

Photonic Technologies to Enable Slow Light Applications

by

Joseph E. Vornehm Jr.

Submitted in Partial Fulfillment of the

Requirements for the Degree

Doctor of Philosophy

Supervised by Professor Robert W. Boyd

The Institute of Optics
Arts, Sciences and Engineering
Edmund A. Hajim School of Engineering and Applied Sciences

University of Rochester
Rochester, New York

2014

To Meta and Lili:

We did it.

Biographical Sketch

Joseph E. Vornehm Jr. was born in Boonton, New Jersey. He attended Northeastern University (Boston, Massachusetts) and graduated in 2001 with a Bachelor of Science degree in electrical and computer engineering. During his undergraduate program he completed internships with IBM Microelectronics (Burlington, Vermont) and The MITRE Corporation (Bedford, Massachusetts). After graduation, he worked for The MITRE Corporation as a software systems engineer in the Signal Processing department from 2001 to 2003. He then attended Northwestern University, graduating in 2005 with a Master of Science degree in electrical and computer engineering. His thesis was titled “Multi-Spectral Raman Gain in Dual-Isotope Rubidium Vapor” and was supervised by Professor M. Selim Shahriar. He began doctoral studies in optics at the University of Rochester in 2005, where he was awarded a Sproull Fellowship from 2005 to 2007. He pursued his research in optics under the direction of Professor Robert W. Boyd.

The following publications were a result of work conducted during doctoral study (reference numbers match those in the References):

- [1] J. E. Vornehm and R. W. Boyd, “Slow and fast light,” in *Tutorials in Complex Photonic Media*, M. A. Noginov, G. Dewar, M. W. McCall, and N. I. Zheludev, eds. (SPIE, Bellingham, WA, USA, 2009), pp. 647–685.
- [2] Z. Shi, A. Schweinsberg, J. E. Vornehm Jr., M. A. Martínez Gámez, and R. W. Boyd, “Low distortion, continuously tunable, positive and negative time delays by slow and fast light using stimulated Brillouin scattering,” *Phys. Lett. A* **374**, 4071–4074 (2010).
- [3] A. Schweinsberg, Z. Shi, J. E. Vornehm, and R. W. Boyd, “Demonstration of a slow-light laser radar,” *Opt. Express* **19**, 15760–15769 (2011).
- [4] A. Schweinsberg, Z. Shi, J. E. Vornehm, and R. W. Boyd, “A slow-light laser radar system with two-dimensional scanning,” *Opt. Lett.* **37**, 329–331 (2012).
- [5] Z. Shi, A. Schweinsberg, J. E. Vornehm, and R. W. Boyd, “A slow-light laser radar (SLIDAR),” *Opt. Photon. News* **23**, 51 (2012).
- [6] J. E. Vornehm, A. Schweinsberg, Z. Shi, D. J. Gauthier, and R. W. Boyd, “Phase locking of multiple optical fiber channels for a slow-light-enabled laser radar system,” *Opt. Express* **21**, 13094–13104 (2013).
- [7] S. Murugkar, I. De Leon, Z. Shi, G. Lopez-Galmiche, J. Salvail, E. Ma, B. Gao, A. C. Liapis, J. E. Vornehm, and R. W. Boyd, “Development of a slow-light spectrometer on a chip,” *Proc. SPIE* **8264**, 82640T (2012).

Acknowledgments

I am deeply grateful for Prof. Robert W. Boyd's mentorship, guidance, and good humor during the course of my graduate study. It has been an exceptional opportunity and a real privilege for me to have him as a research mentor and to observe his example during my studies. Much of science is learned from textbooks, but much more is learned from watching great scientists. Thanks to my committee members, Profs. Andrew Berger, John Howell, and Carlos Stroud, for helpful discussion, thought-provoking questions, and encouragement. I also want to thank our research group members, past and present: Dr. Kam Wai (Cliff) Chan, Dr. Hye Jeong Chang, Dr. Yuping Chen, Prof. Ksenia Dolgaleva, Boshen Gao, Prof. George Gehring, Prof. Anand Jha, Dr. Jerry Kuper, Andreas Liapis, Omar Magaña Loaiza, Prof. Svetlana Lukishova, Mehul Malik, Mohammad Mirhosseini, Colin O'Sullivan, Dr. Giovanni Piredda, Dr. Alex Radunsky, Brandon Rodenburg, Dr. Aaron Schweinsberg, Prof. Zhimin Shi, Dr. Heedeuk Shin, Mahmud Siddiqui, and Dr. Petros Zerom, as well as Ava (Jingwen) Dong and our many summer students and visitors, and our Ottawa group members, especially Dr. Israel De Leon, Prof. Jonathan Leach, Prof. Sangeeta Murugkar, and Dr. Jer Upham.

The Institute of Optics is a singular place, and I am grateful to the faculty for teaching me optics and for the collegial feeling that has been engendered at the Institute. Thanks to the staff for being a welcoming face and a helping hand, including Per Adamson, Betsy Benedict, Joan Christian, Lissa Cotter, Lynn Doescher, Gina Kern, Brian McIntyre, Lori Russell, Barbara Schirmer, Maria Schnitzler, Evelyn Sheffer, Dan Smith, Gayle Thompson, and Noelene Votens, as well as Hajim School staff, and Hugo Bégin, Elvira Evangelista, and Kristelle Lapointe at the University of Ottawa.

The friends I have made in Rochester, and especially at the Institute, will be lifelong friends. I am blessed with too many friends to name, each of whom I cherish. Dr. Kathleen Adelsberger, Dr. Ryan and Bethany Beams, Dr. Amber Beckley, Dr. Brooke Beier, Dr. Luke and Niki Bissell, Dr. Dean and Lynn Brown, Dr. Cristina Canavesi and Andrea Cogliati, Joe and Sora Choi, Dr. Dan and Lisa Christensen, Dr. Eric and Katie Christensen, Dr. Liping Cui, Maj. Jack and Angie DeLong, Dr. Yijing Fu, Dr. Ying (Melissa) Geng, Dr. Tammy Lee, Jordan Leidner, Dr. Suzanne Leslie and Dr. Brad Deutsch, Dr. Ben Masella, Leva McIntire, Dr. Ramkumar Sabesan, Dr. Josh and Jess Schoenly, Dustin and Laura Shipp, Dr. Manuel Guizar Sicaïros and Paloma Ayala Nuñez (and Sebastian Guizar Ayala), Richard Smith, Mike Theisen, Dr. Becky Wilson, Haomin Yao, Yuhong Yao, Aizhong Zhang, Len Zheleznyak, and many, many others: thank you. Chris Todd, you are missed.

The Graduate Writing Project (part of the Writing, Speaking, and Argument program of the College of Arts, Sciences and Engineering) has been immeasurably helpful in completing this dissertation, and I thank Rachel Lee and Liz Tinelli for organizing the program and

offering support, encouragement, and a community for writers. Many thanks to Dean Wendi Heinzelman for supporting this outstanding program.

I am grateful to the OSA Rochester Section for the opportunity to serve as an officer and a council member for several years, particularly past president Jen Kruschwitz for inviting me to volunteer and past presidents Rick Plympton, Julie Bentley, and Chris Palmer for inviting me to stay. It has been a unique opportunity to get to know the optics community and to make a small contribution. For the same reasons, I thank Dirk Fabian for several opportunities with SPIE.

Thanks to my parents, Joseph E. Vornehm Sr. and Marsha Vornehm; to my sister and her family, Becky, Tom, and Patrick Sanderson; and to my in-laws, Sandip and Maida Sengupta and Shona and Ken Armstrong, for teaching me the value of education, cheering me on, and providing a place for me to finish writing.

I cannot thank my wife, Meta, and our daughter, Lili, enough. Neither words nor grateful tears nor happy smiles suffice. This dissertation is for you—we did it. I will always love you.

Lastly, I thank God for life, for countless blessings, for the opportunity to pursue a PhD, and for the miracles that have led to its completion. “With God all things are possible” (Matthew 19:26).

Abstract

Slow light is light that travels at unusual, extreme group velocities—sometimes as slow as walking speed or slower. Light waves can be described by many velocities, but for a narrow-band pulse of light (a carrier frequency modulated by an envelope), the group velocity is the speed of the pulse envelope. The term *slow light* also encompasses other exotic group velocities due to similar techniques, including fast light, stopped light, and backwards light. The science of slow light has been established over the past several years, and research attention is now turning to potential applications of slow light.

One application of slow light is as an all-optical true-time delay. Two slow light methods in optical fibers, stimulated Brillouin scattering (SBS) and dispersive delay, are used to provide a controllable pulse delay in a prototype slow-light phased-array laser radar, called SLIDAR. These slow light methods compensate the group delay mismatch of pulses of 6 ns duration while the phased array is steered in two dimensions. A phase control system is described that maintains phase lock among three signal channels and a reference channel, each containing 2.2 km of optical fiber, while accommodating the demands of the slow light techniques. Residual phase error is kept below $\pi/5$ radians (1/10 wave) RMS.

Slow light can also enhance the spectral sensitivity of spectrometers and interferometers. A design for a slow-light-enhanced nanophotonic spectrometer is presented. One important use of spectrometers is to detect specific chemicals, and I describe an approach to multivariate optical computation, which can be used for automatic chemical spectrum recognition. This technique, which could one day be implemented with a slow-light-enhanced spectrometer as a single-chip chemical detection platform, is explored experimentally in the visible spectrum using a spatial light modulator.

Contributors and Funding Sources

This work was supervised by a dissertation committee consisting of Professors Robert W. Boyd (advisor), Andrew Berger, and Carlos Stroud of the Institute of Optics (Edmund A. Hajim School of Engineering and Applied Sciences), and Professor John Howell of the Department of Physics (College of Arts, Sciences and Engineering). Prof. Boyd also holds an appointment at the University of Ottawa, Ottawa, Canada.

Chapter 1, “Slow and Fast Light,” is an extensive revision of a 2009 slow light review listed in the Biographical Sketch [1]. I wrote the 2009 review, under the supervision of Prof. Boyd. Text is reused from that publication by permission of SPIE Press. Chapter 1 briefly describes two projects that are covered by other publications in the Biographical Sketch, namely a slow-light laser radar called SLIDAR [2–6] and a slow-light-enhanced nanophotonic spectrometer [7]. SLIDAR was a systems engineering project, and interlocking contributions were made by Prof. Boyd, George Gehring, Andreas Liapis, Aaron Schweinsberg, Zhimin Shi, and myself. I developed the phase control system described in chapter 2. I also contributed to the overall system design and to system-level testing and data collection, including collection of the results shown in Fig. 1.5. The nanophotonic

spectrometer was tested partly using software I wrote; although that contribution is not covered in this dissertation, my work in chapter 3 is connected to that research effort.

Chapter 2, “Phase Control of a Slow-Light Laser Radar,” is taken from a 2013 paper listed in the Biographical Sketch [6], with a new introduction to suit the dissertation. The work in the chapter is my own, with the following exceptions: Prof. Dan Gauthier (Duke University) suggested the use of an integrating op-amp configuration. Aaron Schweinsberg, Zhimin Shi, and I jointly tested the system and collected the data. The dispersion compensating module was loaned by Corning, Inc., and the infrared camera used to record Fig. 2.6 was loaned by Ed Watson (now retired) and Larry Barnes (Wright–Patterson Air Force Base). Text and figures are reused from the 2013 paper by permission of OSA.

Chapter 3, “Multivariate Optical Computing for Spectrum Recognition,” is my own work. Ava (Jingwen) Dong assisted in building the apparatus, and Per Adamson loaned the doublet lens used in the experiment. I acknowledge helpful discussions with Zach Smith (University of California, Davis), and especially Zach’s loan of his computer code and data to help me understand his approach to multivariate optical computing.

Chapter 4, “Conclusions and Future Work,” is my own work. The idea to use a waveguide and a series of resonators for a nanophotonic implementation of a multivariate optical computer came out of a discussion with Andreas Liapis.

Any text or figures in chapter 1 and chapter 2 that were previously published have been reused by permission from the publishers (SPIE Press and OSA) under their standard author

rights clauses. Text sources are as described above; figure sources are cited in figure captions. (Figures without citations are my own work.)

My graduate study was supported by a Sproull Fellowship from the University of Rochester. The work in chapter 2 was supported by the DARPA/DSO Slow Light program. The work in chapter 3 was supported by the US Defense Threat Reduction Agency's Joint Science and Technology Office for Chemical and Biological Defense (DTRA JSTO-CBD) under grant number HDTRA1-10-1-0025.

Table of Contents

Biographical Sketch	iii
Acknowledgments	v
Abstract	viii
Contributors and Funding Sources	x
List of Tables	xviii
List of Figures	xix
1 Slow and Fast Light	1
1.1 The origins of slow light	2
1.1.1 Definition of the phase velocity	2
1.1.2 Definition of the group velocity	3
1.1.3 Material vs. structural slow light	8

1.2	Material slow light	8
1.2.1	Susceptibility and the Kramers–Kronig relations	8
1.2.2	Resonance features in materials	11
1.2.3	Stimulated Brillouin scattering in optical fibers	13
1.2.4	Dispersive delay	16
1.2.5	Other material slow light phenomena	18
1.3	Structural slow light	20
1.3.1	Photonic crystals and photonic bandgap structures	20
1.3.2	Other structural slow light systems	24
1.4	Additional considerations	24
1.4.1	Figures of merit	25
1.4.2	Theoretical limits of slow and fast light	26
1.4.3	Spatial compression and nonlinearity enhancement	27
1.4.4	Causality of fast light	28
1.5	Applications of slow and fast light	30
1.5.1	A slow-light laser radar (SLIDAR)	31
1.5.2	A slow-light-enhanced nanophotonic spectrometer	35
1.5.3	Other applications of slow and fast light	37
1.6	Conclusion	41

2	Phase Control of a Slow-Light Laser Radar	42
2.1	Theory	45
2.1.1	The effect of residual phase error	47
2.1.2	The effect of snapbacks	49
2.2	Experimental apparatus	54
2.2.1	Optical system	54
2.2.2	Electronics	56
2.3	Results and discussion	59
2.4	Conclusion	62
3	Multivariate Optical Computing for Spectrum Recognition	63
3.1	Spectral correlation	66
3.1.1	Measuring concentrations and mixtures	68
3.1.2	Implementation on an SLM	70
3.1.3	MOCs vs. traditional spectrometers	73
3.2	Experimental setup	76
3.2.1	Source and filters	77
3.2.2	Broadband $2F-4F$ system	79
3.2.3	SLM operation	81
3.2.4	Detection	82

3.2.5	Control and processing	83
3.2.6	Optical design tradeoffs	84
3.2.7	Preliminary measurements	86
3.3	Results and discussion	88
3.3.1	Spectrum recognition results	88
3.3.2	Spectral scans	89
3.3.3	Correcting erroneous data	93
3.3.4	Experimental issues	94
3.4	Multiple-output multivariate optical computation	97
3.4.1	Results	99
3.4.2	Experimental issues	100
3.4.3	Limit on number of simultaneous matched filters	101
3.5	Conclusion	103
4	Conclusions and Future Work	104
4.1	Phase control of a slow-light laser radar	104
4.2	Multivariate optical computing for spectrum recognition	107
	References	112

Appendix A	SLM Calibration	128
A.1	Two-pixel SLM calibration	129
A.2	SLM phase model	130
A.3	Multiple-output SLM calibration	134
A.4	Calibration issues	135

List of Tables

3.1	Preliminary experiment correlation results	88
3.2	Spectrum recognition results	89
3.3	Multiple-output spectrum recognition results with a single SLM zone . . .	100
3.4	Multiple-output spectrum recognition results with two SLM zones	100

List of Figures

1.1	Dispersive features of an absorption resonance	12
1.2	SBS generator and amplifier configurations in single-mode fiber	14
1.3	Three-level energy level model for EIT	18
1.4	Photonic crystal structures	21
1.5	SLIDAR: a pulsed, phased-array slow-light laser radar	33
1.6	A design for a slow-light-enhanced nanophotonic spectrometer	36
2.1	Phase noise before and after correction by the phase control system	47
2.2	SLIDAR optical system schematic diagram	55
2.3	Block diagram of the SLIDAR phase control electronics	57
2.4	Block diagram of the proportional-integral (P-I) controller	58
2.5	RMS residual phase error vs. approximate SBS gain	59
2.6	Far-field pattern of the three-channel SLIDAR system	61
3.1	Multivariate optical computer (MOC) experiment diagram	77

3.2	Spectra of the white light source and filters	78
3.3	Preliminary MOC experiment diagram	87
3.4	Spectral scans of the white light source and filters	90
3.5	Transmission spectra of didymium and holmium filters	91
3.6	Matched filter spectra \mathbf{a}_k	93
3.7	Multiple-output MOC experiment diagram	99

Chapter 1

Slow and Fast Light

In early 1999, a news article in the prestigious journal *Nature* led off with the announcement, “An experiment with atoms at nanokelvin temperatures has produced the remarkable observation of light pulses traveling at velocities of only 17 m s^{-1} .” The review continued with the understatement, “Observation of light pulses propagating at a speed no faster than a swiftly moving bicycle. . . comes as a surprise” [8]. These findings marked the beginning of the current wave of interest in the field that has come to be called *slow light* [9].

Since light is a wave, it can be described by many different velocities. The term *speed of light* most commonly means the phase velocity, written as c/n , where c is the speed of light in vacuum and n is the refractive index. Stated precisely, this is the speed at which the phase fronts of a monochromatic plane wave propagate in a uniform, homogeneous medium. But real signals are not monochromatic. A narrow-band signal can be modeled as a central frequency, often called a carrier frequency, with a small spread of frequencies around it. Since optical frequencies are of the order of 10^{15} Hz, nearly all optical signals

can be modeled in this way. In the time domain, a narrow-band signal is a carrier frequency modulated by an envelope (such as a Gaussian pulse). The speed at which the envelope moves is the *group velocity*, meaning the velocity of a group of waves at different frequencies. Other velocities may be defined as well [10–13].

Slow light is a term used to describe light traveling at unusual, extreme group velocities, most commonly at group velocities well below c/n . The term generally also encompasses *fast light*, *stopped light*, and *backwards light*, or group velocities that are respectively greater than c/n , near zero, and negative. While slow light produces surprising and often counterintuitive results, it does not violate causality or Maxwell's equations [12].

1.1 The origins of slow light

1.1.1 Definition of the phase velocity

Consider the (complex scalar) electric field of a monochromatic electromagnetic plane wave of amplitude E_0 propagating in the $+z$ direction,

$$E(z,t) = E_0 e^{i\phi}, \quad \phi = kz - \omega t. \quad (1.1)$$

Here, ω is the angular frequency of the plane wave, k is the wavevector, and t is the time.

(The real electric field is equal to the complex electric field plus its complex conjugate.)

When one speaks of the propagation of the plane wave, one means the motion of the wave's

phase fronts, or surfaces defined by constant values of ϕ . The present goal is to observe the motion of one such phase front; a convenient choice is the phase front located at the origin $z = 0$ at time $t = 0$ (such that $\phi = 0$). Its motion is governed by

$$kz - \omega t = 0. \quad (1.2)$$

To account for dispersion, the wavevector k and the refractive index of the medium n are written as functions of ω ,

$$k(\omega) = \frac{n(\omega)\omega}{c}. \quad (1.3)$$

Equations (1.2) and (1.3) can be used to compute dz/dt , which is the phase velocity v , or the speed of propagation of the phase front:

$$\frac{dz}{dt} = \frac{\omega}{k(\omega)} = \frac{c}{n(\omega)} \equiv v(\omega). \quad (1.4)$$

It is important to note that the phase velocity depends on frequency, meaning that monochromatic plane waves at different frequencies generally travel at different speeds.

1.1.2 Definition of the group velocity

The group velocity of a waveform is defined as

$$v_g(\omega) = \frac{d\omega}{dk}. \quad (1.5)$$

When the wavevector is expressed as $k = n(\omega)\omega/c$, the derivative can be rewritten as

$$v_g(\omega) = \left(\frac{dk}{d\omega} \right)^{-1} = \frac{c}{n(\omega) + \omega \frac{dn(\omega)}{d\omega}} = \frac{c}{n_g(\omega)}, \quad (1.6)$$

where the quantity

$$n_g(\omega) = n(\omega) + \omega \frac{dn(\omega)}{d\omega} \quad (1.7)$$

is called the *group index*, by analogy to the refractive index (since $v_g = c/n_g$ just as $v = c/n$).

Now that the group velocity has been defined, it is important to see the role that the group velocity plays in the propagation of a waveform. The Fourier theorem can be applied here to represent the waveform as a sum of monochromatic plane waves. (The constraints of the Fourier theorem are neglected here, since they are satisfied for any situation of interest.) For simplicity, assume that the waveform is propagating in the $+z$ direction through a linear, homogeneous, isotropic, dispersive and dissipative medium. The complex electric field of the waveform is then

$$E(z, t) = \sum_j E_j \exp \left\{ -i\omega_j \left[t - \frac{n(\omega_j)z}{c} \right] - \frac{\alpha(\omega)z}{2} \right\}. \quad (1.8)$$

Each plane wave has a frequency ω_j and a complex amplitude E_j . The quantities $n(\omega)$ and $\alpha(\omega)$ are respectively the real refractive index and intensity absorption coefficient of the

medium. (The sum in Eq. (1.8) could also be written as a Fourier transform integral, but the conceptual points that follow remain the same.)

What is required for the waveform to maintain its shape while it propagates? If the waveform propagates a distance Δz in a time Δt , its shape is preserved *only* if $E(z, t)$ and $E(z + \Delta z, t + \Delta t)$ are related by a complex scaling constant $re^{i\phi}$,

$$E(z + \Delta z, t + \Delta t) = re^{i\phi} E(z, t), \quad (1.9)$$

where r and ϕ are real parameters representing the amplitude change and phase change due to propagation. Combining the description of the waveform in Eq. (1.8) with the constraint of Eq. (1.9) gives

$$\begin{aligned} \sum_j E_j \exp \left\{ -i\omega_j \left[(t + \Delta t) - \frac{n(\omega_j)(z + \Delta z)}{c} \right] - \frac{\alpha(\omega_j)(z + \Delta z)}{2} \right\} \\ = re^{i\phi} \sum_j E_j \exp \left\{ -i\omega_j \left[t - \frac{n(\omega_j)z}{c} \right] - \frac{\alpha(\omega_j)z}{2} \right\}. \end{aligned} \quad (1.10)$$

Rewriting gives

$$\begin{aligned} \sum_j E_j \exp \left\{ -i\omega_j \left[t - \frac{n(\omega_j)z}{c} \right] - \frac{\alpha(\omega_j)z}{2} \right\} \exp \left\{ -i\omega_j \left[\Delta t - \frac{n(\omega_j)\Delta z}{c} \right] - \frac{\alpha(\omega_j)\Delta z}{2} \right\} \\ = re^{i\phi} \sum_j E_j \exp \left\{ -i\omega_j \left[t - \frac{n(\omega_j)z}{c} \right] - \frac{\alpha(\omega_j)z}{2} \right\}. \end{aligned} \quad (1.11)$$

Note that Eq. (1.11) has the form $\sum_j A_j B_j = c \sum_j A_j$. In order for such an equation to hold true under the most general conditions, each of the B_j must be equal to c . In other words,

$$\exp \left\{ -i\omega_j \left[\Delta t - \frac{n(\omega_j)\Delta z}{c} \right] - \frac{\alpha(\omega_j)\Delta z}{2} \right\} = r e^{i\phi}, \quad (1.12)$$

and this must be true for all ω_j —for a wave at any frequency. (The j subscript is omitted going forward.) Since all of the parameters in Eq. (1.12) are real, the equation can be separated by magnitude and phase, giving

$$r = \exp \left[-\frac{\alpha(\omega)\Delta z}{2} \right] \quad (1.13)$$

and

$$e^{i\phi} = \exp \left\{ -i\omega \left[\Delta t - \frac{n(\omega)\Delta z}{c} \right] \right\}. \quad (1.14)$$

In order for the waveform to propagate unchanged and maintain its shape, r and $e^{i\phi}$ must not depend on ω . These constraints lead to a number of important insights. First, the absorption $\alpha(\omega)$ may not vary with frequency. Second, the derivative of $e^{i\phi}$ with respect to ω must be zero. This can only be true when $d\phi/d\omega$ is equal to zero,

$$\frac{d}{d\omega} \left\{ \omega \left[\Delta t - \frac{n(\omega)\Delta z}{c} \right] \right\} = \Delta t - \frac{n(\omega)\Delta z}{c} - \omega \left[\frac{\Delta z}{c} \frac{dn(\omega)}{d\omega} \right] = 0. \quad (1.15)$$

Some manipulation gives

$$\frac{\Delta z}{\Delta t} = \frac{c}{n + \omega \frac{dn(\omega)}{d\omega}} = \frac{c}{n_g(\omega)} = v_g(\omega). \quad (1.16)$$

This is the speed at which the waveform propagates, namely the group velocity $v_g(\omega)$. The propagation time Δt is also referred to as the *group delay*, τ_g . Note that for the same Δz and Δt to apply to all frequencies, group velocity and group index may not vary with frequency, meaning group velocity dispersion $dv_g/d\omega$ must be zero.

To ensure that an arbitrary waveform propagates in a linear medium with its shape unchanged, α and n_g must not vary with frequency. However, all physical systems have spectral variations in both the absorption coefficient and the refractive index; the variations in the refractive index inevitably lead to spectral variations in the group index as well. As a result, all systems, including slow and fast light systems, face some signal distortion. This distortion may range from simple pulse broadening to complex pulse breakup. The amount of distortion allowed sets the limits of the slow light system, usually in terms of maximum achievable pulse delay or maximum operating bandwidth. When designing a slow light system, the designer must first choose an acceptable level of distortion, and then adjust the system design to stay within this level.

1.1.3 Material vs. structural slow light

Slow light has been achieved in a wide array of optical media and systems [1]. In all cases, practitioners create slow light by carefully controlling dispersion. But the different kinds of dispersive systems can be broadly classified into two categories: material slow light and structural slow light. In material slow light systems, the medium can be described by a spatially uniform (but frequency-dependent) refractive index. Typically, the medium's constituent atoms or molecules provide the dispersive properties. Structural slow light systems have a periodically varying refractive index, with a period close to the optical wavelength. The refractive index variation is usually due to a repeating physical structure, such as regularly spaced air holes in a dielectric. Both approaches lead to slow light, but with some subtle differences [14]. Each approach has its merits and applications.

1.2 Material slow light

1.2.1 Susceptibility and the Kramers–Kronig relations

Both the refractive index $n(\omega)$ and the absorption coefficient $\alpha(\omega)$ of a medium have their origin in the susceptibility of the medium $\chi(\omega)$. When an electric field is applied to the medium, the charged particles in the medium (the electrons and protons) shift their positions in response to the field. This shift in the positions of the charges creates an additional electric field, represented by the polarization density P , or simply the polarization, measured in units of electric dipole moment per unit volume (C m m^{-3} , or C m^{-2}). Some materials are more

susceptible than others to being polarized by an incident electric field. The degree to which the material may be polarized by a given electric field is known as the electric susceptibility χ and is defined by

$$P(\omega) = \epsilon_0 \chi(\omega) E(\omega), \quad (1.17)$$

where $E(\omega)$ is the strength of the electric field at the frequency ω , and ϵ_0 is the permittivity of free space. The permittivity $\epsilon(\omega)$ of the medium is defined as

$$\epsilon(\omega) = \epsilon_0 \epsilon_r(\omega) = \epsilon_0 [1 + \chi(\omega)], \quad (1.18)$$

with $\epsilon_r(\omega)$ being the relative permittivity, and the refractive index and absorption coefficient are defined respectively as

$$n(\omega) = \text{Re} \{ \sqrt{\epsilon_r} \} = \text{Re} \{ \sqrt{1 + \chi(\omega)} \}, \quad (1.19)$$

$$\alpha(\omega) = \frac{2\omega}{c} \text{Im} \{ \sqrt{\epsilon_r} \} = \frac{2\omega}{c} \text{Im} \{ \sqrt{1 + \chi(\omega)} \}. \quad (1.20)$$

When $\chi(\omega)$ is small, such as for dilute gases, Eqs. (1.19) and (1.20) simplify to

$$n(\omega) \approx 1 + \frac{1}{2} \text{Re} \{ \chi(\omega) \}, \quad (1.21)$$

$$\alpha(\omega) \approx \frac{\omega}{c} \text{Im} \{ \chi(\omega) \}. \quad (1.22)$$

One can thus think of the real part of the susceptibility as corresponding to the refractive index $n(\omega)$ and the imaginary part as corresponding to the absorption coefficient $\alpha(\omega)$. The

correspondence is a useful conceptual device, but it is only fully accurate for media with weak responses (small susceptibilities).

The medium's electromagnetic response must be causal (it must obey causality). Any change in P at time t must be caused by changes in E that happen *before* time t . In other words, the cause must precede the effect. This may seem obvious, but the causality requirement has important consequences. It can be shown that the electromagnetic susceptibility of any causal medium obeys the Kramers–Kronig relations

$$\text{Im}\{\chi(\omega)\} = \frac{-2\omega}{\pi} \int_0^\infty \frac{\text{Re}\{\chi(\omega')\}}{\omega'^2 - \omega^2} d\omega', \quad (1.23)$$

$$\text{Re}\{\chi(\omega)\} = \frac{2}{\pi} \int_0^\infty \frac{\omega' \text{Im}\{\chi(\omega')\}}{\omega'^2 - \omega^2} d\omega'. \quad (1.24)$$

These relations lead to several important results. First, any material that exhibits absorption must also possess dispersion. Conversely, any dispersive medium must also possess some spectral variation in absorption, meaning that $d\alpha/d\omega$ cannot be zero for all ω . Thus, distortion is ever-present in slow light media. Additionally, the Kramers–Kronig relations dictate that $n(\omega)$ will be nearly linear in the neighborhood of a smooth peak or valley in the absorption spectrum. One important slow light technique is to create resonances that cause peaks in the absorption (or gain) spectrum. (For further discussion of the Kramers–Kronig relations, see section 1.7 of Ref. 15.)

1.2.2 Resonance features in materials

Many of the spectral features of a material's optical response come from material resonances. In many instances, the motion of bound charged particles in a material (such as electrons bound to atoms or molecules, or nuclei within a crystal lattice) is constrained to the form of a damped harmonic oscillator, similar to a mass on a spring. In this model, often called the Lorentz model, the charged particle tends to oscillate at a resonance frequency ω_0 . The equation of motion of the charged particle can then be written as

$$\frac{d^2x}{dt^2} + 2\gamma\frac{dx}{dt} + \omega_0^2x = \frac{eE}{m}, \quad (1.25)$$

where x is the particle's displacement from its equilibrium position, e is the charge carried by the particle ($e < 0$ for an electron), E is the magnitude of the applied electric field, m is the charged particle's mass, and γ is a damping coefficient. It can be shown that, under these conditions, the susceptibility χ of the medium due to the resonance has the form

$$\chi(\omega) \propto \frac{1}{\omega_0^2 - \omega^2 - 2i\omega\gamma}, \quad (1.26)$$

which gives the absorption spectrum $\alpha(\omega)$ a Lorentzian line shape centered at ω_0 with linewidth γ . The Kramers–Kronig relations dictate that such a line shape will cause $n(\omega)$ and $\alpha(\omega)$ to have the forms shown in Fig. 1.1. Of course, real materials have many different resonances, each with its own center frequency, linewidth, and relative strength. The total material response (susceptibility) is the sum of the responses due to the individual

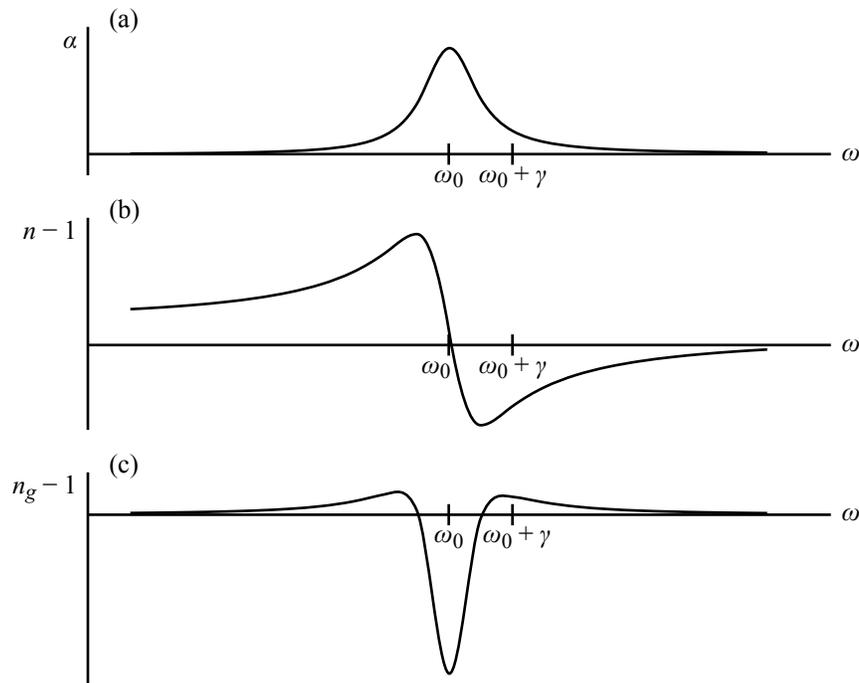


Figure 1.1. Dispersive features of an absorption resonance. (a) The absorption α of a resonance at a frequency ω_0 with linewidth γ . (b) The refractive index $n - 1$ of the same resonance. Near the point $\omega = \omega_0$, α reaches a peak and $n - 1$ crosses the axis. (c) Group index of the resonance. Note that fast light behavior ($|n_g| < 1$) is exhibited near $\omega = \omega_0$, while slow light behavior is exhibited around $\omega = \omega_0 \pm \gamma$. (From Ref. 1; used by permission.)

resonances. (For further discussion of material resonances, see sections 1.4 and 3.5 of Ref. 15 or section 5.5 of Ref. 16.)

Resonances of a similar Lorentzian form can also be induced by certain optical processes. Lasing, for example, consists of creating an inverted population, such that a certain atomic or molecular transition (a resonance) experiences gain. In that case, ω_0 is the frequency of the lasing transition, γ is its linewidth, and the value of $\alpha(\omega_0)$ is negative, indicating gain rather than absorption. The Kramers–Kronig relations then dictate a reversed slope for $n(\omega)$

near the center frequency, giving a $n_g(\omega)$ curve that is flipped vertically relative to the one shown in Fig. 1.1.

Figure 1.1 also gives clues about how to limit distortion in slow and fast light systems. Near $\omega = \omega_0$, $n(\omega)$ varies nearly linearly with ω and $n_g(\omega)$ is nearly flat, making group velocity dispersion nearly zero. However, $n_g(\omega)$ and $\alpha(\omega)$ change significantly at frequencies farther away from ω_0 . One must be careful that the pulse spectrum does not extend too far from the central frequency. In the case of fast light, the effect occurs in a region of strong absorption; the fast light experimenter may accept this absorption, mitigate it somehow, or resort to alternative fast light methods that avoid absorption, such as working at frequencies in between adjacent gain resonances [2].

1.2.3 Stimulated Brillouin scattering in optical fibers

In Stimulated Brillouin scattering (SBS), a strong pump field at frequency ω and a counterpropagating field at the Stokes frequency $\omega_S = \omega - \Omega$ are applied to a material. The electrostrictive effect, whereby materials experience a slight increase in density in response to an applied optical field, induces an acoustic wave (a pressure wave or traveling density modulation) at the beat frequency Ω . The density modulation creates a refractive index grating, and light from the pump field scatters off this grating and into the Stokes field, which thereby experiences gain. The gain resonance leads to steep dispersion in the vicinity of ω_S , resulting in slow light and fast light, similar to the resonance shown in Fig. 1.1. In optical fibers, conservation of momentum (also known as phase matching) requires that the

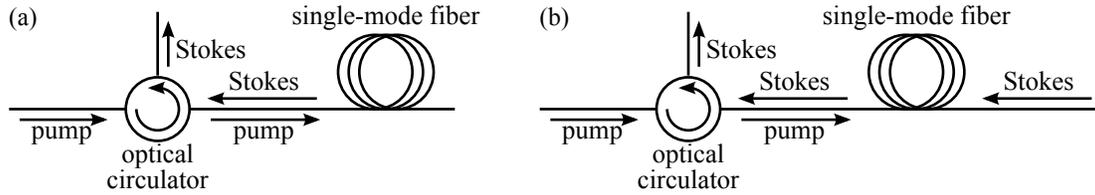


Figure 1.2. SBS generator and amplifier configurations in single-mode fiber. (a) SBS generator configuration; the Stokes field originates from noise. (b) SBS amplifier configuration; a Stokes seed field is injected opposite the pump field.

pump field propagate in the opposite direction to the Stokes field and the acoustic wave [17, 18]. (For more details on SBS, see chapters 8 and 9 of Ref. 15.)

There are two primary configurations for achieving SBS. When both the pump field at ω and the Stokes field at ω_S are applied to the material, the Stokes field acts as a seed for the SBS process, and the configuration is known as SBS amplification. If no Stokes field is applied, a pump photon scatters off a thermal phonon at Ω , creating a Stokes photon to seed the SBS process; this configuration is known as SBS generation. These two configurations are shown in Fig. 1.2.

In the approximation that the SBS pump maintains constant intensity (known as the undepleted-pump approximation), the Stokes field in the fiber propagates according to

$$I(L) = I_0 \exp(gI_P L). \quad (1.27)$$

Here, $I(L)$ is the intensity of the Stokes field after propagating a length L in the fiber, g is the SBS gain, and I_P is the pump intensity. The SBS gain g is given by

$$g = g_0 \frac{(\Gamma_B/2)^2}{(\Omega_B - \Omega)^2 + (\Gamma_B/2)^2}, \quad (1.28)$$

where Ω_B and Γ_B are the Brillouin shift and linewidth of the medium, respectively. The gain at line-center ($\Omega = \Omega_B$) is

$$g_0 = \frac{\gamma_e^2 \omega^2}{nvc^3 \rho_0 \Gamma_B}, \quad (1.29)$$

where γ_e is the electrostrictive coefficient, n is the refractive index, v is the speed of the acoustic wave, and ρ_0 is the mean density of the medium. (The undepleted-pump approximation is useful for intuition. It applies when only a little of the pump intensity is converted to the Stokes frequency—usually when the Stokes field is weak and g_0L is not particularly large. For more complicated situations, coupled wave equations must be solved, although the expressions for g and g_0 remain the same; see Ref. 15.)

The SBS group index is [18]

$$n_g = n_{g,0} + \frac{cg_0I_P}{\Gamma_B} \frac{1 - [2(\Omega_B - \Omega)/\Gamma_B]^2}{\{1 + [2(\Omega_B - \Omega)/\Gamma_B]^2\}^2}. \quad (1.30)$$

Here, $n_{g,0}$ is the group index of the medium in the absence of the SBS effect. Increasing I_P , the SBS pump intensity, increases the group index linearly, making SBS a highly controllable form of slow light. In standard telecommunications-type single-mode optical fibers with

the pump near 1550 nm, $\Omega_B/2\pi$ is about 10.8 to 11 GHz and $\Gamma_B/2\pi$ is about 35 to 70 MHz [18, 19]. However, a modulated, spectrally broadened pump produces a broader Stokes gain line, which allows a broader-bandwidth resonance for the slow light effect, as well as control over the spectral profile of the group velocity [20]. Broadening and shaping the pump spectrum is an important method both for enhancing the slow light bandwidth [21, 22] and for controlling pulse distortion [2, 23]. Slow light based on SBS in optical fibers has achieved 13.4 ns delays of 5.5 ns pulses and 10.9 ps delays of 37 ps pulses [24, 25].

1.2.4 Dispersive delay

Slow light is often induced through some effect that creates extreme dispersion, such as SBS, but it need not be an induced effect. Systems with inherently high dispersion also exhibit slow light. A practical example is highly dispersive fiber, such as dispersion compensating fiber (DCF), in which different wavelengths of light travel at significantly different speeds. DCF is characterized by its dispersion parameter D , expressed in units of $\text{ps km}^{-1} \text{nm}^{-1}$, or picoseconds of additional group delay per kilometer of fiber length and per nanometer of wavelength change. The change in group delay (propagation time) of a pulse through the fiber is then $\Delta\tau_g = DL\Delta\lambda$, where L is the length of the fiber in km, and $\Delta\lambda$ is the wavelength change in nm. For instance, in a 1 km length of fiber with $D = -100 \text{ ps km}^{-1} \text{nm}^{-1}$, increasing the optical frequency by 3 nm will decrease the group delay by 300 ps. Thus, tuning the optical wavelength controls the group delay. This is one of a class of techniques

referred to as *dispersive delay*. It is sometimes called the fiber prism effect, in reference to the high dispersion of prisms [26, 27].

A related technique called *conversion and dispersion* adds wavelength conversion before and after the highly dispersive fiber, which preserves the signal wavelength. The first wavelength conversion selects the wavelength that achieves the desired relative group delay, and the second wavelength conversion returns the signal to the original operating wavelength. Although it is more complex and can have higher loss than simple dispersive delay, the conversion and dispersion technique has achieved delays as long as 1200 pulse widths with 3.5 ps pulses and can operate at data rates well over 40 Gbit/s [28–31].

There has been some debate over whether these methods may properly be termed *slow light*. The group delay through such systems is dominated by simple propagation through the fiber, and long group delays are achieved in part by using long lengths of fiber. Only the differential group delay between wavelengths is controllable, and this differential delay is caused by the differences in refractive index at widely different frequencies (separated by several THz), not by an extreme group index at any particular frequency. While the techniques were discovered (or rediscovered) as part of efforts to discover new slow light methods, and they provide tunable all-optical delay just as slow light does, they may be less controversially referred to as all-optical delay methods.

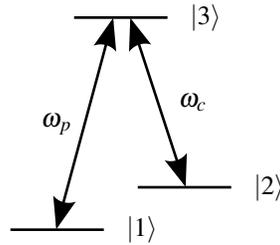


Figure 1.3. Three-level energy level model for EIT. A weak incoming probe beam at frequency ω_p is nearly resonant with the transition between energy levels $|1\rangle$ and $|3\rangle$, and a strong pump beam at frequency ω_c is nearly resonant with the transition between energy levels $|2\rangle$ and $|3\rangle$.

1.2.5 Other material slow light phenomena

Researchers have demonstrated material slow light in a number of other systems. Perhaps the most dramatic of these is electromagnetically induced transparency (EIT), and its counterpart, electromagnetically induced absorption (EIA). EIT has produced some of the most remarkable group velocities, including the 1999 experiment by the Hau group that achieved a group velocity of 17 m s^{-1} . In EIT, a strong pump beam at frequency ω_c creates a quantum mechanical coherence between the two states of an atomic or molecular transition. A probe beam at frequency ω_p , coherent with the pump beam, is tuned to a different transition involving the same excited state as the pump transition; transitions between the two ground states are forbidden, and the transitions form a so-called lambda system, shown in Fig. 1.3. The probe beam experiences a very narrow window of reduced absorption, also called a spectral hole, within the transition [32]. The narrowness and depth of the spectral hole lead to a large group index, producing slow light. EIT typically requires special media and special environments to reduce decoherence, such as cryogenic temperatures or ultrahigh vacuum conditions. Despite these restrictions, it has been a popular experimental method for

achieving slow light in a variety of systems, including Bose–Einstein condensates (BECs) [9], alkali vapors [33–35], crystals [36], semiconductor quantum wells [37–39] and quantum dots [40, 41], and vapor confined within a photonic bandgap fiber [42]. Certain transparency effects similar to EIT have been demonstrated in resonator systems [43–48] and in plasmas [49–53].

Coherent population oscillation (CPO) occurs when a pump beam at frequency ω and a probe beam at frequency $\omega + \delta$ are applied to the same atomic (or molecular) transition. If ω and $\omega + \delta$ both lie within the natural linewidth $1/T_1$ of the transition, a portion of the atomic population oscillates at the beat frequency δ between the two energy levels of the transition. The oscillating population produces a narrow hole in the absorption line centered at frequency ω . By the Kramers–Kronig relations, the narrow spectral hole results in a rapid index variation, producing slow light. (Of course, if the atomic population is initially in the excited state, the CPO effect produces a hole in the gain spectrum, giving a fast light effect [54, 55].) Equivalently, CPO can be viewed as a time-dependent saturable absorption or saturable gain effect; the optimal pulse bandwidth is of the order of δ , so approximately one complete population cycle occurs during the interaction [56–59]. CPO is much easier to achieve experimentally at room temperature than EIT, and CPO has similarly narrow linewidths to EIT, resulting in similarly extreme group velocities. However, CPO typically suffers from a higher degree of residual absorption than EIT. CPO has been achieved in a variety of experimental setups, including in crystals [55, 60], erbium-doped optical fiber [61], semiconductor waveguides [62], and quantum wells and quantum dots [63–65].

A variety of other slow light techniques have been implemented successfully. Picosecond pulses were delayed by as many as 80 pulse widths by operating at the center frequency between two absorption lines (hyperfine ground states) of cesium; the absorption lines provide some residual dispersion over the broad band between the two lines, with a fairly flat group index over a wide bandwidth and little absorption [23, 66]. Slow light in semiconductors has also been achieved using a number of different mechanisms, including the gain of semiconductor optical amplifiers [67, 68] and several excitonic mechanisms [69–71]. Stimulated Raman scattering (SRS) has also been used; the process bears a resemblance to SBS, except the ground states are atomic or molecular vibrational sublevels separated by a vibrational frequency Ω (see chapter 10 of Ref. 15). Slow light based on SRS gain has been observed both in solids [72, 73] and in optical fibers [74].

1.3 Structural slow light

1.3.1 Photonic crystals and photonic bandgap structures

Photonic bandgap devices are formed by introducing periodic changes in the refractive index of a dielectric medium. Often, the periodic index modulation is due to a regular geometric structure in the dielectric, such as regularly spaced ridges or air holes. Because of the periodic index modulation, light within certain wavelength bands is unable to propagate within the device. Specifically, these forbidden wavelength bands lie near $\lambda/n = ma/2$, where n is the effective refractive index, a is the period of the index modulation, and m

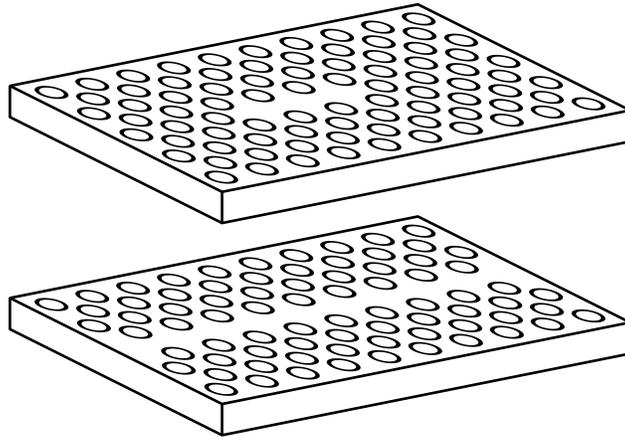


Figure 1.4. Photonic crystal structures. Top: An L3 cavity, formed by removing three consecutive holes. Bottom: A W1 waveguide, formed by removing an entire row of holes.

is any positive integer. The forbidden bands are called *photonic bandgaps*, by analogy to the bandgap of a semiconductor crystal lattice. Wavelengths just outside of the bandgap experience steep dispersion, which results in slow light [75–81].

One of the most common kinds of photonic bandgap devices is the *two-dimensional photonic crystal*, often just called a photonic crystal. It is a dielectric slab with holes etched into it in a hexagonal or sometimes rectangular lattice. Photonic crystals are particularly versatile. They may be used to design many different optical devices and may be fabricated out of virtually any dielectric media. They may also be created from highly nonlinear optical media, for example by using silicon at telecommunications wavelengths [82, 83], or by filling the holes of the photonic crystal lattice with a highly nonlinear fluid [84]. Three-dimensional photonic crystals have been fabricated, though they are challenging to make [85–87]. One-dimensional photonic bandgap structures, such as fiber Bragg gratings, are sometimes called one-dimensional photonic crystals.

The flexibility of photonic crystals comes from the ability to introduce defects in the periodic array, by removing holes, shifting them, or changing their size. For instance, a waveguide may be created by removing a row of holes, or a cavity may be created by removing three, five, or seven consecutive holes. These are respectively called a W1 waveguide and an L3, L5, or L7 cavity [80, 88], and some examples are shown in Fig. 1.4. (Of course, the holes are not “removed”; they are simply not created.) These defects in the photonic crystal lattice allow light to propagate or resonate in specific ways in the photonic crystal device. Light is confined to the defects because it cannot propagate within the photonic crystal lattice. The refractive index contrast between the dielectric slab and the air above and below it confines the light vertically, through total internal reflection, just like a slab waveguide. The exact size, shape, and location of the holes near the defects allow further control of light within the device, especially to reduce the pulse distortion caused by the slow light system—leading to an area that some have termed *dispersion engineering* [89–92].

Photonic crystals are made by nanofabrication, a term that includes several different approaches. One common approach is electron-beam lithography. The process typically begins with a silicon-on-insulator (SOI) wafer, a silicon wafer on which a layer of silicon dioxide is grown and another layer of silicon, around 200 nm thick, is deposited. The SOI wafer is coated in a layer of *resist*, a protective chemical that changes when exposed to an electron beam. An electron beam is then used to write a particular pattern into the resist, and the resist is chemically developed to remove the areas of the resist that were exposed to

the electron beam (if a positive resist is used; with negative resist, the unexposed areas are removed during development). The wafer with the patterned resist is then exposed to an etching chemical, such as hydrofluoric acid. The exposed areas are etched away, while the areas protected by the resist are not etched. In this way, a pattern of holes or other shapes and structures can be written into the wafer. Often, an additional undercutting step will remove the silicon dioxide from underneath the top silicon layer, creating a membrane-like structure; light is vertically confined more tightly in a membrane because the refractive index contrast between silicon and air is greater than between silicon and silicon dioxide. Electron-beam lithography can write extremely small features with sizes on the order of 1 nm. However, photonic crystals are highly sensitive to errors in the placement of the holes and to roughness in the sidewalls of the holes, and the resulting disorder causes both increased scattering and loss [93–95]. The group index in photonic crystal waveguides and similar devices is generally below 1000, which is more modest than atomic-resonance-based systems [89, 96, 97]. When the device is designed for constant group velocity across a wide bandwidth, the group index is often of the order of 30 to 50 [91]. But the increased design flexibility of photonic crystals as well as the transparency of silicon near 1550 nm and the good vertical confinement of SOI membrane structures make this an attractive platform for designing practical devices.

1.3.2 Other structural slow light systems

Slow light effects have also been explored in coupled-resonator structures, often called coupled-resonator optical waveguides (CROWs) or coupled-cavity waveguides (CCWs). CROWs consist of a series of optical resonators placed near each other. Low group velocities are observed in the propagation of light across the CROW, as a result of weak coupling and feedback between the resonators [98]. Here is a conceptual model of how the device works: Light couples evanescently into the first resonator. As the light resonates there, it couples evanescently into the second resonator, where it also resonates. It then couples evanescently into the third resonator, and so forth, until it has “leaked” across the entire waveguide [99]. Any kind of resonator may be used, including Fabry–Pérot cavities, ring or disk microresonators, and photonic crystal defect resonators.

Slow light has also been explored in certain optical filters, including fiber Bragg gratings [100, 101] and Moiré fiber gratings [102]. The dispersion and slow light effects in optical filters are similar to those of coupled-resonator structures.

1.4 Additional considerations

The phenomena of slow and fast light include a number of fascinating features. Several of these deserve mention here, even though they are not part of the remainder of the dissertation.

1.4.1 Figures of merit

Several figures of merit are in common use among workers in slow light. The group index n_g may be inferred from experimental data and reported. When one does not wish to ascribe a homogeneous refractive index to the material (especially with structural slow light systems), alternative quantities may be defined, such as the slowdown factor $S = v/v_g$ [14] or the slowing factor $S = c/v_g$ [103]. Perhaps the most common figure of merit is the group delay τ_g , the time delay induced by propagation through the slow light medium. The group delay is nearly always the experimental quantity that is measured directly, so it is simple to report. However, it is generally easier to produce longer delays for longer pulses. Thus, a more meaningful measure is the fractional delay, or the delay normalized by the pulse width [104]. Fractional delay coincides more closely with the particular application of slow light delay lines, and it is a measure of the number of bits that can be stored by a delay line. Fast light systems may be evaluated in terms of fractional advancement, or negative fractional delay. Fractional delay or advancement is often quoted along with pulse width. Perhaps the most useful single figure of merit for optical delay lines is the delay–bandwidth product (DBP), which is also equal to the maximum possible fractional delay in a given slow light system [105]. The delay–bandwidth product must also be quoted with the bit rate to be a definitive performance measure. The maximum possible delay can be represented in other ways, such as the length of a waveguide required to achieve a given time delay [79] or the ratio of a slow-light quantum memory’s maximum storage time to the input pulse length [106].

Other figures of merit often include some measure of the absorption experienced by the pulse, such as the ratio of the group index or the delay–bandwidth product to the absorption coefficient [107, 108], or the time a signal can propagate in a slow light buffer before needing regeneration or amplification [71]. Pulse distortion can be measured in several different ways, including the input–output pulse width ratio [109], degree of dispersion near an absorption feature or a band edge [79], or group velocity dispersion (GVD). The effects of pulse distortion on a telecommunications system are often the ultimate concern, so some experimenters use commercial telecommunications test equipment to test the bit-error rate (BER) [110] or the eye opening [111]. Many more figures of merit have been defined, and the best choice of figure of merit is application-specific.

1.4.2 Theoretical limits of slow and fast light

The most general theoretical limits of the performance of slow light systems were already mentioned in section 1.1.2: group velocity dispersion, frequency-dependent absorption, and higher-order dispersion and absorption terms must be sufficiently small that the pulse is not distorted too much (though the degree of acceptable distortion is often application dependent). Generally, total linear absorption (αL) must also be sufficiently small that the signal can be detected. More specific limits than these depend on the particular slow light technique in question. Some results are quoted here without further comment.

For many slow light techniques, group velocity and bandwidth are proportional, requiring a tradeoff between the two parameters [108, 112]. Under many circumstances, the minimum

spatial extent occupied by a single optical bit in a slow light medium is roughly one vacuum wavelength [108]. In slow light systems using stimulated Brillouin scattering (SBS) in optical fibers, there is a tradeoff between increased bandwidth and reduced pulse distortion [111].

1.4.3 Spatial compression and nonlinearity enhancement

It is clear that the reduced group velocity in slow light leads to spatial compression of the pulse by a factor equal to the group index. If a pulse of duration τ decelerates from a speed c to a speed c/n_g , its length L must likewise decrease by a factor of n_g , from $L = c\tau$ to $L' = c\tau/n_g$. Conservation of energy then dictates that if the pulse energy was distributed over length L but is compressed down to length L' , the energy density u must increase by the same factor to $u' = n_g u$.

Interestingly, in material slow light systems, the intensity $I = uv_g$ of a pulse is unchanged upon entering the medium, because the increase in u is canceled exactly by the decrease in v_g . Likewise, the electric field strength $E \propto \sqrt{I}$ is unaffected by changes in the group velocity. Thus, although the pulse energy is spatially compressed, its peak electric field strength is unchanged [14, 113, 114]. In contrast to this result, slow light in structural slow light systems such as photonic crystals is accompanied by an increase in electric field strength, and therefore an increase in intensity [84]. A helpful mental picture is that the field inside a structural slow light system undergoes multiple reflections, like in a Fabry–Pérot

cavity, and this leads both to slowed propagation (due to increased effective path length) and higher field strength (due to the presence of multiply reflected fields in the medium) [14].

Beer–Lambert–Bouguer absorption is enhanced in structural slow light but not in material slow light, since Beer–Lambert–Bouguer absorption relates to the intensity of the light, which is only enhanced in structural slow light [14, 115, 116]. Further, optical nonlinearities, which typically scale as some exponent of the field strength, are enhanced in structural slow light only, and not in material slow light [14, 82, 84, 99, 117]. The exact nonlinear enhancement is somewhat complicated, but when the slowdown factor $S = v/v_g$ is large and the medium has a Kerr-type nonlinearity, the enhancement is proportional to S^2 [83, 117, 118].

EIT is also associated with an enhancement of the medium's optical nonlinearity [114]. When the fields applied to a resonant medium are tuned to a resonance, the optical nonlinearity of the medium reaches a local maximum. In the absence of EIT, linear absorption also reaches a local maximum, making the nonlinearity unusable. EIT allows access to these resonant nonlinearities that would otherwise be precluded by absorption. It would be incorrect to say that the nonlinearity enhancement in EIT is caused by the slow light effect; however, the slow light effect and the nonlinearity enhancement in EIT are inseparable.

1.4.4 Causality of fast light

Fast light ($v_g > c$) and backwards light ($v_g < 0$) seem at first to violate causality. However, careful analysis shows that this is not so. Causality is the requirement that any effect must

be preceded by its cause. When combined with the special theory of relativity, causality requires that no information travel faster than the speed of light. (Otherwise, it would be possible to violate causality in certain frames of reference.) What does this mean for a group velocity greater than the speed of light? Extensive discussion of this question has occurred in the literature over the last century; see, for instance, section 5.2 of Ref. 113, section 2.5 of Ref. 12, Ref. 119, and their references. A brief overview is presented here.

From 1907 to 1914, Sommerfeld and Brillouin examined the propagation of a discontinuous jump (like a step function) in the electric field. They examined the front velocity, or the speed of propagation of the first non-zero value of the electric field. They found that the front velocity can never exceed c and that no part of the waveform can overtake the front [10, 11]. Their result was later extended to nonlinear media and to all functions with compact support, meaning functions that are zero except over a finite range [120].

Many fast light experiments and theories use Gaussian-like pulses with long leading and trailing tails. The group velocity can then be used to describe the motion of the pulse envelope or the pulse peak. Using the presence or absence of a pulse to represent one bit of information (as in on-off keying), one may be tempted to think of the peak as carrying the information associated with the pulse, and hence conclude that information is propagating superluminally (faster than c). However, the presence or absence of the long leading edge of the pulse carries the same information as the presence or absence of the peak. A true Gaussian pulse has infinite extent; in a sense, the pulse and its information have already

arrived everywhere, irrespective of the motion of the peak. The superluminal peak velocity is therefore not indicative of superluminal information transfer.

For the more realistic case of a truncated Gaussian pulse, the peak of the pulse may travel superluminally for a time, but the front of the pulse still propagates at or below the speed of light (since the pulse front is a discontinuity). None of the pulse energy can overtake the pulse front. For example, in an on-off keyed binary signal, fast light may shift the peak of a pulse within its bit slot but cannot advance the peak past the beginning of the bit slot. As the peak approaches the front, the pulse becomes highly distorted, often breaking up into a series of peaks or some other irregular shape. In short, attempts to violate causality lead to pulse breakup. Ultimately, this should not come as much of a surprise; fast light comes about entirely from Maxwell's equations combined with the assumption of a causal optical response from the medium (often in the form of the Kramers–Kronig relations), and causality is a natural consequence of those conditions.

1.5 Applications of slow and fast light

Slow and fast light allow researchers to conduct many exciting fundamental studies of physics and light propagation, but they also have several potential practical applications. These applications are the primary focus of this dissertation. Two broad categories of applications exist. Perhaps the most obvious use for a slow light medium is as a way to delay an optical signal. Many slow light methods are tunable, meaning the group delay can

be controlled, whether optically or by some other means. Tunable optical delay can serve as a true-time delay element for phased-array laser radar, and this application is discussed in section 1.5.1. Tunable delay could also be useful in telecommunication networks, optical coherence tomography (OCT), ultrafast pulse metrology, and various kinds of optical signal processing [28, 29].

Slow light can also be used to enhance interferometers and spectrometers. The spectral sensitivity of interferometers and spectrometers depends on the group index, and a large group index makes spectrometers more sensitive. While this does not involve slowing down a pulse of light, the same set of dispersive techniques can be used to increase the resolution of spectrometers. Since many kinds of spectrometers have a direct relationship between their size and resolution (larger spectrometers are generally more precise), this increased sensitivity can be used to shrink the size of a spectrometer while maintaining spectral resolution. A design for an on-chip, slow-light-enhanced nanophotonic spectrometer is described in section 1.5.2, along with an approach to optical computation that would improve the performance of such a device. Additional possible applications are presented in section 1.5.3.

1.5.1 A slow-light laser radar (SLIDAR)

Laser radar, also called lidar or ladar, is a form of sensing in which a laser illuminates a target and the reflected light is analyzed. Perhaps the most common use of laser radar is to measure the distance to a target by sending a single optical pulse to the target and measuring

the round-trip time for the pulse to return [121]. Two key parameters for this use are the longitudinal (or range) resolution, which is determined by the duration of the optical pulse, and the transverse resolution, which is governed by the size of the emitting aperture. Shorter pulses and larger apertures lead to better (finer) resolution. However, large-aperture optics are bulky, expensive to fabricate, and unwieldy to steer.

An alternative approach is to use an array of small apertures, working in concert. Each small aperture, or emitter, produces a pulse at the same time, and the pulses overlap in the far field (i.e., at the target). If the optical phases of the emitters are kept in the appropriate relationship, then the pulses interfere at the target, producing a much narrower far-field spot than that of the individual emitters. Ideally, the baseline of the array (or its total extent) determines the transverse resolution of the spot in the far field. This is known as a phased array [122].

In general, the far-field spot of a phased array can be moved by changing the phases of the individual emitters, rather than steering a single large-aperture optic. (In certain configurations, the individual emitters may need to be steered in-place as well.) However, if a phased-array system with short pulses is steered far off-axis as in Fig. 1.5(a), the emitters see different path lengths to the target, and the pulses from different emitters may not reach the target at the same time, degrading both the transverse and longitudinal resolution of the system. In order for the pulses to overlap, each emitter must include a variable group delay element, called a *true-time delay* [123, 124]. Slow light provides an excellent true-time delay for a phased-array laser radar.

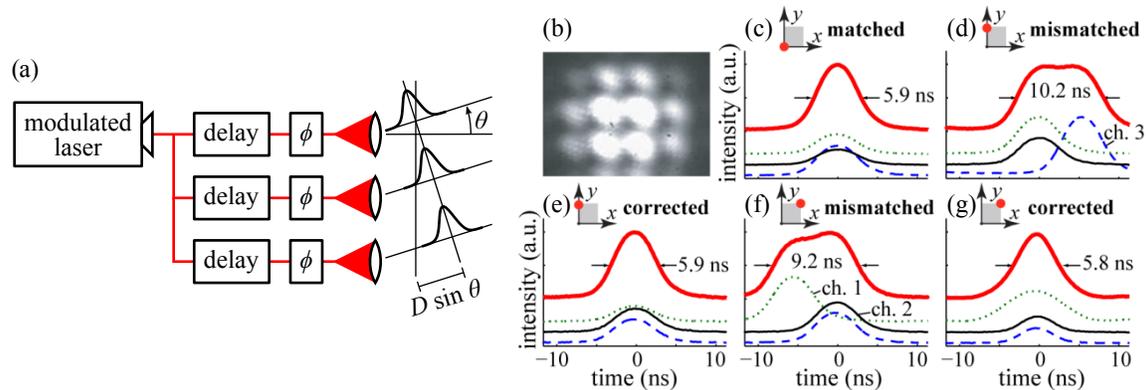


Figure 1.5. SLIDAR: a pulsed, phased-array slow-light laser radar. (a) Block diagram. ϕ : phase control; D : emitter baseline length; θ : steering angle. The diagram shows the beam steered off-axis; for the pulses to overlap at the target, each emitter must include a variable group delay. (b) SLIDAR far-field intensity pattern, formed by three emitters arranged in a right triangle. (c)–(g) Normalized time traces of the returned signal when the beam is steered in three directions (indicated by the red dot at the top of each sub-ispell-figure), with or without group delay compensation. The top thick red trace is the combined signal, and the other three traces correspond to signals emitted individually from each of the three channels. The traces are shifted vertically for clarity. (Adapted from Ref. 3; used by permission.)

In 2011 and 2012, Schweinsberg and coworkers (including myself) demonstrated a prototype slow-light laser radar system, called SLIDAR [3–6]. Figure 1.5(a) shows the system diagram. The system has three emitters arranged in the shape of a right triangle, allowing the beam to be steered in both the horizontal and vertical directions; the resulting far-field intensity pattern is shown in Fig. 1.5(b). To allow delay compensation in both transverse dimensions (horizontally and vertically), two types of slow light are used, namely SBS and dispersive delay; the SBS pump power and the optical wavelength determine the amount of delay. The system is fiber-based, operating at standard telecommunications wavelengths near 1550 nm. It employs both dispersion-shifted fiber (low dispersion) for

SBS and dispersion-compensating fiber (high dispersion) for dispersive delay. Pulses of about 6 ns duration allow a far-field resolution of about 1.8 m.

The two slow light methods used in SLIDAR provide a true-time delay that compensates the path length differences due to beam steering, as shown in Fig. 1.5, sub-figures (c) through (f). When the beam points on-axis, the pulses overlap in the far field, as shown in sub-figure (c). When the beam is pointed upward, the pulses no longer overlap, shown in (d). One slow light method is then used to compensate the group delay mismatch, shown in (e). Next, the beam is steered in the horizontal direction, leading to another group delay mismatch, shown in (f). Finally, the second slow light method is used to compensate the additional group delay mismatch, shown in (g). The slow light true-time delay elements preserve the pulse duration and longitudinal resolution even while the beam is steered far off-axis in both transverse dimensions.

To maintain the transverse resolution, the emitters must be properly phased, in addition to being synchronized in time. The SLIDAR phase control system maintains phase lock among all three emitting channels and a non-emitting phase reference channel. Each channel contains 2.2 km of optical fiber, and gain due to the SBS slow light process results in a varying signal level in some of the channels. The phase control system compensates for these issues while maintaining phase lock, with a root-mean-square (RMS) phase error of $\pi/5$ radians, or 1/10 wave. The SLIDAR phase control system is covered in chapter 2.

1.5.2 A slow-light-enhanced nanophotonic spectrometer

Much research is currently aimed toward the idea of a *lab on a chip*, in which some set of laboratory functions is performed by a tiny, integrated device. Since spectrometers are particularly useful for detecting and analyzing chemicals, an on-chip spectrometer is highly desirable for lab-on-a-chip applications. On-chip photonic devices, often called integrated photonics or nanophotonics, are already a reality. Several designs exist for on-chip spectrometers, but they suffer from an undesirable tradeoff between device size and spectral resolution [125–131]. However, in many types of interferometers and spectrometers, the spectral sensitivity (spectral resolution) of the device depends on the group index, rather than the traditional refractive index [7, 132–135]. Slow light methods can enhance the spectral resolution of these spectrometers. Alternately, since spectrometer resolution usually improves with increasing device size, slow light can be used to shrink the device while maintaining high resolution.

An arrayed waveguide grating (AWG) is a nanophotonic device that disperses input light into a series of output waveguides, each of which carries a separate wavelength band (see Fig. 1.6). An input waveguide connects to a free propagation region, where light enters an array of waveguides. Each waveguide in the array is longer than the next by an increment Δl , and light propagating through the waveguide acquires a phase that is proportional to Δl and inversely proportional to the wavelength. The outputs of the waveguide array are spaced equally at the entrance to another free propagation region. The waveguide array spacing is of the order of a wavelength of light, so light coming out of the waveguide array acts

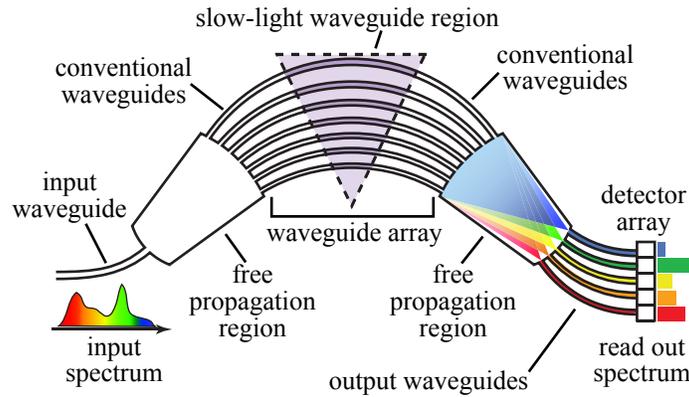


Figure 1.6. A design for a slow-light-enhanced nanophotonic spectrometer. The design is based on an arrayed waveguide grating (AWG), with a slow-light waveguide region that enhances the spectral resolution. (From Ref. 7; used by permission.)

in some sense like light diffracting off a grating. Each successive waveguide output has a linearly increasing, wavelength-dependent phase, meaning the waveguide outputs act like a blazed diffraction grating, but with a different blaze angle at each wavelength. The array outputs are also arranged in a concave pattern, which acts to focus the light. The net effect is that light at each wavelength is focused onto a different output waveguide at the exit of the second free propagation region. AWGs originated as wavelength-separation devices for wavelength-division multiplexing (WDM) optical networking [136].

Figure 1.6 shows a conceptual design for an AWG-based spectrometer whose spectral resolution has been enhanced by slow-light waveguides. The device design is the same as a standard AWG, but the waveguides in the array include a slow-light waveguide section [7, 132, 135]. The slow-light waveguides could be made from photonic crystal waveguides. It is important that the spectral profile of the group index be optimized (and nearly flat) over the operating bandwidth of the spectrometer. One proposed design for a slow-light-enhanced

AWG spectrometer has a resolution of 0.04 nm over a bandwidth of 12.5 nm near a central wavelength of 1550 nm [137].

At the outputs of the slow-light-enhanced AWG, one could place a detector array to record the power spectrum of the input light. However, the goal is often not to record a spectrum but to detect the presence or absence of some particular chemical, such as to detect the presence of chemical agents or toxins. In this case, the spectrum is only a means to an end, and recording the entire spectrum is inefficient. One can achieve much higher signal-to-noise ratios by performing optical computation, a technique discussed in chapter 3.

1.5.3 Other applications of slow and fast light

Controllable delays for communications and networking

Much research is directed toward the possibility of all-optical networking, in which some or all processing is performed optically rather than electronically so that one avoids the so-called optical-electronic-optical (OEO) conversion at routing points between network links [138]. But while the controllable optical delay of slow light seems perfect for an optical buffer in an all-optical network router, in practice nearly all slow light systems have far too little delay—about two pulse widths or less, rather than the thousands or millions of pulse widths of delay needed in a network router. The fractional delay is often limited by distortion or loss. In fact, loss may be a fundamental limitation to all forms of optical buffering [108].

However, slow light delay lines are ideal for all-optical jitter compensation [139, 140]. In modern data networks, transmissions are synchronized to bit slots, or regularly spaced time windows during which either a zero or one value is transmitted. For example, in on-off keyed transmission, logical zero or one is indicated respectively by the absence or presence of a light pulse in each bit slot. In a 10 Gbit/s transmission link, bit slots are 100 ps each. Every 100 ps, the transmitter transmits a pulse of light to indicate a one or no pulse to indicate a zero.

Transmitters and receivers must agree on the duration and start times of bit slots. However, during the transmission process, the optical signal may become slightly stretched, compressed, or delayed. These changes can be caused by temperature-dependent changes in the refractive index or polarization mode dispersion, bit-pattern-dependent nonlinearities in the optical fiber, or environmental effects such as vibration. The consequence is that the data bits become randomly misaligned relative to their bit slots. This effect is known as jitter, and it can cause data corruption and increase the bit-error rate (BER). However, if the receiver can re-synchronize to accommodate these random changes, network throughput can be restored. Tunable slow and fast light can be used for jitter compensation.

Optical memory, optical processing, and coherent control

The group velocity of light can in fact be adjusted to zero, leading to so-called stopped light [34–36]. One stopped-light technique is to use electromagnetically induced transparency (EIT) to “map” a light pulse onto the spin coherence of a medium, effectively storing the

pulse. The pulse can later be retrieved by performing the reverse operation. Other proposed techniques include using solitons in coupled-resonator structures [99, 141]. A stopped-light system could be useful as an optical memory for storing pulses of light.

Many stopped-light techniques also preserve quantum coherence properties. EIT-based techniques, for example, can be used to preserve and store entanglement, or the coherence of two quantum-mechanical systems (such as two photons or a photon and an atom). A so-called quantum memory, one which can store and retrieve entangled photons, could find applications in quantum computing, quantum key distribution, and other technologies that depend on entanglement [106].

One drawback of stopped-light memories is the finite lifetime of the memory. EIT media gradually undergo decoherence (or dephasing). Other kinds of optical memories have different decay mechanisms, but all memories decay and lose their data over time. But decay only needs to be slowed, not eliminated. A memory cell can be refreshed by reading out its value and rewriting the value into the cell, starting the decay cycle anew. In two 2001 stopped-light experiments using EIT, the coherence lifetime was $500 \mu\text{s}$ [35, 36]. Thus, a stopped-light memory cell using similar techniques would need to be refreshed faster than every $500 \mu\text{s}$. For comparison, modern electronic memory (DRAM) cells require a refresh every $7.8 \mu\text{s}$ [142].

Slow light media can buffer not just pulses of light but in fact entire images [143]. Both amplitude and phase information in an image are preserved by slow light media. Such an image buffer could have applications in optical image processing. Further, slow light

can be used to enhance image rotation in a rotating medium [144, 145]. Slow-light optical delay lines have also been proposed to implement certain kinds of optical signal processing elements, particularly equalization filters and other filters that can be implemented as tapped delay lines [139].

Spectroscopic enhancement of interferometers and cavities

The spectral properties of slow and fast light can be used to enhance the spectral performance and robustness of spectroscopic interferometers and spectrometers. In such interferometers, a change in laser frequency causes a change in the optical path length difference (OPD) between the two interferometer paths. A slow light medium placed in one path increases the sensitivity by a factor of the group index [133]. For example, if a slow light medium of length L slows a pulse of light to a group velocity of $c/10^6$, it has the same effect on frequency sensitivity as would an OPD of 10^6L . Of course, by using fast light, one could correspondingly decrease the frequency sensitivity of the interferometer. In theory, using stopped light with zero group velocity and zero group index, one could remove all sensitivity of the phase to frequency shifts, at least over the bandwidth of zero group velocity. Slow- and fast-light enhancements have also been demonstrated in Fourier-transform interferometry [134] and proposed for interferometric rotation sensing [146, 147]. Similar sensitivity increases are seen when a slow- or fast light medium is placed inside a Fabry–Pérot cavity: the cavity linewidth is changed but the cavity storage time remains unaffected [148–150].

Fast light may also be used to construct a white light cavity [151]. A white light cavity is a Fabry–Pérot cavity that contains a fast light (anomalously dispersive) element. In a normal Fabry–Pérot cavity, a slight detuning away from resonance will reduce the cavity transmission drastically; each round trip acquires a slight phase shift, and the multiple round trips add destructively. However, in a white light cavity, the fast light element compensates for this slight phase shift, such that the cavity resonates across a range of wavelengths [147].

1.6 Conclusion

Slow light is an exciting and rapidly maturing field. Control of the group velocity produces startling and sometimes counter-intuitive results, and the study of slow light has enhanced researchers' understanding of some important nuances of the propagation of light. The science of slow light is well established, and while slow light has its limits, some exciting applications of slow light are being pursued.

Chapter 2

Phase Control of a Slow-Light Laser Radar

Laser radar, also called lidar or ladar, has a number of practical applications. In its simplest form, a laser radar consists of sending a single pulse of light from a laser toward a target, detecting when the reflected pulse returns, and using the time of flight to determine the distance (or range) to the target. This is sometimes called range-finding, and it is commonly used for surveying on land and underwater. A number of other laser radar techniques exist, and there are numerous applications beyond range-finding, such as measuring the speed of a moving target, imaging, and even using laser radar with spectroscopy to detect chemicals remotely, such as pollutants in the atmosphere [121, 122, 152].

The precision of a laser radar for range-finding has two components. The longitudinal resolution, or the precision in measuring the range to the target, depends on the pulse duration, and shorter pulses mean better longitudinal resolution. The transverse resolution, or the size of the laser spot in the far field (i.e., at the target), determines how precisely a single target can be picked out of the far-field scene. Transverse resolution is limited by

the size of the aperture that emits the laser pulse. Due to the laws of diffraction, a larger aperture produces a smaller far-field spot. For the best transverse resolution, it is desirable to use the largest possible emitting aperture.

Large optics, however, are expensive and unwieldy. An alternative approach is known as a *phased array*. An array of small emitters all emit a laser pulse at the same time. The pulses are coherent, meaning they are in phase with each other and can interfere (add constructively or destructively). When the pulses reach the target, they produce an interference pattern that has a brighter, smaller spot than any of the small emitters could have produced on its own. The far-field spot can be just as small as the spot from a single large aperture that is the same size as the array, but the small emitters are less expensive and easier to control than a single large aperture. As an added benefit, the far-field spot can be steered simply by changing the phase relationship of the emitters, which is a fast electronic operation, rather than by turning a single large, heavy lens or mirror, which is a slow mechanical operation. (In certain configurations, the individual emitters may need to be steered in-place as well.)

In order for the emitted pulses to interfere in the far field, they need to arrive at the target at the same time. If the emitter array is large enough, the pulses are short enough, and the far-field spot is steered far enough to one side, the pulses won't arrive simultaneously at the target. Slow light can provide a tunable time delay, known as a *true-time delay*, so that each pulse is emitted at the right time and the pulses are synchronized when they arrive at the target [2–5, 20, 22, 140, 153]. Slow light can enhance the precision of the laser radar

by allowing shorter pulses to be used. It can also increase the laser radar's field of view, by allowing the beam to be steered farther off-axis.

The pulses still need to be kept in phase with each other, so that they are coherent in the far field. If the pulses are not coherent, they won't interfere to create a single small, bright spot. Instead, they will add incoherently, like multiple flashlights shining on a single spot: the far-field spot will remain large (poor transverse resolution). An array of incoherent emitters does not have any better transverse resolution than a single small emitter [6, 152, 154–159].

Loss of coherence among the emitters also leads to reduced brightness in the far field, as the same amount of optical power is distributed over a wider far-field spot. The Strehl ratio is the ratio of the far-field spot's actual brightness to its theoretical maximum brightness, and it serves as a measure of the coherence of the pulses. When the pulses are fully coherent, the Strehl ratio is one (the maximum). As the pulses become less coherent with each other, the Strehl ratio is reduced, until it becomes the reciprocal of the number of emitters [6]; for example, if there are ten emitters and they are fully incoherent, the Strehl ratio will be 0.1. A reduced Strehl ratio means the target is being illuminated with less intense light, so the target reflects a lower intensity and a lower intensity reaches the optical detector. This can reduce the signal-to-noise ratio, which in turn can reduce the maximum range or longitudinal resolution of the laser radar [152].

A slow-light laser radar called SLIDAR has been demonstrated [3–6]. The system consists of three emitters arranged in a right triangle, to allow two-dimensional steering. Two

different slow light mechanisms allow compensation of the group delay in both transverse dimensions. In this chapter, I present the phase locking techniques used in SLIDAR. The phase locking system is simple, inexpensive, and reasonably effective, and it allows the proper phasing of each emitter. An electro-optic phase modulator (EOM) in each emitter's signal channel acts as the phase controller, and the signal is monitored using a single-quadrature phase detector in a heterodyne configuration. A feedback control circuit maintains phase lock with a root-mean-square (RMS) phase error of $\pi/5$ radians (1/10 wave), giving a Strehl ratio of 0.8. The EOM has only a finite phase compensation range, while the phase errors of the system can accumulate indefinitely; this problem is overcome using a novel fast $2n\pi$ phase snapback circuit, which rapidly changes the EOM phase by $2n\pi$ radians when the EOM nears the end of its phase compensation range.

2.1 Theory

In a phased-array laser radar with a fixed emitter configuration and constant power, the intensity of the received optical signal is directly proportional to the Strehl ratio, or the ratio of the actual peak combined intensity to the theoretical maximum in the absence of phase noise [154, 155]. A hypothetical phase locking system with infinite bandwidth (instantaneous response) would have no residual phase error, and its Strehl ratio would be 1. Of course, all causal phase locking systems will have some finite residual phase error, so the actual Strehl ratio will be less than 1. In the total absence of phase lock, the Strehl ratio drops to $1/N$ for an N -emitter system, and the transverse far-field resolution is degraded

to the resolution of a single emitter, negating nearly all of the benefits of using multiple emitters. The importance of phase lock to the proper operation of a coherently combined phased-array system cannot be overstated.

The Strehl ratio can be estimated based on the measured phase noise of the system. In this section, I model the phase errors of SLIDAR and find their influence on the Strehl ratio. In doing so, I follow the derivation given by Nabors [154] for the grating Strehl ratio (neglecting the impact of the emitter field profile).

The system has two main contributions to the phase error, both of which can be seen in Fig. 2.1. The first is the residual phase error, or the small amount of phase noise that remains after correction by the phase control system. The second is the error due to $2n\pi$ phase snapback events. Both are random processes, and their statistics depend strongly on environmental conditions, such as temperature fluctuations and vibrations. To gain some insight into their respective effects, I make some simplifying assumptions. First, I take the residual phase noise to have a normal distribution with zero mean and variance σ^2 , and I assume that the phase noise at one moment is independent of the phase noise at a different time (meaning the residual phase noise is a stationary process). Second, I model the snapback process as a linear $2n\pi$ phase ramp of duration τ that occurs at random times, with the snapback times having a Poissonian distribution with rate r . Finally, I assume each emitter's noise processes are statistically independent of those of the other emitters.

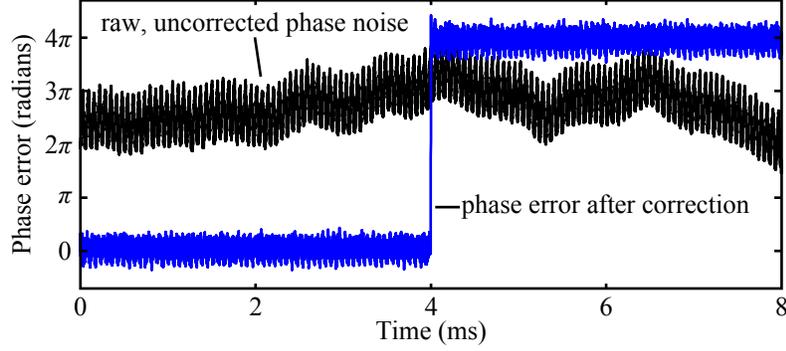


Figure 2.1. Phase noise before and after correction by the phase control system. In this 8 ms sample, the phase error consists of two parts: the zero-mean residual phase error and the 4π snapback event that occurs at 4 ms. Only the phase error modulo 2π affects system performance. (From Ref. 6; used by permission.)

2.1.1 The effect of residual phase error

Assume that each of N emitters have identical field profiles $u_0(x, y)$, each with identical amplitude, relative position (x_j, y_j) , and random phase ϕ_j . Then the total field is given by

$$u(x, y) = u_0(x, y) * \sum_{j=1}^N \delta(x - x_j, y - y_j) \exp(i\phi_j) = u_0(x, y) * g(x, y), \quad (2.1)$$

where the asterisk indicates convolution, and $\delta(x, y) = \delta(x)\delta(y)$ is the two-dimensional Dirac delta. The function $g(x, y)$ is called the near-field grating function. In the Fraunhofer regime, the far field $U(f_x, f_y)$ is equal to $U_0(f_x, f_y)G(f_x, f_y)$, where

$$U_0(f_x, f_y) = \iint u_0(x, y) \exp[i2\pi(f_x x + f_y y)] dx dy \quad (2.2)$$

is the far-field pattern produced by a single emitter, and

$$G(f_x, f_y) = \sum_{j=1}^N \exp[i2\pi(f_x x_j + f_y y_j) + i\phi_j]. \quad (2.3)$$

is the far-field grating pattern. The far-field intensity can then be expressed as

$$I(f_x, f_y) = |U(f_x, f_y)|^2 = |U_0(f_x, f_y)|^2 |G(f_x, f_y)|^2. \quad (2.4)$$

Assuming the far-field intensity has a peak at $I(0, 0)$, the Strehl ratio can be defined as

$$S = \frac{\langle I(0, 0) \rangle}{I_{\max}(0, 0)} = \frac{|U_0(0, 0)|^2 \langle |G(0, 0)|^2 \rangle}{|U_0(0, 0)|^2 [|G_{\max}(0, 0)|^2]} = \frac{\langle |G(0, 0)|^2 \rangle}{|G_{\max}(0, 0)|^2}, \quad (2.5)$$

where angle brackets indicate the expected value. Evaluating $|G(0, 0)|^2$ gives

$$G(0, 0)G^*(0, 0) = \left[\sum_{j=1}^N \exp(i\phi_j) \right] \left[\sum_{j=1}^N \exp(-i\phi_j) \right] = \sum_{j=1}^N \left[1 + 2 \sum_{k=j+1}^N \cos(\phi_j - \phi_k) \right], \quad (2.6)$$

which reaches its maximum value of N^2 when the ϕ_j are equal. Then, $|G_{\max}(0, 0)|^2 = N^2$,

and

$$S = \frac{1}{N^2} \langle |G(0, 0)|^2 \rangle = \frac{1}{N^2} \sum_{j=1}^N \left[1 + 2 \sum_{k=j+1}^N \langle \cos(\phi_j - \phi_k) \rangle \right]. \quad (2.7)$$

Considering only the effects of the zero-mean (normally distributed) residual phase error, the innermost term of Eq. (2.7) becomes

$$\langle \cos(\phi_j - \phi_k) \rangle = \iint \frac{1}{2\pi\sigma^2} \exp\left(-\frac{\phi_j^2 - \phi_k^2}{2\sigma^2}\right) \cos(\phi_j - \phi_k) d\phi_j d\phi_k = \exp(-\sigma^2). \quad (2.8)$$

Then

$$S = \frac{1}{N^2} \sum_{j=1}^N \left[1 + 2 \sum_{k=j+1}^N \exp(-\sigma^2) \right] = \frac{1}{N} + \frac{N-1}{N} \exp(-\sigma^2). \quad (2.9)$$

(This expression matches the expression given by Nabors for the case of uncorrelated phase errors [154].) One can see this expression is the sum of incoherent and coherent contributions. In the limit of complete incoherence among the emitters (large σ), the Strehl ratio is $1/N$, as would be expected for an incoherent sum. In the case of perfect coherence and perfect phase synchronization ($\sigma = 0$), the Strehl ratio becomes 1.

2.1.2 The effect of snapbacks

A snapback event happens when the EOM reaches the end of its range and is quickly shifted, or snaps back, by $2n\pi$ in order to maintain phase lock. Here, I model the snapback event as a linear phase ramp, shifting the phase by $2n\pi$ at a constant rate over an interval τ . I define the total phase error of each emitter as the sum of two stochastic processes

$$\phi_j(t) = \phi_{R,j}(t) + \phi_{S,j}(t), \quad (2.10)$$

where $\phi_{R,j}(t)$ is the (normally distributed) residual phase error of emitter j , as described above. Here, $\phi_{S,j}(t)$ represents the snapback process in emitter j as follows: If a snapback event occurs at time t_0 , then $\phi_{S,j}(t) = 2n\pi(t - t_0)/\tau$ for $t_0 \leq t < t_0 + \tau$, and τ is the duration of a snapback phase ramp. For times at which no snapback occurs, $\phi_{S,j}(t) = 0$. Snapback events occur at random times, following a Poisson process; that is, they occur with an average rate r (measured in events per unit time), and the probability that m snapback events occur in an interval T is a Poisson random variable with dimensionless parameter rT

$$P(m; T) = \exp(-rT) \frac{(rT)^m}{m!}. \quad (2.11)$$

Note that the average rate r of snapback events does not change with time, although the number of snapbacks in any interval depends on the length of the interval and will be rT on average.

The snapback duty cycle, or average fraction of time spent in a snapback process, is equal to $r\tau$. Intuitively, the snapback duty cycle must be kept low [160], so one may assume $r\tau \ll 1$. A further simplifying assumption is that snapback events never overlap in time, either on the same emitter or on different emitters. (It can be shown that this assumption is slightly pessimistic, but in the limit of low snapback duty cycle it is accurate.)

To understand how snapback events impact system performance, consider the time-averaged Strehl ratio. Define the time average of the Strehl ratio over an interval from t_0 to

t_1 as

$$\langle S(t) \rangle_{t_0, t_1} = \frac{1}{t_1 - t_0} \int_{t_0}^{t_1} S(t) dt. \quad (2.12)$$

By substituting the definition of the Strehl ratio from Eq. (2.7),

$$\begin{aligned} \langle S(t) \rangle_{t_0, t_1} &= \frac{1}{N^2} \langle \langle |G(\mathbf{0}, \mathbf{0}; t)|^2 \rangle \rangle_{t_0, t_1} = \left\langle \left\langle \sum_{j=1}^N \left[1 + 2 \sum_{k=j+1}^N \cos(\phi_j(t) - \phi_k(t)) \right] \right\rangle \right\rangle_{t_0, t_1} \\ &= \frac{1}{N^2} \sum_{j=1}^N \left\{ 1 + 2 \sum_{k=j+1}^N \langle \langle \cos[\phi_{R,j}(t) - \phi_{R,k}(t) + \phi_{S,j}(t) - \phi_{S,k}(t)] \rangle \rangle_{t_0, t_1} \right\}. \end{aligned} \quad (2.13)$$

The subscripted angle brackets indicate a time average, while angle brackets without a subscript indicate a statistical average.

Next, consider the impact of a single snapback event on the expression in angle brackets in Eq. (2.13). Assume that emitter j experiences a $2n\pi$ snapback event of duration τ , such that $\phi_{S,j}(t) = 2n\pi(t - t_0)/\tau$ over the interval t_0 to $t_0 + \tau$, while emitter k experiences no snapback event. Then

$$\begin{aligned} &\langle \langle \cos[\phi_{R,j}(t) - \phi_{R,k}(t) + 2n\pi(t - t_0)/\tau] \rangle \rangle_{t_0, t_0+\tau} \\ &= \left\langle \left\langle \int \int \frac{1}{2\pi\sigma^2} \exp\left(\frac{-\phi_j^2 - \phi_k^2}{2\sigma^2}\right) \cos[\phi_j - \phi_k + 2n\pi(t - t_0)/\tau] d\phi_j d\phi_k \right\rangle \right\rangle_{t_0, t_0+\tau} \\ &= \langle \exp(-\sigma^2) \cos[2n\pi(t - t_0)/\tau] \rangle_{t_0, t_0+\tau} = 0. \end{aligned} \quad (2.14)$$

In other words, an emitter undergoing a snapback makes no coherent contribution to the time-averaged Strehl ratio during the time that the snapback is occurring. Conversely, if

there is no snapback occurring, the coherent contribution to the time-averaged Strehl ratio is as before,

$$\langle \langle \cos[\phi_{R,j}(t) - \phi_{R,k}(t)] \rangle \rangle_{t_0, t_0+\tau} = \langle \exp(-\sigma^2) \rangle_{t_0, t_0+\tau} = \exp(-\sigma^2). \quad (2.15)$$

One can think of each snapback (of duration τ) occurring in an interval T as reducing the coherent contribution to the time-averaged Strehl ratio by an amount τ/T . If m snapback events happen during the interval T , the coherent contribution will be reduced by $m\tau/T$. (For simplicity, I neglect edge cases, or snapbacks that happen very near the end of the interval.) The probability that m snapbacks occur in an interval T is given by Eq. (2.11). The expected value of the time-averaged Strehl ratio over an interval T with a given average snapback rate r and residual phase noise variance σ^2 is then

$$\begin{aligned} \langle S(t) \rangle_{t_0, t_0+T} &= \frac{1}{N^2} \langle \langle |G(0, 0; t)|^2 \rangle \rangle_{t_0, t_0+T} \\ &= \frac{1}{N^2} \sum_{m=0}^{\infty} P(m; T) \sum_{j=1}^N \left[1 + 2 \sum_{k=j+1}^N \left(1 - \frac{m\tau}{T} \right) \exp(-\sigma^2) \right] \\ &= \frac{1}{N} + \frac{N-1}{N} (1 - r\tau) \exp(-\sigma^2). \end{aligned} \quad (2.16)$$

By comparing Eq. (2.16) to Eq. (2.9), one can see that the effect of snapbacks is to reduce the coherent contribution to the time-averaged Strehl ratio by the factor $(1 - r\tau)$.

As one might expect, the result given in Eq. (2.16) does not depend on t_0 or T , only on the snapback duty cycle $r\tau$. To maximize the Strehl ratio, one must minimize both the rate r

at which snapback events occur and their duration τ (or minimize the snapback duty cycle, coinciding with the assumption $r\tau \ll 1$). Interestingly, n , the number of 2π phase cycles in each snapback event, influences the time-averaged Strehl ratio only indirectly. Increasing n may decrease r , since larger jumps may mean fewer snapbacks, but it will generally also increase τ , since a larger phase change will take longer to complete. The balance of these effects depends on the actual parameters of the phase control system and the statistics of the (uncorrected) phase noise. Of course, maximizing the Strehl ratio also means minimizing σ .

It bears repeating that the phase noise in a given system is highly dependent on several variables, particularly temperature and vibration. The assumption of a stationary normal (Gaussian) noise process is inaccurate for many models of phase noise, which often have an RMS value that increases with time (and are therefore non-stationary processes). The rate at which snapbacks occur depends highly on the phase noise in the system and may well be non-Poissonian. Nevertheless, the above derivation gives a useful picture of how the residual phase error and the snapback process affect the Strehl ratio, and in particular that the snapback process does not significantly affect the Strehl ratio if the duty cycle is kept low.

2.2 Experimental apparatus

2.2.1 Optical system

The layout of the optical system is shown in Fig. 2.2. A fiber-coupled tunable laser generates about 1 mW of continuous-wave (cw) light near 1550 nm. The laser feeds both a phase reference arm and a pulse carver followed by several signal channels (three channels in this system). The pulse carver consists of an intensity modulator driven by an arbitrary function generator, as well as an erbium-doped fiber amplifier (EDFA) and fiber polarization controllers. The reference arm is frequency-shifted by 55 MHz using an acousto-optic modulator (AOM), and the frequency-shifted reference is used for heterodyne detection to monitor the phase of each signal channel and maintain phase lock. The local oscillator (LO) used by the AOM to frequency-shift the reference is also used by the phase locking electronics in each channel.

Each signal channel contains a slow light delay module consisting of 2.2 km of single-mode fiber (SMF), either dispersion-compensating fiber (DCF) or dispersion-shifted fiber (DSF). The ratio of DCF to DSF in each module creates a relative dispersive delay between each signal arm [30]. One of the delay modules also contains a counterpropagating pump (not shown in Fig. 2.2) to induce stimulated Brillouin scattering (SBS) slow light [2, 22]. By controlling the wavelength of the optical field and the pump power of the SBS module, one can achieve independent group delay compensation in two orthogonal transverse dimensions. Details of the group delay modules have been reported previously [3, 4].

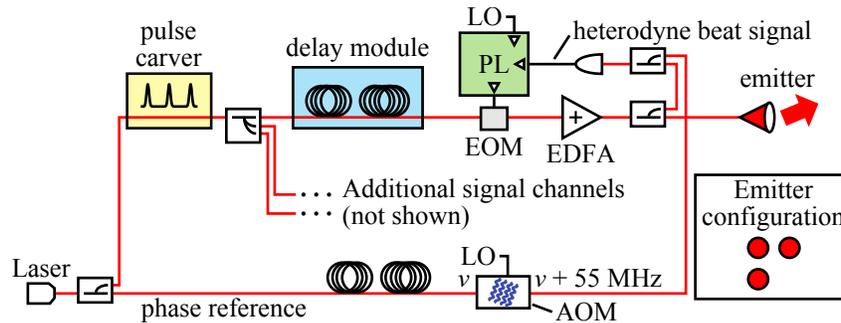


Figure 2.2. SLIDAR optical system schematic diagram. The output from a continuous-wave laser is split between a pulse carver, which feeds three signal channels containing delay modules, and a phase reference arm with an acousto-optic modulator (AOM) to shift the laser frequency, driven by a 55 MHz local oscillator (LO). A fraction of each signal channel's output signal is heterodyned with the frequency-shifted phase reference, and phase locking electronics (PL) use the heterodyne beat signal and the LO to determine the phase error and drive an electro-optic modulator (EOM) to maintain phase lock. Inset: Emitter configuration, as seen looking into the emitters. (From Ref. 6; used by permission.)

After the delay module, each signal channel contains an EOM to compensate for phase errors, an EDFA to maintain constant output amplitude (particularly important to compensate for the variable gain of SBS slow light), and an emitter (a fiber output coupler), as well as a splitter to tap off a fraction of the output signal, a heterodyne detection setup, and phase locking electronics, described below. All channels are phase locked to the same reference. The system uses three signal channels, with the three emitters in a right triangle or L-shaped arrangement that allows two-dimensional steering (shown in the inset of Fig. 2.2). Note that the short length of fiber in each channel between the last splitter and the emitter cannot be monitored for phase noise; this uncompensated fiber length is kept as short as possible to minimize the phase noise it contributes to the system.

The pulse carver creates 6.5 ns pulses (FWHM) on a cw background of about 30% of the peak power. The pulse duration is shorter than the response time of the phase control

electronics, so the phase control circuit does not “see” the pulse. Only the cw background is used by the phase control electronics.

2.2.2 Electronics

A block diagram of the phase control electronics is shown in Fig. 2.3. In essence, a phase-sensitive detection is performed on each signal channel’s heterodyne signal, using a phase-shifted version of the 55 MHz local oscillator as a reference; the result is a voltage proportional to the phase difference between the heterodyne and the LO, which is passed to a proportional-integral (P-I) controller. The phase-shifting of the LO allows each channel to be phased independently, allowing full control over the far-field pattern. The P-I controller consists of a loop filter with gain and a fast $2n\pi$ phase snapback circuit. It filters the error signal and drives the EOM. The design of the P-I controller is shown in Fig. 2.4. PSpice circuit simulations and laboratory measurements of the loop filter transfer function show the loop filter bandwidth to be about 2.3 MHz.

The phase error accumulates rapidly in the P-I controller, and the control signal to the EOM cannot exceed the supply voltages of the loop filter. The half-wave voltage (V_π) of the EOM used in this experiment is nominally 4 V, and the loop filter uses ± 15 V supply voltages, giving a phase actuation range of somewhat over 6π (allowing some margin). To maintain phase lock over many optical cycles, a fast $2n\pi$ phase snapback circuit is implemented, which rapidly adds or subtracts $2n\pi$ radians from the phase error being tracked by the P-I controller whenever the controller’s output voltage approaches a voltage

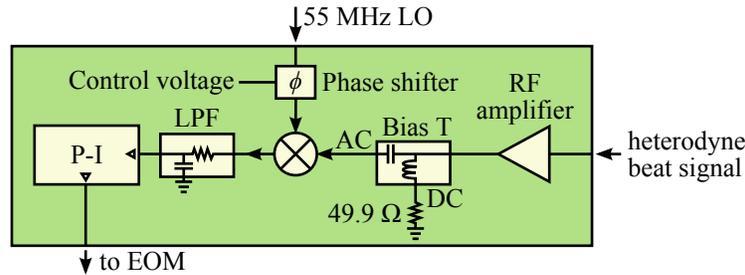


Figure 2.3. Block diagram of the SLIDAR phase control electronics. The 55 MHz local oscillator (LO) is phase-shifted to allow independent, arbitrary phasing of each channel. The heterodyne signal is amplified, its DC bias is removed by the bias T, and it is mixed with the LO and then low-pass filtered (LPF) to perform a phase-sensitive detection of the heterodyne signal. The resulting signal is then fed to the proportional-integral (P-I) controller. (From Ref. 6; used by permission.)

supply level. In this way, the P-I controller is able to track an unbounded amount of phase error. These brief (about $1.2 \mu\text{s}$) phase snapback events do not impact performance as long as the snapback duty cycle is kept low, as discussed in section 2.1. The operation of the snapback circuit is as follows: When the loop filter's output V_{EOM} exceeds a certain threshold (nominally $3V_{\pi}$ or $+12 \text{ V}$), an analog switch engages and strongly drives V_{EOM} toward negative voltages. Once V_{EOM} is below a second, lower threshold (nominally V_{π} or $+4 \text{ V}$), the analog switch disengages, and the loop filter resumes tracking, having shifted its operating point by 2π radians of phase error. A corresponding process occurs for negative values of V_{EOM} . In practice the exact threshold voltages (as well as the resistor and capacitor values) are tuned empirically to account for several parameters such as the drive strength (maximum output current) of the analog switch and the response times of the switch and the loop filter, and the circuit is typically adjusted to provide a 4π phase snapback. (It was verified experimentally that the op-amp never enters saturation because the switch is disengaged as soon as V_{EOM} crosses the appropriate threshold.)

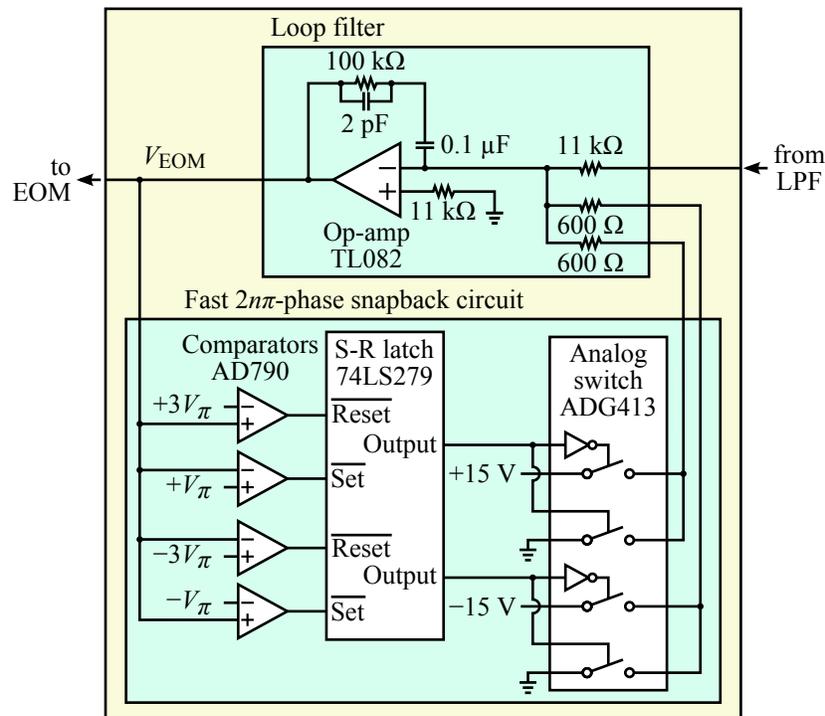


Figure 2.4. Block diagram of the proportional-integral (P-I) controller. Includes the loop filter and the fast $2n\pi$ phase snapback circuit. The signal from the low-pass filter (LPF) in Fig. 2.3 is the primary input to an integrating op-amp that tracks the phase error and drives the EOM. When the op-amp output voltage V_{EOM} exceeds $\pm 3V_\pi$ (nominally ± 12 V), a set-reset (S-R) latch engages an analog switch that quickly drives V_{EOM} toward the opposite polarity; the switch disengages once V_{EOM} is within the range $\pm V_\pi$ (nominally ± 4 V). The process typically takes about $1.2 \mu\text{s}$. In practice, the comparator bias voltages and the resistor and capacitor values are tuned for optimal snapback performance and are usually selected to provide a 4π snapback. Some circuit elements are not shown, including buffering op-amps, diodes, and a power-on reset circuit. (From Ref. 6; used by permission.)

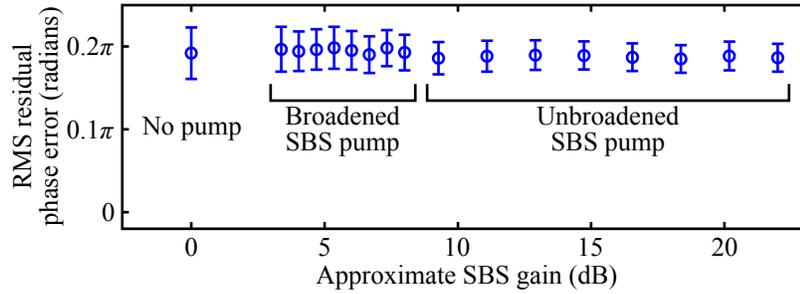


Figure 2.5. RMS residual phase error vs. approximate SBS gain. The RMS residual phase error is about $\pi/5$ radians for all SBS gain values, with the SBS pump both broadened and unbroadened. The dispersion-compensating fiber used to achieve dispersive slow light performs similarly to the dispersion-shifted fiber used here to achieve SBS slow light, including under changes in wavelength. (From Ref. 6; used by permission.)

2.3 Results and discussion

Figure 2.5 shows the phase locking performance of the phase control system in a single SLIDAR channel containing 2.2 km of single-mode DSF, averaged over 20 s. Over this time window, the RMS residual phase error is approximately $\pi/5$ radians (1/10 wave). According to Eq. (2.16), the theoretically predicted Strehl ratio due to this RMS residual phase error is about 0.8, which is comparable to results obtained in several other optical phase locking systems [160, 161].

Phase lock is maintained in the presence of both SBS slow light and dispersive slow light over the full range of operation. The varying SBS gain is compensated by the final EDFA in each signal channel, which ensures a constant heterodyne signal level. Some slight pulse shape distortion due to the SBS process is evident but does not affect the phase locking performance. Broadening of the SBS pump, a common technique in SBS slow light [2, 22], is found not to induce any differences in phase locking behavior. DCF, which is used only

for dispersive slow light and not SBS slow light, sees the same phase locking performance as unpumped DSF; further, the performance does not vary with wavelength.

The residual phase error was obtained by using an oscilloscope in a fast acquisition mode to measure the RMS jitter between the local oscillator and the heterodyne signal. Multiplying the RMS jitter (measured in picoseconds) by 2π times the LO frequency gives the RMS phase error. The measurements were taken while the phase lock was operating, and they include both optical and electronic contributions. The oscilloscope's fast acquisition mode allows repeated measurements of the RMS jitter over any desired duration, but at the expense of losing the time traces, which cannot be directly recorded. The residual phase error was observed not to change appreciably between acquisition times of 1 s and 20 s.

The uncorrected phase noise cannot be measured directly in this experiment, but adding the residual phase error and the applied phase shift due to the EOM gives an estimate of the phase noise. A sample trace of the phase noise obtained in this way was presented earlier, in Fig. 2.1. The applied phase shift is estimated from a time trace of the EOM control voltage (V_{EOM}). The residual phase error is computed by performing software heterodyne detection (mixing and low-pass filtering) on time traces of the local oscillator and the heterodyne signal. Because the time traces must be sampled at more than twice the LO frequency of 55 MHz, the oscilloscope buffer is limited to traces of 8 ms duration.

Figure 2.6 shows an image of the far-field pattern. The pattern is observed to shift rapidly without phase locking and to maintain its position with phase locking. (A real-time video of the far-field pattern is available; see Ref. 6.) In this particular laboratory environment,



Figure 2.6. Far-field pattern of the three-channel SLIDAR system. This is a representative image from a real-time video showing the far-field pattern, both with and without phase locking (image taken when phase locking is enabled). When phase locking is disabled, the bright lobes in the far-field pattern shift rapidly, both horizontally and vertically, since the points of constructive interference change with the emitter phases. When phase locking is enabled, the lobes lock into place, with a small but still visible amount of jitter. (The video is available online; see Ref. 6.)

the primary sources of phase noise are believed to be temperature changes and vibration [160]. When observed for several minutes, the phase noise tends to drift continuously in one direction (either positive or negative); this is believed to be due to diurnal temperature changes. Under these conditions, using a 4π snapback rather than 2π tends to reduce the rate of snapback events. The snapback rate r fluctuates, but the measured average is about two snapback events per second. A single snapback event takes $1.2 \mu\text{s}$, so the snapback duty cycle $r\tau$ is about 2.4×10^{-6} . At this low level, the snapback process has a negligible effect on the Strehl ratio; even a duty cycle several orders of magnitude larger would not impact the Strehl ratio noticeably. In this system, the maximum achievable snapback size is 4π , limited by the drive strength of the analog switch, the response times of various components, and other factors, as discussed in section 2.2.2. While a larger snapback size might further

reduce the impact of the snapback cycles on the residual phase error (see section 2.1.2), the present system provides adequate performance for the purposes of SLIDAR.

Other researchers have developed several quite versatile phase locking techniques, such as optical phase-locked loops (OPLL) [158, 159, 162] and LOCSET [163, 164]. However, the slow light system precludes the use of tuning the source laser wavelength, as in an OPLL, and requires that all signal channels operate at the same frequency, unlike LOCSET. Further, the phase locking system presented here does not require the complex and expensive signal processing electronics used by such feedback control methods as stochastic parallel gradient descent (SPGD) optimization [165–167]. (Of course, a fast $2n\pi$ snapback process can certainly be implemented in a digital signal processing system.)

2.4 Conclusion

The phase locking system presented here is simple and effective, using an electro-optic phase modulator and a feedback circuit with a fast $2n\pi$ phase snapback to allow unlimited phase actuation range. The system maintains phase lock among three channels of 2.2 km each of single-mode fiber while utilizing two different slow light techniques, with the RMS residual phase error around $\pi/5$ radians and the resulting Strehl ratio around 0.8. The phase locking approach is scalable to many more channels, as each channel is independently locked to an optical reference. The approach presented here should also be applicable to coherent beam combining, in which fiber lengths are typically limited to tens of meters.

Chapter 3

Multivariate Optical Computing for Spectrum Recognition

Spectroscopy is a proven way to identify solids, liquids, and gases. Every substance has a unique spectrum, like a fingerprint. *Molecular identification*, *chemical identification*, and *spectrum recognition* are names for the task of analyzing a spectrum to determine what substance produced it. Usually, to identify a substance, a spectrometer digitally records a spectrum, and then a computer algorithm analyzes the digital spectrum and matches its features against a database of known spectra. The most common algorithms are part of a class of computationally intensive methods called multivariate regression techniques, such as partial least squares (PLS) and principal component analysis (PCA).

But even before the spectrum reaches the computer, it must first be recorded, and that can take quite a bit of time. Many spectral signatures are very weak. Weak signals require a long time to detect, usually known as integration time. For example, in Raman spectroscopy, one of the most popular modes of spectroscopy for molecular identification,

it is not uncommon to record for several minutes before obtaining an adequate spectrum, compared to milliseconds for an absorption spectrum with high brightness.

Optical processing can alleviate these problems. A traditional optical spectrometer records all of the features of a spectrum at once, dispersing light over an array of detectors. But when one knows what one is looking for, it can be much more efficient to create a set of custom optical elements, called spectral matched filters, that match the spectra of the substances of interest. Light passing through a spectral matched filter produces a bright signal when the incoming spectrum matches the filter or a dark signal when they are mismatched. The output of the spectral matched filter is focused onto a single detector, instead of being spread out over a whole array of detectors, so the signal strength much higher (known as the Fellgett advantage, or multiplex advantage [168]). The integration times can be shorter and the detection process is faster and more sensitive—and because the spectral matched filters are customized to the spectra of interest, the system can be just as selective as a traditional spectrometer. This approach has been called by many names, including *multivariate optical computing*: *multivariate* because of its relationship to multivariate regression, and *optical computing* because the computation is done optically rather than digitally. Other names include compressive detection, predictive spectroscopy, optical regression, and information-efficient spectroscopy. Multivariate optical computing has other potential benefits, such as improved environmental robustness and reduced need for digital computer power. It may even be possible to miniaturize multivariate optical

computers and make them part of on-chip spectrum recognition systems. These are important advantages for field-deployed molecular identification systems.

A number of multivariate optical computer (MOC) designs have been proposed in the past. Some approaches use a fixed spectral matched filter (sometimes called a multivariate optical element), made for instance from a thin-film interference filter [169, 170] or a custom holographic grating [171–174]; however, such filters are difficult to make and are not reconfigurable. Another type of MOC disperses the light spatially, such as with a grating, and then uses spatial filters to modify the spectrum [175]. When the filter is implemented using a computer-controlled spatial light modulator (SLM) or digital micromirror device (DMD), the filter is easy to make and easy to reconfigure [176–182]. Related techniques have been proposed for use in hyperspectral imaging [183–188]. (In this chapter, the term *SLM* is used to refer only to liquid crystal SLMs and not to DMDs, although some of the principles presented below also apply to DMDs.)

In this chapter, I present a MOC that uses a spatial light modulator in a configuration that is potentially more efficient than existing designs, allowing chemicals to be identified more rapidly. The liquid crystal SLM can implement arbitrary filter coefficients natively; DMD mirror states are binary and arbitrary coefficients must be implemented by toggling the mirrors and averaging the results, resulting in increased complexity and potentially slower measurements. Previous SLM-based designs have implemented only a single spectral matched filter at a time on the SLM, while the present design can have more than one simultaneous filter. The advantages of this new design could allow even faster molecular

identification. (Additionally, this design uses a commercial, off-the-shelf SLM, unlike some previous designs based on custom-made SLMs.) The experimental demonstration of the technique in this chapter uses absorption spectroscopy, but it is equally applicable to Raman spectroscopy and other emission-type modes of spectroscopy.

The primary improvement over previous techniques is in the use of the SLM: A vertical input slit is imaged onto a reflective SLM. A grating in the imaging path disperses the slit images horizontally across the SLM array. The SLM pixel columns impart a tilt onto the reflected light using a blazed phase grating; the period of the blazed grating controls the tilt angle, and the reflectance can be varied by controlling the efficiency of the blazed grating. In this way, the SLM implements an arbitrary set of spectral filter coefficients (reflectances). Through two methods covered in section 3.4, the SLM can be divided into multiple zones, each of which directs light toward a different detector. This allows multiple spectral matched filters to be implemented simultaneously on the SLM, which speeds up the molecular identification task.

3.1 Spectral correlation

Spectral matched filtering, or spectral correlation, is often used to determine the resemblance of a sample spectrum to a set of known spectra. It is the fundamental optical computation operation performed by a MOC. An input spectrum s_j is one of N_S different spectra, each representing a different substance or chemical to be identified. The spectrum s_j is

passed separately through each of several different matched filters \mathbf{a}_k (these matched filters are implemented on the SLM). The total transmission through each filter is equal to the correlation (inner product) $c_{j,k}$ of the filter's spectral response and the input spectrum

$$c_{j,k} = \mathbf{s}_j^T \mathbf{a}_k. \quad (3.1)$$

(Here, lowercase italicized letters represent scalar quantities, lowercase bold letters represent column vectors, uppercase bold letters represent matrices, the superscript T indicates a vector or matrix transpose, and discretized spectra and filters are represented as N -element vectors.) If the filters are properly chosen, the filter corresponding to \mathbf{s}_j will have the highest total transmitted power, and other filters will have no transmitted power. The ideal results are $c_{j,k} = \delta_{j,k}$, where $\delta_{j,k}$ is the Kronecker delta ($\delta_{j,k} = 1$ if $j = k$ and 0 otherwise). The spectrum recognition operation consists of performing spectral correlations of the spectra \mathbf{s}_j with the matched filters \mathbf{a}_k that result in correlation coefficients $c_{j,k} = \delta_{j,k}$.

What are the optimal filters \mathbf{a}_k that identify the different chemicals? If the source spectra and filters are written as rectangular $N \times N_S$ matrices $\mathbf{S} = (\mathbf{s}_1 \ \mathbf{s}_2 \ \dots)$ and $\mathbf{A} = (\mathbf{a}_1 \ \mathbf{a}_2 \ \dots)$, then the ideal response for the spectrum recognition task is

$$\mathbf{C} = \mathbf{S}^T \mathbf{A} = \mathbf{I}, \quad (3.2)$$

where \mathbf{C} is a correlation matrix whose elements are $c_{j,k}$ (j is the column index, and k is the row index), and \mathbf{I} is the $N_S \times N_S$ identity matrix. There are many possible matrices \mathbf{A} that

satisfy Eq. (3.2), assuming the spectra that make up the columns of \mathbf{S} are linearly independent. One solution is found using the Moore–Penrose pseudoinverse. The pseudoinverse of \mathbf{S}^T , written as $(\mathbf{S}^T)^+$, has the property that

$$(\mathbf{S}^T)^+\mathbf{S}^T = \mathbf{I}. \quad (3.3)$$

Left-multiplying Eq. (3.2) by the pseudoinverse and substituting Eq. (3.3) yields

$$(\mathbf{S}^T)^+\mathbf{C} = (\mathbf{S}^T)^+\mathbf{S}^T\mathbf{A} = (\mathbf{S}^T)^+\mathbf{I}. \quad (3.4)$$

Simplifying Eq. (3.4) yields

$$\mathbf{A} = (\mathbf{S}^T)^+. \quad (3.5)$$

In other words, the ideal matched filters \mathbf{a}_k are the columns of $(\mathbf{S}^T)^+$. (In fact, the \mathbf{a}_k may be the columns of *any* left pseudoinverse of \mathbf{S}^T , not just the Moore–Penrose pseudoinverse. For simplicity the Moore–Penrose inverse is used exclusively here, but there is the possibility that a different left pseudoinverse may have experimental advantages.)

3.1.1 Measuring concentrations and mixtures

Smith and coworkers showed in Ref. 180 that these same filters \mathbf{a}_k can be used to measure linear combinations, or mixtures, of the \mathbf{s}_j . If a spectrum $c_j\mathbf{s}_j + c_{j'}\mathbf{s}_{j'}$ is incident on these filters, then the inner product of the spectrum with \mathbf{a}_j is c_j , and the inner product with $\mathbf{a}_{j'}$ is

$c_{j'}$. For Raman spectroscopy, or other modes of spectroscopy in which the spectral intensity is proportional to the chemical concentration, the coefficients c_j and $c_{j'}$ are proportional to the molar concentrations of the two components. (This assumes the components do not react when mixed, and no nonlinear or other optical interactions occur between the components.) If one is careful to observe proper scaling when recording the s_j , then the spectral matched filters \mathbf{a}_k will produce correlation coefficients c_j and $c_{j'}$ exactly equal to the molar concentrations.

However, absorption spectroscopy is governed by the Beer–Lambert–Bouguer Law, which states that the light transmitted through an absorbing chemical is proportional to the negative exponential of the concentration:

$$I/I_0 = e^{-\epsilon Lc}, \quad (3.6)$$

where I/I_0 is the transmitted intensity fraction, ϵ is the molar extinction coefficient, L is the optical path length through the chemical, and c is the molar chemical concentration. (Equivalently, ϵ times c is equal to the absorption cross-section times the number density.) Due to this exponential relationship, changes in concentration lead to a change in the shape of the absorption spectrum, not a simple linear scaling of the spectral intensities. At low concentrations, the intensity transmission is nearly linear with the concentration. To measure higher chemical concentrations using absorption spectroscopy, one must effectively linearize the exponential, by taking the spectra of several different concentrations of a single chemical and treating these naively as independent “components,” then interpolating

the actual concentration [170]. Multiple \mathbf{s}_j are used to represent different concentration ranges of a single chemical. If mixtures of multiple chemicals are used, then in general, both the individual chemical spectra and the spectra of their mixtures must be represented among the \mathbf{s}_j , including mixtures with varying chemical concentrations [189]. Despite these difficulties, absorption spectroscopy is a convenient laboratory alternative for testing the optical computation methods described here, and the methods themselves apply equally well to other forms of spectroscopy.

3.1.2 Implementation on an SLM

The matched filters \mathbf{a}_k consist of real numbers, both positive and negative, of any scale. The SLM, though, can only implement reflectances (matched filter values) between zero and one. The most common approach is to separate the positive and negative values into separate filters, and scale each filter so its maximum value is unity (and scale the detected power by the corresponding factor).

Define two filters \mathbf{a}_k^+ and \mathbf{a}_k^- containing respectively the positive and negative elements of \mathbf{a}_k , such that the elements of each filter are

$$a_{k,i}^+ = \begin{cases} a_{k,i} & \text{if } a_{k,i} > 0, \\ 0 & \text{if } a_{k,i} \leq 0 \end{cases} \quad \text{and} \quad a_{k,i}^- = \begin{cases} 0 & \text{if } a_{k,i} \geq 0, \\ -a_{k,i} & \text{if } a_{k,i} < 0. \end{cases} \quad (3.7)$$

Then, find the maximum value for each filter

$$m_k^\pm = \max_i |a_{k,i}^\pm|, \quad (3.8)$$

where the \pm symbol represents either sign. Finally, define the scaled filters

$$\mathbf{b}_k^\pm = \mathbf{a}_k^\pm / m_k^\pm. \quad (3.9)$$

The scaled filters \mathbf{b}_k^\pm have values on a scale of zero to one and can be implemented on the SLM.

When the spectrum \mathbf{s}_j is illuminating the SLM, the power arriving at the detector can be expressed as

$$P = \mathbf{s}_j^T \mathbf{r} + P_{j,\text{bg}}. \quad (3.10)$$

The vector \mathbf{r} represents the SLM reflectance, normalized to a scale from zero to one. The SLM response is controlled by changing \mathbf{r} . When the SLM reflectance for a given wavelength is at a maximum, the corresponding element of \mathbf{r} is unity; when the SLM reflectance for a given wavelength is at a minimum, the corresponding element of \mathbf{r} is zero. When all elements of \mathbf{r} are zero, the power measured is the background power $P_{j,\text{bg}}$. (Each source has a different background power, hence the j subscript.)

To compute the correlation coefficient $c_{j,k}$ for a source \mathbf{s}_j and a filter \mathbf{a}_k , a series of three measurements must be made. First, the background power $P_{j,\text{bg}}$ must be measured, by setting

$\mathbf{r} = 0$. Next, \mathbf{r} must be set to the normalized filters \mathbf{b}_k^+ and \mathbf{b}_k^- , producing measurements $P_{j,k}^+$ and $P_{j,k}^-$, respectively. Finally, the correlation coefficient can be computed by subtracting the background, scaling the results, and subtracting the negative part from the positive part, as follows:

$$c_{j,k}^+ = (P_{j,k}^+ - P_{j,\text{bg}})m_k^+, \quad (3.11)$$

$$c_{j,k}^- = (P_{j,k}^- - P_{j,\text{bg}})m_k^-, \quad (3.12)$$

$$c_{j,k} = c_{j,k}^+ - c_{j,k}^-. \quad (3.13)$$

It is straightforward to show that this $c_{j,k}$ is equal to the one defined in Eq. (3.1). Thus, Eq. (3.13) shows that the spectral correlation of \mathbf{s}_j with \mathbf{a}_k can be achieved with three measurements and a simple calculation.

The values of \mathbf{s}_j are recorded using a particular mode of operation that I have termed a *spectral scan*. A spectral scan of the source with index j works as follows: First, \mathbf{r} is set to $(1 \ 0 \ 0 \ \dots)$, and the first element of \mathbf{s}_j is set to the recorded power (minus the background power $P_{j,\text{bg}}$). Then, \mathbf{r} is set to $(0 \ 1 \ 0 \ \dots)$, and the next element of \mathbf{s}_j is recorded. This continues until all the elements of \mathbf{s}_j have been recorded. The elements of \mathbf{s}_j are essentially the spectral powers of source j , as recorded by the experimental apparatus. The spectral vectors \mathbf{s}_j must be measured using the MOC before the spectral matched filters \mathbf{a}_k can be computed. This can be seen as a form of training data, and once the \mathbf{s}_j are

measured and the \mathbf{a}_k have been computed, spectrum recognition measurements can be made repeatedly and rapidly.

The source intensity can fluctuate, and the fluctuations must be accounted for in the measurement process. If the fluctuations are rapid, the source itself must be monitored. If the fluctuations are slow, an alternative procedure can normalize the data. When the spectral powers \mathbf{s}_j are first measured, the background powers $P_{j,\text{bg}}$ are recorded and stored with the data. When the spectrum recognition powers $P_{j,k}^\pm$ are taken, the background powers are again measured; call them $P'_{j,\text{bg}}$. After these background powers are subtracted from $P_{j,k}^\pm$, the residues are then scaled by $P_{j,\text{bg}}/P'_{j,\text{bg}}$; equivalently, in Eqs. (3.11) and (3.12), the scaling constants m_k^\pm are replaced by $m_k^\pm(P_{j,\text{bg}}/P'_{j,\text{bg}})$. A similar normalization procedure can be used during the measurement of \mathbf{s}_j , by measuring the background power occasionally and fitting the change in background over time to a low-order polynomial or other slowly changing function, then normalizing \mathbf{s}_j by the fitted background function.

3.1.3 MOCs vs. traditional spectrometers

Multivariate optical computing offers a better signal-to-noise ratio (SNR) than traditional spectrometers; this is true for both scanning and non-scanning spectrometers. Scanning spectrometers, including the Czerny–Turner design and Fabry–Pérot interferometers, consist of some sort of monochromator followed by a single detector. The monochromator scans over the spectrum, and as it scans, the spectrum is recorded by the detector. The monochromator blocks most of the input light, allowing only a narrow wavelength range to pass. This

results in a low SNR. Non-scanning spectrometers disperse light over an array of detectors and read out the entire spectrum simultaneously, so they are much faster than scanning spectrometers. But each detector element still only sees a reduced amount of light, meaning the SNR is the same as for a scanning spectrometer of the same spectral resolution.

In contrast, a MOC combines multiple parts of the spectrum and directs them toward a single detector. Much more of the input light reaches the detector with each measurement, and the higher light levels lead to a higher SNR. This is sometimes known as the Fellgett advantage or multiplex advantage [168]. The increased SNR can mean more precise measurements as well as reduced integration times. Only one or a few readings are taken to identify a sample, so a MOC can be faster even than a non-scanning spectrometer, due to reduced integration times. If signal strengths are particularly low, as with Raman spectroscopy and certain other methods, a MOC may be necessary simply to have a high enough SNR to take a reliable measurement.

The choice of using a traditional spectrometer or a MOC to build a system for molecular identification depends on several tradeoffs. MOCs are probably best suited for repeated identification problems in which a few known substances are distinguished from each other. MOCs are optimized for a specific molecular identification task, and they perform that task well but need to be reconfigured or replaced in order to be used for a different task. Spectrometers are more general and more flexible (though a MOC using a reconfigurable element such as a spatial light modulator can be nearly as flexible).

All MOCs require some initial training data, either from an existing spectroscopic database or from taking a series of training measurements with the MOC itself. It wouldn't make sense to obtain all this training data only to perform a single measurement. But once the training data is in hand, it may be simpler to build a molecular identification system based on a MOC than one based on a traditional spectrometer. And a MOC can't identify new substances that it hasn't been designed to recognize—essentially no automated system can do this, beyond perhaps indicating the dissimilarity between a sample spectrum and a set of known spectra. In a philosophical sense, this is indicative of one of the main features (and drawbacks) of molecular identification: when one is searching for something, one must know not only what one is looking for but also what other things may be present that one does not wish to find.

It is not clear whether a MOC or a traditional spectrometer is the best tool for distinguishing among millions of different chemicals at the same time; a MOC might still provide an SNR improvement but would require a unique measurement for each possible chemical. There may well be some fundamental limit on the number of different spectra that can be distinguished for a given SNR or for a given signal strength, but the discovery of such a limit is not pursued here.

3.2 Experimental setup

An experiment diagram is shown in Fig. 3.1. A white light source is coupled into a vertical slit, which is followed by a polarizer. The white light passes through an optional absorption filter, and light passing through the filter (or the unfiltered white light) constitutes the input spectrum. The optical system projects the input spectrum onto the columns of a reflective phase-only SLM; for example, a spectrum consisting of two discrete spectral lines would illuminate two columns of the SLM. Each of the SLM columns imparts a tilt on the reflected spectrum, directing the reflected light downward towards a detector. The tilt is implemented as a blazed phase grating, and the efficiency of the blazed grating controls the amount of light reaching the detector, on a normalized scale from 0 (minimum efficiency) to 1 (maximum efficiency). The reflectances of the SLM columns are represented as the normalized vector \mathbf{r} and are used to implement the coefficients of the scaled matched filters (the \mathbf{b}_k^\pm). The optical system recombines all of the deflected light into a single beam arriving at the detector. The result is an optically computed correlation between the input spectrum and the matched filter on the SLM.

First, the SLM response is calibrated according to a procedure described in Appendix A. Next, spectral scans are taken of the white light source with no filter and with two doped glass filters, one doped with didymium (a combination of neodymium and praseodymium) and one with holmium. These three spectral scans are the \mathbf{s}_j and are used to generate the scaled filters \mathbf{b}_k^\pm that perform the spectrum recognition task. The \mathbf{b}_k^\pm are each put on the SLM in turn, and power measurements are taken as the source spectrum is changed (no

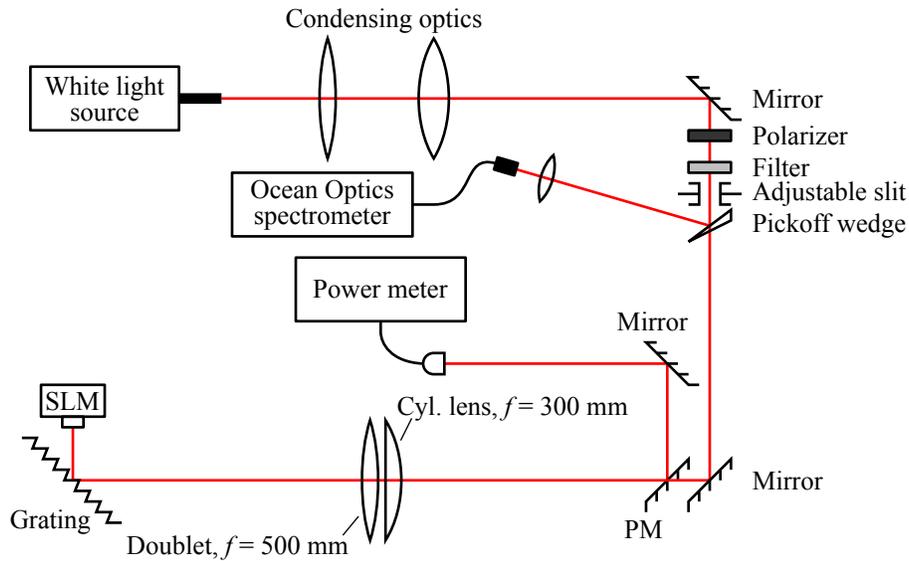


Figure 3.1. Multivariate optical computer (MOC) experiment diagram. PM: pickoff mirror, positioned below the input beam. The path length from the slit to the SLM matches the path length from the SLM to the detector. The polarizer is oriented vertically.

filter, didymium filter, holmium filter). These power measurements are used to compute the spectral correlation coefficients $c_{j,k}$.

3.2.1 Source and filters

The source is a Fostek microscope white light source with a fiber-coupled output. A stabilized 150 W power supply drives an EKE-type halogen bulb, and the source is operated at about 20% of its maximum output power. The source is allowed to warm up for 30 to 60 minutes before experimentation begins. The source spectra are created by two rare-earth-doped glass filters, one doped with didymium (a combination of neodymium and praseodymium) and one doped with holmium. The didymium filter is an FGB67 filter (Thorlabs, Inc.), 2 mm thick, 25 mm in diameter, made from Schott BG36 glass. The

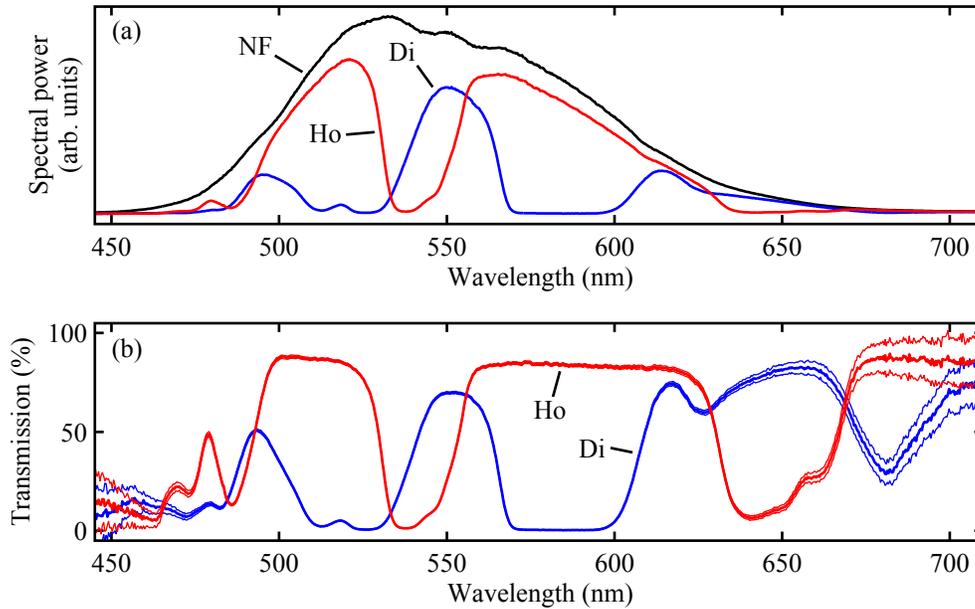


Figure 3.2. Spectra of the white light source and filters. (a) White light spectrum with no filter (NF), didymium filter (Di), and holmium filter (Ho), as measured by the Ocean Optics USB2000 spectrometer. Standard deviations are too small to be discernible and are not shown. (b) Transmission spectra of the didymium (Di) and holmium (Ho) filters, computed from the data in (a). Central lines show averages across multiple recordings, and thin outer lines show averages plus and minus one standard deviation.

holmium filter is an NT66-918 filter (Edmund Optics, Inc.), 3 mm thick, 25 mm in diameter, made from Schott S-8047 glass. The spectra of these two filters include some overlapping and some non-overlapping features across the visible spectrum. The white light source itself, with no absorption filter, is used as a third input spectrum. All three spectra are shown in Fig. 3.2(a), and the transmission spectra of the two filters are shown in Fig. 3.2(b). The two filters are mounted on a rotating filter wheel, allowing easy selection and consistent positioning. (An empty position on the filter wheel is used for the no-filter case.)

The input light is collected and directed through the filter into an adjustable vertical slit. A pickoff wedge directs a small portion of the input light towards a USB2000 commercial spectrometer (Ocean Optics, Inc.). The remainder of the optical system images the slit onto the SLM and re-images the reflected light from the SLM onto the detector. Because of this, the width of the input slit partially determines the spectral resolution of the system, and a narrower slit means a narrower slit image on the SLM and finer spectral control. However, a narrower input slit also means that less light is allowed into the system. The slit does not have a gauge to indicate the size of its opening, but it is about 2.5 mm wide when fully open, and during the experiment it is closed to allow only about 5% to 10% of the light to pass.

The entire experiment is enclosed in a custom-crafted black foam-core box, and additional sub-enclosures help further reduce stray light within the experiment. Black felt skirts and shrouds protect edges and seams both inside and outside the enclosure. With the enclosure and shrouds in place, turning the room lights on or off produces no change in the reading of the detector, which has a sensitivity of about 0.6 pW. The first mirror, filter, slit, pickoff wedge, and second mirror are mounted on a dovetail rail for collinearity and experimental ease. The pickoff wedge causes the beam to deflect slightly, but slight adjustments to the other elements compensate for this.

3.2.2 Broadband $2F-4F$ system

In principle, the input slit is imaged through a grating onto the SLM, and light reflected from the SLM is re-imaged through the grating onto the slit. The SLM imposes a vertical tilt on

some columns (wavelength bins), and at those wavelengths, the image appears below the slit. A pickoff mirror located below the input beam allows the lower image to be directed to a detector. The path length from the input slit to the SLM is equal to the path length from the SLM to the detector.

However, the SLM is a phase-only device. In order to impart wavefront tilt, it must be in the Fourier plane of a canonical $4F$ optical processor. To make this a possibility, the system is anamorphic, meaning it has different focal lengths in the x and y directions. (Here, x refers to the horizontal axis that is transverse to the direction of light propagation, y is vertical, meaning normal to the optical table, and z is the direction of propagation.) In the x - z (horizontal) plane, the system is an imaging system, with the input slit located in the object plane, the SLM located in the intermediate image plane, and the detector in the final image plane. However, in the y - z (sagittal) plane, the system's focal length is longer, and the SLM is located in the Fourier plane of the input slit. The vertical tilt imparted by the SLM becomes a vertical deflection in the detector plane. (The pickoff mirror in Fig. 3.1 is not exactly in the detector plane, but it is positioned where the downward-deflected image has separated from the undeflected image and where it does not obstruct the input beam.) This system is a traditional imaging system in one axis and a $4F$ system in the other axis, and I refer to it as a $2F$ - $4F$ system.

The $2F$ - $4F$ imaging system consists of a 50.8 mm spherical achromatic doublet lens with a 500 mm focal length and a 60×62 mm cylindrical lens with a 300 mm focal length. The cylindrical lens is oriented to focus in the x - z plane (i.e., the cylindrical axis is along the

y axis). The lenses are positioned to demagnify the slit image slightly (enhancing spectral resolution) and adjusted to reduce image aberrations, and they are mounted on a dovetail rail for easy adjustment and for collinearity. The effective focal lengths are 500 mm in the y - z plane and about 250 mm in the x - z plane.

A 50 mm square, 300 line/mm blazed (ruled) grating disperses the spectrum spatially into separate slit images on the SLM, ensuring different wavelengths land on different SLM columns. The operating bandwidth of the system is determined by the size of the SLM, the pitch of the grating, and the distance between the two, and the spectral resolution of the system is limited by the slit width, the cylindrical lens line width (the cylindrical analog of spot size), and the SLM pixel pitch. The SLM is positioned near enough to the grating that the SLM covers the visible range. The flat grating in the imaging path causes a small amount of field curvature; to correct this, one can use a flat-field grating (which is a grating with a curved surface) or a field-flattening cylindrical lens, but in this experiment the curvature is small enough to be neglected.

3.2.3 SLM operation

The SLM is a computer-controlled, 1024×768 -pixel, reflective, phase-only liquid-crystal-on-silicon (LCOS) SLM from Cambridge Correlators, Inc., and it is controlled using a computer's VGA port. This SLM has a phase range of about 0.8π radians at 633 nm, and the SLM pixel pitch is 9 μm in both x and y . Its columns are used to tilt the wavefront downward, with a different (controllable) efficiency at each wavelength. The wavefront tilt

is achieved using a blazed grating, which acts like a phase ramp. For maximum efficiency, the period of the blazed phase grating is chosen to be an integer number of pixels. In this experiment, a 2-pixel blazed grating is used, because it was found experimentally to have the highest efficiency and to give the best separation from the front-surface reflection of the SLM. (A 2-pixel blazed grating is functionally equivalent to a 2-pixel sinusoidal grating, due to the Nyquist limit.)

In the normal mode of operation, a scaled matched filter is put on the SLM (i.e., $\mathbf{r} = \mathbf{b}_k^\pm$). However, a spectral scan can be taken by setting $\mathbf{r} = (1 \ 0 \ 0 \ \dots)$, taking a reading from the detector, then setting $\mathbf{r} = (0 \ 1 \ 0 \ \dots)$, taking a reading, and so forth. In this spectral scan mode, the MOC operates as a scanning spectrometer, without the advantages of a MOC but with the ability to record an entire spectrum. The spectral scan mode is used to measure the s_j , which are then used to compute the \mathbf{a}_k .

In order to speed data acquisition in the spectral scan mode, each group of 16 columns of SLM pixels is manipulated as a single wavelength “bin,” giving 64 wavelength bins total. The entire experiment is operated using 64 wavelength bins.

3.2.4 Detection

A pickoff mirror located before the cylindrical lens and below the input beam redirects the deflected output beam toward the detector. The detector is placed such that the path length from the input slit to the SLM equals the path length from the SLM to the detector; this is necessary for an image to form. (The exact placement of the detector is adjusted

experimentally to reduce aberrations.) A Newport 2936-R power meter with a Newport 918D-SL-OD3R calibrated silicon sensor is used as the detector; this power meter/sensor combination is sensitive down to about 0.6 pW and takes readings about once or twice per second.

3.2.5 Control and processing

Virtually the entire experiment is automated. The filter wheel must be turned manually to select a new source filter, but otherwise, the experimental enclosure never needs to be opened once the experiment is aligned. Extensive software was written in the LabVIEW environment (National Instruments, Inc.) to control the SLM and acquire data from the power meter and USB2000 spectrometer. The LabVIEW software is robust and fault-tolerant, running for weeks at a time without interruption or failure. An advanced parallel architecture makes the software particularly responsive. A command file tells the LabVIEW software what images to display on the SLM and how many readings to take for each image. Software was written for MATLAB (The MathWorks, Inc.) to generate command files for the LabVIEW software and to analyze the resulting data. The LabVIEW software totals over 20 000 nodes, and the MATLAB software totals over 40 000 lines of code, all custom-written for this experiment.

3.2.6 Optical design tradeoffs

A number of design tradeoffs must be made for any system utilizing a grating and SLM. The tradeoffs primarily concern the physical size of the system, the operating bandwidth, the spectral resolution, and the amount of optical power available (which directly determines integration time).

The pixel pitch of the SLM determines the pitch of the vertical blazed grating on the SLM, which in turn determines the deflection angle of the output image. The output image cannot overlap the entrance slit (or the SLM front surface reflection, which is re-imaged onto the entrance slit). So once the deflection angle is determined, the focal length of the optical system is also determined.

The grating pitch and the distance between the grating and the SLM determine the total resolvable bandwidth of the system. Ideally, these parameters should be chosen so that the spectrum of interest exactly fills the SLM, in order to maximize the spectral resolution achievable.

The slit width is an important tradeoff between the total power reaching the detector and the spectral resolution of the system. A wider slit increases the power to the detector but also broadens the width of the image on the SLM columns, decreasing spectral resolution. Arranging the optics to demagnify the slit image at the SLM can improve spectral resolution somewhat. Optical aberrations can also reduce the spectral resolution by broadening the line spread function. The number of SLM pixel columns in a wavelength bin is also a tradeoff between spectral resolution and amount of light per column. One would like ideally to use

each column of SLM pixels independently, but optical power limitations may require the use of multi-pixel columns. The ultimate limit here is the dynamic range of the detector; the detector must be able to register the change in output power due to switching a single column on and off, particularly in spectral scan mode. The minimum detectable difference determines how wide the columns must be. Data acquisition time may also dictate using wider columns. The number of SLM pixel columns puts an ultimate limit on the number of resolvable spectral lines. In this experiment, however, the spectral resolution is limited by the slit width and optical aberrations; no improvement in resolution was evident when the SLM was divided into a larger number of wavelength bins, and the noise increased.

One of the most important parameters of the system is the detection limit of the detector, which is related to the signal to noise ratio (SNR) of the system. The vertical size of the entrance slit influences the total optical power that reaches the detector (assuming a filled slit; an under-filled slit is effectively the same as a smaller slit). Some additional sensitivity can be gained by increasing the averaging time or integration time of the power meter, though of course this increases the time required to make measurements. The clear aperture of the lenses, grating, and SLM can all limit the amount of light reaching the detector. The spectral radiance of the source should be as high as possible to maximize the SNR. One possibility is to use a high-radiance source to take the training data and then use a weaker source to perform the spectrum recognition task.

As mentioned in section 3.2.2, the flat grating in the imaging path induces field curvature. Because of the relatively long horizontal focal length of the system (250 mm) and the

relatively broad pitch of the grating (300 line/mm), the field curvature is not found to be significant in this experiment, and a field-flattening cylindrical lens inserted near the SLM is not observed to improve the image quality noticeably. However, in a system with a shorter focal length or a finer-pitched grating, field curvature may need to be corrected. The effect of field curvature would be to reduce the spectral resolution of the system due to blurring of the optical image on the outermost columns of the SLM.

3.2.7 Preliminary measurements

To test the concept of performing spectral correlations using an SLM, a simplified setup was used, shown in Fig. 3.3. The collimated beam from a tunable quasi-CW laser (a Newport Tsunami Ti:sapphire laser operating in the non-mode-locked regime) is expanded and illuminates a 50 mm square, 1800 line/mm holographic grating. The cylindrical lens focuses the diffracted light into a vertical line on the SLM. The SLM is in the focal plane, or Fourier plane, of the cylindrical lens. As the laser wavelength is scanned, the focused laser line sweeps horizontally across the SLM. The SLM deflects certain wavelengths using a blazed grating as before (but with a three-pixel blaze period in this case, oriented to deflect the beam upward instead of downward). The signal reflected from the SLM is re-collimated by the cylindrical lens and diffracted by the grating back into the optical path. A pickoff mirror, located above the input beam, directs the deflected light toward a detector, a Coherent Fieldmaster-GS power meter. Two irises and a shroud (not shown) reduce the stray light

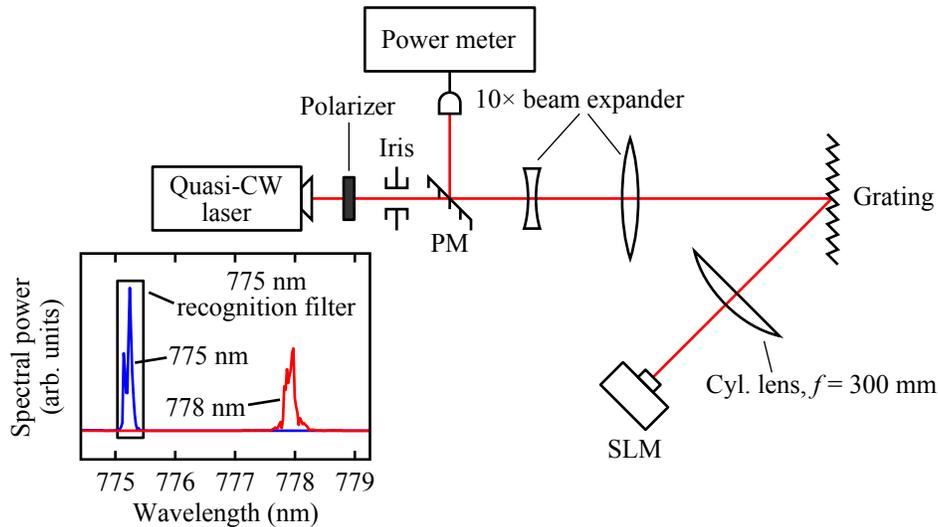


Figure 3.3. Preliminary MOC experiment diagram. PM: pickoff mirror, positioned above the input beam. The SLM is in the focal plane (Fourier plane) of the cylindrical lens. Inset: Quasi-CW laser spectra and filter for the preliminary experiment, with the laser tuned to approximately 775 nm and 778 nm. The rectangle marks the spectral region used to detect the 775 nm laser line; the matched filter has a value of one in this region and zero outside of it. The inverse filter is used to detect the 778 nm line.

reaching the detector. In the preliminary experiment, the SLM is a Boulder Nonlinear Systems XY Series nematic 512×512 -pixel SLM, with a pixel pitch of $15 \mu\text{m}$.

The simplified setup was used at two wavelengths, approximately 775 nm and 778 nm. The spectra of the two quasi-CW modes were scanned using the SLM, and the results are shown in the inset of Fig. 3.3. The instantaneous spectra fluctuated rapidly, presumably due to mode competition, but the spectra were found to be visually similar to those recorded by an Ocean Optics USB2000 spectrometer. Rather than generating the optimal filters \mathbf{a}_k as described in section 3.1, two simple binary filters were used, the first of which is shown in the inset of Fig. 3.3, and the second of which is the complement of the first (i.e., $\mathbf{a}_2 = 1 - \mathbf{a}_1$). The correlation results, shown in Table 3.1, are within 0.003 of the ideal.

Table 3.1. Preliminary experiment correlation results. Ideal results are the identity matrix.

Filter	775 nm	778 nm
\mathbf{a}_1	1.0000	0.0000
\mathbf{a}_2	0.0030	1.0000

3.3 Results and discussion

3.3.1 Spectrum recognition results

The main results of the spectrum recognition experiment are shown in Table 3.2. The spectrum recognition method performed adequately, with all values being within 0.07 of the ideal values of zero and one. The standard deviation of each value in Table 3.2 is less than 0.0025. The standard deviation of the correlation value is computed from the standard deviation of the power measurements using propagation of uncertainty principles, with the assumption that all deviations follow independent normal distributions. (The presence of negative values in Table 3.2 is due to the subtraction of the positive and negative parts of the correlation coefficients, and due to background subtraction; see Eqs. (3.10)–(3.13)). To avoid some calibration-related errors, the calculation of the \mathbf{a}_k and the \mathbf{b}_k^\pm excluded wavelength bins with too little spectral power, namely those bins with less than 8% of the power of the maximum bin. (An experiment performed with these bins included was substantially similar; the correlation values changed by less than 0.025, and maximum standard deviation increased to 0.0045.)

Table 3.2. Spectrum recognition results. Ideal results are the identity matrix.

Filter	White light	Didymium	Holmium
\mathbf{a}_1	1.0043	0.0020	0.0107
\mathbf{a}_2	0.0300	0.9591	0.0327
\mathbf{a}_3	-0.0225	-0.0089	0.9337

3.3.2 Spectral scans

Figure 3.4(a) shows the spectra of the white light source with no filter, with the didymium filter, and with the holmium filter, as recorded in spectral scans using the SLM. (These three spectra are the \mathbf{s}_j in section 3.1.) The spectral data from the Ocean Optics USB2000 spectrometer are shown in Fig. 3.4(b), for reference; this is the same plot as Fig. 3.2(a). The spectral resolution (instrument function) of the experiment is broader than that of the USB2000. The instrument function of the experiment is estimated to be 15 nm wide and asymmetric in shape (details below). Additionally, there are some subtle differences between the spectra; I believe these are due to optical system aberrations (especially those presumed to cause the asymmetric instrument function), and due to the presence of additional diffracted orders from the ruled grating (discussed in greater detail in section 3.3.4). These features do not greatly affect the spectrum recognition operation or the calculation of the doped glass filter transmission spectra, since they are accounted for by the calibration and matched filter generation procedures.

From the data in Fig. 3.4, the transmission spectra of the didymium and holmium filters can be calculated, and these transmission spectra are shown in Fig. 3.5(a). For reference,

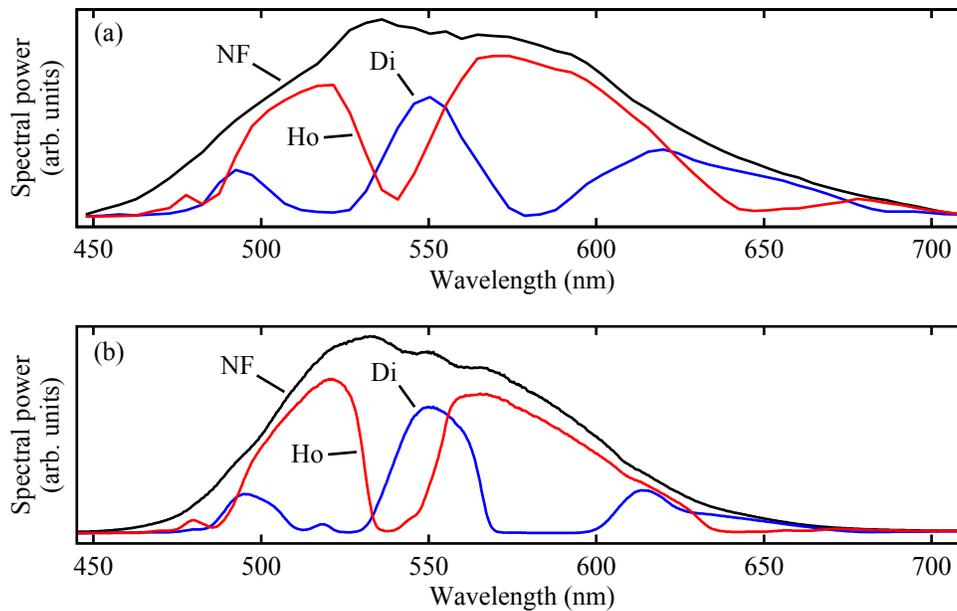


Figure 3.4. Spectral scans of the white light source and filters. (a) Spectral scan of the white light spectrum with no filter (NF), didymium filter (Di), and holmium filter (Ho), as measured by the experiment. These data constitute the s_j . Lines between data points are guides for the eye. (b) Data from the Ocean Optics USB2000 spectrometer (same data as in Fig. 3.2). Standard deviations in both plots are too small to be discernible and are not shown.

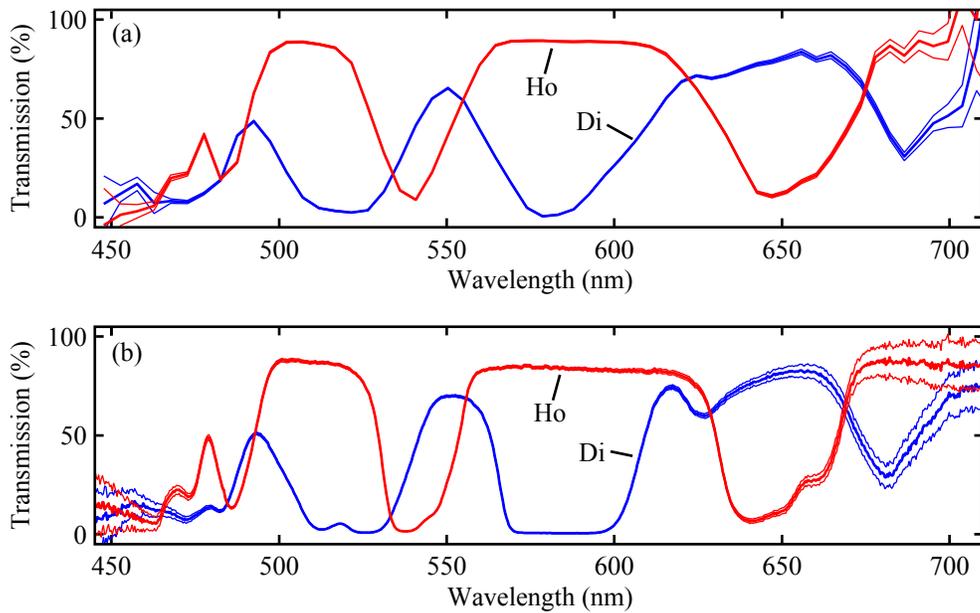


Figure 3.5. Transmission spectra of the didymium (Di) and holmium (Ho) filters, as measured by (a) the experiment and (b) the Ocean Optics USB2000 spectrometer (same data as in Fig. 3.2). Central lines show averages across multiple recordings, and thin outer lines show averages plus and minus one standard deviation.

the corresponding data from the USB2000 are shown in Fig. 3.5(b); this is the same plot as Fig. 3.2(b). The white light source has reduced power in the wings of the spectrum, and the resultant increase in the standard deviation can be seen in both plots. The thick central lines show the averages across multiple recordings, and the thin outer lines show one standard deviation above and below the averages. In both Fig. 3.4 and Fig. 3.5, a good qualitative match between the commercial spectrometer and the experimental system is evident.

The transmission spectra shown in Fig. 3.5 are used to calibrate the wavelength correspondence of the SLM pixels (see Appendix A). The transmission data are also used to estimate the spectral resolution (instrument function) of the system at about 15 nm, which is significantly broader than that of the USB2000 (specified as 1.5 nm by the manufacturer). The estimation was performed by numerically convolving the transmission spectrum from the USB2000 with a Gaussian line shape of varying widths; the 15 nm linewidth provided the most visually similar transmission spectrum. It is estimated that the system has a slightly finer resolution toward the blue end of the spectrum, possibly due to a slight misalignment between the grating and the SLM. Based on the appearance of the transmission curves, the instrument function may be asymmetric, which is likely caused by aberrations in the optical system. The linewidth of 15 nm is generally well-matched to spectral features of the rare-earth-doped filters. A brief experiment with the number of SLM pixels per wavelength bin showed that using smaller wavelength bins did not improve the spectral resolution, meaning that the resolution is determined primarily by the slit width and the optical aberrations.

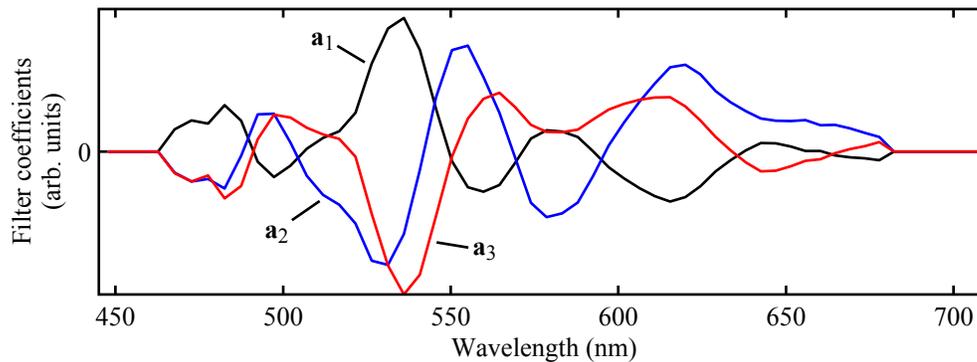


Figure 3.6. Matched filter spectra \mathbf{a}_k . \mathbf{a}_1 corresponds to no filter, \mathbf{a}_2 to the didymium filter, and \mathbf{a}_3 to the holmium filter. Lines between data points are guides for the eye.

The data in Fig. 3.4, after background subtraction, are used as the s_j to compute the matched filters \mathbf{a}_k . Figure 3.6 shows the three filters \mathbf{a}_k . Each matched filter generally emphasizes the wavelengths for which its associated spectrum is stronger than the other spectra. While the filters themselves are not mutually orthogonal, each filter correlates with only one of the three spectra. The filter coefficients for the outermost wavelengths, those for which the white light spectral power is less than 8% of the maximum, are set to zero; this was found to reduce the error slightly.

3.3.3 Correcting erroneous data

The raw power meter readings occasionally contain outliers, caused perhaps by dust in the imaging path, electrical noise, or other problems. These outliers typically are single readings that are much greater or much smaller than the other readings for the same SLM image. The LabVIEW software uses an averaging and outlier-removal procedure, as follows: Once the SLM image display has stabilized, the software takes a specified number of power

meter readings (usually 10). The readings are averaged to produce a single data point. Upon completing the of readings, the LabVIEW software checks the relative standard deviation of the readings to ensure that it is less than 2×10^{-4} ; if not, the reading with the largest deviation from the mean is discarded, and another reading is taken. This process is continued until the relative deviation falls below the threshold or a maximum number of readings has been taken (usually 50). (The threshold was chosen by empirical observation that data without outliers nearly always had a relative deviation below 2×10^{-4} .) The outlier-removal process works, though it slows down data acquisition.

The outlier-removal process helps correct an additional problem: sometimes the SLM switches slowly. After sending a new image to the SLM to display, sometimes the detected power follows an exponential curve with a time constant of several seconds, rather than switching immediately. The outlier-removal process sees this as a large standard deviation and continues taking data until the relative standard deviation is sufficiently low (or the maximum number of readings have been taken).

3.3.4 Experimental issues

Several experimental issues impacted the fidelity of the spectrum recognition operation. The primary issues are optical power inefficiencies, ghosting from the blazed diffraction grating, and possibly optical aberrations.

The white light source, which has a broad emission angle (about 70° full angle), is extremely poorly matched to the MOC's optical system, which has an acceptance angle of

less than 15° and a mostly closed slit at the entrance. About 0.75 W of white light leaves the source's output fiber bundle, but no more than a few microwatts pass through the slit. Further, many of the optical components are lossy. Neither the lenses nor the pickoff wedge are coated, so 4% reflections at each surface lead to a 33% loss (considering that the light goes through the lenses twice). The grating has a loss of about 40%, according to the manufacturer's data. The SLM's light efficiency is reduced by the imperfect fill fraction of the pixels (the less active area each pixel has, the lower the diffraction efficiency of the SLM), compounded by the inefficiency of the short-period blazed grating and any phase calibration errors. (The SLM also has some optically damaged spots on its surface, occupying no more than about 10% of the area of the array, which further reduce its efficiency slightly.) In addition, the back reflections from each optical surface, the extra diffraction orders from the grating, and the stray light inside the enclosure (due to the large amount of source light that is not coupled into the optical system) contribute to a comparably high level of background light at the detector.

Blazed (ruled) gratings, such as the one used in this experiment, have higher efficiencies than holographic gratings, but they produce multiple diffraction orders, leading to an effect called ghosting. When the input light reaches the detector, the strongest diffraction order is directed to the SLM, but part of the extra diffracted orders are reflected back into the experiment and reach the detector, acting as background light. When the light returning from the SLM reaches the diffraction grating, it is again diffracted into multiple orders, and some of these couple into the optical system and land on the detector, but these stray orders act to

increase the signal strength, since their intensity is controlled by the SLM. But both of these effects should be accounted for in the way the SLM is used; any increase in background is subtracted, and any increase in signal strength is accounted for in normalization. However, there is a possibility that multiple diffracted input orders reach the SLM. This would create a kind of crosstalk, in which changing one wavelength bin on the SLM also affects light at a second wavelength due to the extra diffracted orders. The SLM did not appear visually to be receiving more than one diffracted order, but an experimental verification would have been quite complicated. A holographic grating would reduce this stray light, but at a cost of reduced optical efficiency, since holographic gratings are typically less efficient than blazed gratings. The exact geometry of the stray diffraction orders may help explain the small spectral irregularities visible in Fig. 3.4(a). (These irregularities may also be partly due to aberrations in the optical system, which can contribute to spectral crosstalk and reduced spectral resolution.)

Some improvements could be made by using equipment that is better tailored to the experiment. For instance, using antireflection-coated optics and custom lenses could improve light efficiency (and therefore spectral resolution, due to the tradeoffs mentioned in section 3.2.6). The input and output image paths are common in this experiment, but one can conceive of a system in which they are separated: upon reflection from the SLM, the light would enter a second, complementary optical system leading to the detectors. Such a system would be more complex but might have fewer aberrations, less spectral crosstalk, and possibly a lower background level due to better control of stray light. A custom SLM

could improve both light efficiency and spectral resolution [179]; however, this makes the system significantly more expensive, and it can complicate the SLM control.

An increase in optical power would enable the use of a faster power meter, one with a kilohertz-level data rate. The power meter used in the experiment, which was chosen for its high sensitivity, takes about half a second to a full second to record a single reading. The data shown in Table 3.2 require about 40 minutes to obtain, with a minimum of 100 readings per data point (reducing this to a minimum of 10 readings per data point would cut the time by about a factor of 10). The calibration data required for the experiment take an additional 27 hours to record, not counting time spent adjusting the experiment.

3.4 Multiple-output multivariate optical computation

As mentioned in the introduction to chapter 3, with a slight change to the experimental layout, multiple spectral matched filters can be implemented simultaneously on the SLM. An SLM image with multiple blaze periods produces multiple output images; for instance, by creating an SLM image with different blaze periods in two areas of the image, and adding an additional pickoff mirror and an additional detector, two half-correlations (the $c_{j,k}^{\pm}$) can be computed simultaneously. This is a speed advantage when performing a spectrum recognition task. In particular, a positive and a negative half-filter (\mathbf{b}_k^+ and \mathbf{b}_k^-) can be implemented simultaneously, allowing the correlation coefficient for a single matched filter (\mathbf{a}_k) to be computed in real time. In one possible setting, this would allow the spectral

correlation to be computed continuously as samples are brought in front of the light source for interrogation. An additional method could also allow multiple \mathbf{a}_k to be implemented simultaneously. Both methods are studied experimentally below.

A two-output experiment is shown in Fig. 3.7. Note that the experimental layout is almost identical to Fig. 3.1, with the final mirror replaced by a pickoff mirror, and another mirror and detector added after that. (The Newport 2936-R optical power meter supports a second Newport 918D-SL-OD3R sensor, allowing the two detectors to be read more or less simultaneously.) To implement the filters on the SLM, the vertical blazed grating periods must be chosen for minimum overlap of the output image. Empirically, 2-pixel and 5-pixel gratings are found to have good optical power and minimum overlap. Since the positive and negative scaled filters \mathbf{b}_k^\pm don't overlap, wavelength bins with positive coefficients are implemented with a 2-pixel blaze, and negative coefficients are implemented with a 5-pixel blaze. A second implementation is pursued in which the SLM is divided into two vertical zones, with positive coefficients implemented on the top zone and negative coefficients on the bottom zone. This configuration has the advantage of being able to implement two arbitrary filters and is extensible to a higher number of outputs. (Due to optical aberrations, the vertical zones are not made strictly rectangular; instead, the boundary between the zones is chosen individually for each wavelength bin to ensure the correct power distribution reaches each detector.)

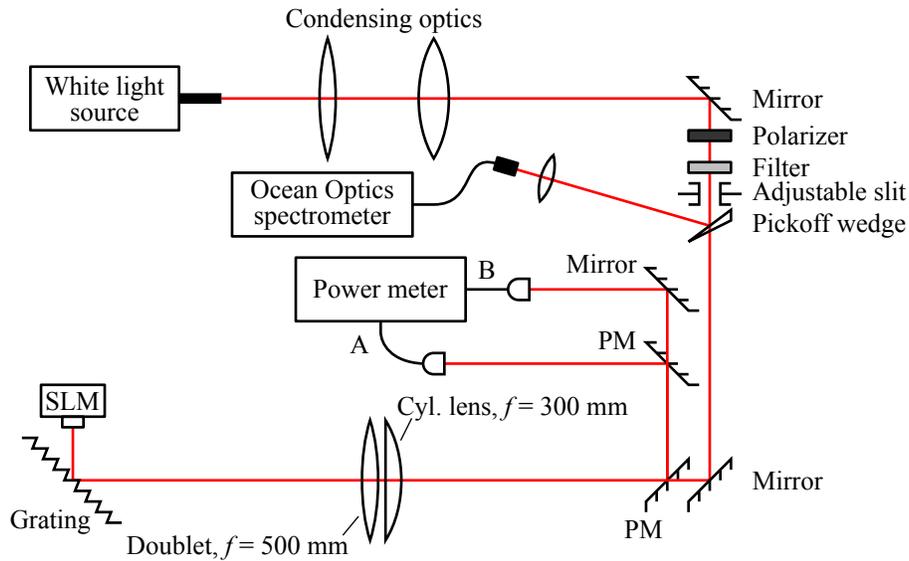


Figure 3.7. Multiple-output MOC experiment diagram. PM: pickoff mirror, positioned below the beam. As before, the path length from the slit to the SLM matches the path length from the SLM to the detector, and the polarizer is oriented vertically.

3.4.1 Results

The results of the one-zone, two-output spectral recognition experiment are listed in Table 3.3. The performance was nearly as good as the one-output experiment, with all but one value within 0.075 of the ideal values and the outlier being 0.164 from the ideal. The slight degradation in performance is probably attributable to calibration errors (see Appendix A). The standard deviation of each value in Table 3.3 is less than 0.0047, again computed using propagation of uncertainty principles from the standard deviation of the power measurements. The wavelength bins with less than 8% of the maximum spectral power were again excluded. (If these bins are included, the correlation values change by less than 0.089 and the maximum standard deviation increases to 0.015.)

Table 3.3. Multiple-output spectrum recognition results with a single SLM zone. Ideal results are the identity matrix.

Filter	White light	Didymium	Holmium
\mathbf{a}_1	0.9929	-0.0065	0.0114
\mathbf{a}_2	0.1640	0.9392	0.0650
\mathbf{a}_3	-0.0540	-0.0914	0.9257

Table 3.4. Multiple-output spectrum recognition results with two SLM zones. Ideal results are the identity matrix.

Filter	White light	Didymium	Holmium
\mathbf{a}_1	1.0661	-0.0141	0.0389
\mathbf{a}_2	0.0297	0.9066	0.0628
\mathbf{a}_3	-0.1575	-0.1304	0.8903

The results of the two-zone output spectral recognition experiment are listed in Table 3.4. This experiment performed as well as the previous one. The standard deviation of each value in the table is less than 0.0060. (Inclusion of low-power wavelengths changes correlation values by up to 0.028; the maximum standard deviation does not increase.) Note that this method is scalable to more than two spectral matched filters, meaning more than one chemical can be recognized simultaneously.

3.4.2 Experimental issues

Ghosting effects from the blazed (ruled) grating and aberrations in the optical system may be causing small spectral anomalies, as mentioned in section 3.3.4. Notably, the vertical distribution of optical power reaching the SLM is slightly different for each wavelength bin,

indicating the presence of optical aberrations. In the two-zone experiment, the SLM zones were made irregular to accommodate this fact. Further, the vertical power distributions seen by each detector are slightly different, such that in the two-zone experiment, each detector receives just over half as much the power that it receives from the single-zone experiment. (Naively, one would expect half or just under half the power to reach each detector.) SLM calibration issues also likely play a role (discussed in Appendix A). As noted in section 3.3.4, the power meter is slow, and a faster power meter would allow further improvements; in the current configuration, it takes about 7 minutes to record the data shown in Tables 3.3 and 3.4 (minimum of 10 readings per data point), with an additional 48 hours to take the calibration data needed (beyond the 27 hours needed for the calibration data in section 3.3, and not including time spent adjusting the experiment).

3.4.3 Limit on number of simultaneous matched filters

One may ask how many filters may be implemented simultaneously on the SLM. Obviously, optical power is one limiting factor. Due to the experiment geometry, crosstalk between the filters is another limit.

This experiment used two SLM zones, to implement two spectrum recognition filters simultaneously. A higher number of filters allows more spectral processing to happen simultaneously. The obvious tradeoff is with optical power; as a fixed amount of optical power is divided across an increasing number of SLM zones, the signal-to-noise ratio of each filter is reduced, or the required integration time is increased. For instance, if adding a

second filter requires doubling the integration time, it is not clear that using multiple outputs offers an advantage. An entire filter (positive and negative values) can be implemented without reconfiguring the SLM, but there is the possibility of crosstalk degrading the results slightly.

Preventing crosstalk between the different filters is important and becomes an increasing problem with an increasing number of filters, due to the experiment geometry. The vertical size of the input slit determines the vertical size of the output images, and the number of possible simultaneous filters is limited by the size of the output images and the need for the images not to overlap. All of the output images must fit within the range of the main SLM diffraction order, which is determined by the SLM pixel pitch. More filters can be fit in by reducing the vertical size of the input slit, but this may be at the cost of reduced spectral power. If the source light can be collimated, a non-imaging experimental configuration such as that in Fig. 3.3 may be more appropriate and could lead to a larger number of filters.

In my opinion, if the experimental issues described above can be resolved, the number of simultaneous filters that can reasonably be implemented on the SLM in this manner is at least five, and possibly as high as 20 with very careful design and alignment. Even implementing two filters simultaneously, as demonstrated in this section, is a two-fold improvement over previous work, allowing for much simpler processing and potentially faster acquisition of spectrum recognition results.

3.5 Conclusion

Spectroscopy is a proven way to identify substances, and multivariate optical computation is an experimentally successful approach to speeding the identification task. A multivariate optical computer (MOC) has been implemented using a reflective liquid-crystal-on-silicon spatial light modulator (LCOS SLM), using a broadband white light source and a unique optical configuration. The SLM is able to implement multiple filters simultaneously with arbitrary coefficients, an advantage over other MOCs that offers faster processing. A commercial, off-the-shelf SLM was used, unlike other approaches that have required custom-built SLMs. Several experimental issues reduced the fidelity of the results, but these issues did not prevent successful application of the method, and this MOC shows great promise for allowing rapid identification of chemicals.

Chapter 4

Conclusions and Future Work

Slow light is a fascinating topic. The science behind slow and fast light has been investigated by many researchers, and attention is turning towards finding applications for slow light techniques. Two potential applications for slow light have been presented here, as well as some critical system components needed to implement those applications. Below, I summarize the findings and discuss possible future directions for research.

4.1 Phase control of a slow-light laser radar

Chapter 2 describes a phase control system that maintains phase lock in a slow-light phased-array laser radar. The RMS residual phase error is about $\pi/5$ radians (1/10 wave), and the corresponding Strehl ratio is about 0.8. The phase control system performed adequately and enabled the demonstration of slow light for true-time delay in laser radar.

The statistical model of the phase error presented in chapter 2 gives a good indication of the impact the snapback process has on the phase control system's RMS phase error.

However, the model is based on an assumption that the phase error has a normal distribution. While this makes the analysis much easier, it is not a realistic model, as mentioned in section 2.1.2. For instance, a small vibration of a single optical component (such as a mirror) at a single frequency should produce an approximately sinusoidal phase fluctuation, which gives a very different statistical distribution than the one in the model. Similarly, the snapback events are not necessarily Poissonian-distributed, and the occurrence of snapback events is influenced by the characteristics of the phase noise (since snapback events occur when the phase noise exceeds a certain threshold). A more realistic statistical model of the phase error would allow more precise analysis of this class of phase control problems and could provide additional insight into the best approach to phase control. Additionally, the precise effect of the feedback circuit and electro-optic modulator (EOM) on the optical phase was not included in the model; this is a possible further refinement.

The phase control system uses a phase-sensitive heterodyne detector, a loop filter (the proportional-integral controller), and an EOM, which is essentially a voltage-controlled optical phase shifter. In the purely electronic domain, a phase-locked loop (PLL) is a well-known circuit that is used to keep two signals in phase lock with each other. A PLL consists of a phase-sensitive detector, a loop filter, and a voltage-controlled oscillator (VCO). The key difference is a VCO changes its output frequency—not its phase—in response to an applied voltage. The frequency shift allows the output signal to accumulate phase, or “catch up” to the reference signal. As the output signal approaches the reference signal, the phase error diminishes and the VCO output shifts closer to the reference frequency. Because the

phase shifting is accomplished by a frequency difference, the PLL can in principle correct for an unlimited amount of phase error, if one is willing to wait long enough for the output signal's phase to catch up to the reference signal's phase. In the optical domain, a VCO could be emulated using an acousto-optic modulator (AOM), which shifts the optical signal's frequency by the amount of an applied radio frequency. The AOM would be driven by a VCO, or another form of tunable driver. Using an AOM instead of an EOM would create an optical PLL (or OPLL), which would have no need of a snapback circuit. Additional AOMs were not available for the SLIDAR experiment, but it would be interesting to implement this kind of OPLL. Similar approaches have already been demonstrated with diode lasers acting as the frequency-shifting elements [158, 159]. However, an AOM requires a good deal of radio frequency (RF) power, and it typically is used with a specialized RF driver; tunable RF drivers for AOMs can be expensive.

As mentioned, the phase control problem for laser radar is very similar to the one for high-power coherent beam combining, and the demonstrated phase control system could be used for a coherent beam combining system. Coherent beam combining systems generally use much shorter lengths of fiber (tens of meters), but the phase control is complicated by the presence of nonlinear optical effects in the fibers.

The SLIDAR phase control system uses analog electronics for processing, which makes the system simple and inexpensive. However, there is great power in digital signal processing systems, and it is possible that a more effective feedback algorithm could be found, especially if a more complex statistical model of the phase noise were developed as described above.

The snapback algorithm has a very natural digital implementation: simply compare the phase shift desired from the EOM with the snapback thresholds, and subtract $2n\pi$ radians from the phase if needed.

4.2 Multivariate optical computing for spectrum recognition

The multivariate optical computer described in chapter 3 can distinguish two rare-earth-doped glass filters from each other and from an unfiltered white light source with reasonably high fidelity. The resulting improvement in signal-to-noise ratio, known as the Fellgett advantage or multiplex advantage, allows accurate recognition of chemicals even from weak spectra. Multiple spectral matched filters can be implemented simultaneously using the methods described in chapter 3, which can speed up the molecular identification process when the spectrum is sufficiently intense.

There is some theoretical work yet to be done in the area of multivariate optical computing. As mentioned at the end of section 3.1.3, an upper limit can probably be derived on the maximum number of distinguishable spectra N_S for a given signal-to-noise ratio or signal strength. Other quantities for which limits would be helpful include the spectral resolution needed to distinguish two similar spectra, the expected spectral similarity between different classes of chemicals, and the number of different reflectance levels needed to implement

multivariate optical computing efficiently (Wilcox and coworkers argue that this last number is two [182]).

A number of experimental issues were identified in chapter 3. More work is needed to find an accurate broadband calibration approach for the spatial light modulator (SLM). An improved optical system with reduced aberrations could offer improvements in the spectral resolution of the system; a holographic grating and antireflection-coated optics would reduce the background light level. The imaging configuration described in chapter 3 is required because of the incoherent, extended source (non-point source) used for white light. A collimated source could allow a more efficient configuration, similar to the one used in Fig. 3.3 to take preliminary measurements. A coherent supercontinuum source is one example of a coherent broadband light source.

The multiple-output MOC with multiple SLM zones described in section 3.4 can be extended to implement more than two spectral matched filters at once, meaning two or more chemical species may be recognized simultaneously. It would be intriguing to explore the maximum number of chemicals that can be recognized simultaneously, as well as what limits this number, both theoretically and experimentally. Crosstalk between the outputs of the matched filters would likely be an important consideration.

One exciting area of future work is the investigation of multivariate optical computing for recognition of chemicals of practical interest. For instance, hydrogen cyanide (HCN) is a highly useful industrial chemical, but it is also extremely poisonous and can be made into a chemical weapon. The absorption spectrum of hydrogen cyanide includes a number

of absorption lines near 1550 nm, an experimentally convenient wavelength [190]. One could build a multivariate optical computer to distinguish between, for example, HCN and acetylene, another chemical with an interesting absorption spectrum near 1550 nm [191]. One could even build a small library of similar chemical spectra in this wavelength range and use a multivariate optical computer to distinguish them.

Another intriguing possibility is applying the principles of multivariate optical computing to terahertz spectroscopy. It is well-known that certain high explosives include strong resonances in the terahertz domain, and detecting traces of such explosives is one use for terahertz technology [192, 193]. A multivariate optical computer operating in this new regime is likely to offer the same signal-to-noise ratio improvements, potentially making this a highly sensitive method of detecting explosives. Because terahertz wavelengths are so much longer than optical wavelengths, the spatial light modulator (SLM) pixels must be made much larger than the pixels in an SLM intended for use at visible or near-infrared wavelengths; also, the SLM will likely be made from something other than a liquid-crystal-on-silicon design [194]. Alternatively, an existing commercially-available digital micromirror device (DMD) can be used with its pixels grouped together to approximate a smaller number of larger-area pixels [195].

The multivariate optical computer (MOC) described in chapter 3 could also be used for a spectral form of *compressive sensing* [185, 196, 197]. As with multivariate optical computing, compressive sensing takes advantage of the *sparsity* of the quantities to be measured. In this case, where the number of wavelength bins N is known to be much

greater than the number of target spectra N_S , the spectrum is said to be N_S -sparse. Although multivariate optical computing has also been called compressive detection [179, 182], compressive sensing differs from multivariate optical computing as follows: In place of the N_S spectral matched filters generated from the power spectrum data, compressive sensing generates $M \geq N_S$ random filters and makes a measurement for each filter. Reconstruction algorithms recreate the N spectral bins from the M measurements and the known random filters. For certain classes of random filters, the reconstruction succeeds with high probability [196].

Compressive sensing provides a distinct advantage over multivariate optical computing, namely that no “training data” is needed (the N_S spectral scans of N points each described in chapter 3). This can significantly reduce the time required to get an initial concentration measurement, although since $M > N_S$ there is a slight speed disadvantage when performing many repeated concentration measurements (such as in a production-line inspection setting). Compressive sensing may be able to handle larger libraries of target chemicals than multivariate optical computing, and compressive sensing methods can be implemented with only one detector (but are also amenable to the multiple-vertical-zone, multiple-output method described in section 3.4). However, unlike multivariate optical computing, compressive sensing requires extensive post-measurement computations in order to reconstruct the spectral information, and then the chemical concentration information would need to be computed through further postprocessing using traditional multivariate regression techniques. The M random measurement filters of compressive sensing act as a random basis for

the measurement, and it is possible that an intelligent choice of random filters would allow measurement of chemical concentrations more directly. Alternatively, it may be possible to develop a reconstruction algorithm that directly reconstructs the relative concentrations, rather than the raw spectral data.

The slow light effect might be used to enhance the spectral resolution of the MOC described in chapter 3. The grating shown in Fig. 3.1 could be replaced by an atomic prism [198] or possibly by a photonic crystal superprism with appropriate input and output coupling [127, 128]. An even more exciting possibility presents itself in nanophotonics: It may be possible to implement multivariate optical computation methods in an on-chip device [199]. For instance, using the slow-light-enhanced arrayed waveguide grating (AWG) structure discussed in section 1.5.2, the coupling efficiency of the output waveguides could be used to implement the spectral matched filter coefficients (corresponding to the SLM reflectance in the experiment in chapter 3). Alternatively, a waveguide on a chip could be evanescently coupled to a series of microresonators of varied resonance frequencies, with the coupling coefficient between the waveguide and each resonator acting as the spectral matched filter coefficient at each wavelength. An on-chip MOC coupled with a slow-light-enhanced nanophotonic spectrometer, as well as other integrated, on-chip optical devices such as sources and detectors, could lead to a single-chip solution for detection of hazardous chemicals, biological agents, or other substances.

References

- [1] J. E. Vornehm and R. W. Boyd, “Slow and fast light,” in *Tutorials in Complex Photonic Media*, M. A. Noginov, G. Dewar, M. W. McCall, and N. I. Zheludev, eds. (SPIE, Bellingham, WA, USA, 2009), pp. 647–685.
- [2] Z. Shi, A. Schweinsberg, J. E. Vornehm Jr., M. A. Martínez Gámez, and R. W. Boyd, “Low distortion, continuously tunable, positive and negative time delays by slow and fast light using stimulated Brillouin scattering,” *Phys. Lett. A* **374**, 4071–4074 (2010).
- [3] A. Schweinsberg, Z. Shi, J. E. Vornehm, and R. W. Boyd, “Demonstration of a slow-light laser radar,” *Opt. Express* **19**, 15760–15769 (2011).
- [4] A. Schweinsberg, Z. Shi, J. E. Vornehm, and R. W. Boyd, “A slow-light laser radar system with two-dimensional scanning,” *Opt. Lett.* **37**, 329–331 (2012).
- [5] Z. Shi, A. Schweinsberg, J. E. Vornehm, and R. W. Boyd, “A slow-light laser radar (SLIDAR),” *Opt. Photon. News* **23**, 51 (2012).
- [6] J. E. Vornehm, A. Schweinsberg, Z. Shi, D. J. Gauthier, and R. W. Boyd, “Phase locking of multiple optical fiber channels for a slow-light-enabled laser radar system,” *Opt. Express* **21**, 13094–13104 (2013).
- [7] S. Murugkar, I. De Leon, Z. Shi, G. Lopez-Galmiche, J. Salvail, E. Ma, B. Gao, A. C. Liapis, J. E. Vornehm, and R. W. Boyd, “Development of a slow-light spectrometer on a chip,” *Proc. SPIE* **8264**, 82640T (2012).
- [8] J. Marangos, “Slow light in cool atoms,” *Nature (London)* **397**, 559–560 (1999).
- [9] L. V. Hau, S. E. Harris, Z. Dutton, and C. H. Behroozi, “Light speed reduction to 17 metres per second in an ultracold atomic gas,” *Nature (London)* **397**, 594–598 (1999).
- [10] A. Sommerfeld, “Über die Fortpflanzung des Lichtes in dispergierenden Medien,” *Ann. Phys.* **349**, 177–202 (1914).
- [11] L. Brillouin, “Über die Fortpflanzung des Lichtes in dispergierenden Medien,” *Ann. Phys.* **349**, 203–240 (1914).

- [12] P. W. Milonni, *Fast Light, Slow Light and Left-Handed Light* (Institute of Physics Publishing, Bristol, UK, 2004).
- [13] S. Glasgow, M. Meilstrup, J. Peatross, and M. Ware, “Real-time recoverable and irrecoverable energy in dispersive-dissipative dielectrics,” *Phys. Rev. E* **75**, 016616 (2007).
- [14] R. W. Boyd, “Material slow light and structural slow light: similarities and differences for nonlinear optics [invited],” *J. Opt. Soc. Am. B* **28**, A38–A44 (2011).
- [15] R. W. Boyd, *Nonlinear Optics* (Academic Press, Burlington, MA, USA, 2008), 3rd ed.
- [16] B. E. A. Saleh and M. C. Teich, *Fundamentals of Photonics* (Wiley–Interscience, New York, 2007), 2nd ed.
- [17] K. Y. Song, M. Herráez, and L. Thévenaz, “Observation of pulse delaying and advancement in optical fibers using stimulated Brillouin scattering,” *Opt. Express* **13**, 82–88 (2005).
- [18] Y. Okawachi, M. S. Bigelow, J. E. Sharping, Z. Zhu, A. Schweinsberg, D. J. Gauthier, R. W. Boyd, and A. L. Gaeta, “Tunable all-optical delays via Brillouin slow light in an optical fiber,” *Phys. Rev. Lett.* **94**, 153902 (2005).
- [19] A. Loayssa, R. Hernández, D. Benito, and S. Galech, “Characterization of stimulated Brillouin scattering spectra by use of optical single-sideband modulation,” *Opt. Lett.* **29**, 638–640 (2004).
- [20] M. González Herráez, K. Y. Song, and L. Thévenaz, “Arbitrary-bandwidth Brillouin slow light in optical fibers,” *Opt. Express* **14**, 1395–1400 (2006).
- [21] Z. Zhu, A. M. C. Dawes, D. J. Gauthier, L. Zhang, and A. E. Willner, “12-GHz-bandwidth SBS slow light in optical fibers,” in *Proceedings of the Optical Fiber Communication Conference*, (IEEE, Piscataway, NJ, USA, 2006), paper PDP1.
- [22] Z. Zhu, A. M. C. Dawes, D. J. Gauthier, L. Zhang, and A. E. Willner, “Broadband SBS slow light in an optical fiber,” *J. Lightwave Technol.* **25**, 201–206 (2007).
- [23] M. D. Stenner, M. A. Neifeld, Z. Zhu, A. M. C. Dawes, and D. J. Gauthier, “Distortion management in slow-light pulse delay,” *Opt. Express* **13**, 9995–10002 (2005).
- [24] T. Sakamoto, T. Yamamoto, K. Shiraki, and T. Kurashima, “Low distortion slow light in flat Brillouin gain spectrum by using optical frequency comb,” *Opt. Express* **16**, 8026–8032 (2008).
- [25] K. Y. Song and K. Hotate, “25GHz bandwidth Brillouin slow light in optical fibers,” *Opt. Lett.* **32**, 217–219 (2007).

- [26] R. D. Esmann, M. Y. Frankel, J. Dexter, L. Goldberg, M. Parent, D. Stilwell, and D. Cooper, "Fiber-optic prism true time-delay antenna feed," *IEEE Photon. Technol. Lett.* **5**, 1347–1349 (1993).
- [27] S. Johns, D. Norton, C. Keefer, R. Erdmann, and R. Soref, "Variable time delay of microwave signals using high dispersion fibre," *Electron. Lett.* **29**, 555–556 (1993).
- [28] J. Sharping, Y. Okawachi, J. van Howe, C. Xu, Y. Wang, A. Willner, and A. Gaeta, "All-optical, wavelength and bandwidth preserving, pulse delay based on parametric wavelength conversion and dispersion," *Opt. Express* **13**, 7872–7877 (2005).
- [29] Y. Okawachi, J. E. Sharping, C. Xu, and A. L. Gaeta, "Large tunable optical delays via self-phase modulation and dispersion," *Opt. Express* **14**, 12022–12027 (2006).
- [30] Y. Okawachi, R. Salem, and A. L. Gaeta, "Continuous tunable delays at 10-Gb/s data rates using self-phase modulation and dispersion," *J. Lightwave Technol.* **25**, 3710–3715 (2007).
- [31] O. F. Yilmaz, L. Christen, X. Wu, S. R. Nuccio, I. Fazal, and A. E. Willner, "Time-slot interchange of 40 Gbits/s variable length optical packets using conversion–dispersion-based tunable delays," *Opt. Lett.* **33**, 1954–1956 (2008).
- [32] S. E. Harris, J. E. Field, and A. Imamoglu, "Nonlinear optical processes using electromagnetically induced transparency," *Phys. Rev. Lett.* **64**, 1107–1110 (1990).
- [33] H. Kang, G. Hernandez, and Y. Zhu, "Superluminal and slow light propagation in cold atoms," *Phys. Rev. A* **70**, 011801 (2004).
- [34] M. D. Lukin, S. F. Yelin, and M. Fleischhauer, "Entanglement of atomic ensembles by trapping correlated photon states," *Phys. Rev. Lett.* **84**, 4232–4235 (2000).
- [35] F. D. Phillips, A. Fleischhauer, A. Mair, R. L. Walsworth, and M. D. Lukin, "Storage of light in atomic vapor," *Phys. Rev. Lett.* **86**, 783–786 (2001).
- [36] A. V. Turukhin, V. S. Sudarshanam, M. S. Shahriar, J. A. Musser, B. S. Ham, and P. R. Hemmer, "Observation of ultraslow and stored light pulses in a solid," *Phys. Rev. Lett.* **88**, 023602 (2001).
- [37] M. Phillips and H. Wang, "Spin coherence and electromagnetically induced transparency via exciton correlations," *Phys. Rev. Lett.* **89**, 186401 (2002).
- [38] C. J. Chang-Hasnain, P.-C. Ku, J. Kim, and S.-L. Chuang, "Variable optical buffer using slow light in semiconductor nanostructures," *Proc. IEEE* **91**, 1884–1897 (2003).
- [39] S.-W. Chang, S. L. Chuang, C. J. Chang-Hasnain, and H. Wang, "Slow light using spin coherence and V-type electromagnetically induced transparency in [110] strained quantum wells," *J. Opt. Soc. Am. B* **24**, 849–859 (2007).

- [40] J. Kim, S. L. Chuang, P.-C. Ku, and C. J. Chang-Hasnain, “Slow light using semiconductor quantum dots,” *J. Phys.: Cond. Matt.* **16**, S3727–S3735 (2004).
- [41] S. Marcinkevičius, A. Gushterov, and J. P. Reithmaier, “Transient electromagnetically induced transparency in self-assembled quantum dots,” *Appl. Phys. Lett.* **92**, 041113 (2008).
- [42] S. Ghosh, J. E. Sharping, D. G. Ouzounov, and A. L. Gaeta, “Resonant optical interactions with molecules confined in photonic band-gap fibers,” *Phys. Rev. Lett.* **94**, 093902 (2005).
- [43] L. Maleki, A. B. Matsko, A. A. Savchenkov, and V. S. Ilchenko, “Tunable delay line with interacting whispering-gallery-mode resonators,” *Opt. Lett.* **29**, 626–628 (2004).
- [44] D. D. Smith, H. Chang, K. A. Fuller, A. T. Rosenberger, and R. W. Boyd, “Coupled-resonator-induced transparency,” *Phys. Rev. A* **69**, 063804 (2004).
- [45] D. D. Smith, N. N. Lepeshkin, A. Schweinsberg, G. Gehring, R. W. Boyd, Q.-H. Park, H. Chang, and D. J. Jackson, “Coupled-resonator-induced transparency in a fiber system,” *Opt. Commun.* **264**, 163–168 (2006).
- [46] J. D. Franson and S. M. Hendrickson, “Optical transparency using interference between two modes of a cavity,” *Phys. Rev. A* **74**, 053817 (2006).
- [47] Y. Dumeige, T. K. Ngân Nguyễn, L. Ghişa, S. Trebaol, and P. Féron, “Optical measurement of the phase shift introduced by a slow-light medium based on coupled erbium-doped fiber resonators,” *Proc. SPIE* **6904**, 690407 (2008).
- [48] Y. Dumeige, T. K. Ngân Nguyễn, L. Ghişa, S. Trebaol, and P. Féron, “Measurement of the dispersion induced by a slow-light system based on coupled active-resonator-induced transparency,” *Phys. Rev. A* **78**, 013818 (2008).
- [49] S. E. Harris, “Electromagnetically induced transparency in an ideal plasma,” *Phys. Rev. Lett.* **77**, 5357–5360 (1996).
- [50] G. Shvets and J. S. Wurtele, “Transparency of magnetized plasma at the cyclotron frequency,” *Phys. Rev. Lett.* **89**, 115003 (2002).
- [51] M. Tushentsov, G. Shvets, A. Yu. Kryachko, and M. D. Tokman, “Undulator-induced transparency of magnetized plasma: New approach to electromagnetic energy compression,” *IEEE Trans. Plasma Sci.* **33**, 23–31 (2005).
- [52] R. Gad, J. G. Leopold, A. Fisher, D. R. Fredkin, and A. Ron, “Observation of magnetically induced transparency in a classical magnetized plasma,” *Phys. Rev. Lett.* **108**, 155003 (2012).

- [53] E. Kawamori, W.-J. Syugu, T.-Y. Hsieh, S.-X. Song, and C. Z. Cheng, “Experimental identification of electromagnetically induced transparency in magnetized plasma,” *Phys. Rev. Lett.* **108**, 075003 (2012).
- [54] S. E. Schwarz and T. Y. Tan, “Wave interactions in saturable absorbers,” *Appl. Phys. Lett.* **10**, 4–7 (1967).
- [55] M. S. Bigelow, N. N. Lepeshkin, and R. W. Boyd, “Observation of ultraslow light propagation in a ruby crystal at room temperature,” *Phys. Rev. Lett.* **90**, 113903 (2003).
- [56] A. Schweinsberg, N. Lepeshkin, M. S. Bigelow, R. W. Boyd, and S. Jarabo, “Observation of superluminal and slow light propagation in erbium-doped optical fiber,” *Europhys. Lett.* **73**, 218–224 (2006).
- [57] N. G. Basov, R. V. Ambartsumyan, V. S. Zuev, P. G. Kryukov, and V. S. Letokhov, “Propagation velocity of an intense light pulse in a medium with inverse population,” *Sov. Phys. Dokl.* **10**, 1039 (1966).
- [58] A. C. Selden, “Pulse transmission through a saturable absorber,” *Brit. J. Appl. Phys.* **18**, 743–748 (1967).
- [59] A. C. Selden, “Analysis of the saturable absorber transmission equation,” *J. Phys. D* **3**, 1935–1943 (1970).
- [60] M. S. Bigelow, N. N. Lepeshkin, and R. W. Boyd, “Superluminal and slow light propagation in a room-temperature solid,” *Science* **301**, 200–202 (2003).
- [61] G. M. Gehring, A. Schweinsberg, C. Barsi, N. Kostinski, and R. W. Boyd, “Observation of backwards pulse propagation through a medium with a negative group velocity,” *Science* **312**, 895–897 (2006).
- [62] J. Mørk, R. Kjøer, M. van der Poel, and K. Yvind, “Slow light in a semiconductor waveguide at gigahertz frequencies,” *Opt. Express* **13**, 8136–8145 (2005).
- [63] P. Palinginis, S. Crankshaw, F. Sedgwick, E.-T. Kim, M. Moewe, C. J. Chang-Hasnain, H. Wang, and S.-L. Chuang, “Ultraslow light (< 200 m/s) propagation in a semiconductor nanostructure,” *Appl. Phys. Lett.* **87**, 171102 (2005).
- [64] H. Su and S. L. Chuang, “Room temperature slow and fast light in quantum-dot semiconductor optical amplifiers,” *Appl. Phys. Lett.* **88**, 061102 (2006).
- [65] P.-C. Ku, F. Sedgwick, C. J. Chang-Hasnain, P. Palinginis, T. Li, H. Wang, S.-W. Chang, and S.-L. Chuang, “Slow light in semiconductor quantum wells,” *Opt. Lett.* **29**, 2291–2293 (2004).

- [66] R. M. Camacho, M. V. Pack, J. C. Howell, A. Schweinsberg, and R. W. Boyd, "Wide-bandwidth, tunable, multiple-pulse-width optical delays using slow light in cesium vapor," *Phys. Rev. Lett.* **98**, 153601 (2007).
- [67] X. Zhao, P. Palinginis, B. Pesala, C. J. Chang-Hasnain, and P. Hemmer, "Tunable ultraslow light in vertical-cavity surface-emitting laser amplifier," *Opt. Express* **13**, 7899–7904 (2005).
- [68] P. C. Ku, C. J. Chang-Hasnain, and S. L. Chuang, "Slow light in semiconductor heterostructures," *J. Phys. D* **40**, R93–R107 (2007).
- [69] K. Ogawa, T. Katsuyama, and H. Nakamura, "Time-of-flight measurement of excitonic polaritons in a GaAs/AlGaAs quantum well," *Appl. Phys. Lett.* **53**, 1077–1079 (1988).
- [70] N. Shaw, W. J. Stewart, J. Heaton, and D. R. Wight, "Optical slow-wave resonant modulation in electro-optic GaAs/AlGaAs modulators," *Electron. Lett.* **35**, 1557–1558 (1999).
- [71] P.-C. Ku, C. J. Chang-Hasnain, and S. L. Chuang, "Variable semiconductor all-optical buffer," *Electron. Lett.* **38**, 1581–1583 (2002).
- [72] K. Lee and N. M. Lawandy, "Optically induced pulse delay in a solid-state Raman amplifier," *Appl. Phys. Lett.* **78**, 703–705 (2001).
- [73] F. L. Kien, J. Q. Liang, and K. Hakuta, "Slow light produced by far-off-resonance Raman scattering," *IEEE J. Sel. Topics Quantum Electron.* **9**, 93–101 (2003).
- [74] J. Sharping, Y. Okawachi, and A. Gaeta, "Wide bandwidth slow light using a Raman fiber amplifier," *Opt. Express* **13**, 6092–6098 (2005).
- [75] E. Yablonovitch and T. J. Gmitter, "Photonic band structure: The face-centered-cubic case," *Phys. Rev. Lett.* **63**, 1950–1953 (1989).
- [76] S. John and J. Wang, "Quantum optics of localized light in a photonic band gap," *Phys. Rev. B* **43**, 12772–12789 (1991).
- [77] J. M. Bendickson, J. P. Dowling, and M. Scalora, "Analytic expressions for the electromagnetic mode density in finite, one-dimensional, photonic band-gap structures," *Phys. Rev. E* **53**, 4107–4121 (1996).
- [78] M. Scalora, R. J. Flynn, S. B. Reinhardt, R. L. Fork, M. J. Bloemer, M. D. Tocci, C. M. Bowden, H. S. Ledbetter, J. M. Bendickson, J. P. Dowling, and R. P. Leavitt, "Ultrashort pulse propagation at the photonic band edge: Large tunable group delay with minimal distortion and loss," *Phys. Rev. E* **54**, R1078–R1081 (1996).

- [79] M. L. Povinelli, S. Johnson, and J. D. Joannopoulos, “Slow-light, band-edge waveguides for tunable time delays,” *Opt. Express* **13**, 7145–7159 (2005).
- [80] J. D. Joannopoulos, S. G. Johnson, J. N. Winn, and R. D. Meade, *Photonic Crystals: Molding the Flow of Light* (Princeton University Press, Princeton, NJ, 2008), 2nd ed.
- [81] T. Baba, “Slow light in photonic crystals,” *Nature Photon.* **2**, 465–473 (2008).
- [82] C. Monat, B. Corcoran, M. Ebnali-Heidari, C. Grillet, B. J. Eggleton, T. P. White, L. O’Faolain, and T. F. Krauss, “Slow light enhancement of nonlinear effects in silicon engineered photonic crystal waveguides,” *Opt. Express* **17**, 2944–2953 (2009).
- [83] C. Monat, M. de Sterke, and B. J. Eggleton, “Slow light enhanced nonlinear optics in periodic structures,” *J. Opt.* **12**, 104003 (2010).
- [84] M. Soljačić and J. D. Joannopoulos, “Enhancement of nonlinear effects using photonic crystals,” *Nature Mater.* **3**, 211–219 (2004).
- [85] S. Y. Lin, J. G. Fleming, D. L. Hetherington, B. K. Smith, R. Biswas, K. M. Ho, M. M. Sigalas, W. Zubrzycki, S. R. Kurtz, and J. Bur, “A three-dimensional photonic crystal operating at infrared wavelengths,” *Nature (London)* **394**, 251–253 (1998).
- [86] A. Blanco, E. Chomski, S. Grabtchak, M. Ibisate, S. John, S. W. Leonard, C. Lopez, F. Meseguer, H. Miguez, J. P. Mondia, G. A. Ozin, O. Toader, and H. M. van Driel, “Large-scale synthesis of a silicon photonic crystal with a complete three-dimensional bandgap near 1.5 micrometres,” *Nature (London)* **405**, 437–440 (2000).
- [87] S. Noda, K. Tomoda, N. Yamamoto, and A. Chutinan, “Full three-dimensional photonic bandgap crystals at near-infrared wavelengths,” *Science* **289**, 604–606 (2000).
- [88] M. Notomi, A. Shinya, S. Mitsugi, E. Kuramochi, and H. Ryu, “Waveguides, resonators and their coupled elements in photonic crystal slabs,” *Opt. Express* **12**, 1551–1561 (2004).
- [89] D. Mori and T. Baba, “Dispersion-controlled optical group delay device by chirped photonic crystal waveguides,” *Appl. Phys. Lett.* **85**, 1101–1103 (2004).
- [90] A. Yu. Petrov and M. Eich, “Zero dispersion at small group velocities in photonic crystal waveguides,” *Appl. Phys. Lett.* **85**, 4866–4868 (2004).
- [91] J. Li, T. P. White, L. O’Faolain, A. Gomez-Iglesias, and T. F. Krauss, “Systematic design of flat band slow light in photonic crystal waveguides,” *Opt. Express* **16**, 6227–6232 (2008).

- [92] S. A. Schulz, L. O’Faolain, D. M. Beggs, T. P. White, A. Melloni, and T. F. Krauss, “Dispersion engineered slow light in photonic crystals: A comparison,” *J. Opt.* **12**, 104004 (2010).
- [93] S. Hughes, L. Ramunno, J. F. Young, and J. E. Sipe, “Extrinsic optical scattering loss in photonic crystal waveguides: Role of fabrication disorder and photon group velocity,” *Phys. Rev. Lett.* **94**, 033903 (2005).
- [94] L. O’Faolain, T. P. White, D. O’Brien, X. Yuan, M. D. Settle, and T. F. Krauss, “Dependence of extrinsic loss on group velocity in photonic crystal waveguides,” *Opt. Express* **15**, 13129–13138 (2007).
- [95] L. O’Faolain, S. A. Schulz, D. M. Beggs, T. P. White, M. Spasenović, L. Kuipers, F. Morichetti, A. Melloni, S. Mazoyer, J. P. Hugonin, P. Lalanne, and T. F. Krauss, “Loss engineered slow light waveguides,” *Opt. Express* **18**, 27627–27638 (2010).
- [96] Yu. A. Vlasov, M. O’Boyle, H. F. Hamann, and S. J. McNab, “Active control of slow light on a chip with photonic crystal waveguides,” *Nature (London)* **438**, 65–69 (2005).
- [97] R. Jacobsen, A. Lavrinenko, L. Frandsen, C. Peucheret, B. Zsigri, G. Moulin, J. Fage-Pedersen, and P. Borel, “Direct experimental and numerical determination of extremely high group indices in photonic crystal waveguides,” *Opt. Express* **13**, 7861–7871 (2005).
- [98] J. K. S. Poon, J. Scheuer, Y. Xu, and A. Yariv, “Designing coupled-resonator optical waveguide delay lines,” *J. Opt. Soc. Am. B* **21**, 1665–1673 (2004).
- [99] J. Scheuer, G. T. Paloczi, J. K. S. Poon, and A. Yariv, “Coupled resonator optical waveguides: Toward the slowing and storage of light,” *Opt. Photon. News* **16**, 36–40 (2005).
- [100] G. Lenz, B. J. Eggleton, C. K. Madsen, and R. E. Slusher, “Optical delay lines based on optical filters,” *IEEE J. Quantum Electron.* **37**, 525–532 (2001).
- [101] S. Longhi, M. Marano, P. Laporta, M. Belmonte, and P. Crespi, “Experimental observation of superluminal pulse reflection in a double-Lorentzian photonic band gap,” *Phys. Rev. E* **65**, 045602 (2002).
- [102] J. B. Khurgin, “Light slowing down in Moiré fiber gratings and its implications for nonlinear optics,” *Phys. Rev. A* **62**, 013821 (2000).
- [103] J. K. S. Poon, L. Zhu, G. A. Rose, and A. Yariv, “Polymer microring coupled-resonator optical waveguides,” *J. Lightwave Technol.* **24**, 1843–1849 (2006).

- [104] R. W. Boyd, D. J. Gauthier, A. L. Gaeta, and A. E. Willner, “Maximum time delay achievable on propagation through a slow-light medium,” *Phys. Rev. A* **71**, 023801 (2005).
- [105] R. S. Tucker, P.-C. Ku, and C. J. Chang-Hasnain, “Delay-bandwidth product and storage density in slow-light optical buffers,” *Electron. Lett.* **41**, 208–209 (2005).
- [106] M. Fleischhauer and M. D. Lukin, “Quantum memory for photons: Dark-state polaritons,” *Phys. Rev. A* **65**, 022314 (2002).
- [107] P. Ginzburg and M. Orenstein, “Slow light and voltage control of group velocity in resonantly coupled quantum wells,” *Opt. Express* **14**, 12467–12472 (2006).
- [108] R. S. Tucker, P.-C. Ku, and C. J. Chang-Hasnain, “Slow-light optical buffers: Capabilities and fundamental limitations,” *J. Lightwave Technol.* **23**, 4046–4066 (2005).
- [109] H. Shin, A. Schweinsberg, G. Gehring, K. Schwertz, H. J. Chang, R. W. Boyd, Q.-H. Park, and D. J. Gauthier, “Reducing pulse distortion in fast-light pulse propagation through an erbium-doped fiber amplifier,” *Opt. Lett.* **32**, 906–908 (2007).
- [110] B. Zhang, L. Yan, I. Fazal, L. Zhang, A. E. Willner, Z. Zhu, and D. J. Gauthier, “Slow light on Gbit/s differential-phase-shift-keying signals,” *Opt. Express* **15**, 1878–1883 (2007).
- [111] E. Shumakher, N. Orbach, A. Nevet, D. Dahan, and G. Eisenstein, “On the balance between delay, bandwidth and signal distortion in slow light systems based on stimulated Brillouin scattering in optical fibers,” *Opt. Express* **14**, 5877–5884 (2006).
- [112] M. D. Lukin and A. Imamoglu, “Controlling photons using electromagnetically induced transparency,” *Nature (London)* **413**, 273–276 (2001).
- [113] R. W. Boyd and D. J. Gauthier, “‘Slow’ and ‘fast’ light,” in *Progress in Optics*, vol. Volume 43 (Elsevier, 2002), chap. 6, pp. 497–530.
- [114] S. E. Harris and L. V. Hau, “Nonlinear optics at low light levels,” *Phys. Rev. Lett.* **82**, 4611–4614 (1999).
- [115] S. Chin, I. Dicaire, J.-C. Beugnot, S. Foaleng-Mafang, M. Gonzalez-Herraez, and L. Thévenaz, “Material slow light does not enhance Beer–Lambert absorption,” in *Slow and Fast Light*, (Optical Society of America, 2009), OSA Technical Digest (CD), paper SMA3.
- [116] I. Dicaire, S. Chin, and L. Thévenaz, “Structural slow light can enhance Beer–Lambert absorption,” in *Slow and Fast Light*, (Optical Society of America, 2011), OSA Technical Digest (CD), paper SLWC2.

- [117] Y. Chen and S. Blair, “Nonlinearity enhancement in finite coupled-resonator slow-light waveguides,” *Opt. Express* **12**, 3353–3366 (2004).
- [118] N. A. R. Bhat and J. E. Sipe, “Optical pulse propagation in nonlinear photonic crystals,” *Phys. Rev. E* **64**, 056604 (2001).
- [119] D. J. Gauthier and R. W. Boyd, “Fast light, slow light and optical precursors: What does it all mean?” *Photon. Spectra* **41**, 82–90 (2007).
- [120] A. Içsevçi and W. E. Lamb, “Propagation of light pulses in a laser amplifier,” *Phys. Rev.* **185**, 517–545 (1969).
- [121] P. S. Argall and R. J. Sica, “Lidar,” in *Encyclopedia of Imaging Science and Technology*, J. P. Hornak, ed. (John Wiley & Sons, Inc., 2002), pp. 869–889.
- [122] P. F. McManamon, T. A. Dorschner, D. L. Corkum, L. J. Friedman, D. S. Hobbs, M. Holz, S. Liberman, H. Q. Nguyen, D. P. Resler, R. C. Sharp, and E. A. Watson, “Optical phased array technology,” *Proc. IEEE* **84**, 268–298 (1996).
- [123] I. Frigyes and A. J. Seeds, “Optically generated true-time delay in phased-array antennas,” *IEEE Trans. Micro. Theory Tech.* **43**, 2378–2386 (1995).
- [124] M. I. Skolnik, *Radar Handbook* (McGraw-Hill, New York, 2008), pp. 13.1–13.51, 3rd ed.
- [125] S. Janz, A. Balakrishnan, S. Charbonneau, P. Cheben, M. Cloutier, A. Delâge, K. Dossou, L. Erickson, M. Gao, P. A. Krug, B. Lamontagne, M. Packirisamy, M. Pearson, and D.-X. Xu, “Planar waveguide echelle gratings in silica-on-silicon,” *IEEE Photon. Technol. Lett.* **16**, 503–505 (2004).
- [126] P. Cheben, J. H. Schmid, A. Delâge, A. Densmore, S. Janz, B. Lamontagne, J. Lapointe, E. Post, P. Waldron, and D.-X. Xu, “A high-resolution silicon-on-insulator arrayed waveguide grating microspectrometer with sub-micrometer aperture waveguides,” *Opt. Express* **15**, 2299–2306 (2007).
- [127] H. Kosaka, T. Kawashima, A. Tomita, M. Notomi, T. Tamamura, T. Sato, and S. Kawakami, “Superprism phenomena in photonic crystals,” *Phys. Rev. B* **58**, R10096–R10099 (1998).
- [128] B. Momeni, M. Chamanzar, E. Shah Hosseini, M. Askari, M. Soltani, and A. Adibi, “Strong angular dispersion using higher bands of planar silicon photonic crystals,” *Opt. Express* **16**, 14213–14220 (2008).
- [129] A. Nitkowski, L. Chen, and M. Lipson, “Cavity-enhanced on-chip absorption spectroscopy using microring resonators,” *Opt. Express* **16**, 11930–11936 (2008).

- [130] V. D. Nguyen, B. I. Akca, K. Wörhoff, R. M. de Ridder, M. Pollnau, T. G. van Leeuwen, and J. Kalkman, “Spectral domain optical coherence tomography imaging with an integrated optics spectrometer,” *Opt. Lett.* **36**, 1293–1295 (2011).
- [131] Z. Xia, A. A. Eftekhar, M. Soltani, B. Momeni, Q. Li, M. Chamanzar, S. Yegnanarayanan, and A. Adibi, “High resolution on-chip spectroscopy based on miniaturized microdonut resonators,” *Opt. Express* **19**, 12356–12364 (2011).
- [132] O. M. Matos, M. L. Calvo, P. Cheben, S. Janz, J. A. Rodrigo, D.-X. Xu, and A. Delage, “Arrayed waveguide grating based on group-index modification,” *J. Lightwave Technol.* **24**, 1551–1557 (2006).
- [133] Z. Shi, R. W. Boyd, D. J. Gauthier, and C. C. Dudley, “Enhancing the spectral sensitivity of interferometers using slow-light media,” *Opt. Lett.* **32**, 915–917 (2007).
- [134] Z. Shi, R. W. Boyd, R. M. Camacho, P. K. Vudyasetu, and J. C. Howell, “Slow-light Fourier transform interferometer,” *Phys. Rev. Lett.* **99**, 240801 (2007).
- [135] Z. Shi and R. W. Boyd, “Slow-light enhanced spectrometers on chip,” *Proc. SPIE* **8007**, 80071D (2011).
- [136] M. K. Smit and C. Van Dam, “PHASAR-based WDM-devices: Principles, design and applications,” *IEEE J. Sel. Topics Quantum Electron.* **2**, 236–250 (1996).
- [137] A. C. Liapis, Z. Shi, and R. W. Boyd, “Optimizing photonic crystal waveguides for on-chip spectroscopic applications,” *Opt. Express* **21**, 10160–10165 (2013).
- [138] J. R. Lowell and E. Parra, “Applications of slow light: A DARPA perspective,” *Proc. SPIE* **5735**, 80–86 (2005).
- [139] C. Yu, T. Luo, L. Zhang, and A. E. Willner, “Data pulse distortion induced by a slow-light tunable delay line in optical fiber,” *Opt. Lett.* **32**, 20–22 (2007).
- [140] B. Zhang, L. Zhang, L.-S. Yan, I. Fazal, J. Yang, and A. E. Willner, “Continuously-tunable, bit-rate variable OTDM using broadband SBS slow-light delay line,” *Opt. Express* **15**, 8317–8322 (2007).
- [141] S. Mookherjea and A. Yariv, “Kerr-stabilized super-resonant modes in coupled-resonator optical waveguides,” *Phys. Rev. E* **66**, 046610 (2002).
- [142] “DDR3 SDRAM standard,” JEDEC Standard JESD79-3F, JEDEC Solid State Technology Association, Arlington, VA (2012).
- [143] R. M. Camacho, C. J. Broadbent, I. Ali-Khan, and J. C. Howell, “All-optical delay of images using slow light,” *Phys. Rev. Lett.* **98**, 043902 (2007).

- [144] S. Franke-Arnold, G. Gibson, R. W. Boyd, and M. J. Padgett, “Rotary photon drag enhanced by a slow-light medium,” *Science* **333**, 65–67 (2011).
- [145] E. Wisniewski-Barker, G. Gibson, S. Franke-Arnold, Z. Shi, R. W. Boyd, and M. J. Padgett, “Evidence of slow-light effects from rotary drag of structured beams,” *New J. Phys.* **15**, 083020 (2013).
- [146] M. S. Shahriar, G. S. Pati, R. Tripathi, V. Gopal, M. Messall, and K. Salit, “Ultrahigh enhancement in absolute and relative rotation sensing using fast and slow light,” *Phys. Rev. A* **75**, 053807 (2007).
- [147] M. Salit, G. S. Pati, K. Salit, and M. S. Shahriar, “Fast-light for astrophysics: Super-sensitive gyroscopes and gravitational wave detectors,” *J. Mod. Opt.* **54**, 2425–2440 (2007).
- [148] G. S. Pati, M. Salit, K. Salit, and M. S. Shahriar, “Demonstration of a tunable-bandwidth white-light interferometer using anomalous dispersion in atomic vapor,” *Phys. Rev. Lett.* **99**, 133601 (2007).
- [149] G. S. Pati, M. Salit, K. Salit, and M. S. Shahriar, “Demonstration of tunable displacement-measurement-sensitivity using variable group index in a ring resonator,” *Opt. Commun.* **281**, 4931–4935 (2008).
- [150] M. S. Shahriar, G. S. Pati, M. Messall, and K. R. Salit, “Slow-light based control of interferometer sensitivity,” in *Frontiers in Optics/Laser Science XXII*, (OSA, Washington, DC, USA, 2006), paper LTuJ4.
- [151] A. Wicht, K. Danzmann, M. Fleischhauer, M. Scully, G. Müller, and R.-H. Rinkleff, “White-light cavities, atomic phase coherence, and gravitational wave detectors,” *Opt. Commun.* **134**, 431–439 (1997).
- [152] P. F. McManamon, “Review of ladar: A historic, yet emerging, sensor technology with rich phenomenology,” *Opt. Eng.* **51**, 060901 (2012).
- [153] F. Öhman, K. Yvind, and J. Mørk, “Slow light in a semiconductor waveguide for true-time delay applications in microwave photonics,” *IEEE Photon. Technol. Lett.* **19**, 1145–1147 (2007).
- [154] C. D. Nabors, “Effects of phase errors on coherent emitter arrays,” *Appl. Opt.* **33**, 2284–2289 (1994).
- [155] S. J. Augst, T. Y. Fan, and A. Sanchez, “Coherent beam combining and phase noise measurements of ytterbium fiber amplifiers,” *Opt. Lett.* **29**, 474–476 (2004).
- [156] T. Y. Fan, “Laser beam combining for high-power, high-radiance sources,” *IEEE J. Sel. Topics Quantum Electron.* **11**, 567–577 (2005).

- [157] N. J. Miller, M. P. Dierking, and B. D. Duncan, "Optical sparse aperture imaging," *Appl. Opt.* **46**, 5933–5943 (2007).
- [158] W. Liang, N. Satyan, F. Aflatouni, A. Yariv, A. Kewitsch, G. Rakuljic, and H. Hashemi, "Coherent beam combining with multilevel optical phase-locked loops," *J. Opt. Soc. Am. B* **24**, 2930–2939 (2007).
- [159] W. Liang, A. Yariv, A. Kewitsch, and G. Rakuljic, "Coherent combining of the output of two semiconductor lasers using optical phase-lock loops," *Opt. Lett.* **32**, 370–372 (2007).
- [160] S. J. Augst, J. K. Ranka, T. Y. Fan, and A. Sanchez, "Beam combining of ytterbium fiber amplifiers (Invited)," *J. Opt. Soc. Am. B* **24**, 1707–1715 (2007).
- [161] P. Zhou, Z. Liu, X. Wang, Y. Ma, H. Ma, X. Xu, and S. Guo, "Coherent beam combining of fiber amplifiers using stochastic parallel gradient descent algorithm and its application," *IEEE J. Sel. Topics Quantum Electron.* **15**, 248–256 (2009).
- [162] A. M. Marino and C. R. Stroud, "Phase-locked laser system for use in atomic coherence experiments," *Rev. Sci. Instrum.* **79**, 013104 (2008).
- [163] T. M. Shay and V. Benham, "First experimental demonstration of phase locking of optical fiber arrays by RF phase modulation," *Proc. SPIE* **5550**, 313–319 (2004).
- [164] T. M. Shay, V. Benham, J. T. Baker, B. Ward, A. D. Sanchez, M. A. Culpepper, D. Pilkington, J. Spring, D. J. Nelson, and C. A. Lu, "First experimental demonstration of self-synchronous phase locking of an optical array," *Opt. Express* **14**, 12015–12021 (2006).
- [165] M. A. Vorontsov, G. W. Carhart, and J. C. Ricklin, "Adaptive phase-distortion correction based on parallel gradient-descent optimization," *Opt. Lett.* **22**, 907–909 (1997).
- [166] L. Liu and M. A. Vorontsov, "Phase-locking of tiled fiber array using SPGD feedback controller," *Proc. SPIE* **5895**, 58950P (2005).
- [167] X. Wang, P. Zhou, Y. Ma, J. Leng, X. Xu, and Z. Liu, "Active phasing a nine-element 1.14 kW all-fiber two-tone MOPA array using SPGD algorithm," *Opt. Lett.* **36**, 3121–3123 (2011).
- [168] P. Fellgett, "A contribution to the theory of the multiplex interferometric spectrometer. I.—les principes généraux des méthodes nouvelles en spectroscopie interférentielle—A propos de la théorie du spectromètre interférentiel multiplex," *J. Phys. Radium* **19**, 187–191 (1958).

- [169] M. P. Nelson, J. F. Aust, J. A. Dobrowolski, P. G. Verly, and M. L. Myrick, "Multivariate optical computation for predictive spectroscopy," *Anal. Chem.* **70**, 73–82 (1998).
- [170] O. Soyemi, D. Eastwood, L. Zhang, H. Li, J. Karunamuni, P. Gemperline, R. A. Synowicki, and M. L. Myrick, "Design and testing of a multivariate optical element: The first demonstration of multivariate optical computing for predictive spectroscopy," *Anal. Chem.* **73**, 1069–1079 (2001).
- [171] L. Cao, X. Ma, Q. He, H. Long, M. Wu, and G. Jin, "Imaging spectral device based on multiple volume holographic gratings," *Opt. Eng.* **43**, 2009–2016 (2004).
- [172] Z. Li, D. Psaltis, W. Liu, W. R. Johnson, and G. Bearman, "Volume holographic spectral imaging," *Proc. SPIE* **5694**, 33–40 (2005).
- [173] L. Cao and C. Gu, "Matched spectral filter based on reflection holograms for analyte identification," *Appl. Opt.* **48**, 6973–6979 (2009).
- [174] C. Gu, X. Yang, J. Zhang, R. Newhouse, and L. Cao, "Fiber sensors for molecular detection," *Proc. SPIE* **7851**, 785105 (2010).
- [175] M. Bottema, W. Plummer, and J. Strong, "Water vapor in the atmosphere of Venus," *Astrophys. J.* **139**, 1021 (1964).
- [176] E. P. Wagner, B. W. Smith, S. Madden, J. D. Winefordner, and M. Mignardi, "Construction and evaluation of a visible spectrometer using digital micromirror spatial light modulation," *Appl. Spectrosc.* **49**, 1715–1719 (1995).
- [177] N. Uzunbajakava, P. de Peinder, G. W. 't Hooft, and A. T. M. van Gogh, "Low-cost spectroscopy with a variable multivariate optical element," *Anal. Chem.* **78**, 7302–7308 (2006).
- [178] N. T. Quyen, E. Da Silva, N. Q. Dao, and M. D. Jouan, "New Raman spectrometer using a digital micromirror device and a photomultiplier tube detector for rapid on-line industrial analysis. part I: Description of the prototype and preliminary results," *Appl. Spectrosc.* **62**, 273–278 (2008).
- [179] B. M. Davis, A. J. Hemphill, D. Cebeci Maltaş, M. A. Zipper, P. Wang, and D. Ben-Amotz, "Multivariate hyperspectral Raman imaging using compressive detection," *Anal. Chem.* **83**, 5086–5092 (2011).
- [180] Z. J. Smith, S. Strombom, and S. Wachsmann-Hogiu, "Multivariate optical computing using a digital micromirror device for fluorescence and Raman spectroscopy," *Opt. Express* **19**, 16950–16962 (2011).

- [181] C. W. Freudiger, W. Min, G. R. Holtom, B. Xu, M. Dantus, and X. S. Xie, "Highly specific label-free molecular imaging with spectrally tailored excitation-stimulated Raman scattering (STE-SRS) microscopy," *Nature Photon.* **5**, 103–109 (2011).
- [182] D. S. Wilcox, G. T. Buzzard, B. J. Lucier, P. Wang, and D. Ben-Amotz, "Photon level chemical classification using digital compressive detection," *Anal. Chim. Acta* **755**, 17–27 (2012).
- [183] W. C. Sweatt, C. A. A. Boye, S. M. Gentry, M. R. Descour, B. R. Stallard, and C. L. Grotbeck, "ISIS: An information-efficient spectral imaging system," *Proc. SPIE* **3438**, 98–106 (1998).
- [184] F. Melgani and L. Bruzzone, "Classification of hyperspectral remote sensing images with support vector machines," *IEEE Trans. Geosci. Remote Sens.* **42**, 1778–1790 (2004).
- [185] R. M. Willett, M. E. Gehm, and D. J. Brady, "Multiscale reconstruction for computational spectral imaging," *Proc. SPIE* **6498**, 64980L (2007).
- [186] N. Goldstein, P. Vujkovic-Cvijin, M. Fox, B. Gregor, J. Lee, J. Cline, and S. Adler-Golden, "DMD-based adaptive spectral imagers for hyperspectral imagery and direct detection of spectral signatures," *Proc. SPIE* **7210**, 721008 (2009).
- [187] S. P. Love, "Programmable matched filter and Hadamard transform hyperspectral imagers based on micromirror arrays," *Proc. SPIE* **7210**, 721007 (2009).
- [188] A. Wuttig and R. Riesenberger, "Sensitive Hadamard transform imaging spectrometer with a simple MEMS," *Proc. SPIE* **4881**, 167–178 (2003).
- [189] D. M. Haaland and E. V. Thomas, "Partial least-squares methods for spectral analyses. 1. Relation to other quantitative calibration methods and the extraction of qualitative information," *Anal. Chem.* **60**, 1193–1202 (1988).
- [190] H. Sasada and K. Yamada, "Calibration lines of HCN in the 1.5- μm region," *Appl. Opt.* **29**, 3535–3547 (1990).
- [191] S. Sudo, Y. Sakai, H. Yasaka, and T. Ikegami, "Frequency-stabilized DFB laser module using 1.53159 μm absorption line of C_2H_2 ," *IEEE Photon. Technol. Lett.* **1**, 281–284 (1989).
- [192] H.-B. Liu, Y. Chen, G. J. Bastiaans, and X.-C. Zhang, "Detection and identification of explosive RDX by THz diffuse reflection spectroscopy," *Opt. Express* **14**, 415–423 (2006).
- [193] J. Chen, Y. Chen, H. Zhao, G. J. Bastiaans, and X.-C. Zhang, "Absorption coefficients of selected explosives and related compounds in the range of 0.1–2.8 THz," *Opt. Express* **15**, 12060–12067 (2007).

- [194] W. L. Chan, H.-T. Chen, A. J. Taylor, I. Brener, M. J. Cich, and D. M. Mittleman, “A spatial light modulator for terahertz beams,” *Appl. Phys. Lett.* **94**, 213511 (2009).
- [195] D. Shrekenhamer, C. M. Watts, and W. J. Padilla, “Terahertz single pixel imaging with an optically controlled dynamic spatial light modulator,” *Opt. Express* **21**, 12507–12518 (2013).
- [196] R. G. Baraniuk, “Compressive sensing [lecture notes],” *IEEE Signal Process. Mag.* **24**, 118–121 (2007).
- [197] M. E. Gehm, S. T. McCain, N. P. Pitsianis, D. J. Brady, P. Potluri, and M. E. Sullivan, “Static two-dimensional aperture coding for multimodal, multiplex spectroscopy,” *Appl. Opt.* **45**, 2965–2974 (2006).
- [198] D. J. Starling, S. M. Bloch, P. K. Vudyasetu, J. S. Choi, B. Little, and J. C. Howell, “Double Lorentzian atomic prism,” *Phys. Rev. A* **86**, 023826 (2012).
- [199] W.-C. Lai, S. Chakravarty, X. Wang, C. Lin, and R. T. Chen, “On-chip methane sensing by near-IR absorption signatures in a photonic crystal slot waveguide,” *Opt. Lett.* **36**, 984–986 (2011).
- [200] S.-T. Wu, “Birefringence dispersions of liquid crystals,” *Phys. Rev. A* **33**, 1270–1274 (1986).

Appendix A

SLM Calibration

A liquid-crystal-on-silicon spatial light modulator (LCOS SLM), like the one used in chapter 3, consists of a silicon electronic substrate with an array of electrodes, a layer of liquid crystal that can be rotated electrically, a top transparent electrode, and a cover glass or another protective covering. The liquid crystal is birefringent; horizontally and vertically polarized light propagating through the liquid crystal will experience different refractive indices. The electrodes apply a voltage that rotates the liquid crystal. As the liquid crystal molecules rotate, light of a fixed polarization experiences a varying refractive index and accumulates a varying optical phase. In this way, the SLM imparts an electronically controlled amount of phase to each pixel of the light field. (An SLM that operates in this manner is often called a phase-only SLM; other liquid-crystal-based SLMs exist that can modulate amplitude or both amplitude and phase, but they are not covered in this dissertation.)

The SLM is controlled by a computer-generated 8-bit grayscale image, and the 8-bit grayscale values determine the voltages applied to the liquid crystal. The relationship between the 8-bit values and the actual phase imparted to the optical field depends in a nonlinear way on the birefringence and voltage-dependent rotation of the liquid crystal, the wavelength of the field, and the mapping of 8-bit values to the electrode voltage. Further, the liquid crystal is dispersive, particularly across the visible wavelengths (the range used in the experiment in chapter 3). In order to use the SLM effectively, its response must be calibrated, and the calibration must be performed at many wavelengths.

A.1 Two-pixel SLM calibration

The two-pixel blazed grating is straightforward to calibrate. The SLM is configured within the experiment, in the same layout shown in Fig. 3.1. As described in section 3.2.3, the SLM is divided into 64 wavelength bins, and each bin is calibrated independently. A single wavelength bin is set to a two-pixel blazed grating (described in section 3.2.3), and pixels in all other wavelength bins are set to zero. The two-pixel blazed grating has two alternating pixel values, and the maximum deflection efficiency should be achieved when the pixel phases differ by π radians. One of the pixel values is kept at zero, and the other is varied from 0 to 255 (i.e., all possible 8-bit values). An output power measurement is taken for each value. The process is repeated for all 64 wavelength bins, giving a total of 16384 measurements. Once the 256 powers for each wavelength bin are recorded, they are used as a lookup table, and the filter reflectance coefficients are computed on a relative scale

from the minimum to the maximum power. (The choice of zero as the “background value” for the SLM blazed gratings was made after setting the entire SLM—all pixels at once—successively to the values 0 through 255 and measuring the output power. A very small but repeatable change was observed in the detected power, with the lowest power occurring for the value zero.)

To compute the correspondence between each wavelength bin and an actual wavelength, the transmission spectra of the didymium and holmium filters shown in Fig. 3.5 are used in a nonlinear least-squares optimization algorithm to generate a second-order polynomial fit between the wavelength bin indices and the wavelengths reported by the Ocean Optics USB2000 spectrometer. Essentially, the algorithm scales and stretches the horizontal axis of Fig. 3.5(a) until both curves match the curves in Fig. 3.5(b) as closely as possible. (The actual metric is the difference between the experiment data curve and the USB2000 curve interpolated at the same points, weighted by the reciprocal of the standard deviation, so that high-noise data points contribute less to the metric than low-noise data points. Both the didymium and the holmium curves are used in the metric.)

A.2 SLM phase model

The 2-pixel blazed gratings can be calibrated exhaustively, as there are only 256 possible gratings with one of the pixel values is fixed at zero. However, for an n -pixel grating, even with one of the pixel values fixed at zero, there are 256^{n-1} possible gratings (times 64

wavelength bins), and brute-force calibration is intractable. In order to calibrate the SLM for the multiple-output experiment in section 3.4, a model of the SLM phase response must be developed. The phase model used in the SLM calibration in chapter 3 is presented here. (The model developed here is fit to the 2-pixel calibration data, and then the fitted model is extrapolated to produce 5-pixel values as discussed in section A.3.)

Light travels into the liquid crystal, is reflected off the substrate, and travels back through the liquid crystal. For a liquid crystal depth of L (total path length of $2L$), the phase imparted by the SLM is given by

$$\phi = \frac{2\pi\Delta n(2L)}{\lambda}, \quad (\text{A.1})$$

where λ is the wavelength and Δn is the change in refractive index due to the liquid crystal's rotation (from 0 to some maximum change). The form of Eq. (A.1) is easily recognized as $\phi = k(2L)$, except the overall refractive index of the liquid crystal is ignored, since it only adds an overall phase to the light. (That overall phase is different for different wavelengths, but the relative phases of different wavelengths are not important in this experiment.)

The input light is linearly polarized, and the refractive index change Δn is due to the liquid crystal's rotation. The refractive index change is

$$\Delta n = \sin(\theta)\delta n(\lambda), \quad (\text{A.2})$$

where $\delta n(\lambda)$ is the birefringence of the liquid crystal (the difference between the ordinary and extraordinary refractive indices), and θ is the difference between the polarization angle

of the input light and the liquid crystal rotation angle. The angle θ depends in some way on the 8-bit SLM pixel value v , probably nonlinearly. To simplify, the $\sin(\theta)$ term and several other coefficients are replaced with a general function of v . Then, the phase change imposed by the SLM for a given wavelength and a given pixel value is

$$\phi(\lambda, v) = \frac{f(v)\delta n(\lambda)}{\lambda}, \quad (\text{A.3})$$

where $f(v)$ represents the influence of the pixel value on the phase and includes the factor of $2\pi \cdot 2L$ from Eq. (A.1). Note that the influence of the pixel value on the phase (i.e., the rotation of the liquid crystal) is independent of the wavelength. The birefringence of the liquid crystal is given by Wu in Eq. (8) of Ref. 200 as

$$\delta n(\lambda) = G \frac{\lambda^2(\lambda^*)^2}{\lambda^2 - (\lambda^*)^2}, \quad (\text{A.4})$$

where λ^* is a resonance wavelength of the order of 150 nm to 300 nm, and G is a coefficient of the order of 10^{-6} nm^{-1} . (Wu specifies G as a function of temperature, but the temperature is assumed here to be constant.)

The power diffracted by a pixelated (2-pixel) blazed phase grating can be calculated from Fourier optics; the result is used here without derivation. For a wavelength bin centered at wavelength λ , with the white light source as the input, the expected output power from the experiment is

$$P_{\text{cal}}(\lambda, v) = P_{1,\text{bg}} + s_{1,\lambda} \left\{ \frac{1}{2} - \frac{1}{2} \cos \phi(\lambda, v) \right\}. \quad (\text{A.5})$$

$P_{1,\text{bg}}$ is the background power of the white light source (i.e., output power when all SLM pixels are set to zero), and $s_{1,\lambda}$ is the spectral power of the white light source at the wavelength bin corresponding to λ .

To fit the data to the model, Eq. (A.5) must essentially be inverted. The limited phase range of the SLM presents a problem, in that the value of $s_{1,\lambda}$ cannot be determined uniquely for wavelengths at which the SLM has a phase range of less than π . (For reference, the Cambridge Correlators SLM is specified to have a maximum phase range of about 0.8π at 633 nm.)

To resolve this problem, only wavelength bins for which the SLM phase range exceeds π radians are used. The range is determined to exceed π when $P_{\text{cal}}(\lambda, \nu)$ reaches a maximum for some ν and then decreases for increasing values of ν . For these bins, the value of $s_{1,\lambda}$ can be determined, and then the arccosine is computed to find $\phi(\lambda, \nu)$. (Wavelength bins with too little spectral power are excluded, namely those with less than 8% of the maximum power. For these bins, the calibration results were unpredictable, and including the bins made the results unreliable.)

With data for $\phi(\lambda, \nu)$ in hand (for the wavelengths for which the SLM phase range exceeds π), the parameters G and λ^* and the function $f(\nu)$ are estimated using a nonlinear least-squares optimization procedure. The nonlinear optimization procedure works as follows: Beginning with representative values of G and λ^* , the birefringence $\delta n(\lambda)$ is calculated, and an 8th-order polynomial is fit to $f(\nu)$. Then, with the known values of $P_{1,\text{bg}}$ and $s_{1,\lambda}$ and the estimates of G , λ^* , and $f(\nu)$, the model is used to compute the expected

output powers. Based on the difference between these predicted output powers and the measured output powers $P_{\text{cal}}(\lambda, \nu)$, new values of G and λ^* are chosen, and the process is repeated. The end result is a model of $\phi(\lambda, \nu)$ that is a best fit of the available data.

A.3 Multiple-output SLM calibration

Once an estimate of the SLM phase mapping $\phi(\lambda, \nu)$ has been obtained, it is used to generate 5-pixel blazed gratings of 256 different efficiencies for each wavelength bin. (The number 256 was chosen for consistency with the number of different 2-pixel gratings.) Then, using the multiple-output experimental configuration in Fig. 3.7, all the 2-pixel and 5-pixel gratings are put on the SLM, one at a time, and the resulting powers are recorded on both the detectors in order to measure output efficiency and crosstalk. For the two-zone implementation, the optimal sizes of the two vertical SLM zones are determined empirically to balance the power between the two detectors, and all the 2-pixel and 5-pixel gratings are displayed in their respective zones, with efficiency and crosstalk measurements recorded. (In principle, the vertical zones should be rectangular and of equal sizes, but misalignments, aberrations, and other experimental considerations dictate otherwise.) In both implementations, crosstalk is compensated in the representations of the scaled filters \mathbf{b}_k^\pm . In other words, the SLM blazed grating efficiencies are chosen such that the power reaching each detector is as close as possible to the value that would be expected if there were no crosstalk, and Eqs. (3.11) and (3.12) are still used to compute the half-correlation coefficients $c_{j,k}^\pm$.

A.4 Calibration issues

Despite the extensive efforts to calibrate the SLM precisely, several issues are evident from the calibration data, and these issues likely impact the fidelity of the spectrum recognition operation. For instance, the 5-pixel gratings are generated to have a linear spread of efficiencies, meaning for example that grating numbers 0, 10, 20, and 30 should result in linearly increasing power. However, the 5-pixel grating efficiencies are found to increase super-linearly for low efficiencies and then linearly for high efficiencies. (For comparison, the 2-pixel grating efficiencies increase linearly at first, then sub-linearly.) The reason for this discrepancy is not known, but clearly, the phase model is inadequate in some way (barring errors in implementation of the algorithm).

It is also possible that temperature fluctuations have an impact on calibration. The room temperature is assumed not to fluctuate significantly, and the factor G in Eq. (A.4) is assumed to be constant. This assumption may be incorrect. However, including the effects of temperature in the calibration of the SLM phase response would make the calibration much more difficult, both mathematically and experimentally.

As noted in section 3.4.2, the boundary between the vertical zones is irregular, indicating the presence of optical aberrations, or possibly calibration errors.

While these issues with the SLM calibration are real, and they do reduce the fidelity of the spectrum recognition results, they do not entirely prevent the spectrum recognition operation from working (except possibly for the multiple-zone experiment, as discussed).