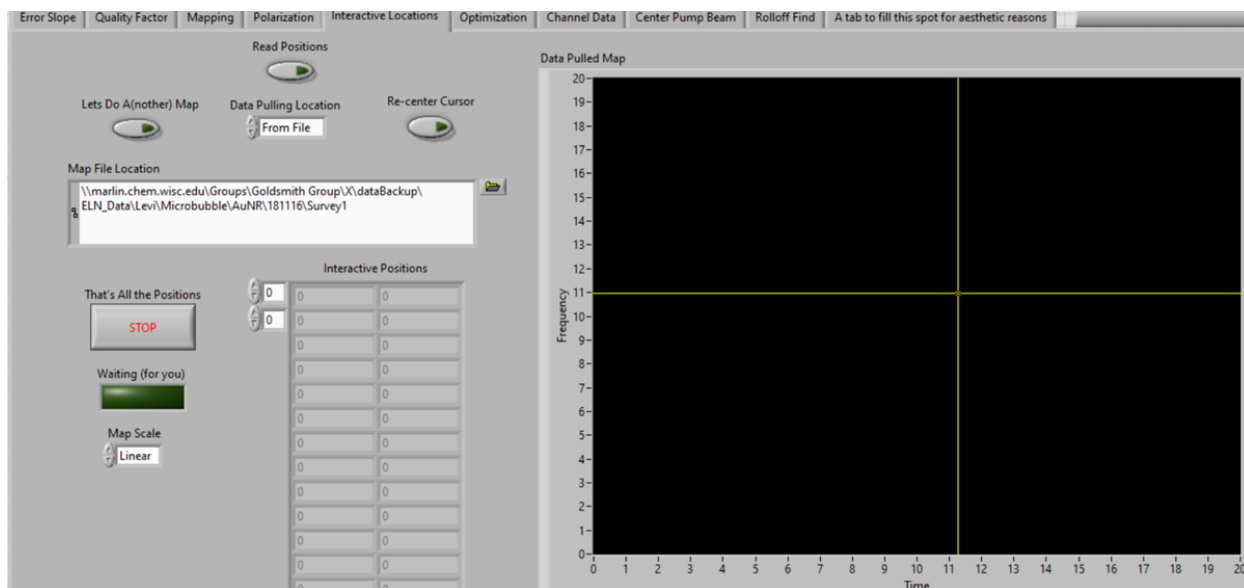


Figure 3-29. The “Interactive Locations” tab of the photothermal master code (version 3g).

3.11.7. Optimization tab

The Optimization tab (**Figure 3-30**Figure 5-5) allows the user to set the galvo coordinates to park the pump beam on an object on the resonator, and then optimize (almost always maximize, unless that gives too much signal for the electronics to handle) the photothermal signal being measured from that object. This tab, initially designed by Erik Horak and subsequently added to by me, uses an **event structure**. This

is not quite a state machine (and is definitely not a queued state machine), however, it has some very similar elements to a state machine coding paradigm.

The most frequently adjusted parameters for this tab are the galvo voltages, the focusing-piezo voltage, and the pump power. The galvos and the piezo are iteratively optimized, and both parameters have multiple ranges of optimization from coarse to fine. The pump power is adjusted as needed, since sometimes for a desired lock-in sensitivity setting, an acceptable input voltage becomes too high once the signal is optimized.

Another useful parameter to adjust is the LC voltage, which consequently changes the angle of the linearly polarized light at the resonator. This is not typically useful for dust, but may be useful for determining what polarization angle at which to interrogate an analyte with a defined excitation transition dipole.

The tab also contains some control capability of the optical chopper frequency, which could be useful for finding the “sweet spot” between increasing amplitude modulation frequency to reduce $1/f$ noise and maintaining the necessary level of heat transport in the resonator. This component to the Optimization tab is currently only partially working, and needs some improvement. In the future, **adding M-Squared wavelength tunability would be useful for this tab.**

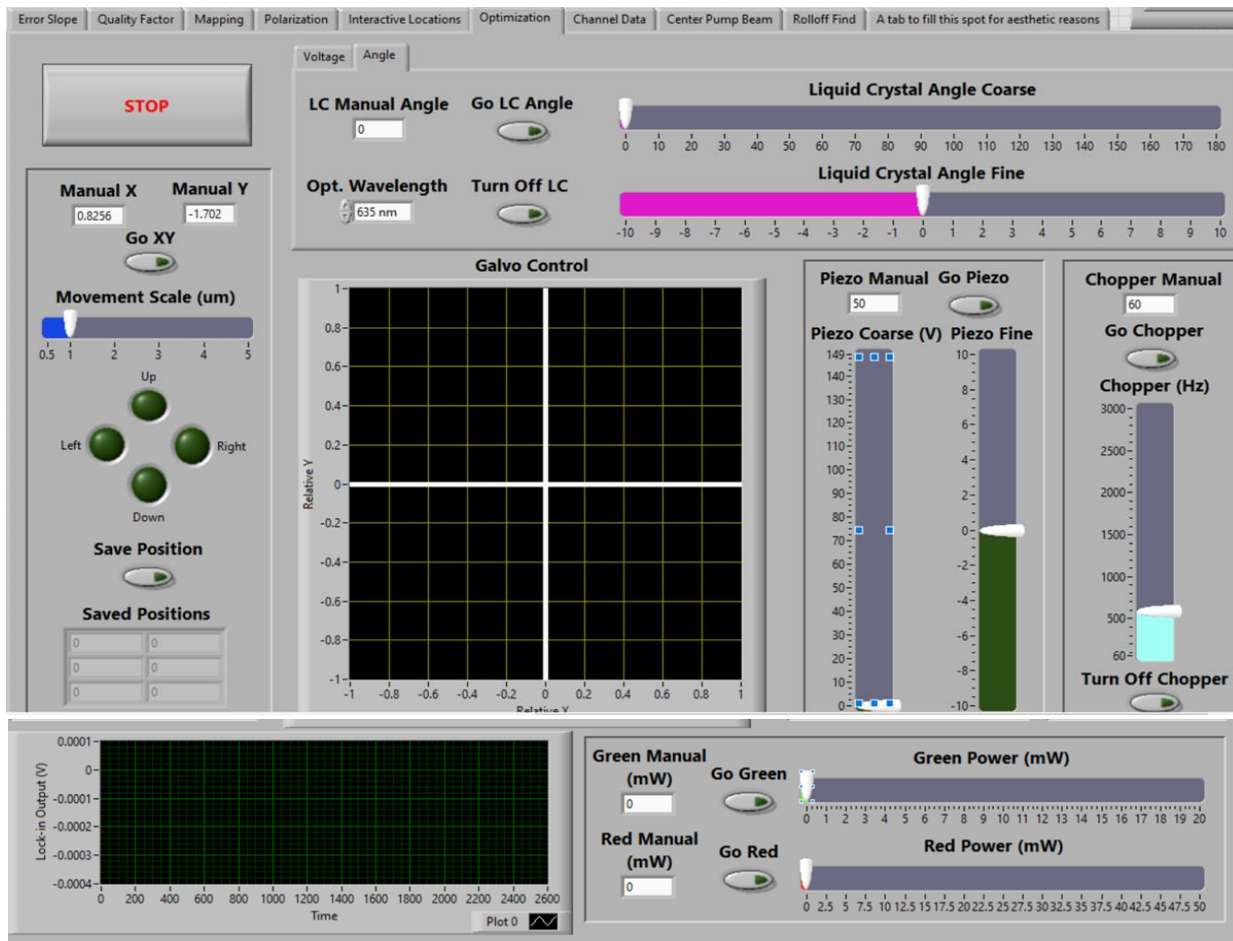


Figure 3-30. The “Optimization” tab of the photothermal master code (version 3g).

3.11.8. Time series collection tab

The “Time Series” tab is designed to take data from one or two DAQ input channels for a prescribed time at a prescribed temporal resolution. Channels typically measured are those associated with the DC Tap, Error Signal, and the Lock-in output. Typically this tab is used when something interesting or unusual is observed in one more data channels, and the user wishes to record it for later. This tab also includes a MATLAB script, which saves a figure and a data structure of the obtained data, if the “Save” option is selected.

3.11.9. Pros/cons of transitioning to other control software

LabVIEW isn't the only game in town for instrumentation control, even if it is what almost all of our photothermal measurements use in our group, and even if it is very common in academic circles. MATLAB can also be used to control instrumentation, and although this would come with the same academic-applicable constraints as LabVIEW, it might lead to more efficient integration of instrument control and data processing. Python is undoubtedly the most ubiquitous, open-source programming language, and with some significant effort, our photothermal experiments could be transitioned to this platform. A C-language could also be used. C++ is common, though C# drivers are becoming more commonly available, at least from Thorlabs, for various pieces of equipment. If very fast acquisition and/or feedback is necessary, perhaps a FPGA (Field-programmable gate array) would be worth the time investment to design. Lastly, sometimes asynchronous control of different components in an experiment is good enough, and thus multiple GUIs or third-party scripts based on different coding language can be employed. While there is no apparent need currently to transition to other software platforms, if the photothermal code is getting significantly revamped or expanded anyway, transitioning to another language (likely Python) may be worth the effort!

3.11.10. Quirks and tips for various drivers, GUIs, etc.

Here, I outline a few particularly nasty issues I have faced in coding the master code and calibration codes for the current microresonator spectrometer set-up:

(1) **Power meter drivers:** A particularly annoying quirk is encountered when using Thorlabs Optical Power Meters. These meters can either run on legacy drivers (PM100D) or relatively newer drivers (TLPM). We have various home-written scripts that use each type of driver. The important thing to note is that the software download for these meters comes with a **driver switching** application ("Power Meter Driver Switcher")¹, which can be used to switch between the two driver types for a power

¹ Though I will leave finding this downloaded application as a scavenger hunt for the reader.

meter box. Even so, we have still sometimes found that certain box/meter pairs just do not want to talk with certain computers, and we are not sure why this is the case.

(2) **COM Ports:** COM ports (communication ports) are one standard way of talking with equipment via a USB interface.¹ Importantly, the identity of a COM port (if its “COM1” or “COM2”, etc.) can switch with computer reboots, and reconfiguring USB connections, especially if USB hubs are being used (which is a given, considering how many components make up this microscope). Thus, attention must be paid in code to either hard-code the serial number of the particular for each piece of equipment (the less difficult but more rigid solution), or to carefully sort through what equipment each COM port is talking to when a code is executed (the more difficult but more flexible solution).

(3) **Simultaneous control of Chopper and Liquid Crystal:** Control of the variable liquid crystal controller through the master photothermal code is crucial. Control of the optical chopper is not crucial, though I have found it convenient! Unfortunately, these two very similar grey boxes that Thorlabs sells seem to communicate with the computer so similarly via COM ports that LabVIEW can get quite confused about which grey box it should be communicating with. While serial numbers for equipment can be hard-coded as described above, the more general, challenging route is to have LabVIEW perform an exercise that identifies what COM port is talking to what type of device, and then correctly assigns the COM port for each. We have some prototype code that I wrote for doing the latter that could be developed further if desired, but generally the former method suffices.

(4) **APT vs. Kinesis:** APT is an older version of motion control software from Thorlabs, and it is still employed in many of our codes. However, the APT interface has been legacy software for quite a while at this point,² in lieu of the newer Kinesis interface. Thus, whenever writing new code or overhauling older code significantly, it would be wise to continue migrating LabVIEW codes to operating

¹ Though I think this is a legacy term from the true “COM” ports that used to be common for interfacing equipment with computers.

² Perhaps for as long as I have been in graduate school.

via the Kinesis interface, as it will be better supported in the future, and is less glitchy than APT in my experience. Thorlabs software manuals for various Kinesis-compatible equipment, and existing group codes, can be consulted for representative frameworks.¹

(5) **Elliptic Software:** A new addition to the photothermal instrument, since the M Squared system was added, is a piezo-electric resonator motor device (Thorlabs ELL14K). The goal of this device is to allow for automated switching of ND filters for the EMM1950 module once the fiber-launch path is covered by a safety enclosure. The Elliptic software GUI is enough for current needs, but while the device seems to be engineered to be controlled with C# based code, luckily some other folks in our group have already written LabVIEW code for this line of products that can be borrowed from as needed.

3.11.11. Future code improvements

Beyond the above-suggested improvements to the photothermal master code, the most pressing of which being control of the M Squared laser system for spectral acquisitions, there are a couple of other possible future improvements I would like to suggest.

If the code were to be completely overhauled, I suggest writing it as a queued state machine, which will employ event structures, while loops, etc. so that the code keeps running once it has been started, and the user can toggle between different options, with items clicked on joining a queue for the code to perform. Besides the attractive process flow for this, the way in which different processes would be folded into an event structure would make the block diagram far more compact, and hopefully more understandable and easier to edit.

Whether or not the code is overhauled in such a way, there are a variety of processes that could be further automated, for example the “optimization” tab could have an “auto-optimize” feature. Also, acquiring additional data from the DAQ during acquisition of resonances shifts (ex. from the error and

¹ And there are many example Kinesis codes available if you can find the download folder on the Thorlabs website. I have been satisfied with the package, except with regards to the motorized polarization paddles, which I have yet to get to work via LabVIEW using Kinesis.

DC Tap channels) could be helpful for post-correcting data, ex. in the situation where an error slope is varying over time.

3.12. Taking Photothermal Data

3.12.1. Aligning the microresonator, taper, camera, and pump beam

With probe beam path connected, pump beam path built, electronics up and running, calibrations done, and code understood, we're ready to actually **use the microscope to take some photothermal data!** So...how shall we go about doing that? Here I'll give a suggested set of strategies, specifically geared towards microbubble resonators as I currently use them. Note that I **will not detail how to work microresonators themselves**, and will leave that training for the reader to seek elsewhere.

First, we need to align a few more components of our experiment. There are five things that must be aligned: (1) the pump beam focus piezo, (2) the pump beam position, (3) the witness camera, (4) the tapered optical fiber, and (5) the microbubble resonator. They should be aligned in this order. Once a given taper/resonator system is aligned, re-aligning it in the subsequent days will be much easier.

(1) Pump beam piezo: Once you (later) hit the stage of optimizing the focus of the pump beam to an object on the resonator, you'll need to tune the voltage of the focusing-objective piezo. The range of this piezo is 0-140 V. Thus, it's a good idea to center this voltage somewhere in the middle, ex. 70 V, so that when adjustments are needed later, they can be in either direction.¹ Based on how you align your system, you may find a different voltage more ideal for centering the piezo to start with. While it is possible to make course adjustments with the focusing knob for the objective, this is difficult to do reliably and precisely. Additionally, changing the height of the objective too much from the initial alignment may impact the transmission and quality of the beam scanning system, as the back aperture of the objective deviates from the focus of the second lens in the relay.

¹ Seriously, if you get in the habit of this, you will save yourself many sad-realignments

(2) Pump beam position: The pump beam has a “center” position, i.e. the point in the middle of the FOV where maximum transmission is (as determined by the FOV calibration). Thus, the next step for beam alignment is to turn on a pump beam (there may be multiple pump beams co-aligned) and set the galvanometer voltages so that the beam is at its center coordinates.

(3) Witness camera: While this step isn’t critical, centering and focusing the witness camera (above the microresonator in the current geometry, using a low N.A. objective) can be helpful. Setting the crosshairs in the ThorCam software (if you’re using a Thorlabs camera as witness), you can fairly easily align the camera to the beam. The beam may be out of focus. Typically, you’ll want to change the camera position along the optical axis, **not** the focusing objective position, to get the camera imaging plane matched to the focal plane of the focusing objective.

(4) Taper: Next, I suggest aligning the tapered optical fiber. Eventually you gain a pretty good intuition of where the taper needs to start to easily align it, but this comes with hard-fought experience. To do this alignment, set the taper height to be a bit above the focusing objective (knowing the working distance of your focusing objective can help you pick a starting point). It’s important to be aware of the spatial location of the taper relative to other items (ex. microbubble capillary, focusing objective, witness objective, etc.), as a wrong move can easily break the taper.¹

Next, sweep the taper sideways so that it must pass under the imaging objective. Watch the camera output on a computer monitor to see if you detect the passing through of the taper. Usually, you can also see the taper by eye when it passes through the white-light focus, and this can be a useful alignment aid. If you did not observe the taper, change the height (most likely towards the objective if you started at a conservative distance from it), and do the sideways sweep backwards. Eventually, you’ll find the taper, and then should adjust the **taper height** (not the above camera, as that is already set to

¹ Which will not lead to positive feelings.

the pump beam!) so that the taper is in-focus. There are three particularly useful tricks I've found for this process if it is proving difficult:

- Start at a thicker region of the taper. Once you find and focus the taper, you can then translate the taper to its center while being aligned with the witness camera center (likely needing to use all three axes to do this, unless the taper is perfectly glued down to its mount).
- Turn off the white-light source from above and turn on the pump beam from below (if you have turned it off already). Auto-adjust the contrast of the witness camera. Now, when the taper passes beneath the witness objective as you sweep it, it should be noticeable even when significantly more out-of-focus than when using the white light. This makes sense, as the laser source has a far longer coherence length than the white light source. You might think of this like the “memory” of the taper staying in the laser beam for much longer than in white light.
- If you pass probe light through the taper, any dust on the thin portions of the taper should “light up”, which can be a useful alignment aid. This trick will require a sufficiently thin region of the taper such that it is lossy.

Lastly, it's worth considering what portion of the taper you want to center on the pump beam. This can be adjusted later as needed, but getting in the ballpark will be helpful to start with. This process will take some intuition. Essentially, you want a region of the taper that is sufficiently thin for the microbubble you're using. Whether or not the bubble is air- or water-filled will also impact this choice, as thinner regions of taper more readily phase-match into the water-contained optical modes. In general, it will be a better choice to couple to the resonator closer to the APD-side of the taper (output side), as then you put as many optical losses from the taper *before* the resonator as possible (which can usually be compensated by turning probe power up), as opposed to having more losses after the resonator, where the photons have much more value, as they carry resonance information.

(5) The microbubble: Aligning the microbubble is the last step! I usually start by translating the bubble using the coarse translation stage it is mounted on (in-between there is also a 3D piezo stage (Attocube), though this is best used for the last fine alignments). Keeping it just above the taper and translating the bubble, it should be evident once the capillary crosses beneath the white light. The resonator mount can then be translated until the capillary lights up more, indicating that bubble is now underneath the white light.

Typically, I back-off the witness objective, and translate it to find the bubble where it is. Then, I iteratively lower the objective, and lower the bubble.¹ Eventually, the taper will come back into view on the camera, and then I switch over to using the Attocube stage to move the bubble until it is centered on the camera. At this point, it is worth double checking that the pump beam, witness camera, taper, and resonator are all aligned. If the alignment is proving difficult, employing the tricks listed above for taper-alignment (i.e. start at a thicker capillary region and/or use laser light instead of white light) may prove useful. Then, you can “stick” the bubble to the resonator, and if the probe beam is on, you should see the resonator “light up” (from light scattering out, presuming that a visible probe wavelength is being used).

3.12.2. Finding and coupling to resonances

It’s now time to find some resonances! To do this, I suggest first turning on the “sweep”, sometimes referred to as the “span”. This is a triangle wave that gets applied to the back of the Velocity laser controller (at the “wavelength modulation” input, which drives the piezo within the laser that fine-tunes the resonance). This modulation dithers the wavelength of the probe beam back and forth quickly, which allows for seeing any resonances within the dithered span. Centering either the bottom or top of the triangle wave on an oscilloscope will be very helpful when locking to resonances.

¹ Iterating this process is the best, as accidentally ramming the resonator and taper together can break either or both of these nanophotonic devices. If you do this, it will be sad, but don’t worry, as those before you have also done it, and it is somewhat a rite of passage.

For the microbubble set-up, typical values for the Sweep are 1 Vpp (for the current iteration, this involves actually setting the amplitude to 2 Vpp, as it gets teed off in a way that halves the voltage), 40 Hz, and 50% ramp symmetry (such that up-sweep and down-sweep are identical). These values contribute to the correct calculation of Error Slope in the code, and the Error Slope is then used to calculate resonance shift from the lock-in output voltage during experiments. The values for the Sweep may be adjusted as needed, especially the amplitude of the wave (which sets the sweep width in wavelength-space), for given experiments as needed (just make sure that the code-calculations know what the actual value is!).

3.12.3. Error slope measurement

For initially optimizing the system as described below, having a correct error slope isn't necessary, since the value monitored is simply being maximized. However, as soon as a quantitative resonance shift is desired, it is necessary to have this value accurately measured. For a given experimental set-up, you'll develop an intuition of what error slopes can be manageably locked to. Thus you may end up deciding to employ more or fewer attenuators between the PDH mixer and the lockbox, and between the lockbox and the probe laser controller, to hit the steepest error slope that is reasonably lockable.

The variables for the Sweep described above also need to be correctly entered into the front panel of the LabVIEW code for the measurement, and the correct laser-piezo calibration parameter in the block diagram. The front panel may also have multiple probe lasers to choose from, so the correct one much be chosen.

3.12.4. Having a robust and trustworthy resonance lock

Having a resonance lock that is both robust (i.e. resilient to perturbations) and trustworthy (i.e. locked to the error slope that you wanted to lock to) is critical for quantitative measurements. Below I outline a suggested procedure for this, which had intellectual contributions from other group members, especially Feng Pan.

In our experiments, we usually monitor at least three data channels, namely the “DC Tap” (the resonance channel), the “Error Signal” (the signal being sent from the lockbox to the lock-in), and the “Lock-In” channel (output voltage of the lock-in). A fourth channel often also included is the sweep signal being sent to the probe laser controller, as this channel is helpful when locking.

There are two key relations to know about, between the unlocked state of the microscope, and the locked state. **In the locked state:**

- (1) The DC Tap should be at the **same voltage level as the bottom of the target resonance was** (the “dip depth”).
- (2) The Error Slope should be **centered at 0 V**, presuming that you adjusted the DC offset to set the center of the slope at 0 V prior to locking.

Figure 3-31 shows traces of the DC Tap and Error Signal channels while applying a voltage sweep, with the resonance-of-interest noted by a gold star, and the corresponding positions of the two channels while locked correctly.




Figure 3-31. Robust lock assurance. *Left:* The DC Tap (top, green) and error signal (bottom, blue) while UNLOCKED, with the desired locking point shown by the yellow stars. The channel data is shown twice, due to the back-and-

forth nature of the probe beam sweep. *Right:* The two channels while LOCKED. Importantly, the DC tap level is at the same voltage as the bottom of the unlocked resonance, and the error signal is at 0 V.

3.12.5. Locking procedure

We now know some indications of a good resonance lock, but how do we actually go about locking to a resonance? Different members of our research group all have their own trips and tricks, and the necessary skills will depend largely on the Q-factor (and thus error slope steepness) and the stability of your resonances.

Firstly, a desirable resonance should be found by course-tuning the Velocity laser, using the laser controller, software GUI, or custom software. Iteratively, the second set of polarization paddles (after the phase modulator) may be adjusted to optimize resonances or reveal previously unseen resonances. This is often best done with the PDH electronics off.

Once a resonance is found, the **PDH electronics should then be turned on**. Now, adjusting the first paddle set (before the phase modulator) may be necessary to optimize the error slope (which may then require adjustment of the second paddle set as well). Another likely adjustment for optimizing the error slope is tuning the frequency of the VCO which drives the phase modulator. I suggest Dr. Erik Horak's thesis for further discussion on how this tuning mathematically impacts the error slope.⁷ These paddle and VCO adjustments aren't always necessary, but will likely be necessary at least after installing a new taper or resonator!

Once the error slope is optimized, the offset knob on the lockbox can be used to fine-tune the error slope that will be locked to such that it overlaps with the "zero point" of the sweep (halfway between the bottom point and top point of the triangle wave). Next, the DC offset knob of the lockbox should be adjusted to move the error slope up or down on the oscilloscope such that it centers around 0 V.

Now comes the sometimes-tricky part, which is getting the lockbox actually locked onto the correct error slope. To do this with the Vescent lockbox, we need to turn *off* the sweep (delivered from a function generator), then flip the relevant lockbox switch to “lock”. Sometimes the lock will be correct on the first try. More often than not, the lock will not be correct, and the lockbox will be locked to some other zero-point on a slope nearby. There are a few ways to go about remedying this situation:

- **Adjusting the DC Offset knob** (the one used to raise or lower the error slope). Thinking back to when the Sweep was on, what will eventually happen is that the top or bottom of the slope locked to will be passed over, and the lockbox will then grab onto the next slope available. Eventually, you may find the correct error slope this way.
- **Changing the gains in the electronics** (in truth, mostly attenuations, which are just negative gains). This is quite important, as if the gain is *too* attenuated, then the lockbox will not be able to keep hold of the error slope, but if the gain isn’t attenuated enough, the feedback will be too strong, and the lock will immediately break (or have enough noise in it that the measurement quality suffers). This gain adjustment is usually done by adjusting the gain knobs (course and fine) on the lockbox, but sometimes, adding or subtracting SMA attenuators between the mixer and the lockbox, or between the lockbox and the probe laser controller may be necessary. See below for more discussion on the SMA attenuators.
- **Changing polarity of the error slope.** While this could be done by adjusting the VCO frequency, the Vescent lockbox provides a convenient switch to flip the direction of the signal (thus a negative error slope becomes a positive one, and vice versa). Sometimes a lock may be more easily achievable with one polarity as compared to the other.
- Unlock, turn the sweep back on, and **decrease the sweep voltage**. Iteratively adjusting the lockbox and oscilloscope parameters while decreasing this sweep will allow for more confident alignment of the error slope zero-point to that of the sweep voltage. If you retake an error slope

measurement, just make sure to account for the changed sweep voltage in the code, so that the calculated slope is correct!

Locked oscillations: Something you may observe once locked are oscillations in the DC Tap and Error Signal channels of the oscilloscope. Dr. Kevin Heylman in his thesis stated that these oscillations come from a mechanical resonance of the laser cavity at 2.1 kHz,¹ and what I've seen (albeit with a different Velocity laser and controller) are oscillations at around that frequency. Newport suggests using these oscillations as a diagnostic that you are in fact locked to a resonance, though in their example it is to the side of an atomic resonance, rather than the zero-crossing of an error slope. They suggest increasing gain (decreasing attenuation) until these oscillations show up, and then backing off slightly on the gain.

However, once the current iteration of the microbubble microscope was built, I have found that I can't have a stable lock *without* these oscillations (**Figure 3-32B,C**)! If I "back off" on the gain to where the oscillations disappear, the lock becomes unstable most of the time, such that minor perturbations (such as a loud voice in the room, or a large object on the resonator getting photothermally mapped) will break the lock, or at least contribute to significant background signal.

I have tested parameter sweeps measuring the same object photothermally at different gain settings (data not shown), which correspond to different amplitudes of oscillations. I have found that the unstable regime with no oscillations leads to worse SBR in the measurement. Alternatively, if the gain is so high that the oscillations are in a chaotic pattern, the SBR also suffers. Thus there is a sweet-spot where stable oscillations of moderate amplitude are ideal given the current set-up. However, it is unclear

¹I have yet to verify this from an independent source, but I trust that Kevin properly vetted this claim

if these oscillations are adding some background compared to a stable, non-oscillating locked regime, which we have achieved before in the group with certain photothermal experiments.¹

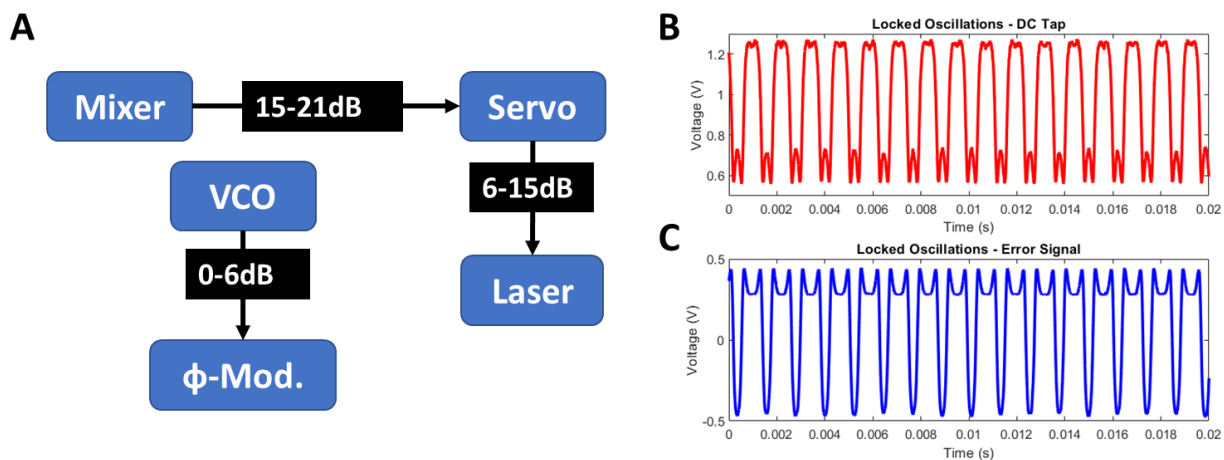


Figure 3-32. Locked-resonance oscillations. (A) A diagram showing the range of attenuators put into different sections of the electronic paths. Not shown are the additional attenuation options provided by the Servo itself. (B) Oscillations in the DC Tap signal. (C) Oscillations in the Error Signal.

While attaining stable, oscillation-free locking may be a worthwhile future pursuit, the oscillations may have some utility of their own. For example, I've observed that the amplitude of the oscillations changes during experiments, which I attribute to the error-slope changing while locked to it. While these changes would ideally be avoided altogether, it is perhaps possible that by simultaneously monitoring the oscillation amplitude, error slope drift could be corrected for to obtain a more trustworthy measurement of resonance shift as an experiment progresses. Another use for oscillations, as Newport suggest, is to make sure that you can produce them by increasing the gain while locked. If oscillations can't be achieved, the system likely isn't locked to the desired resonance, and probably won't be useful for photothermal measurements.

¹ Additionally, while I have no evidence to indicate that the oscillations are damaging to the laser cavity, I don't know for a fact that this isn't bad for the laser's health. If the oscillations are bad enough (or sometimes in other circumstances), the Velocity controller will produce a loud whine when the lock is on. In this case, I immediately turn the lock off.

A note on SMA attenuators: We have found it is typically necessary to put attenuators at various points in the PDH electronics (**Figure 3-32A**). In some spots, especially between the mixer and the lockbox, this attenuation reduces the error slope! We have found this attenuation to be necessary for locking many resonances. What this means is that for a particularly high Q-factor resonances, the error slope is **forced to become shallower**. Though we haven't rigorously modeled this, I believe this effectively lowers the sensitivity of the measurement (see *Section 2.7*), as there is less voltage shift in the error slope per resonance shift, thus making it as if a lower-Q resonance is locked to.

Finding ways to reliably lock to higher error slopes should improve the limit of detection of our experiment, though this improvement may not be needed for many experiments. It is also a bit unclear how detrimental this added attenuation is, as the lockbox itself employs attenuations in processing the feedback for the laser controller. Ultimately, diagnosing the minimal voltage change detectable by the lock-in amplifier should give the answer to at what point reduction in error slope will start to negatively impact photothermal experiments.

3.12.6. Initial mapping

Once locked to a resonance (or honestly anything you can lock to that's photothermally responsive to begin with), it is time for some initial photothermal maps. Usually, doing a very large map that is expected to cover most or all of the resonator is a good first step. The setting will likely be rather coarse (something like 1 μm steps), and high pump beam powers should typically be employed. In fact, if you have multiple pump lasers co-aligned, having all of them running at once can help expedite finding an object to optimize to (though then the responsible laser(s) then need to be ascertained for optimization). Once an object is found, preferably near where on the resonator you want to perform more studies (typically near the center of the pump-incident resonator face), you can move on to the step of optimization.

3.12.7. Optimizing signal

Once one or more objects have been found on the microresonator during initial mapping, we can optimize to an object. The “Optimization Tab” is designed just for this purpose. The center-coordinates, per photothermal mapping, can be input as the starting galvo coordinates (in voltage measurements). Then, the X and Y coordinates and the piezo focus can be iteratively adjusted to maximize the lock-in output reading (which is shown within the LabVIEW tab), along with pump power and other settings.

3.12.8. Taking “real” data

Now that everything is optimized and calibrated, it is time to take “real” photothermal data. What this means may vary, but almost always it will involve first finding objects of interest. Thus, doing more of the big maps that span most or all of the bubble, perhaps at multiple polarizations and/or wavelengths to ensure that all objects are found, is a good first step. For the big maps, pump power may need to be lowered to allow the largest objects to be mapped without losing lock.

Chapter 4 (on gold nanorod studies) and Chapter 5 (on photothermal background) of this thesis are good resources for considering the variety of experiments that might be pursued. Once objects are found, they may be mapped, swept using a variety of parameters including pump polarization and pump wavelength.

3.13. Understanding Photothermal Data

3.13.1. Calculating resonance shift

The only thing actively measured during our experiments is the output voltage of the lock-in amplifier.¹ Quite a simple readout for a fairly complicated experiment! Now it is our task to convert that voltage into a resonance shift. First, that *output DC* voltage needs to be converted to the *input AC* voltage that the lock-in received (the component at the amplitude modulated frequency). This is done using the “Sens.” parameter. Ex. a 4 V output measured at the DAQ, and a set sensitivity of 20 mV, the input AC

¹ Though in the future, simultaneous measurement of parameters such as pump beam power and probe beam polarization state may be useful to incorporate into the instrument’s hardware and software.

voltage is calculated to be 8 mV. This 8 mV itself comes from the error signal, as the error output goes up and down, as the resonance oscillates back and forth due to the amplitude modulation. A factor of 2.83 is needed to convert from the V_{rms} (rms = root mean squared) measured to the V_{pp} (pp = peak-to-peak) of the oscillation. The next step is to use the measured error slope and the known probe laser piezo calibration parameter to convert the lock-in input to a resonance shift! Considering all of the above, and starting from the measured lock-in output, resonance shift is calculated using **Equation 3-5**, with units for each parameter italicized in parentheses. Note that the conversion from an error slope including picometers to a resonance shift of attometers comes from a factor of 10^6 that is somewhat hidden in the $\mu V/V$ unit division in the numerator (for lock-in sensitivity setting over output).

Equation 3-5

$$\text{Res. Shift (am)} = V_{DAQ}(V) \frac{2.83 * \frac{\text{Sens. (}\mu V)}{10 (V)}}{\text{Error Slope (V/pm)}}$$

The error slope in the above equation is already in our preferred units of $V/\mu m$. During the error slope measurement step (or during the resonance shift calculation in older version of the code, using the V/s error slope), **Equation 3-6** is used to get that parameter, with units for each parameter italicized in parentheses.

Equation 3-6

$$\text{Error Slope (V/pm)} = \frac{\text{Error Slope (V/s)}}{V_{pp} (V) * 2f_{\text{sweep}}(Hz) * \text{Piezo Cal. (pm/V)}}$$

3.13.2. Polarization results

Chapter 4 is a good resource for understanding how to interpret polarization data in these experiments, as are other photothermal microresonator papers from the Goldsmith group, so I will not elaborate further here.

3.13.3. Post-acquisition application of correction factors

Sometimes it may be realized after taking valuable data that there was an incorrect value input into the code. For example, having an incorrectly recorded probe piezo sweep amplitude would lead to an incorrectly calculated error slope, propagating to an incorrectly calculated resonance shift. Another parameter that might lead to needed corrections of data is an incorrect calibration value for the probe piezo. Luckily, errors such as this may be fixable as long as one knows what the true parameter was.

Here, I've put some an example of how to proportionally correct certain variables based on a change in another variable, with the example being how error slope, quality factor, and resonance shift are impacted if a correction is needed for the probe cavity piezo calibration parameter. This is a real-world example that I had to go through when calibrating the 635nm Velocity laser piezo and correcting for data assuming the calibration parameter for the 780nm Velocity.

If the probe piezo calibration parameter is 20 pm/V instead of 60 pm/V, and these parameters were initially calculated with the latter:

- **Q-factors** are actually 3x *higher* than initially calculated because the linewidth is actually 3x narrower than calculated.
- **Error slopes** in the "new" units of pm/V are 3x *higher* than initially calculated. Error slope in the "old" units of V/s are as initially calculated (as the probe piezo parameter in the old calculations did not get used until the resonance shift calculation).
- **Resonance shifts** are 3x *lower* than initially counted because a given error voltage is actually 1/3 the calculated correction.

3.13.4. COMSOL modeling to obtain absorption cross-section

All of the work and calibration that we've put in so far results in a measured resonance shift for a given pump beam perturbation. However, "resonance shift" is not very informative physically or chemically for a given analyte, and thus further conversion of this parameter is needed. Ultimately the parameter we'll

want to extract is an absorption cross-section, σ_{abs} , of our analyte. This is somewhat analogous to the bulk parameter ϵ , the extinction coefficient, however ϵ includes both absorption and scattering terms, the latter of which may or may not be significant depending on the system under study. To calculate σ_{abs} , some modeling will be necessary. We use COMSOL simulations, both thermal and optical, to calculate this conversion. See the supporting information in *Chapter 4* of this thesis for details on COMSOL modeling.

3.14. Photographs of the Photothermal Setup

Below are several pictures of the current version of the second-generation microresonator photothermal instrument, with further descriptions in the captions.




Figure 3-33. Pictures of microscope setup. (A) A view of the entire table, with the M Squared system on the left. From front to back on the right are two solid-state diodes in the foreground (red and green rectangles), the Velocity probe laser (red octagon), the pump-beam combining optics, the optical chopper, and in the very back the galvanometer mirror which sends the pump beam to the left to the “microscope box”. (B) The interior of the “microscope box”, with the attocube on the left holding a microbubble mount, the tapered fiber held by the stage on the right, the witness objective above, the focusing objective below, and beneath that the LCPR optics above the 45-degree mirror. A trash bag can be seen on the back side of the box, for protection against air currents. (C) An alternative view showing the LCPR and a focusing objective. The power meter shown on the top of the image is there for calibration purposes temporarily. (D) A zoom-in, where from top to bottom are the focusing objective,

the tapered fiber, the microbubble resonator (with pump light focused on it, as seen by the red dot), and the witness objective.




Figure 3-34. Pictures of microscope setup. (A) The “paddle tree”, holding both sets of polarization paddles. (B) Two laser beam paths are co-aligned using three HWPs and a polarizing beam splitter (here, a combiner). (C) Three co-aligned beams (532 nm, 561 nm, 635 nm) are shown on a notecard, indicating good co-alignment. The green spot on the lens holder is photographic artefact. (D) The red pump beam bounces off the galvo mirrors and passes through the lens relay on its way to the 45-degree mirror.

4. Toward Real-Time Monitoring and Control of Single Nanoparticle Properties with a Microbubble Resonator Spectrometer^{III}

4.1. Abstract

Optical microresonators have widespread application at the frontiers of nanophotonic technology, driven by their ability to confine light to the nanoscale and enhance light-matter interactions. Microresonators form the heart of a recently developed method for single-particle photothermal absorption spectroscopy, whereby the microresonators act as microscale thermometers to detect the heat dissipated by optically pumped, non-luminescent nanoscopic targets. However, translation of this technology to chemically dynamic systems requires a platform that is mechanically stable, solution compatible, and visibly transparent. We report microbubble absorption spectrometers as a versatile platform that meets these requirements. Microbubbles integrate a two-port microfluidic device within a Whispering Gallery Mode (WGM) microresonator, allowing for the facile exchange of chemical reagents within the resonator's interior while maintaining a solution-free environment on its exterior. We first leverage these qualities to investigate the photo-activated etching of single gold nanorods by ferric chloride, providing a method for rapid acquisition of spatial and morphological information about

^I Authors: Levi T. Hogan^{+,‡}, Erik H. Horak^{+,‡}, Jonathan M. Wards[§], Cassandra A. Knappert⁺, Síle Nic Chormaic^s, Randall H. Goldsmith^{+,*}

[‡] These authors contributed equally.

[†] Department of Chemistry, University of Wisconsin-Madison

[§] Light-Matter Interactions Unit, Okinawa Institute of Science and Technology Graduate University

The manuscript was written through contributions of all authors. All authors have given approval to the final version of the manuscript. The published version of this manuscript can be found at reference 20.

^{II} This chapter is presented almost entirely as originally published, with the exception that I have streamlined the formatting of section headings, figures, etc. to flow with the rest of the thesis. I also cut the last section of the supporting information, which was just a few pages of MATLAB code.

nanoparticles as they undergo chemical reactions. We then demonstrate the ability to control nanorod orientation within a microbubble through optically exerted torque, a promising route toward the construction of hybrid photonic-plasmonic systems. Critically, the reported platform advances microresonator spectrometer technology by permitting room-temperature, aqueous experimental conditions, which may be used for time-resolved single-particle experiments on non-emissive, nanoscale analytes engaged in catalytically and biologically relevant chemical dynamics.

4.2. Main Text: Introduction

Optical microresonators, devices that confine light to microscopic volumes, have found widespread application within chemistry, biology, physics, and engineering.^{8, 39-42} A broad class of optical microresonators, whispering-gallery mode (WGM) resonators, have exhibited superb sensitivity including the detection of single nanoparticles,^{28, 43} single molecules^{32, 44-46} and even single metal ions.²⁶ However, the ability to perform spectroscopy on adsorbed objects would not only allow for label-free chemical identification, but also allow the interrogation of single object properties, free from the static and dynamic blurring of typical ensemble measurements. To this end, we recently employed microtoroid resonators as single-particle absorption spectrometers, whereby the heat dissipated by optically pumped nano-objects such as gold nanorods (AuNRs),^{10, 47-49} carbon nanotubes,⁵⁰ or conductive polymers¹⁹ is detected *via* small shifts in the WGM resonance condition. However, to harness the sensitivity of this method for chemically dynamic systems, a platform easily compatible with solution-phase measurements is necessary. Here, we report such a platform, the microbubble resonator, and use it to study the photo-activated chemical etching and reorientation of single AuNRs.

AuNRs⁵¹ have important chemical and biological⁵² applications such as bioimaging,⁵³ treatment of cancer⁵⁴⁻⁵⁵ and infection,⁵⁶ label-free biosensing⁵⁷ down to single molecules,⁵⁸ surface-enhanced Raman spectroscopy,⁵⁹⁻⁶⁰ fluorescence enhancement,⁶¹ drug delivery,⁶² and light harvesting to drive catalytic reactions.⁶³ These applications heavily rely on tuning the morphology-dependent optical

features of AuNRs, necessitating precise tailoring of their dimensions. This result can be achieved during AuNR fabrication, where seed-mediated synthesis⁶⁴ can often tame the polydispersity that typically plagues samples. However, in many cases polydisperse AuNR samples are still common and post-synthetic modifications offer an attractive route to achieve a desired morphology. Furthermore, significant particle-to-particle variations of key optical properties of AuNRs both on surfaces⁶⁵ and in solution⁶⁶ highlight the heterogeneity within a population of nanoparticles, underscoring the need for single-particle inspection, including during nanoparticle synthesis and modification. A variety of optical methods exist for probing non-luminescent single nanoparticles and molecules *via* photothermal,⁶⁷⁻⁷⁰ scattering,⁷¹⁻⁷³ and other techniques.⁷⁴⁻⁷⁷ Observation of the chemical etching of single AuNRs has recently been accomplished with one-photon luminescence,⁷⁸⁻⁷⁹ dark-field scattering,⁸⁰⁻⁸⁶ and liquid transmission electron microscopy (TEM).⁸⁷ However, a highly sensitive absorption technique for monitoring such chemical dynamics is needed to compliment these methods, and would be extremely valuable for accessing targets that are not luminescent or are too small for scattering experiments. WGM resonators are perfectly poised to fill this gap in methodology.

Various WGM microresonator geometries have been employed for sensing in solution, including microspheres,⁴⁴⁻⁴⁶ microrings,^{42, 88} microtoroids,^{28, 89} microbubbles,⁹⁰⁻⁹¹ microdroplets,⁹² microtubes,⁹³ and microbottles.⁹⁴⁻⁹⁵ In particular, the variations and capabilities of hollow microresonators for sensing have been reviewed in detail elsewhere.⁹⁶ To adapt a microresonator for in-solution, visible-wavelength photothermal spectroscopy, three requirements must be met: (i) high sensitivity for interrogating nanoscopic analytes, (ii) resonator transparency at visible wavelengths to mitigate photothermal background, and (iii) robust performance in solution. Employing silica-on-silicon (SiO₂-Si) microtoroids for photothermal spectroscopy, one can resolve attometer shifts of the WGM resonant wavelength from thermal fluxes of target nano-objects.¹⁰ High backgrounds in SiO₂-Si toroids can be mitigated with all-glass microtoroids, which can be used for visible spectroscopy.⁴⁸⁻⁴⁹ However, immersing a WGM

microresonator in water mandates the use of larger microresonators to avoid bending losses,⁹⁷ with consequent lower photothermal sensitivity. Furthermore, although tapers⁹⁸ and prisms⁴⁴ can be optically coupled to WGM microresonators in water, immersion of such couplers in solution may reduce mechanical stability and also result in fouling, particularly as more caustic reagents are employed for chemical studies. Therefore, an alternative platform is preferable for in-solution experiments. The microbubble WGM resonator, **Figure 4-1B**, which possesses a hollow, solution-accessible interior, while maintaining an air-glass exterior interface, meets the requirements for in-solution, visible spectroscopy of nanoscopic analytes.

Microbubbles are fabricated from glass capillaries, resulting in low background signals at visible wavelengths, tunable fabrication, and two-port connectivity through which it is easy to flow reagents. Compared to a solid resonator immersed in solution, a microbubble maintains an air-glass interface on its exterior, enabling a higher refractive index contrast and allowing for smaller diameter resonators before bending losses occur. Additionally, the tapered optical fiber used for coupling light into the resonator can approach in air, reducing noise from the instability of coupling in-solution and eliminating solution contamination of the taper. Furthermore, the unique, thin-walled structure of the microbubble allows for high-order optical modes that exist almost entirely within the liquid-core of the resonator, a situation termed the “quasi-droplet regime”. Operating in the quasi-droplet regime, microbubbles have proven exceptional sensors, most recently for detecting polystyrene nanoparticles in aqueous solution with a sensitivity ~280 times larger than similar experiments using microsphere resonators.⁹⁰ Together, these factors make microbubble resonators ideal for time-resolved spectroscopy of single-particle chemical reaction dynamics when exposed to solution. In this paper, we introduce microbubble absorption spectrometers for probing and controlling the chemical and rotational dynamics associated with the photo-activated chemical etching of AuNRs. This platform holds promise for elucidating mechanistic insights into nanoparticle reactions with a method orthogonal to existing techniques that

rely on scattering or luminescence, and is an attractive candidate for future single-particle and single-molecule studies.

4.3. Main Text: Experimental Design

WGM resonators operate *via* total internal reflection, wherein light propagates around a closed geometric loop, resulting in resonance conditions where only specific wavelengths propagate constructively. WGM resonances are interrogated by the “probe beam”, provided by a continuous wave (CW), narrow-linewidth, tunable laser coupled through a tapered optical fiber⁹⁹ (**Figure 4-1A**). Transmitted light through the tapered fiber is collected and the probe beam is actively locked to a resonance by a Pound-Drever-Hall locking system.^{10, 28-31} The hollow core of the microbubble is filled with the desired reagents by attaching the microbubble capillary to a syringe pump. Two microbubbles are pictured in **Figure 4-1B**, highlighting the tunability of geometric parameters, and consequent versatility on optofluidic properties.

A second beam, the “pump beam”, is focused onto the microbubble surface to excite analytes. CW diode lasers at 532 nm, 635 nm, and 785 nm are coaligned using dichroic mirrors, permitting interrogation at different wavelengths. The linearly polarized pump beam is amplitude modulated at 433 Hz using an optical chopper, encoding the photothermal signal at this frequency and allowing for use of lock-in amplification to drastically lower the experimental noise floor.¹⁰ Two galvanometer mirrors steer the pump beam through a relay lens system to a 40x objective with piezo-controlled focus, providing spatial control of the pump beam on the microbubble resonator. This spatial control is leveraged to photothermally map the interior surface of microbubble resonators, at low resolution for an entire resonator, and high resolution for single diffraction-limited objects. **Figure 4-1C** shows two photothermal maps of a microbubble resonator at different objective foci, with the out-of-focus map indicating the curvature of the microbubble from the varied PSFs across the map. The polarization angle of the linearly polarized pump beam is rapidly scanned using a voltage-controlled liquid crystal, which is sandwiched

between a polarizer and a quarter-waveplate (see **Methods**). This combination of wavelength, spatial, and polarization control permits thorough characterization of individual analytes bound to the resonator, realized at exquisite sensitivity due to the double-modulation scheme.

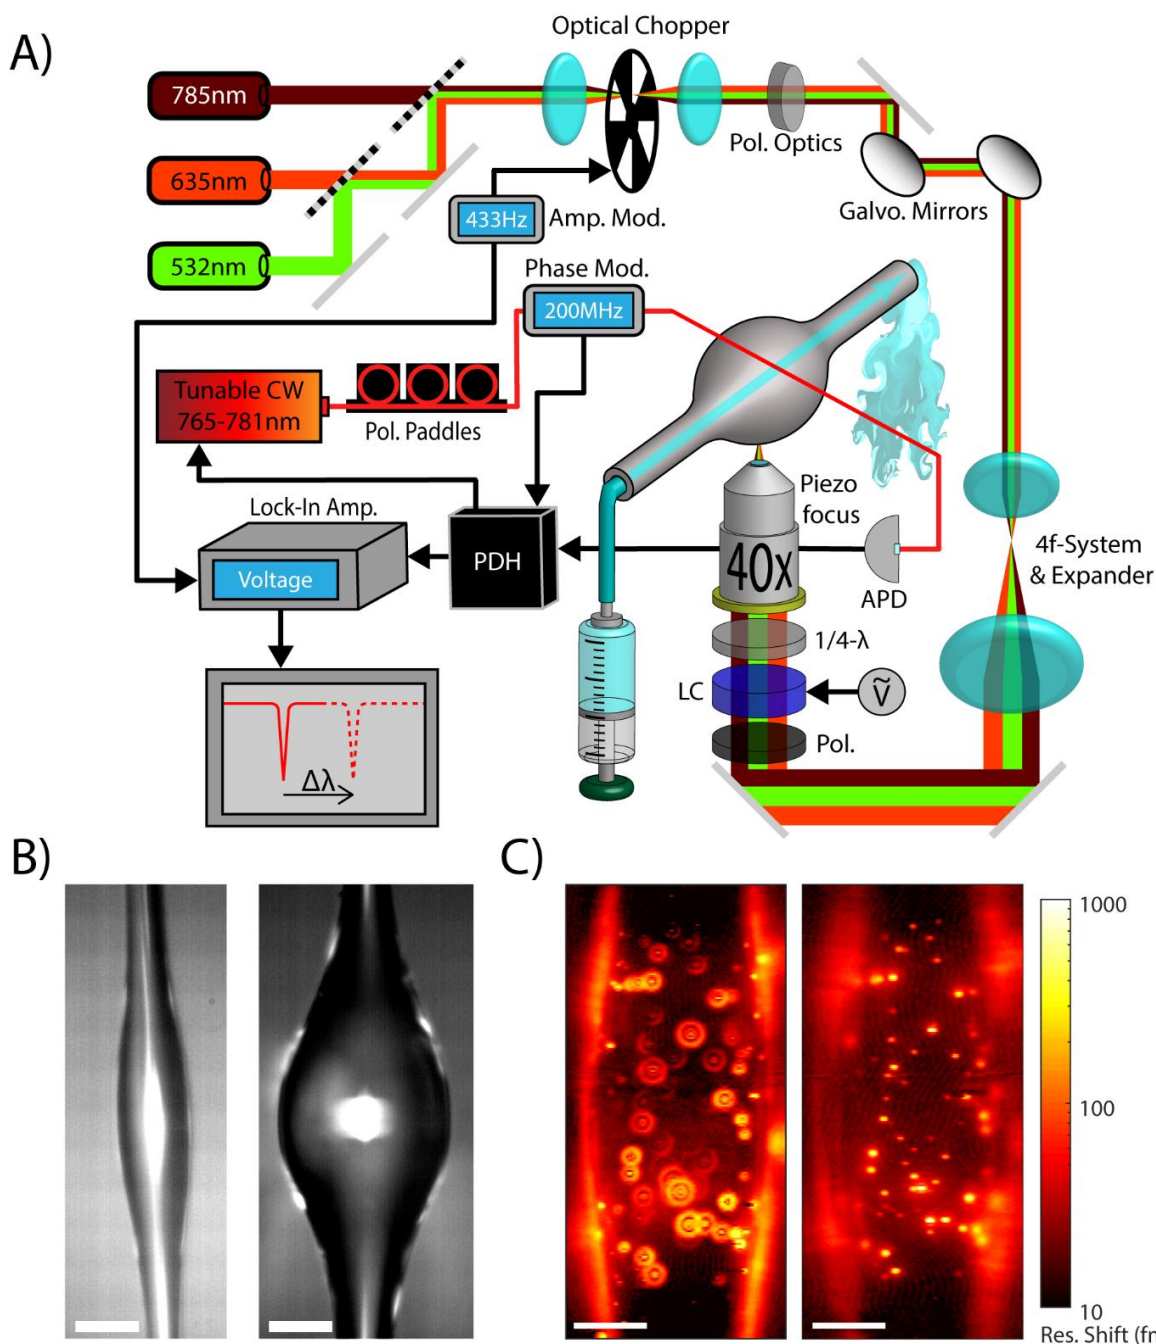


Figure 4-1. Microbubble Absorption Spectroscopy. (A) Cartoon of instrumentation. PDH=Pound Drever Hall. LC=Liquid Crystal. APD=Avalanche Photodiode. (B) Optical micrographs of two microbubble resonators with different geometries. Scale bars 20 μm. (C) Photothermal maps of a microbubble resonator similar in geometry to the left microbubble in Figure 1B, both out-of-focus (left), and in-focus (right). Scale bars 20 μm.

4.4. Main Text: Operation of Single-Particle Microresonator Spectrometers

Microresonator sensing schemes generally rely on the reactive mechanism,⁴⁵ whereby binding of an analyte imparts a small refractive index change, shifting the resonance wavelength. Instead, microresonator-based photothermal spectroscopy relies on a resonance shift resulting from the heat plume generated by optically pumping a non-emissive object bound to the resonator surface.⁵⁰ The temperature rise accompanying this heat plume alters the resonator's refractive index according to its thermo-optic coefficient (dn/dT), changing the WGM optical path length and shifting the resonance condition. The ability to detect this resonance shift is related to the figure of merit Q/V , the ratio of the resonator's quality factor (Q) and mode volume (V). A resonator with minimized absorption, bending, and scattering losses allows photons to repeatedly circulate the resonator, resulting in a high Q , narrow linewidth resonance. This narrow linewidth increases the visibility of minute resonance shifts. A smaller resonator with consequent tighter confinement of light produces a smaller V , increasing the overlap between the thermal plume of the analyte and the optical mode. This increased overlap contributes a larger effective refractive index change and thus a larger resonance shift.⁵⁰ To properly examine the microbubble photothermal response we employed finite-element simulations (COMSOL) of both the optical modes and the thermal properties of the microresonator. Simulated optical modes for a particular microbubble geometry are shown in **Figure 4-2A**. Varying mode numbers, defined in the traditional spherical geometric indices (polar, azimuthal, and radial, **Supporting Information**), clearly show the complicated mode structure inherent in the microbubble resonators. This complex mode structure gives rise to several important experimental considerations.

First, the efficient excitation of high-order modes leads to incredibly congested mode spectra. An illustrative 180 pm window of a water-filled microbubble's resonance landscape is shown in **Figure 4-2B**. This high mode density stems from the highly prolate resonator geometry lifting the polar mode

degeneracy relative to an ideal spherical resonator,¹⁰⁰ leading to varying effective resonator sizes for modes, as well as differing free-spectral ranges.¹⁰¹ The differing free-spectral ranges cause spectral overlap of modes of different azimuthal mode order,¹⁰² an effect that is compounded by the disparate dielectric environments experienced by different-order radial modes, which have different fractions of the electric field contained in glass, water, and air. Second, the burrowing of higher-order radial modes into the water-filled interior not only changes the effective refractive index of the mode but also yields tremendous variations in dn/dT . This varied dn/dT , which can even switch signs, produces very different thermal responsivities for modes. The combined congested mode spectrum and differential shifting from dn/dT variation amplifies experimental challenges, as photothermal heating or ambient temperature drifting can cause modes to shift through each other. Therefore, modes that are both thermally responsive and spectrally isolable are desirable.

Choosing a high thermal responsivity mode requires delving into the expected thermal response with finite element simulations. As described above, the radial mode order drastically alters the effective dn/dT , as glass has a small positive dn/dT of $9 \times 10^{-6} \text{ K}^{-1}$ and water a large negative dn/dT of $-91 \times 10^{-6} \text{ K}^{-1}$. Thus, while glass-contained modes in a water-filled resonator offer Q values over 10^6 and show small positive resonance shifts upon heating, higher-order, water-contained modes offer Q values of mid- 10^5 and show large negative resonance shifts captured both experimentally (**Supporting Information**) and in our simulations in **Figure 4-2C**. The “shark fin” shape results from the pump beam amplitude modulation. Interestingly, this modulation rides atop a rising baseline magnitude as heat builds over many modulation cycles (left panel) before thermal equilibrium is reached (right panel). This baseline stems from the lack of an effective, proximal heat sink in microbubbles with equilibration reached only after sufficient heat dissipation to the air, yielding a baseline shift about ten times larger than the modulating shift, both experimentally and theoretically (**Supporting Information**).

Finding high-order water-contained radial modes to leverage their larger thermal response requires careful consideration of the coupling geometry. Specifically, the tapered fiber diameter¹⁰¹ greatly impacts mode selectivity through phase matching conditions and evanescent field overlap. By translating along the length of the tapered fiber this diameter was tuned until these water-contained modes were suitably excited. Then a thermally sensitive and spectrally isolated mode was selected by wavelength scanning. Importantly, the precise identity of this mode was not discerned, precluding direct relation of a resonance shift with an absolute absorption cross-section as in our previous experiments.¹⁰ ⁵⁰ Identification is possible,¹⁰³⁻¹⁰⁵ particularly when implementing procedures to simplify the mode structure,^{102, 106} but difficult in practice and was not pursued here.

Importantly, this lack of mode identification precludes the selection of the maximally thermally responsive mode. While the optimal mode is required for ideal sensor response, use of a less responsive resonance was sufficient for examination of AuNRs. The limit-of-detection of our system was investigated by optically pumping a single AuNR inside a resonator. The photothermal signal, averaged for 30 s with a time constant of 1 s (**Figure 4-2D**), was monitored at decreasing powers until it was indistinguishable from the signal obtained with the pump beam blocked. As the inset in **Figure 4-2D** shows, the signal remains linear over multiple orders of magnitude, flattening out at low powers as the noise floor is reached. The detection limit for this platform is in the low tens of attometers of wavelength shift (a comparison with microtoroids is made in the **Supporting Information**). For context, the typical photothermal response of a single AuNR at our pump fluxes is in the range of 10-100 fm, easily resolvable by many orders of magnitude. Additionally, this detection limit surpasses the expected femtometer photothermal shift for measuring a single chromophore.⁵⁰ As a first step towards monitoring reaction dynamics of molecules, we show below that microbubbles are well-suited for probing the chemical and spatial dynamics of single AuNRs.

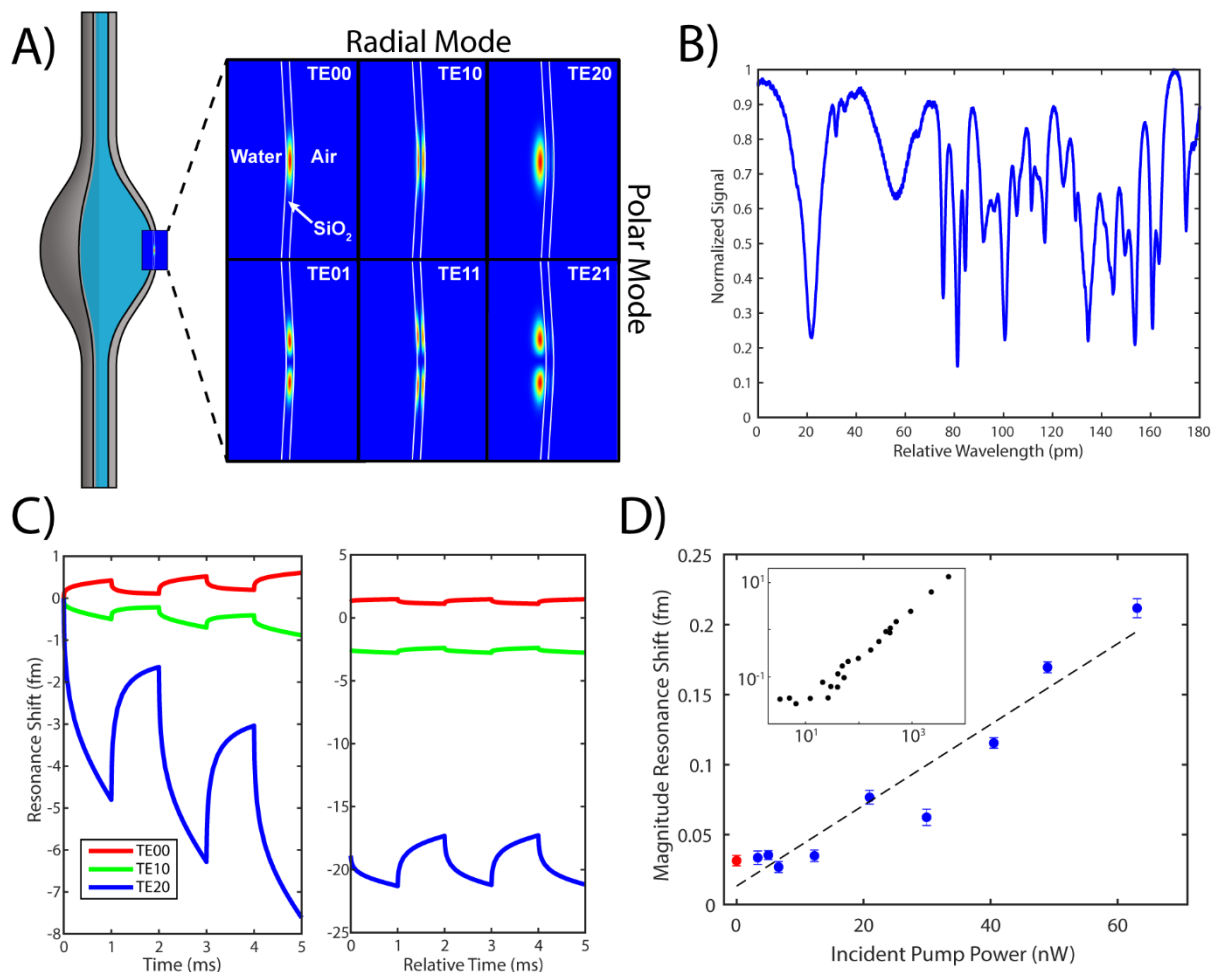


Figure 4-2. Optical resonances in microbubble resonators. (A) Simulated electric field distributions at 780 nm for first, second, and third order radial modes, for both first and second order polar modes. All modes shown are Transverse Electric (TE). White curves are added to clearly indicate the position of the microbubble walls. (B) A 180 pm span of the mode spectrum of a microbubble resonator. (C) Left: The signal at the beginning of analyte pumping. Right: Signal once the resonator has reached a thermal equilibrium with its surroundings (theoretical). (D) Resonance shift from pumping a single gold nanorod with the 635 nm beam at decreasing powers (blue points). The red point indicates the signal for pump beam off. The inset is a zoom-out, showing signal linearity over orders of magnitude in pumping power. Further details in main text. Error bars are standard deviation of the mean.

4.5. Main Text: Probing Photophysical Features of Single AuNRs

AuNRs exhibit optical features known as localized surface plasmon resonances (LSPRs), which result from light exciting collective oscillations of conduction band electrons. Two orthogonal LSPRs exist in AuNRs: the longitudinal plasmon band (LPB) and the transverse plasmon band (TPB), oriented parallel and perpendicular, respectively, to the long axis of the rod (**Figure 4-3A**). The LPB is at the longer wavelength in the bulk extinction spectra of the AuNRs used in this report (80×40 nm), **Figure 4-3B**. These

spectral features are probed with the microbubble platform detailed in **Figure 4-1A** at the single AuNR level, at specific pump beam wavelengths (solid vertical lines **Figure 4-3B**). The LPB central wavelengths will likely be red-shifted compared to the bulk due to interaction with the glass surface.^{71, 107-108}

After depositing AuNRs inside of a microbubble resonator (**Methods**), the resonator is photothermally mapped to find objects. To confirm successful deposition of single AuNRs, photothermal maps are acquired with the pump beam, linearly polarized from 0 to 180 (**Figure 4-3C**) at three different pump wavelengths. These maps are fit to extract an intensity at each polarization (**Methods**), shown as data points in **Figure 4-3D**. These polarization dependent intensities are fit (dotted lines) to give a depth-of-modulation, M (**Methods**). The 635 nm and 785 nm traces, which probe the LPB, have a value of M close to unity for a single AuNR. A criterion of $M \geq 0.98$ was used for classifying an object as a single AuNR. Small well-ordered aggregates could also exhibit high M values, but these are unlikely due to the presence of CTAB during deposition. Alternatively, a much faster method of probing AuNR orientation is to rapidly rotate the pump beam's linear polarization while centered on an object, and fit the results to extract M . Although this method (**Figure 4-3E**) lacks background subtraction, it allows for hundreds of data points to be collected in a few seconds, resulting in quick determination of AuNR orientation and relative absorption cross-section at the pumping wavelength during reactions.

For both of the above polarization methods, the traces for the 635 nm and 785 nm pump beams align in peak angle (**Figure 4-3D** and **Figure 4-3E**) because both of these wavelengths excite the LPB. In contrast, the 532 nm trace has a peak angle orthogonal to the other two traces because this wavelength excites the TPB at a pump polarization orthogonal to the LPB excitation. In addition, the 532 nm trace does not go to zero, because at that wavelength the pump beam is not only pumping the polarization-dependent TPB, but also the interband transitions of gold, which are independent of pump beam polarization. Notably, the use of multiple wavelengths means that AuNRs can not only be localized, but

also studied spectroscopically. Herein, these capabilities are employed to study the etching of AuNRs in

real-time.

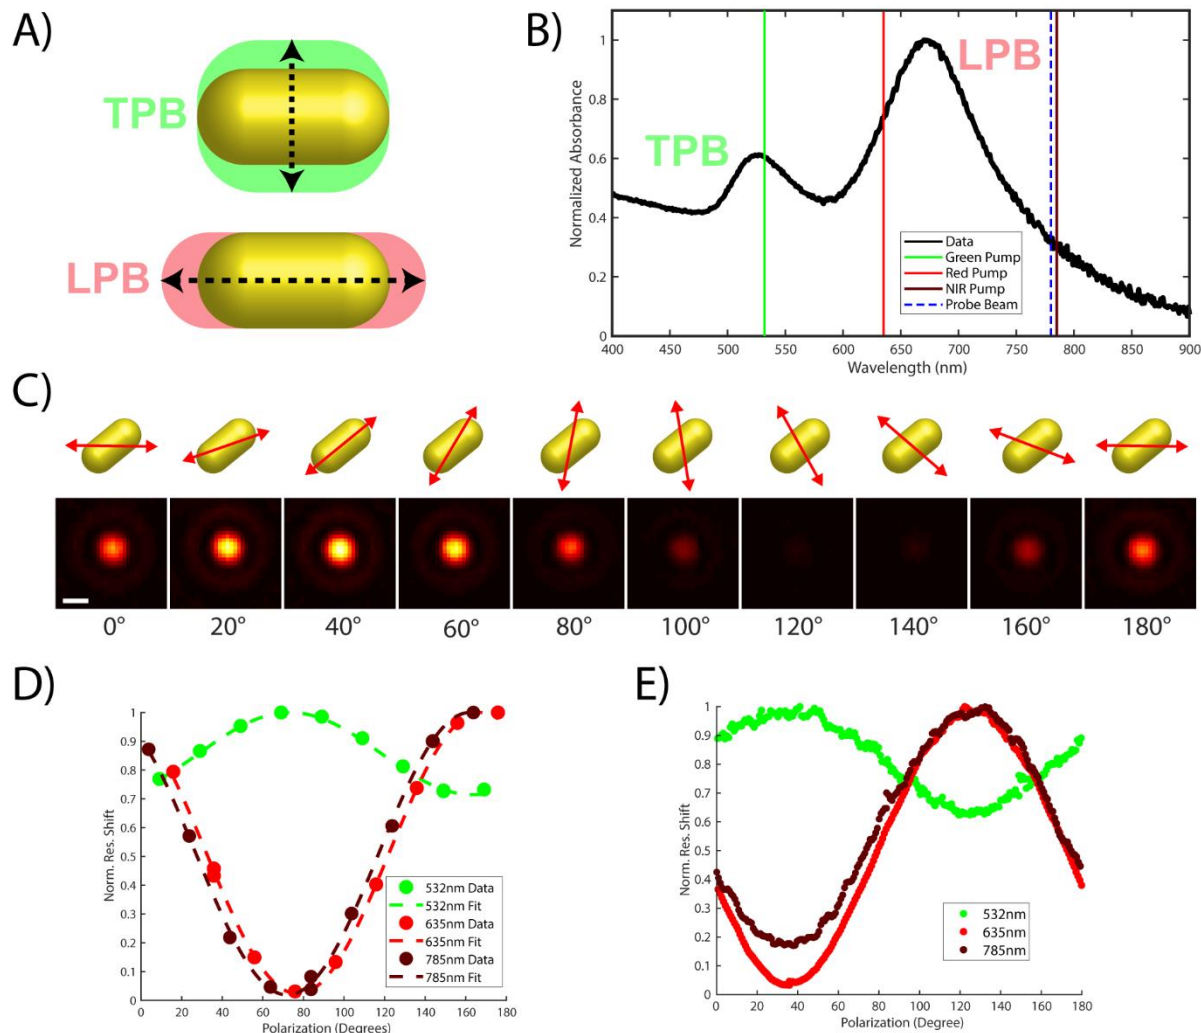


Figure 4-3. Probing photophysical features of single AuNRs. (A) Cartoon illustrating the photophysical features of a AuNR. LPB=Longitudinal Plasmon Band. TPB=Transverse Plasmon Band. (B) Bulk absorption spectrum of AuNRs, with the various laser beams in our experiment indicated by vertical lines. LPB and TPB indicated. (C) Example photothermal maps of a nanorod as pump polarization is varied in increments of 20 degrees, as shown by the red arrow in the cartoon above the photothermal maps. Scale bar 1 μm . (D) Polarization fits for three different pump beams acquired using photothermal mapping. (E) Polarization traces for three different pump beams, acquired by recording photothermal signal as the linear pump polarization is quickly rotated 180 degrees (~ 10 seconds).

4.6. Main Text: Selecting an etchant

The photophysical properties of AuNRs are well understood as a function of geometry¹⁰⁹ and post-synthetic modifications are extremely useful for exerting control over these properties. Since 2002, when Jana and coworkers observed anisotropic etching of gold spheroids in both cyanide and persulfate

solutions,¹¹⁰ at least 20 other reagents have been reported to etch or accelerate the etching of AuNRs, often with spatial selectivity (see **Supporting Information**). Such reports include assays for facile detection of analytes at ultra-low concentrations in both environmental and biological samples, indicating the utility of morphological control of AuNRs both in-the-field and at points-of-care.

Most of the aforementioned reports used spectrophotometers to study ensembles of AuNRs. Although some studies used TEM intermittently to verify nanorod morphology, this approach is limited in time resolution, and generally requires stopping reactions for analysis. Optical monitoring of reactions of single AuNRs can also be accomplished *In Situ*. Dark-field spectral imaging has been used to study anisotropic etching of individual AuNRs by hydrogen peroxide,⁸¹ potassium iodide/iodine,⁸² and gold(III).⁸⁵ In a different experimental design, luminescence was employed to study the cyanide etching of AuNRs.⁷⁹ Additionally, dissolution of AuNRs *via* substrate voltage tuning has been monitored using dark-field hyperspectral imaging.^{80, 86} Perhaps the most commonly reported reagent for etching single AuNRs in recent years is iron(III) chloride, starting with bulk studies in 2009.¹¹¹ Since then, FeCl₃ etching of single AuNRs has been reported using dark-field monitoring, sometimes utilizing Le Chatelier's principle to drive the reaction.⁸³⁻⁸⁴ Ferric etching of single AuNRs using an electron beam, monitored by liquid-TEM, has also been reported.⁸⁷ Excepting the electron beam study, these reports evoked purely chemical mechanisms to explain their reported chemistries. However, the light-induced etching of AuNRs using FeCl₃ has also been reported, both in bulk studies,¹¹² and single AuNR experiments using one-photon luminescence.⁷⁸ Due to significant interest in ferric etching of AuNRs and the intriguing mechanistic parameter space, FeCl₃ was employed for etching in this report.

4.7. Main Text: Single AuNR Reactions

After single AuNRs deposited in the microbubble resonator were identified, they were chemically etched using FeCl₃. The etching solution, ranging between 250 μM and 2 mM FeCl₃ dissolved in dilute hydrochloric acid (pH~1.3) to prevent hydrolysis of the oxidant, was flowed into the microbubble. Due

to differences in sensitivity resulting from microbubble geometries, mode selection, and even nanorod location within the same microbubble, the relative photothermal signals between nanorods cannot be directly compared. However, the relative signal of one AuNR reacting over time, using the same resonance, directly maps onto a change in absorption cross-section of the nanorod at the pump wavelength, and thus its etching progress. Conveniently, nanorod etching was found to be photo-activated by the pump beam illumination (discussed further below), allowing controlled reaction initiation. The AuNRs were monitored by repeatedly rotating the linearly polarized pump, interrogating the relative absorption and orientation of the AuNR as it is etched. Importantly, before each polarization trace is taken, a beam-centering algorithm is used to mitigate any false signal decrease from spatial drift of the bubble. The centering also serves as a “dosing” period to enable AuNR etching between polarization traces. Each polarization data point required only 50 ms of data acquisition time, suggesting that fast chemical dynamics can be followed with our approach. A further discussion of time resolution and imaging of small AuNPs is presented below.

Figure 4-4A (i-iii) features three exemplary traces of a single AuNR reaction (additional examples in the **Supporting Information**). These three reactions were taken in different microbubbles on different days, confirming reproducibility of the experiment. A logarithmic version of reaction (i) (**Figure 4-4B**) readily shows the late stage continued reaction progression along with AuNR rotation. This behavior is better illustrated in the extracted maximum signal and angle traces from reaction (ii) seen in **Figure 4-4C** and **Figure 4-4D** respectively. “Exposure Time” refers to the total time that the AuNR has been exposed to the laser beam, which doesn’t include the “switching time” at the dotted lines, where the laser was turned off so that the power could be increased. The reactions slowed as they progressed, as seen by the plateauing effect in the maximum signal. This plateauing is a direct result of the photo-activation mechanism: as a nanorod shrinks, its absorption cross-section decreases, resulting in less light absorption and thus slower etching. This photo-activation is further confirmed by incrementally increasing the pump

power (dotted vertical lines), quickening the reactivity before plateauing once again. Although the three reactions in **Figure 4-4A** were taken at ferric chloride concentrations spanning almost an order of magnitude, the time scales of reaction vary by much less (discussion in **Supporting Information**). It is also evident that AuNRs sometimes rotate as they etch, especially late in reactions, as seen in **Figure 4-4D** and discussed later (see **Single AuNR Rotations**).

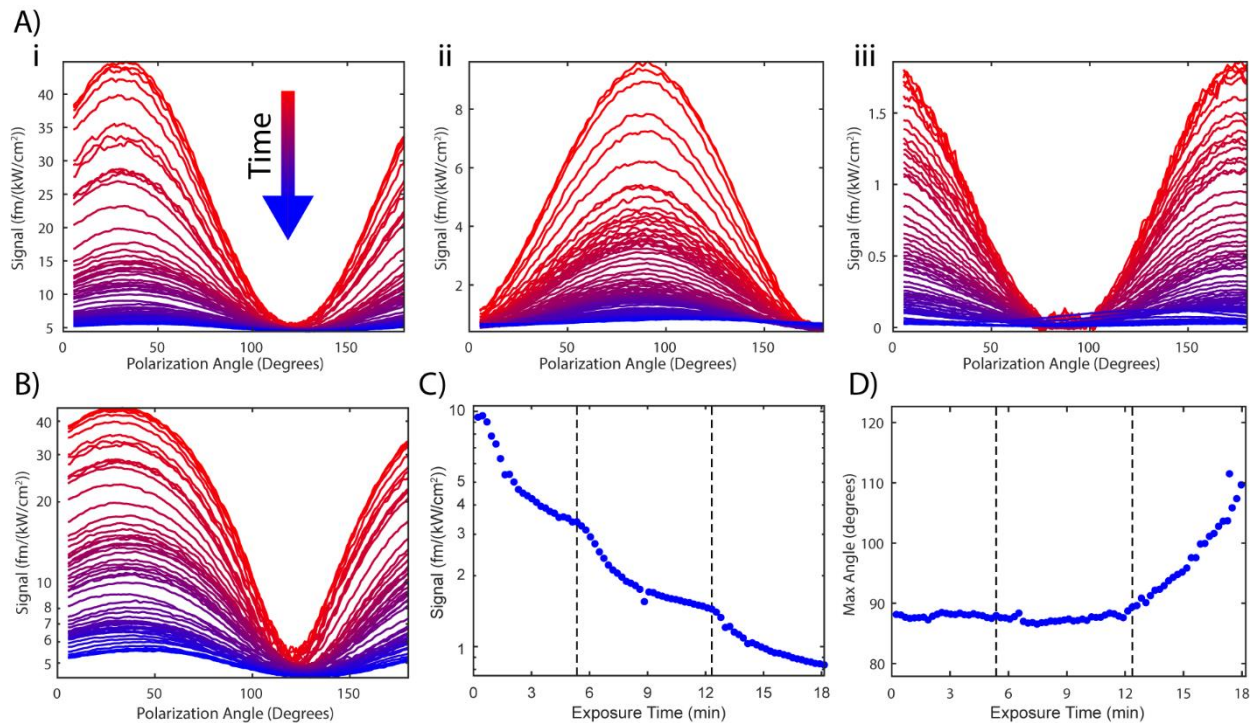


Figure 4-4. Etching single AuNRs. (A) Reaction series of polarization traces for three different reactions, progressing in time from red traces to blue traces. Maximum signal is normalized by pump flux. Each trace was taken over the course of 10 seconds (0.05 s per point, 200 different angles), with a 1 second delay between traces (except when switching power), and 3 seconds for beam-centering between each trace. (B) The data for reaction (ii), but with signal shown logarithmically. (C) Maximum signal of polarization traces over the course of reaction (ii), showing the decrease in relative absorption cross section at 635 nm. Dashed lines indicate points in time at which pump power was increased. (D) Maximum angle of polarization traces over the course of reaction (ii), showing nanorod orientation. Dashed lines indicate points in time at which pump power was increased. Reaction conditions: dilute aqueous HCl (pH~1.3), room temperature, varied FeCl₃ concentrations (i) 1mM, (ii) 250uM, (iii) 2mM. Pump fluxes for reaction (i) were 2.7, 6.7, 11.4, 21.0, and 34.5 kW/cm². Pump fluxes for reaction (ii) were 4.1, 11.9, 35.4 kW/cm². Pump fluxes for reaction (iii) were 6.4, 15.9, 31.6, 57.3, 121 kW/cm².

A control experiment was performed with a nanorod-containing resonator filled with dilute hydrochloric acid (pH~1.3) without FeCl₃, confirming that the acid alone is not enough to etch the nanorods under illumination. Additionally, when nanorods are left in etching solution for multiple days

without laser illumination, they do not observably react, supporting a photo-activated etching mechanism. AuNRs exposed to etching solution for hours within a microbubble, with the probe beam on and locked to water-dominated modes but no pump beam, also did not undergo significant etching, indicating that the probe beam is not sufficient to drive etching. Therefore, we hypothesize a photo-activated mechanism resulting from hot electrons generated from LPB decay (see **Mechanistic Discussion**). We also note that, occasionally, nanorods were “impervious” to photo-activated etching, as discussed further in the **Supporting Information**.

To exemplify the spectroscopic versatility of this platform, nanorod etching was also induced with the 532 nm pump beam. Because this wavelength pumps the TPB (as well as direct interband transitions), nanorod orientation can still be tracked. In **Figure 4-5A**, a single AuNR reaction time series is shown for 532 nm-driven (TPB-driven) conditions, with a logarithmic version of the data shown in **Figure 4-5B** for clear visualization of late-stage etching data. The polarization traces are conspicuously different than in **Figure 4** because of the presence of polarization-independent interband transitions. In **Figure 4-5C**, maximum signals are extracted for this reaction and one other TPB-driven reaction in the same resonator. For direct comparison, a different nanorod in the same microbubble was reacted using the 635 nm pump beam (LPB-driven), with the extracted maximum signals shown in **Figure 4-5D** (note, this is the same reaction shown in **Figure 4-4A (iii)**). Overall, three such reactions were performed for each color in the same resonator to confirm reproducibility (**Supporting Information**). Although quantitative comparison of reaction rates between experiments is difficult due to the decreasing rate as absorption cross-section diminishes, it is clear from the extracted maximum signals that while the etching rate of the 532 nm-induced reactions is faster than the rate of the 635 nm-induced reactions, it is not multiple orders of magnitude faster, in contrast to previously reported bulk measurements, discussed further below.¹¹² The shapes of maximum signal traces for the LPB-driven and TPB-driven reactions are

noticeably different, with the LPB-driven reactions yielding a concave up shape, and the TPB-driven reactions yielding a concave-down (**Figure 4-5c (i)**) or even sigmoidal shape (**Figure 4-5c (ii)**).

We also use data from a TPB-driven reaction to discuss imaging of small AuNPs and compare to darkfield microscopy. Typical dark field techniques can image AuNPs as small as 30 nm in diameter,¹¹³ with more optimized approaches pushing down to 10 nm.¹¹⁴ The threshold at which the absorption cross-section overtakes the scattering cross-section in magnitude is at a diameter of around 80 nm for gold nanospheres.⁷¹ At diameters of 20 nm, this ratio of $\sigma_{\text{abs}}/\sigma_{\text{scat}}$ is ~ 100 .⁷¹ The initial sizes of the AuNRs in this report (approximate volume $4 \times 10^5 \text{ nm}^2$, as calculated for an $80 \times 40 \text{ nm}$ cylinder) are already well below the volume of an 80 nm nanosphere ($2 \times 10^6 \text{ nm}^2$). Furthermore, after etching the AuNR the volume is much lower, and we can estimate that volume. For this calculation, we use a 532 nm-driven reaction because the TPB is not expected to shift significantly in its center wavelength regardless if the etching is isotropic or anisotropic, meaning that the measured signal should scale proportionally with nanorod volume.⁷¹ The inset of **Figure 4-5c(i)** shows a polarization trace from late in the reaction, with its corresponding data point indicated. Although later data points show even smaller signals, this data point was selected because the trace still clearly shows the polarization-dependence associated with the TPB. The signal at 180 degrees in the inset is lower than the signal at 0 degrees because the nanorod is reacting rapidly at the high pump flux (as indicated by the slope of **Figure 4-5c(i)**). The selected data point has a signal 34 times smaller than the initial signal, equivalent to a volume of $1.2 \times 10^4 \text{ nm}^2$, or a sphere with a diameter of 14 nm. Thus, even with time resolution of 50 ms per point, microbubble spectrometers are well-suited to study AuNPs of a size that is challenging to reach with dark-field measurements. Significant

further increases to our sensitivity can be achieved with use of media with higher thermo-optic coefficients^{69, 115-117} or use of more optimized optical modes (see discussion above).

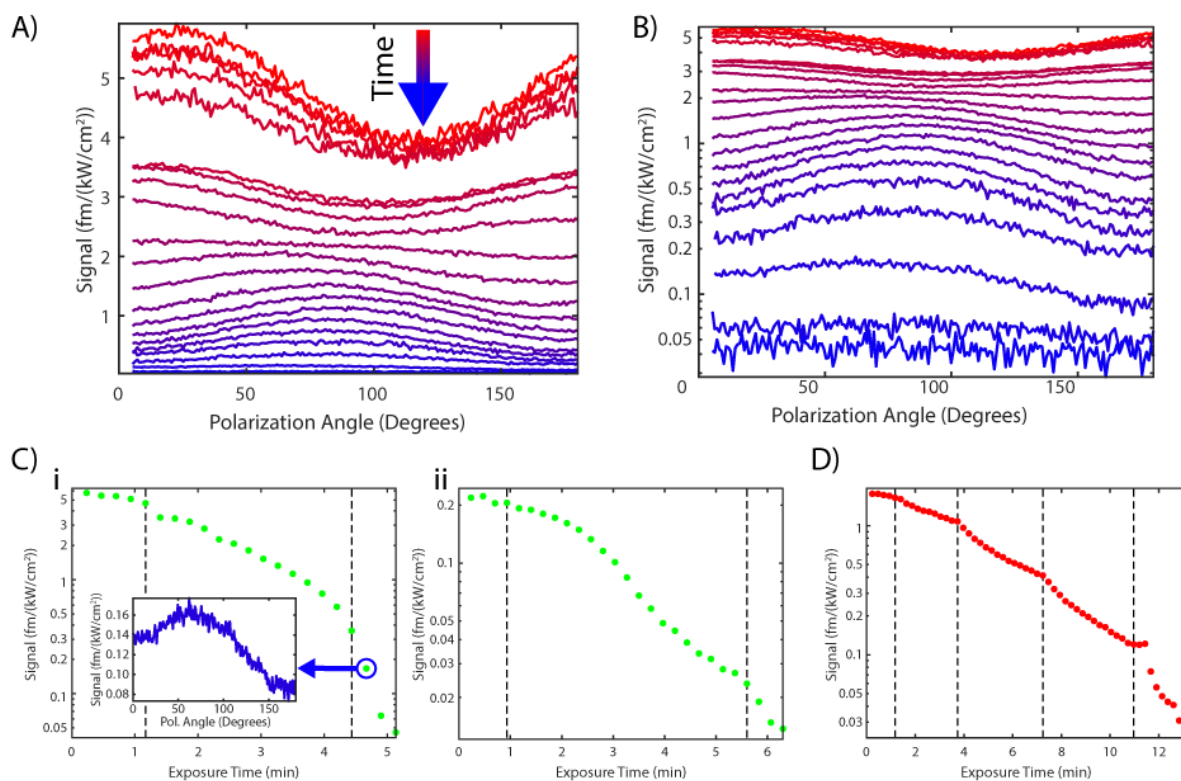


Figure 4-5. Etching reactions driven at two different pump wavelengths. (A) The reaction of a single AuNR being driven with the 532 nm pump beam, progressing in time from red traces to blue traces. Maximum signal is normalized by pump flux. Each trace was taken over the course of 10 seconds (0.05 s per point, 200 different angles), with a 1 second delay between traces (except when switching power), and 3 seconds for beam-centering between each trace. (B) The same data as in (A), but with signal shown logarithmically. (C) (i) Maximum signal of polarization traces for the reaction shown in (A). Polarization trace for indicated data point in inset. (ii) A similar trace for a different nanorod reacted in the same bubble using the 532 nm beam. Dashed lines indicate time points at which pump power was increased. (D) Maximum signal of polarization traces for the reaction of a nanorod in the same bubble, but using the 635 nm pump beam to drive the reaction. Dashed lines indicate time points at which pump power was increased. Reaction conditions: dilute aqueous HCl (pH~1.3), room temperature, 1mM FeCl₃. Pump fluxes for reaction (Ci) 16.1, 39.2, 63.3 kW/cm². Pump fluxes for reaction (Cii) were 6.4, 31.6, 57.3 kW/cm². Pump fluxes for reaction (D) were 6.4, 15.9, 31.6, 57.3, 121 kW/cm².

4.8. Main Text: Mechanistic Discussion

Our proposed mechanism relies on the generation of hot electrons, as in previous reports of photo-activated AuNR etching.^{111, 118} Interest in such hot carrier chemical processes has exploded in recent years in a variety of applications,¹¹⁹ especially photocatalysis.¹²⁰⁻¹²² The mechanisms of hot carrier generation and transfer have been extensively studied,¹²³⁻¹²⁶ including efforts towards untangling the

contributions of hot carrier effects and photothermal effects in nanoparticle synthesis¹²⁷ and plasmonic photocatalysis,¹²⁸ and mapping hot carrier driven catalytic reactions nanoscopically.¹²⁹⁻¹³⁰ Although there are still unresolved questions regarding plasmon-driven chemistry, hot carrier transfer can generally be predicted by using an energy overlap model.¹³¹

In the reported reactions, we expect a nominal temperature rise at the nanorod's surface of less than 1K as calculated from the known absorption cross-section and excitation intensity (**Supporting Information**), meaning that the observed etching mechanism should not be significantly influenced by photothermal heating.⁷⁸ With no noticeable dark reaction rate, and a photothermal mechanism ruled out, a hot carrier mechanism, whereby the decay of LSPRs or excitation of interband transitions results in hot electrons that can transfer to ferric ions on the nanorod surface, must be invoked to explain the observed reactions. Hot electrons lower the Gibbs free energy of the etching reaction,⁸⁴ and reduce the thermal activation barrier for electron transfer,¹²⁸ thus modifying both the thermodynamics and kinetics of the reaction. Photoexcitation of plasmonic NPs has been experimentally shown to lower the activation enthalpy for transferring electrons from gold nanoparticles to Fe³⁺,¹³ and to lower the energy barrier for reaction at the surface of plasmonic NPs by affecting ligand-NP interactions.¹³²

When the LSPRs of an AuNR decay, hot electrons are generated from the conduction band of the nanorod. These hot electrons are on average at lower energy than those generated by interband transitions, though a few carriers will be hotter in the LSPR decay case.¹¹² Although it was previously reported that interband pumping can drive etching orders of magnitude faster than LPB pumping,¹¹² this was not observed in our reported reactions. CTAB concentration and halide concentration play significant roles in AuNR etching^{111, 118} and we attribute this observed discrepancy to different CTAB-mediated mechanisms. Our nanorod deposition procedure likely removes a significant part but not all of the CTAB (see **Supporting Information**).¹³³⁻¹³⁴ Thus, the AuNRs in our study have significantly lower CTAB coverage than the report where pumping the interband transition resulted in order of magnitude increase in etch

rate, where etching occurred in a medium with CTAB concentration greater than the CMC.¹¹² Ultimately, the reaction relies on ferric ions binding to the nanorod before hot electron transfer can take place. Thus, a plausible mechanism for the reported reactions entails a two-step process, whereby slow ligand exchange of CTAB with Fe^{3+} , or intercalation of the Fe^{3+} through the residual CTAB, is followed by fast photoactivated etching. A relative increase in the rate of the photoactivated step, as seen previously,¹¹² is then ultimately masked in the observed rate due to the slowness of the first ligand exchange step, as shown **Figure 4-6**. In the limit where ligand exchange is slow, increases in the rate of the photo-induced step would give a somewhat muted effect on the overall reaction rate, as observed.

To further understand the shape of the reaction profiles, we modeled the spectral changes of the TPB and LPB for a nanorod being etched for etching schemes ranging from tip-only etching to side-only etching (**Supporting Information**). Modelling was able to reproduce the concave-up shape observed in the LPB-driven reactions (**Figure 4-4C**). However, no combination of variables was able to capture the concave down or sigmoidal trend seen in some TPB-driven reactions (**Figure 4-5C**), or fully reproduce the threshold behavior observed. Thus, with a completely linear reaction mechanism ruled out, we can speculate on possible origins of non-linearity in the etching mechanism. One possible origin stems from the evolution of nanorod morphology over time. For example, it was observed previously that for certain laser powers, nanorod LPBs would red-shift, then blue-shift.⁷⁸ Another possible origin derives from the changing surface concentration of CTAB, with a relatively dense coverage providing competitive inhibition for ferric ion binding at early times but AuNR etching resulting in easier access at later times. Future studies utilizing multiple pump beams could be valuable in studying these complex kinetics, as the evolution of the LPB and TPB spectra would yield important insight into the reaction mechanism.

Thus, the reported microbubble platform may be used for spectroscopic, mechanistic studies into the wavelength-dependence of hot-carrier driven chemical dynamics in single plasmonic nanoparticles.




Figure 4-6 . Proposed mechanistic explanation for etching rates. A slow initial step requiring CTAB dissociation before ferric ions can bind determines the overall rate for the reaction, muting the effect of the higher rate for TPB excitation, even though hot electrons are more efficiently generated. Eventually, etching stops when the absorbed light falls below a threshold necessary for hot-electron-driven etching.

4.9. Main Text: Single AuNR Rotations

Beyond using our microbubble spectrometer to monitor and control nanoparticle size, we can also use it to monitor and control nanoparticle orientation, adding significant utility to the microbubble spectrometer platform. Indicated by the shift in peak polarization during etching reactions (**Figure 4-4** and **Figure 4-5**), AuNRs can rotate while etching. Alternatively, active rotation can be induced with the pump beam, allowing for control over nanorod orientation. This control results from the optical torque exerted by linearly polarized light on an anisotropic, absorbing plasmonic structure (**Figure 4-7a**), a phenomenon that has been demonstrated experimentally¹³⁵ and theoretically,¹³⁶⁻¹³⁷ and is discussed further in the **Supporting Information**.

Theory predicts that the optical torque acting on the AuNR from the 635 nm pump beam will align the AuNR perpendicular to the polarization of the incident beam. Indeed, this perpendicular alignment is exhibited upon sufficient excitation power. In this way nanorod orientation can be controlled to within approximately 10 degrees (**Figure 4-7b**) as AuNR orientation is stepped through a ~180 degree rotation. This control was accomplished by monitoring the photothermal signal, dislodging the nanorod with a large optical torque above some threshold incident power, and dithering the polarization until the photothermal signal was minimized at the desired polarization angle. This thresholding behavior was demonstrated further by a stepwise ramping of the pump laser intensity, resulting in an upward staircase of photothermal signal, until rotation was finally induced. As can be seen in the examples in **Figure 4-7C**, two different nanorods required significantly different pump powers to dislodge them from the microbubble surface. Following dislodgment, the signal quickly stabilized to around 70% of the maximum signal for the left nanorod trace, whereas it behaved semi-stochastically for the right trace, before settling at <40% of the maximum signal. These differences in orientational dynamics highlight the differences in the local environments around the two nanorods, including both Coulombic effects and refractive index differences. These staircase experiments were performed repeatedly for both nanorods in **Figure 4-7C** to confirm reproducibility (**Supporting Information**).

Light-induced rotation during reactions was observed more frequently as AuNR etching progressed. Likely, as Coulombic attractions between the AuNR and the resonator's surface were weakened, hydrodynamic or optical torques were allowed to rotate the AuNR. Although rotation events could be forced in water-filled resonators, higher pump thresholds were generally required, and AuNRs were immune to rotation at powers that would result in rotation in ferric chloride solution. Therefore, it appears that the presence of etchant reduces the Coulombic attraction between nanorods and the resonator wall, possibly through charge screening, permitting facile rotation. Though optical rotation of

nanorods has been seen in previous experiments, the coupling between evolving surface chemistry and propensity for rotation has not been explored to our knowledge.

Radiation pressure and optical gradient forces could also influence AuNRs, affecting the rotation power threshold. However, varying the pump beam focus position, which would change forces along the optical axis, did not significantly impact rotation thresholds. Therefore, it appears that the torque described above is the dominant driver of nanorod rotation. Although scattering forces in three-dimensional trapping can orient anisotropic plasmonic nanoparticles parallel to the optical axis of the excitation beam,¹³⁸ one would expect this to result in a highly stochastic signal over time in the polarization traces of **Figure 4-3** and **Figure 4-4**, as well as a revival of signal upon shutting off and turning back on the pump beam. Such behavior was not observed. Therefore, AuNRs likely remain parallel to the plane of the resonator's surface during rotation, rotating only in two dimensions.

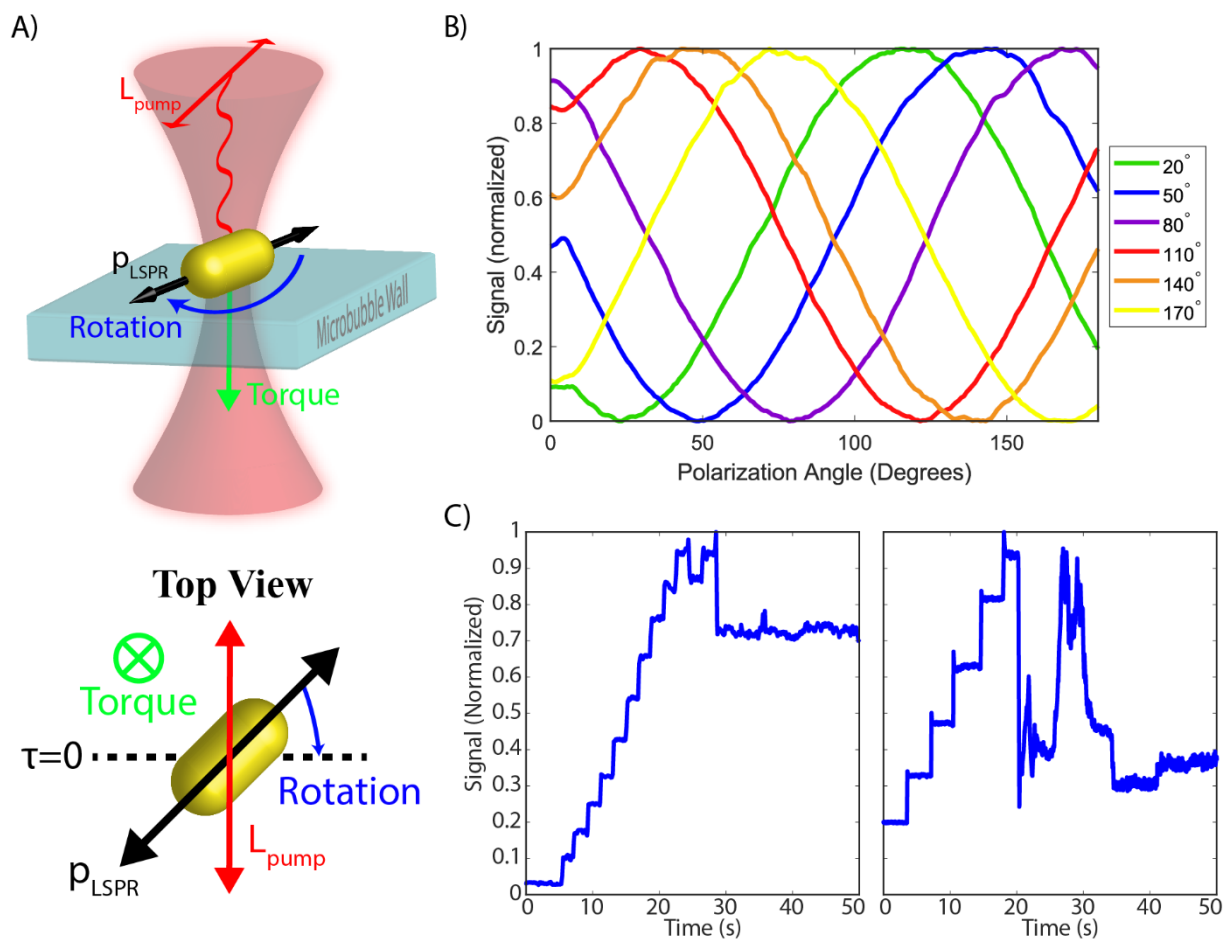


Figure 4-7. Orientation control of single AuNRs. (A) A cartoon illustrating the optically induced torque that a AuNR experiences under illumination with linearly polarized light, both from side-view (*top*) and top-view (*bottom*). (B) A series of pumping experiments showing optical control of nanorod orientation. (C) Trace showing the photothermal signal as two different AuNRs are pumped at increasing laser powers until the AuNRs dislodge slightly from the resonator wall and rotate, eventually settling down off-axis of the polarization.

4.10. Main Text: Conclusions

We have demonstrated microbubble resonators as a robust platform for studying chemical dynamics in solution *via* single-particle absorption spectroscopy. We used a microbubble spectrometer to observe changes in the optical properties of AuNRs as they were controllably etched by ferric chloride *via* a photo-induced mechanism. The sizes of the etched AuNRs push the limit of what can be imaged *via* dark-field scattering. The hypothesis that ligand-exchange-limited etching is the reason for muted wavelength dependence of etching will require more experiments in the future to confirm and

understand. Additionally, we monitored and controlled the orientation of the AuNRs using optical torque.

With this demonstration, we lay the groundwork for studying more complex reaction dynamics of single particles and molecules. Thus, this technique provides a complementary measurement to the luminescence and dark-field methods previously used to observe similar reactions as reported here. In particular, the demonstrated exquisite sensitivity offers prospects of examining non-emissive objects inaccessible with fluorescence and too small to observe with scattering, which scales as $1/\text{Volume}^2$, a more severe penalty than in absorption measurements, which scale more favorably as $1/\text{Volume}$. Furthermore, rotational control could be used to estimate Coulombic forces attaching deposited objects to the resonator, helping to understand the interface between nanoparticles and the surface. This knowledge, combined with structured light field manipulation of nanoparticles,¹³⁹ might be used to arrange arrays of plasmonic nanoparticles as desired. The optical control of plasmonic nanoparticles within a microbubble resonator may allow for in-solution, photonic-plasmonic assembly, and live-control of emergent optical properties in such coupled systems. Additionally, by providing a direct thermal readout, our method could be used to untangle the respective contributions of photothermal heating and hot carrier generation for nanoparticle reactions, aiding the design of improved nanocatalysts. Overall, there is a compelling case for the use of microbubbles in materials studies, sensing, and chemical kinetics, and even hybridizing them with plasmonic or acoustic sensing schemes for further applications. Microbubble absorption spectrometers thus hold great potential for pushing the frontiers of absorption spectroscopy at the nanoscale.

4.11. Main Text: Methods/Experimental

4.11.1. Microbubble fabrication

Microbubble resonators were fabricated according to the method reported by Yang and coworkers.¹⁴⁰

First, a glass capillary (Polymicro Technologies, TSP250350) is tapered using a heat-and-pull method, until

it is approximately 25 μm in diameter. Next, counter-propagating CO_2 laser beams are focused onto the capillary while positive pressure is applied from the inside of the capillary using an inert gas. The heat from the laser beams softens the capillary, allowing for the local expansion of the capillary to 50-100 μm in diameter, depending on the experimental parameters. Eventually, radiative cooling from the expanded glass outcompetes the expansion process, and the bubble's size stops increasing. To operate in the quasi-droplet regime, a wall thickness close to the wavelength of the laser beam used for WGM excitation is desirable. Microbubble wall thickness is determined by an equation reported by Henze and coworkers,¹⁴¹ and validated separately by others.¹⁴²

4.11.2. Tapered optical fiber fabrication

Single-mode optical fiber was purchased from Corning (HI 780C). Tapered fibers are made by removing the polymer sheath, cleaning the fiber, and tapering using a heat-and-pull method with a hydrogen torch and motorized actuators (Thorlabs Z825B) until the fiber returned to single-mode, as determined using a 785nm diode laser (Thorlabs LPS-785-FC) and optical power meter.

4.11.3. Instrumentation for photothermal spectroscopy

A tunable, ultra-narrow linewidth, fiber-coupled CW laser (Newport TLB-6712) with a wavelength range of 765-781 nm was used for coupling into resonators. Pound-Drever-Hall locking electronics were constructed as previously reported,¹⁰ except for the use of a different voltage-controlled oscillator (Mini Circuits ZX95-310A+) and different lithium niobate phase-modulator (EOSPACE PM-0S5-01-PFA-PFA-765/782). PDH feedback was applied to the tunable laser using high-speed servo controller (Newport LB1005). The optical output from the experiment was collected using an APD (Thorlabs APD430A), and the photothermal signal was extracted using a lock-in amplifier (Ametek 7265). The resulting signal was collected using a Data Acquisition (DAQ) card (National Instruments BNC2120) for later processing. Custom LabVIEW code was used for instrumentation control. For photothermal mapping, a lock-in time constant of 20 ms or 50 ms was used. For polarization traces, a lock-in time constant of 50 ms was used.

Diode lasers were used for pump beams, with the wavelengths 532nm (FTEC2 532-20), 635 nm (FTEC2 635-50), and 785 nm (Thorlabs LPS-785-FC). The pump beam was amplitude modulated using an optical chopper system (Thorlabs MC200B), and steered using galvanometer mirrors (Thorlabs GVS212) run by outputs from the DAQ mentioned above, modified using custom electronics. The pump beam was focused using a piezo-controlled (Thorlabs DRV517) objective (Nikon Plan 40x, 0.65 NA). Pump beam polarization was controlled using a three-optic system, comprised of a linear polarizer (LPVISE100-A), followed by a liquid crystal variable retarder (Thorlabs LCC1423-A, LCC25) with its fast axis set 45 degrees relative to the polarization axis, followed by a zero-order achromatic quarter-wave plate (Thorlabs AQWP05M-600) with its fast axis set 45 degrees relative to the liquid crystal's fast axis. In this design, tuning of the liquid crystal voltage results in a rotation of linearly polarized light at the output of the three-optic system.

4.11.4. Polarization plots

Shown in Figure 3C and Figure 3D, AuNRs are first identified by photothermally mapping them at different pump polarizations, and processing these maps with a 2D-Gaussian fit, which results in background subtraction. Then, the maximum signals for each plot are together fit to provide a depth-of-modulation (M), using Equation 4-1, which also gives the maximum signal (σ_{\max}) plotted in Figure 4C, and the polarization angle of the maximum signal (θ_{\max}) plotted in Figure 4D.

$$\text{Equation 4-1} \quad \sigma(\theta) = \sigma_{\max} \left(1 - M * \sin^2(\theta - \theta_{\max}) \right)$$

For the data in Figure 3E, the liquid crystal is used to rapidly collect many data points that are then fit to equation (1) to obtain M . This method does not include background subtraction.

4.11.5. Bulk/UV-Vis studies

The bulk absorption spectrum in Figure 3B was taken using a UV-Visible Spectrophotometer (Varian Cary 50). Additionally, studies were conducted to confirm the effects of CTAB concentration and FeCl_3

concentration on bulk nanorod etching. The effects of added NaCl were studied to further examine the impacts of chloride concentration. Results and further discussion are in the **Supporting Information**.

4.11.6. Nanorod deposition in microbubbles

All chemicals were purchased through Sigma-Aldrich unless otherwise noted. To deposit nanorods in a microbubble, a 500x serial dilution is made of AuNRs (Nanopartz A12-40-650-CTAB-DIH-1-25, size: 80 x 40 nm, zeta potential: 35 mV, stock pH: 7, stock CTAB concentration: 5mM) in a solution of 200x diluted HCl and 25 μ M CTAB in water. The low CTAB concentration prevents nanorod aggregation during deposition, but keeps the CTAB concentration well below the CMC of \sim 1mM. Dilute hydrochloric acid, which results in a pH of around 1.3, encourages binding of the nanorods on the resonator interior by enhancing Coulombic interactions.⁴⁴ For deposition, water is first flowed through the resonator using a syringe pump attached to the first port of the resonator's capillary. Then, dilute HCl is flowed through the resonator to prime the glass surface for deposition. Next, deposition solution is backfilled through the second capillary port, which is cut to a much shorter length to reduce deposition of AuNRs to the capillary's interior walls. Following this, dilute HCl is flowed through the resonator, followed by water, through the first port of the capillary to push out the deposition solution while maintaining a pH gradient. Water is flowed through the resonator for at least several minutes to ensure the removal of non-bound objects, and remove excess CTAB from the nanorod surfaces.

4.11.7. Reactions in microbubbles

All chemicals were purchased from Sigma-Aldrich. Reaction mixtures are made by dissolving and serial diluting ferric chloride hexahydrate in 1/200 dilute hydrochloric acid, resulting in a solution pH of around 1.3. While reaction solution was flowed into the bubble, resonances shift as the refractive index being probed by the WGM changed. Complete filling of the microbubble with reaction mixture is indicated when the resonances have stopped shifting. Following this stabilization, the syringe pump pressure is released, resulting in a microbubble primed for etching experiments.

4.12. Supporting Information: Geometric Parameters of Microbubbles

An equation from Henze and coworkers¹⁴¹ (Equation 4-2) is used to calculate the wall thickness of the bubble, w :

Equation 4-2

$$w = a - \sqrt{a^2 - \frac{1}{4}D_1^2(1 - \xi^2)}$$

where a is the microbubble radius (as measured by optical micrograph), D_1 is the outer diameter of the tapered capillary (as measured by optical micrograph), and ξ is the constant ratio of inner capillary diameter to outer capillary diameter (taken from specification sheet of original capillary, or measured by cross-section of non-tapered capillary), which for the bubbles used in this report is $250 \mu\text{m}/350 \mu\text{m}=0.714$. For the bubbles used for the reported data in the main text, parameters are extracted from this equation using the optical micrographs shown below (**Figure 4-8**), and are reported in the following table.

Table 4-1. Geometric parameters of microbubbles used in this paper.

Experiment	Tapered Capillary Width (μm)	Microbubble Radius (μm)	Bubble Wall Thickness (μm)
Red Reaction (i) (Figure 4)	20	33	0.75
Red Reaction (ii)	23	47.5	0.73
Red Reaction (iii) & Red/Green Experiments	22	48.5	0.61

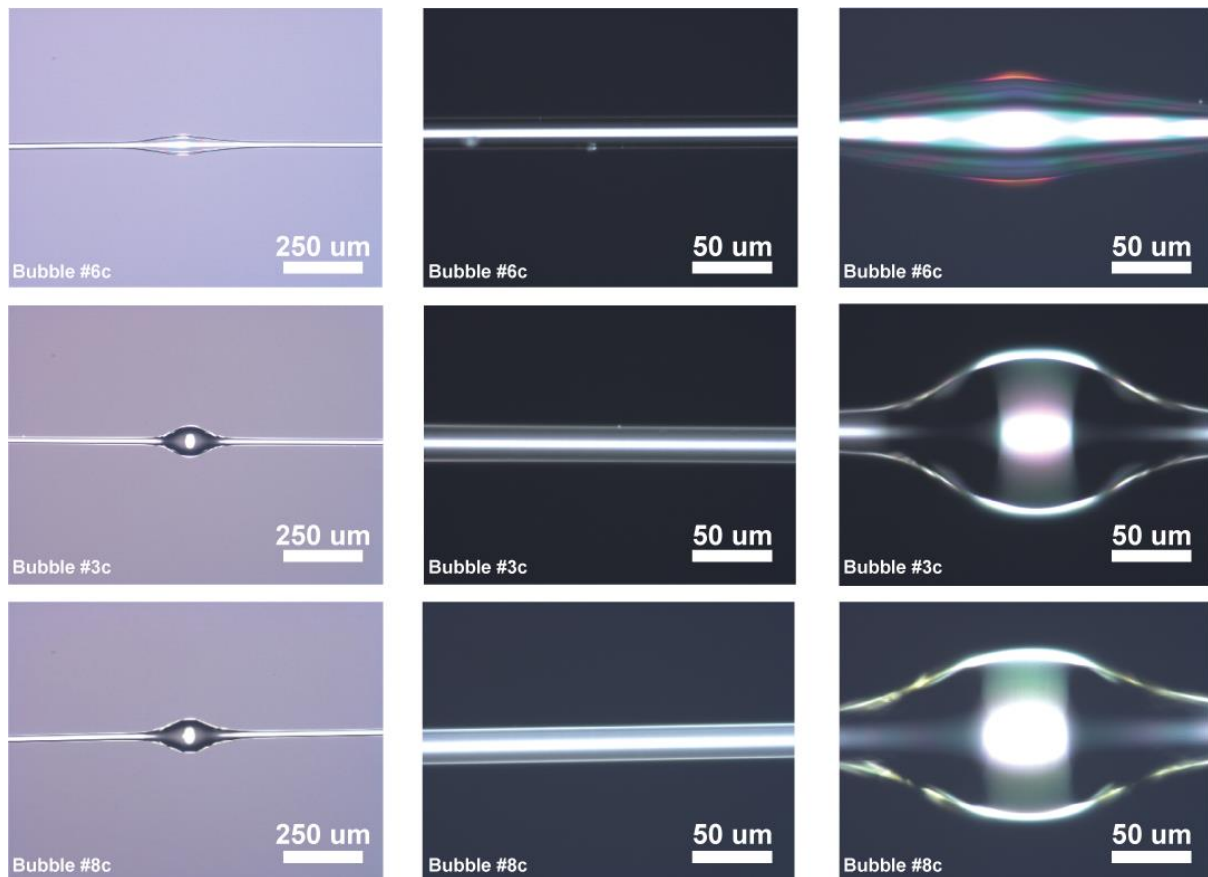


Figure 4-8 Optical micrographs of the microbubbles used to collect the data in the main text.

4.13. Supporting Information: COMSOL Simulations

As in our previous report on microtoroid spectrometers,¹⁰ we used COMSOL to model the expected resonance shift from optically pumping a nano-object on the microbubble. The five main steps for this are described below, and shown in **Figure 4-9**.

- **Step 1:** The microbubble geometry is constructed. Importantly, the wall thickness is calculated from the bubble radius and tapered capillary diameter, as described by Henze and coworkers.¹⁴¹ The bubble itself is modeled as fused silica, the core is modeled as water, and air is included for a prescribed radius around the microbubble

- **Step 2:** A point heat source for a given power is modeled for a prescribed time of resonator heating using a thermal simulation. This can be done with a modulating envelope to represent amplitude modulation of the pump beam. Steady state temperatures can also be calculated; this is described more below.
- **Step 3:** Eigenvalues (energies/wavelengths for resonances) and eigenvectors (spatial profiles of the electric field for resonances) are calculated for the same geometry as used for the thermal simulation, using an optical simulation developed by Oxborrow.⁹
- **Step 4:** The two simulations are then overlapped for the optical modes of interest.

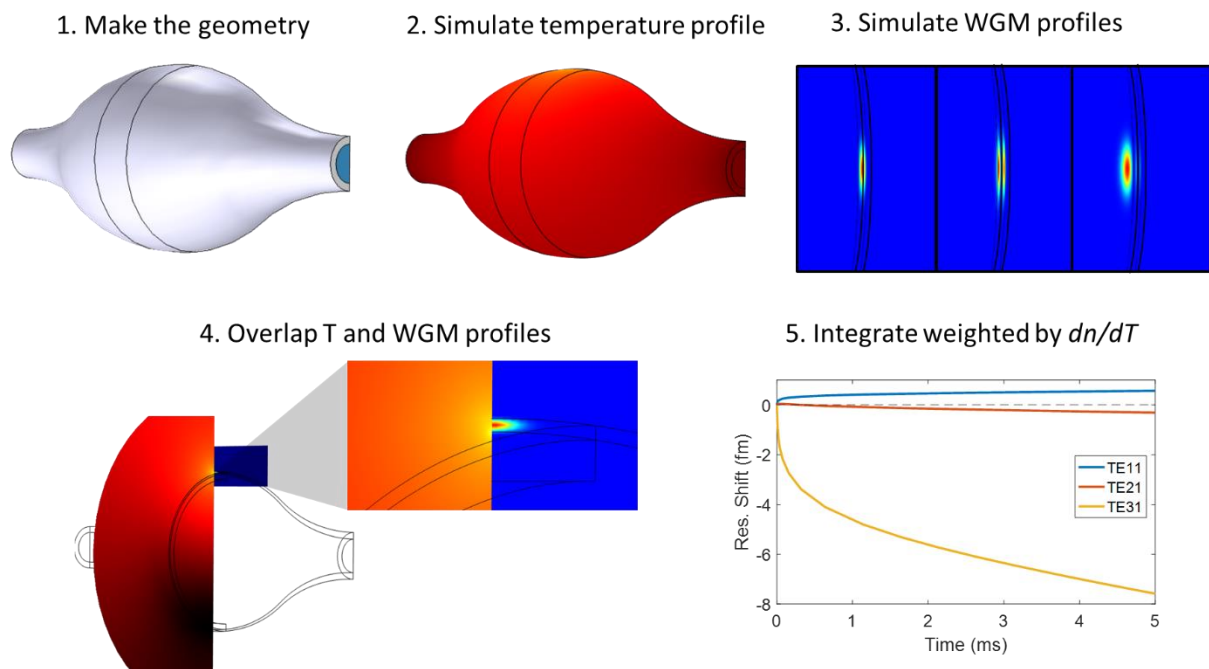


Figure 4-9. Steps in a simulating the resonance shift from a point-heat source.

- **Step 5:** The resonance shift is calculated for each mode by using an overlap integral, weighted by the dn/dT for each point calculated.

Geometric parameters for simulations in the main text and this Supporting Information are given in the table below. Wall thickness for each model was calculated according to the equation reported by Henze and coworkers.¹⁴¹ All simulations were run around a probe wavelength 780 nm.

Table 4-2. Microbubble thermal heat-sink simulation results.

#	Simulation	Bubble Diameter (μm)	Bubble Length (μm)	Capillary Diameter (μm)	Capillary Length (μm)*	Air Diameter (μm)^	Heat Power (nW)
1	Figure 2A	45	100	12	N/A	200	N/A
2	Figure 2C	45	100	12	N/A	200	10
3	Figure S3	45	100	12	Variable	Variable	10
4	Figure S4C	40	120	10.95	200	560	1000

*On each side, beyond the bubble.

^Around axis of microbubble/capillary (see **Figure S4**).

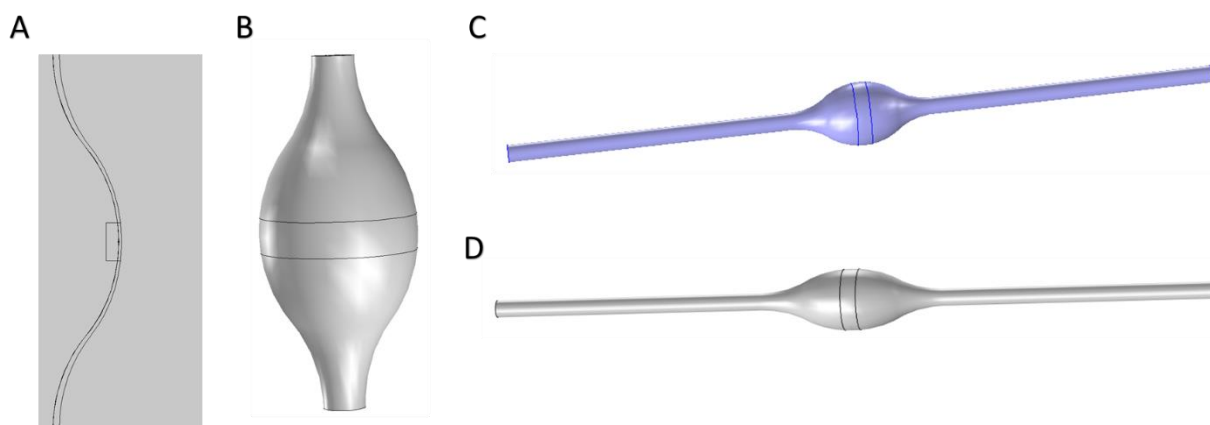


Figure 4-10. Pictures of microbubble models used for the COMSOL models in this report. Air cylinders omitted for clarity in revolved images.

Figure 4-10 shows example portions of COMSOL models. **Figure 4-10A** shows the geometry used for simulation table entry #1, with the revolved version of the bubble used for thermal simulations in entry #2 shown in **Figure 4-10B**. A version of this with longer capillary and added air around the resonator was

used for entry #3 of the above table (**Figure 4-10C**, air cylinder omitted for clarity). These results are further described below (**Figure 4-104**). A thinner, longer microbubble as used for table entry #4 (**Figure 4-10D**, air cylinder omitted for clarity).

In **Figure 4-9**, step 5, the microbubble clearly continues to heat for longer than the simulation was run. A high-mesh simulation with many time steps is too intensive to model for more than a few milliseconds. Thus, a steady-state thermal simulation at $\frac{1}{2}$ heating power (corresponding to the average of the modulated pump beam) can first be run to calculate the equilibrium temperature of the system, followed a simulation with amplitude modulation around the equilibrium temperature (**Figure 4-2C** in main text). This temperature was found to rely on both the length of additional capillary beyond the microbubble, and the radius of air included around the microbubble and capillary. Results from these steady state simulations are shown below in **Figure 4-11**. In **Figure 4-11A**, the parameters representing capillary length and radius of air are shown. In **Figure 4-114C**, the steady state resonance shift is calculated for different ratios of air-radius (R_{air}) and microbubble-radius (R_m), for a value of L_{cap} of 200 μm . It is evident that the solution approaches asymptotically to a resonance shift of around -60 fm for the given simulation. In **Figure 4-114B**, the results of the R_{air}/R_m model are calculated for an increasingly longer capillary. Only the data points for $R_{air}/R_m=2$ and sufficiently larger R_{air}/R_m ratio such that increased R_{air} does not alter the results are shown. It is evident that for these large R_{air}/R_m ratios that adding capillary length diminishes in effect as length increases, reaching a reasonably stable threshold at around 200 μm). Green circles indicate a data point shared between the two plots; green triangles indicate a different data point shared between the two plots. Experimental confirmation of this offset, and

experimental and theoretical details about its magnitude relative the modulation amplitude, is described in the following section.

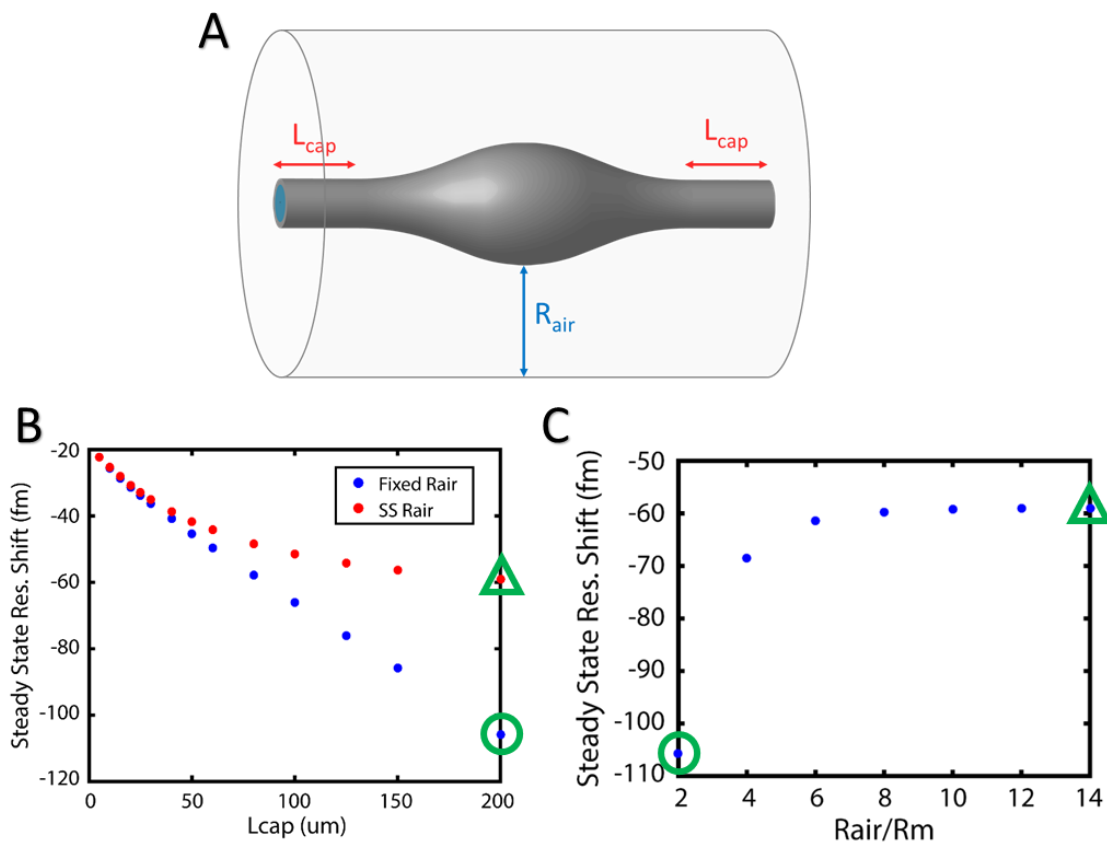


Figure 4-11 Steady-state COMSOL simulations. (A) A cartoon describing the two parameters varied in these simulations. (B) The steady state resonance shift for increasing capillary length, shown for both $R_{air}/R_m=2$ (blue dots), and $R_{air}/R_m=14$ (Red dots). (C) Steady state resonance shift as a function of air surrounding the microbubble, for 200 um of additional capillary. Green circles indicate a data point shared between the two plots; green triangles indicate a different data point shared between the two plots.

4.14. Supporting Information: Comparison with Microtoroids

Microbubble and microtoroid absorption spectrometers¹⁰ have different figures of merit that will inform resonator selection for a given experiment. A finite element simulation of the photothermal signal for a microbubble is given in **Figure 4-2C**. The “shark fin” shape results from the pump beam amplitude modulation. This modulation rides atop a rising baseline as heat builds up over many modulation cycles (left panel) before thermal equilibrium is reached (right panel), a phenomenon not observed with toroidal resonators.¹⁰ This difference arises from the lack of an effective and proximal heat sink for

microbubbles, in comparison to the silicon pillar underneath the center of a toroid, which acts as an excellent heat sink to permit thermal equilibrium within a single amplitude modulation. Further simulations to understand this steady-state offset, which is typically around 10x the modulation amplitude (for a modulation speed near 500 Hz), are included in the **Supporting Information**. As previously reported,¹⁰ thermal expansion of the resonator is expected to impact the resonance shift much less than the refractive index change (**Supporting Information**), and thus this effect is ignored for the remainder of this report.

To compare the sensitivity of microtoroid and microbubble spectrometers, the Q , V , and dn/dT of each platform must be compared. Microbubble resonance Q 's depend on what modes are being excited. At 785 nm pump wavelength, low-order radial modes propagating within the glass wall experience Q -factors as high as 5×10^7 ,^{90,143} and maintain mid- 10^6 Q -factors when the bubble is filled with water, while higher-order radial modes propagating in the water core typically have Q -factors of mid- 10^5 .¹⁴⁴ Comparatively, Q -factors at 1550 nm for SiO_2 -Si toroids reside at mid- to high 10^7 ,^{10,36} and for all-glass toroids at mid- 10^6 .⁴⁸ To compare the mode volumes of each platform, the refractive index contrast guiding modes must be considered. A toroid in air has a mode-guiding interface of air-glass, while a toroid immersed in water has a water-glass interface, yielding lower refractive index contrast. Conversely, a microbubble filled with water maintains air on its exterior, keeping the air-glass interface, and allowing for smaller-diameter (and thus smaller mode volume) microbubbles without reducing Q factor. Although such high-order water modes spread out further in space than the glass modes, smaller microbubbles will mitigate this factor, making the mode volumes of the two platforms comparable. The dn/dT of fused silica is $9 \times 10^{-6} \text{ K}^{-1}$, while the dn/dT of water is $-91 \times 10^{-6} \text{ K}^{-1}$. Thus, while most microtoroid modes have the same magnitude of resonance shift for a given temperature rise, the modes of a microbubble will vary wildly in both the direction and magnitude of their shifts, with the most thermally responsive modes experiencing a shift 10x that of a toroid mode. This variation in resonance shifts is shown both

theoretically (**Figure 4-2C**) and experimentally (**Supporting Information**). In the future, solvents with even higher thermo-optic coefficients may be exploited to enhance thermal responsivity, as has been done previously for single-molecule studies.⁶⁹

For toroidal microresonators, a limit-of-detection of less than 1 am of wavelength shift (100 Hz in frequency) at 1550 nm has been demonstrated,¹⁰ well below the calculated shift of 2 fm for a single chromophore pumped at near-saturating power.⁵⁰ To experimentally determine the limit-of-detection of this microbubble platform, a single AuNR inside a resonator was optically pumped, with data averaged for 30 seconds at a 1-second lock-in time constant (**Figure 4-2D**). This procedure was repeated for decreasing levels of power until the signal was indistinguishable from the signal obtained with the pump beam off. As the inset in **Figure 4-2D** shows, the signal remains linear over multiple orders of magnitude, flattening out at low powers near the limit-of-detection as expected. The water-mode used to take this data had a Q-factor of around 10^5 . Other resonances that experience different values of Q, V, and dn/dT will have different thermal responsivities and detection thresholds. We report that the detection limit for this platform is in the low tens of attometers of wavelength shift, surpassing the threshold for photothermally measuring a single chromophore.

4.15. Supporting Information: Static Offset vs. Modulated Signal

In our experiments, the size of the static offset is large in comparison to the modulated signal. This is due to the lack of an effective, proximal heat sink (see main text). Below in **Figure 4-12**, both experimental and theoretical results are presented to determine the approximate ratio of static offset to modulated signal. In **Figure 4-12A**, the laser feedback while optically pumping an analyte reveals a static offset of approximately 10.3 pm. In **Figure 4-12B**, the photothermal signal reveals a modulation amplitude of approximately 1.3 pm for the same experiment. Thus, in this experiment, modulation amplitude was about 12% of the static offset in magnitude. In comparison, a photothermal simulation (**Figure 4-12C**) for a microbubble with the length (200 μm) and air radius ($R_{\text{air}}/R_{\text{m}}=14$) (described above as the parameters

that reasonably approximate the steady state offset), with a $1 \mu\text{W}$ heat source, gives a comparable value of around 7%. The magnitude of the y-axes is different for the theory and experiment shown due to differences in the sizes in heat sources between the experiment and theory. However, because the photothermal effect is very linear with respect to thermal power in the range of our experiments, this comparison is still valid. This and other experimental observations confirm that the steady state photothermal static offset is approximately 10x larger than the modulating signal for a microbubble resonator with an amplitude modulation of around 500 Hz. Although the lock-in measurement should be agnostic to the static offset, mode congestion can result in different modes shifting through each other as the resonator heats up, resulting in experimental challenges (see main text).




Figure 4-12. Comparison of photothermal static offset to modulated signal. (A) Experimental offset for optically pumping an analyte, determined from the laser feedback signal. (B) Experimental photothermal signal, with the modulation amplitude indicated. (C) COMSOL simulation for a water mode, with a $1 \mu\text{W}$ heat source.

4.16. Supporting Information: Mode Shifting for Different dn/dT Values (Experimental)

Data is shown in **Figure 4-13** for pumping an analyte in a water-filled microbubble using a taper with appropriate thickness so-as to couple into both water and glass WGMs. The data shown is the error signal, rather than the transmission channel. Therefore, the resonance center is at the middle of the error curves, rather than the minimum of resonance dips. In practice, a thinner taper than the one used for this experiment helps to selectively excite the water-modes, however this experiment shows the

different directions and magnitudes that different modes shift based upon their thermo-optic coefficients.




Figure 4-13. Error signal shown for multiple resonances. Optically pumping an analyte on a microbubble (bottom panel) and with for the pump beam off (top panel). Dotted lines guide the eye to show that the resonances shift in different directions and with different magnitudes, as qualitatively shown by the arrows at the top. The full x-axis span is 180 pm in wavelength.

Three modes are indicated by three different dotted lines, passing through the top panel (pump beam off), and the bottom panel (pump beam on). Arrows indicate direction and qualitative magnitude of resonance shift. It is evident that the leftmost resonance (green dotted line) shifts the most. The significant blue-shift experienced by this resonance is indicated by the green arrow. In contrast, the rightmost resonance (red dotted line), shifts in the opposite direction (red shift), and with a much smaller magnitude than the green shift. The slopes of the middle (pink dotted line) and right modes are also much steeper than the slope of the left mode, consistent with the expectation that glass-modes will have higher quality factors (corresponding to a steeper slope in the error signal) than water modes.

4.17. Supporting Information: Modeling of LPB and TPB During Etching

To further understand the shape of the etching time traces, we modeled the spectral changes of the TPB and LPB for a nanorod being etched. This was done for a variety of etching schemes, from tip-only etching to side-only etching. For each set of etching/reaction parameters, the reaction was modeled separately for being driven by the LPB pump and the TPB pump. A “threshold” was set at 0.5 for all of the runs here described, so that the code would stop running at the point at which ratio of current signal to initial signal had hit 0.5. This threshold was **not** used to influence the reaction rate. For a given etching scheme, both length and width were prescribed “rate constants” (k) in nm/s of etching. The following rate equations were then used for each etching step to calculate the change in length and width of the nanorod:

Equation 4-3

$$r_{\text{length}} = k_{\text{length}} * \sigma_{\text{pump}}^n$$

Equation 4-4

$$r_{\text{width}} = k_{\text{width}} * \sigma_{\text{pump}}^n$$

where r is the rate of etching, k is the etching rate constant, σ is the relative signal of the plasmon band being used to drive the reaction at the pumping wavelength, and n is the reaction order with respect to “hot electron concentration” (which is the photothermal signal of the current iteration over the signal from the first iteration). Once the rates were calculated for a given iteration, the new size parameters were calculated for the nanorod. Then, the spectra for both the LPB and TPB were calculated. The assumptions that were made in constructing this model are listed below (and also listed at the end of the Matlab code, which is included in this Supporting Information).

Assumptions:

1. For tip-only etching, the LPB max amplitude linearly follows a trend starting at amplitude “1” at 650nm, and ending at amplitude “0.1” at 520nm. This results from taking an initial ratio of

TPB/LPB of 0.4 from bulk Nanopartz data, and dividing the ratio by two, because our experiment only probes one axis of the TPB, making the ratio 0.2 to start with. Then, when the LPB has met with the TPB in central wavelength (due to etching), the volume should have halved (for tip-only etching), meaning that TPB=LPB=0.1. This linear trend is for bounding the maximum LPB amplitude; side-etching will result in quicker decrease in LPB amplitude.

2. The LPB central wavelength linearly follows a trend starting with aspect ratio=2 being at 650nm, and at aspect ratio=1 being at 520nm.
3. The TPB max amplitude linearly follows a trend that's directly proportional to the current volume over the initial volume of the nanorod (approximated as a rectangular prism).
4. The TPB central wavelength does not shift.
5. Interband transitions can be ignored, as they should scale the same as the TPB with volume, and thus have an identical impact on rate of reaction.
6. Transverse etching, if it is happening, is equivalent in both transverse directions.
7. The two pump channels used to extract photothermal signal are agnostic to the plasmon band which they're not probing, as the two polarizations for plasmon band excitation are orthogonal.
8. The effects of scattering on the bulk spectra data used to make these assumptions are negligible, and thus the above conclusions may be drawn.
9. Linewidths are constant for both plasmon bands (an approximation).
10. "Hot electron concentration" is proportional to the photothermal signal (i.e. the relative absorption cross-section of the nanorod at the pumping wavelength).

The "rate constants" for length and width tested were: $(k_{\text{Length}}, k_{\text{Width}}) = (1,0), (1,0.1), (1,0.25), (1,0.5), (1,0.75), (1,0.9), (1,1), (0.9,1), (0.75,1), (0.5,1), (0.25,1), (0.1,1), (0,1)$. Shown in **Figure 4-14** are the modeled spectra evolving over time for four examples from this set, for the "TPB-driven" model (meaning that the 532 nm laser is being used to drive the reaction). These examples represent four different

regimes of etching, (1) tip-only etching, (2) tip-preferred etching, with some side-etching, (3) pure isotropic etching (identical rate constants for length and width), and (4) side-only etching.

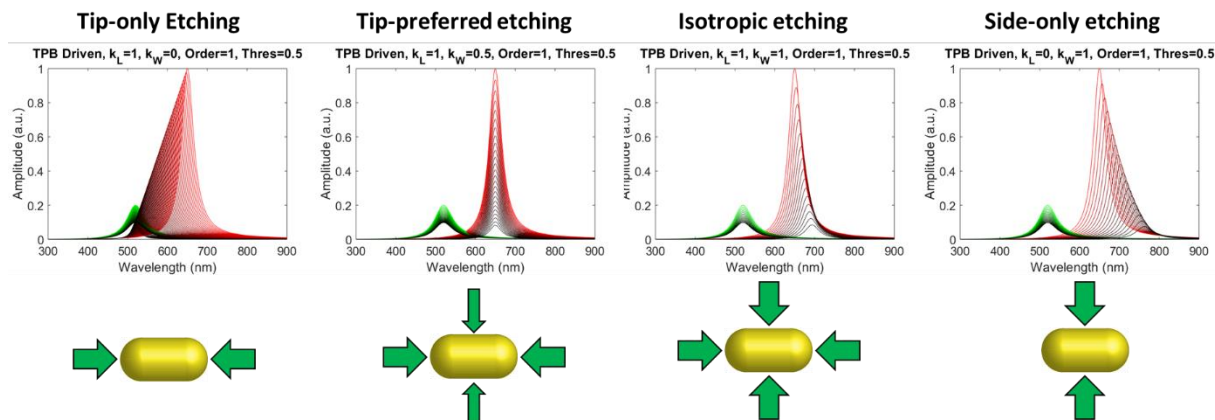


Figure 4-14. Modeled spectra for the TPB and LPB. For four different etching regimes, with TPB-driven etching, a reaction order of 1 with respect to hot electron concentration, and a threshold of 0.5.

The model also outputs a few other parameters. These are shown in **Figure 4-15** for the tip-only, reaction order=1 scenario for both TPB-driven and LPB-driven cases. The leftmost column shows the spectral evolutions. In the second column from the left are the length, width, volume, and aspect ratio as etching progresses. In the third column from the left are the *rates* of reaction over time for length and width etching (in nm/s). In the rightmost column are the **predicted photothermal signals**. For the top row (LPB-driven), the signal is only shown for the 635 nm pump. Likewise, for the TPB-driven reaction, only the 532 nm signal is shown. The possibility of probing with other wavelengths (ex. drive reaction at high LPB power but also measure signal at low TPB power) is discussed both in the main text, and further below. For a given threshold, it is expected that the TPB will have many more “steps” calculated in the code than the LPB. This is true because (1) the LPB drops in signal more quickly than the TPB as volume changes, and (2) the LPB can shift, whereas the TPB cannot, causing a significantly faster drop-off in LPB signal, triggering the threshold condition sooner.




Figure 4-15. For TPB and LPB driven, tip-only etching. With reaction # order 1 and threshold 0.5: modeled spectra (leftmost), nanorod dimensions (second from left), reaction rates (third from left), and predicted photothermal signal (rightmost).

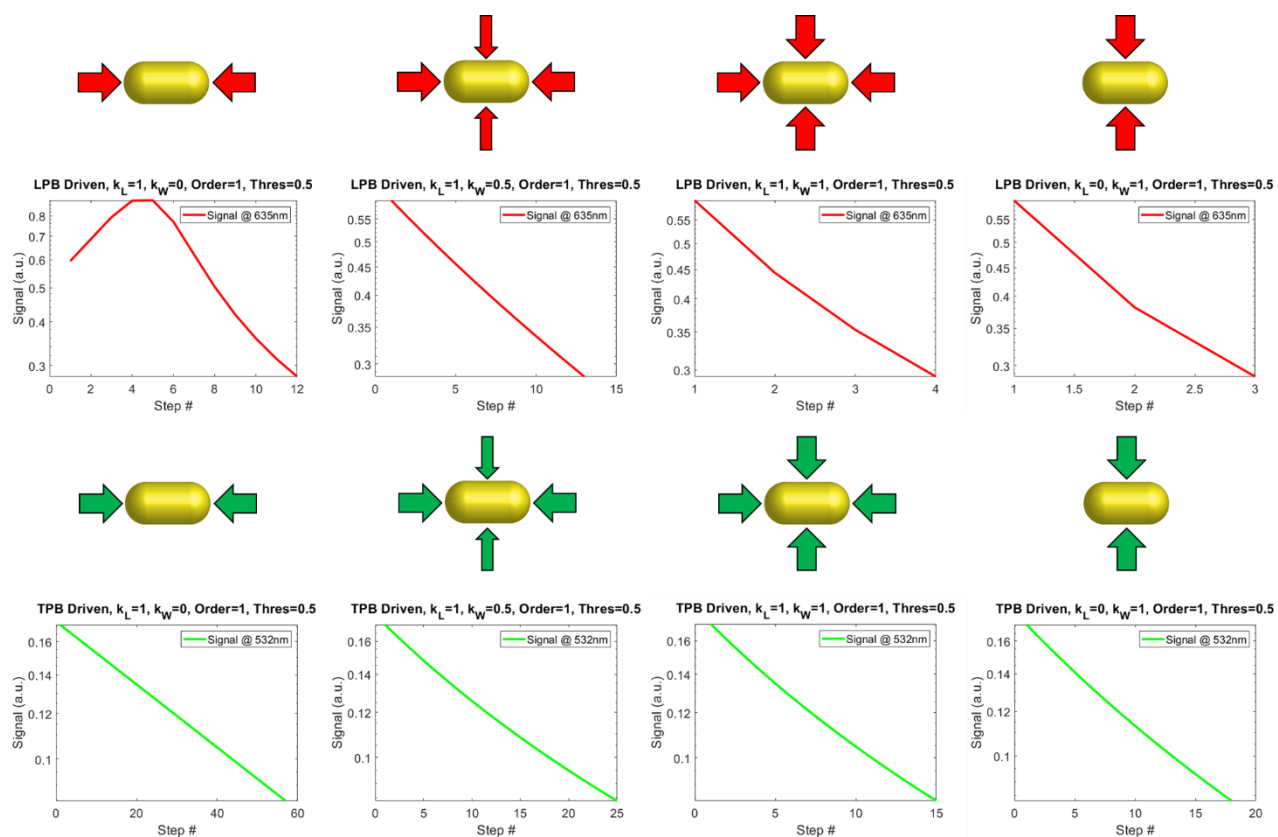


Figure 4-16. LPB and TPB driven etching. For LPB (top) and TPB (bottom) driven etching, reaction order 1 and threshold 0.5, the predicted photothermal signal is shown for different etching regimes.

The predicted photothermal signal over the course of etching is shown for the four different etching regimes from **Figure 4-16**, below for both TPB and LPB-driven reactions, all for reaction order=1 .

Reaction orders of 1, 2, 3, 4, and 5 (with respect to “hot electron concentration”) were modeled (Figure 4-17). The first four orders are shown for both an LPB and a TPB driven reaction. As expected, higher-order kinetic models show faster reactivity at the beginning, but slower reactivity later on, because the reaction rate is more dependent upon the “hot electron concentration”.

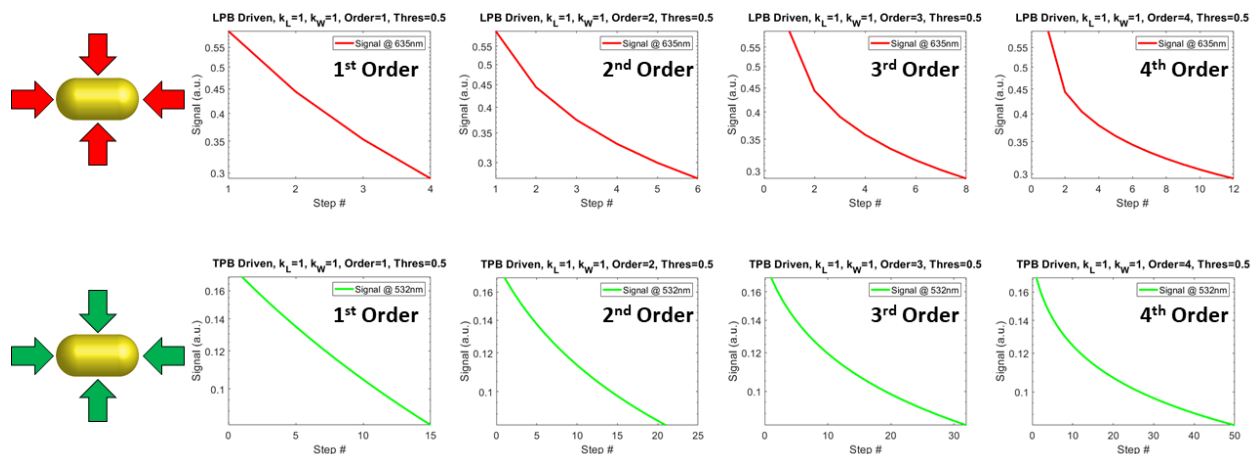


Figure 4-17. The isotropic etching scenario for both TPB and LPB driven reactions. The predicted photothermal signal is shown for different orders of reaction with respect to hot electron concentration.

In the tip-only etching, it is predicted that an *increase* in photothermal signal will initially be observed, as the plasmon shifts toward the pumping wavelength. Notably, we have never observed this significant signal increase experimentally. This could be due to either (1) the plasmon being red-shifted enough on the surface that the signal increase from shifting does not increase fast enough to overcome the signal drop from etching, or (2) the etching not being tip-only. Even in tip-preferred regime, it may be that the signal drops quickly enough to mask the fact that the LPB approaches the pump beam wavelength (blue-shifts).

To better deduce the morphology changes from the photothermal signal, it would be advantageous to do multi-beam experiments, where the absorption cross-section of the nanorod is examined at multiple wavelengths. The reaction could still be driven at relatively high power with the desired wavelength, while the cross-section at different wavelengths could be acquired at lower powers, below the threshold for reactivity, using a longer lock-in time constant. This has not yet been realized experimentally, but is easy to do theoretically. In fact, the code here described automatically generates the signal at both pumping wavelengths for each plot of photothermal signal, although this has been removed from the above figures for clarity, and to better align with our experimental results. In **Figure 4-18**, the signal at *both* 532 nm and 635 nm is shown for first-order, isotropic etching for both TPB and LPB driven reactions.




Figure 4-18. The isotropic etching scenario for both TPB and LPB driven reactions. The predicted photothermal signal is shown for first-order kinetics, at both 532 nm and 635 nm.

It is evident from the *TPB-driven, 635 nm signal* that for TPB-driven reactions, 635 nm signal curves can have a shape other than concave-up. Nonetheless, such shapes were only theoretically observed at 635 nm, not at 532 nm, where we experimentally observe concave-down and sigmoidal reaction traces. This is not very surprising, given that the TPB is not allowed to shift during etching in this model, a fairly accurate approximation given the bulk data from Nanopartz, for which the TPB only ranges 510-520 nm for nanorods of different geometries.¹⁴⁵ Additionally, more significant “threshold” behavior is observed

experimentally than our model can recapitulate; the model will eventually react the nanorod to nothing, while experimentally the reaction slows down to almost zero rate once enough reaction has occurred for a given pumping power. Thus, we speculate that there is a non-linear reaction mechanism that evolves alongside nanorod morphology. We also speculate that there may be significant mechanistic differences between the LPB-driven and TPB-driven reactions. This is discussed further in the main text. Future experimental studies with multiple colors, as described above, may also allow for fitting the concave-up shapes of the LPB reactions to extract more kinetic information.

4.18. Supporting Information: Thermal Expansion

As in our previous report,¹⁰ we use the values for fused silica's thermo-optic coefficient and linear thermal expansion coefficient¹⁴⁶ ($dn/dT=8.6*10^{-6} K^{-1}$, $\alpha_L=5.5*10^{-7} K^{-1}$) to estimate that the contribution of thermal expansion should **at most** only be about 6% of the total signal ($\alpha_L/(dn/dT+ \alpha_L)$), with the other 94% coming from the change in refractive index. This is for an all-glass mode. When considering a water-mode, the relative contribution from thermal expansion should be far less, because the dn/dT coefficient goes up significantly in magnitude. We do not consider the expansion coefficient of water in this analysis, as the pressure waves due to water expansion should travel much faster than the heat transfer. This means that water can push its way out of the microbubble as it expands, rather than contributing to resonator expansion. Because the contribution of thermal expansion to the signal is far less than that of refractive index change, we do not correct for the impact of thermal expansion in this report.

4.19. Supporting Information: Diagram for Mode Indices

The mode-number indices used in the main text are radial, polar, and azimuthal. In **Figure 4-2b**, only the first two types of modes are shown. The azimuthal mode number, not shown in **Figure 4-2b**, indexes the number of electric-field nodes around the waist of the microbubble. A cartoon illustrating these three mode numbers for a spherical system is shown in **Figure 4-19**. A cylindrical model is also shown, where the "polar" mode number has been replaced by the "axial" mode number. In some bottle microresonator

literature, the cylindrical symmetry terminology is used. However, we've used the nomenclature of spherical symmetry in this report, which we think better captures the geometry of the microbubbles, modeled as prolate-spheroids.




Figure 4-19. Mode number nomenclature. The nomenclature used to describe mode numbers (the eigenvalues satisfying Maxwell's equations for constructive propagation of light through the resonator). *Left:* Spherical symmetry nomenclature. *Right:* Cylindrical symmetry nomenclature.

4.20. Supporting Information: Additional Single AuNR Etching Data

Additional data sets for single nanorod reactions are shown below. Some data traces show an artifact of rising signal in the first few degrees, due to the data collection starting before the liquid crystal had returned completely to the retardation necessary for zero degrees polarization. *LPB-Driven reactions* are shown in **Figure 4-20**. *TPB-Driven reactions* are shown in **Figure 4-21**.

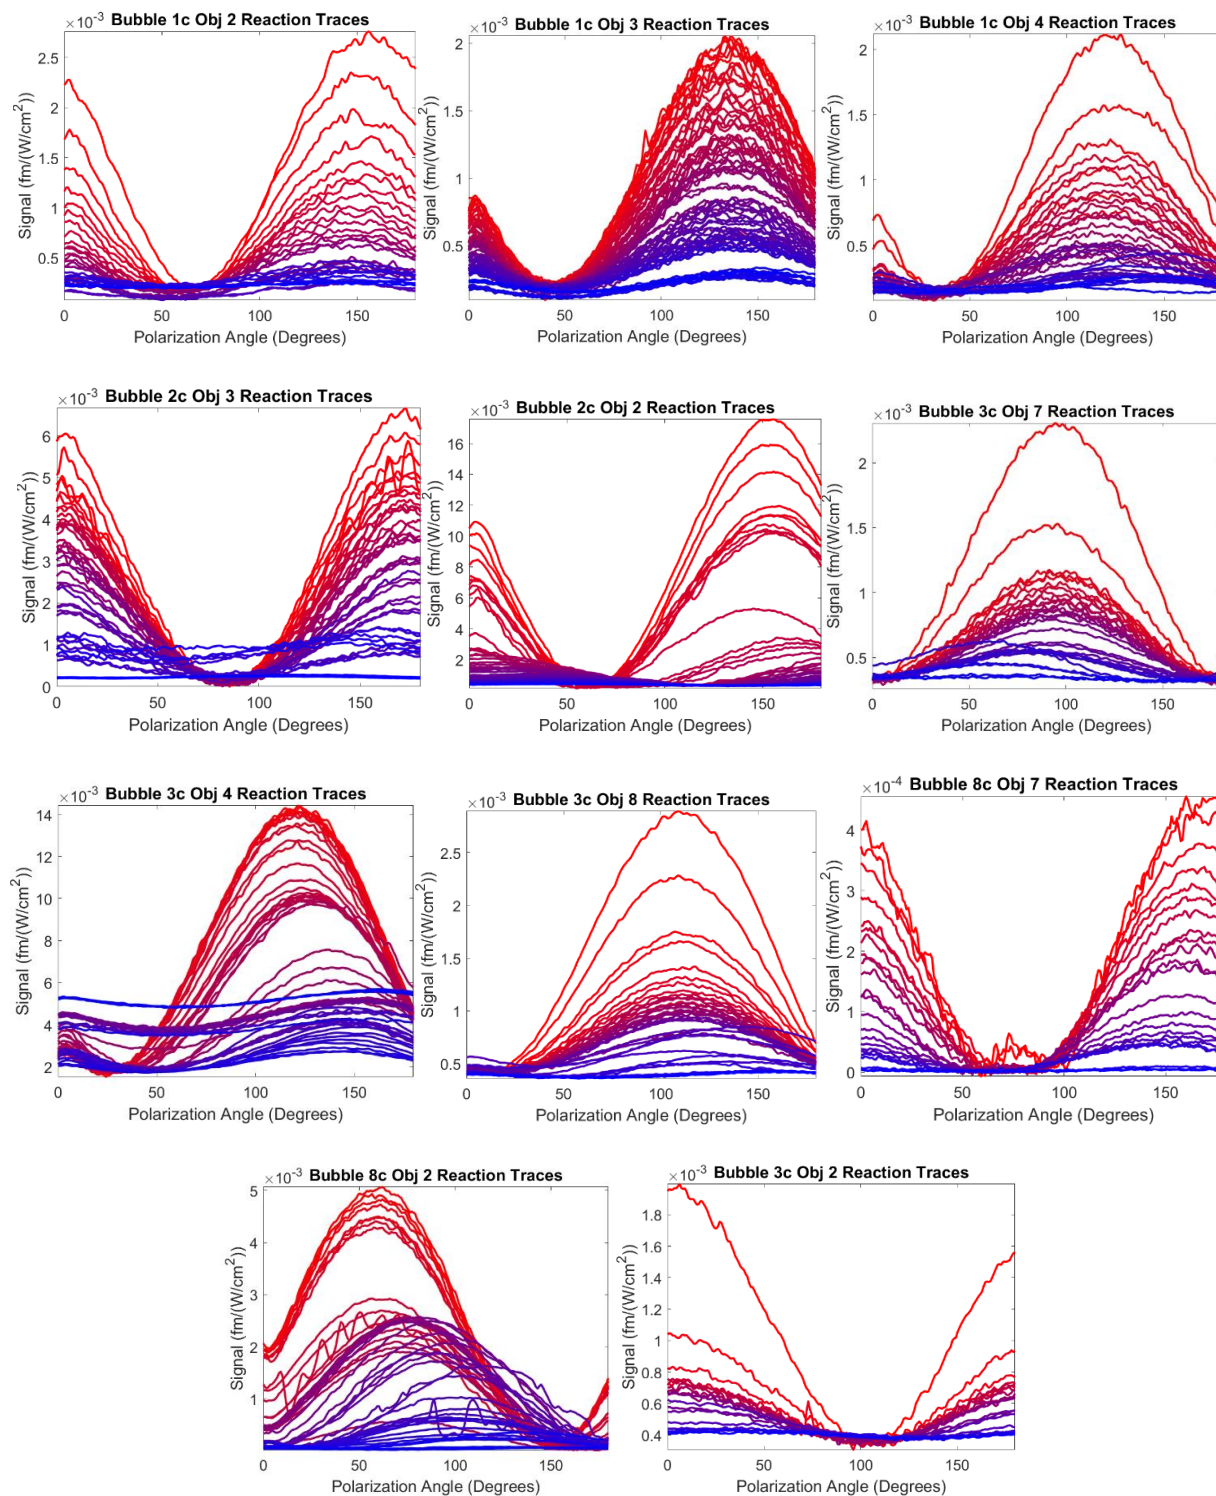


Figure 4-20. Additional examples of LPB-driven reactions.




Figure 4-21. TPB-driven reactions. All three TPB-driven reactions (including the two shown in main text (**Figure 4-5**)). Two LPB-driven reactions from the same bubble are in the dataset in **Figure 4-20**.

4.21. Supporting Information: Additional Rotation Data

Additional data sets for single nanorod rotations are shown in **Figure 4-22**. These confirm reproducibility in probing the local environment of the single nanorods (see main text for details).

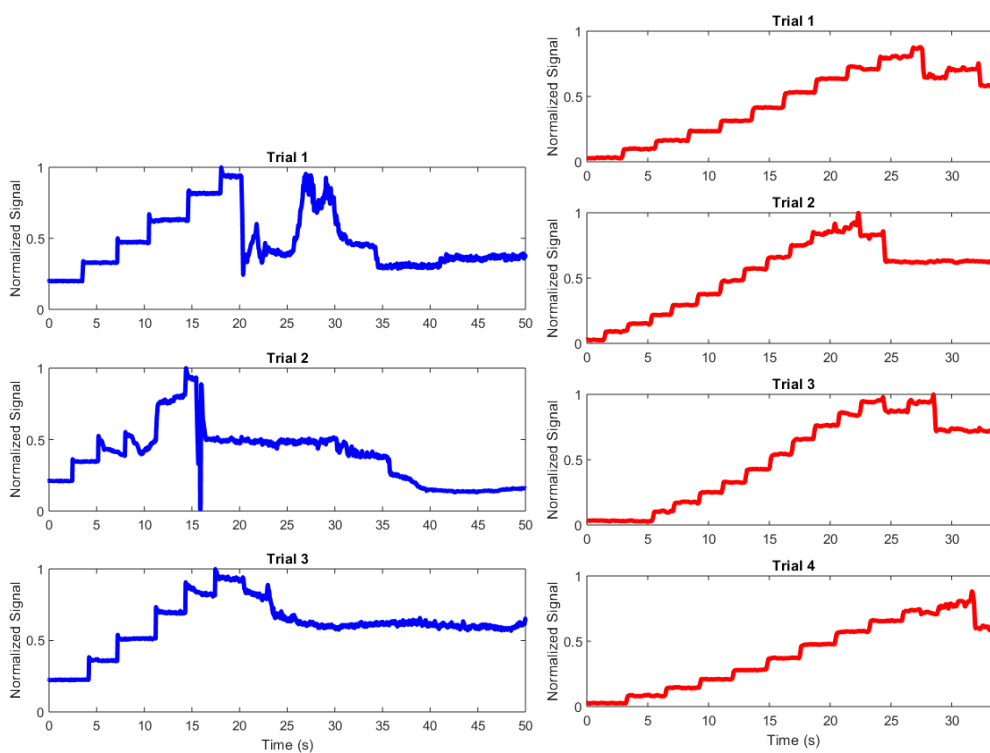


Figure 4-22. Reproducible staircase plots. The left side (blue staircases) is for the left staircase in **Figure 7c** of the main text. The right side (red staircases) is for the right staircase in **Figure 7C** of the main text.

4.22. Supporting Information: Bulk Reaction Results

To verify the similarity of AuNR etching with FeCl_3 in our own hands, we conducted a variety of bulk experiments. Nanopartz nanorods of dimensions 80 nm x 40 nm (aspect ratio=2), from a different lot but of very similar parameters to the single nanoparticle experiments, were reacted with ferric chloride, monitored by UV-Vis spectrophotometry. These experiments were run at high oxidant concentration (100 mM), low oxidant concentration (20 mM), low oxidant concentration with increased CTAB, and low oxidant concentration with added NaCl. For all reactions, $[\text{Au}]=166\mu\text{M}$, and $\text{pH}\sim 1.3$ (dilute HCl). All nanorod solutions were made by centrifuging a stock solution at 5000 RPM for 10 minutes, pulling off the supernatant, and resuspending the nanorods in a solution of desired CTAB concentration and pH. The results of these experiments are shown in **Figure 4-23**.

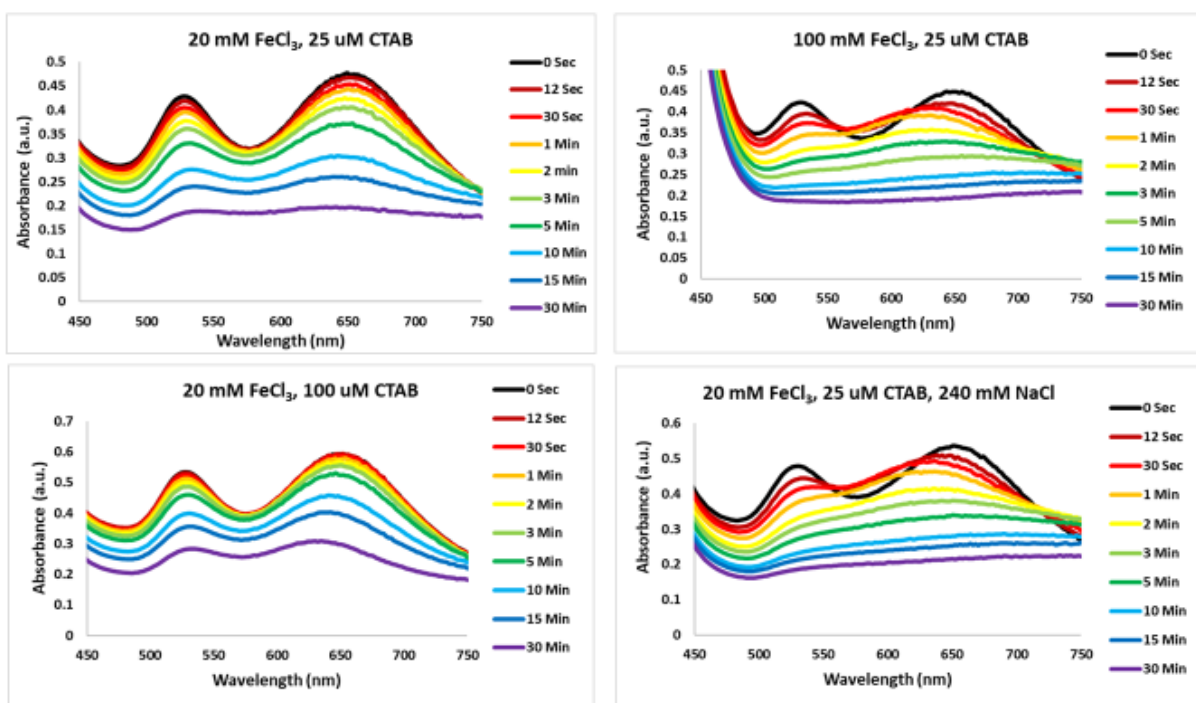


Figure 4-23 Bulk reactions with large nanorods. monitored by UV-Visible spectrophotometry.

It is clear from above that the dependence on ferric chloride is important: the reaction slows significantly when ferric chloride is lowered from 100 mM to 20 mM. In our experiments, for the low-oxidant concentrations, CTAB played a less significant role in reactivity than in previous reports,¹¹¹

although those reports were in a different concentration regime. Although our conditions are somewhat different, our results on halide dependency concur with Zou and coworker's conclusions,¹¹¹ as adding salt revived much of the reactivity seen at higher oxidant concentrations. Thus, our nanorods likely undergo similar bulk reactions to those in previous reports. However, it is clear from the above data that the reactions do not effect prominent blue-shifting of the LPB, although they may effect some. Also, the faster reactions show an interesting feature where the two plasmons quickly merge, although this may be due to aggregation occurring from nanorod destabilization as reactions progress, as these solutions are at well below the CMC of CTAB.

An additional factor worth noting is that surface curvature may affect reactivity, as some etching reagents have been hypothesized to have surface-curvature-dependent etching mechanisms, such as Au(III) etching.¹¹⁸ With an interest in using smaller nanorods in future single nanoparticle experiments, we also sought to understand the reactivity of such smaller nanorods. In Figure 4-24 we show two examples of smaller-nanorod reactivity. Notably, these nanorods are of a different aspect ratio (3.8), and much smaller dimensions (10 nm x 38 nm) than the larger nanorods. The two reactions shown are for high oxidant concentration (200 mM) at two different CTAB concentrations.

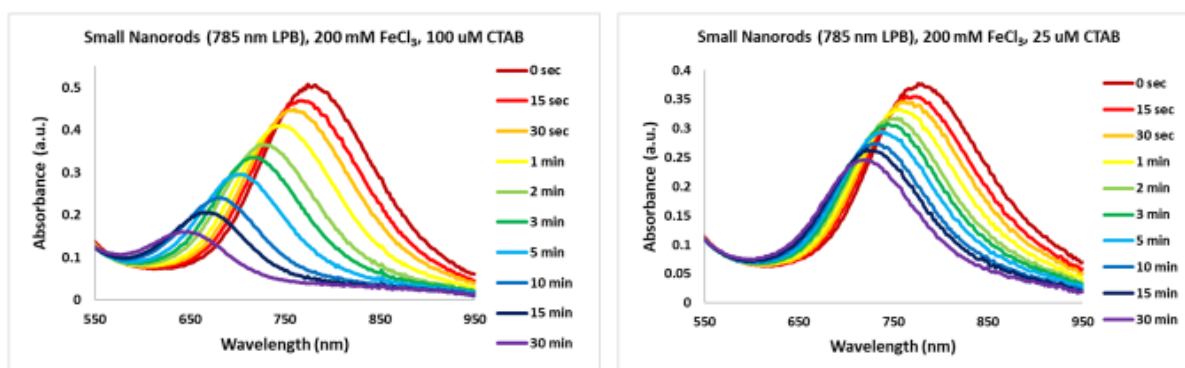


Figure 4-24 Bulk reactions with small nanorods. Monitored by UV-Visible spectrophotometry.

Interestingly, these results more closely match those previously reported,¹¹¹ although these nanorods are smaller and the CTAB concentrations are lower. The smaller nanorods clearly show a blue-shifting in their LPB as etching occurs, indicative of a spatially-selective, tip-preferential etching mechanism that may be reliant on surface-curvature effects. Notably, the effect of CTAB appears more significant in this case. The 100 μM example shown in Figure 4-24 proceeds much further in 30 minutes than the 25 μM CTAB case. Similar trends as discussed above were also observed for small nanorods, such as recovered reactivity at lower oxidant concentrations upon the addition of sodium chloride (data not shown).

Overall, these bulk reactions reinforce previous reports on the reactivity of AuNRs to FeCl_3 . As noted in the main text, however, the photo-activated mechanism observed in the single-particle reactions, which are run at ferric chloride concentrations approximately two orders of magnitude below the bulk studies, may be very different in its dependencies on reactants than in bulk reactions. Nonetheless, the bulk reactions are still worthwhile to consider when speculating mechanisms of nanorod etching.

4.23. Supporting Information: Brief Discussion on Extracting Reaction Kinetics

As shown above in the plasmon band modeling, different etching mechanisms will alter the LPB in different ways. Because the etching reported here is photo-activated, the etching rate changes as the reaction progresses. Therefore, distinguishing between the LPB decreasing in amplitude and the LPB also shifting in wavelength space is difficult, and detailed kinetic parameter extraction isn't possible without further information. However, it's worth noting that although the three reactions in **Figure 4-4D** were taken at ferric chloride concentrations spanning almost an order of magnitude, the time scales of reaction vary by much less than an order of magnitude. This could possibly be due to saturation of oxidant on the nanorod surface at these concentrations. Future kinetic studies are required to better understand

such reagent concentration dependencies. The LPB/TPB modeling described earlier in this Supporting Information may also help to elucidate the morphology changes in future etching experiments.

4.24. Supporting Information: Background on AuNR etchants

Although the photophysics of AuNRs is well understood as a function of geometry,¹⁰⁹ morphological control is not always as accurate and precise as desired, resulting in necessary post-synthetic modifications for some applications. Since 2002, when Jana and coworkers observed anisotropic etching of gold spheroids in both cyanide and persulfate solutions,¹¹⁰ a variety of reagents for etching AuNRs has been reported, including hydrogen peroxide,¹⁴⁷⁻¹⁴⁸ gold(III),¹¹⁸ silver ions with potassium tetrachloroplatinate(II),¹⁴⁹ potassium permanganate,¹⁵⁰ cysteamine,¹⁵¹ and aqueous O₂ under acidic conditions.¹⁵² AuNR etching has also been widely employed for sensing, where small changes in nanorod aspect ratio due to etching are detected *via* colorimetric readout in bulk solution, even with the unaided eye. Analytes reported to selectively etch or accelerate the etching of AuNRs in recent years for sensing include mercury,¹⁵³ chromium(VI),¹⁵⁴ lead(II),¹⁵⁵ nitrite,¹⁵⁶ copper(II),¹⁵⁷⁻¹⁵⁸ cobalt(II),¹⁵⁹ molybdate,¹⁶⁰ iron(III),¹⁶¹ and iodide.¹⁶² Furthermore, AuNR etching is useful in biologically-relevant schemes, such as sensing glucose,¹⁶³⁻¹⁶⁴ detecting telomerase activity,¹⁶⁵ performing plasmonic ELISAs for detecting alkaline phosphatase,¹⁶⁶ detecting organophosphorus pesticides,¹⁶⁷ and visualizing superoxide dismutase in serum.¹⁶⁸ Mechanistic studies on the roles of halides, surfactants, surface curvature, hydroxyl radicals, and dissolved oxygen in the etching of AuNRs have been published,¹⁶⁹⁻¹⁷² although there are still open mechanistic questions within the field.

4.25. Supporting Information: Impervious nanorods

Reactions: As mentioned in the main text, we occasionally observed “impervious” nanorods, that were visible photothermally using both the 532 nm and 635 nm pump beams, and had $M \sim 1$, yet didn’t undergo photo-activated etching. Sometimes, this was observed for sets of a few nanorods deposited in the same bubble. Therefore, we hypothesize that this lack of reactivity may be due to specific deposition conditions

that led to unreactive nanorods, such as nanorods with too much CTAB coating, or perhaps small aggregates of well-ordered, parallel nanorods, which could still give a depth-of-modulation near unity. Further studies would be necessary to support either of these hypotheses. Interestingly, ferric chloride has also been shown to stabilize gold particles, although this was a very different context, in which ablating a gold surface in a solution of ferric chloride formed stable, iron-complex coated gold nanoparticles.¹⁷³ However, it remains a possibility that a mechanism exists whereby some AuNRs are impervious to reaction in our system.

Rotations: The two staircases shown in **Figure 4-7** of the main text were from “impervious” nanorods as well. The left staircase was for an impervious nanorod that had sat in solution for multiple days without reacting when pumped with either the 635 nm or 532 nm beams. The right staircase in **Figure 4-7** is from a single nanorod in a water-filled microbubble. However, this nanorod didn’t react at a later date when a reaction was attempted. Given the strong depth of modulation of these objects, and their visibility using both the 635nm and 532 nm pump beams, we expect that both objects were gold nanorods. As above, we hypothesize that these objects were chemically inert due to either being CTAB-coated and incapable of etching under these conditions, or actually being small sets parallel nanorods. We have also observed rotation of AuNRs towards perpendicular to our pumping polarization in a number of etching reactions, supporting our statement of control over nanorod rotation. Future studies will allow for insight into why some nanorods do not react, and also allow for insights into the relationship between chemical etching and proclivity for rotation.

4.26. Supporting Information: Discussion of Role of CTAB in Etching Mechanism

For many reported etching reagents, a simple redox analysis predicts that AuNR etching should not be spontaneous. Although the oxidation potential for gold nanorods can be significantly lower than that of bulk gold,¹⁷⁴ it has nonetheless been demonstrated that additional reagents are typically required to

modulate reactivity. We go to significant lengths (in both the **Main Text** and this **Supporting Information**) to consider CTAB's role in the reported chemistry, largely due to the literature precedence for CTAB playing a critical role in AuNR etching. For example, CTAB at concentrations above the critical micelle concentration (CMC) facilitates Au(III) etching of AuNRs, presumably by complexing with Au(III) and preferentially aggregating near AuNR tips.¹¹⁸ CTAB dependence was also observed in the first report of FeCl₃ etching of AuNRs, including at CTAB concentrations well below the CMC, and was hypothesized to be due to the complexation of bromide from CTAB and chloride from FeCl₃ with gold, lowering its oxidation potential.¹¹¹ In the single AuNR reactions in this report, oxidant concentration is two orders of magnitude lower than bulk studies showing light-independent etching (**Supporting Information**). This means that any halide dependence from CTAB or FeCl₃ is negligible in comparison to the much larger chloride concentration from the dilute hydrochloric acid. In conjunction with the studies discussed below, it is evident that CTAB plays an important role beyond halide donation.

In the bulk study by Toste and co-workers,¹¹² etching was multiple orders of magnitude faster when driven with interband excitation than with LPB excitation, a phenomenon attributed to differences in the generated hot electron distributions. This chemistry was accomplished with CTAB concentrations above the CMC. Toste and coworkers found that the larger population of lower-energy hot electrons generated from interband transitions drove ferric chloride etching of AuNRs orders of magnitude more efficiently than those generated from the LPB.¹¹² However, given that the TPB-driven reactions in this report proceed at a rate that is comparable to the rate of the LPB-driven reactions (**Figure 4-5**), we conclude that the etching mechanism observed here is fundamentally different than that observed by Toste.¹¹² In contrast, in the single AuNR etching reported by Khatua and co-workers,⁷⁸ CTAB was thoroughly stripped from the surface of the nanorods using ozone. Using only a green laser to drive etching, which excites both the TPB and interband transitions, the authors then exhibited bidirectional

control over the AuNR aspect ratio, attributed to a balance of thermally driven and hot electron driven mechanisms.

In our work, CTAB was removed from solution and reduced at the surface of deposited nanorods by flowing water following deposition (**Methods**). However, due to the gentle nature of this approach, we hypothesize that our conditions are at an intermediate surface CTAB concentration between reports from Khatua and Toste. We hypothesize that this intermediate CTAB concentration results in a hot electron etching mechanism that is different than in the above reports, as detailed in the **Main Text**.

4.27. Supporting Information: Estimation of CTAB Remaining on AuNRs after Deposition

The nature of the CTAB ligand density after nanoparticle washing is relevant to the etching mechanism. While CTAB adsorption/desorption kinetics have been studied on powdered activated carbon¹⁷⁵ and silica¹⁷⁶ surfaces, we are unaware of such kinetics studies using AuNPs. Chemical removal of CTAB from AuNRs with sodium borohydride has also been studied using Fourier-transform infrared spectroscopy (FTIR), surface-enhanced Raman spectroscopy (SERS), and x-ray photoelectron spectroscopy (XPS),¹³⁴ however it is difficult to gain inside into the desorption of CTAB in pure water from this data. SERS has also been employed for studying the structural transitions of the CTAB surfactant bilayer on AuNRs.¹⁷⁷ However, this report only went down to concentrations of 0.2 mM CTAB, still above the necessary concentrations for understanding what occurs during the washing step of our deposition. Additionally, this method likely preferentially probes the tips of the AuNRs (where the electric field is most concentrated by LPB), make it hard to extrapolate the data to the entire nanostructure. Adsorption kinetics of other ligands, including thiols¹⁷⁸⁻¹⁷⁹ and 1,2-dimyristoyl-sn-glycero-3-phosphocholine (DMPC),¹⁸⁰ have been examined using AuNPs, but the differences in chemical functionality precludes direct extrapolation to CTAB kinetics. Alternatively, there have been at least two reports with binding

isotherm data for CTAB adsorption on a flat gold surface.¹⁸¹⁻¹⁸² With no known AuNR data for CTAB at the necessary surfactant concentrations, we will use the data from these studies for the following discussion.

Both of these studies give isotherm data for adsorbed CTAB mass vs. concentration of CTAB. We extracted this data using WebPlotDigitizer,¹⁸³ and then fit the data points in Matlab to Equation 4-5 (equivalent to a Langmuir isotherm),

Equation 4-5

$$\Theta = \frac{K_a * C}{1 + K_a * C}$$

where, by using Θ , the fraction of surface sites occupied by adsorbate, and C , the concentration of adsorbate in solution, the adsorption equilibrium constant, K_a , may be extracted from the data. We only fit data points occurring at CTAB concentrations below the CMC, as this point also corresponds to the plateau observed in the isotherm data, presumably signifying surface-saturation of CTAB on the gold. We normalized the data so that the maximum value (at the start of the plateau/at the CMC) is “1”, making our dependent variable (Θ) fractional. The concentration (C) was converted to molarity as necessary.

Figure 4-25 shows data and fits from the two reports. **Figure 4-25** is the data from the report by Knag and coworkers,¹⁸² specifically from Figure 5 of their report. Data was extracted for two of their data sets, namely the pure CTAB, and the 3% hexanol added. The reason for this is because in their report, they suggest that their pure CTAB is adsorbing as aggregates, and the mixed system as a bilayer. Since we expect pure CTAB on AuNRs to adsorb as a bilayer,¹⁸⁴ we wanted to extract the parameters for each system to get a range of possible values. **Figure 4-25** and **Figure 4-25** show the fits and extracted K values for pure CTAB and the mixed system, respectively. **Figure 4-25** shows the data extracted from the report by Kawasaki and coworkers, specifically Figure 1b in their report.¹⁸¹ **Figure 4-25** shows the corresponding fit and extracted K for their data.

From the isotherms in the referenced studies, K_a values of approximately $5e3-1.5e4$ were obtained. Assuming complete coverage of a nanorod with CTAB prior to the rinse step (and approximating this coverage as spatially homogenous across the nanorod), a nanorod starts with $7.31e4$ CTAB molecules covering it. This was calculated by estimating two CTAB molecules per 55 \AA^2 (based on the assumption of a bilayer, and a number previously reported for area-per-molecule⁷), and dividing this from the calculated nanorod surface area of $2.01e6 \text{ \AA}^2$ (from a model of a cylinder 80 nm long and 40 nm wide). To estimate the fraction of CTAB coverage remaining after the rinse cycle, we ask an analogous question: **given the molecules needed to cover a nanorod, dissolved freely in the exchange solution volume, what is the equilibrium coverage of CTAB molecules on the AuNR?**

Using an estimated exchange volume of 0.1 mL, and the above calculated $7.31e4$ CTAB molecules, we arrive at a solution concentration of 1.2 fM. Using a value of K_a of 10^4 (middle of the calculated values), and the calculated concentration of 1.2 fM, we arrive at a fractional coverage estimate of $1.2e-11$. Given the expected $7.31e4$ CTAB molecules needed to cover the surface of the nanorod, this means that we have estimated that **zero CTAB molecules** will remain adsorbed to the nanorod, which corresponds to the conclusion that **all CTAB will desorb from the nanorod during the rinse step.**

However, this conclusion that CTAB should be entirely removed from the surface from a simple water wash is not in-line with the observations by others that complete removal of CTAB from AuNRs can be difficult to accomplish¹³³ In particular, it has been noted that residual CTAB is often left on the surface if repeatedly washed, and aggregation of AuNRs in solution will occur before this residual CTAB can be fully removed.¹³⁴ Of course, our extracted values are for CTAB on a flat gold surface, not a gold nanorod's surface. The curvature of a gold nanorod is expected to significantly impact ligand dynamics and the equilibrium state. Additionally, this analysis used a Langmuir model to estimate the equilibrium constant for binding. Critically, this model (1) neglects both direct and indirect adsorbate-adsorbate interactions, (2) assumes that all surface sites are created equal, and (3) is meant for modeling monolayer

adsorption. Given that CTAB certainly has intermolecular interactions among itself (ex. hydrophobic interactions), the nanorod likely doesn't adsorb equally on all surfaces, and the CTAB is thought to form a bilayer,¹⁸⁴ the Langmuir model likely does not fully capture the CTAB adsorption/desorption behavior on AuNRs. It also may be that flowing of 0.1 mL of solution through the microbubble (which typically takes 1 hour or less) cannot be equivalently modeled by a solution surrounding the nanorod of equivalent volume if desorption occurs too slowly to equilibrate with the solution as it passes through.

Other isotherm models might be used to account for some of these factors. However, such models would require additional thermodynamic data, and/or more complex fitting, than is attainable with the extracted data. Therefore, in-depth studies of the CTAB interactions with AuNRs are needed in the future. Very recently, quantitative imaging of CTAB on single AuNRs was demonstrated.¹⁸⁵ The study used electron energy loss spectroscopy in a scanning transmission electron microscope (STEM-EELS) to accomplish this imaging. Although this study did not report dependencies on CTAB concentration, such studies might be accomplished in the future to obtain desired thermodynamic parameters for CTAB adsorption/desorption of AuNRs. Notably, the STEM-EELS study reported heterogeneity in CTAB coverage between AuNRs, and even across a single AuNR, indicating the need for single-particle inspection, and nanoscopic spatial resolution.

Other factors that have been shown to influence ligand structure on AuNRs include the addition of a salt to solution¹⁸⁶ and plasmonic heating.¹⁸⁷ Thus, the addition of etchant (i.e. a salt), and any plasmonic heating present would also have to be considered for a rigorous model of the ligand structure on the AuNRs in our report. Theoretical work has also been done to understand the adsorption and exchange kinetics of CTAB and PEG on different faces of AuNRs,¹⁸⁸ and future work in this area may fortify understanding of adsorbate-nanoparticle interactions. Overall, understanding both the static structure and dynamic behavior of ligands on AuNRs, as well as their impact on etching, is worth future research that is outside the scope of this report.

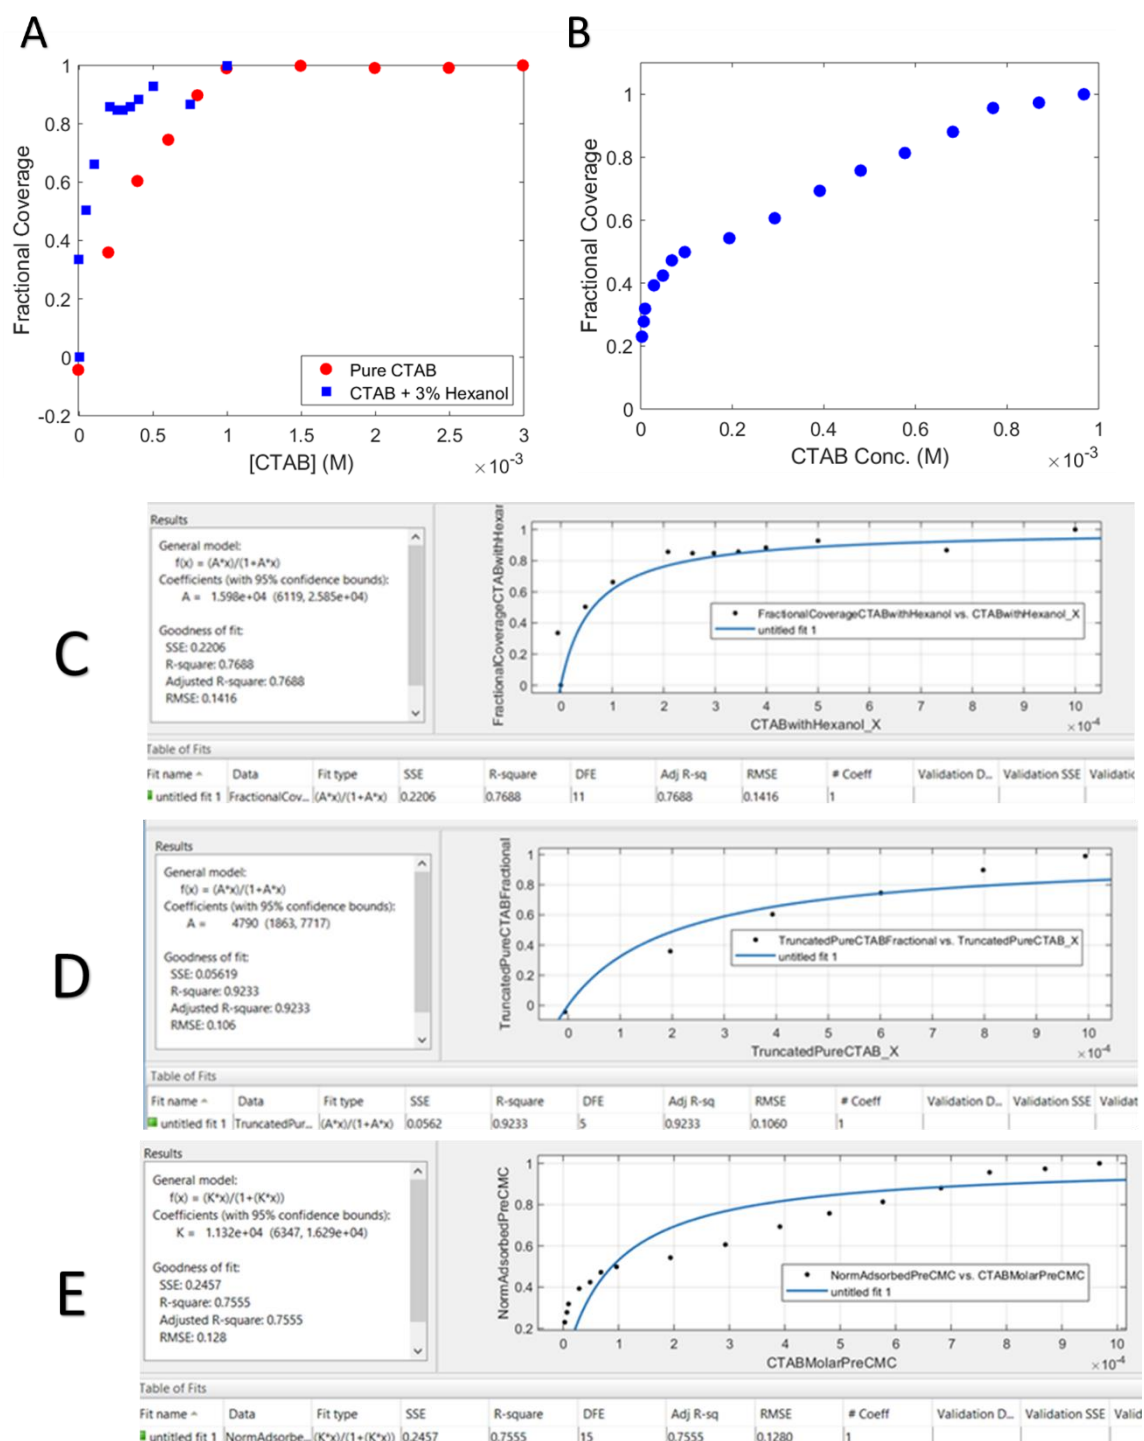


Figure 4-25. Langmuir isotherm data for CTAB. (A) Data extracted from report by Knag and coworkers. (B) Data extracted from the report by Kawasaki and coworkers. (C) Matlab fit for data from Figure S18A, “pure CTAB”. (D) Matlab fit for data from Figure S18A, “CTAB + 3% Hexanol”. (E) Matlab fit for data from Figure S18B.

4.28. Supporting Information: Estimation of Nanorod Temperature Increase

To calculate the anticipated temperature increase of the nanorods undergoing reaction in this report, we use an equation provided in the supporting information of a report by Thambi and coworkers,⁷⁸ based on work by Yorulmaz and coworkers,¹⁸⁹ where ΔT is the change in temperature, σ_{abs} is the absorption cross-section of a single AuNR, I_{exc} is the pump flux, κ is the thermal conductivity of the medium, and R is the equivalent radius of a spherical nanoparticle as calculated from σ_{abs} .

Equation 4-6

$$\Delta T = \frac{\sigma_{\text{abs}} I_{\text{exc}}}{4\pi\kappa R}$$

Using the extinction coefficient (of the LPB) from the Nanopartz specification sheet of $\epsilon=1.9\text{E}10 \text{ cm}^{-1} \text{ M}^{-1}$, we calculate a single nanorod σ_{abs} of $3.16\text{E}-11 \text{ cm}^2$ (by multiplying ϵ by 1000, and dividing by Avogadro's number). Plugging this into **Equation 4-6**, and using an excitation power of 6.4 kW cm^{-2} (taken from the starting power of the reaction in **Figure 4-4a(iii)**), a thermal conductivity value of $6.145\text{E}-4 \text{ kW m}^{-1} \text{ K}^{-1}$ (we use a value for water near room temperature, since the value for water is lower than the thermal conductivity of glass, leading to us calculating a maximum bound on the thermal increase),¹⁹⁰ and a value for R of $3.17\text{E}-8 \text{ m}$ (calculated for a sphere with the above σ_{abs}), we arrive at a calculated temperature rise of just **0.83 K**, well below the threshold needed for thermal contributions to etching to be of the same order or higher than hot electron contributions.⁷⁸

4.29. Supporting Information: Theory of Nanorod Rotation

The optically-induced torque experienced by anisotropic plasmonic structures under illumination with linearly polarized light has been theoretically described by considering the interaction of the induced dipole of the LSPR and the incident electromagnetic (EM) field.¹³⁸ Equivalently, the torque can be described as a minimization of optical potential energy within an energy well. This potential energy well

results from unequal polarizability of orthogonal axes of the nanoparticle, making orientation in one direction more favorable than in the other direction.¹³⁸ However, the pure-dipole model assumes that the optical field is constant across the nanoparticle, which is an inaccurate assumption. Therefore, a more robust theory has since been developed, using the Maxwell stress tensor. This theory predicts that the torque changes sign when the excitation wavelength passing through the LSPR,¹⁹¹⁻¹⁹² manifesting in the nanorod aligning perpendicular to the linear pumping polarization on the blue side of the LSPR, but parallel on the red side of the LSPR. Furthermore, additional theory predicts that this crossing point in torque doesn't align with the LSPR peak wavelength,¹³⁷ because in addition to the extinction component of the torque that changes sign as it passes through the LSPR, there is also a scattering component to the torque that does not change sign. Interestingly, this theory means that both polarization and wavelength can be used to control nanorod orientation, though we do not showcase the latter control mechanism in this report.

We will briefly expound the theory of torque generation here for the reader, largely as described by Xu and coworkers.¹³⁷ The coherent oscillating electron density comprising the LSPR generates a current density within the nanorod. This current density results in a vector potential, which is then used to describe the scattering of the incident EM field by the rod. This scattered field generates a scattering torque, which interacts with the incident field to result in an extinction torque. Together these two torques comprise the total torque. This torque is calculated using the Maxwell stress tensor, which represents the total angular momentum flux density. When multiplied by the unit vector, the net angular momentum density flowing through the nanorod is achieved, which is then used to calculate the torque experienced by the nanoparticle. For larger aspect ratio nanorods, multipole resonances should also be considered.^{137, 192} Interestingly, circularly polarized light (CPL) has also been used to induce rotation through a different mechanism, whereby spin angular momentum (SAM) from individual photons is converted into a mechanical torque upon light absorption by a nanorod or nanowire. This has been used

to spin silver nanowires and gold nanoparticle dimers at a few Hz experimentally¹³⁸ and theoretically,¹⁹³ and to spin gold nanorods at tens of kHz experimentally.¹⁹⁴ Yet more recently, this extremely fast rotation of AuNRs has been further explored in conjunction with induced photothermal heating,¹⁹⁵ as plasmonic heating is important to consider when modeling the light-matter interactions that govern optically-induced rotation of plasmonic nanostructures.¹⁹²

5. Towards Single Molecules: Reducing Photothermal

Background

5.1. Abstract

Our microtoroid and microbubble spectrometers can measure resonance shifts smaller than 1 am and 50 am respectively,^{10, 20} orders of magnitude below the photothermal signal we expect from a single chromophore (>1 fm). However, our platform has thus far been practically limited from single-molecule experiments due to photothermal background. I have diagnosed two sources of this photothermal background in microbubbles, and have made considerable progress in reducing both sources.

The first source of background comes from electronics. In particular, I discovered that channel bleed-through in our lock-in amplifier was increasing background. By optimizing amplifier settings, I lowered this electronic-background by over an order of magnitude to <100 am of apparent resonance shift. Importantly, this background is **independent of pump laser power**, meaning signal-to-background ratio (SBR) can be improved by increasing pump power!

The second source of background comes from materials absorbing pump beam light. While we anticipate negligible absorption of visible pump light by the silica of our microbubbles (calculation in *Section 5.12*), contaminants from either the fabrication process or the resonator collecting dust will cause unwanted light absorption. Critically, this materials-based photothermal background **scales linearly with pump power**, and thus more power will not improve SBR. I have reduced materials-background through an oxygen plasma treatment, with one microbubble achieving a 5x reduction in photothermal background. I have further improved SBR by a theoretical factor of 2x by switching to a higher numerical aperture (NA) objective, which provides a more tightly focused beam that permits reduction in pump beam power. We have also explored alterations to the microbubble fabrication process to reduce contaminant incorporation into the resonators.

Through the reduction of electronic-based and materials-based background, we now are **in a regime where we calculate single molecule measurements to be just-feasible.**

Additional reduction of background through more rigorous cleaning procedures and fabrication alterations will improve the feasibility of single molecule photothermal studies. Further improvement will also be achieved by using microring resonators instead of microbubbles. Microrings will permit an inverse-excitation geometry, allowing the switch from an air objective (0.9 NA) to an oil objective (1.45 NA) if the resonator substrate is thin enough, providing a concomitant increase in SBR by a predicted factor of over 2.5. Thus, with our photothermal experiments at the threshold for single-molecule capabilities, and multiple avenues for additional reductions in background in the future, **the prospects are promising for single molecule absorption spectroscopy with microresonator spectrometers.**

5.2. About This Chapter

This story spans many months, with a pandemic, building construction disasters, building a new lab space, and other setbacks in-between. Thus, the totality of data collected will not be discussed in this chapter, particularly since comparing one dataset to another from this time period may not be a good comparison. I will, however, pull in illustrative data as I tell the story of what we have learned so far about photothermal background in microresonator absorption spectroscopy, where open questions remain, and where there is room for further improvement.

A note to the reader: This chapter employs significant technical language regarding microresonator photothermal spectrometers. Chapters 2 and 3 of this thesis walk through much of this jargon in significant detail.

5.3. Background on Photothermal Background

Our research group has been saying for years that our photothermal microresonator spectrometers (initially of the toroidal variety, and more recently, the microbubble variety, and in the future, probably multiple other varieties) have the necessary sensitivity to allow us to study single molecules. And while we have done the limit-of-detection tests to confidently support that claim, we have yet to

experimentally fulfill it with single-molecule measurements, unless you count single polymer studies!¹⁹

What's holding us back from doing true single molecule spectroscopy?

First, let's take a hard look at those limit-of-detection measurements mentioned above. In both instances, we parked the pump beam on a single gold nanorod (AuNR). Then, we measured the photothermal signal at various pump beam powers, decreasing power until the signal was indistinguishable from the "pump off" resonance shift. Importantly, we measured the signal at a lock-in time constant of 1 second, and averaged that data for 30 seconds. While this might be a perfectly acceptable averaging time for plenty of interesting studies, it is pretty long if you want to look at a dynamic system, for example a molecule chemically reacting. And while many chemical reactions can be tuned to proceed at a desired speed by adjusting reagent concentrations, eventually instrumental sources of noise will start to negatively impact the limit-of-detection if acquisition time is too long.

Another very important aspect to our experiments is that we have to *find* the nano-objects of interest in the first place! This involves discriminating against the photothermal background of the resonator (signal-to-background), as well as against the variation in photothermal background from pixel-to-pixel while mapping. While this variation could be interpreted as "signal-to-noise", it is probably wise to avoid this assignment, as this variation is not noise, but a static and reproducible pattern in photothermal background. In essence, to find an object, we need it to have enough "hot pixels" compared to the pixels around it, such that we can robustly identify the object, usually with a 2D-Gaussian fit.

Thus, our signal from a nano-object has to be robust enough to overcome both the background (real + instrumental), as well as the noise on top of the background, *and* any variations in materials-based

¹ Which may count as single molecules from an absorption cross-section perspective, but still fall short of the gratification that would come from single dye molecule measurements, for example.

photothermal background. The less “flat” the background is, the harder it will be to see an object on top of it. In summary, **we want to minimize background and noise, and also have a *flat* background.**

5.4. Origins of Photothermal Background

Where does the photothermal background come from in our microresonator experiments? Firstly, having light-absorbing contaminants (i.e. **dust**) on or within a resonator will increase photothermal background. Even for a pristine resonator, however, there may still be materials-based sources of photothermal background coming from the **microresonator itself**. For silica-on-silicon toroids, with a visible pump beam, the most obvious culprit is the silicon pillar and substrate.⁴⁸ For the same toroids but with a NIR pump beam, our group has evidence implicating residual water (perhaps within the resonator from fabrication and/or on the surface of the resonator).⁷ For the all-glass toroids that invented within the group,⁴⁸ background is far lower in the visible light regime. However, background is *still* likely too high to accomplish single molecule studies, perhaps due to residual silicon that was not fully oxidized during the fabrication process. Such residual silicon (as well as water) has also been observed in another type of toroids fabricated within our group, based on polysilicon.³⁶

Given the above difficulties with toroids, we decided to try out a new platform: microbubble resonators. In addition to their various benefits, discussed in Chapter 4 of this thesis, we anticipated that having a silica-sourced microresonator that did not involve us oxidizing silicon at any step should give us negligible photothermal background in the visible regime. Unfortunately, this turned out not to be the case! Since microbubbles are an attractive platform for single molecule studies, we decided to do an in-depth study to unearth the origins of photothermal background in microbubble resonators, and to see if we could reduce background enough to make single molecule experiments viable.

Before we start our deep-dive into microbubble photothermal background, it is worth noting that photothermal background has also been noted elsewhere in the literature. In particular, work from the Orrit group has made interesting observations in photothermal background.^{69, 196} While such

precedent is worth being aware of, we won't discuss these literature precedents here in-depth, as they are not particularly relevant to the conclusions drawn in this chapter.

5.5. Two Types of Background

After initial rounds of data collection on photothermal background of microbubble resonators (methods are described later in this chapter), it became apparent that the background we see can be lumped into two primary categories. The first category is “real” photothermal background, which arises from the absorption of the pump light by the resonator itself (or solution within it, or very tiny or out-of-focus objects on it, etc.) I call this **materials-based** background. This sort of background may be reducible with altering the fabrication procedures and/or source materials for the resonator, or perhaps even with cleaning methods, as discussed later in this chapter.

The second type of background actually has nothing to do with the photothermal effect¹. Instead, this background comes from the various electronic instrumentation itself, thus I often refer to it as **instrumental background**.² This background was discovered when we realized that there was a non-zero, and sometimes significant (relative to an anticipated signal molecule signal) “resonance shift” measured (i.e. lock-in amplifier voltage output) even when the pump beam was off! The relative magnitude of each type of background, materials or instrumental, depends significantly on the circumstances, as discussed in detail in this chapter. But before characterizing how much background we have, we should first take stock of how low we need background to be to make single molecule experiments feasible.

¹ I suppose a small amount would fundamentally arise from thermo-refractive noise near the modulation frequency, if that exists, but that's rather irrelevant in this particular discussion.

² I colloquially call this “Fake News Background”.

5.6. How Low Does Background Need to Be for Single Molecule Experiments?

This question has three important answers for a given molecule-of-interest:

- (1) The **instrumental background** must **be less** than the expected signal for the single molecule.
- (2) The **materials background** must **scale less** with pump laser power than the single molecule.
- (3) The pump power employed cannot go beyond the **saturation intensity** of the single molecule.

5.6.1. Toroid calculations

Previous calculations from our group for single molecule feasibility with toroidal microresonators used a single chromophore with $4.6 \times 10^{-17} \text{ cm}^2$ absorption cross-section ($\epsilon = 12,000 \text{ L mol}^{-1} \text{ cm}^{-1}$) pumped with an intensity of $1.9 \times 10^9 \text{ W/cm}^2$ of light (calculated with a 120 mW, 556 nm wavelength beam and a numerical aperture (NA) of 0.95 for the focusing objective). Importantly, for the molecule used for these calculations a very short excited state lifetime of 2 ps means that even this high intensity of light would have the molecule only experiencing 10% of its saturation intensity. The importance of this saturation intensity will be discussed shortly below. Given the above variables, one arrives at an expected heat output of **10 nW** from the molecule. This 10 nW of heat was then fed into COMSOL simulations (see an explanation of how these optical and thermal simulations work, with microbubble resonators, in *Section 4.13*). All said and done, a resonance shift of **1.1 fm** is anticipated. Considering complications due to the thermal transport properties of water, a worst-case scenario was assumed to be a 10x reduction in resonance shift, i.e. **110 am**.

Thus, for the toroidal system just outlined, in-water (we want solution-exposed molecules to study chemical dynamics), the resonance shift as a function of pump beam power is on the orders of **1-10 am/mW** (using the best- and worst-case scenarios above).

5.6.2. Our molecule of interest

The target for single molecule studies in this chapter is a dye called Black Hole Quencher 3 (BHQ-3; Biosearch Technologies), the structure of which is shown in **Figure 5-1** along with some relevant spectral information.¹ BHQs are essentially very good chromophores (high absorption cross-section) and very poor fluorophores (abysmal quantum yield). Because non-radiative decay channels are much faster transitions for molecules than radiative decay channels, this means that the excited state lifetime of the molecule is very short compared to a fluorophore. Thus, BHQs both (1) efficiently transduce absorbed photon energy into heat output, and (2) do so very quickly. This property of quickly turning photons into heat is essential for picking an initial single molecule target for photothermal microresonator spectroscopy.² I have not seen an excited state lifetime reported for BHQs, but given their very poor quantum yield (and thus minimal use of radiative transitions to return from excited to ground state), and considering a typical good fluorophore has an excited state lifetime of around 1 ns, we estimate that BHQs will have an excited state lifetime of low tens of picoseconds. Approximating it somewhat conservatively as 20 ps, we should expect the saturation intensity of this system to be about 10% of that outlined in the toroidal calculations above. Thus, using a similar objective and wavelength, we anticipate that we will saturate our BHQs at around 120 mW of pump power, and should **target to hit them with something on the order of 12 mW** of pump power to stay well below saturation.

¹ Goldsmith group members can find other suggested single molecule targets that we've previously talked about in my files on databackup.

² Why make this harder even harder than it already is?




Figure 5-1. BHQ-3 structure, spectral parameters, and solubility.

5.6.3. Microbubbles calculations

How does the **microbubble** system studied in this chapter perform in a similar calculation (to the toroidal calculations above) for BHQ-3? We will start with the fact that in our microbubble COMSOL calculations (Section 4.13), we learned that a 10 nW heat source would yield about 4 fm of resonance shift (see Figure 4-2), and recall that it's the amplitude of shift oscillation that matters, as our lock-in detection is blind to the thermal static offset. Compared to the toroid calculation, we get a boost of x3.33 from the relative cross-sections being considered (BHQ absorption cross-section of 1.5×10^{-16} , from $\epsilon = 40,000 \text{ L mol}^{-1} \text{ cm}^{-1}$). We pay a small penalty of x0.88 for going to a redder wavelength of 635nm (spot size increases). When combining all of these adjustments, we arrive at a thermal response of **9.7-97 am/mW** for water-modes of a microbubble when maintaining the original order-of-magnitude range.

5.6.4. Impact of dn/dT

A primary reason for the increased responsivity, as compared to the toroidal system, is that for microbubbles the optical mode for some resonances lives **within the water** inside the resonator. This gives a x10 boost to the thermo-optic coefficient, dn/dT (Section 2.6) compared to optical modes contained within glass. The COMSOL calculation of 4 fm of resonance shift (for 10 nW heat output) isn't exactly 10x that of the toroid estimate (1 fm), likely due to the increase in mode volume compared to

the well-confined fundamental mode of a toroid, as well as the other thermal properties (ex. heat capacity) of the microbubble system, which the COMSOL thermal simulations should account for.

5.6.5. Impact on Q-factors

Another thing to note is that the Q-factors for the water-modes of a microbubble will be about $1/10^{\text{th}}$ the Q-factors for glass modes of a toroid (we've observed this experimentally). This will impact the limit-of-detection of the system, as lower Q-factor means shallower error slope, which means the lock-in amplifier received less voltage change per resonance shift. Thus, *sensitivity* is lowered (*Section 2.7*) However, this drop in quality factor should not impact the thermal *responsivity* of the resonator (*Section 2.7*). Additionally, with our current PDH system we actually attenuate our error slope quite a lot, so it is unclear if the limit-of-detection will actually be that negatively impacted by this loss in Q-factor.

5.6.6. Other considerations

A last consideration is that the microbubbles have very poor heat sinks (air and thin glass capillaries). As described in *Section 4.13*, the consequent limitations of thermal transport in the microbubble system impact the amplitude modulation speeds that can be employed. The slower modulation speed for microbubbles, as compared to toroids, will lead to a higher $1/\text{frequency}$ noise floor, meaning that instrumental background may suffer compared to the fast toroidal modulation.

Overall, we've obviously made a lot of assumptions in the above chain of calculations. However, the order-of-magnitude range given above should be enough for us to assess if an observed photothermal background puts us in a realm of single-molecule feasibility. The rest of this chapter details progress towards reaching the single-molecule limit, initial attempts at single molecule experiments using microbubble resonators,¹ and recommended next steps for future work.

¹ Spoiler Alert: We aren't there yet, but I think we're getting quite close.

5.7. Measuring Photothermal Background

Before we delve into the details about each type of photothermal background, we'll first discuss here the experimental protocols used to measure background. This is both to explain how the data was taken that led to the conclusions reached in this chapter, but also for future application in microresonator experiments, as I suggest making photothermal background measurement a **standard procedure** for every new resonator (and I suggest making this as automated as possible!). In this section, we will cover instructions and considerations for **taking** the data. The **analysis** of the data will be discussed in the following sections.

The first step to characterizing photothermal background is to get a microresonator experiment fully set-up. Get the taper, resonator, and pump beam all aligned to each. Get all of the electronics on and connected. Get light coupled into your resonator from the taper and confirm that you have resonances of acceptable quality. Measure an error slope and lock to that resonance. For help with all of this preparation, consult Chapter 3 of this thesis.

5.7.1. Pump off map

The first step to take is “pump beam off” data.¹ This is typically done while the locked to a resonance (requiring phase-modulation to be on), and with the galvo mirrors scanning. Importantly, amplitude modulation must be on, and the lock-in parameters must be considered (as described below). That is, **the experiment is identical to a normal photothermal map, except the pump beam is off or blocked.**¹¹ This provides the **instrumental background** of the system, which is a pump-power-independent baseline of resonance shift, which is the “noise floor” that must be overcome in signal to see an object. Given that the instrumental background will follow something like a Gaussian profile, it may be that a very poor SBR

¹ Well, the order doesn't technically matter, but you might think of this like your “blank” sample, to ensure that the system is operating properly before you dig into further experiments.

¹¹ It might be that you also want to test out other experimental components turned off in isolation (unless they prohibit the measurement from working), to try and further isolate noise sources, though hopefully this chapter gives you a good idea of where instrumental background likely originates.

will actually be sufficient for identifying objects, however ideally, I recommend aiming to get your desired signal at least three times higher than the average instrumental background.

In the way I have prescribed it, the instrumental background is taken while in “mapping mode.” Once could instead just take a time series of lock-in output values and calculate apparent resonance shift from this. However, I think doing the mapping provides a more-realistic experience, as it makes sure that all electronics, such as DAQ outputs and galvo mirrors, are operating, as they would be when acquiring an actual photothermal map. Thus, if for example the galvos’ loud ringing noises are causing vibrations through the air or the table that get transduced into background signal, this will be included in the measurement.¹

5.7.2. Pump on Maps

In my photothermal background studies, I’ve mostly incorporated on two types of materials-based background studies. The first is what I call an “object map”, named so because the map includes some object, typically a piece of dust, within it. The second is called a “region map” which deliberately targets an area of the bubble that includes no known pieces of dust within it. I suggest taking data both types of data, as this gives a good variety of information to analyze. Next we will cover more details on each kind of map, with instructions for analysis relegated to the following sections of this chapter.

5.7.2.1. Object maps

Object maps are a convenient way to determine the beam size in a resonator area that is being mapped, which for a sufficiently small map, should be about constant. If the map is too large, or near the edge of the resonator, then the curvature will mean that the beam size at the resonator will change as a function of lateral position, unless active control of the objective-focusing piezo is employed (i.e. a by-pixel calibration matrix is used, a capability not currently included in our custom control software). If the

¹ I have yet to observe any increase in instrumental background from running the galvos, but I still think it to be good practice for these measurements to have all possible instrument components operating.

mapped object is a known analyte, such as a gold nanorod, it can also serve as a useful quantitative check to make sure that the experimental resonance shifts match theory, and to give a good indicator of what the smallest object is that could be seen, given a certain resonator and experimental conditions.

5.7.2.2. Region maps

Region maps are a good way to ascertain photothermal background patterns. After initially mapping the entire bubble to identify where pieces of dust are, one can select a region in which no dust is apparent. This allows one to crank up the pump power to very high values (tens, perhaps hundreds of mW) without worrying about a sudden spike in resonance shift (due to a piece of dust) knocking the servo off of the lock. Additionally, higher powers more readily reveal the materials background pattern! As noted above, this background may be excitation wavelength-dependent as well. Sometimes after turning the power up (and likely increasing the spatial scanning resolution, from something like 1 μm for course mapping to something like 100 nm for fine mapping) very small pieces of dust are suddenly discovered that were not seen in the larger maps! These tiny dust objects also may give an idea of what scale of resonance shift is readily resolvable above the photothermal background, provided that one is in a position to trust the absolute scale of resonance shift.

5.7.3. Mapping parameters to consider

There are a few important mapping parameters to consider when taking these background maps. If you're uncertain about the optimal value for a parameter, or want to assess the impact of that parameter, you may want to do a **parameter sweep** where you take background maps at multiple values for that parameter, with all other parameters held constant. Of course, one has to be judicious with parameter sweeps, as sweeping every possible parameter against each other typically leads to an unrealistic quantity of data to acquire and process due to the massive parameter space. Such mapping parameters to consider include:

- **Excitation Wavelength:** Whatever materials/contaminants are absorbing pump light likely due so with wavelength-dependence. Thus, whatever wavelength(s) you want to use for sensitive experiments should be tested for photothermal background. Also, using multiple wavelengths may help you figure out exactly what materials/contaminants are giving rise to photothermal background!
- **Excitation Power:** The materials-based background will scale with excitation power. Thus, a high excitation power is likely desirable to confidently get the shift/power parameter, and to reveal spatial variation in the background. We expect the relationship between materials background and power to be linear, unless whatever is giving rise to the background becomes saturated, which we do not anticipate at the powers at which we operate.
- **Excitation intensity:** This parameter considers both the set power, as well as the laser spot size at the resonator surface. Some of the code operation discussed below allows for determining excitation intensity so that resonance shift can be plotted against it. However, unless the background is saturated (unlikely in our case), the intensity trend should be the same as the excitation trend. Thus in practice, I rarely analyze power vs. intensity. Even so, the spot size of the laser beam is quite important when considering SBR!
- **Resolution and size of maps:** Larger maps give a more robust picture of the overall photothermal background of the resonator. However, for larger maps either resolution must be compromised, or experiment duration must be increased. Additionally, larger maps run risks of (1) the error slope changing or lock breaking, and (2) encountering an object, which can cause problems if operating at a high pump power. High resolution (ex. 100 nm steps) gives an excellent picture of the local variation in photothermal background on the spatial scale relevant for single molecule studies.

- **Location of maps:** It is wise to map out multiple regions of a resonator to get an adequate sample size for assigning an average background to a resonator, and understanding the variability in background across it. Additionally, this may help see how the curvature of the resonator, and the distance from the taper, impact background. However, these parameters will likely **also** impact the signal from a single molecule, thus an **internal standard** (such as AuNRs) should be employed if taking this route. **Mode selection** may also be important in this case, as the location of the optical mode is expected to play an important role in the spatial-location-dependence of photothermal signal!

5.7.4. More instrumental parameters to consider

Now that we've discussed a few important mapping parameters, let's also consider some important electronic parameters of the instrument, especially with regards to the lock-in amplifier. There is further discussion on lock-in amplification as it relates to background later in this chapter.

- **Time Constant:** The time constant of the lock-in amplifier sets how long the amplifier averages in the time-domain. Perhaps intuitively, the longer one averages their data, the "better" that data will be.¹ This is because if the lock-in takes longer to make a measurement, the bandwidth of the measurement decreases, meaning that in the frequency domain, a narrower slice of frequencies are sampled. Consequently, the **noise** picked up in that frequency range, and down-converted to DC "background", decreases. Thus, **increasing the time constant is expected to decrease instrumental background.**
- **Sensitivity:** The "Sensitivity" parameter of the lock-in amplifier sets the *dynamic range* that the amplifier can operate at. No matter the Sensitivity, the output of the lock-in will always be a voltage between 0 V and 10 V. For example, if the Sensitivity is set to 100 mV, and the input

¹ Until the averaging time gets long enough that some other sort of instrumental noise starts to creep in. An "Allen Deviation" plot is often used to identify what the optimal TC is for minimizing noise.

signal is oscillating has an amplitude of 10 mV (at the reference frequency), the output of the lock-in will be 1 V. This sensitivity should be set to maximize the dynamic range of the measurement (i.e. be set as low as possible).

5.8. Processing Photothermal Background Data

Here I'll walk through how the codes I've written process the photothermal background data. Exactly how the data is processed will depend on the type of data taken, namely if it was a "pump off", an "object" map, or a "region" map, as described above, with some options that the user can choose to employ for further analysis.

5.8.1. Pump off maps

To quantitatively analyze the instrumental background, I take all the pixel values of resonance shift from the pump off map, and combine them into a histogram. An example of pump on, pump off, and histogram data from the latter is shown in **Figure 5-2**. From this histogram is extracted resonance shift mean, median, mode, standard deviation, and % standard deviation. Each of these parameters has its own utility when considering how photothermal background will impact experiment. For example, the mean vs. the median will give an indication on if the distribution is normal (which, while qualitatively apparent when looking at the histogram, is not clear from only a mean that might be input into a larger table summarizing multiple data sets). In the case of instrumental background, mean and mode are likely to be very close to one another due to the normal distribution of pixels. This pump off processing, while written in codes for standalone processing,¹ is also built into many of the larger processing codes for photothermal background, as it allows for a sort of "background" subtraction of instrumental background from materials background.

¹ Goldsmith group members: These will be called "OneHistogram" codes (MATLAB) in my files.




Figure 5-2. Representative pump on map, pump off map, pump off histogram. *Left*: A photothermal map of a piece of dust. *Middle*: A “map” taken, with the pump beam off. *Right*: The same data as in the pump off map, but shown in histogram form with extracted statistics.

5.8.2. Region maps

For the “Region” map, the pixels are also combined to form a histogram and give statistics. The standard deviation is quite useful here as it gives the variance experienced in the background. The % standard deviation is perhaps even more useful, as it indicates how flat the background is, and thus how low of an SBR can be tolerated for a given experiment. In *Section Error! Reference source not found*. I suggest parameters that could be developed to convey more information than just a standard deviation.

This data is then divided by the **pump power** used give **resonance shift/pump power** (often in fm/mW or am/mW), normalizing the resonance shift to the set pump power, which is helpful for comparing backgrounds across experiments with different pump powers. This normalized power is also useful because it can then be multiplied by a pump power planned for a single molecule experiment to predict the SBR of the measurement. This does assume that the shift/power trend is linear across a large range of powers, which is something that we have observed both in our published work looking at gold nanorods,^{10, 20} and in photothermal background measurements in this study (*Section 5.10.1*).

This pump on region data may, if the user chooses, have pump off data subtracted from it (i.e. the mean value of a pump off map). Also as described below, a calibration map of FOV is also typically employed for correcting disparities in transmission efficiency across the FOV.

5.8.3. Object maps

“Object” data is the most-complicated to process. The code requires some user input and quality-control throughout the processing, as the user will be asked to select a few MATLAB **data structure** files from a directory, and for each file, the code will map out the data and request the user to verify that the file selected was indeed the intended data.¹

First the code asks the user to select “pump off” data. Then, the code asks the user to select a “FOV map” (FOV = field-of-view). The user selects the most recent calibration for the microscope’s transmission for the appropriate pump wavelength. Thus, the pixels of the map (in voltages saved during mapping) can be correctly divided by the pump power for this object’s data set (pulled from the “pump on” data structure), and then multiplied by the transmission % for that given for each pixel. This by-pixel strategy provides a superior normalization to pump power than just a single power value.

The code then asks the user to select “pump on” data of an object, which is the data that we care the most about in this processing. Once the user confirms that the data is correct, the code finds the maximum point in the pump on map (which must be at roughly the center of the single largest object in the map), and performs a 2D-Gaussian fit on the object (which, if diffraction-limited, should just look like the pump laser’s profile at the resonator surface). This 2D-fit serves two purposes:

- (1) **Excising the object** from the map, because what we care about is the *background* data, not the resonance shift due to the object in the map. The code takes all pixels within a prescribed distance from the calculated center-point of the object, and removes this disc of points from the dataset being processed for photothermal background. The excised data is plotted in its own subfigure alongside the remaining pixels for quality assurance (**Figure 5-3**). Within the code, the user is able to change the width of the excised data disc by adjusting a single integer that gets

¹ I borrowed heavily from some data processing code that Dr. Erik Horak wrote (for the work in Chapter 4 of this thesis) for the user-interface components of this code.

multiplied by the FWHM from the 2D-Gaussian fit. In the **future**, this width parameter could be changed to automatically scale, for example based on when the width of the excision increasing no longer changes the average background of the map by a certain threshold.

- (2) **Calculating intensity:** The 2D-fit also can be used to normalize the photothermal background resonance shift by intensity, rather than by power as described above for the “region” mapping. Although this is a purpose that our data has suggested may not be necessary for the analyses in this chapter, the intensity normalizations have proven vital before, such as when processing the gold nanorod etching data discussed in Chapter 4. With the intensity calculated, the histogram statistics of the resonance shift for the map, *sans object* (post-excision), can be normalized. This normalization to intensity can **also** be done for a “region map”, without excision of an object, by changing a binary in the code so that it knows that the “pump on” map is *not* the map to be processed for photothermal background, in which case the code prompts for an **additional map** to be processed. This works out fine as long as the spot size calculated should still be fairly accurate, i.e. for a region reasonably close to the object, or somewhere else on the resonator in the same focal plane.

Drawbacks: The added complexity of doing an object map compared to a region map is both an experimental and a data processing burden. Additionally, the background from an object sometimes extends quite a distance from that object. While most of the object’s signal seems to have died off within around two FWHM’s, if one plots the background data in log-scale, it is apparent that an object of high SBR has photothermal background contributions well past what an eye-check would imply when looking at linearly-plotted data. Thus, often quite a few pixels must be excised for the remaining portion of the map to be indicative of “true” bubble background.

The practical approach: Considering all of these things, while I still always first find the dust on a resonator and get focused to the resonator surface using a piece of dust, I usually then only take and

process region data, as this seems to contain all of the necessary information for quantifying the photothermal background of a resonator. This said, if one is looking for single molecules, then the amount of extended spatial impact that dust objects have becomes relevant for ascertaining how much of a resonator can be searched for molecules, and how optimally focused the pump beam will be in the regions being searched.




Figure 5-3. Excision map. *Left:* Remaining data. *Right:* Excised data.

5.9. Instrumental Background

5.9.1. Reference bleed through in the lock-in amplifier

Instrumental background was significantly reduced after conversations with Rob McClain¹, as well as correspondence with Signal Recovery/Ametek, the manufacturer of the lock-in amplifier we use (DSP 7250) for the microbubble resonator microscope. The advancement was a procedure that I refer to as the “auto-offset” procedure, which is used to remove any reference signal that is leaking from the reference channel into the signal channel of the lock-in amplifier. Thus, any reference signal that was previously being down-converted to DC as output signal gets removed from the lock-in output. In practice, this procedure involves three steps:

1. **Auto-offset:** With the **pump beam off/blocked** (and thus no “real” incoming photothermal signal to the lock-in), run the “auto-offset” procedure from the front panel of the lock-in.

¹ The Analytical Lab Director in the UW-Madison Chemistry Department.

2. **Auto-phase:** With the **pump beam on**, and hitting the resonator such that real photothermal signal is being detected by the lock-in, run the “auto-phase” procedure from the front panel of the lock-in.
3. **Auto-offset:** With the **pump beam off/blocked**, repeat step #1 by running the auto-offset procedure from the lock-in front panel.

Step #1 tells the lock-in how much apparent “signal” is not genuine, and needs to be subtracted. It is a bit unclear if step #2 is actually needed in our experiment, since we operate in quadrature mode, where our measurement should be insensitive to the phase-difference between reference and lock-in channel. Nonetheless, I have typically performed all three steps, as suggested by the lock-in manual.

So, how well does the auto-offset procedure work? In the microbubble experiments I performed, it **lowered my instrumental background from the order of 1 fm to the order of 100 am**, and sometimes to the order of 10am or less! Thus, instrumental background was reduced by at least one order of magnitude, and sometimes more. Given the signal we anticipate from a single molecule (*Section 5.6.3*), and the laser power at our disposal, the instrumental background just went from being high enough to obscure a single molecule signal to instead being a minor factor.

How often should the auto-offset procedure be run? It seems like the benefits mostly carried from one resonator to the next in my experiments, but just to be safe, anytime a new resonator is installed, or the resonator environment has changed (ex. by flowing in a solution), it might be helpful to run the auto-offset procedure. If anything about the instrument is changed, ex. a BNC cord is swapped out (which might change the phase between the signal and reference channels), it would also be advisable to run the procedure again. This procedure could also likely be automated and included in the master-code for the photothermal microscope if desired.

Error! Reference source not found. shows another interesting result from the auto-offset procedure, as well as how the lock-in sensitivity parameter impacts the output signal. Four data sets are

shown, all of which were taken **with no “real” photothermal signal**. The output voltage of the lock-in was measured as a function of the “Sensitivity” parameter (here called “dynamic range”), for both an “open” lock-in input, and for having an input that is locked to a resonance. The auto-offset was performed, and then both of these datasets were taken again. Observations from this dataset include:

- (1) For all of the traces, **output scales linearly with sensitivity** past a certain point (which I call the “sensitivity corner”), and below this point, **output voltage is constant**, and thus independent of the Sensitivity setting.
- (2) The auto-offset procedure did **not impact the background before the corner**, but **lowered the background above the corner**, by almost an order of magnitude, as indicated by “a”.
- (3) The corner **shifted to the left** (lower sensitivity) after the auto-offset, as indicated by “b”. The amount of this shift may be dependent on at what Sens. setting the auto-offset was performed, though I have yet to rigorously test that. For this data, the auto-offset was done while Sens. was set to 20 mV.
- (4) For both data set pairs, having the lock-in input open has **lower background when left of the corner**, by over an order of magnitude (as indicated by “c”), as compared to having a locked-resonance input. This implies that **PDH electronics may be contributing to instrumental background**.
- (5) The absolute level of the locked-resonance background varies. Here, it was actually a bit higher after the auto-offset than before, but this could be due to the variation in error slope between the two data sets and/or variation in locking stability.

¹ Doesn't that seem more fitting?




Figure 5-4. Lock-in amplifier measurements, before and after the post-offset procedure. “Optimization” means auto-offset procedure, and “Dynamic Range” is the “Sensitivity” parameter. Data is shown for both open-input and locked-resonance input with no pump beam. Optimization lowered background at higher dynamic ranges where we typically operate (“a”). The left-shift in the “corner” where background starts increasing (“b”) also indicates lowered background. Noise from other electronics is also observed (“c”). An increase in the post-optimized signal floor (“d”) compared to the pre-optimized floor indicates noise variation in locking to different resonances.

While the auto-offset procedure reduced instrumental background to an almost-negligible impact for our planned single molecule studies, if materials background is lowered enough, then the instrumental background may again become a dominant source of background. In such a regime, the instrumental background might prohibit the study of weaker analytes or faster dynamics. Next we discuss possible routes for lowering instrumental background further.

5.9.2. Ground loops

Another area that could maybe still yield noticeable improvement, per another suggestion from Rob McClain, would be to search for and eliminate **ground loops** in the electronic paths of the system, which may occur when different parts of an instrument reference a different “ground” voltage, rather than the same true earth. These would be especially troubling if occurring in circuits involving the lock-in amplifier. One suggestion to try was the inner part of the “B” input on the lock-in (which we typically leave open) to the grounded outer part of the “A” port/BNC leading into it, and then using the A-B setting on the lock-in rather than just reading the A channel. I have yet to try this, but it is worth exploring in the

future. Other ways that reference channels can leak into signal channels are inductive and capacitive routes, but I have no reason to believe that these are happening in our system to any significant degree, and that checking for resistive routes (i.e. ground loops) is the most likely avenue for improvements to be made.

5.9.3. Other pieces of equipment

Are any of the other many components to this instrument contributing to the instrumental background? For example, does acoustic noise from the chopper get transduced into additional background? Does having the oscilloscope on, the LC on, or any other piece of equipment on create electrical noise (whether resistive, inductive, or capacitive) that somehow impacts the instrumental background? Does 50 Ω termination of various paths matter? Does the 3D piezo stage for the microresonator (attocube) add in background through the mechanical frequencies of its piezos, even when at rest? The list goes on. While I certainly haven't tried every single permutation I can think of (there are many of them!) and probably haven't thought of other possibilities, **I have not seen any significant impact to anything I've tried.** Perhaps occasionally a decrease of 10-20% in instrumental background was observed, which is not particularly impactful, even if real.

However, even after the auto-offset procedure, there is (usually) still some instrumental background left! Whether or not we can get rid ourselves of this remaining background, it would be nice to know where exactly it is coming from. We haven't fully diagnosed this remaining background yet, but we do have some ideas and a couple of preliminary data sets! There are at least four questions worth asking:

- (1) What background and noise are the lock-in and DAQ causing?
- (2) How do acoustics and air currents play a role?
- (3) What extra background and noise are being produced as a result of the rest of the PDH system?

Can we isolate one or more specific components as the primary culprit?

(4) How much noise is fundamental and/or inevitable?

5.9.3.1. Background from the DAQ-input and lock-in output

What if we are limited by the components measuring our signal, and not components upstream? We measure the lock-in output with a DAQ (Data AcQuisition card¹; part numbers from National Instruments are PCIe-6353 for the card, NI BNC-2120 for the measurement block, and SHC68-68-EPM for the shielded cable). Thus, it is worth determining what the **minimum voltage** that can be measured by the DAQ is. For a given error slope and lock-in setting, this voltage can then be backed out to what I call the **“DAQ-limited resonance shift.”** The lock-in amplifier’s output, with no input signal, should also be measured (if it is above the DAQ’s minimum resolvable voltage), to determine what I call a **“lock-in limited resonance shift”**.

For the DAQ-limited input, one parameter that I explored was the grounding setting for the BNC input used to measure signal from the lock-in output. The inputs of the DAQ block have two settings: floating source (FS), and ground-referenced source (GS). The latter connects directly to ground, while the former connects to ground through a 4.99 k Ω resistor (per the installation manual for the BNC-2120). It was not intuitive to me which setting should be used in our experiment, so I tried out both settings to see if one gave lower voltage threshold and/or lower measurement noise (which will impact how small the resolution is between baseline and the lowest measurable signal).

Below, and in **Figure 5-5**, a few useful results from the grounding tests are summarized in terms of voltages measured.¹¹ Further below we also discuss the conversion of these voltages into resonance shifts.

¹ Sometimes abbreviate as DAC, though this runs the risk of confusion with a digital-to-analog converter (though the DAQ perhaps arguably functions in this capacity since it takes/gives analog voltages from inputs/outputs and talks to a computer on the back end).

¹¹ Group members can find the raw data and further analysis under Microbubble/Photothermal Background/210715

- **Lock-in off:** With the output of the lock-in still connected to the DAQ, but the lock-in off, essentially zero voltage was measured at the DAQ for either the GS or FS setting, indicating that the **DAQ will not be a limiting factor**.
- **Lock-in On, input open:** This test should indicate what noise comes from the lock-in amplifier itself, while a reference is applied but no signal is input. This was tested for a variety of Sensitivity settings of the lock-in (this is essentially the same experiment, in briefer form, as done for in Error! Reference source not found.). The result was that for higher Sens. settings of 20 mV and 1 V, GS and FS channels each measured about **80 mV** of signal. For a much lower Sens. setting of 2 uV, the measured voltage was about 4 V for both grounding settings, though this is somewhat unsurprising given the results in Error! Reference source not found., and not relevant to our experiments, as we usually operate in the tens of mV input range.

So, it looks at least for these experiments that the **grounding setting chosen did not matter**. This is helpful moving forward, but if instrumental background looks high for some reason, it is probably still worth toggling the grounding setting to see if that helps!

I also took measurements with other pieces of equipment in shared electrical paths turned on or off, namely the lockbox (probe laser servo controller), and the oscilloscope. None of these measurements, including ones where I placed a 50 Ω terminator at various spots in the electrical paths, had any noticeable impact on instrumental background.

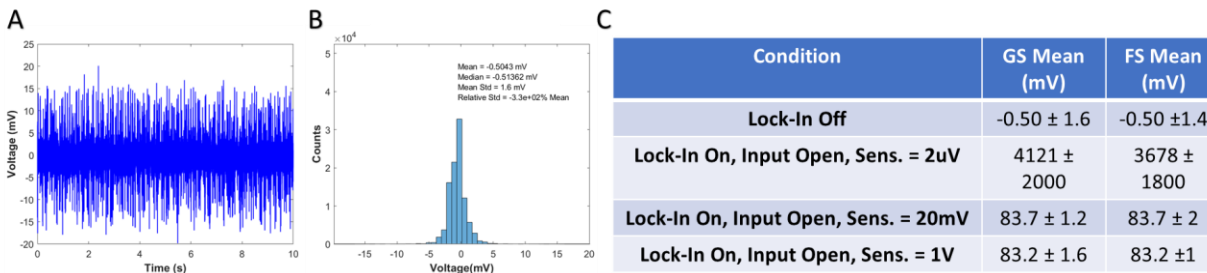


Figure 5-5. DAQ input and lock-in output tests. (A) An example data trace from this study, where the voltage of a DAQ channel was measured for 10 seconds. (B) A histogram showing the data from the trace, with extracted statistics. (C) Mean voltage measured (with standard deviation of the mean) for the four most relevant electronic configurations, for both DAQ grounding options.

Converting to resonance shift: Let us work with a value of 80 ± 2 mV of DAQ-measured, lock-in originating instrumental background, as this was a consistently measured value. Using a typical error slope that we lock to of 5 V/pm, and the 20 mV Sens. setting, this translates to **90 ± 2 am of instrumental shift from the lock-in itself**. This might be an unsurprising results, as ~ 100 am is the value I quoted for remaining instrumental background after the auto-offset procedure. However, this is also very useful. It tells us that **perhaps almost all remaining instrumental background in our experiments comes from the lock-in amplifier**.

Channel cross-talk: I also tested if measuring two DAQ input channels simultaneously gave cross-talk (something we've observed before when simultaneously measuring multiple channels) that inflated the background. With my brief tests, this does not impact the measurements, but further work could be done under various circumstances to see if this always holds true.

5.9.3.2. Background from acoustics and air currents

While the DAQ measurements imply that perhaps almost all remaining instrumental background originates from the lock-in amplifier itself, it is still worth considering the other above-hypothesized routes for instrumental background, as they could still perhaps find their way into future experiments.

The current permutation of the microbubble microscope employs what look like a number of different diving-board geometries for the tapered fiber, the microresonator, the excitation objective, and the polarization rotation optics. These diving boards must have their own *mechanical* resonances. If any of these happen to overlap with the amplitude modulation frequency (within the bandwidth of the lock-in detection) then they may be transduced into background signal, of the “instrumental” variety. Even without such mechanical resonances, vibrational noise at the amplitude modulation frequency might seep through into signal. Additionally, any vibrations in the air (i.e. sound) that impinge upon the resonator/taper system can directly impact experimental stability and noise floor. In particular, such noise sources may impact the **stability of taper-resonator coupling**.

We have yet to see any impact from table-transduced noise. For example, tests with the optical chopper off, but reference still being sent to the lock-in, did **not** result in a change in instrumental background. However, we have seen impacts of **acoustic noises** in the lab, such as noise from doors shutting and loud talking show up in our signal. Future work in installing acoustic shielding around the photothermal set-up will be very helpful in remedying this issue!¹

Later in this chapter (*Section 5.9.5*) **resonance stability** will be discussed. Both the stability of the resonance and the **stability of the resonance lock** are crucial to reducing the impact of noise! Thus, lockbox settings, attenuation of the error slope, etc. may all be employed to obtain a more robust lock that is less-perturbed by acoustic noises. I have also occasionally noticed a significant increase in acoustic-sensitivity for **water-filled** microbubbles as compared to air-filled microbubbles. Performing acoustic simulation experiments (ex. in COMSOL) could be enlightening to understanding this phenomenon and

¹ An earlier version of the microbubble microscope in a different lab space did in-fact have such a padded box. The current setup is in a much cleaner airspace in which we don't want to inhibit air flow. Additionally, the introduction of the M-Squared laser system has made a full-table box impractical. These considerations in-mind, localized acoustic shielding around the resonator seems like a viable and worthwhile thing to try.

seeking to mitigate it, ex. by engineering the aforementioned diving boards to avoid certain mechanical resonances.

5.9.3.3. Background from PDH electronics

As described earlier, it is reasonable to anticipate that some of the instrumental photothermal background arises from noise within the PDH electronics. However, given how integrated the PHD system is, isolating the dominant sources of instrumental background from PDH may prove rather difficult, though perhaps an electronic spectrum analyzer could aid in this endeavor. Future efforts into optimizing PDH components with photothermal background might be worthwhile!

5.9.3.4. Background from fundamental noise sources

It is important to recall that our technique uses lock-in amplification. Thus, only the noise happening near the amplitude modulation frequency will impact the output of the lock-in as “background”. Therefore we should consider what sorts of noise may occur, at an unavoidable and fundamental limit, which will consequently get transduced into background signal.

In toroidal experiments, it was found that error signal had noise following a Gaussian profile, with a standard deviation of 217 kHz over 50ms.¹⁶ Given that Newport specifies a linewidth of 200 kHz over 50ms, this strongly indicates that at least for the toroidal photothermal experiments, **probe laser frequency jitter accounts for the large majority of remaining noise in the experiment.**

Beyond probe laser jitter, it is worth considering other sources of noise in our experiments, especially any types of noise that would give fundamental floors to our limits of detection. In his thesis, Dr. Kevin Heylman elaborated on these for the toroidal resonator system:¹⁶

- **Shot noise:** He calculated that at 1 uW probe beam power, and typical resonator and resonance parameters, the shot noise limit would be 0.008 attometers, many orders of magnitude less than the expected resonance shift from a single molecule (~1 fm).

- **Thermorefractive noise:** In our experiments, thermorefractive noise is the more relevant noise floor to consider. Fundamentally, there are statistical fluctuations of temperature within a material.¹ These fluctuations will then map into refractive-index fluctuations, and consequently some degree of measured resonance shift for a given frequency and lock-in detection bandwidth. It was calculated that for 1 s time constant and typical toroidal resonator parameters, a value of **0.2 am** of resonance shift could be anticipated due to thermorefractive noise.¹⁰ This is close to, but below, the limit of detection measured for toroidal resonators, and luckily orders of magnitude below the signal we expect to obtain from a single chromophore.
- **Other noise sources:** It was indicated that other noise sources such as radiation pressure and ponderomotive noise can be ignored.

While the above numbers likely carry over reasonably well to other photothermal platforms in our group, for future pursuits with different resonator systems, it may be worth revisiting these calculations for the different systems we use.

5.9.4. Ensuring a correct lock

PDH locking (*Section 3.3.3*) is vital for achieving the sensitivity that our photothermal microscopes possess. However, locking to resonances is not always the most straightforward. An important parameter for backing-out a resonance shift from the lock-in amplifier's voltage output is the *error slope*, which is measured for each resonance before locking to it. Thus, if the user records one error slope, but in reality ends up locked to a different error slope, this will lead to an incorrect resonance shift calculation, off by the ratio between the two error slopes. Procedures to ensure a proper lock and possible future improvements are discussed in Chapter 3 and will not be recapitulated here. However, it's worth noting

¹ Statistical fluctuations on a microscopic level. This is different from large scale fluctuations, such as drift of the ambient temperature.

that if measuring photothermal background, both materials-based and some types of instrumental-based background require an accurate lock for quantitatively correct results.

5.9.5. Resonance stability

Similarly to having a correct lock, the stability of resonances is very important in whether you can trust the signal recorded in a photothermal data set. Critically, if the error slope changes underneath you (i.e. while you're locked to it), this will result in an incorrectly calculated resonance shift for subsequent measurements that use the now-incorrect error slope in resonance shift calculations. Checking the Error Signal and DC Tap channels on the oscilloscope throughout an experiment is helpful in ensuring that calculated resonance shifts can be quantitatively trusted.

5.10. Materials Background

Above, we've discussed in great detail the instrumental sources of background. And while there is still room for further improvement in lowering that background, all indications are that instrumental background is *not* likely to inhibit us from single molecule studies at this point. Thus, we shall turn our attention to the second type of background: materials background, which arises from **transduction of pump beam photons into heat** in the resonator.

Importantly, photothermal background will likely be impacted significantly by what's *inside* of the microresonator. The majority of the photothermal background studies discussed in this chapter were done using air-filled microbubbles so that we could focus on the microresonator itself.¹ However, moving forward, photothermal background studies including the solvent (likely water¹¹) to be used in single molecule studies would provide an even more realistic diagnosis of photothermal background. In *Section 5.12* I walk through some calculations on expected and experimentally observed impact of water being

¹ There are a few data sets I have taken on water-filled microbubbles instead, which are included in various photothermal background updates in my Goldsmith group files.

¹¹ Though a project led by Dr. Feng Pan, in our group, used chloroform-filled microbubbles. That work is currently under review for publication.

present. The upshot is that while initial calculations imply that absorption of pump light by water might lead to significant materials background, in practice we have not observed this.

When thinking about materials photothermal background, it is useful not only to understand what materials are absorbing light and transducing it into heat, but **how that heat impacts resonance shift**. I suggest the reader familiarize themselves with the concept of “photothermal strength”, as discussed by Orrit *et al.*¹⁹⁶ The important takeaway here is that the thermo-optic coefficient is **not the only parameter** that matters for how heat is transported through the resonator, and this will in fact have consequences for the resonance shift produced for a given heat output. While COMSOL simulations largely take care of the impact of material properties (Ex. the glass microbubble and the water within), such parameters merit consideration. For example, the lack of a heat sink in our microbubbles as they currently are mounted has significant impacts on the range of amplitude modulation speeds that can be selected (*Section 4.15*).

However, whatever the photothermal strength of medium employed, it can be expected that the photothermal signal will scale the same for both analyte and materials photothermal background unless one of those two things is saturated with light. Thus, while employing a material of higher photothermal strength will give more signal overall (helping to get away from instrumental background!), and while the thermal properties of the resonator system’s materials may impact parameters such as modulation speed, the SBR of a given measurement should remain the same as far as material background is concerned.

5.10.1. Background patterns: reproducibility and chromaticity

One thing observed early in my photothermal background experiments was that **background patterns were reproducible**, particularly when using high pump powers! Another thing observed was that background patterns were **excitation wavelength dependent**. This is shown in **Figure 5-6**, where the same

region of a microbubble¹ gives very different patterns of photothermal background for excitation at 532 nm and 635 nm pump beams. While beam size is expected to be smaller for the green beam (provided similar beam sizes incident on the back of the focusing objective), this difference in mapping resolution cannot explain the different in pump patterns observed, which is a phenomenon we have seen across many regions on several different resonators.




Figure 5-6. Wavelength-dependent background mapping patterns. The same $100 \mu\text{m}^2$ region of a microbubble, mapped with the red pump beam (left) and green pump beam (right).

Reproducibility in background pattern arising from materials is not particularly surprising, and neither perhaps is the wavelength dependence, given that most materials absorb different wavelengths of light to different extents. In future work, these patterns of photothermal background and their chromatic-dependence could provide insight into from where materials background arises. The new M Squared laser system (*Section 3.6*), with its large range of tunable wavelength, could help in this endeavor by providing true *absorption spectra* of the materials background, which might indicate the chemical nature of the background absorbers!

As detailed earlier in this chapter, the **flatness** of such background patterns is very important for deciding if a given average photothermal background will be tolerable for single molecule experiments.

¹ Microbubble 1H, for internal group records.

Suggestions for future work in quantifying flatness of background are discussed in *Section Error! Reference source not found.*

5.10.2. Pump power dependence

I have alluded to the fact that materials background scales linearly with pump power a couple of times now. In **Figure 5-7** is data to back up this claim. This data series involved taking photothermal maps of the same region of a microbubble¹ at different pump powers, and then plotting the average resonance shift for each map, and dividing it by the pump power to acquire it. We indeed observe a linear trend of shift vs. power. For this particular region, we also observe that green light is absorbed more than red light! This is not particularly surprising (though it is still worth noting), given that many materials that could be contributing likely absorb more the further blue pump light goes.

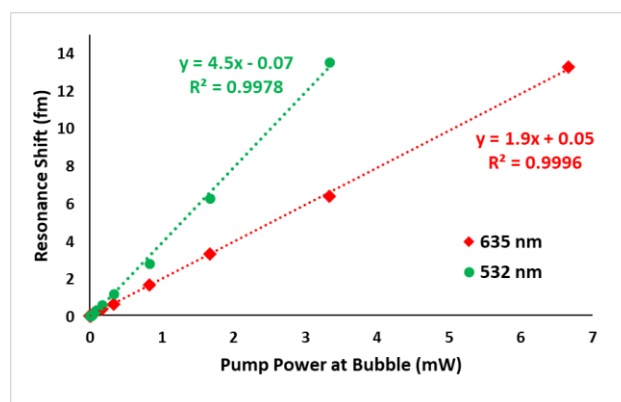


Figure 5-7. Materials background is linear with pump power. Each data point is the average background shift for a “Region” map. The dotted lines are linear fits (with fit parameters included) for each pump power series. Two series are shown, one for the red beam and one for the green beam. Notably, the slope of the green line is over twice that of the slope of red line.

The **slope** of the lines in **Figure 5-7** is the most important parameter that we are after, as this can be multiplied by a given pump power to predict the photothermal background that will arise at that power. Additionally, the y-intercept could in-theory give the instrumental background (if allowed to be a fit parameter, as for the above figure). However, it is fairly straightforward to measure instrumental

¹ Bubble 1H for internal group records.

background independently, so I suggest always doing so, and perhaps including that value as a data point in the materials-background fit if desired.

5.10.3. The carbon black hypothesis

This hypothesis is actually what initiated the microbubble photothermal background project in the first place. Our collaborators in Okinawa, who produce fantastic microbubbles for us, have a step in the fabrication process that is a bit concerning from a single molecule perspective. The “traditional” microbubbles that they make start as a glass capillary. That capillary is heated up by a CO₂ laser while being pulled by motors. This process tapers the capillary from around 300 um in diameter to around 20 um in diameter. Then, the capillary is pressurized by an inert gas, while the CO₂ laser is once again shone upon the capillary. The softening of the walls due to the laser heat causes a glass bubble to expand in a self-limiting process that produces the microbubble resonator.

The concerning portion of this procedure is that to remove the polymer sheath from the original capillary (the sheath is necessary to make the capillary flexible enough to wind up on a spool, similar to fiber optics), a flame is employed. Our suspicion is that when burning the polymer sheath off, some residual carbon black is left on the surface of the capillary. Then, when the microbubble is fabricated, perhaps that carbon black gets incorporated *within* the walls of the microbubble, in a place where no cleaning procedure is going to remove it. Carbon black, as its name suggests, is great at absorbing all colors of visible light, therefore we are concerned that if it is incorporated into the very structure of our microbubbles, it will lead to increased materials photothermal background.

It's important to note that **this is only a hypothesis!** We don't know if carbon black is actually incorporated, or to what extent. To set out on testing this, we asked our collaborators to fabricate bubbles that did not involve the sheath-burning procedure. It turns out that they had also previously fabricated microbubbles from glass capillary *without* the polymer sheath on it. Thus, we set out to compare the previously used bubble type (ex. for the work in Chapter 4), fabricated with sheath-burning

and termed by us as “once-sheathed” (OS), with bubbles that had never had a polymer sheath on them, termed “never-sheathed” (NS). A primary reasoning for using the OS bubbles for most fabrications is that the flexibility imbued by the polymer sheath permits spooling of the capillary, whereas for NS capillaries, only short, <1 m lengths of capillary are purchasable.

With a small sample size^I of both NS and OS bubbles, we did **not observe a significant decrease in photothermal background**. Future investigations into NS vs. OS bubbles are still probably worthwhile, however, at least if microbubbles as a platform are chosen to pursue single molecule studies further.

5.10.4. The glass-stress/in-coupling hypothesis

One interesting thing that I have *sometimes* observed is what seems like an increase in photothermal background at the walls of the resonator, as shown in **Figure 5-8**.^{II} One possible explanation for this is an actual increase in photothermal background due to an **increase in path length** of the resonator experienced by the pump beam. However, this seems very unlikely to provide the magnitude of increased signal observed.^{III}

An alternative explanation is that we’re somehow **scattering pump light into the resonator**, and then outcoupling that light and detecting it at our APD. And while our group has done studies looking at light scattered into resonators through objects (specifically, gold nanorods optoplasmonically coupled to whispering gallery modes),¹⁸ the detection path in this chapter’s experiments should be blind to any scattered light. This is because what we extract through our PDH process is the out-of-phase (mathematically imaginary) component of the signal produced by beating of the carrier frequency and

^I Various factors such as Covid-19, building closures, etc. hampered efforts at more substantially sized dataset. The current data can be found by group members on databackup.

^{II} This was Bubble 5E, for internal group records.

^{III} Actually doing a calculation to determine the volume of glass eclipsed by the pump beam, as a function of position on the microresonator, would be a worthwhile future pursuit. My intuition says that perhaps as long as the beam is fully on the microresonator on both the incident and opposing side that this value would be constant, but I’m not sure if this is true. The geometry of the resonator (which could be at least initially approximated as a sphere) and the focusing beam geometry make this calculation rather unintuitive for me.

sideband frequencies (see *Section 3.3.3* for further discussion of PDH, and these resources^{7, 29} for deeper discussion). After the APD converts photons-in-taper into electrons-in-wire, the electrical signal is passed through a diplexer that splits the DC and RF components. The RF components, while they are next down-converted to DC, are the parts that travel to become the error signal that gets fed to the lock-in amplifier. Since the pump beam does not carry any phase-modulation, but only the amplitude modulation frequency, this modulation should not be down-converted to the DC error signal, and thus should not show up in the lock-in output!

Nonetheless, given what we have observed, further exploration into this phenomenon is needed. If we really are coupling light into the resonator via free-space, this could have very useful implications for the types of experiments we do with WGM resonators, including our photothermal experiments! Another interesting observation in **Figure 5-8** is that for both colors of pump light shown, the right side of the resonator shows much more photothermal background than the left side of the resonator. Perhaps this observation even further supports the in-coupling of light hypothesis, as the directionality of edge-coupling would likely significantly influence the quantity of light that would reach the APD at the end of the tapered optical fiber.




Figure 5-8. Higher photothermal background at bubble edges. Shown for both red pump beam (left) and green pump beam (Right) is a full-bubble map. Interestingly, the photothermal background is quite high relative to the rest of the bubble at the walls, and is-fact asymmetrically higher on the right side of the bubble. Light couples into

the bubble from left to right in the geometry employed for this experiment, on the top-side (imaged side) of the bubble.

There are a couple of rather *indirect* ways that come to mind to test the in-coupling hypothesis. These are (1) changing **the resonance coupled to**, and (2) changing the **polarization angle of the pump beam**. Interestingly, each approach changes a parameter in a different one of the two beams in our experiment. The rationale for (1) is that if the spatial distribution of the mode is varied (by changing resonances) then the efficiency of scattering light *via* each resonance into the taper may differ. The rationale for (2) is similar, in that the polarization state of the *probe* beam impacts coupling into the resonator, thus it stands to reason that the polarization state of the *pump* beam, if scattered into the resonator, might impact coupling out of the resonator! With these two thoughts in mind, I took photothermal maps of a resonator region using three different resonances, and for each resonance, three different pump polarizations. These results are shown in **Figure 5-9**.

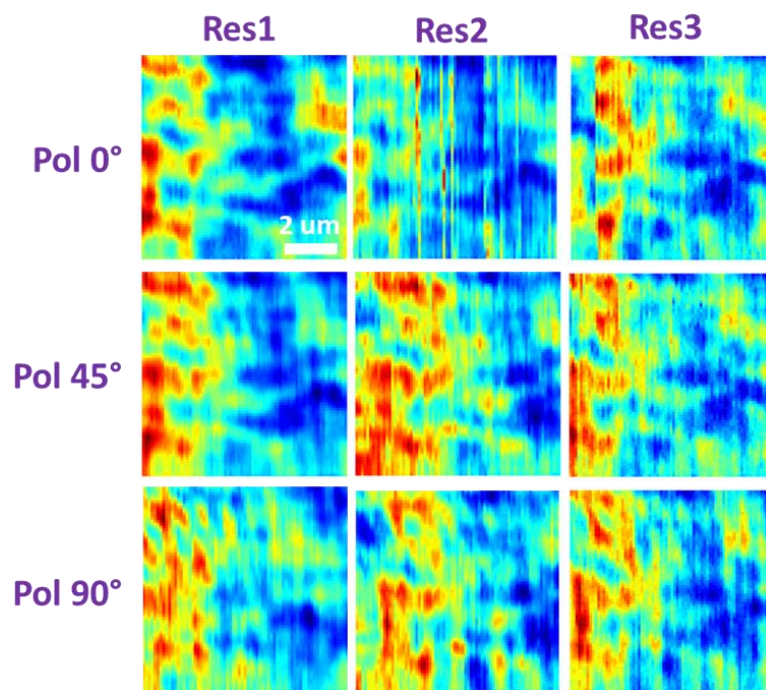


Figure 5-9. Background pattern as a function of resonance choice and pump polarization. This data is for the red pump beam. The same region was also mapped at these same 9 condition pairs with the green beam (data not

shown). The pump pattern was different for the green beam, which is no surprise, but it was constant across that set of 9 maps as well.

While there are some variations between the maps in the figure (ex. the vertical striations, which are probably due to instability in the resonance lock), for the most part the maps look quite the same! This implies that **neither pump polarization nor resonance choice impact materials background**, at least qualitatively! Thus, while coupling light into the glass cannot be 100% ruled out, given its unlikelihood in the first place, plus the resonance and polarization independence of background, we can set it aside is very unlikely to be the main source of materials photothermal background.¹

5.10.5. The dusty bubble hypothesis

Another hypothesis to explain the origin of materials photothermal background is that there are small (low light absorbing) dust particles on either the “incident” surface of a resonator (the side where the beam is focused) contribute to an overall background pattern without being strong enough and/or spatially isolated enough to have resolvable profiles. We also hypothesize that larger dust particles on the *opposite* side of the resonator to the focused pump beam may contribute to these photothermal background patterns, with the idea being that although the beam will have diverged significantly after the ~100 um diameter of a microbubble, that light could then hit multiple dust objects on the other side, which would give a small amount of photothermal signal that varies with pump beam position.

Due the randomness of dust location on the resonator, the above-two hypothesis could, by themselves of combined, create a seemingly random pattern of photothermal background. It is thus difficult to confirm this hypothesis directly with how our experiment is currently designed, as focusing to the opposite side of the resonator is challenging. While it is geometrically feasible, and more accomplishable with an air-filled resonator, once water is in the resonator we have observed very substantial lensing effects, which radically distort the PSFs seen for objects on the opposite side of the

¹ Though I still have my suspicions, and recommend that this hypothesis stay in-mind going forward.

resonator. Experimentally, the resonator could perhaps be mounted with rotational capabilities to examine all sides of it (or mounted such that it could at least be flipped upside down and still be compatible with the rest of the microscope's constraints). Theoretically, ray optics simulations plus COMSOL simulations could also perhaps give an indication to if this hypothesis of opposite-side dust contributing to photothermal background has merit. So far, we attempted neither these experimental nor the theoretical suggestions to directly how small pieces of dust randomly distributed on (or in) the resonator would impact materials photothermal background. However, we have tried some cleaning methods that could remove such dust, which will be discussed in the next two subsections.

5.10.5.1. Plasma cleaning

A common method for getting glassware such as coverslips clean for single molecule experiments is to plasma clean the glassware. While this likely wouldn't remove any contaminants embedded in the glass (such as the hypothesized carbon black from within a microbubble) it should scourge the glassware of any surface-contaminants, at least on the outside of the microbubble. Thus, plasma cleaning microbubbles seems like a worthwhile method to try for reducing photothermal background!

Plasma cleaning was first tested on a somewhat older and dusty NS microbubble.¹ Before plasma cleaning, the air-filled microbubble was measured to give a photothermal background in the range of 300-800 am/mW at 635 nm excitation, and 1-2 fm/mW for 532 nm excitation. After plasma cleaning, the photothermal background was measured to be <100 am/mW for 635 nm excitation, and around 200 am/mW for 532 nm excitation (data not shown). Thus, **plasma cleaning reduced photothermal background by up to a factor of five!** Additionally, this data indicated that the % standard deviation of background reduced significantly, meaning that the **background got flatter**. The Q-factors of this

¹ Bubble "5E", for our group's records.

microbubble's resonances also improved remarkably, as did the visible appearance of the bubble by the witness camera, as shown in **Figure 5-10**.

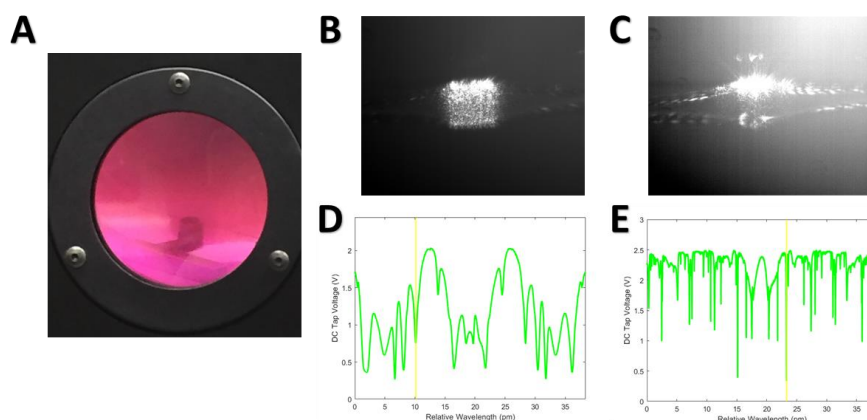


Figure 5-10. Results of plasma cleaning a dusty microbubble. (A) A picture of a microbubble (still epoxied onto an aluminum mount) being plasma-cleaned. (B) An optical micrograph of light coupled into the microbubble before plasma-cleaning. (C) An optical micrograph of light coupled into the microbubble after plasma-cleaning. (D) Resonances in the pre-plasma cleaned bubble, with the Q-factor of the resonance overlapped by the vertical yellow line being $Q \sim 10^6$. (E) Resonances in the pre-plasma cleaned bubble, with the Q-factor of the resonance overlapped by the vertical yellow line being $Q \sim 10^7$.

As the reader may have already made note of, I stated that the above bubble was older and already dusty before attempting plasma cleaning. Thus the above experiment may **overestimate** the reduction in photothermal background for a more pristine bubble. Nonetheless, I strongly suggest plasma cleaning microbubbles in the future if they are to be used for single molecule experiments, and did in fact clean plasma clean microbubbles (at “high” RF power, for ~15 minutes) for the attempted single molecule experiments discussed in *Section 5.11.2*. Another possible benefit of plasma cleaning is that oxygen-plasma should leave the glass wall of the resonator negatively charged, and thus more coulombically favorable for binding an analyte of positive charge (a downside if an analyte of interest has a negative charge).

5.10.5.2. Caustic liquid cleaning

It is quite possible that plasma cleaning will not scourge all surface contaminants from the resonator's glass. We hypothesized that employing a highly caustic solvent clean, as we do for many samples to get them single-molecule clean, might lower photothermal background further beyond that achieved with only plasma cleaning. For this, we decided to use one step of an RCA clean (RCA standing, not usefully, for Radio Corporation of America). A typical RCA clean involves a piranha bath (concentrated H_2SO_4 mixed with concentrated H_2O_2),¹ a basic step (diluted ammonium hydroxide and H_2O_2), and another acid step (HCl and H_2O_2). To start with, we decided to only employ an SC1 clean, which is the basic (high pH) cleaning step. While as far as dust is concerned we are mostly concerned about the exterior of the resonator, it is also worth cleaning the interior of the resonator as well, to see if this lowers photothermal background, which we planned to do via a syringe attachment to pull caustic solution into the capillary.

SAFETY: While less hazardous than piranha solution, the SC1 is still a significant hazard, so safety is paramount when using it! Keep in mind that these solutions are often used to scourge organic material very well from glassware, and you yourself are, in fact, made of much organic material.

To fully submerge the bubble, we anticipated needing to **remove the epoxy** holding the bubbles to the mounts (given that it is an organic material), which is there due to the way in which we have collaborators fabricate and ship the microbubbles. Additionally, we did not know how the caustic solution would interact with the aluminum mount, imbedded magnet, or polymer sheath on the OS microbubble capillaries. It turns out that Norland Optical Adhesive 81 (the epoxy holding the resonator to the mount) is quite solvent-resistant. Leaving a mounted resonator sitting in acetone overnight, for example, did not noticeably degrade the epoxy. However, per a recommendation from the epoxy manufacturer, I found that hot chloroform (55 °C; BP is 61.2 °C) works quite well. Dichloromethane would also likely work, but this solvent is substantially more of a health hazard than chloroform.

¹ Some nasty business, that.

It became evident that the procedure is risky for the fragile bubble and capillary, and we never got the procedure fully vetted, though some initial attempts had promise. However, we have perhaps found a much easier route for caustic microbubble cleaning than the above procedure. With the challenges described above in unmounting the bubble (and then having to have the bubble survive cleaning, and then remounting), we decided to test how the mounted bubble in entirety would respond to SC1. In particular, the materials tested needed to be the epoxy holding the capillary down, the aluminum mount (with magnet embedded), and the polymer-sheath remaining on the untampered capillary regions. Thus, I prepared a solution of SC1 following our lab's Standard Operating Procedure, and tested small pieces of each material type (ex. small piece of cured epoxy, ex. small piece of aluminum foil). Initial tests indicated that the materials are compatible with the SC1 solution, and thus that the full mounted bubble can survive the SC1 clean (**SAFETY: not necessarily other RCA cleans! Don't dunk one of these bubbles-on-mount in piranha!**).

Whether or not the bubble will emerge cleaner than before is unknown. While we have tested the epoxy removal process (which hopefully is not needed), and some brief SC1 materials tests, we have yet to truly try cleaning a microbubble in these conditions and then test its background. Thus, while some bubbles are clean enough not to mandate such cleanings, **such cleanings are worth future pursuit** if microbubbles are to be further explored as a platform for single-molecule studies.

5.10.6. Other mapping experiments

5.10.6.1. Mapping the tapered fiber

Interestingly, microbubble photothermal maps *sometimes* indicate the presence of the tapered optical fiber. What's even more interesting is that sometimes the taper shows up as a higher photothermal signal than the rest of the resonator (a "positive" signal, I call it), and sometimes the taper shows up lower photothermal signal than the rest of the resonator (a "negative" signal). While I can rationalize both things happening, ex. perhaps a negative signal comes from the taper blocking pump light, such that it

can't be absorbed by the resonator, ex. perhaps a positive signal comes from the taper absorbing light itself, and given that the taper is physically touching the microbubble, this heat gets transduced as background. However, I have yet to figure out a hypothesis on why **both things happen**, and why they happen in certain circumstances. Does the mode coupling to effect this? Does this have to do with taper fabrication? Taper thickness? **Future work**, experimental and/or theoretically, into this might be interesting, as it could perhaps help us understand more about both photothermal background, and maybe even the impact on transfer of energies in our system when the taper is in contact with the resonator.

5.10.6.2. Mapping the tapered capillary

Another area that is interesting to photothermally map in microbubble resonator systems is the tapered capillary portions of the device! This is because capillaries (and in fact, tapered optical fibers) also harbor WGM resonances. I have mapped the tapered capillary portions a few times (data not shown) to prove that it can in fact be done. However, nothing interesting enough to include here has arisen. Nonetheless, I suggest that future group members remember that the tapered capillary is **also** a viable resonator for mapping, though (1) the optical modes will likely be very spread-out bottle modes, and (2) the modes will be contained entirely within the glass in those regions, leading to poor mode-analyte overlap.

5.10.7. Changing the focusing objective

5.10.7.1. Overview

As noted by previous group members,¹ perhaps the lowest hanging fruit to making our photothermal absorption experiments single-molecule-ready is using a focusing objective with **higher numerical aperture (N.A.)**. By switching to an objective with a higher N.A., the pump beam can then be focused to a smaller spot size, increasing the intensity of the beam. With an increased pump intensity, a sub-

¹ Especially Kevin Heylman, several times in his thesis.

diffraction limited object will see more photons/time, and thus will absorb more photons and create more photothermal signal (provided that the beam intensity is below the saturation threshold of the target molecule). In contrast, the microresonator overall will continue seeing the same amount of photons, and thus background signal remains the same. Therefore, signal-to-background increases as N.A. of the objective increases! Another way to think about this situation is that if you want a specific amount of signal from a nano-object, increasing the N.A. permits you to use lower pump power to achieve that signal. This can be important, since you may not have enough photons available to reach that signal for a larger spot size. Other ways to shrink the spot size \include overfilling the objective's back aperture, and/or moving to a bluer pump wavelength (though in practice, the pump wavelength chosen will be determined by the analyte under study).

For the data shown in Chapter 4, the focusing objective used was a 40x Nikon Plan with 0.65 N.A. While the use of microbubbles does not permit¹ the use of immersion objectives (ex. a 1.45 NA oil objectives), there are air objectives with substantially higher NA than 0.65. Thus, I tested two higher, air-use NA Nikon objectives (both infinity corrected) that seemed like good candidates based on their specifications: the (1) Plan Apo Lambda 60XC, and the (2) CFI Plan NCG 100X.

The 60x objective has the benefit of being slightly higher NA. It also has a potentially substantial benefit regarding achromaticity, as the Plan Apo is a line of objectives designed to provide fewer chromatic aberrations than the Plan line. However, the 60x objective (the same model that our group's NIR microtoroid experiments have used for years, in fact) has the drawback of having cover glass correction built into the design, which we suspect inhibits us from reaching the smallest theoretical spot size, given that we do not use a coverslip in our experiments. The working distance is also quite short, at 110-210 μm , depending on where the cover glass correction collar is set to. The 100x objective, while it may suffer from enhanced chromatic aberrations, has a much longer working distance (1 mm) and does

¹ As far as I can imagine, anyway.

not have the cover-glass correction. It was not immediately apparent which objective would perform better for us, so I set out to test both of them experimentally.




Figure 5-11. Testing different pump objectives. *Left:* The three objects side by side, from left to right, the 40x objective used in previous experiments, the 60x objective and the 100x objective to be tested. *Right:* The back aperture of the 100x objective is quite small, which is likely related to the long working distance of the objective. This back aperture requires careful consideration of beam size and expansions in the pump beam path.

5.10.7.2. Experiment

How tightly will each objective focus the same laserbeam (i.e. which objective will yield the smallest beam waste at its focal plane)? A good way to measure this beam waste (spot size) is to use the beam to photothermally map a sub-diffraction-limited object¹ by raster-scanning the pump beam. Because the object is smaller than the pump beam, its photothermal map will look the same as the beam's PSF. This PSF can then be fit to a 2D gaussian equation to provide a beam size that can be converted to a full-width half-maximum (FWHM).

A microbubble was photothermally mapped using each of the three objectives (40x, 60x, 100x), and the beam focus was optimized to ensure that the smallest portion of the beam waist was used for mapping. Then, a piece an object (i.e. piece of dust) was found on the resonator and mapped. Notably, the object analyzed with each objective may have been a different object, but as long as the object is

¹ Ex. a small piece of dust, or an intentionally-deposited gold nanorod, if you want to get fancy.

smaller than the diffraction-limited beam (as almost all dust is in our photothermal experiments), this will not impact the resulting analysis.

5.10.7.3. Results

The results of the photothermal tests for each objective are summarized in **Figure 5-12**. Notably, the **100x objective gives the smallest spot size**, with a FWHM diameter of 413 nm for a 532 nm beam.




Figure 5-12. Comparison of beam spot size for three different objectives. Left: The measured beam diameter for each objective. Right: The photothermal maps used to determine the beam diameters (using a 2D-Gaussian fit). The scale bars are different for each map.

How does this result compare to the diffraction limit? Using the equation for the Abbe diffraction limit (**Equation 5-1**), where **d** is the beam diameter, **λ** is the wavelength of light, and NA is the numerical aperture of the objective, one calculates that for this wavelength and NA, the smallest achievable spot size is 295 nm. So, we still seem to be a fair bit away from the best case scenario, but 413 nm is a great improvement on >800 nm, which was the measured spot size for the 40x objective!

Equation 5-1

$$d = \frac{\lambda}{2 * NA}$$

5.10.7.4. Impact on beam scanning

An important consequence of the new objective is its impact on the **galvo calibration parameter**. The acquisition of this parameter, and its use in our codes, is discussed in detail in *Section 3.10.2*. This

calibration parameter was collected for all three objectives. As can be seen in **Figure 5-13**, this parameter scales inversely with objective magnification, as one would likely expect. For the 40x objective, 1 V of change for the galvo mirror would result in 145 μm of change in the focused beam's position, while for the 100x objective, the position would only change by about 58 μm . This factor of 2.5 difference in parameter is the same as the ratio of objective magnifications, which makes good sense.¹




Figure 5-13. Galvo calibration for different objectives.

One interpretation of these results is that higher magnification objectives might provide better mapping resolution, if the limitation on spatial resolution stems from the resolution in voltage change. This is because the higher magnification objective essentially “spreads out” the voltage span needed to achieve a certain spatial span. However, we have never found resolution to be an issue before with the galvo mirrors (typical “fine” scan distance is 100nm/step, which is easily achievable with our current system), so this consideration is not particularly important in our case.

However, there is at least one major drawback to the above results, namely that at higher galvo voltages, the beam is more likely to clip various parts of the beam relay, ultimately **limiting our FOV** to a

¹ It is important to note that the absolute values of these calibration parameters are also a function of the lens parameters in the beam relay system used for this particular experiment. So while the ratio of parameters may hold true for other alignments using different lenses, the absolute values may not.

smaller region of a microresonator. This is demonstrated by the data in **Figure 5-13**, where the FOVs for the 40x and 100x objective are compared for the red pump beam. Also of important note is the fact that the max transmission for the 100x beam was less than half that for the 40x beam! This is in large part due to the back aperture of the 100x objective being so small (**Figure 5-11**), such that it clipped the beam severely. This is of course a remediable situation, but one worth keeping in mind. FOV and transmission issues aside, the higher N.A. objective is likely critical to us reaching single molecule studies, so these drawbacks will have to be accepted.




Figure 5-14. FOV calibrations for the red beam, for two different objectives.

5.10.7.5. Chromatic aberrations

With the M Squared laser system (**Figure 3-13**) now integrated into the photothermal microscope, future experiments will determine how significant chromatic aberrations are for the 100x objective. Various calibrations could be used to account for such chromatic effects, and may indeed be highly necessary for doing quantitative, broadband absorption spectroscopy on single molecules and nanoparticles.

5.10.8. Future work: distance-dependence of background variation

It may prove quite useful to quantitate how materials-based photothermal background varies spatially across a range of different length scales. This is because the **flatness** of the background is crucial when trying to identify objects on top of it. This analysis might take the form of a plot of “average background change” vs. distance, where for each map pixel (the pixels near the edge may need to be truncated), the

background value for that pixel is absolute-value subtracted from every pixel in a ring at a prescribed distance, and all of those calculated values are then averaged to become the value for the new 2D map. Once this is done, all pixels in that map would then be averaged to give a single value for the **expected background variation against that length distance**. The **standard deviation** could also be included as a sort of error bar.¹ This analysis can be done (in an automated fashion, hopefully) for a range of distances, from “1 pixel away” to much greater distances, to produce the plot I initially suggested. This method involves a lot of averaging, and thus still may lose out on significant spatial information, but perhaps it would hold good utility, as it would help identify what the average variation of background would be across length scale.

I have not designed this analysis beyond the conception mentioned here, and there may be much smarter and/or more computationally efficient ways to get this or a similar parameter. However, I do recommend considering the development of such an analysis for future photothermal background characterizations.

It’s also possible that employing the other tools our kit, namely polarization and wavelength, will allow for seeing objects more clearly against a non-flat background. However, this will require taking several maps for different parameters, which significantly increases experiment length, and for single molecules, may increase risk of photobleaching molecules before their dynamics can be studied.

5.11. So, are We Good Enough to Look at Single Molecules?

Considering everything that we have discussed so far regarding photothermal background...are we in a position to succeed in single molecule absorption spectroscopy experiments? As discussed earlier in this chapter, for a single black hole quencher dye, we anticipate a thermal response of **9.7-97 am/mW** for the water modes of a microbubble. And as also discussed in this chapter, a plasma-cleaned, NS bubble

¹ Though whether it is better to take the standard deviation of fully-calculated pixels, or to propagate deviation from each pixel is not clear to me.

gave had around **100 am/mW** of photothermal background when air-filled. Thus, if we take the upper end of our estimated signal, and assume similar photothermal background when the bubble is water or air filled, perhaps we could achieve a SBR of unity. And while an SBR of just over 1 might be usable provided a *very* flat background, this experiment would realistically be overly difficult.

Luckily (though sometimes frustratingly), it turns out that the photothermal background of bubbles can vary quite a lot. A new set of OS microbubbles arrived from collaborators, and one of the bubbles¹ was exceptionally clean! In fact, out-of-the-box, it seemed that instrumental background was the limiting factor, at around **90-120 am**, while the materials background was very low, something like **6-15 am/mW** at 635 nm pump. Thus, the BHQ molecule of interest should give more signal than the materials background! And with the range of 9.7-97 am/mW for the molecule, somewhere between 1 mW and 10 mW of pump light should be enough to overcome instrumental background. Therefore, it was worth trying to do our first single molecule absorption spectroscopy experiment!

5.11.1. The first single molecule attempt

5.11.1.1. Calculations

For this experiment, a 100x 0.9 NA objective (Nikon MRL03901) was used. The pump laser was a 635 nm single wavelength diode (BlueSky FTEC2635-M50PFQ). Given the transmission efficiency of the system, and relatively low starting power of 40 mW, it was expected that only around 5 mW of light at most would reach the single molecule. Thus, getting enough signal to get above instrumental background may have proven difficult, but it was worth attempting the experiment nonetheless.¹¹ BHQ-3 was selected as the molecule for study.

¹ Bubble "2i", for group records.

¹¹ In the more recently built version of the microscope, the high-powered, tunable M Squared laser system will afford *much* higher powers, such that instrumental background is of no concern.

5.11.1.2. Experimental

Prior to deposition, the OS microbubble was air-plasma cleaned for around 20 minutes, still adhered to its aluminum mount. All glassware (namely the beaker used for solvent and the vials needed for the dilution series of the analyte) was rinsed multiple times with HPLC acetone, then HPLC methanol, then HPLC water, then HPLC methanol from a glass bottle (a standard procedure in our group for getting glassware single-molecule clean), and the glassware was also plasma cleaned.

A dilution series of BQH-3 in methanol was made (after dissolving the powder of BQH-3, stored at -20 °C freezer), ending in a solution of approximately 50 fM BQH-3 in methanol. This concentration was chosen because it is about the same as the estimated concentration of AuNRs as used in our gold nanorod depositions (see Chapter 4). Of course nanorods and single molecules are vastly different in size and composition, charge, surfactant presence, etc., but this seemed like a reasonable starting concentration.

Estimating a volume of 1.5 nL in the microbubble (assuming a cylinder 100 um wide and 200 um long), around 100 molecules should be in the resonator at any given point, assuming that a significant percentage have not been lost by being bound to the capillary wall on their way to the microbubble, or in the vial holding the molecules. This could lead to a huge amount of molecules binding (making a chromophore carpet, so to speak), or perhaps none binding...it all depends on what the wall-molecule-solution interactions are, which is difficult to predict *a priori*.

After the solution was flowed through for several “drops” exiting the opposing side of the capillary, the microbubble was purged with milliQ water. The water backfill had three motivations (1) hopefully the somewhat hydrophobic BQH-3 molecules, if already adhered to the resonator inner wall, would stay in place with water as the liquid in the resonator, (2) milliQ water is the cleanest solvent (3) water is what all of our other solvent-filled microbubble experiments had previously used.

5.11.1.3. Results

Unfortunately, no single molecules were observed in the inaugural initial attempt. This may be due to significant resonance instabilities, which I hypothesize arose with the water-filled resonator being somehow more acoustically sensitive. It is also very possible that no molecules deposited in the resonator! Based on the above calculations, it is also conceivable that there was not enough laser power to overcome instrumental background.

5.11.2. Additional single molecule attempts

Additional single molecule experiments were attempted, using both BHQ-3, as well as a different molecule, SY-7 (Fluoroprobes). SY-7 studies were attempted using 561 nm excitation light, as this should give a factor of about two enhancement in absorption relative to BHQ-3 with 635 nm light. However, issues with the fragility of the remaining microbubble supply and other challenges have so far hampered further progress in these experiments.

5.12. Fundamental Materials Absorption

It is worth considering what the lowest resonance shift is we can possibly expect to achieve for a water-filled microbubble. For this calculation, we simply need to consider the absorption coefficients for both water and glass, as well as the corresponding path length experienced by the laser beam through each material (as we will assume a linear absorption relationship in this situation, which should be valid given the very small % of light that will be absorbed).

Resonance shift from glass: Finding the absorption coefficient for glass in the visible wavelength regime is actually a bit tricky, given that glass is so transparent that its absorption of light is very often approximated to be zero. After considering a few approaches (such as extrapolating from fiber optic losses), I pulled a value from a paper published by Pinnow and coworkers in 1973¹⁹⁷. In their paper, the

¹ I was directed to this paper by the thesis of Dr. Kevin Heylman.

value for “total intrinsic loss” of fused silica at 635 nm wavelength is approximately **6 dB/km**. If we then take an overestimate that the wall thickness of a microbubble is **1 μm** (1×10^{-9} km), and assume the full laser beam passes through the wall once on each side, we calculate an expected loss of **1.2×10^{-8} dB**, or 2.76 parts-per-billion. This corresponds to **2.76 pW** of heat for every **1 mW** of pump light. Given our anticipated value for microbubbles of 4 fm resonance shift per 10 nW of heat output, we calculate an expected resonance shift of **1.1 am/mW_{pump}** due to glass absorption.

Resonance shift from water: For our starting calculation, we will use an absorbance coefficient of 0.3 m^{-1} for water at 635 nm wavelength.¹⁹⁸ For our initial calculation, let us assume a pathlength the diameter of an average bubble, 100 μm (which we will discuss as a questionable assumption shortly below). Using **Equation 5-2**, where **L** is the pathlength and **a** is the absorbance coefficient, we calculate that **P**, the percentage of light absorbed, is **0.003%**.

Equation 5-2

$$P = 100 * \left(1 - \frac{1}{e^{L*a}}\right)$$

This corresponds to **30 nW** of heat output (assuming all absorbed light goes into heat) for every **1 mW** of incident light. That sounds very troublesome, given our previous calculations for a water-filled microbubble yielded an anticipated resonances shift **0.4 fm/nW_{heat}**, this means that we would anticipate something like **12 fm/mW_{pump}**, which is far above the **10-100 am/mW_{pump}** we expect from a single BHQ3 molecule!

Have we observed such a drastic resonance shift in practice? Let us look to some experimental data, specifically photothermal background data for the same microbubbles, air-filled and then water-filled.¹ These measurements were done with at least 3 mW of pump power at the resonator surface, so given the above calculation, we would expect a water-absorbance resonance shift of at least 36 fm. In

¹ Microbubble 1J, for internal group records.

Figure 5-15, resulting histograms are shown for four conditions, namely the instrumental background (“pump off”) for both air-filled and water-filled, and “pump on” maps for both conditions. The regions mapped for the pump on data are not necessarily the same regions, but are representative of the bubbles overall.

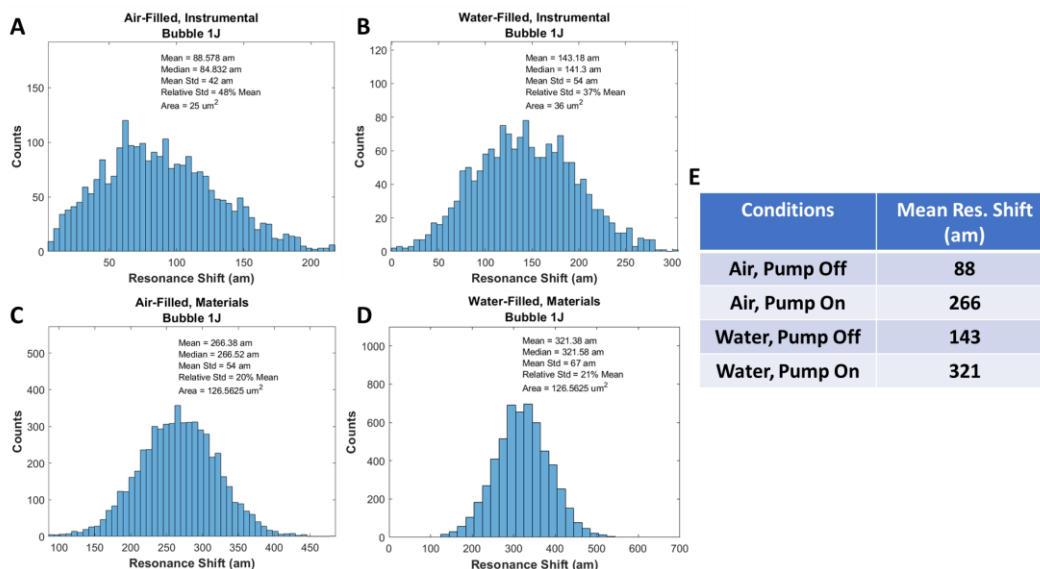


Figure 5-15. Comparing the instrumental and materials background for an air- vs. water-filled bubble. (A) Pump Off data for air-filled bubble. (B) Pump Off data for the same bubble when water-filled. (C) Pump On data for the air-filled bubble, for a region. (D) Pump On data for the same bubble when water-filled. (E) Summary of mean values from the four measurements. Pump on is total resonance shift measured, and does not include background subtraction of the Pump Off data.

It appears that the instrumental background (“pump off”) increased from 88 am to 143 am. This is perhaps due to the presence of the water changing vibrational noise transduction, or perhaps just because the lock was less stable; it would be hard to know. However, if we background subtract the pump off data from the pump on data, we get a materials background value of 178 am for **both** the water and air conditions, which used the same pump power! Thus, it seems like we observe **no additional background** for having the water present.

So, it looks like something must be quite wrong with our above calculation for anticipated resonance shift due to water absorption. Barring there being something fundamentally wrong with the

calculation,¹ there are a couple of suspicions that arise about that naïve calculation. The first is the fact that we used the entire diameter of the bubble to calculate pump-length. Our group has documented both experimentally and computationally (via COMSOL) that objects need to be **near the optical mode** to have significant impact on the resonances. If we lower our estimate of path length for water to something more like **2 μm** (1 μm on each side, around where the mode resides), this lowers our expected resonance shift to only 1.2 fm/mW_{pump}. This is still very high, and should have been observed in the above data!

An additional problem with the 100 μm pathlength originally used in calculation is that it ignores the divergence of the pump beam. For a crude calculation, if we consider that the 100x objective used has a working distance of 1 mm, and we assume a beam of 2 mm radius at the beginning of focusing, then the angle of convergence is about 62 degrees. At an angle of divergence of 62 degrees, a beam focused to one side of the resonator will be almost 190 μm wide by the other side of the resonator, with much light escaping before traversing a full 100 μm . Thus, Gaussian optics should probably be consulted to determine a realistic average path length experienced by the pump light.

For now, luckily, the addition of water into our resonators does *not* seem to increase photothermal background noticeably. Further work needs to be done, perhaps with COMSOL simulations, to try and prescribe a more authentic expectation, which hopefully matches experimental results, for photothermal background due to water-absorption in microbubbles.

5.13. Getting to Single Molecules in the Future

There are a few particularly fruitful routes I see in the future for finally getting to single molecules with our photothermal microresonator absorption spectroscopy platform. These include:

¹ I don't think there is, but it wouldn't hurt for the reader to do this calculation entirely on their own.

- (1) Employing the new M Squared laser system, which will provide both high power and many wavelengths to study both analytes and background with.
- (2) Trying out different microresonator platforms, such as ring resonators.
- (3) Moving to a higher N.A. objective, such as an oil-immersion objective, should the proper resonator/substrate permit this.
- (4) Improving further electronic sources of noise for reducing instrumental background.
- (5) Moving to higher modulation speeds for resonator platforms that have a better heat sink, which will move further away from $1/f$ noise.
- (6) Continuing to explore caustic cleaning steps for resonators, and perhaps other RCA components or a hydrofluoric acid dip. These steps might also be useful to try on substrates prior to resonator fabrication (Ex. on the pre-tapered capillary before microbubbles are fabricated).

5.14. Conclusion

The combined reduction of instrumental background via lock-in amplifier optimization, and the reduction of materials background through plasma cleaning have gotten us tantalizingly close to what calculations indicate is the single-molecule absorption spectroscopy threshold. While we have yet to succeed in single-molecule experiments with microbubble resonators, there are promising routes forward for further lowering photothermal background in both this and other resonator platforms.

6. Photoluminescence Studies of Single Carbon Dots¹

6.1. Motivation

So far in this thesis, we have almost exclusively discussed photothermal microresonator absorption spectroscopy. This chapter, however, is an ongoing story about photoluminescence! Why did we suddenly change topics?

The initial motivation for the photothermal microresonator method was to develop a new tool for studying nano-objects that are *not* emissive. However, there are many nanoparticles and molecules out there which *are* emissive which could still be usefully examine using an absorption spectroscopy method. In particular, being able to measure *both* absorption and fluorescence information, simultaneously on single nanoparticles or molecules, would be a fantastic way to learn deeply about the photophysics of excitonic systems. Doing such a simultaneous measurement is currently out of the question for a single fluorescence dye molecule, as regardless of the percentage of absorbed energy that goes into the non-radiative decay channel, the prominence of the radiative decay channel will raise the average excited state lifetime by orders of magnitude, effectively throttling the ability to do a photothermal measurement. However, doing such a simultaneous measurement on a **nanoparticle**, with significantly higher absorption cross-section than a molecule, could certainly be feasible. Thus, we have been exploring how to augment our photothermal measurements with simultaneous fluorescence detection.

¹ This chapter contains work in an ongoing investigation that will soon be prepared into a manuscript. While they are not yet co-authors (since there has yet to be a submitted manuscript), I want to thank Lisa-Maria Needham for her vast help in sample preparation, data acquisition, and discussions. The carbon dots in this chapter were synthesized by Riley Lewis, in the group of Christy Haynes at the University of Minnesota-Twin Cities. This work has been funded largely through the *Center for Sustainable Nanotechnology*, in which this collaboration arose.

For a first target, we selected an intriguing, relatively new class of nanoparticle, Carbon Dots (CDs).¹ However, as will be described in detail below, there are still many outstanding questions about carbon dots that fluorescence-only studies could answer. Thus, learning to walk before we run, we set out on a photoluminescence exploration of carbon dots.

6.2. Introduction to Carbon Dots

Carbon Dots are a relatively new and exciting class of emissive nanoparticle, heralded as a possible replacement for inorganic semiconductor quantum dots (QDs). While QDs typically contain heavy metals that may pose an environmental hazard once QD-containing objects are disposed of, it is anticipated that CDs may be much more environmentally benign. The toxicity of CDs and their degradation products in the environment is an active area of research.¹⁹⁹⁻²⁰⁰

While there were a few earlier reports regarding luminescence of carbon fragments, “carbon dots” were first reported by Sun and coworkers in 2006.²⁰¹ Their dots were produced using laser ablation of a graphene source, producing CDs ~5 nm in diameter, and interestingly were only strongly luminescence after post-synthetic modification with PEG (poly-(ethylene glycol)). Since then, the field of CDs has vigorously proliferated, with many different kinds of synthetic strategies employed for both top-down and bottom-up synthesis. While the astounding number of synthesis routes to reach carbon dots speaks to the variety of applications they may find, it also makes it quite challenging to draw general conclusions about carbon dots, given that so many studies follow their own synthetic recipe. Even within microwave-assisted solvothermal syntheses (the type of syntheses most relevant to this chapter), a variety of precursor materials, reaction temperatures, pressures, and times have resulted in carbon dots of varying sizes and emissive properties. This indicates that even for one classification of dots, multiple chemical groups and photophysical mechanisms may be achieved via small changes in synthetic

¹ My apologies to the circular dichroism folks who may not be a fan of their go-to acronym being co-opted, but alas this is the widely accepted acronym for this class of nanoparticles.

procedure. Some efforts have been made to correlate photophysical parameters with synthesis routes,²⁰² and there are a variety of review articles from throughout recent years summarizing the many different synthetic routes, chemical characterizations, and photophysical studies that have been accomplished.²⁰³⁻²⁰⁶ Despite all of the above efforts, there is still uncertainty for many types of CDs as to what chemical and physical features govern their photophysical behavior.

In the last couple of years, a particular subclass of CDs, “multicolor polymeric carbon dots”, has been quickly gaining attention.²⁰⁷ One exciting thing about this subclass of CDs is that the one-pot microwave-assisted synthesis employed simultaneously generates carbon dot subpopulations that emit anywhere from violet to red light. Through extensive experimental and computational work, Zhi and coworkers evaluated their multicolor carbon dots (the same type used in this work) for three fluorescence mechanisms, namely quantum size effects, molecular fluorophores, and surface state defects. They determined that the blue fluorescent CDs likely have emission arising from small molecules trapped in carbonaceous cages within the CDs, while other wavelengths of emission, especially red emission, likely originate from pyrene analogs.²⁰⁷ Even given these insights, there is still much to learn about the chemical and photophysical properties of these CDs, thus motivating the work in this chapter.

6.3. Background

In contrast to inorganic quantum dots, CD emission often **is not dependent on particle size**, though there are exceptions, and this depends on the type of CD being studied. In the literature, a three-class classification has been suggested for CDs, listed below with some relevant characteristics.²⁰⁴

1. **Graphene Quantum Dots:** Disc-shaped, such that while they may be a several nanometers across, the thickness is only a few to several layers of graphene. Synthesized using a top-down approach (such as laser ablation of a carbon source, for example).
2. **Carbon Nanodots:** Spheroid-shaped, and containing ordered sp^2 carbon domains within. Can be generated from top-down or bottom-up synthetic strategies.

3. **Polymer Dots:** Spheroid-shaped, containing no ordered sp^2 domains. Generated via reaction of small molecules into polymers with cross-linked backbones. Synthesized with bottom-up synthetic strategies.

Broadly speaking, there are three different hypothesized mechanisms for carbon dot photoluminescence. Any combination of these might be used to explain the emissive properties for a certain class or subclass of CDs:

1. **Surface state:** Chemical defects on the surface of the nanoparticles govern the HOMO-LUMO gap of the lowest energy states of the nanoparticle, which in turn govern emission wavelength.
2. **Carbon core state:** A crystalline core is responsible for the emissive properties. If this is the dominant emission mechanism for a type of carbon dot, it would be expected that there could be size-dependence of emission, though notably it is the size of sp^2 carbon domains that would matter, and not necessarily the overall diameter of the CD.
3. **Molecule state:** This emission mechanism would be due to chromophores/fluorophores embedded within the carbon dot. Different molecules and their surrounding nano-environments would govern the color of emission.

So how are we going to go about studying our multicolor polymer CDs, given the above information? While previous bulk experiments (and simulations) have indeed provided deep insights into the structure and behavior of these CDs,²⁰⁷ the fact is that some information is only going to be gleaned via single-particle studies.

One example of where bulk studies fall short is that there have been reports of “excitation-dependent emission” of carbon dots. Such behavior at the single CD level would indeed be very interesting, as usually with fluorescence, the excitation wavelength does not change the emission spectrum, but rather the amplitude of emission. However, almost all such reports of excitation-

dependent emission for CDs were for **bulk** samples (as an aside, there is one report that interestingly saw a small fraction of their CDs emitting at two wavelengths for a single excitation wavelength²⁰⁸). Thus, subpopulations of CDs in-ensemble are likely responsible for the observed emission behavior. In this particular example, it is only fair to note that separation of samples through various chromatographic methods can perhaps make the bulk studies more relevant to the question of excitation-dependent emission. However, it still stands that the most thorough way to identify if subpopulations exist with an ensemble is to perform single-particle measurements.

Now that we have decided on single-particle methods, what sorts of behavior should we look for? There are a few standard parameters one might choose to study when looking at single emissive nanoparticles:

1. **Blinking behavior:** Which carries information about the transitions between fluorescence and non-fluorescence states, and can be used to glean photophysical information.
2. **Photobleaching behavior:** The nature of which, ex. digital (single-step) or analog (steady decline) will yield information about how the nanoparticle absorbs light.
3. **Excitation polarization dependence:** Which will also inform on if the nanoparticles are comprised of one or multiple absorbers.
4. **Emission polarization dependence:** Which will inform on if there are one or more emitters in the particle. The relative alignment of the emission transition dipole angle to the excitation transition dipole angle (if there is one) also can provide information about energy transfer within the particle.
5. **Single CD emission spectra:** Which may lead to understanding the chemical nature of the nanoparticle.

Not included in the list above are fluorescence lifetime measurements, as those will not be explored in this chapter, though they can also be of value in learning about photophysics. Lastly, I will

point out that a few ultrafast spectroscopic techniques have also been applied for studying carbon dots of a variety of class and emission properties in the past, though most studies (save one example discussed briefly below). have been done on bulk samples.²⁰⁹⁻²¹¹ Next, we will discuss the variety of already-existing reports on single CD experiments.

6.4. Previous Single CD Studies

There have been a number of previous reports that examined single carbon dots, with the first example to my knowledge being published in 2013 by Xu and coworkers.²¹² Single CD studies have also been performed at cryogenic temperatures, which can have utility in elucidating photophysics.²¹²⁻²¹³ And in the vein of applications, the blinking characteristics of certain CDs has proven utility in *in vivo* super-resolution imaging,²¹⁴ and other CDs have found use in super-resolution fluorescence methods such as STED (stimulated emission depletion) microscopy²¹⁵ and SOFI (super-resolution optical fluctuation bioimaging).²¹⁶ In a more recent report than many of the above, Nguyen and coworkers achieved sub-particle resolution in single CD studies using single-molecule absorption detected by scanning tunneling microscopy (SMA-STM). They determined that the large bandgap of graphitic cores in their CDs was inconsistent with fluorescence wavelengths, which could instead be explained by local oxygen defects.²¹⁷ Even more recently, that same set of authors reported a technique called time-resolved SMA-STM (trSMA-SMT), an ultrafast technique, to study the energy flow within and between single carbon dots.²¹⁸

However, similar to the aforementioned bulk studies, there is great variability in what carbon dots have been studied at the single CD level, as well as the resulting conclusions drawn. For example, while some groups have reported CDs with both blinking and single-step photobleaching,^{208, 213, 219-220} others have observed CDs without any blinking behavior and without single-step photobleaching.^{212, 221} In yet another report, multi-state blinking was observed!²²² This particular report suggested that a Dexter-type electron transfer between surface groups on the CDs is responsible for the switching from a fluorescence state to a dark or gray state, while they claim that the temperature-independence of the

rate of dark-to-bright states (CDs “turning back on”) can be attributed to electron tunneling from the CD core to its surface. And in another interesting study by Fathi and coworkers in 2019, both single-step and multi-step photobleaching were observed for CDs.²⁰²

While blinking of CDs is important to understand their photophysics, it can also be deleterious for making extended measurements. Recently, Wei and coworkers published an exciting report detailing how with the addition of Trolox (6-hydroxy-2,5,7,8-tetramethylchroman-2-carboxylic acid), they were able to entirely arrest the blinking of single CDs.²²⁰ Furthermore, they observed a seven-fold increase in fluorescence time before photobleaching, as well as an increase in fluorescence intensity. Based on this work, we have also employed Trolox for experiments in this chapter.

Two very relevant optical parameters to inspect are the excitation transition dipole (or lack thereof) and the emission transition dipole. For example, a carbon dot acting as a single absorber would be expected to have a well-defined excitation transmission dipole, while a carbon dot comprising many absorbers acting independently would lead to a poorly defined angle for the excitation dipole. Additionally, if the excitation and emission dipoles are parallel or not, and if they are constant or not through the lifetime of a CD are also interesting questions, which may help answer how energy transfer occurs within single nanoparticles.^{213, 219} To date, there have been a few examples of studying such dipoles in CDs using azimuthally polarized light. These examples notably all studied rather different types of CDs, which both adds the challenges of comparing across reports, but hints at photophysical properties that may be broadly true across multiple types of CDs. More specifically, these polarization studies (all from the same research group) have been applied to learn about electron-phonon coupling and its impact on CD photoluminescence,²¹⁹ the aforementioned Dexter-energy transfer study,²¹³ as well as to study carbon dots inside a polymer matrix.²²² Given the above, measuring polarization-dependence of excitation and emission for our CDs will be worthwhile.

A particularly informative piece of data one can acquire from single nanoparticles is their emission spectra. A few of the aforementioned reports have done so,^{213, 219, 221-222} including monitoring of spectra over time.²²² Thus, emission spectra will also be a parameter we want to measure in this study.

Given all of the above, it is apparent that there have already been several reported single CD studies. However, the large variety of synthesis methods employed, and characteristics of dots produced makes it impossible to tell a cohesive story that can explain all observed photophysics across the different classes of CDs. Additionally, many of the above papers are for graphene quantum dots in-particular. And while such reports are interesting, and will have bearing on our understanding the whole landscape of CD photophysics, these reports have limited applicability to the CDs we are most interested in, namely multicolor polymeric CDs. To date, we know of **no reports on the polarization behavior or spectral characteristics of single multicolor polymeric carbon dots**. Given the wide range of potential applications for this subclass of CDs, and outstanding questions regarding their photophysics and chemical nature, there is strong impetus for single particle studies on these nanoparticles.

6.5. Results and Discussion

The CDs were synthesized as previously reported, namely through a microwave-assisted solvothermal treatment of citric acid and urea in formamide by our collaborators.²⁰⁷ Samples of our multicolor carbon dots are shown in **Figure 6-1A**, where the “unpure” sample is as-synthesized, and the “pure” sample has been dialyzed in water for three days. **Figure 6-1B** shows a spuncoat unpure sample (intentionally spuncoat at high concentration) being excited at 488 nm, with a carpet of CDs emitting yellow light. **Figure 6-1C** shows a similar sample, but being excited at 561 nm, with a carpet of CDs emitting red light. These CD-films won’t provide us any single-particle insights, but do confirm that we have fluorescent species for both of these excitation wavelengths in our samples.[!]

[!] Also, they look really pretty!




Figure 6-1. Carbon dot samples. (A) To the left is shown what we call the “unpure” sample, which is red-brown-colored and very viscous. To the right is the “pure” sample (dialyzed), which is blue and much less viscous and residue-leaving than the unpure sample. (B) A picture of a layer of diluted, spuncoat unpure carbon dot sample on a coverglass, being excited with 488 nm light. (C) A picture of a similar sample, but being excited with 561 nm light. Scales bars are excluded as exact distance is not known (due to the pictures being acquired with a cell phone camera).

I also obtained bulk data for the unpure and pure CD samples, to ensure that they behaved as expected and as reported in the literature, before going too far with single particle studies. A subset of the results are shown **Figure 6-2** in as examples. We saw the anticipated emission behavior for our CD samples.




Figure 6-2. Bulk data for CDs. Some representative bulk data for the CD samples above. (A) UV-Vis measurements on both pure and unpure CD samples (at different concentrations for the two samples shown). (B) Spectrofluorimetry measurements for the pure CD sample, but at two different excitation wavelengths. The peaks on the left are excitation light making it through the monochromator. (C) An example Excitation-Emission Matrix (EEM) for a Pure CD sample. Given that our sample is a mixture of CD fractions, excitation-dependent emission is

not surprising to observe. Additionally, artefacts are present (namely the slanted lines of signal) due to excitation light leaking into the emission path of the monochromator.

At one point in our experiments, the “pure” CD sample seemed to have gone bad, as diagnosed by reduced fluorescence, and what looked like possible aggregates of fluorescence material when spuncoat (data not shown). Thus, we obtained a newly synthesized sample from our collaborates and dialyzed it in-house per their procedure. The new “pure” sample was then spuncoat at different concentrations, with some illustrative results shown in **Figure 6-3** that indicate the purification was a success. The degradation, particularly of the “pure” sample, matches with expectations based on conversations with our collaborators, and future effort should be made to both keep samples fresh, and to **diagnose for how long they can be used** before sample degradation and/or aggregation sets in.




Figure 6-3. Spuncoat CD Dilution Series. Freshly dialyzed “pure” CD sample spuncoat from three different orders of magnitude of concentration in water (data here is air-exposed CDs). The density of bright spots scales as expected.

We have observed that many of the carbon dots, particularly when exposed to air, blink quite rapidly! Shown in **Figure 6-4** is a subset of blinking trajectory analysis, at a rather long acquisition speed of 200 ms/frame to average over some of the fastest blinking. Nonetheless, significant heterogeneity in blinking dynamics between nanoparticles is still evident. Additionally, both digital and analog photobleaching behavior is observed, which is quite interesting.

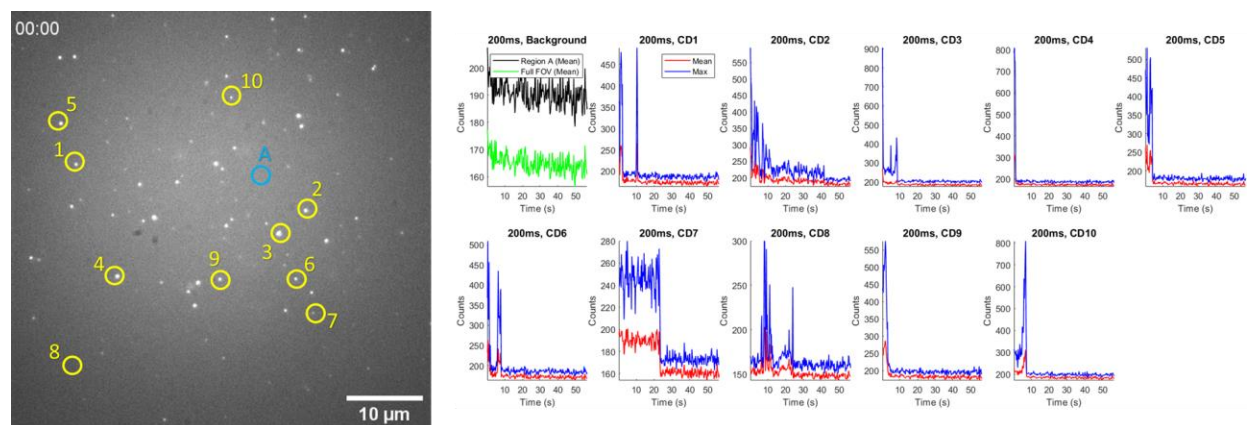


Figure 6-4. Fluorescence trajectories for single CDs. Shown are trajectories (taken at 200 ms acquisition speed) for several single CDs, as compared to background. There is significant heterogeneity in the traces observed, including both analog and digital bleaching behavior.

With several blinking trajectories analyzed, we decide to next attempt a few excitation-polarization studies. These were performed via the HWP-rotating method, for which calibration data and a microscope image are shown in **Figure 6-5**.

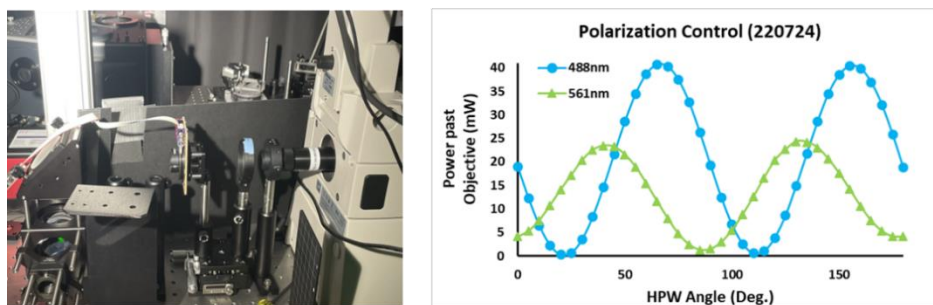


Figure 6-5. Calibration of HWP for excitation polarization studies. (Left) The HWP automated rotation mount is followed by the widefield lens, and next by a calibration polarizer. (Right) The extinction results for the 561 nm laser beam (Coherent Sapphire) and the 488 nm beam (LaserTack). The calculated extinction ratio for the green beam was 18, and for the blue beam was 243. Further improvements to extinction ratios are achievable through alignment adjustments, if needed.

Representative data from the first set of extinction polarization data is shown in **Figure 6-6**. Notably, some of the carbon dots display a deep depth-of-modulation, which is indicative of single excitation transition dipoles! This could be expected for single-absorbers, which is what we anticipate are the objects displaying the digital blinking and photobleaching behavior observed for some CDs earlier

in our experiments. However, we would **not** expect such excitation polarization behavior for multi-absorber objects, which is what we hypothesize the analog-bleaching traces arise from. Thus, more polarization experiments were in order.

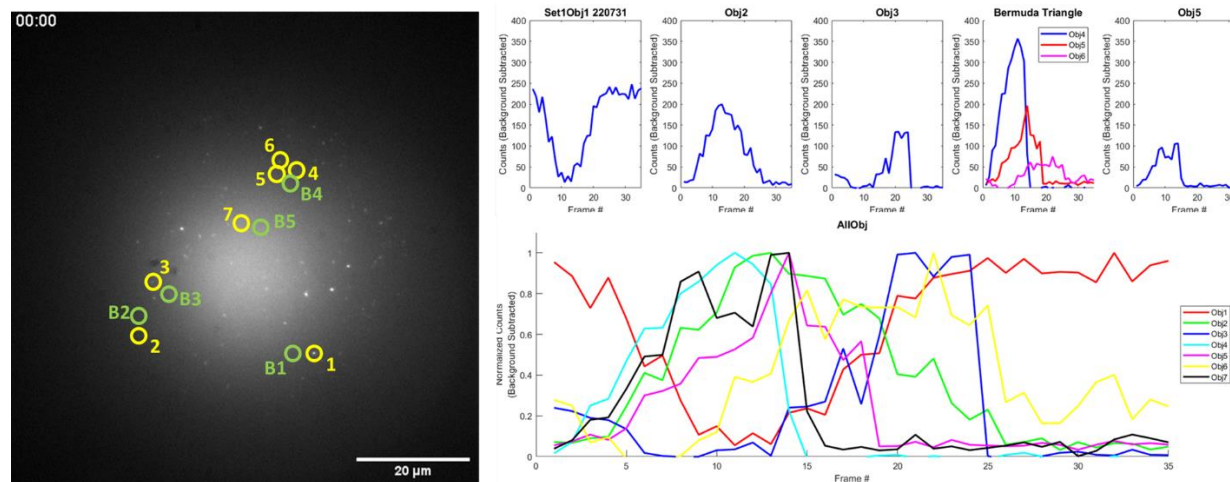


Figure 6-6. Excitation polarization results for one set of CDs using the rotating HWP method. A variety of interesting behaviors are observed. Notably, some objects do appear to possess good depth-of-modulation, which is indicative of a well-defined excitation transition dipole. The data traces show emission counts plotted against “frames”, as the nature of the rotation mount we are using makes correlating frame directly to HWP angle challenging. Future, more refined analyses will require this conversion.

For the next round of excitation polarization experiments, I wanted to acquire data for a longer period of time, and also study how the polarization traces evolved for objects that slowly photobleached in an analog fashion. The designed apparatus for doing this performs four “polarization sweeps” (HWP rotations), with a beam-shutter movement in-between to mark time stamps (and allow for rotation stage homing). This apparatus is pictured and further described in **Figure 6-7**.

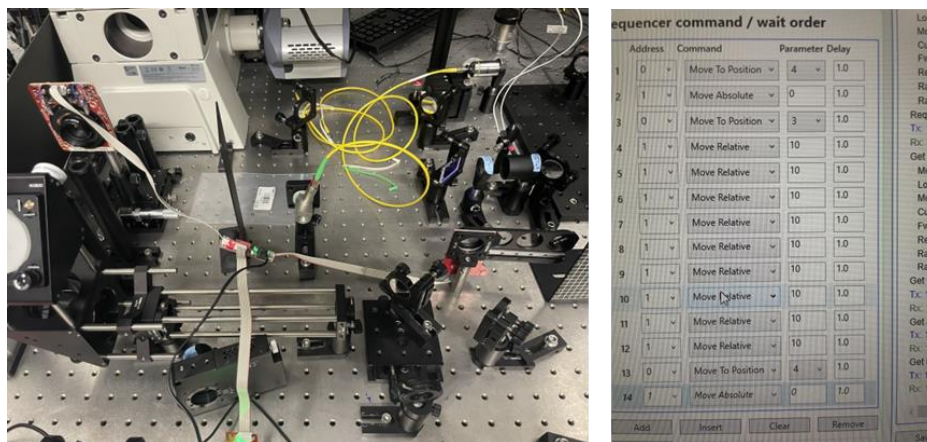


Figure 6-7. The quadruple polarization experiment. *(Left)* To acquire excitation polarization information, a piezo-driven rotation mount is employed. It is connected via a USB bus, which also connects to the piezo-driven shutter. *(Right)* The USB bus allows for the Elliptic software GUI (Thorlabs) to run both devices, thus automating shuttering and unshuttering of the beam, polarization steps, homing of devices, and delay times.

In addition to the above instrumental modification, the arresting of fluorescence blinking in our experiments would be useful, and even more importantly, **extending the lifetime of the CDs**, which seem to bleach rather quickly, would have great utility. In particular, the extended CD lifetimes would be of great assistance in expanding our budget of photons enough to acquire CD spectra, which will disperse the photons over many camera pixels. Based on literature precedence discussed earlier in this chapter,²²⁰ we decided to try the addition of Trolox to our solution to see if we could arrest blinking and extend the length of emission time before photobleaching for single CDs.

The first step in this process was to make sure that CDs would stay adhered to the coverglass surface when water was added. This initially turned out to be challenging, as CDs were observed to immediately desorb from plasma-cleaned glass upon water addition, which is perhaps unsurprising since they were spuncoat from water in the first place. Thus, we applied a poly-lysine treatment, where 0.01% poly-L-lysine (Millipore Sigma) is incubated on top of plasma-cleaned coverslips for at least 15 minutes, removed and then rinsed. This worked, as carbon dots now would stick to the surface and remain there after being “dropcast” onto the surface (where the CD solution was not allowed to dry, but incubated for

10 minutes, and subsequently removed and replaced with supernatant water). The typical Pure CD concentration used for dropcasting has been a stock solution diluted by a factor of 10,000.

With water-exposed CD deposition solved, we then set out to add Trolox to the solution. After several attempts, results are **mixed but promising**. One difficulty is that Trolox readily seems to crash out or crystallize out of the 20% ethanol/ 80% water solution, even at a concentration of 4 mM Trolox. In future attempts, adjusting the pH to a higher value will help keep Trolox dissolved. Even with the above difficulties, the fact is that we have **successfully observed arresting of blinking** during at least a few different experiments, presumably due to the presence of Trolox.

Additionally, some preliminary data hinted that **excitation polarization-dependence might be different for CDs exposed to Trolox vs. air or only water**. Thus, we wanted to see explore this phenomenon further! More specifically, I wanted to measure the excitation polarization behavior under the following four conditions:

1. Spuncoat CDs, from water, air-exposed on microscope
2. Spuncoat CDs, from 4mM Trolox dilution, air-exposed on microscope
3. Dropcast CDs, from water, water-exposed on microscope
4. Dropcast CDs, from 4mM Trolox dilution, 4mM Trolox-exposed on microscope

Single frames from the quadruple excitation polarization experiment described above are shown in **Figure 6-8** for each of the four conditions. In each image, 10 objects are marked with circles and labeled with a number. These numbers will correspond to data in following single-CD trajectories.




Figure 6-8. Single frames from movies for Trolox experiments. (A) CDs spuncoat onto a coverslip from water, and exposed to air. (B) CDs spuncoat onto a coverslip from a solution of 4 mM Trolox in 20% ethanol/ 80% water, and exposed to air. (C) CDs dropcast onto a poly-lysine treated coverglass, diluted in water and exposed to water. (D) CDs dropcast onto a poly-lysine treated coverglass, diluted in Trolox solution and exposed to the same solution.

The next four figures show data for the forty objects selected above, with each figure corresponding to one of the images (i.e. one of the Trolox/no-Trolox conditions) as noted in the captions. It is worth noting that the “plateaus” observed in many of the traces may be artefacts of the background-subtraction method described in the figure captions. Thus, a more robust background subtraction method will be needed for final analysis of the data, however for this initial data analysis (mostly to look at depth of modulation trends, or lack there-of), these figures suffice.

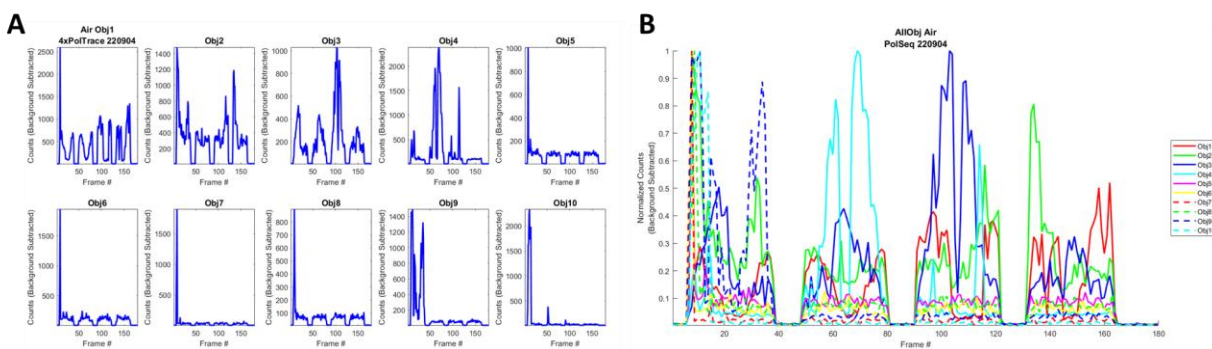


Figure 6-9. Excitation polarization results for CDs in air, spuncoat from water. (A) Quadruple polarization traces for 10 different selected objects (identified in Figure 6-8A). Traces, which are the maximum signal of the object ROI

by frame, have been background subtracted by subtracting the average intensity of a nearby background ROI. **(B)** All ten of the traces from **(A)**, but each normalized to itself.

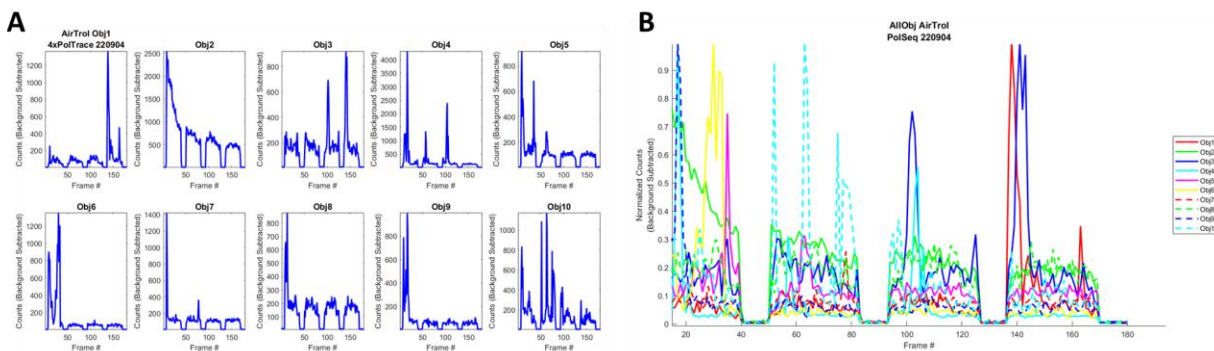


Figure 6-10. Excitation polarization results for CDs in air, spuncoat from Trolox solution. **(A)** Quadruple polarization traces for 10 different selected objects (identified in Figure 6-8B). Traces, which are the maximum signal of the object ROI by frame, have been background subtracted by subtracting the average intensity of a nearby background ROI. **(B)** All ten of the traces from **(A)**, but each normalized to itself.

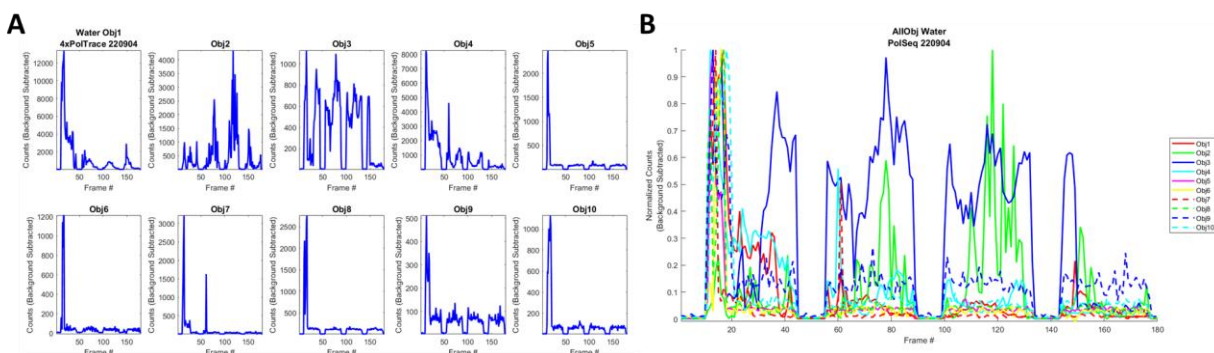


Figure 6-11. Excitation polarization results for CDs in water. **(A)** Quadruple polarization traces for 10 different selected objects (identified in Figure 6-8C). Traces, which are the maximum signal of the object ROI by frame,

have been background subtracted by subtracting the average intensity of a nearby background ROI. (B) All ten of the traces from (A), but each normalized to itself.

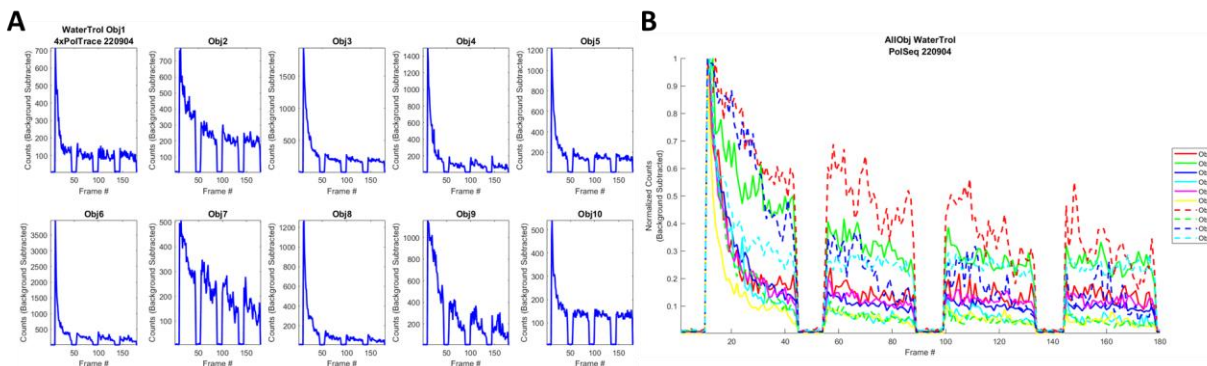


Figure 6-12. Excitation polarization results for CDs in Trolox solution. (A) Quadruple polarization traces for 10 different selected objects (identified in Figure 6-8D). Traces, which are the maximum signal of the object ROI by frame, have been background subtracted by subtracting the average intensity of a nearby background ROI. (B) All ten of the traces from (A), but each normalized to itself.

Further analysis of the data from the previous four figures (and the much larger data set of objects not selected for the representations here) is **ongoing**, and thus firm conclusions cannot yet be drawn. However, two things are immediately noticeable:

- There are some conspicuous, at least partial depth-of-modulation traces in the non-Trolox cases. The fluorescence intermittency (blinking) throughout the measurement will have to be factored into further analysis. Additionally, the fact that four of the ten objects analyzed were selected intentionally for their longevity (time till photobleaching) from the overall sample may bias this result.
- There is **no modulation observed in the Trolox-solution case**, and additionally, many of the objects show **prolonged, analog bleaching**. While this has also been observed for some air-exposed, non-Trolox-treated CDs, it is not the typical observation as with the Trolox-solution CDs. Possible hypotheses for this result are discussed later in this chapter.

The last two pieces of data we want to acquire for single CDs are (1) emission polarization data, and (2) emission spectra. Issues with sample degradation have led to challenges in reliably obtaining deposited single CDs, thus we do not have any emission polarization data yet. However, I have taken a small quantity of spectral data, an example of which is shown in **Figure 4-13**. The data shown here may or may not be on single carbon dots (this was for a sample that had degraded somewhat), thus I have not attempted to extract emission spectra for the objects in the images. The important take-away here is that **0th order information** is needed **in addition to 1st order** diffraction information, so that the relative emission spectra (relative due to the different starting carbon dot spatial positions) can be shifted to absolute values of wavelength.




Figure 6-13. Carbon Dot spectral acquisition. (*Left*) Using a custom lens relay, the iXon camera, and the Kymera spectrograph set to “0 nm”, the zeroth diffraction order is collected a sample of CDs. (*Right*) For the same field of view, data is collected with the Kymera spectrograph grating centered at 600 nm spectrally, to provide the first order diffraction image.

6.6. Conclusions & Future Work on Carbon Dots

Investigations are still actively ongoing within our lab, with a current focus on expanding the Trolox vs. no Trolox excitation polarization dataset. Interestingly, it appears that the presence of Trolox changes not only the blinking behavior, but may also change the excitation polarization behavior. This could perhaps indicate that the Trolox is chemically interacting with the CDs in a way that activates additional

absorbers or emitters. Further data acquisition and analysis is needed to support this hypothesis, however.

Next steps for this project include:

- Further analysis of excitation polarization data
- Acquisition of emission polarization data, by exciting with circularly polarized light, and using a rotating linear polarizer in the emission path.
- Acquisition of single carbon dot spectra.

Depending on how the story evolves based on the above experiments, taking data at a **second excitation wavelength** (probably 488 nm, given the current instrumental setup) for many of the above experiments could be very interesting, as this second wavelength would probe a **different fraction** of the multicolor carbon dots. Whether in this investigation or a later one, such multi-wavelength excitation data can help answer outstanding questions regarding the similarities and differences between the different carbon dot fractions from the one-pot synthesis.

With the above experiments performed, we will be able to garner new insights into the photophysical mechanisms of multicolor polymeric carbon dots. Such insights will help steer researchers in the rational design, chemical modification, and application space for this up-and-coming class of nanoparticles.

6.7. Experimental Methods

Equipment: All single CD data shown above was acquired with a 561 nm laser (Coherent Sapphire) set to 100 mW, fiber launched to clean up the mode quality, in widefield mode, along with a 561 nm long-pass filter (Semrock) and dual-band 488/561 nm dichroic (Semrock). All data shown in this chapter was taken using an inverted Nikon microscope and a 100x, 1.4 NA Nikon oil objective. Single particle fluorescence

images were taken with either an Andor Sona camera or an Andor iXon camera. Single particle spectral information was taken using the iXon camera attached to an Andor Kymera spectrograph.

Sample preparation: All coverslips and dilution glassware were plasma-cleaned in an air-plasma cleaner at high RF power for ~30 minutes prior to use. For the aqueous studies of CDs, coverslips were first incubated in 0.01% poly-L-lysine (Millipore Sigma) for at least 15 minutes, and then rinsed with milliQ water. Carbon dot samples were prepared via serial dilution in with milliQ water or freshly prepared 4 mM aqueous Trolox solution (20% ethanol), and were then dropcast for 10 minutes on the coverslips, followed by CD solution removal and replacement of the relevant supernatant. Spuncoat samples were not poly-lysine treated, and were spun at 4000 RPM for 30 seconds with 50 μ L of the desired solution on top.

7. Thesis Recap & Outlook

In this thesis, we have explored an interesting merger of the worlds of nanophotonics and chemistry in the form of microbubble photothermal absorption spectrometers. We started with a set of analogies to build up an intuitive understanding for microresonator experiments based on everyday experiences. We then dove into the technical background necessary for understanding topics such as thermo-optic coefficient, phase-matching, and the polarization of light. Next, we looked into the many different components of our microresonator spectrometer, and how to assemble, calibrate, and operate it. Following construction of the instrument, we then applied it to study solution-exposed gold nanorods undergoing chemical dynamics, a story that included interesting and unexpected turns such as photo-activated etching and controllable rotational dynamics added in. We then took a hard look at why our platform has yet to reach the single-molecule threshold, despite being more than sensitive enough to do so, and we discussed progress towards reaching the single molecule limit via lowering of photothermal background. Lastly, we zoomed out a bit to explore a different system, Carbon Dots, using photoluminescence spectroscopy, which will hopefully prove an exciting and more completed story soon.

While important progress was made in reducing photothermal background through the studies in Chapter 5, the fact remains that the majority of microbubbles as it stands aren't cut out for single molecule studies due to persistent, materials-based photothermal background. Additionally there are other difficulties with microbubbles, such as still needing a tapered optical fiber, and having to figure out what mode we are coupled to if we want to be quantitative with our measurements. Thus, while microbubbles may yet have a future in single molecule studies, and while they definitely still have stories to tell when it comes to light-matter interactions such as optoplasmonic coupling, greener pastures should be sought for single molecule studies. In particular, a variety of other microresonator platforms are in active development within our group for application in photothermal experiments. While each

platform will come with its own pros and cons that will be best-suited to different studies, hopefully one of these designs will enable single molecule absorption spectroscopy at room-temperature in solution.

8. References

1. Fowles, G. R., *Introduction to modern optics*. Courier Corporation: 1989.
2. Hecht, E., *Optics*. Global Edition ed.; Pearson: 2017.
3. Gaponenko, S. V., *Introduction to nanophotonics*. Cambridge University Press: 2010.
4. Lakowicz, J. R., *Principles of Fluorescence Spectroscopy*. Third Edition ed.; Springer: 2006.
5. Introduction to Polarization. <https://www.edmundoptics.com/knowledge-center/application-notes/optics/introduction-to-polarization/> (accessed September 9th).
6. Rea, M. Exploring Plasmonic-Photonic Coupled Systems Using Microcavity-Based Single-Particle Spectroscopy. University of Wisconsin - Madison., Madison, Wisconsin., 2020.
7. Horak, E. H. Optical Microresonators as Single-Particle Photothermal Absorption Spectrometers. University of Wisconsin-Madison., Madison, Wisconsin., 2019.
8. Heylman, K. D.; Knapper, K. A.; Horak, E. H.; Rea, M. T.; Vanga, S. K.; Goldsmith, R. H., Optical Microresonators for Sensing and Transduction: A Materials Perspective. *Advanced Materials* **2017**, *29* (30).
9. Oxborrow, M., Traceable 2-D finite-element simulation of the whispering-gallery modes of axisymmetric electromagnetic resonators. *Ieee Transactions on Microwave Theory and Techniques* **2007**, *55* (6), 1209-1218.
10. Heylman, K. D.; Thakkar, N.; Horak, E. H.; Quillin, S. C.; Cherqui, C.; Knapper, K. A.; Masiello, D. J.; Goldsmith, R. H., Optical microresonators as single-particle absorption spectrometers. *Nature Photonics* **2016**, *10* (12), 788-+.
11. Knight, J. C.; Cheung, G.; Jacques, F.; Birks, T. A., Phase-matched excitation of whispering-gallery-mode resonances by a fiber taper. *Opt. Lett.* **1997**, *22* (15), 1129-1131.
12. Polarization (Waves). [https://en.wikipedia.org/w/index.php?title=Polarization_\(waves\)&oldid=1108868677](https://en.wikipedia.org/w/index.php?title=Polarization_(waves)&oldid=1108868677) (accessed September 9th).
13. Kim, Y.; Torres, D. D.; Jain, P. K., Activation Energies of Plasmonic Catalysts. *Nano Letters* **2016**, *16* (5), 3399-3407.

14. Bigelow, J. E.; Kashnow, R. A., Poincaré sphere analysis of liquid crystal optics. *Appl. Opt.* **1977**, *16* (8), 2090-2096.
15. O'Neill, E. L., *Introduction to statistical optics*. Courier Corporation: 2003.
16. Heylman, K. D. Toroidal Optical Microresonators as Single-Particle Absorption Spectrometers. University of Wisconsin, Madison., Wisconsin, Madison., 2016.
17. Knapper, K. A. Advanced fabrication of toroidal optical microresonators for label-free photothermal imaging and spectroscopy. University of Wisconsin-Madison, Madison, Wisconsin, 2019.
18. Pan, F. Exploring Plasmonic-Photonic Coupled Systems Using Microcavity-Based Single-Particle Spectroscopy. University of Wisconsin - Madison, Madison, Wisconsin., 2022.
19. Horak, E. H.; Rea, M. T.; Heylman, K. D.; Gelbwaser-Klimovsky, D.; Saikin, S. K.; Thompson, B. J.; Kohler, D. D.; Knapper, K. A.; Wei, W.; Pan, F.; Gopalan, P.; Wright, J. C.; Aspuru-Guzik, A.; Goldsmith, R. H., Exploring Electronic Structure and Order in Polymers via Single-Particle Microresonator Spectroscopy. *Nano Letters* **2018**, *18* (3), 1600-1607.
20. Hogan, L. T.; Horak, E. H.; Ward, J. M.; Knapper, K. A.; Nic Chormaic, S.; Goldsmith, R. H., Toward Real-Time Monitoring and Control of Single Nanoparticle Properties with a Microbubble Resonator Spectrometer. *ACS Nano* **2019**, *13* (11), 12743-12757.
21. Carmon, T.; Yang, L.; Vahala, K. J., Dynamical thermal behavior and thermal self-stability of microcavities. *Opt. Express* **2004**, *12* (20), 4742-4750.
22. Jiang, X.-F.; Xiao, Y.-F.; Zou, C.-L.; He, L.; Dong, C.-H.; Li, B.-B.; Li, Y.; Sun, F.-W.; Yang, L.; Gong, Q., Highly Unidirectional Emission and Ultralow-Threshold Lasing from On-Chip Ultrahigh-Q Microcavities. *Advanced Materials* **2012**, *24* (35), OP260-OP264.
23. Lu, T.; Yang, L.; Carmon, T.; Min, B., A Narrow-Linewidth On-Chip Toroid Raman Laser. *IEEE Journal of Quantum Electronics* **2011**, *47* (3), 320-326.
24. Gold, D. C.; Bhadkamkar, A. S.; Carpenter, S.; Hogan, L. T.; Dwyer, M.; Beede, M.; Goldsmith, R. H.; van der Weide, D.; Yavuz, D. D., High-power near-CW Raman lasing in mm-sized glass disks. *Opt. Lett.* **2022**, *47* (16), 4171-4174.
25. Shao, L.; Jiang, X.-F.; Yu, X.-C.; Li, B.-B.; Clements, W. R.; Vollmer, F.; Wang, W.; Xiao, Y.-F.; Gong, Q., Detection of Single Nanoparticles and Lentiviruses Using Microcavity Resonance Broadening. *Advanced Materials* **2013**, *25* (39), 5616-5620.
26. Baaske, M. D.; Vollmer, F., Optical observation of single atomic ions interacting with plasmonic nanorods in aqueous solution. *Nature Photonics* **2016**, *10* (11), 733-+.
27. Spillane, S. Fiber-coupled ultra-high-Q microresonators for nonlinear and quantum optics. California Institute of Technology, Pasadena, California., 2004.
28. Swaim, J. D.; Knittel, J.; Bowen, W. P., Detection of nanoparticles with a frequency locked whispering gallery mode microresonator. *Applied Physics Letters* **2013**, *102* (18).

29. Black, E. D., An introduction to Pound-Drever-Hall laser frequency stabilization. *American Journal of Physics* **2001**, *69* (1), 79-87.
30. Barnes, J. A.; Gagliardi, G.; Loock, H. P., Absolute absorption cross-section measurement of a submonolayer film on a silica microresonator. *Optica* **2014**, *1* (2), 75-83.
31. Carmon, T.; Kippenberg, T. J.; Yang, L.; Rokhsari, H.; Spillane, S.; Vahala, K. J., Feedback control of ultra-high-Q microcavities: application to micro-Raman lasers and microparametric oscillators. *Opt. Express* **2005**, *13* (9), 3558-3566.
32. Kim, E.; Baaske, M. D.; Schuldes, I.; Wilsch, P. S.; Vollmer, F., Label-free optical detection of single enzyme-reactant reactions and associated conformational changes. *Science Advances* **2017**, *3* (3).
33. Kim, E.; Baaske, M. D.; Vollmer, F., In Situ Observation of Single-Molecule Surface Reactions from Low to High Affinities. *Advanced Materials* **2016**, *28* (45), 9941-9948.
34. Wu, F. C.; Wu, Y. Q.; Niu, Z. W.; Vollmer, F., Integrating a DNA Strand Displacement Reaction with a Whispering Gallery Mode Sensor for Label-Free Mercury (II) Ion Detection. *Sensors* **2016**, *16* (8).
35. Vollmer, F.; Arnold, S.; Keng, D., Single virus detection from the reactive shift of a whispering-gallery mode. *Proceedings of the National Academy of Sciences of the United States of America* **2008**, *105* (52), 20701-20704.
36. Knapper, K. A.; Horak, E. H.; Vollbrecht, C. H.; Heylman, K. D.; Goldsmith, R. H., Phase-sensitive photothermal imaging of ultrahigh-Q polyoxide toroidal microresonators. *Applied Physics Letters* **2018**, *113* (23).
37. Connolly, E. W. Experiments with toroidal microresonators in cavity QED. California Institute of Technology., Pasadena, CA., 2009.
38. Lei, F.; Murphy, R. M. J.; Ward, J. M.; Yang, Y.; Nic Chormaic, S., Bandpass transmission spectra of a whispering-gallery microcavity coupled to an ultrathin fiber. *Photon. Res.* **2017**, *5* (4), 362-366.
39. Kim, E.; Baaske, M. D.; Vollmer, F., Towards next-generation label-free biosensors: recent advances in whispering gallery mode sensors. *Lab on a Chip* **2017**, *17* (7), 1190-1205.
40. Bozzola, A.; Perotto, S.; De Angelis, F., Hybrid plasmonic-photonic whispering gallery mode resonators for sensing: a critical review. *Analyst* **2017**, *142* (6), 883-898.
41. Li, Y.; Jiang, X.; Zhao, G.; Yang, L. Whispering gallery mode microresonator for nonlinear optics *arXiv e-prints* [Online], 2018. <https://arxiv.org/abs/1809.04878> (accessed September 13, 2018).
42. Wade, J. H.; Bailey, R. C., Applications of Optical Microcavity Resonators in Analytical Chemistry. *Annual Review of Analytical Chemistry, Vol 9* **2016**, *9*, 1-25.
43. Zhi, Y. Y.; Yu, X. C.; Gong, Q. H.; Yang, L.; Xiao, Y. F., Single Nanoparticle Detection Using Optical Microcavities. *Advanced Materials* **2017**, *29* (12).

44. Baaske, M. D.; Foreman, M. R.; Vollmer, F., Single-molecule nucleic acid interactions monitored on a label-free microcavity biosensor platform. *Nature Nanotechnology* **2014**, *9* (11), 933-939.
45. Dantham, V. R.; Holler, S.; Barbre, C.; Keng, D.; Kolchenko, V.; Arnold, S., Label-Free Detection of Single Protein Using a Nanoplasmonic-Photonic Hybrid Microcavity. *Nano Letters* **2013**, *13* (7), 3347-3351.
46. Yu, W. Y.; Jiang, W. C.; Lin, Q.; Lu, T., Cavity optomechanical spring sensing of single molecules. *Nature Communications* **2016**, *7*.
47. Thakkar, N.; Rea, M. T.; Smith, K. C.; Heylman, K. D.; Quillin, S. C.; Knapper, K. A.; Horak, E. H.; Masiello, D. J.; Goldsmith, R. H., Sculpting Fano Resonances To Control Photonic-Plasmonic Hybridization. *Nano Letters* **2017**, *17* (11), 6927-6934.
48. Knapper, K. A.; Heylman, K. D.; Horak, E. H.; Goldsmith, R. H., Chip-Scale Fabrication of High-Q All-Glass Toroidal Microresonators for Single-Particle Label-Free Imaging. *Advanced Materials* **2016**, *28* (15), 2945-2950.
49. Knapper, K. A.; Pan, F.; Rea, M. T.; Horak, E. H.; Rogers, J. D.; Goldsmith, R. H., Single-particle photothermal imaging via inverted excitation through high-Q all-glass toroidal microresonators. *Opt. Express* **2018**, *26* (19), 25020-25030.
50. Heylman, K. D.; Knapper, K. A.; Goldsmith, R. H., Photothermal Microscopy of Nonluminescent Single Particles Enabled by Optical Microresonators. *Journal of Physical Chemistry Letters* **2014**, *5* (11), 1917-1923.
51. Chen, H. J.; Shao, L.; Li, Q.; Wang, J. F., Gold nanorods and their plasmonic properties. *Chemical Society Reviews* **2013**, *42* (7), 2679-2724.
52. Dreaden, E. C.; Alkilany, A. M.; Huang, X. H.; Murphy, C. J.; El-Sayed, M. A., The golden age: gold nanoparticles for biomedicine. *Chemical Society Reviews* **2012**, *41* (7), 2740-2779.
53. Huang, X. H.; El-Sayed, I. H.; Qian, W.; El-Sayed, M. A., Cancer cell imaging and photothermal therapy in the near-infrared region by using gold nanorods. *Journal of the American Chemical Society* **2006**, *128* (6), 2115-2120.
54. Yin, D. Y.; Li, X. L.; Ma, Y. Y.; Liu, Z., Targeted cancer imaging and photothermal therapy via monosaccharide-imprinted gold nanorods. *Chemical Communications* **2017**, *53* (50), 6716-6719.
55. Ali, M. R. K.; Wu, Y.; Ghosh, D.; Do, B. H.; Chen, K.; Dawson, M. R.; Fang, N.; Sulchek, T. A.; El-Sayed, M. A., Nuclear Membrane-Targeted Gold Nanoparticles Inhibit Cancer Cell Migration and Invasion. *Acs Nano* **2017**, *11* (4), 3716-3726.
56. Meeker, D. G.; Chen, J. Y.; Smeltzer, M. S., Could targeted, antibiotic-loaded gold nanoconstructs be a new magic bullet to fight infection? *Nanomedicine* **2016**, *11* (18), 2379-2382.
57. Cao, J.; Sun, T.; Grattan, K. T. V., Gold nanorod-based localized surface plasmon resonance biosensors: A review. *Sensors and Actuators B-Chemical* **2014**, *195*, 332-351.

58. Taylor, A. B.; Zijlstra, P., Single-Molecule Plasmon Sensing: Current Status and Future Prospects. *Acs Sensors* **2017**, *2* (8), 1103-1122.
59. Lin, K. Q.; Yi, J.; Hu, S.; Liu, B. J.; Liu, J. Y.; Wang, X.; Ren, B., Size Effect on SERS of Gold Nanorods Demonstrated via Single Nanoparticle Spectroscopy. *Journal of Physical Chemistry C* **2016**, *120* (37), 20806-20813.
60. Gao, Z.; Burrows, N. D.; Valley, N. A.; Schatz, G. C.; Murphy, C. J.; Haynes, C. L., In solution SERS sensing using mesoporous silica-coated gold nanorods. *Analyt* **2016**, *141* (17), 5088-5095.
61. Khatua, S.; Paulo, P. M. R.; Yuan, H. F.; Gupta, A.; Zijlstra, P.; Orrit, M., Resonant Plasmonic Enhancement of Single-Molecule Fluorescence by Individual Gold Nanorods. *Acs Nano* **2014**, *8* (5), 4440-4449.
62. Nima, Z. A.; Alwbari, A. M.; Dantuluri, V.; Hamzah, R. N.; Sra, N.; Motwani, P.; Arnaoutakis, K.; Levy, R. A.; Bohliqa, A. F.; Nedosekin, D.; Zharov, V. P.; Makhoul, I.; Biris, A. S., Targeting nano drug delivery to cancer cells using tunable, multi-layer, silver-decorated gold nanorods. *Journal of Applied Toxicology* **2017**, *37* (12), 1370-1378.
63. Wang, F.; Li, C. H.; Chen, H. J.; Jiang, R. B.; Sun, L. D.; Li, Q.; Wang, J. F.; Yu, J. C.; Yan, C. H., Plasmonic Harvesting of Light Energy for Suzuki Coupling Reactions. *Journal of the American Chemical Society* **2013**, *135* (15), 5588-5601.
64. Gole, A.; Murphy, C. J., Seed-mediated synthesis of gold nanorods: Role of the size and nature of the seed. *Chemistry of Materials* **2004**, *16* (19), 3633-3640.
65. Baida, H.; Christofilos, D.; Maioli, P.; Crut, A.; Del Fatti, N.; Vallee, F., Surface plasmon resonance spectroscopy of single surfactant-stabilized gold nanoparticles. *European Physical Journal D* **2011**, *63* (2), 293-299.
66. Li, Z. M.; Mao, W. Z.; Devadas, M. S.; Hartland, G. V., Absorption Spectroscopy of Single Optically Trapped Gold Nanorods. *Nano Letters* **2015**, *15* (11), 7731-7735.
67. Yorulmaz, M.; Nizzero, S.; Hoggard, A.; Wang, L. Y.; Cai, Y. Y.; Su, M. N.; Chang, W. S.; Link, S., Single-Particle Absorption Spectroscopy by Photothermal Contrast. *Nano Letters* **2015**, *15* (5), 3041-3047.
68. Berciaud, S.; Cognet, L.; Tamarat, P.; Lounis, B., Observation of intrinsic size effects in the optical response of individual gold nanoparticles. *Nano Letters* **2005**, *5* (3), 515-518.
69. Gaiduk, A.; Yorulmaz, M.; Ruijgrok, P. V.; Orrit, M., Room-Temperature Detection of a Single Molecule's Absorption by Photothermal Contrast. *Science* **2010**, *330* (6002), 353-356.
70. Chien, M. H.; Brameshuber, M.; Rossboth, B. K.; Schutz, G. J.; Schmid, S., Single-molecule optical absorption imaging by nanomechanical photothermal sensing. *Proceedings of the National Academy of Sciences of the United States of America* **2018**, *115* (44), 11150-11155.
71. Crut, A.; Maioli, P.; Del Fatti, N.; Vallee, F., Optical absorption and scattering spectroscopies of single nano-objects. *Chemical Society Reviews* **2014**, *43* (11), 3921-3956.

72. Zrimsek, A. B.; Wong, N. L.; Van Duyne, R. P., Single Molecule Surface-Enhanced Raman Spectroscopy: A Critical Analysis of the Bianalyte versus Isotopologue Proof. *Journal of Physical Chemistry C* **2016**, *120* (9), 5133-5142.
73. Young, G.; Kukura, P., Interferometric Scattering Microscopy. *Annual Review of Physical Chemistry* **2019**, *70* (1), 301-322.
74. Celebrano, M.; Kukura, P.; Renn, A.; Sandoghdar, V., Single-molecule imaging by optical absorption. *Nature Photonics* **2011**, *5* (2), 95-98.
75. Chong, S. S.; Min, W.; Xie, X. S., Ground-State Depletion Microscopy: Detection Sensitivity of Single-Molecule Optical Absorption at Room Temperature. *Journal of Physical Chemistry Letters* **2010**, *1* (23), 3316-3322.
76. Maley, A. M.; Lu, G. J.; Shapiro, M. G.; Corn, R. M., Characterizing Single Polymeric and Protein Nanoparticles with Surface Plasmon Resonance Imaging Measurements. *Acs Nano* **2017**, *11* (7), 7447-7456.
77. Jiang, D.; Jiang, Y. Y.; Li, Z. M.; Liu, T.; Wo, X.; Fang, Y. M.; Tao, N. J.; Wang, W.; Chen, H. Y., Optical Imaging of Phase Transition and Li-Ion Diffusion Kinetics of Single LiCoO₂ Nanoparticles During Electrochemical Cycling. *Journal of the American Chemical Society* **2017**, *139* (1), 186-192.
78. Thambi, V.; Kar, A.; Ghosh, P.; Khatua, S., Light-Controlled in Situ Bidirectional Tuning and Monitoring of Gold Nanorod Plasmon via Oxidative Etching with FeCl₃. *Journal of Physical Chemistry C* **2018**, *122* (43), 24885-24890.
79. Carattino, A.; Khatua, S.; Orrit, M., In situ tuning of gold nanorod plasmon through oxidative cyanide etching. *Physical Chemistry Chemical Physics* **2016**, *18* (23), 15619-15624.
80. Al-Zubeidi, A.; Hoener, B. S.; Collins, S. S. E.; Wang, W.; Kirchner, S. R.; Hosseini Jebeli, S. A.; Joplin, A.; Chang, W.-S.; Link, S.; Landes, C. F., Hot Holes Assist Plasmonic Nanoelectrode Dissolution. *Nano Letters* **2019**, *19* (2), 1301-1306.
81. Cheng, J.; Liu, Y.; Cheng, X. D.; He, Y.; Yeung, E. S., Real Time Observation of Chemical Reactions of Individual Metal Nanoparticles with High-Throughput Single Molecule Spectral Microscopy. *Analytical Chemistry* **2010**, *82* (20), 8744-8749.
82. Sun, S. S.; Gao, M. X.; Lei, G.; Zou, H. Y.; Ma, J.; Huang, C. Z., Visually monitoring the etching process of gold nanoparticles by KI/I₂ at single-nanoparticle level using scattered-light dark-field microscopic imaging. *Nano Research* **2016**, *9* (4), 1125-1134.
83. Wang, J.; Zhang, H. Z.; Liu, J. J.; Yuan, D.; Li, R. S.; Huang, C. Z., Time-resolved visual detection of heparin by accelerated etching of gold nanorods. *Analyst* **2018**, *143* (4), 824-828.
84. Zhang, H. Z.; Li, R. S.; Gao, P. F.; Wang, N.; Lei, G.; Huang, C. Z.; Wang, J., Real-time dark-field light scattering imaging to monitor the coupling reaction with gold nanorods as an optical probe. *Nanoscale* **2017**, *9* (10), 3568-3575.

85. Xie, T.; Jing, C.; Ma, W.; Ding, Z. F.; Gross, A. J.; Long, Y. T., Real-time monitoring for the morphological variations of single gold nanorods. *Nanoscale* **2015**, *7* (2), 511-517.
86. Flatebo, C.; Collins, S. S. E.; Hoener, B. S.; Cai, Y.-y.; Link, S.; Landes, C. F., Electrodeposition Inhibition of Gold Nanorods with Oxoanions. *The Journal of Physical Chemistry C* **2019**, *123* (22), 13983-13992.
87. Ye, X. C.; Jones, M. R.; Frechette, L. B.; Chen, Q.; Powers, A. S.; Ercius, P.; Dunn, G.; Rotskoff, G. M.; Nguyen, S. C.; Adiga, V. P.; Zettl, A.; Rabani, E.; Geissler, P. L.; Alivisatos, A. P., Single-particle mapping of nonequilibrium nanocrystal transformations. *Science* **2016**, *354* (6314), 874-877.
88. Sun, Y. Z.; Fan, X. D., Optical ring resonators for biochemical and chemical sensing. *Analytical and Bioanalytical Chemistry* **2011**, *399* (1), 205-211.
89. Lu, T.; Lee, H.; Chen, T.; Herchak, S.; Kim, J. H.; Fraser, S. E.; Flagan, R. C.; Vahala, K., High sensitivity nanoparticle detection using optical microcavities. *Proceedings of the National Academy of Sciences of the United States of America* **2011**, *108* (15), 5976-5979.
90. Ward, J. M.; Yang, Y.; Lei, F. C.; Yu, X. C.; Xiao, Y. F.; Chormaic, S. N., Nanoparticle sensing beyond evanescent field interaction with a quasi-droplet microcavity. *Optica* **2018**, *5* (6), 674-677.
91. Barucci, A.; Berneschi, S.; Giannetti, A.; Baldini, F.; Cosci, A.; Pelli, S.; Farnesi, D.; Righini, G. C.; Soria, S.; Conti, G. N., Optical Microbubble Resonators with High Refractive Index Inner Coating for Bio-Sensing Applications: An Analytical Approach. *Sensors* **2016**, *16* (12).
92. Giorgini, A.; Avino, S.; Malara, P.; De Natale, P.; Gagliardi, G., Liquid Droplet Microresonators. *Sensors* **2019**, *19* (3).
93. Madani, A.; Harazim, S. M.; Quinones, V. A. B.; Kleinert, M.; Finn, A.; Naz, E. S. G.; Ma, L. B.; Schmidt, O. G., Optical microtube cavities monolithically integrated on photonic chips for optofluidic sensing. *Opt. Lett.* **2017**, *42* (3), 486-489.
94. Han, K. W.; Kim, J.; Bahl, G., High-throughput sensing of freely flowing particles with optomechanofluidics. *Optica* **2016**, *3* (6), 585-591.
95. Stoian, R.-I.; Bui, K. V.; Rosenberger, A., Silica hollow bottle resonators for use as whispering gallery mode based chemical sensors. *Journal of Optics* **2015**, *17* (12), 125011.
96. Ward, J. M.; Dhasmana, N.; Chormaic, S. N., Hollow core, whispering gallery resonator sensors. *European Physical Journal-Special Topics* **2014**, *223* (10), 1917-1935.
97. Kippenberg, T. J.; Spillane, S. M.; Vahala, K. J., Demonstration of ultra-high-Q small mode volume toroid microcavities on a chip. *Applied Physics Letters* **2004**, *85* (25), 6113-6115.
98. Armani, A. M.; Vahala, K. J., Heavy water detection using ultra-high-Q microcavities. *Opt. Lett.* **2006**, *31* (12), 1896-1898.
99. Armani, D. K.; Kippenberg, T. J.; Spillane, S. M.; Vahala, K. J., Ultra-high-Q toroid microcavity on a chip. *Nature* **2003**, *421* (6926), 925-928.

100. Murugan, G. S.; Petrovich, M. N.; Jung, Y.; Wilkinson, J. S.; Zervas, M. N., Hollow-bottle optical microresonators. *Opt. Express* **2011**, *19* (21), 20773-20784.
101. Nasir, M. N. M.; Murugan, G. S.; Zervas, M. N., Spectral cleaning and output modal transformations in whispering-gallery-mode microresonators. *Journal of the Optical Society of America B-Optical Physics* **2016**, *33* (9), 1963-1970.
102. Ding, M.; Murugan, G. S.; Brambilla, G.; Zervas, M. N., Whispering gallery mode selection in optical bottle microresonators. *Applied Physics Letters* **2012**, *100* (8).
103. Schunk, G.; Furst, J. U.; Fortsch, M.; Strekalov, D. V.; Vogl, U.; Sedlmeir, F.; Schwefel, H. G. L.; Leuchs, G.; Marquardt, C., Identifying modes of large whispering-gallery mode resonators from the spectrum and emission pattern. *Opt. Express* **2014**, *22* (25), 30795-30806.
104. Ward, J. M.; Yang, Y.; Nic Chormaic, S., Highly Sensitive Temperature Measurements With Liquid-Core Microbubble Resonators. *Ieee Photonics Technology Letters* **2013**, *25* (23), 2350-2353.
105. Hall, J. M. M.; Francois, A.; Afshar, S.; Riesen, N.; Henderson, M. R.; Reynolds, T.; Monro, T. M., Determining the geometric parameters of microbubble resonators from their spectra. *Journal of the Optical Society of America B-Optical Physics* **2017**, *34* (1), 44-51.
106. Murugan, G. S.; Wilkinson, J. S.; Zervas, M. N., Selective excitation of whispering gallery modes in a novel bottle microresonator. *Opt. Express* **2009**, *17* (14), 11916-11925.
107. Davletshin, Y. R.; Lombardi, A.; Cardinal, M. F.; Juve, V.; Crut, A.; Maioli, P.; Liz-Marzan, L. M.; Vallee, F.; Del Fatti, N.; Kumaradas, J. C., A Quantitative Study of the Environmental Effects on the Optical Response of Gold Nanorods. *Acs Nano* **2012**, *6* (9), 8183-8193.
108. Ni, W. H.; Chen, H. J.; Kou, X. S.; Yeung, M. H.; Wang, J. F., Optical fiber-excited surface plasmon resonance spectroscopy of single and ensemble gold nanorods. *Journal of Physical Chemistry C* **2008**, *112* (22), 8105-8109.
109. Park, K.; Biswas, S.; Kanel, S.; Nepal, D.; Vaia, R. A., Engineering the Optical Properties of Gold Nanorods: Independent Tuning of Surface Plasmon Energy, Extinction Coefficient, and Scattering Cross Section. *Journal of Physical Chemistry C* **2014**, *118* (11), 5918-5926.
110. Jana, N. R.; Gearheart, L.; Obare, S. O.; Murphy, C. J., Anisotropic chemical reactivity of gold spheroids and nanorods. *Langmuir* **2002**, *18* (3), 922-927.
111. Zou, R. X.; Guo, X.; Yang, J.; Li, D. D.; Peng, F.; Zhang, L.; Wang, H. J.; Yu, H., Selective etching of gold nanorods by ferric chloride at room temperature. *Crystengcomm* **2009**, *11* (12), 2797-2803.
112. Zhao, J.; Nguyen, S. C.; Ye, R.; Ye, B. H.; Weller, H.; Somorjai, G. A.; Alivisatos, A. P.; Toste, F. D., A Comparison of Photocatalytic Activities of Gold Nanoparticles Following Plasmonic and Interband Excitation and a Strategy for Harnessing Interband Hot Carriers for Solution Phase Photocatalysis. *Acs Central Science* **2017**, *3* (5), 482-488.
113. Zijlstra, P.; Orrit, M., Single metal nanoparticles: optical detection, spectroscopy and applications. *Reports on Progress in Physics* **2011**, *74* (10).

114. Weigel, A.; Sebesta, A.; Kukura, P., Dark Field Microspectroscopy with Single Molecule Fluorescence Sensitivity. *Acs Photonics* **2014**, *1* (9), 848-856.
115. Chang, W. S.; Link, S., Enhancing the Sensitivity of Single-Particle Photothermal Imaging with Thermotropic Liquid Crystals. *Journal of Physical Chemistry Letters* **2012**, *3* (10), 1393-1399.
116. Parra-Vasquez, A. N. G.; Oudjedi, L.; Cognet, L.; Lounis, B., Nanoscale Thermotropic Phase Transitions Enhancing Photothermal Microscopy Signals. *Journal of Physical Chemistry Letters* **2012**, *3* (10), 1400-1403.
117. Ding, T. N. X.; Hou, L.; van der Meer, H.; Alivisatos, A. P.; Orrit, M., Hundreds-fold Sensitivity Enhancement of Photothermal Microscopy in Near-Critical Xenon. *Journal of Physical Chemistry Letters* **2016**, *7* (13), 2524-2529.
118. Rodriguez-Fernandez, J.; Perez-Juste, J.; Mulvaney, P.; Liz-Marzan, L. M., Spatially-directed oxidation of gold nanoparticles by Au(III)-CTAB complexes. *Journal of Physical Chemistry B* **2005**, *109* (30), 14257-14261.
119. Brongersma, M. L.; Halas, N. J.; Nordlander, P., Plasmon-induced hot carrier science and technology. *Nature Nanotechnology* **2015**, *10* (1), 25-34.
120. Zhang, Y. C.; He, S.; Guo, W. X.; Hu, Y.; Huang, J. W.; Mulcahy, J. R.; Wei, W. D., Surface-Plasmon-Driven Hot Electron Photochemistry. *Chemical Reviews* **2018**, *118* (6), 2927-2954.
121. Kale, M. J.; Avanesian, T.; Christopher, P., Direct Photocatalysis by Plasmonic Nanostructures. *Acs Catalysis* **2014**, *4* (1), 116-128.
122. Linic, S.; Aslam, U.; Boerigter, C.; Morabito, M., Photochemical transformations on plasmonic metal nanoparticles. *Nature Materials* **2015**, *14* (6), 567-576.
123. Besteiro, L. V.; Kong, X. T.; Wang, Z. M.; Hartland, G.; Govorov, A. O., Understanding Hot-Electron Generation and Plasmon Relaxation in Metal Nanocrystals: Quantum and Classical Mechanisms. *Acs Photonics* **2017**, *4* (11), 2759-2781.
124. Hartland, G. V.; Besteiro, L. V.; Johns, P.; Govorov, A. O., What's so Hot about Electrons in Metal Nanoparticles? *Acs Energy Letters* **2017**, *2* (7), 1641-1653.
125. Wu, K.; Chen, J.; McBride, J. R.; Lian, T., Efficient hot-electron transfer by a plasmon-induced interfacial charge-transfer transition. *Science* **2015**, *349* (6248), 632-635.
126. Minutella, E.; Schulz, F.; Lange, H., Excitation-Dependence of Plasmon-Induced Hot Electrons in Gold Nanoparticles. *Journal of Physical Chemistry Letters* **2017**, *8* (19), 4925-4929.
127. Kamarudheen, R.; Castellanos, G. W.; Kamp, L. P. J.; Clercx, H. J. H.; Baldi, A., Quantifying Photothermal and Hot Charge Carrier Effects in Plasmon-Driven Nanoparticle Syntheses. *Acs Nano* **2018**, *12* (8), 8447-8455.

128. Zhou, L. A.; Swearer, D. F.; Zhang, C.; Robotjazi, H.; Zhao, H. Q.; Henderson, L.; Dong, L. L.; Christopher, P.; Carter, E. A.; Nordlander, P.; Halas, N. J., Quantifying hot carrier and thermal contributions in plasmonic photocatalysis. *Science* **2018**, *362* (6410), 69-+.
129. Wu, C. Y.; Wolf, W. J.; Levartovsky, Y.; Bechtel, H. A.; Martin, M. C.; Toste, F. D.; Gross, E., High-spatial-resolution mapping of catalytic reactions on single particles. *Nature* **2017**, *541* (7638), 511-515.
130. Cortes, E.; Xie, W.; Cambiasso, J.; Jermyn, A. S.; Sundararaman, R.; Narang, P.; Schlucker, S.; Maier, S. A., Plasmonic hot electron transport drives nano-localized chemistry. *Nature Communications* **2017**, *8*.
131. Schather, A. E.; Manjavacas, A.; Lauchner, A.; Marangoni, V. S.; DeSantis, C. J.; Nordlander, P.; Halas, N. J., Hot Hole Photoelectrochemistry on Au@SiO₂@Au Nanoparticles. *Journal of Physical Chemistry Letters* **2017**, *8* (9), 2060-2067.
132. Smith, J. G.; Jain, P. K., The Ligand Shell as an Energy Barrier in Surface Reactions on Transition Metal Nanoparticles. *Journal of the American Chemical Society* **2016**, *138* (21), 6765-6773.
133. Burrows, N. D.; Lin, W.; Hinman, J. G.; Dennison, J. M.; Vartanian, A. M.; Abadeer, N. S.; Grzincic, E. M.; Jacob, L. M.; Li, J.; Murphy, C. J., Surface Chemistry of Gold Nanorods. *Langmuir* **2016**, *32* (39), 9905-9921.
134. He, J.; Unser, S.; Bruzas, I.; Cary, R.; Shi, Z. W.; Mehra, R.; Aron, K.; Sagle, L., The facile removal of CTAB from the surface of gold nanorods. *Colloids and Surfaces B-Biointerfaces* **2018**, *163*, 140-145.
135. Ruijgrok, P. V.; Verhart, N. R.; Zijlstra, P.; Tchebotareva, A. L.; Orrit, M., Brownian Fluctuations and Heating of an Optically Aligned Gold Nanorod. *Physical Review Letters* **2011**, *107* (3).
136. Trojek, J.; Chvatal, L.; Zemanek, P., Optical alignment and confinement of an ellipsoidal nanorod in optical tweezers: a theoretical study. *Journal of the Optical Society of America a-Optics Image Science and Vision* **2012**, *29* (7), 1224-1236.
137. Xu, X. H.; Cheng, C.; Zhang, Y.; Lei, H. X.; Li, B. J., Scattering and Extinction Torques: How Plasmon Resonances Affect the Orientation Behavior of a Nanorod in Linearly Polarized Light. *Journal of Physical Chemistry Letters* **2016**, *7* (2), 314-319.
138. Tong, L. M.; Miljkovic, V. D.; Kall, M., Alignment, Rotation, and Spinning of Single Plasmonic Nanoparticles and Nanowires Using Polarization Dependent Optical Forces. *Nano Letters* **2010**, *10* (1), 268-273.
139. Yan, Z. J.; Sweet, J.; Jureller, J. E.; Guffey, M. J.; Pelton, M.; Scherert, N. F., Controlling the Position and Orientation of Single Silver Nanowires on a Surface Using Structured Optical Fields. *ACS Nano* **2012**, *6* (9), 8144-8155.
140. Yang, Y.; Saurabh, S.; Ward, J. M.; Nic Chormaic, S., High-Q, ultrathin-walled microbubble resonator for aerostatic pressure sensing. *Opt. Express* **2016**, *24* (1), 294-299.
141. Henze, R.; Seifert, T.; Ward, J.; Benson, O., Tuning whispering gallery modes using internal aerostatic pressure. *Opt. Lett.* **2011**, *36* (23), 4536-4538.

142. Cosci, A.; Quercioli, F.; Farnesi, D.; Berneschi, S.; Giannetti, A.; Cosi, F.; Barucci, A.; Conti, G. N.; Righini, G.; Pelli, S., Confocal reflectance microscopy for determination of microbubble resonator thickness. *Opt. Express* **2015**, *23* (13), 16693-16701.
143. Watkins, A.; Ward, J.; Wu, Y. Q.; Nic Chormaic, S., Single-input spherical microbubble resonator. *Opt. Lett.* **2011**, *36* (11), 2113-2115.
144. Ward, J. M.; Yang, Y.; Chormaic, S. N., Highly Sensitive Temperature Measurements With Liquid-Core Microbubble Resonators. *Ieee Photonics Technology Letters* **2013**, *25* (23), 2350-2353.
145. Nanopartz Gold Nanorods. https://www.nanopartz.com/bare_gold_nanorods.asp (accessed June 8th).
146. Fernandes, L. Birefringence and Bragg grating control in femtosecond laser written optical circuits. 2012.
147. Ni, W.; Kou, X.; Yang, Z.; Wang, J. F., Tailoring longitudinal surface plasmon wavelengths, scattering and absorption cross sections of gold nanorods. *Acs Nano* **2008**, *2* (4), 677-686.
148. Chandrasekar, G.; Mougin, K.; Haidara, H.; Vidal, L.; Gnecco, E., Shape and size transformation of gold nanorods (GNRs) via oxidation process: A reverse growth mechanism. *Applied Surface Science* **2011**, *257* (9), 4175-4179.
149. Xie, F.; Ye, W.; Sun, H. Y.; Kou, S. F.; Guo, X., Silver Ions Induce Lateral Etching of Gold Nanorods by K₂PtCl₄. *Langmuir* **2015**, *31* (24), 6823-6828.
150. Monga, A.; Pal, B., Influence of Oxidative Etching of Au Nanostructures by KMnO₄ on its Surface Morphology, Electro-kinetic Properties and Improved Catalytic Activity. *Journal of Industrial and Engineering Chemistry* **2015**, *31*, 223-230.
151. Szychowski, B.; Leng, H. X.; Pelton, M.; Daniel, M. C., Controlled etching and tapering of Au nanorods using cysteamine. *Nanoscale* **2018**, *10* (35), 16830-16838.
152. Tsung, C. K.; Kou, X. S.; Shi, Q. H.; Zhang, J. P.; Yeung, M. H.; Wang, J. F.; Stucky, G. D., Selective shortening of single-crystalline gold nanorods by mild oxidation. *Journal of the American Chemical Society* **2006**, *128* (16), 5352-5353.
153. Rex, M.; Hernandez, F. E.; Campiglia, A. D., Pushing the limits of mercury sensors with gold nanorods. *Analytical Chemistry* **2006**, *78* (2), 445-451.
154. Li, F. M.; Liu, J. M.; Wang, X. X.; Lin, L. P.; Cai, W. L.; Lin, X.; Zeng, Y. N.; Li, Z. M.; Lin, S. Q., Non-aggregation based label free colorimetric sensor for the detection of Cr (VI) based on selective etching of gold nanorods. *Sensors and Actuators B-Chemical* **2011**, *155* (2), 817-822.
155. Chen, Y. Y.; Chang, H. T.; Shiang, Y. C.; Hung, Y. L.; Chiang, C. K.; Huang, C. C., Colorimetric Assay for Lead Ions Based on the Leaching of Gold Nanoparticles. *Analytical Chemistry* **2009**, *81* (22), 9433-9439.

156. Chen, Z. P.; Zhang, Z. Y.; Qu, C. L.; Pan, D. W.; Chen, L. X., Highly sensitive label-free colorimetric sensing of nitrite based on etching of gold nanorods. *Analyst* **2012**, *137* (22), 5197-5200.
157. Sreeprasad, T. S.; Samal, A. K.; Pradeep, T., Body- or tip-controlled reactivity of gold nanorods and their conversion to particles through other anisotropic structures. *Langmuir* **2007**, *23* (18), 9463-9471.
158. Zhang, Z. Y.; Chen, Z. P.; Qu, C. L.; Chen, L. X., Highly Sensitive Visual Detection of Copper Ions Based on the Shape-Dependent LSPR Spectroscopy of Gold Nanorods. *Langmuir* **2014**, *30* (12), 3625-3630.
159. Zhang, Z. Y.; Chen, Z. P.; Pan, D. W.; Chen, L. X., Fenton-like Reaction-Mediated Etching of Gold Nanorods for Visual Detection of Co²⁺. *Langmuir* **2015**, *31* (1), 643-650.
160. Zhang, Z. Y.; Chen, Z. P.; Chen, L. X., Ultrasensitive Visual Sensing of Molybdate Based on Enzymatic-like Etching of Gold Nanorods. *Langmuir* **2015**, *31* (33), 9253-9259.
161. Liu, J. M.; Wang, X. X.; Jiao, L.; Cui, M. L.; Lin, L. P.; Zhang, L. H.; Jiang, S. L., Ultra-sensitive non-aggregation colorimetric sensor for detection of iron based on the signal amplification effect of Fe³⁺ catalyzing H₂O₂ oxidize gold nanorods. *Talanta* **2013**, *116*, 199-204.
162. Liu, J. M.; Jiao, L.; Cui, M. L.; Lin, L. P.; Wang, X. X.; Zheng, Z. Y.; Zhang, L. H.; Jiang, S. L., highly sensitive non-aggregation colorimetric sensor for the determination of I⁻ based on its catalytic effect on Fe³⁺ etching gold nanorods. *Sensors and Actuators B-Chemical* **2013**, *188*, 644-650.
163. Liu, X.; Zhang, S. Y.; Tan, P. L.; Zhou, J.; Huang, Y.; Nie, Z.; Yao, S. Z., A plasmonic blood glucose monitor based on enzymatic etching of gold nanorods. *Chemical Communications* **2013**, *49* (18), 1856-1858.
164. Saa, L.; Coronado-Puchau, M.; Pavlov, V.; Liz-Marzan, L. M., Enzymatic etching of gold nanorods by horseradish peroxidase and application to blood glucose detection. *Nanoscale* **2014**, *6* (13), 7405-7409.
165. Wang, D. N.; Guo, R.; Wei, Y. Y.; Zhang, Y. Z.; Zhao, X. Y.; Xu, Z. R., Sensitive multicolor visual detection of telomerase activity based on catalytic hairpin assembly and etching of Au nanorods. *Biosensors & Bioelectronics* **2018**, *122*, 247-253.
166. Zhang, Z. Y.; Chen, Z. P.; Wang, S. S.; Cheng, F. B.; Chen, L. X., Iodine-Mediated Etching of Gold Nanorods for Plasmonic ELISA Based on Colorimetric Detection of Alkaline Phosphatase. *Acs Applied Materials & Interfaces* **2015**, *7* (50), 27639-27645.
167. Wu, S.; Li, D. D.; Gao, Z. M.; Wang, J. M., Controlled etching of gold nanorods by the Au(III)-CTAB complex, and its application to semi-quantitative visual determination of organophosphorus pesticides. *Microchimica Acta* **2017**, *184* (11), 4383-4391.
168. Lu, S. M.; Zhang, X.; Chen, L.; Yang, P., Colorimetric visualization of superoxide dismutase in serum via etching of Au nanorods from superoxide radical. *Sensors and Actuators B-Chemical* **2018**, *259*, 1066-1072.

169. Zhu, Q. N.; Wu, J.; Zhao, J. W.; Ni, W. H., Role of Bromide in Hydrogen Peroxide Oxidation of CTAB-Stabilized Gold Nanorods in Aqueous Solutions. *Langmuir* **2015**, *31* (14), 4072-4077.
170. Weng, G. J.; Dong, X. J.; Li, J. J.; Zhao, J. W., Halide ions can trigger the oxidative etching of gold nanorods with the iodide ions being the most efficient. *Journal of Materials Science* **2016**, *51* (16), 7678-7690.
171. Yuan, H. F.; Janssen, K. P. F.; Franklin, T.; Lu, G.; Su, L.; Gu, X.; Uji-i, H.; Roeffaers, M. B. J.; Hofkens, J., Reshaping anisotropic gold nanoparticles through oxidative etching: the role of the surfactant and nanoparticle surface curvature. *Rsc Advances* **2015**, *5* (9), 6829-6833.
172. Wen, T.; Zhang, H.; Tang, X. P.; Chu, W. G.; Liu, W. Q.; Ji, Y. L.; Hu, Z. J.; Hou, S.; Hu, X. N.; Wu, X. C., Copper Ion Assisted Reshaping and Etching of Gold Nanorods: Mechanism Studies and Applications. *Journal of Physical Chemistry C* **2013**, *117* (48), 25769-25777.
173. Nouraddini, M. I.; Ranjbar, M.; Dobson, P. J.; Farrokhpour, H.; Johnston, C.; Jurkschat, K., Stabilized gold nanoparticles by laser ablation in ferric chloride solutions. *Applied Physics a-Materials Science & Processing* **2017**, *123* (12).
174. Masitas, R. A.; Zamborini, F. P., Oxidation of Highly Unstable < 4 nm Diameter Gold Nanoparticles 850 mV Negative of the Bulk Oxidation Potential. *Journal of the American Chemical Society* **2012**, *134* (11), 5014-5017.
175. Yalcin, M.; Gurses, A.; Dogar, C.; Sozbilir, M., The adsorption kinetics of Cethyltrimethylammonium Bromide (CTAB) onto powdered active carbon. *Adsorption-Journal of the International Adsorption Society* **2004**, *10* (4), 339-348.
176. Woods, D. A.; Petkov, J.; Bain, C. D., Surfactant Adsorption Kinetics by Total Internal Reflection Raman Spectroscopy. 2. CTAB and Triton X-100 Mixtures on Silica. *Journal of Physical Chemistry B* **2011**, *115* (22), 7353-7363.
177. Lee, S.; Anderson, L. J. E.; Payne, C. M.; Hafner, J. H., Structural Transition in the Surfactant Layer that Surrounds Gold Nanorods as Observed by Analytical Surface-Enhanced Raman Spectroscopy. *Langmuir* **2011**, *27* (24), 14748-14756.
178. Bellino, M. G.; Calvo, E. J.; Gordillo, G., Adsorption kinetics of charged thiols on gold nanoparticles. *Physical Chemistry Chemical Physics* **2004**, *6* (2), 424-428.
179. Park, J. W.; Shumaker-Parry, J. S., Strong Resistance of Citrate Anions on Metal Nanoparticles to Desorption under Thiol Functionalization. *Acs Nano* **2015**, *9* (2), 1665-1682.
180. Kundu, J.; Levin, C. S.; Halas, N. J., Real-time monitoring of lipid transfer between vesicles and hybrid bilayer on Au nanoshells using surface enhanced Raman scattering (SERS). *Nanoscale* **2009**, *1* (1), 114-117.
181. Kawasaki, H.; Nishimura, K.; Arakawa, R., Influence of the counterions of cetyltrimethylammonium salts on the surfactant adsorption onto gold surfaces and the formation of gold nanoparticles. *Journal of Physical Chemistry C* **2007**, *111* (6), 2683-2690.

182. Knag, M.; Sjoblom, J.; Gulbrandsen, E., The effect of straight chain alcohols and ethylene glycol on the adsorption of CTAB on gold. *Journal of Dispersion Science and Technology* **2005**, *26* (2), 207-215.
183. Rohatgi, A. WebPlotDigitizer Version 4.2. <https://automeris.io/WebPlotDigitizer> (accessed August 27th, 2019).
184. Nikoobakht, B.; El-Sayed, M. A., Evidence for bilayer assembly of cationic surfactants on the surface of gold nanorods. *Langmuir* **2001**, *17* (20), 6368-6374.
185. Janicek, B. E.; Hinman, J. G.; Hinman, J. J.; Bae, S. h.; Wu, M.; Turner, J. G.; Chang, H.-H.; Park, E.; Lawless, R.; Suslick, K. S.; Murphy, C. J.; Huang, P. Y., Quantitative imaging of organic ligand density on anisotropic inorganic nanocrystals. *Nano Letters* **2019**.
186. Cheng, X. Y.; Anthony, T. P.; West, C. A.; Hu, Z. W.; Sundaresan, V.; McLeod, A. J.; Masiello, D. J.; Willets, K. A., Plasmon Heating Promotes Ligand Reorganization on Single Gold Nanorods. *Journal of Physical Chemistry Letters* **2019**, *10* (6), 1394-1401.
187. Zhang, Z.; Lin, M. S., Fast loading of PEG-SH on CTAB-protected gold nanorods. *Rsc Advances* **2014**, *4* (34), 17760-17767.
188. Gao, H. M.; Liu, H.; Qian, H. J.; Jiao, G. S.; Lu, Z. Y., Multiscale simulations of ligand adsorption and exchange on gold nanoparticles. *Physical Chemistry Chemical Physics* **2018**, *20* (3), 1381-1394.
189. Yorulmaz, M.; Khatua, S.; Zijlstra, P.; Gaiduk, A.; Orrit, M., Luminescence Quantum Yield of Single Gold Nanorods. *Nano Letters* **2012**, *12* (8), 4385-4391.
190. Water - Thermal Conductivity. https://www.engineeringtoolbox.com/water-liquid-gas-thermal-conductivity-temperature-pressure-d_2012.html (accessed September 9th).
191. Liaw, J. W.; Lo, W. J.; Kuo, M. K., Wavelength-dependent longitudinal polarizability of gold nanorod on optical torques. *Opt. Express* **2014**, *22* (9), 10858-10867.
192. Liaw, J. W.; Lo, W. J.; Lin, W. C.; Kuo, M. K., Theoretical study of optical torques for aligning Ag nanorods and nanowires. *Journal of Quantitative Spectroscopy & Radiative Transfer* **2015**, *162*, 133-142.
193. Liaw, J. W.; Chen, Y. S.; Kuo, M. K., Rotating Au nanorod and nanowire driven by circularly polarized light. *Opt. Express* **2014**, *22* (21), 26005-26015.
194. Shao, L.; Yang, Z. J.; Andren, D.; Johansson, P.; Kall, M., Gold Nanorod Rotary Motors Driven by Resonant Light Scattering. *Acs Nano* **2015**, *9* (12), 12542-12551.
195. Andren, D.; Shao, L.; Lank, N. O.; Acimovic, S. S.; Johansson, P.; Kall, M., Probing Photothermal Effects on Optically Trapped Gold Nanorods by Simultaneous Plasmon Spectroscopy and Brownian Dynamics Analysis. *Acs Nano* **2017**, *11* (10), 10053-10061.
196. Gaiduk, A.; Ruijgrok, P. V.; Yorulmaz, M.; Orrit, M., Detection limits in photothermal microscopy. *Chemical Science* **2010**, *1* (3), 343-350.

197. Pinnow, D. A.; Rich, T. C.; Ostermayer, F. W.; DiDomenico, M., Fundamental optical attenuation limits in the liquid and glassy state with application to fiber optical waveguide materials. *Applied Physics Letters* **1973**, *22* (10), 527-529.
198. Pope, R. M.; Fry, E. S., Absorption spectrum (380–700 nm) of pure water. II. Integrating cavity measurements. *Appl. Opt.* **1997**, *36* (33), 8710-8723.
199. Frank, B. P.; Sigmon, L. R.; Deline, A. R.; Lankone, R. S.; Gallagher, M. J.; Zhi, B.; Haynes, C. L.; Fairbrother, D. H., Photochemical Transformations of Carbon Dots in Aqueous Environments. *Environmental Science & Technology* **2020**, *54* (7), 4160-4170.
200. Zhi, B.; Gallagher, M. J.; Frank, B. P.; Lyons, T. Y.; Qiu, T. A.; Da, J.; Mensch, A. C.; Hamers, R. J.; Rosenzweig, Z.; Fairbrother, D. H.; Haynes, C. L., Investigation of phosphorous doping effects on polymeric carbon dots: Fluorescence, photostability, and environmental impact. *Carbon* **2018**, *129*, 438-449.
201. Sun, Y.-P.; Zhou, B.; Lin, Y.; Wang, W.; Fernando, K. A. S.; Pathak, P.; Mezziani, M. J.; Harruff, B. A.; Wang, X.; Wang, H.; Luo, P. G.; Yang, H.; Kose, M. E.; Chen, B.; Veca, L. M.; Xie, S.-Y., Quantum-Sized Carbon Dots for Bright and Colorful Photoluminescence. *Journal of the American Chemical Society* **2006**, *128* (24), 7756-7757.
202. Fathi, P.; Khamo, J. S.; Huang, X.; Srivastava, I.; Esch, M. B.; Zhang, K.; Pan, D., Bulk-state and single-particle imaging are central to understanding carbon dot photo-physics and elucidating the effects of precursor composition and reaction temperature. *Carbon* **2019**, *145*, 572-585.
203. Li, H.; Kang, Z.; Liu, Y.; Lee, S.-T., Carbon nanodots: synthesis, properties and applications. *Journal of Materials Chemistry* **2012**, *22* (46), 24230-24253.
204. Zhu, S.; Song, Y.; Shao, J.; Zhao, X.; Yang, B., Non-Conjugated Polymer Dots with Crosslink-Enhanced Emission in the Absence of Fluorophore Units. *Angewandte Chemie International Edition* **2015**, *54* (49), 14626-14637.
205. Zhu, S.; Song, Y.; Zhao, X.; Shao, J.; Zhang, J.; Yang, B., The photoluminescence mechanism in carbon dots (graphene quantum dots, carbon nanodots, and polymer dots): current state and future perspective. *Nano Research* **2015**, *8* (2), 355-381.
206. He, C.; Xu, P.; Zhang, X.; Long, W., The synthetic strategies, photoluminescence mechanisms and promising applications of carbon dots: Current state and future perspective. *Carbon* **2022**, *186*, 91-127.
207. Zhi, B.; Yao, X.; Wu, M.; Mensch, A.; Cui, Y.; Deng, J.; Duchimaza-Heredia, J. J.; Trerayapiwat, K. J.; Niehaus, T.; Nishimoto, Y.; Frank, B. P.; Zhang, Y.; Lewis, R. E.; Kappel, E. A.; Hamers, R. J.; Fairbrother, H. D.; Orr, G.; Murphy, C. J.; Cui, Q.; Haynes, C. L., Multicolor polymeric carbon dots: synthesis, separation and polyamide-supported molecular fluorescence. *Chemical Science* **2021**, *12* (7), 2441-2455.
208. Das, S. K.; Liu, Y.; Yeom, S.; Kim, D. Y.; Richards, C. I., Single-Particle Fluorescence Intensity Fluctuations of Carbon Nanodots. *Nano Letters* **2014**, *14* (2), 620-625.

209. Ghosh, S.; Ghosh, A.; Ghosh, G.; Marjit, K.; Patra, A., Deciphering the Relaxation Mechanism of Red-Emitting Carbon Dots Using Ultrafast Spectroscopy and Global Target Analysis. *The Journal of Physical Chemistry Letters* **2021**, *12* (33), 8080-8087.
210. Wang, L.; Zhu, S.-J.; Wang, H.-Y.; Wang, Y.-F.; Hao, Y.-W.; Zhang, J.-H.; Chen, Q.-D.; Zhang, Y.-L.; Han, W.; Yang, B.; Sun, H.-B., Unraveling Bright Molecule-Like State and Dark Intrinsic State in Green-Fluorescence Graphene Quantum Dots via Ultrafast Spectroscopy. *Advanced Optical Materials* **2013**, *1* (3), 264-271.
211. Wen, X.; Yu, P.; Toh, Y.-R.; Hao, X.; Tang, J., Intrinsic and Extrinsic Fluorescence in Carbon Nanodots: Ultrafast Time-Resolved Fluorescence and Carrier Dynamics. *Advanced Optical Materials* **2013**, *1* (2), 173-178.
212. Xu, Q.; Zhou, Q.; Hua, Z.; Xue, Q.; Zhang, C.; Wang, X.; Pan, D.; Xiao, M., Single-Particle Spectroscopic Measurements of Fluorescent Graphene Quantum Dots. *ACS Nano* **2013**, *7* (12), 10654-10661.
213. Khan, S.; Li, W.; Karedla, N.; Thiant, J.; Gregor, I.; Chizhik, A. M.; Enderlein, J.; Nandi, C. K.; Chizhik, A. I., Charge-Driven Fluorescence Blinking in Carbon Nanodots. *The Journal of Physical Chemistry Letters* **2017**, *8* (23), 5751-5757.
214. Zhi, B.; Cui, Y.; Wang, S.; Frank, B. P.; Williams, D. N.; Brown, R. P.; Melby, E. S.; Hamers, R. J.; Rosenzweig, Z.; Fairbrother, D. H.; Orr, G.; Haynes, C. L., Malic Acid Carbon Dots: From Super-resolution Live-Cell Imaging to Highly Efficient Separation. *ACS Nano* **2018**, *12* (6), 5741-5752.
215. Leménager, G.; De Luca, E.; Sun, Y.-P.; Pompa, P. P., Super-resolution fluorescence imaging of biocompatible carbon dots. *Nanoscale* **2014**, *6* (15), 8617-8623.
216. Chizhik, A. M.; Stein, S.; Dekaliuk, M. O.; Battle, C.; Li, W.; Huss, A.; Platen, M.; Schaap, I. A. T.; Gregor, I.; Demchenko, A. P.; Schmidt, C. F.; Enderlein, J.; Chizhik, A. I., Super-Resolution Optical Fluctuation Bio-Imaging with Dual-Color Carbon Nanodots. *Nano Letters* **2016**, *16* (1), 237-242.
217. Nguyen, H. A.; Srivastava, I.; Pan, D.; Gruebele, M., Unraveling the Fluorescence Mechanism of Carbon Dots with Sub-Single-Particle Resolution. *ACS Nano* **2020**, *14* (5), 6127-6137.
218. Nguyen, H. A.; Srivastava, I.; Pan, D.; Gruebele, M., Ultrafast nanometric imaging of energy flow within and between single carbon dots. *Proceedings of the National Academy of Sciences* **2021**, *118* (11), e2023083118.
219. Ghosh, S.; Chizhik, A. M.; Karedla, N.; Dekaliuk, M. O.; Gregor, I.; Schuhmann, H.; Seibt, M.; Bodensiek, K.; Schaap, I. A. T.; Schulz, O.; Demchenko, A. P.; Enderlein, J.; Chizhik, A. I., Photoluminescence of Carbon Nanodots: Dipole Emission Centers and Electron-Phonon Coupling. *Nano Letters* **2014**, *14* (10), 5656-5661.
220. Wei, Z.; Wang, B.; Hong, D.; Xie, M.; Wan, S.; Yang, W.; Lu, S.; Tian, Y., Rational Building of Nonblinking Carbon Dots via Charged State Recovery. *The Journal of Physical Chemistry Letters* **2021**, *12* (35), 8614-8620.

221. Das, A.; Gude, V.; Roy, D.; Chatterjee, T.; De, C. K.; Mandal, P. K., On the Molecular Origin of Photoluminescence of Nonblinking Carbon Dot. *The Journal of Physical Chemistry C* **2017**, *121* (17), 9634-9641.
222. Ghosh, S.; Oleksiievets, N.; Enderlein, J.; Chizhik, A. I., Emission States Variation of Single Graphene Quantum Dots. *The Journal of Physical Chemistry Letters* **2020**, *11* (17), 7356-7362.