HIGHLY DOPED SEMICONDUCTORS FOR PLASMONIC WAVEGUIDES AND
FLAT-COMPOSITE GRATINGS

BY

AARON JOSHUA ROSENBERG

THESIS

Submitted in partial fulfillment of the requirements
for the degree of Master of Science in Electrical and Computer Engineering
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2014

Urbana, Illinois

Adviser:

Assistant Professor Daniel Wasserman

Abstract

Doped semiconductors are investigated as a foundation for plasmonic waveguides and gratings in the mid-infrared. The potential applications of plasmonics to the MIR are reviewed, along with a brief derivation of a mathematical formalism for hybrid plasmonic waveguides which utilize both doped semiconductors and noble metals. A COMSOL Multiphysics model for these waveguides is demonstrated, with a high degree of automation. In addition, I demonstrate lateral control of carrier concentration in doped silicon, by patterning commercially available spin-on dopants down to the subwavelength scale. Samples are characterized by Fourier transform infrared spectroscopy and microscopy, surface profilometry, and infrared emissivity measurements. Samples show strong diffraction from 1D arrays of 'metal' lines patterned using the spin-on dopants.

To Scott Carney and Kanti Jain, who showed me that grad school is more than just grad school.

Acknowledgements

Contents

1 Introduction

Nanophotonics has developed into a vibrant field of research in the past decade, owing to a desire to manipulate light at very small dimensions. We define nanophotonics as the study of the interaction between light and matter on length scales smaller than the wavelength of the light, where the matter responds differently to electromagnetic radiation than it does at larger length scales (i.e. in bulk). The ability to squeeze electromagnetic (EM) energy into tiny volumes holds great potential for sensing and optical integration. For example, today's microelectronic circuitry has shrunk down to feature sizes as low as 22 nm. Meanwhile, optical components such as telecom lasers and fiber optics operate at much higher frequencies, chiefly in the near-IR, but compared to state-of-the-art transistors they are large and cumbersome. Furthermore, traditional dielectric lenses cannot focus light to smaller sizes than about $\lambda/2$, as explained by the diffraction limit of classical optics. There is a strong interest in integrating optical components with electrical circuits for next generation technologies (e.g. through optical interconnects), but this size mismatch has presented a large stumbling block. To circumvent this roadblock, it is necessary to find a way to conquer the diffraction limit, and one of the most promising approaches is *plasmonics.*

By plasmonics, we do not mean "plasma physics" of the sort that describe the behavior of ionized gases (plasmas), but rather the interaction between EM radiation with free electrons inside a solid material [2, 3]. As is well-known from materials science, the outermost valence electrons in a metal crystal are delocalized. They form a "sea" or "plasma" of electrons that are free to move under the influence of some driving oscillation. When a collection of electrons within the plasma oscillates, the concentration of charge oscillates with it, and thus the density of charge within the material may be described by a wave or quasiparticle, which we label a plasmon. By coupling light energy to these plasmons, it is possible to squeeze the energy into subwavelength volumes. For example, at the interface between a metal and a dielectric, these plasmons can couple to a driving electromagnetic wave to create a hybrid excitation known as a surface plasmon polariton (SPP). "Polariton" refers to the coupling to the traveling EM wave. The modifier "surface" indicates that these oscillations exist only at the surface of the metal, as we shall see later. The hybrid excitation is described by a dispersion relation, relating the mode's propagation constant to its oscillation frequency, effectively giving the SPP wavelength, which can be much shorter than the wavelength of a free-space photon with equal frequency.

Plasmonics offers subwavelength control of light, which holds promise for applications as diverse as enhanced sensing [4, 5], waveguides for integrated optical circuits [6], nano-lasers and light emitters [7, 8, 9, 10], and more generally for their ability to squeeze light into small volumes and thus conquer the diffraction limit [11]. Visible and near-infrared (NIR) wavelengths have been the proving ground for much of the pioneering work in plasmonics, due to both the technological

importance of this range and the plasmonic properties of traditional, well understood metals such as gold and silver. However, the noble metals are inherently lossy [12], which has impeded their use in the development of commercial plasmonic devices. Furthermore, they are not well-suited to many plasmonic applications across other regions of the spectrum where subwavelength confinement is also of interest. In particular, the mid-infrared (MIR) has become a wavelength range of increasing importance for sensing and security and defense applications, and is the primary focus of this paper.

The mid-infrared is generally defined as ranging from 2 to 20 μm (or, equivalently, 62 meV to 620 meV or 500 cm^{-1} to 5000 cm^{-1}), and is a key range for many sensing and defense applications such as infrared countermeasures (IRCM) and cloaking [13]. Within the mid-infrared lies the "molecular fingerprint" region, so-called because it contains many strong and unique absorption lines corresponding to the excitation of the rotational and vibrational modes of many molecular species [14]. The mid-IR also contains what is known as the molecular fingerprint region of the spectrum and is widely used by chemists to elucidate the chemical structure of materials. Also, the thermal (blackbody) radiation from most objects peaks in the MIR over a wide range of temperatures, including room temperature. Given these factors, plasmonic devices in the mid-infrared are highly desirable [15]. However, the size mismatch between the wavelengths of the MIR and the nanophotonic or molecular lengthscales are even larger than those in the visible or NIR. In addition, traditional metals are in many ways poorly suited for plasmonics in the mid-IR. In this region, the permittivity of the metals becomes extremely large and negative, causing them to behave like perfect electrical conductors (PECs). While bound surface plasmon polaritons *have* been demonstrated at the interface between dielectrics and noble metals in the mid-IR [16, 17, 18, 19], these modes penetrate deeply into the dielectric, which defeats the original purpose of nanophotonics, namely, to confine the EM fields to molecular and electronic length scales. Also, localized surface plasmon resonances (LSPRs), which produce scattering resonances from, and local field enhancements near, metallic nanoparticles, cannot be excited at wavelengths where the metal permittivity is so large and negative. By contrast, LSPRs require the metal permittivity to be negative but on the same order of magnitude as the surrounding dielectric [1, 20].

To achieve the same subwavelength control of light at mid-IR frequencies that noble metals allow for at visible wavelengths, another method is needed, one that can create and control the metallic sea of electrons. Doped semiconductors have been demonstrated as low-loss plasmonic materials in this range, with optical properties that resemble (and in some ways improve upon) their noble metal counterparts in the visible range [21, 22, 23, 24]. Whereas gold and silver have fixed material properties and therefore fixed permittivities at given wavelengths, doped semiconductors are subject to human control. By modulating the amount of doping in a semiconductor, we can control the number of free electrons and therefore the optical response of the material. This al-

lows us to create "designer metals" which can be optimized for performance over a large range. In the next section, I will review the theory of the optical response of metals, the theory of propagating surface plasmon polaritons and applications to optical circuitry, and the significance of lateral control of permittivity. Following that, I will describe two major research projects which utilize doped semiconductors. The first, described in section 3, is a theoretical framework for designing plasmonic waveguides in the mid-IR, with computational models performed using COMSOL Multiphysics' finite-element-method (FEM) solver. The end goal of these models is to design an optical circuit using the waveguides. The second project, described in section 4, is an experimental demonstration of composite plasmonic materials with lateral control of permittivity. These materials are completely flat and appear so in the visible spectrum to the human eye, yet they diffract light in the mid-IR as though a diffraction grating had been etched into them. These gratings have potential for use as plasmonic couplers, and more generally, the lateral control of permittivity is a vital step towards creating monolithic, integrated nano-optoelectronics. I shall begin by reviewing the fundamentals of the optical response of metals and the theory of propagating surface plasmon polaritons, and some devices that may utilize them.

2 Review of Theory

2.1 Drude Model

The Drude model is a purely classical approximation of the optical response of metals, but performs very well for our doped semiconductors despite being over 100 years old. The Drude model assumes that [25] electron-electron interactions and coulombic electron-ion interactions are negligible. These are called the independent and free electron approximations, respectively. It further claims that scattering occurs from electrons bouncing off fixed, unmoving ions as in the kinetic theory of gases. While this is a gross oversimplification of the true behavior of electrons in a crystal, it still suffices for modeling conductivity. Indeed, the Drude model sweeps all detailed scattering mechanisms under the rug, and claims that the electron experiences collisions at some rate $\Gamma = 1/\tau$, where τ is called the relaxation time or the scattering time. There is no explicit dependence on a particular mechanism here; τ is simply a phenomenological tool. In addition to the scattering time, the Drude model includes two other parameters [1]: a background dielectric constant (ε_s), and a plasma frequency (ω_p). The permittivity is given by:

$$\varepsilon_m(\omega) = \varepsilon_s \left(1 - \frac{\omega_p^2}{\omega^2 + i\omega\gamma} \right) \tag{1}$$

$$\varepsilon_m = \varepsilon_m' + i\varepsilon_m'' = \varepsilon_s \left(1 - \frac{\omega_p^2}{\omega^2 + \gamma^2} \right) + i\varepsilon_s \left(\frac{\frac{\gamma\omega_p^2}{\omega}}{\omega^2 + \gamma^2} \right) \tag{2}$$

$$\omega_p^2 = \frac{ne^2}{\varepsilon_s \varepsilon_0} \tag{3}$$

The imaginary part of the permittivity determines losses in optical signals. The plasma frequency ω_p (or equivalently the plasma wavelength $\lambda_p = 2\pi c/\omega_p$) marks the approximate point where the sign of the real part of the permittivity switches from positive (where the semiconductor behaves as a dielectric), to negative (where it behaves as a metal). As mentioned previously, the plasmonic effects of greatest interest, such as LSPRs and tightly-bound SPPs, occur at wavelengths just longer than λ_p, which can be controlled by doping. These designer plasmonic films can be grown epitaxially, typically using narrow band gap materials with low effective masses, and have been effectively used as plasmonic constituent materials in hyperbolic metamaterial structures [21, 26, 27]. The epitaxially grown doped semiconductors allow for accurate control of plasma wavelength and the ability to control material properties with atomic-scale precision in the growth direction. These materials can be patterned to show localized or propagating modes [28, 29], and have the potential to be integrated with epitaxially grown optoelectronic devices

[30]. Plasmonic effects have also been demonstrated with silicides or doped silicon, in the former case using sequential Si/metal evaporation followed by a post-deposition anneal [31], or for the latter, commercially available doped Si substrates, ion-implanted dopants, or alternatively, spin-on dopants [32, 33, 34]. Doped silicon has certain advantages over epitaxially grown III-V engineered metals, such as the extensive and mature Si fabrication and processing toolset available, as well as the potential for enabling the integration of mid-IR functionality from plasmonic structures, with well-established silicon photonics architectures. However, silicon tends to have a larger scattering rate, and therefore higher losses, than III-V materials and thus may not be ideally suited for all applications. One such application is plasmonic waveguides, in which metallic structures are used to guide light down a certain path, with excellent confinement. If the metal is very lossy, however, the signal will quickly die out as it travels down the guide. Having seen that doped semiconductors can behave as metals in the mid-IR, we now turn to the theory of plasmonic waveguides.

2.2   2D Plasmonic Waveguides

The dispersion relation for an SPP traveling along an infinite metal-dielectric interface is well-documented [2, 35] and shall not be rehashed here. A real waveguide, however, is never an infinite interface. Indeed, the entire purpose of nanophotonics is to confine light to subwavelength volumes, and an infinite planar interface can only perform this task in one direction. Consider the waveguide shown in figure 1. We seek an SPP which propagates either into or out of the page. We expect that such a mode will be confined in the y-direction to the interface between the air and the metal; that is, the fields should reach a maximum at the interface and decay going into both the air and the metal. In addition, the metal/air interface is sandwiched between two perfect electrical conductors (PECs), whose permittivity is taken to be negative infinity (this is a reasonable approximation for gold in this range). Electromagnetic boundary conditions require that the tangential electric field must be zero at the interface with the PEC. This provides confinement in the x-direction. In general, it is quite difficult to solve Maxwell's equations exactly for this problem. For instance, the perpendicular component E_x is not forced to zero (or "shorted", to borrow a term from circuit theory) at the PEC interface. However, we can make a few simplifying assumptions.
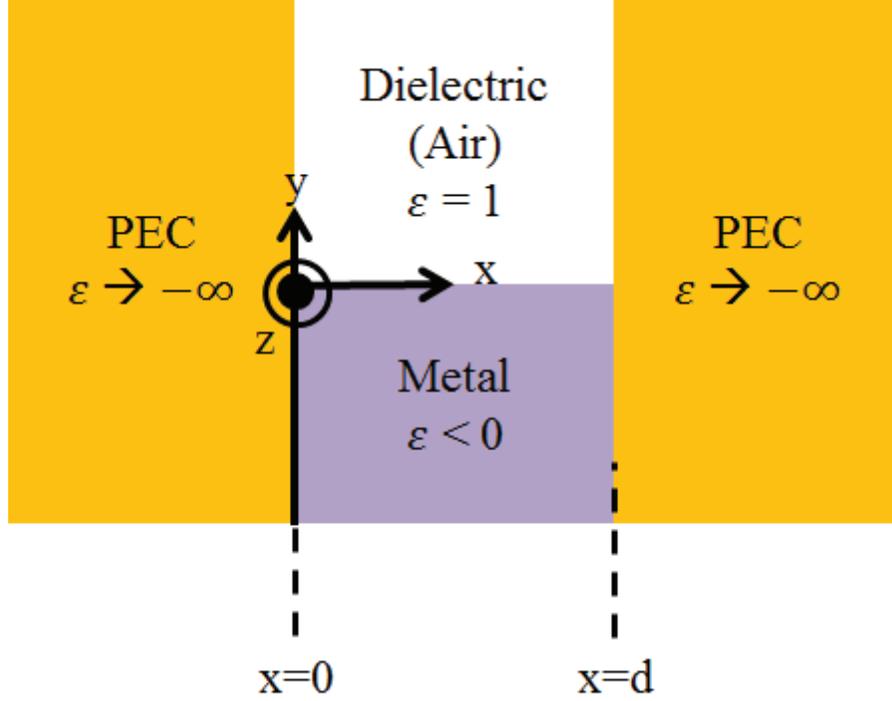
Figure 1: A hybrid plasmonic/PEC waveguide. A surface plasmon polariton can propagate along the metal/air interface if certain conditions are met.

1.) Assume $H_y \approx 0$. This assumption is borrowed from the theory of rectangular waveguides [35].

2.) Assume that all field components have the same z-dependence. Specifically, assume that the field can be described as a plane wave with a factor $e^{i\beta z}$, where β is called the propagation constant. As such, $\partial/\partial z = i\beta$.

3.) Assume that E_y and E_z have an x-dependence of $\sin k_x x$ within the notch and are identically zero elsewhere. To fulfill the continuity boundary condition $E_{||}(x = 0) = E_{||}(x = d) = 0$, this requires $k_x d = m\pi$. The term m can be any integer but for now we will examine only the fundamental mode, with $m = 1$ and $k_x = \pi/d$.

4.) Assume that all fields decay away from the metal/ air interface, i.e. they go as $e^{-\alpha_d y}$ in the dielectric and $e^{\alpha_m y}$ in the metal, where $\alpha_m > 0$ and $\alpha_d > 0$ are to be determined.

However, the remaining components and their magnitudes are unknown. We can only write these components in the following form for now (I will absorb α_d and $-\alpha_m$ into a general, phenomenological α for neatness):

$$E_x = A_x e^{-\alpha y} f_x(x) \cdot e^{i\beta z} \tag{4a}$$

$$E_y = A_y e^{-\alpha y} \sin(k_x x) \cdot e^{i\beta z} \tag{4b}$$

$$E_z = A_z e^{-\alpha y} \sin(k_x x) \cdot e^{i\beta z} \tag{4c}$$

$$H_x = C_x e^{-\alpha y} g_x(x) \cdot e^{i\beta z} \tag{4d}$$

$$H_y = 0 \tag{4e}$$

$$H_z = C_z e^{-\alpha y} g_z(x) \cdot e^{i\beta z} \tag{4f}$$

where the A's and C's are constants, and the f's and g's are unknown functions of x (although we expect them to be sinusoidal). We must avail ourselves of all four of Maxwell's equations:

$$H = \frac{1}{i\omega\mu} \nabla \times E = \frac{1}{i\omega\mu} \left\{ \hat{x} \left(\frac{\partial}{\partial y} E_z - ik_z E_y \right) + \hat{y} \left(ik_z E_x - \frac{\partial}{\partial x} E_z \right) + \hat{z} \left(\frac{\partial}{\partial x} E_y - \frac{\partial}{\partial y} E_x \right) \right\} \tag{5}$$

$$E = \frac{1}{-i\omega\varepsilon} \nabla \times H = \frac{1}{-i\omega\varepsilon} \left\{ \hat{x} \left(\frac{\partial}{\partial y} H_z - ik_z H_y \right) + \hat{y} \left(ik_z H_x - \frac{\partial}{\partial x} H_z \right) + \hat{z} \left(\frac{\partial}{\partial x} H_y - \frac{\partial}{\partial y} H_x \right) \right\} \tag{6}$$

$$\nabla \cdot D = 0 = \frac{\partial}{\partial x} E_x + \frac{\partial}{\partial y} E_y + i\beta E_z = 0 \tag{7}$$

$$\nabla \cdot H = 0 = \frac{\partial}{\partial x} H_x + \frac{\partial}{\partial y} H_y + i\beta H_z = 0 \tag{8}$$

We will begin by using the assumption that $H_y = 0$ to eliminate as many terms as possible. From Faraday's law:

$$H_y = \frac{1}{i\omega\mu} \left(i\beta E_x - \frac{\partial}{\partial x} E_z \right) = 0 \tag{9}$$

$$E_x = \frac{1}{i\beta} \frac{\partial}{\partial x} E_z = \frac{1}{i\beta} A_z e^{-\alpha y} \frac{\partial}{\partial x} \sin(k_x x) \cdot e^{i\beta z} = \frac{A_z k_x}{i\beta} e^{i\beta z - \alpha y} \cos(k_x x) = A_x e^{i\beta z - \alpha y} f_x(x) \tag{10}$$

from which we conclude $f_x(x) = \cos(k_x x)$ and $A_x = A_z k_x / i\beta$. From Ampere's law:

$$E_z = \frac{1}{-i\omega\varepsilon}\left(\frac{\partial}{\partial x}H_y - \frac{\partial}{\partial y}H_x\right) = \frac{1}{i\omega\varepsilon}\frac{\partial}{\partial y}H_x = \frac{1}{i\omega\varepsilon}(-\alpha)C_x e^{-\alpha y}g_x(x)\cdot e^{i\beta z} = A_z e^{-\alpha y}\sin(k_x x)\cdot e^{i\beta z}$$

$$\tag{11}$$

from which we conclude $g_x(x) = \sin(k_x x)$ and $C_x = -i\omega\varepsilon A_z/\alpha$. Ampere's law also gives

$$E_x = \frac{1}{-i\omega\varepsilon}\left(\frac{\partial}{\partial y}H_z - ik_z H_y\right) = \frac{\alpha}{i\omega\varepsilon}H_z = \frac{\alpha}{i\omega\varepsilon}C_z e^{-\alpha y}g_z(x)\cdot e^{i\beta z} = \frac{A_z k_x}{i\beta}e^{i\beta z-\alpha y}\cos(k_x x) \tag{12}$$

so we conclude $g_z(x) = \cos(k_x x)$ and $C_z = A_z k_x \omega\varepsilon/\alpha k_z$. We can check using Gauss's law for magnetic fields:

$$\frac{\partial}{\partial x}H_x = -i\beta H_z \tag{13}$$

$$H_x = \int \frac{-iA_z k_x \omega\varepsilon}{\alpha}e^{i\beta z-\alpha y}\cos(k_x x)\,dx = \frac{-i\omega\varepsilon A_z}{\alpha}e^{i\beta z-\alpha y}\sin(k_x x) \tag{14}$$

The last remaining component to solve for is E_y. Using Gauss's law for electric fields:

$$\frac{\partial}{\partial x}E_x + i\beta E_z = -\frac{\partial}{\partial y}E_y = \frac{\partial}{\partial x}\frac{A_z k_x}{i\beta}e^{i\beta z-\alpha y}\cos(k_x x) + i\beta A_z e^{i\beta z-\alpha y}\sin(k_x x)$$

$$= \frac{-A_z k_x^2}{i\beta}e^{i\beta z-\alpha y}\sin(k_x x) + i\beta A_z e^{i\beta z-\alpha y}\sin(k_x x) = iA_z\left(\frac{k_x^2}{\beta}+\beta\right)e^{i\beta z-\alpha y}\sin(k_x x) \tag{15}$$

$$E_y = -\int \frac{iA_z}{\beta}\left(k_x^2+\beta^2\right)e^{i\beta z-\alpha y}\sin(k_x x)\,dy = \frac{iA_z}{\alpha\beta}\left(k_x^2+\beta^2\right)e^{i\beta z-\alpha y}\sin(k_x x) \tag{16}$$

Or alternatively, using Ampere's Law:

$$E_y = \frac{1}{-i\omega\varepsilon}\left(ik_z H_x - \frac{\partial}{\partial x}H_z\right) = \frac{1}{-i\omega\varepsilon}\left(i\beta\frac{-i\omega\varepsilon E_0}{\alpha}e^{i\beta z-\alpha y}\sin(k_x x) - \frac{\partial}{\partial x}\frac{A_z k_x \omega\varepsilon}{\alpha\beta}e^{i\beta z-\alpha y}\cos(k_x x)\right)$$

$$= \frac{iA_z}{\alpha\beta}\left(\beta^2+k_x^2\right)e^{i\beta z-\alpha y}\sin(k_x x) \tag{17}$$

We now have:

$$E_x = \frac{E_0 k_x}{i\beta}e^{i\beta z-\alpha y}\cos(k_x x) \tag{18a}$$

8

$$E_y = \frac{iE_0}{\alpha\beta} \left(\beta^2 + k_x^2\right) \sin\left(k_x x\right) \tag{18b}$$

$$E_z = E_0 e^{i\beta z - \alpha y} \sin\left(k_x x\right) \tag{18c}$$

$$H_x = \frac{-i\omega\varepsilon E_0}{\alpha} e^{i\beta z - \alpha y} \sin\left(k_x x\right) \tag{18d}$$

$$H_y = 0 \tag{18e}$$

$$H_z = \frac{E_0 k_x \omega\varepsilon}{\alpha\beta} e^{i\beta z - \alpha y} \cos\left(k_x x\right) \tag{18f}$$

and the dispersion relation is obtained from the wave equation:

$$\nabla^2 E + \omega^2 \varepsilon\mu E = 0 \tag{19}$$

in the dielectric region this gives:

$$\left(\alpha^2 E_z - k_x^2 E_z - \beta^2 E_z + \omega^2 \varepsilon\mu E_z\right) = 0 \tag{20}$$

and we can use a similar argument in the metal region. The dispersion relations are:

$$-\alpha_d^2 + k_x^2 + \beta^2 = \omega^2 \varepsilon_d \varepsilon_0 \mu_0 \tag{21a}$$

$$-\alpha_m^2 + k_x^2 + \beta^2 = \omega^2 \varepsilon_m \varepsilon_0 \mu_0 \tag{21b}$$

in the dielectric and metal respectively. Now, we still have the continuity condition available. Evaluating H_x at y=0 on both sides of the interface gives us:

$$\frac{\varepsilon_d}{\alpha_d} = \frac{\varepsilon_m}{-\alpha_m} \tag{22}$$

as before. Substituting eq. (22) into eq. (21b) gives us:

$$-\left(\frac{\varepsilon_m \alpha_d}{\varepsilon_d}\right)^2 + k_x^2 + \beta^2 = \omega^2 \varepsilon_m \varepsilon_0 \mu_0 \tag{23}$$

Subtracting eq. (21a) from this, we obtain:

9

$$\alpha_d^2 \left(1 - \frac{\varepsilon_m^2}{\varepsilon_d^2} \right) = \omega^2 \left(\varepsilon_m - \varepsilon_d \right) \varepsilon_0 \mu_0 \tag{24}$$

From here the other wavevector components can be obtained. A sample plot of the dispersion relation and the fields E_x and E_z is shown in figure 2.
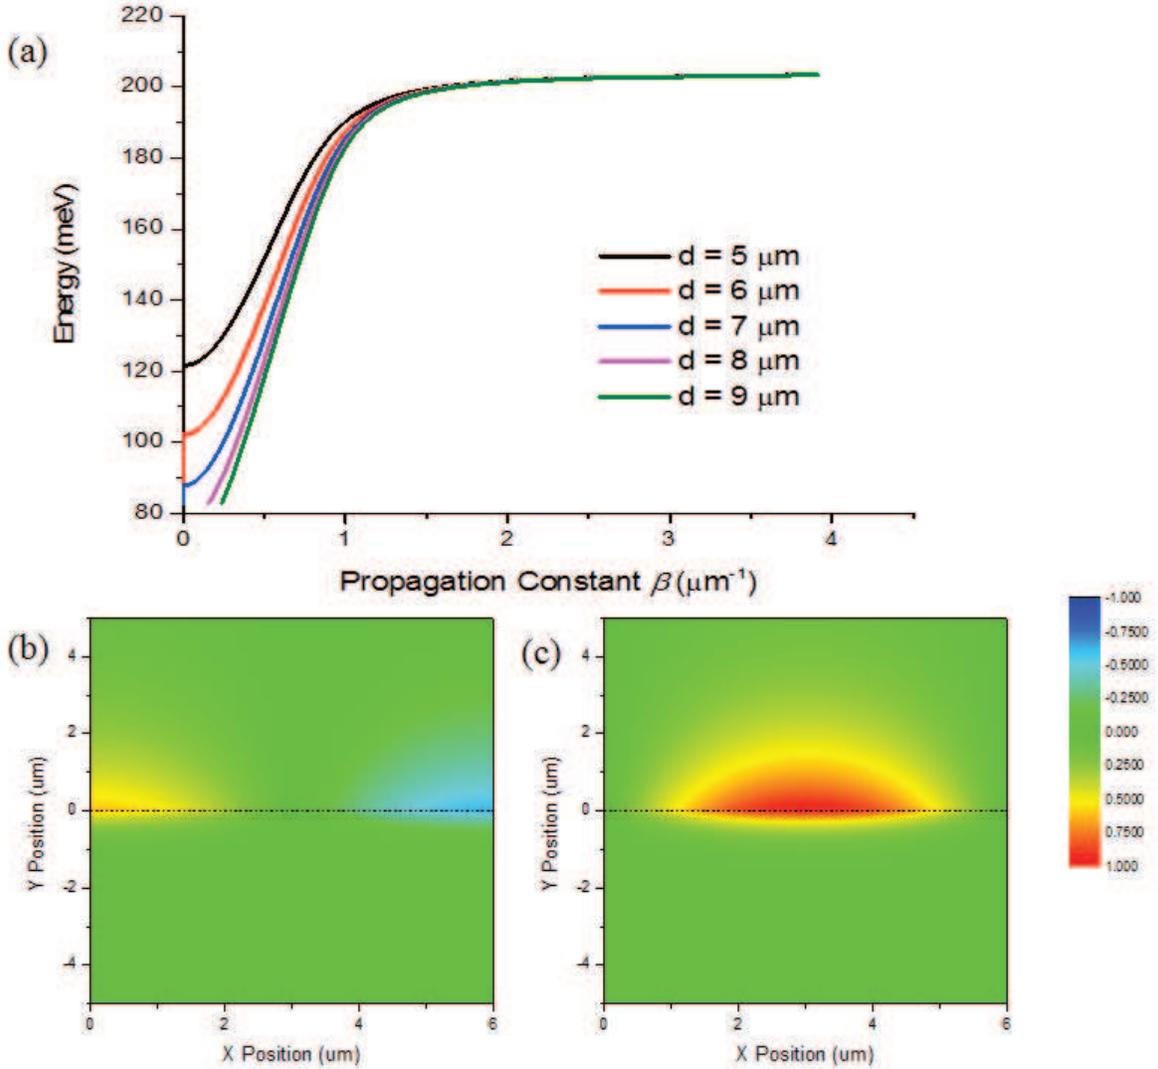

Figure 2: (a) SPP dispersion relation for a waveguide like the one in figure 1, with the metal taken to be doped indium arsenide ($\lambda_p = 5.85$ μm, $\varepsilon_\infty = 12.3$, and assuming no losses such that $\Gamma = 0$). (b) E_x and (c) E_z at a wavelength of 7 μm, plotted using eq. (18). Positive values of position Y are in the air, while negative values are in the metal. E_0 is taken to be 1 V/m.

Finally, we define the effective index $n_{eff} = \beta / k_0$, where k_0 is the wavevector in free space,

that is, $2\pi/\lambda_0$. An ideal simulation of this waveguide would extend infinitely in the vertical direction. However, we can make yet another approximation. For a strongly confined SPP, the field rapidly dies away as one moves into the airgap. If the decay constant is large enough, then the traveling mode should be unable to "see" any perturbations at an appreciable distance away from the interface. Thus, if we put another perfect electrical conductor some distance away from the interface (see figure 3), it should not affect the properties of the SPP at all, and yet it will spare us the inherent challenges associated with infinite domains and perfectly matched layers (PMLs). I will utilize this approximation for COMSOL simulations of plasmonic structures.

According to this theory, the metal-dielectric interface can always sustain a propagating mode as long as the real part of the permittivity of the metal remains negative, but larger in magnitude than the dielectric permittivity. However, as the permittivity, or rather the propagation constant, approaches zero, the waveform takes on some interesting attributes which we shall next consider.

2.3   Optical Circuitry

It is well-known that many analogues of electrical phenomena and technology can be found in optics. Dielectric waveguides confine light as a quantum well confines electrons. Photonic crystals with a periodic refractive index give rise to forbidden photonic bandgaps just as an atomic crystal with periodic Coulombic potential gives rise to an electronic bandgap. Another potential manifestation of this duality is the *optical circuit*: a device which manipulates information as an electrical circuit does, except that it handles the information by controlling the flow of light, rather than the control of electrons. In the framework proposed by Engheta [6], such a circuit would consist of nanoparticles, each with deeply subwavelength dimensions, embedded in an optical waveguide. Such nanoparticles can mimic the "lumped" circuit elements familiar from electronic circuit theory depending on their properties. A nanoparticle with positive real permittivity will behave like a capacitor, one with negative real permittivity will behave like an inductor, and one with positive imaginary permittivity (as is almost always the case) will behave like a resistor. This "lumped" circuit model only holds if the particles are much smaller than the effective wavelength of the light. But as seen in the previous section, we can tweak the dimensions and material properties of a plasmonic waveguide to obtain an extremely small propagation constant, and therefore an extremely large effective wavelength. This, combined with the subwavelength confinement available, makes plasmonic waveguides an attractive foundation for optical circuitry.

2.4 Lateral Control of Permittivity

Almost all realizations of plasmonic devices or structures require some modulation of the plasmonic surface or patterning of the plasmonic material. For the former, the modulation of the surface is often required to match the momentum a bound plasmonic mode to free space photons, while for the latter, the patterned plasmonic material gives either the LSPR response, or is used to engineer the effective permittivity of a composite metamaterial. Thus, the optical properties of the material system result from either (or both) the geometry of the plasmonic material, or the geometry and relative composition of the composite. There have been significant recent efforts in the field of "flat optics" or "metasurfaces" [36, 37, 38, 39, 40], single layer metamaterials of ultra-subwavelength thickness capable of extraordinary manipulation and control of incident light. However, despite their description as "flat" metasurfaces, in fact, these novel structures require deposition of subwavelength metallic structures (of finite thickness) on dielectric substrates. For epitaxially-grown doped semiconductor plasmonic materials, we are able to vertically control the doping profile of the material. However, it is often required, especially for the aforementioned metasurface work, that the composite material have lateral variations in local permittivity. One method of achieving this is with spin-on dopants, which can themselves be patterned so that the dopants only selectively diffuse into certain areas of the substrate. By patterning a grating in the dopant and then performing a diffusion bake, we can introduce a *periodicity* to the permittivity of the substrate. The resulting will thus have varying permittivity like a grating, and behave like one, even though the surface is itself still flat. My group's experimental work in this area is detailed in section 4. For now, we shall put the flat gratings aside and return to the modeling of plasmonic waveguides.

3 Simulations of Plasmonic Waveguides

3.1 Plasmonic Waveguides: Procedure and Philosophy

In practice, a plasmonic waveguide like the one shown in figure 1 is quite hard to fabricate. The PEC sidewalls must be several microns tall to get good confinement, and the doped semiconductor (designer metal) has a finite thickness, limited by dopant diffusion or epitaxial growth constraints. However, my research group has successfully fabricated a structure with many similarities, and it is this structure (shown in figure 3(a) that I shall model. The waveguide structure consists of a gold notch, covered with an epitaxial film of highly-doped InAs and flipped upside down. The doped InAs was prepared by molecular beam epitaxy on a GaAs semi-insulating substrate, with a 0.5 μm-thick buffer layer of undoped InAs in between to relieve strain. In this model, the gold is taken to be a perfect electrical conductor. Technically, the GaAs and both layers of InAs extend infinitely in the x direction (side to side). The traditional approach towards dealing with infinite domains like this is to use perfectly-matched layers (PMLs). However, for these simulations we will assume that the film can be modeled as if it were surrounded by a PEC on all sides. This is admittedly a rather artificial approximation. However, consider this: the mode we seek is an SPP, with most of the electromagnetic power concentrated at the InAs-air interface just inside the gold cavity. As such, we expect that the fields will decay very rapidly in the lateral direction as one moves away from the air box, even inside the epitaxial film (that is, for y < 0). If the fields do indeed decay this fast, then there is nothing to gain by using PMLs to model an infinitely long film. From the SPP's point of view, it makes little difference whether the epitaxial film is a semi-infinite wafer or a thin slab of thickness 4.7 μm surrounded by PECs on all sides. The other major difference between this model and the ideal one from section 2.2 is that we have added a gold ceiling above, so the air column is no longer infinite. Electromagnetic boundary conditions require that $E_x = E_z = 0$ at this ceiling. Again, however, we note that most of the energy in an SPP is confined to the metal/air interface (or in this case, the doped InAs/air interface). If the air box is sufficiently tall, the SPP should not even "know" that the gold ceiling is there, and so our equations from section 2.2 should still hold reasonably well. If the air box is very short, the model will begin to fail, but it remains a useful approximation and a helpful phenomenological tool.
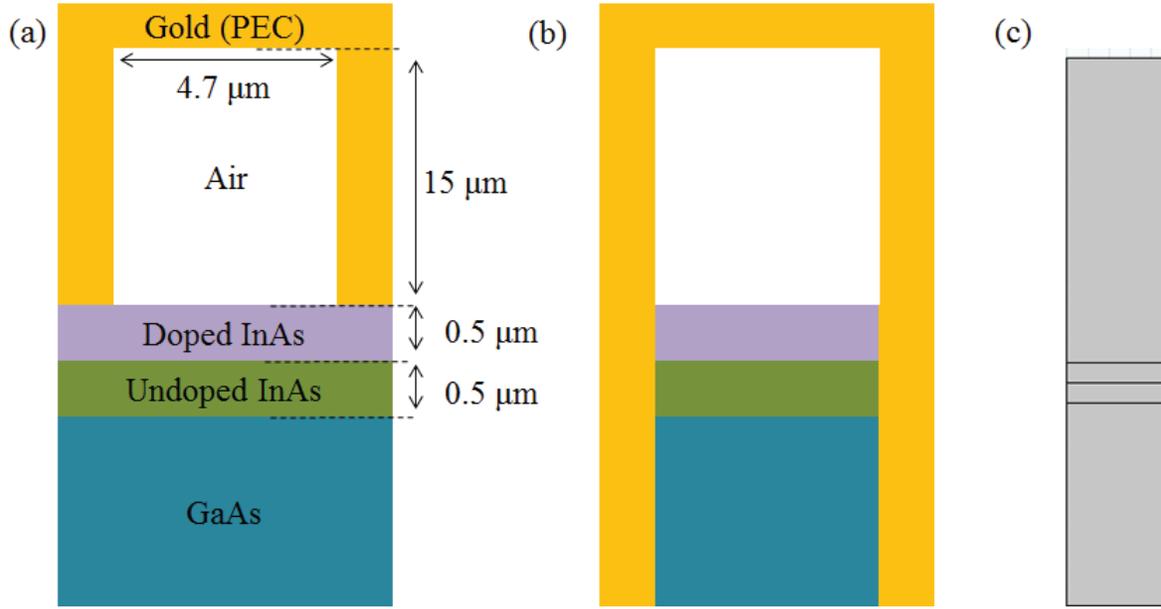
Figure 3: Simulation of a hybrid plasmonic waveguide made using gold (a PEC) and an epitaxial, highly doped InAs film. (a) Schematic of the structure as fabricated. The epitaxial film extends infinitely in the lateral direction. (b) Schematic of the structure as modeled in COMSOL. Extending the PEC boundaries around the sides is a reasonable approximation and obviates the need for PMLs. The boundary conditions at the bottom are set to "scattering." (c) Actual computational geometry in COMSOL, with air as the top layer, then doped InAs, then the undoped InAs buffer layer, then the GaAs substrate. The PECs are accounted for at the boundaries. The left edge is taken to be x=0, and the air / doped InAs interface is taken to be y=0.

To find the SPP, COMSOL's Mode Analysis study from the RF branch is used. After setting up the geometry, the user specifies a frequency of light to excite within the structure, and then COMSOL solves Maxwell's equations numerically to find allowed modes and generate plots of the E and H fields, and the power flow. The user provides an initial guess for the effective index of the mode and the software tries to find the closest matches. However, COMSOL cannot itself recognize whether a mode is a bound SPP or a radiative cavity mode, and in fact, it often computes "junk modes" which satisfy the numerical constraints, but upon inspection are clearly nonphysical. These junk modes present a significant obstacle when one wishes to track the behavior of a true physical mode across a wide range of parameters. For a given set of parameters, one must often instruct COMSOL to search for as many as 5 or 10 modes to find the one of interest. For a large array of inputs, wading through all those modes becomes extremely time-consuming. However, we can reduce this problem by use of the MATLAB Livelink. Consider this: if the geometry of the waveguide is only altered slightly, the effective index of the mode should change only slightly. For example, if we simulate a waveguide with a width of 4 μm and a height of 4 μm, and search for a large number of modes, we can use the field and power plots to find the mode that represents

14

a surface plasmon polariton. We will have only a vague guess as to what the effective index is until this simulation is complete. If we next change the width of our waveguide from 4 μm to 4.01 μm, but take the value of the effective index we found on the previous trial, and input *that* as our initial guess, we can assume that the first mode COMSOL finds for this new geometry will also correspond to an SPP; one with almost the same effective index, in fact. As such, we only need to search for one mode on this second trial and beyond. We can then change the width from 4.01 μm to 4.02 μm, and for our guess use the effective index we obtained for the 4.01 μm case, and so forth. The MATLAB Livelink allows us to change the initial guess in this manner every cycle.

The Livelink also provides insurance against junk modes. When using the GUI, one must wait for COMSOL to finish simulating every combination of input parameters before examining the results of any of them. This is especially problematic because sometimes COMSOL will "hop" to a different mode (e.g. from a bound SPP to a radiative mode or a junk mode), and there is no way to tell until the entire simulation finishes. However, by running COMSOL through the MATLAB Livelink, we can load plots for each trial as it is run. If COMSOL hops to a junk mode, the user can see, and immediately terminate the simulation. However, by some judicious use of save functions, we can retain all the data from the good simulations that were already performed on that run. MATLAB also allows us to save simulation results directly without having to go through the COMSOL GUI's export command. But again, a good initial guess is paramount to make this work efficiently, and so the target mode for the first set of input parameters is always found using the GUI.

To simulate these waveguides, the effective index of a surface plasmon was found for one combination of height and width using the GUI, and then this mode was tracked through MATLAB for different combinations of input parameters. The InAs was modeled as a Drude metal with $\omega_p = 5.85 \times 10^{14} s^{-1}$ and $\gamma = 3 \times 10^{13} s^{-1}$. The frequency of light was chosen to match the lasers available in our laboratory. If COMSOL hopped to a different mode, the increments in the input parameters were decreased until this problem was resolved. Both the real and imaginary parts of the effective index, which determine effective wavelength and losses respectively, were plotted to find the best available combination.

3.2 Results and Discussion

In practice, making a waveguide with a gold notch as tall as the one in figure 3(a) is quite difficult. However, we can still use this geometry as a test of the approximations made in section 2.2. Figure 4 plots the field profiles for an SPP in this structure. The wavelength of light is 9.5 μm, corresponding to a CO2 laser. The fields follow many of the same trends as those predicted by eq. (18), even

though the air box now has a finite height and the doped semiconductor a finite thickness. Specifically, E_x and H_z both have a cosine-like dependence on x, while E_y, E_z, and H_x have a sine-like dependence. Also, H_y is virtually zero. The fields display the hallmark of SPPs: their maxima lie at the metal/air interface and decay in both the positive and negative y directions away from the interface. The effective index of the mode as computed by COMSOL is 0.187, which corresponds to $\beta = 0.1236$, which matches the value of $\beta = 0.12$ computed using eqs. (21-25). These results justify the assumptions we made when deriving the equations in section 2.2. Furthermore, they suggest that the presence of the gold ceiling causes only a minimal disturbance in the behavior and effective index of the SPP, as we expected.
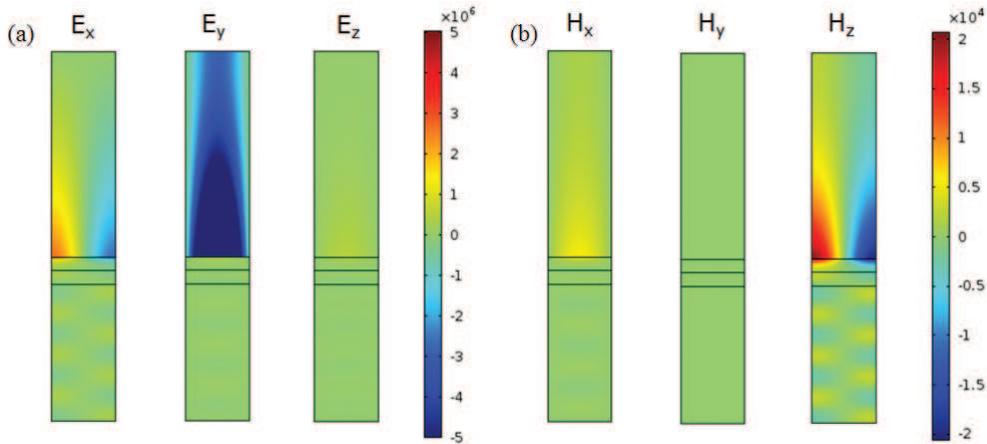

Figure 4: (a) Electric and (b) magnetic field profiles at $\lambda = 9.5$ µm for the structure shown in figure 3, in V/m and A/m, respectively. All field profiles have the x- and y- dependence predicted in section 2.2.

Figure 5 displays results from additional simulations of the same structure, but for a more physically feasible range of heights. Notably, the effective index changes much more rapidly with waveguide width than with height. We expect that this is due to the strong confinement of energy at the air-InAs interface: because all the light is at the interface, it barely even "knows" that the gold ceiling is a few microns above it, and so the height has little impact. However, the width of the waveguide directly determines the transverse wavevector component k_x from eq. (18). As the width of the waveguide increases, less of the light energy gets pushed into the lossy substrate, and as a result, the losses decrease. But this also implies a reduction in confinement. This is emblematic of the major challenge to plasmonics today: the trade-off between loss and size. And in the context of optical circuitry, we can infer a corresponding tradeoff between loss and effective wavelength. For an optical circuit, we seek a mode that has both a very long effective wavelength (i.e. small real part of effective index) and very low losses (i.e. small imaginary part of effective index). Both the real and imaginary parts of n_{eff} are equal to about 0.08 for a width of 4.565 µm and a height of 5.6

µm. In theory, it should be possible to excite an SPP in such a waveguide with a very long effective wavelength yet still observe it at the far end. In many cases, it may prove difficult to fabricate a waveguide to these exact specifications. However, a tunable laser could solve this problem: instead of changing the dimensions of the cavity to suit the wavelength, we can change the wavelength of the laser to get the low-effective index mode we seek, as demonstrated in figure 5(c-d).
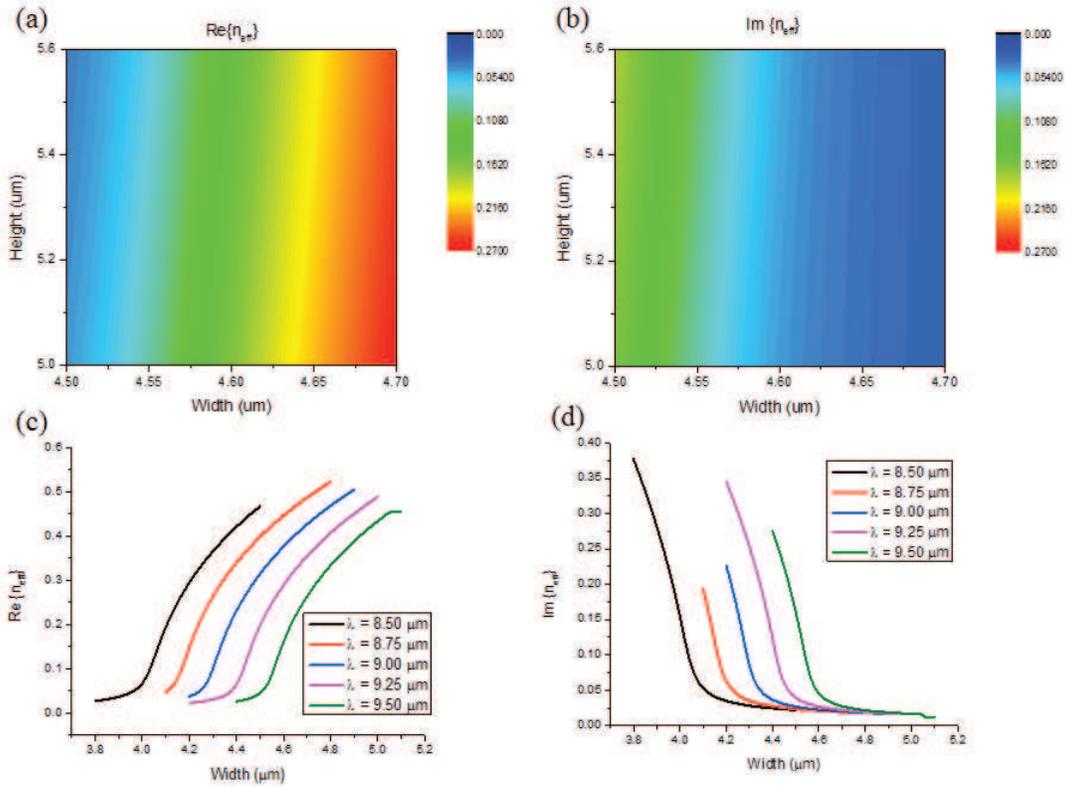

Figure 5: Additional COMSOL simulation results for the structure shown in figure 3, but with revised dimensions. (a) Real and (b) imaginary part of effective index of the fundamental SPP mode at $\lambda = 9.5$ µm as a function of waveguide width and height of the airbox. The ideal mode for ENZ optical circuitry is one with both the real and imaginary parts very small. (c) Real and (d) imaginary part of effective index for the same mode as a function of wavelength and width, for a fixed airbox height of 5 µm.

4 Flat Gratings

4.1 Procedure

As mentioned above, our "flat gratings" are pieces of silicon with periodic doping concentration, and hence periodic permittivity, in the lateral direction. Our initial efforts sought to characterize the spin-doping of silicon substrates, absent any effort to pattern the dopants [1]. This exercise was vital in order to determine the minimum time and temperature required to highly dope the silicon substrate. Because our ultimate goal is a laterally-varying composite metal/dielectric material, we must ensure that we are able to turn the Si metallic, but at the same time, we must limit lateral diffusion of dopants in order to prevent blurring of the lateral variation in permittivity. Samples of undoped silicon are first briefly etched in a buffered oxide etch (BOE) to remove native oxide, cleaned and baked dry. A Filmtronics spin-on dopant with phosphorus concentration of 4.00% was applied to the samples and spun at 5000 RPM for 30s. Next, the samples were baked at 200ºC for 15 minutes to harden the dopant into a solid film. The samples were then inserted into a furnace and baked for varying amounts of time (15 min, 22 min, 30 min, and 38 min) at 900°C, with 2 L/min of oxygen and 7 L/min of nitrogen consistently supplied. After removal from the furnace, the dopant on the surface was etched away in 49% HF. Following bake-in and the subsequent removal of the dopant via a BOE etch, samples were characterized by IR reflection microscopy using a Bruker IRII mid-IR microscope coupled to a Bruker v80v Fourier transform infrared (FTIR) spectrometer. Data were taken across the mid-spectral range from 650 to 7500 cm^{-1}, at a resolution of 16 cm^{-1}. It was determined that 30 minutes was a sufficient bake time to dope the silicon substrate with enough free carriers so that the surface would behave (optically) as a metal at long wavelengths. Following characterization of the uniformly doped substrates, the doping process was modified in order allow for lateral control of permittivity by control of doping concentration. Dopant was spun on to three silicon samples and baked following the process described above. Following the dopant bake, a layer of photoresist was spun over the baked dopant and patterned using standard UV lithography. The samples were then etched, using the patterned photoresist as an etch-mask, with a diluted 1:20 solution of BOE:Water to remove the unprotected dopant. The geometry of the doping patterns fabricated is shown in figure 6(d), and consists of a large, solid rectangular block on one side and a series of grating stripes on the other. Both the solid doped and grating regions were approximately 5mm×10mm in area, for a total sample size of 10 mm × 10 mm. Grating periods of 10, 20 and 30 μm were patterned, and all grating patterns had 50% duty cycle (the width of the unetched dopant stripes was half the period of the grating). To prevent undercutting of the dopant under the resist strips, samples were only dipped in the etching solution for one or two seconds at a time, and then immediately rinsed with water. This process was repeated until the bare silicon surface was

exposed. Next, the photoresist was stripped with acetone and isopropanol, and the samples were blow-dried. Samples were then baked in the furnace for 30 minutes at 900ºC. After the bake, the remaining dopant was etched away in BOE as before. To confirm that silicon surface remained flat following processing, samples were analyzed in an Alpha Step IQ Profilometer.
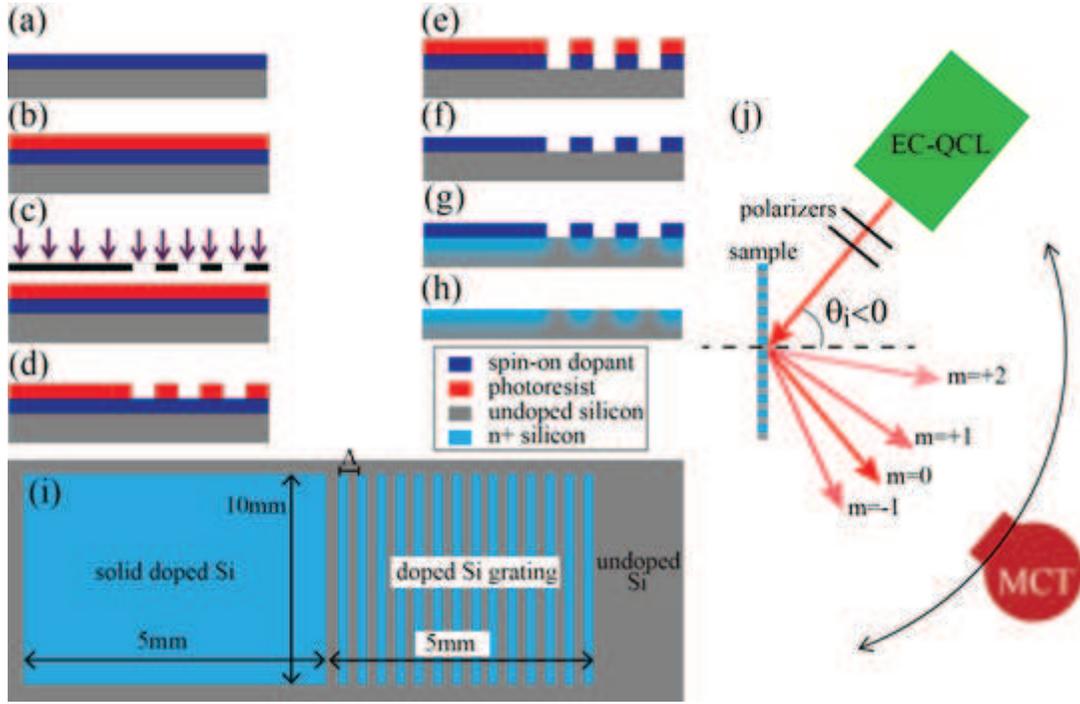

Figure 6: (a-h) Schematic of fabrication process for patterned doped Si surface with legend indicating the spin-on dopant, photoresist, undoped Si, and doped Si color scheme. (i) Top-view schematic of final grating sample showing solid doped, grating, and undoped portions of the surface. (j) Experimental set-up for measurement of diffraction from patterned Si surface [1].

The frequency-dependent reflection of the doped silicon was modeled using a multilayer transfer matrix method (TMM). Such an approach is trivial for a uniformly doped semiconductor sample; however, the sample we are considering will have a vertical variation in material parameters resulting from the diffusion of phosphorus dopants into the Si. Thus our permittivity profile will follow from a calculation of the dopant diffusion into the sample. The profile of dopants was modeled assuming a simple fixed surface concentration of dopants in the spin-on glass. The expression for the concentration as a function of position and time ($C(z,t)$), and the dopant dose ($Q(t)$) are then given by

$$C(z,t) = C_s erfc\left(\frac{z}{2\sqrt{Dt}}\right) \tag{25}$$

19

$$Q(t) = \int_0^\infty C(z,t)\,dz = \frac{2}{\sqrt{\pi}} C(0,t)\sqrt{Dt} \tag{26}$$

respectively, where $C_s = C(0,t)$ is the (assumed) constant surface concentration of dopants, D is the dopant diffusion coefficient, and Dt is the characteristic diffusion length. Knowing the dopant concentration as a function of position, and assuming full ionization, allows us to calculate the permittivity of the Si as a function of position, assuming a doping-independent scattering rate (γ). The doped Si is then divided into thin layers ($\Delta z < \lambda_o/100$) with each layer assigned a frequency-dependent permittivity derived from eqs. (25) and (26). The adjustable parameters in our model are the surface doping concentration (C_s), the scattering rate (γ), and the characteristic diffusion length (Dt). Figure 7(a) shows the reflection spectra from samples baked at 900ºC for 15, 22, 30, and 38 minutes, as well as the modeled reflectivity for the 30 minute sample. By matching the optical properties of the doped samples, we are able to extract dopant concentration as a function of depth (figure 7(b)), as well as a depth and frequency-dependent permittivity for the doped sample (figure 7(c,d)).
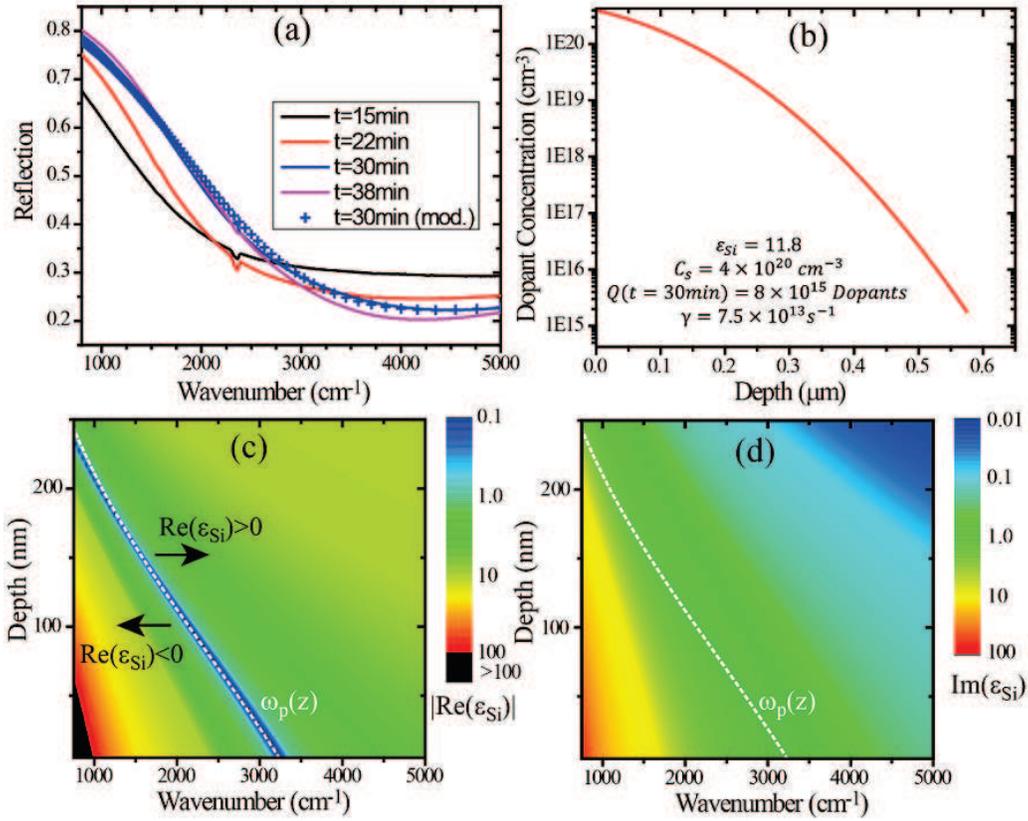
Figure 7: (a) Reflection spectra for doped silicon samples with different bake times, including a TMM model of the 30 minute bake sample. (b) Modeled dopant concentration profile as a function of depth and (c) absolute value of the real part of the permittivity and (d) the imaginary part of the permittivity as a function of depth and wavenumber for the 30 minute bake sample. Both (c) and (d) show the calculated plasma frequency (white dashed line) as a function of depth. In (c), the area to the left of the white dashed line has $Re(\varepsilon_{Si}) < 0$, while the area to the right has $Re(\varepsilon_{Si}) > 0$ [1].

Patterned samples were also measured using the above-described reflection spectroscopy set-up, with reflection data taken from an area that had been covered with a solid, rectangular block of dopant during diffusion, as seen in figure 6(i). In order to accurately resolve reflection in the grating area as a function of both position and wavelength, spatially-resolved reflection data were taken from the 30 μm grating sample using a PerkinElmer Spotlight FTIR microscope coupled to a HgCdTe detector linear array. However, the diffraction-limited optics of the microscope and linear array can give a spatial resolution of only 6.25 μm, leading to the appearance of a weaker spatial modulation of permittivity, an effective "blurring" of our spatially-resolved reflection. In order to improve the spatial resolution of our characterization, we also performed IR-thermal imaging of our grating samples using an InfraScope II thermal imaging system. In this set-up the sample is heated to 80ºC, and thermal emission is collected, through a 15x mid-IR objective, by a 256 x 256

21

element InSb focal plane array (FPA), across a wavelength range from 2-5.5 μm. Though such a set-up cannot give spectral information, the difference between the reflectivity (and thus absorption and emissivity) of the doped Si material, compared to the undoped region, is significant enough to produce a marked difference in the thermal emission. The spatial dependence of this emission can be used to determine the diffusion of dopants in the sample fabrication process. Finally, in order to further characterize the quality of the patterned material, we investigated diffraction from patterned doped samples of different grating periodicities. The experimental set-up is shown in figure 6(j) where the patterned material is mounted on a motorized rotation stage. A Daylight Solutions external cavity tunable quantum cascade laser was used as a tunable infrared light source. Laser light was incident on the sample at 50º with respect to normal. A polarizer pair was placed in front of the QCL in order to control the polarization of the incident light. A HgCdTe detector was mounted on a rotational arm, and rotated around the incidence point of the laser on the sample. Diffracted light intensity as a function of detector angle was measured from each sample (10, 20, 30 μm). For each sample, diffraction was measured for three laser wavelengths (λ = 8.5, 9, 9.5 μm). A simple prediction of the diffraction from the samples can be obtained using the grating equation,

$$\Lambda \left(\sin \theta_i + \sin \theta_m \right) = m\lambda_0 \tag{27}$$

$$m = 0, \pm 1, \pm 2, ...$$

where Λ is the periodicity of the grating, ϑ_i is the angle of the incident light from normal (in the configuration used θ_i = -50°), θ_m is the angle of the mth order diffracted mode (from normal), and λ_o the wavelength of the incident light. Using Eq. 3, we also calculated the expected angles of the first and second order diffraction peaks for the three different wavelengths used (8.5 μm, 9 μm, 9.5 μm).

4.2   Results and Discussion

As mentioned above, figure 7 shows the reflection spectra from our spin-doped Si as a function of bake time. As expected, the silicon becomes more metallic (more reflective at long wavelengths) with increased bake time. However, the reflectivity spectra do not appear to change significantly after 30 minutes of bake. We attribute this to both the finite concentration of phosphorus atoms within the spin-on dopant, as well as the limited penetration depth of light into our doped material. For the latter, once the surface becomes highly reflective over a depth approximately equivalent to the optical penetration depth, the addition of dopants (and thus carriers) deeper in the sample does

little to change the material's reflectivity. In order to obtain good lateral control of permittivity, one must optimize the metallic nature of the doped material, while at the same time minimizing the lateral diffusion of dopants into areas where the film had been etched away. From figure 7, it was determined that 30 minutes of bake time would result in a sufficiently metallic silicon surface layer [$Re(\varepsilon_{Si}) < 0$], for the spectral range of our laser] while at the same time causing minimal lateral diffusion. As can be seen from figure 7 (b), the penetration depth of our dopants is well less than a micron, suggesting that even an equivalent lateral diffusion would still allow for the fabrication of subwavelength "metallic" features. It is likely that spin-on dopants with higher doping concentrations, longer diffusion times, or higher bake temperatures would result in more "optically metallic" Si, while still maintaining a distinct lateral modulation of permittivity. The limits of the technique presented here is the subject of ongoing research. Measurement of diffracted light from our patterned grating samples (figure 8) shows strong and distinct diffraction orders. The peak position of the diffracted orders follows the diffraction angles calculated using the grating equation given in eq. (27) for the $m = 1$ and $m = 2$ diffraction orders as a function of grating period and laser wavelength, as shown in figure 8(d). For the $\Lambda = 10\ \mu$m sample, only the $m = 1$ diffracted order is detected in our experimental set-up, while for $\Lambda = 20\ \mu$m and $\Lambda = 30\ \mu$m samples, strong $m = 1$ and $m = 2$ diffraction is seen. The detector response plotted in figure 3 is normalized to the intensity of the specular reflection peak for each scan, in order to negate the frequency dependent response of the detector and emission intensity of the laser. Our results show strong diffraction for all samples, with some diffraction peaks showing intensities ½ that of the specular reflection peak. This fact, along with the sharp and narrow diffraction peaks, suggest a distinct grating pattern from our laterally patterned doped Si.
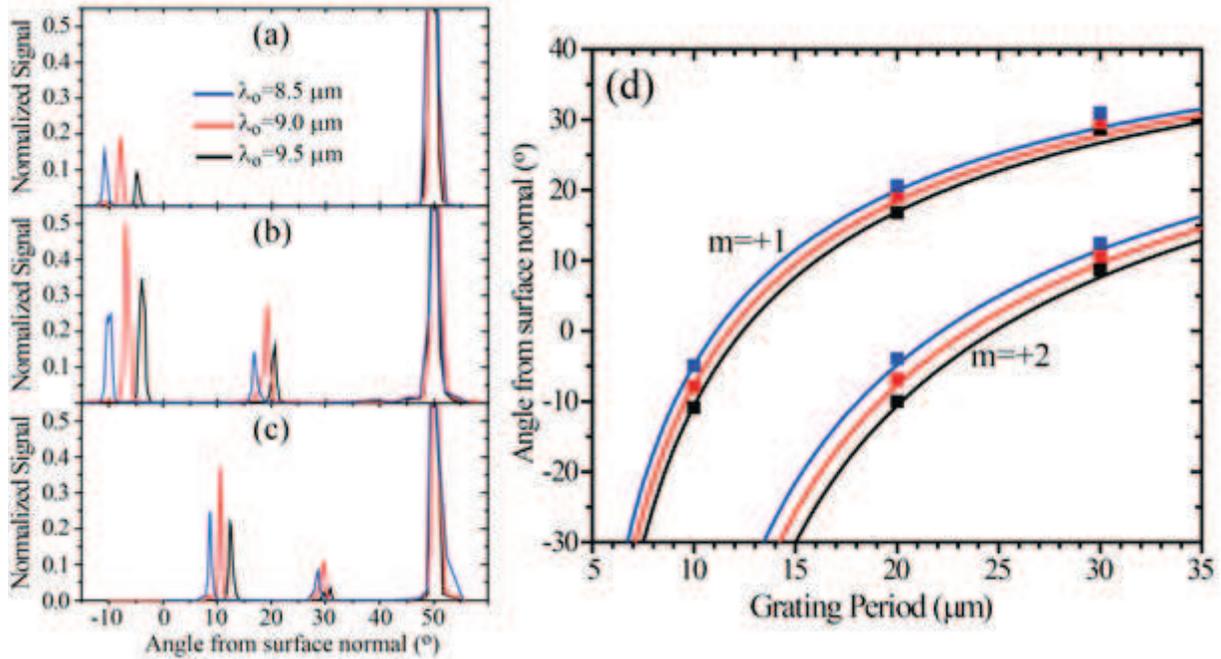
Figure 8: Diffraction from flat gratings with period (a) 10 μm, (b) 20 μm and (c) 30 μm for laser wavelength λ = 8.5 μm (blue), λ = 9 μm (red) and λ = 9.5 μm (black). For the data shown, the incident laser light is polarized with such that the incident magnetic field is parallel to the grating lines. All data is normalized such that the specular reflection peak at 50° has an intensity of 1 for each scan. As can be seen from this data, there is significant coupling into both the m = 1 and m = 2 diffraction orders. (d) Comparison of experimental results with grating equation. Squares show data points from (a-c), solid lines show predicted diffraction angles for the m = 1 and 2 diffraction orders for laser wavelengths λ = 8.5 μm (blue), λ = 9 μm (red) and λ = 9.5 μm (black) as a function of grating period [1].

Diffraction can result, however, not only from a lateral variation in permittivity, but also from a varying surface profile of the sample. In order to demonstrate that our diffraction is in fact caused by the former, we took surface profilometry data of all samples, as shown in figure 9(a). This data confirms that each sample surface is flat, with average roughness on the scale of a few nanometers, limited by the sensitivity of our experimental apparatus. These results clearly show that the optical properties of the samples are a result of the lateral control of Si permittivity, and not a fabrication-induced surface topography.
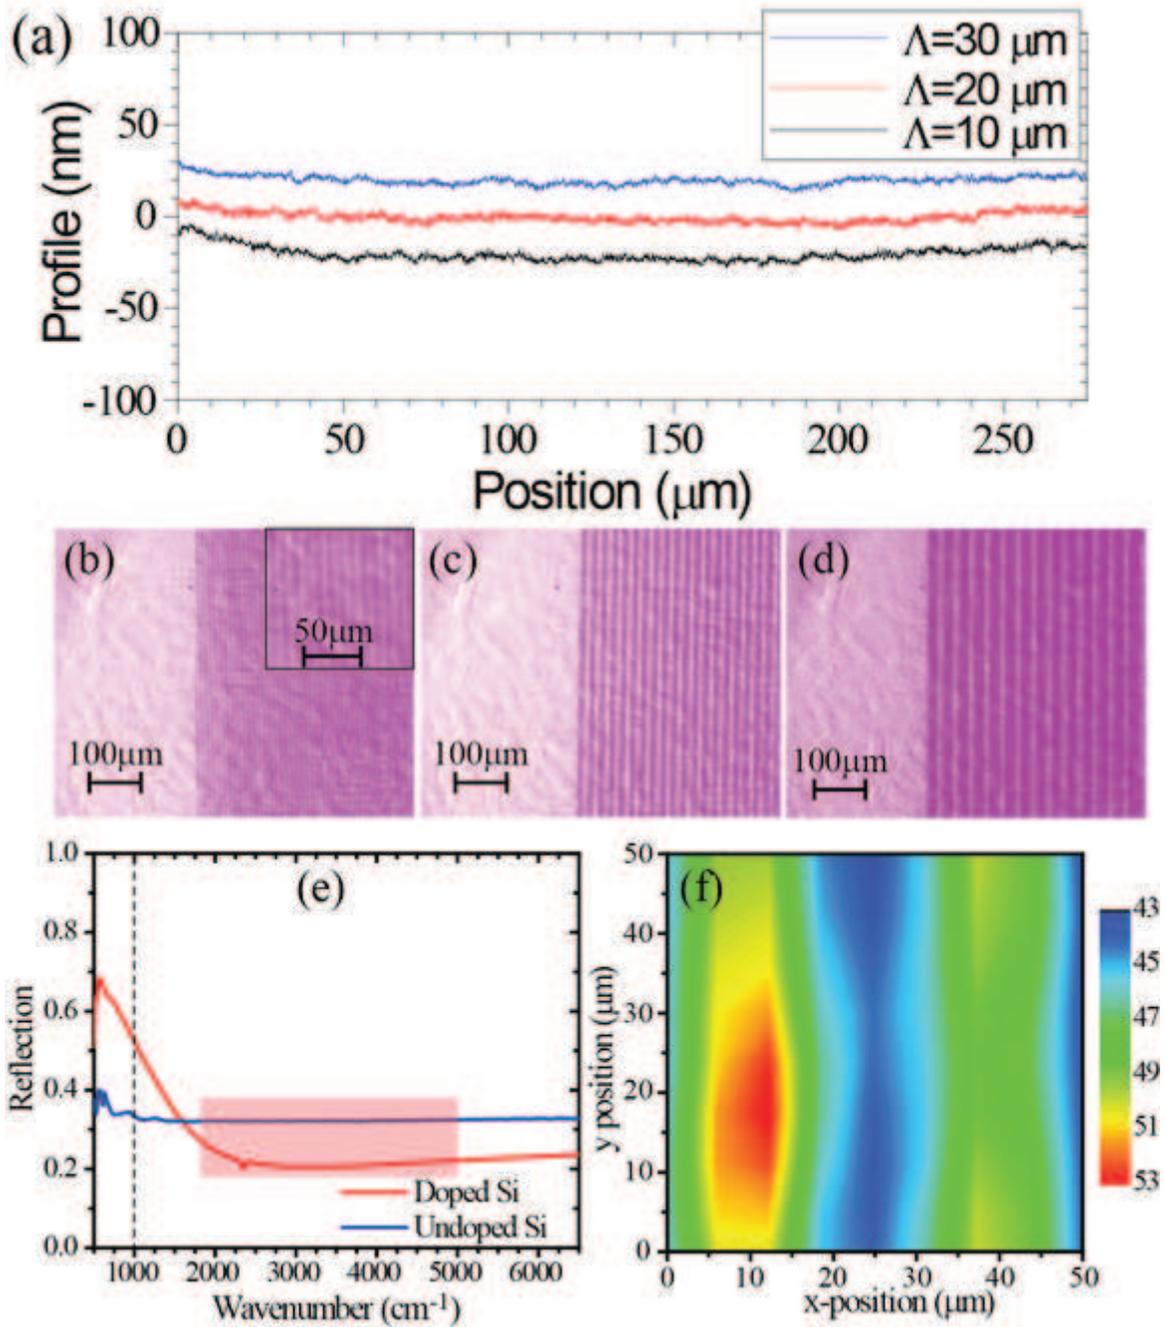
Figure 9: (a) Surface profile of the three samples with flat gratings. IR thermal image of both a rectangular area doped throughout and the flat gratings with period (b) 10 μm, with inset showing an expanded view, (c) 20 μm and (d) 30 μm, showing segregation of dopants. (e) Comparison of reflection spectrum from the rectangular doped area in part (b) with a control sample of undoped, unprocessed silicon. The highlighted area indicates the region of the spectrum used for the images in parts (b-d). (f) Spatially-resolved image of reflection from the grating with period 30 μm in part (d), taken with PerkinElmer Spotlight at 10 μm [1].

In order to determine the lateral diffusion of dopants in our Si material, we studied the thermal

25

emission from the surface of our grating samples. The experimental set-up collects spatially (but not spectrally) resolved thermal emission from the heated surfaces of our gratings. Nonetheless, across the wavelength range collected (2-5.5 μm), it is expected that the doped Si will have a somewhat different reflectivity and absorptivity than the undoped material (the difference in reflection from the doped and undoped material is shown in figure 9(e)). Thus, though the thermal emission does not give spectrally resolved emission, nor does it detect emissivity in the wavelength range where the doped-Si is "metallic", by measuring the spatial dependence of the thermal emission, a qualitative understanding of dopant diffusion can be obtained. The thermal emission images shown in figure 9(b-d) clearly indicate that a lateral variation in emissivity is obtained from our grating samples. More importantly, the lateral variation of emissivity shows a strong modulation between the highly doped (light purple) and undoped (dark purple) regions of the grating sample, with little intermixing of the two regions. These results corroborate our modeled diffusion profile for dopants in our highly-doped Si, which showed diffusion lengths < 1 μm, well below the spatial resolution of the thermal imaging experiment. Finally, figure 9(f) shows the spatially-dependent reflectivity, taken with the PerkinElmer Spotlight system, of the $\Lambda = 30$ μm sample at a wavelength of $\lambda = 10$ μm, where we expect the doped Si to behave as an optical metal. It should be noted that in this figure the reflectivity from the grating only varies from 43% to 53%. This result is consistent with the broad-area reflection from the solid doped portion of the patterned sample and that of an undoped Si wafer, shown in figure 9(e). Ideally, a stronger variation of permittivity would be desired, in order to more clearly differentiate between dielectric and "metallic" portions of a composite material fabricated using the outlined process. However, because all evidence points to weak, sub-micron dopant diffusion lengths, it is expected that the use of higher doping concentration spin-on dopants would allow for even greater modulation of permittivity with minimal effect on diffusion.

5 Conclusion

In summary, I have demonstrated two applications of highly doped semiconductors to mid-infrared photonics: a computational model for plasmonic waveguides using COMSOL, and an experimental demonstration laterally-patterned flat gratings. The results of COMSOL simulations agree with analytical predictions using Maxwell's equations. The code is easily adaptable to different material systems and allows for extensive automation, and could be extended for use in the design of optical circuitry with nanoparticles. I have shown the ability to laterally control the permittivity of silicon using patterned spin-on dopants, while leaving the Si surface completely flat, as measured using surface profilometry. Lateral variation in permittivity has been confirmed using FT-IR spectroscopy and thermal imaging, demonstrating minimal, subwavelength lateral diffusion. In addition, I have shown strong diffraction of incident light from patterned surfaces, with experimental results closely matching the simple grating equation. The smallest grating period fabricated showed clear lateral variation of permittivity on a subwavelength scale, as measured by spatially-resolved thermal emission and IR reflection microscopy. These results are a potential first step towards truly flat composite plasmonic materials with subwavelength constituents having both positive and negative permittivity, with the potential for use as couplers to plasmonic waveguides, flat metasurface architectures, and other mid-infrared nanophotonic structures and devices.

References

[1] A. Rosenberg, J. Surya, R. Liu, W. Streyer, S. Law, L. S. Leslie, R. Bhargava, and D. Wasserman, "Flat mid-infrared composite plasmonic materials using lateral doping-patterned semiconductors," IOP Journal of Optics, submitted for publication.

[2] S. A. Maier, Plasmonics: Fundamentals and Applications, New York: Springer, 2007.

[3] J. A. Schuller, E. S. Barnard, W. Cai, Y. C. Jun, J. S. White, and M. L. Brongersma, "Plasmonics for extreme light concentration and manipulation," Nature Materials, vol. 9, no. 3, pp. 193-204, Feb. 2010.

[4] J. N. Anker, W. P. Hall, O. Lyandres, N. C. Shah, J. Zhao, and R. P. Van Duyne, "Biosensing with plasmonic nanosensors," Nature Materials, vol. 7, no. 6, pp. 442-453, Jun. 2008.

[5] M.E. Stewart, C. R. Anderton, L. B. Thompson, J. Maria, S. K. Gray, J. A. Rogers, and R. G. Nuzzo, "Nanostructured plasmonic sensors," Chemical Review, vol. 108, no. 2, pp. 494-521, Jan. 2008.

[6] N. Engheta, "Circuits with Light at Nanoscales: Optical Nanocircuits Inspired by Metamaterials," Science, vol. 317, no. 5845, pp. 1698-1702, Sep. 2007.

[7] D. J. Bergman and M. I. Stockman, "Surface plasmon amplification by stimulated emission of radiation: quantum generation of coherent surface plasmons in nanosystems," Phys. Rev. Lett., vol. 90, no. 2, p. 027402, Jan. 2007.

[8] M. T. Hill et al. "Lasing in metal-insulator-metal subwavelength plasmonic waveguides," Optics Express, vol. 17, no. 13, pp. 11107-11112, Jun. 2009.

[9] M. A. Noginov et al., "Demonstration of a spaser-based nanolaser," Nature, vol. 460, no. 7259, pp. 1110-1113, Aug. 2009.

[10] R. F. Oulton et al., "Plasmon lasers at deep subwavelength scales," Nature, vol. 461, no. 7264, pp. 629-632, Aug. 2009.

[11] D. C. Adams et al., "Funneling light through a subwavelength aperture with epsilon-near-zero materials," Phys. Rev. Let., vol. 107, no. 13, p. 133901, Sep. 2011.

[12] A. Boltasseva and H. A. Atwater, "Low-loss plasmonic metamaterials," Science, vol. 331, no. 6015, pp. 209-291, Jan. 2011.

[13] S. Law, V. Podolskiy, and D. Wasserman, "Towards nano-scale photonics with micro-scale photons: the opportunities and challenges of mid-infrared plasmonics," Nanophotonics, vol. 2, no. 2, pp. 103-130, Jan. 2013.

[14] R. F. Curl and F. K. Tittel, "Tunable infrared laser spectroscopy," Annu. Rep. Prog. Chem., Sect. C, vol. 98, pp. 219–272, Jul. 2002.

[15] K. Anglin et al., "Voltage controlled active mid-infrared plasmonic devices," J. Appl. Phys., vol 109, no. 12, p. 123103, Jun. 2011.

[16] Y. H. Ye and J. Y. Zhang "Middle-infrared transmission enhancement through periodically perforated metal films," Applied Physics Letters, vol. 97, no. 16, p. 211110, Apr. 2004.

[17] D. Wasserman, E. A. Shaner, and J. G. Cederberg, "Mid-infrared doping tunable extraordinary transmission from sub-wavelength gratings," Applied Physics Letters, vol. 90, no. 19, p. 191102, May 2007.

[18] S. C. Lee, S. Krishna, and S. R. J. Brueck, "Quantum dot infrared photodetector enhanced by surface plasma wave excitation," Optics Express, vol. 17, no. 25, pp. 23160-23168, Dec. 2009.

[19] J. P. Tetienne et al., "Injection of midinfrared surface plasmon polaritons with an integrated device," Applied Physics Letters, vol. 97, no. 21, p. 211110, Nov. 2010.

[20] K. A. Willets and R. P. Van Duyne, "Localized surface plasmon resonance spectroscopy and sensing," Ann. Rev. Phys. Chem., vol. 58, pp. 267-297, May 2007.

[21] S. Law, D. C. Adams, A. M. Taylor, and D. Wasserman, "Mid-infrared designer metals," Optics Express, vol. 20, no. 11, pp. 12155-12165, May 2012.

[22] F. Marquier, K. Joulain, J. P. Mulet, R. Carminati, and J. J. Greffet, "Engineering infrared emission properties of silicon in the near field and the far field," Optics Communications, vol. 237, pp. 379-388, Jul. 2004.

[23] J. Jung and T. G. Pedersen, "Analysis of plasmonic properties of heavily doped semiconductors using full band structure calculations," J. Appl. Phys., vol. 113, no. 11, p. 114904, Mar. 2013.

[24] Y. B. Chen, "Development of mid-infrared surface plasmon resonance-based sensors with highly-doped silicon for biomedical and chemical applications," Optics Express, vol. 17, no. 5, pp. 3130-3140, Mar. 2009.

[25] N. W. Ashcroft and N. D. Mermin, Solid State Physics, Holt, Rinehart and Winston, 1976.

[26] S. Law, L. Yu, and D. Wasserman, "Epitaxial growth of engineered metals for mid-infrared plasmonics," Journal of Vacuum Science and Technology B, vol. 31, no. 3, p. 03C121, May 2013.

[27] A. J. Hoffman et al. "Negative refraction in semiconductor metamaterials," Nature Materials, vol. 6, no. 12, pp. 946-950, Oct. 2007.

[28] S. Law, L. Yu, A. Rosenberg, and D. Wasserman, "All-semiconductor plasmonic nanoantennas for infrared sensing," Nano Letters, vol. 13, no. 9, pp. 4569-4574, Aug. 2013.

[29] S. Law et al. "All-semiconductor negative-index plasmonic absorbers," Phys. Rev. Lett., vol. 112, no. 1, p. 017401, Jan. 2014.

[30] D. Li and C. Z. Ning, "All-semiconductor active plasmonic system in mid-infrared wavelengths," Opt. Express., vol. 19, no. 15, pp. 14594-14603, Jul. 2011.

[31] J. W. Cleary et al., "IR permittivities for silicides and doped silicon," J. Opt. Soc. Am. B., vol. 27, no. 4, pp. 730-734, Apr. 2010.

[32] J. C. Ginn, R. L. Jarecki Jr., E. A. Shaner, and P. S. Davids, "Infrared plasmons on heavily-doped silicon," J. Appl. Phys. vol. 110, no. 4, p. 043110, Aug. 2011.

[33] R. Soref, "Mid-infrared photonics in silicon and germanium," Nature Photonics, vol. 4, no. 8, pp. 495-497, Aug. 2010.

[34] W. Streyer, S. Law, G. Rooney, T. Jacobs, and D. Wasserman, "Strong absorption and selective emission from engineered metals with dielectric coatings," Optics Express, vol. 21, no. 7, pp. 9113-9122, Apr. 2013.

[35] S. L. Chuang, Physics of Photonic Devices. Hoboken: John Wiley & Sons, Inc., 2009.

[36] Y. Zhao and A. Alu, "Manipulating light polarization with ultrathin plasmonic metasurfaces," Physical Review B, vol. 84, no. 20, pp. 205428, Nov. 2011.

[37] N. Yu et al., "Light propagation with phase discontinuities: generalized laws of reflection and refraction," Science , vol. 334, no. 6054, pp. 333-337 (2011).

[38] A. V. Kildishev, A. Boltasseva, and V. M. Shalaev, "Planar photonics with metasurfaces," Science vol. 339, no. 6125, p.1232009, Mar. 2013.

[39] A. Pors, O. Albrektsen, I. P. Radko, and S. I. Bozhevolnyi, "Gap plasmon-based metasurfaces for total control of reflected light," Sci. Rep., vol. 3, Jul. 2013.

[40] Y. Huang et al., "Phase-gradient gap-plasmon metasurface based blazed grating for real time dispersive imaging," Applied Physics Letters, vol. 104, no. 16, p. 161106, Apr. 2014.