

Introduction to Experimental Biophysics

Biological Methods for Physical Scientists, Second Edition

Jay L. Nadeau



CRC Press

Taylor & Francis Group
Boca Raton London New York

CRC Press is an imprint of the
Taylor & Francis Group, an **informa** business

CHAPTER 17

Introduction to Nanofabrication

Orad Reshef

17.1 INTRODUCTION

This chapter aims to give a practical introduction to some of the techniques of nanofabrication, written by someone who recently completed a PhD using these techniques. It begins by discussing patterning and pattern transfer, and then techniques of deposition by spin coating, evaporation, and chemical vapor deposition (CVD). Techniques of characterization specific to nanoscale devices are then discussed in two sections: the first dealing with *metrology* and the second with spectroscopy, with some spectroscopic techniques that were not mentioned in **Chapters 14** and **16**. Finally, the chapter concludes with some practical tips for achieving and maintaining a clean, high-quality sample, noting where cutting corners can have disastrous results.

The planar process

Most nanofabrication techniques in use today were developed to be compatible with the *planar process* (**Figure 17.1**). This process, which was initially developed in the 1950s by Fairchild Semiconductor (a predecessor of Intel), is a top-down approach to fabricating nanostructures. A sample begins with a substrate, a wafer or chip, which acts as the foundation for the sample. A layer of material is deposited onto the substrate, forming a *thin film*. A layer of *resist*, a chemically sensitive polymer, is subsequently deposited onto the thin film. The desired shape is then patterned into the resist using *lithography* techniques. Finally, the design is transferred from the resist to the deposited film using a *pattern transfer* technique. The resist is dissolved away, and a protective layer called a *cladding* is deposited and polished. This procedure may be repeated multiple times to form a vertical stack of two-dimensional structured layers. In this chapter, we will discuss each of these steps in detail to produce a nanostructure composed of a single two-dimensional layer.

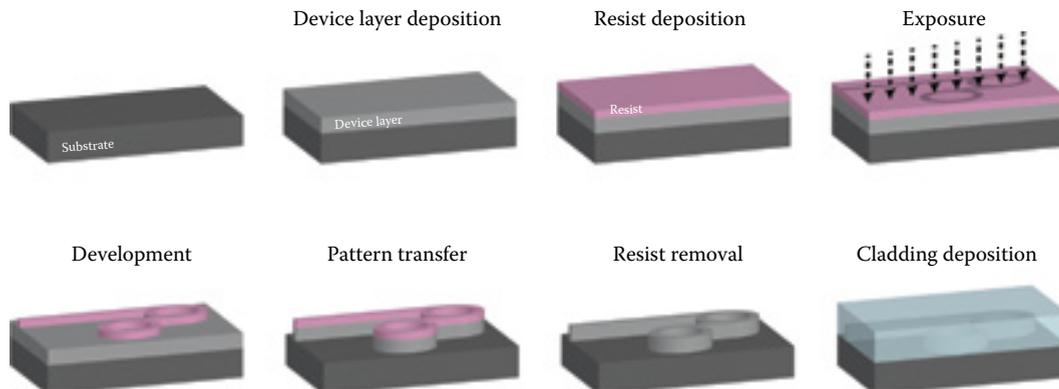


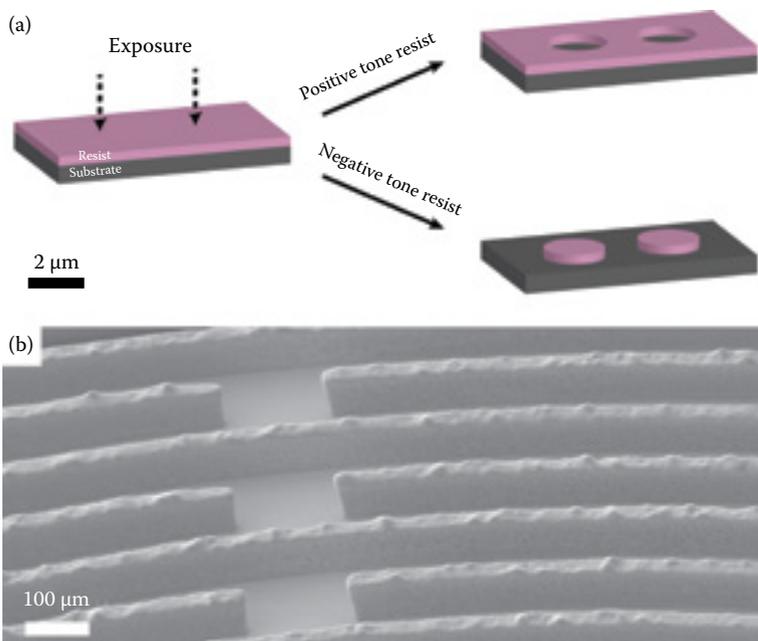
Figure 17.1 The planar process can produce arbitrary shapes that are confined to the 2-D plane of the device layer.

17.2 PATTERNING

Resist

The first step to structuring a pattern into a surface is to deposit a layer of resist. A resist is a soft material that is sensitive to a lithography (i.e., pattern writing) step but is insensitive to subsequent processing steps. A pattern can be written into a resist and then transferred via a pattern transfer method into a surface of choice. Where the resist is exposed, the chemical bonds are changed from the surrounding area, so that it reacts differently to a developer solvent (Figure 17.2). Resist can be positive-tone, where regions that are exposed during the lithography

Figure 17.2 (a) In a positive-tone resist, exposing a circle yields a hole in the polymer, whereas in a negative-tone resist, exposing a circle yields a polymer cylinder. (b) Negative-tone resist (SU-8) structures after development. (SEM courtesy of Victor White.)



step are dissolved away, or negative-tone, where the affected regions are the only ones that remain.

In research contexts, the workhorse resist is polymethyl methacrylate (PMMA) due to its cost and versatility. It is a positive-tone resist for electron beam lithography (EBL). It can be produced at varying concentrations and molecular weights, which yield a broad range of viscosities and exposure sensitivities. Thus, it can be reliably deposited at different thicknesses, ranging from below 100 nm up to over 1000 nm. Its chemical resistivity is fair. We compare this polymer to a proprietary alternative positive-tone resist, ZEP520A (by Zeon Chemicals), which is much more chemically inert and can be exposed using much smaller doses, dramatically reducing exposure times. This resist is known to also yield smoother features with shorter exposure times at the expense of a much higher price.

Lithography

There are many different methods with which a pattern can be defined in a resist. To a certain degree, the desired structure and application will dictate what lithography tool to use. *Photolithography*, where the exposure is provided by light, is the most used lithography method by industry due to its speed. A common variant of photolithography is called *contact lithography*; a wafer is coated with resist, a mask with a desired pattern is pressed against it, and the entire wafer is flooded with light at the same time. In other variants, the mask never physically touches the resist, preventing defects and damage. The image can even be focused on the surface of the resist using a lens following transmission through the mask, which provides an additional reduction in feature size (e.g., by a factor of 5). Advanced photolithography tools are even capable of maskless photolithography by projecting an image directly on the resist. The entire exposure process takes a fraction of a second, regardless of the number of features in the pattern. A whole *wafer* can be up to 300 mm in diameter and is subsequently diced into its individual chips, resulting in hundreds or thousands of chips being fabricated per minute. A *stepper* can be used to repeatedly expose a smaller pattern many times on the surface of a wafer.

Photolithography suffers from many limitations. First, the smallest feature sizes are limited by the exposure wavelength, which is usually fixed by the tool you have available. Thus, without the absolute state-of-the-art equipment, it is challenging to reliably achieve feature sizes smaller than a wavelength (e.g., 500 nm), which is a major obstacle to many applications. Second, before the exposure step, most photolithography tools require a mask to be fabricated, which dramatically slows down the prototyping and iteration cycle.

For these reasons, *EBL* is a popular tool in research and development. EBL can routinely achieve feature sizes down to 10 nm without need of a mask. It works by raster-scanning a collimated beam of electrons onto the resist. The major drawback in EBL is its operating speed—the writing process scales with the exposed area due to the rastering of a focused beam, so the total write time for an entire wafer can grow to be prohibitively long. For this reason, mass production is universally done in photolithography once a design is perfected and ready for the market.

These lithographic techniques produce 2-D structures that are compatible with the planar process. Though the pattern in the plane of the sample can have arbitrary

shapes, the cross-sectional profiles out of the plane are regularly rectangular, with near-right-angled corners. Any three-dimensional structures have to be produced with some ingenuity that takes advantage of anisotropic material-dependent deposition or etching properties. A true arbitrary three-dimensional polymer structure can be produced using grayscale lithography, which makes use of resist that is only partially exposed, or direct laser writing techniques, where a laser beam is focused within the body of the resist. These methods are highly specialized and are beyond the scope of this chapter; however, they are rapidly becoming accessible and commercialized.

Nanoimprint lithography

Nanoimprint lithography is a fast and cheap method to scale up devices that are created using EBL. A negative template is first created using other lithographic methods, and then it is stamped into an imprint resist. This method can ultimately yield the mass production of devices with a resolution on the order of tens of nanometers.

Focused ion beam

The final lithographic technique we will discuss is *focused ion beam* (FIB), where a collimated beam of inert ions (e.g., Ga^+) mills away the material. Using an FIB, the desired structure is directly chiseled out of the material without the need for a resist, making it the most versatile of all lithography techniques. Because it is also a rastering technique, it suffers from the same issues with production speed as EBL. It also produces rough surfaces compared to the other methods and is known to be relatively challenging to master. Nevertheless, it can yield very small features and can structure materials that are often otherwise challenging to etch. It is most often used to prepare extremely thin samples (on the order of 100 nm wide) for characterization, for example, for transmission electron microscopy.

17.3 PATTERN TRANSFER

Wet etching

Etching a material is the process of chemically removing it. In a wet etch, the entire sample is immersed in a reactive solvent. Ideally, this solvent will target a specific material, leaving all the other materials on the sample untouched. For example, to remove a layer of gold sitting on a glass substrate, the gold should be removed completely, while the glass should be kept in pristine condition. This chemical selectivity is thus important to consider during the planning stages of a fabrication process, before even beginning the deposition of any materials. There exist a lot of literature with tables of materials and solvents and their etch rates that have been established experimentally. They should be heavily consulted before proceeding with any etching. Some references are given at the end of the chapter to articles giving these data.

A *wet etch* is typically isotropic, which means the etchant reacts equally in all directions, ignoring any lattice planes in the material. Thus, wet etches are used almost solely to remove a selected material from a sample indiscriminately. However, the isotropy of the etching process could also be exploited to undercut a material and achieve freestanding structures (**Figure 17.3**).



Figure 17.3 A wet etch is used to etch under the resist and form a suspended structure.

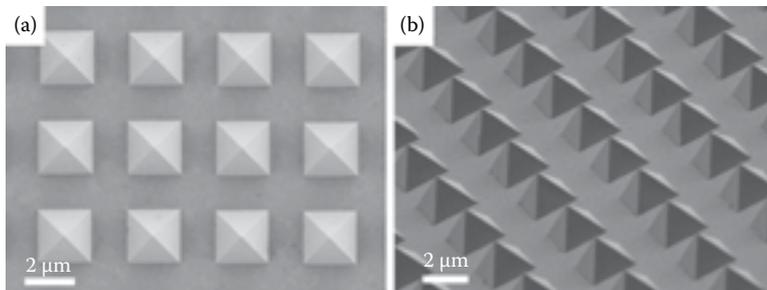


Figure 17.4 Gold nanopillars fabricated using silicon and an anisotropic wet etch. (Figure courtesy of Nabihha Saklayan.)

Besides isotropic wet etches, there also exist anisotropic wet etches that are much more material specific. For example, potassium hydroxide (KOH) etches the (100) silicon surface along the (111)-oriented crystal lattice planes, which can be used to yield sloped walls. These properties have been used to create negative templates for nanopillars (Figure 17.4). Anisotropic etches thus provide a degree of freedom for creating three-dimensional structures.

After completing a wet etch, the sample is rinsed with an inert solvent (e.g., distilled water). Best practices dictate rinsing the sample multiple times in separate flasks of water. This might seem excessive, but there is a sound reason for it—the concentration of the etchant could be high enough that only rinsing in a single cup will leave some residue that is sufficiently concentrated to continue etching the surface.

Dry etching

As the name suggests, a *dry etch* is similar to a wet etch, but instead of using liquids, gases are used to selectively remove materials. It turns out that this small change complicates things dramatically. First, a dry etch procedure cannot be done in the open air or a fume hood; at the very least, a vacuum chamber is necessary. Second, since the gases must be flowed at a consistent rate, extra equipment is necessary to control their flow and the chamber pressure. All of these considerations drastically increase the cost of the etch.

However, this increased cost and complexity afford a finer control of etching properties and more sophistication in the etch. A voltage applied on the gases can excite them to form a plasma, which is accelerated toward the sample. This bombardment adds a unidirectional physical component to the etch, in addition to the chemical component provided by the choice of reacting gases. Together, standard dry etching tools provide enough flexibility to obtain perfectly vertical sidewalls or sidewalls with tailored angles (Figure 17.5).

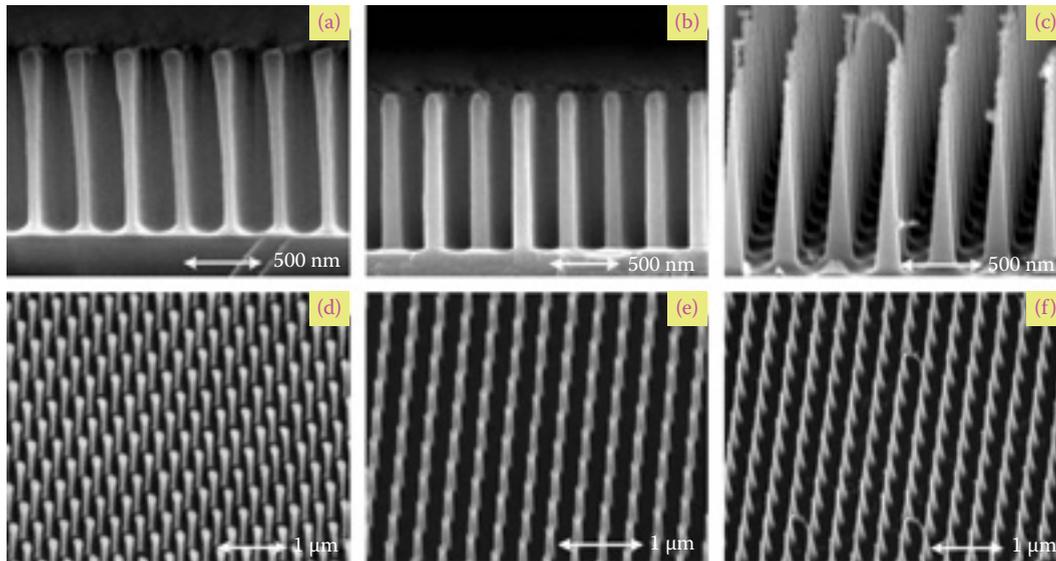


Figure 17.5 Dry etching can produce sidewalls with (a, d) positive or (b, e) negative slope angles. (c, f) With some optimization, even vertical sidewalls can be achieved. (Figure courtesy of Yung-Jr Hung and Brian Thibeault.)

There are dozens of variables that can be tuned while etching, which include the pressure of the chamber, the flow rate of the gases, as well as various temperatures, powers, and voltages. References are given at the end of the chapter to some key articles summarizing these variables. Though finding an acceptable etch recipe sounds daunting, the manufacturer always delivers the tool with a set of standard recipes for the most commonly used materials, such as Si or GaAs. However, even with the same equipment and process parameters, results will be variable and will have to be tweaked to optimize the etch to exhibit the desired properties. Beyond this, a design of experiments (i.e., the Taguchi method) could be employed to find an ideal etch without exploring the entire parameter space of variables.

Two more practical points on etches are apropos. First, there needs to be excellent thermal contact between the sample and the sample holder to allow for the dissipation of heat accumulated by the sample due to ion bombardment. This is paramount to maintaining a consistent etch between uses. We recommend using a thermal grease instead of thermal paste; grease or oil is much cleaner because it dissolves completely in a solvent and does not redeposit on the sample. The appropriate amount of grease for a given sample size will need to be determined empirically. Too much grease will seep out from under the sample, and the ions in the etch will sputter grease onto the face of the sample. Too little grease will not be enough to spread out under every part of the sample, and only the small region of the sample on top of the grease will etch properly.

Second, the etch in the very center of the chip will be a little different from the etch on the edges of the sample. This is because the complete physics of a dry etch process includes contributions from secondary ions that recoil from reacting elsewhere on the chip in addition to the primary ions that come directly from above. Since the regions on the edges have access to half as many of these secondary ions, the etch rate on the etch can vary significantly from the center. If all the features of interest lie in the center of the sample, perhaps this is not an

important consideration. However, if there are any structures that reach the very edge of the chip, they will be etched to a different depth and at a different quality. A quick fix to this problem is to make the starting substrate larger. Alternatively, a secondary set of dummy samples may be placed surrounding the sample. As these samples are etched away, they can produce the secondary ions to mimic the center of the chip on its edges.

Lift-off

Many materials are very challenging or even impossible to etch reliably. Often, there are practical reasons not to etch, such as the lack of the necessary gases for a given reaction, or a ban on certain materials by the owner of the tool for fear of contamination. The layer *under* the one being etched might also need to be protected at all costs, and so etching is too risky. Whatever the reason, there exists an entirely different pattern transfer procedure called a *lift-off* process, illustrated in **Figure 17.6**. The negative image of a pattern is written into a resist, and a layer of the desired material is deposited directly on top of it. In the regions where the resist is dissolved away, the deposited material makes direct contact with the substrate. All remaining resist is then dissolved away, leaving behind the desired material. Using this method, virtually any material that can be evaporated can be structured.

There are many reasons a dry etch is typically the preferred method to structure a material. First, lift-off processes are relatively slow and unscalable when compared to dry etches. Next, using a negative-tone resist directly will always result in smaller features than lifting off material from a positive-tone resist. Features yielded in an etch can be optimized to be smoother than those lifted off, as well. Finally, when etching, the starting material can be atomically smooth and epitaxial, and can also be thoroughly characterized prior to being structured.

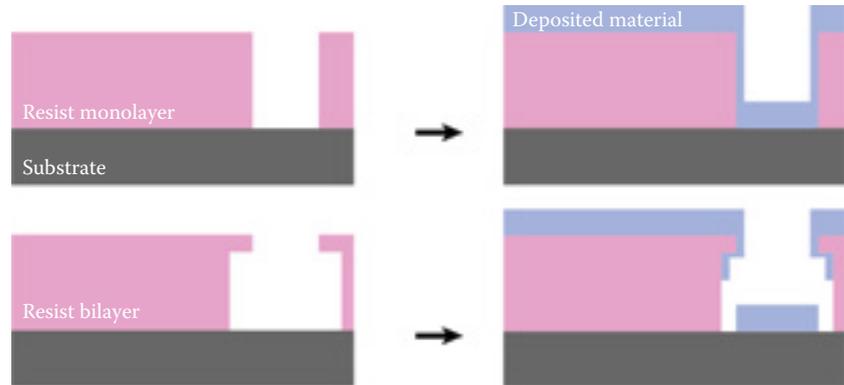
As in etch procedures, the lift-off procedure is highly material specific, and so a recipe needs to be developed for every individual scenario. There are certain rules of thumb:

- The deposition method needs to be nonconformal, that is, the material should not coat the sides of the resist in the exposed areas. Otherwise, the structured region remains in contact with the region designed to be lifted off. This causes the entirety of the deposited material to be removed with the resist, including the desired structures.
- The thickness of the resist needs to be at least two or three times thicker than the desired thickness of the final material.
- During the lift-off procedure, the sample needs to be agitated gently in the solvent. Anything more aggressive that uses stir bars or ultrasonication



Figure 17.6 In a standard material lift-off procedure, material is deposited on top of the resist. The resist is subsequently removed, leaving the desired structure on the substrate.

Figure 17.7 Cross section of a sample, comparing a resist monolayer to a resist bilayer in the lift-off process. When using a bilayer with an overhang, the desired structure is isolated from the material being lifted off.



should be avoided if possible. However, sometimes, the adhesion between the deposited material and the substrate is really good, and sonication can still be employed to accelerate the process.

- It is preferable to set the sample to soak in the solvent on a hot plate overnight if there is time. This method yields excellent reproducible results for most cases.

It sometimes remains challenging to successfully lift off a material. A contributing factor could be that the deposition method is inherently a little too conformal, making it impossible to avoid continuous material deposition along the resist sidewalls, despite following the rules of thumb outlined previously. To avoid this, a resist bilayer may be used (Figure 17.7). Instead of depositing only a single layer of resist, two layers are deposited; however, the bottom layer is chosen to be a resist that is slightly more sensitive to the exposure. This produces an overhang in the resist, which helps isolate the desired structure on the substrate. Using a resist bilayer is not necessary; however, it does act as a cure-all to most lift-off problems.

Template stripping

Another way to structure a material for which there is not an etcher or an etch recipe available is called *template stripping*. It is a slow, manual process that is not particularly scalable, but it can achieve excellent results for research purposes. It begins with a structured surface called a template. The pattern in the template needs to be the negative pattern of what is desired in the final result. A layer of the desired material (e.g., gold) is deposited on top of the template. This is followed by an adhesion layer, which is subsequently pressed into a separate sample (e.g., a glass slide), which will act as the new substrate. Finally, the template is peeled off of the sample, yielding the desired structure (Figure 17.8). Template stripping is notable for producing smooth structures in any material that can be deposited

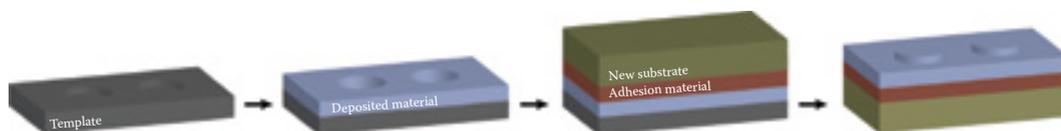


Figure 17.8 The template-stripping process.

since the roughness of the final surface is dictated only by the roughness of the template. If the procedure is performed carefully, the master template can even be reused to produce multiple samples.

17.4 MATERIAL DEPOSITION

The deposition method to use is in part dictated by the material in question and the desired film properties. Not every method is capable of depositing every material, and the properties of the material once deposited (e.g., the crystal structure, grain size, conformality, surface conductivity, etc.) vary dramatically depending on the tool that is used. Additionally, just because an available facility owns a particular tool does not mean that all of the necessary raw materials (e.g., the gases or pure material target) are available. Finally, no single technique is optimal for all materials—every material has a deposition method it “prefers” for a particular application—some materials can only be grown epitaxially (i.e., in a perfect crystal lattice) using only one particular deposition method. For these reasons, a lot of homework needs to be done when selecting the deposition technique for a process. In this section, we will present the most commonly used deposition techniques.

For most of these devices, the deposition rate can be monitored *in situ* using a quartz crystal microbalance (see [Chapter 10](#)) or an interferometer. In newer tools, a feedback loop can be set up so that the deposition rate is held constant, yielding a consistent film within a few nanometers (i.e., a few atomic layers) of the desired thickness. However, not all tools have these capabilities. In these cases, a test run is performed on a so-called witness sample (e.g., an unstructured substrate like a glass slide) to establish a deposition rate.

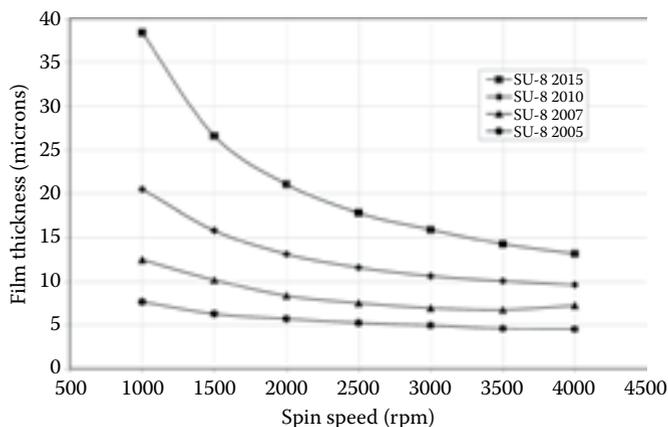
Spin coating

The cheapest and simplest tool to deposit a thin film of material is the *spin-coater*. A small amount of viscous solvent is pipetted onto a sample or a wafer, which is then spun at several thousand revolutions per minute. The faster the spin, the thinner the layer produced. This is almost always followed by a baking step, e.g., 60 s on a hot plate set at 180°C. It is also good practice to bake a sample *before* spinning anything on it (e.g., 5 min above 100°C) to remove any moisture that has adsorbed onto the surface and to ensure a consistent spin.

This is the method by which resist is deposited. Commercial resists are shipped with a spin curve, a recipe that predicts the film thickness as a function of the spin speed. It is recommended to use the speed in the middle of the curve for the best consistency. For example, for the spin curve in [Figure 17.9](#), this corresponds to 2500 rpm. A spun-on layer is usually only uniform in the center of the sample, with a raised ridge at the edges of the sample, known as an edge bead. The more symmetric the sample, the larger the uniform region in the middle of the sample can be. Circular samples such as wafers produce the most uniform layers with the smallest edge beads, the corners in a square sample are typically unusable, and rectangular samples will have even larger unusable areas.

As a final note, we recommend doing a spin test before using a sample in the spin-coater. A spinner has a vacuum safety lock that will prevent it from spinning if it does not sense a sample. The spin should also stop if the sample is launched

Figure 17.9 Spin curve for SU-8 2000, a negative-tone photoresist.



off of the spinner. Thus, the standard spinner test is as follows: with no sample mounted, start the spinner. It should not spin at this point. Place your gloved finger on the spinner to block the suction hole. This should prompt the spinner to start spinning. When you remove your finger, it should stop spinning. If any of these parts of the test fail, it is likely because there is some blockage in the spinner caused by some resist that has leaked into its pipes, and you have just saved your sample.

Sputter deposition

To *sputter* a material, a sample is placed facing a pure target of a material in a vacuum chamber. An inert gas (e.g., argon) is flowed into the chamber, bombarding the target, which in turn emits charged ions of the target material. These ions are accelerated toward the sample with either a direct current (DC) or radiofrequency (RF) bias similar to what was discussed previously in etching systems, redepositing the ions on the sample and forming a thin film. The sample can be rotated on its axis to achieve a more uniform film. The mechanism is entirely physical and does not depend on any material-dependent properties, meaning most materials can be deposited using sputter deposition provided there is a high-quality target available.

The sample can be heated during deposition to encourage the formation of a particular crystal structure. There is often the option of flowing an additional gas to react with the target material. For example, oxygen gas can be flowed alongside argon in a chamber with a titanium target to produce titanium dioxide. This is called reactive sputtering, and it is a suitable method to obtain a compound material when a particular target for that material (in this case, a titanium dioxide target) is unavailable. A chamber can also come with multiple targets. For example, a titanium target and a gold target could be installed in the chamber, and both could be sputtered simultaneously. This makes it relatively straightforward to deposit a stack of materials (e.g., titanium/gold/titanium/gold, etc.).

Beyond this, the physics of the sputtering process is very complex and can lead to counterintuitive effects such as hysteresis. For example, a higher-powered plasma does not necessarily produce a faster deposition rate. More so than for any other of the other deposition methods described in this chapter, we highly recommend reading some literature if serious about developing a sputtering process.

Thermal and electron beam evaporation

Evaporation techniques are relatively simple and affordable. A sample is placed facedown above a crucible (often called a “boat”) filled with the desired material in a vacuum chamber. The boat is then heated up until the material evaporates, leaving a thin film on the sample. This is called *thermal evaporation* (TE). Alternatively, the material is placed in a crucible and is melted with a focused electron beam. This is called *electron beam evaporation* (EBE). Both techniques are similar and yield films of comparable quality. However, the region heated in the crucible is much smaller when using EBE, leading to a much more directional deposition. EBE can achieve much higher temperatures than TE as well, allowing for the deposition of materials with higher boiling points. A more important consideration is that the material quality depends on the tool that is selected, and it is not always straightforward to predict which tool will perform better for a given material. Finally, compared to sputtering, there are fewer parameters available to control when using evaporation tools. Thus, there are fewer degrees of freedom available to optimize a deposition.

Chemical vapor deposition

In CVD, gases are flowed over a sample. These gases react on the surface of the sample to yield a film. The CVD tool itself is in fact almost identical to a dry etcher, save for the choice of gases and the recipes employed. Thus, it is just as complicated to develop a reliable deposition recipe because of the large number of tunable parameters during deposition as it is to obtain a reliable etch recipe. In this case, just as in the case for the dry etcher, the manufacturer ships the tool with a few recipes for commonly used materials. In general, it is a very complex tool with many sensitive components, making it more fragile and prone to downtime than the other deposition tools listed so far. It is also capable of depositing a shorter list of materials since a recipe needs to be developed for each one individually. However, once a good recipe is established, the CVD is a reliable turnkey machine with repeatable depositions, constant deposition rates, and little room for human error.

Atomic layer deposition

Atomic layer deposition (ALD) is a relatively new deposition technique, which is derived from CVD. The principle behind ALD is the same as the one behind CVD, with one key difference—instead of flowing the gases continuously, the gases are injected in intermittent spurts. This allows for the deposition of individual atomic layers upon the surface, making ALD an extremely precise and controllable deposition technique, albeit an incredibly slow one. ALD also has the distinction of being the most conformal (i.e., least directional) deposition method. Using ALD, the sidewalls of a structure are uniformly coated as well as the top. Trenches with very high aspect ratios (50:1) can be filled with ease. ALD conforms so well that even the backside of the sample gets coated during a deposition!

Just as for CVD, there is but a short list of materials that can be deposited (mostly oxides). However, the extreme conformality and accuracy of this deposition technique has caused ALD to become an attractive tool for industrial applications, and expanding this list of materials has become an active area of research. For example, in 2016, researchers managed to deposit gold for the first time, a material without a naturally occurring oxide.

17.5 METROLOGY

Once a structure has been fabricated, it needs to be characterized. For feature sizes on the nanoscale, this is something surprisingly nontrivial to accomplish. Most features are too small to resolve using optical microscopy. Light microscopes are also incapable of measuring heights or thicknesses. To this end, the entire field of metrology has been developing advanced microscopy and profilometry techniques. In this section, we discuss the most useful and accessible metrology tools that can be expected to be available in every cleanroom.

Scanning electron microscopy

Scanning electron microscopy (SEM) is the simplest and most reliable way to image a nanostructured surface. It produces high-quality, high-resolution images of features down to tens of nanometers, far smaller than what is physically possible using an optical microscope. (Every picture in this chapter was taken using an SEM.) Turnkey commercial SEMs are now so reliable that anyone can learn how to load samples and take stunning images of nanoscale structures with very little practice or preparation.

An SEM works by focusing a collimated beam of electrons (often the very same beam that is used in EBL) onto the surface of a sample in a vacuum chamber. These electrons raster-scan the surface, reflecting off of it as well as sometimes ejecting a secondary electron from the material on the surface of the sample. Detectors sense these electrons to reconstitute an image. Different types of detectors can be employed in an SEM. Some are more sensitive to relative heights on a sample, while others target the secondary electrons and so are more capable of distinguishing constituent materials. Most SEMs have access to multiple detectors that can be toggled while imaging. The beam rasters the surface so quickly that the image appears in a fraction of a second; the scan speed and averaging methods can also be altered to find a reasonable balance between a short wait time and an increased signal-to-noise ratio.

There is another important consideration when imaging a sample using an SEM. The electron beam is constantly pumping electrons onto the sample. In principle, the beam could damage the sample, although this is rarely an issue save for specific cases. The real problem is that these electrons need an escape route from the sample. Otherwise, the surface of the sample collects these charges, deflecting subsequent electrons and severely distorting the resulting image. This charging effect happens on the order of a fraction of a second, making it challenging to obtain a well-focused image.

One solution to this problem is just to set the acceleration voltage of the electron beam very low (~1 kV). This often works at the cost of a lower-resolution, grainier image. Some SEMs can flow gas onto the sample during imaging, which acts to compensate for the charging effect. An alternative solution is to deposit a conductive coating onto the surface of the sample. Usually, sample preparation rooms will have a small sputterer for precisely this purpose, with just the ability to deposit a thin (5–10 nm) layer of gold. This surface needs to have a direct connection between the surface of the sample and the “ground,” i.e., the conductive sample holder. Often, this deposited surface will spoil the sample—for example, gold can diffuse into silicon and disrupt its conductive and optical properties. It also cannot be deposited and removed from a biological sample in

a biocompatible manner. Thus, if a sample is being prepared to be imaged in the SEM, it can be considered a destructive measurement.

Profilometry

At its heart, a *profilometer* is meant to supplement the SEM in characterizing microscopic devices. While an SEM excels at measuring lateral dimensions, it falls short at measuring heights and depths. There is a vast library of profilometry tools available. The simplest profilometer is a cheap tool consisting of a needle on a cantilever that operates using the same principle as a record player or gramophone. A profilometer can make coarse measurements for a range of step heights, down to tens of nanometers and up to tens of microns. It has very poor lateral resolution, so it cannot measure the width of any features on a sample with any reliability. Also, since the stylus makes contact with the surface of the sample, it will often blemish that region of a sample. If this is to be avoided, there exist noncontact profilometers that work on the principle optical interference, but these depend on certain optical properties of the materials that might be beyond control.

Atomic force microscopy

An *atomic force microscope* (AFM) is a high-resolution stylus profilometer, one type of *scanning probe microscope* (SPM). It has the capability to raster-scan a surface and provides stunning true three-dimensional images with the highest available resolution, down to a fraction of a nanometer (**Figure 17.10**). AFMs also come with a noncontact mode, which leaves the sample surface untouched and clean. The main drawback of the AFM is how slow it is: because it scans by physically moving around a needle, it can take several minutes to complete an image. Also, unlike an SEM, an AFM has no way of distinguishing materials and is only capable of measuring relative heights. Either way, commercial AFMs have come a long way to becoming turnkey and user friendly, just like the SEM.

The AFM tip itself is always user replaceable because of its fragility. Tips come in a range of sizes and qualities (**Figure 17.11**). It is critical to know the properties of the tip used during a measurement since the results will be heavily influenced by tip choice. Tips also come with different conductive and mechanical properties. They range in price from a few dollars to several hundreds of dollars each, so knowing the required parameters is essential before purchasing a set.

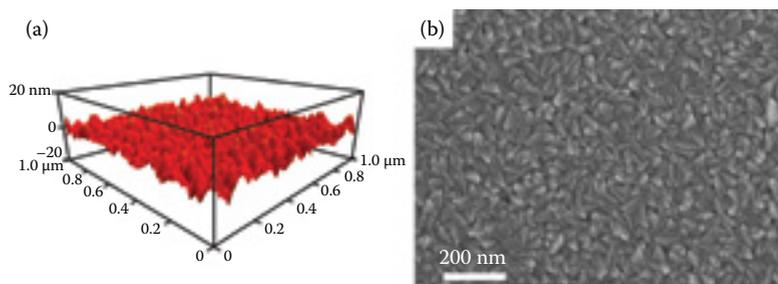


Figure 17.10 (a) AFM is used to supplement (b) SEM to measure profile data and extract a surface roughness measurement. (Figure courtesy of Jonathan D. B. Bradley.)

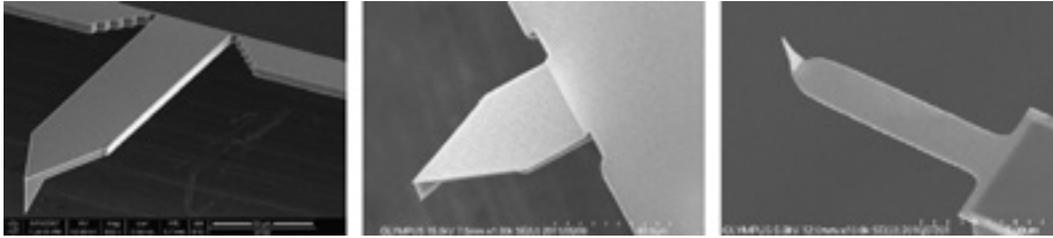


Figure 17.11 SEM images of different kinds of commercially available AFM tips.

17.6 THIN-FILM CHARACTERIZATION

Just as fabricated surfaces should be characterized once they are structured, thin films should also be evaluated to make sure that they possess the desired properties. In fact, it should not even be taken for granted that a tool will deposit precisely the material that it claims to! Here, we present several metrology techniques that are specifically tailored to investigating unstructured thin films.

Ellipsometry

An *ellipsometer* is a metrology tool that is used to characterize thin films. It works by focusing a laser on the surface of a sample and then measuring the coherent reflection for different input and output polarizations of light. This information could be used to back-calculate the film thicknesses and optical constants (e.g., absorption and refractive index information). In a fabrication setting, they are most often used to determine the thickness of films for the optical properties have already been established. For example, most resists are delivered with optical constants to be used with an ellipsometer to characterize your resist thickness before exposure.

The simplest ellipsometer only performs a measurement at a single optical wavelength and is known as a single-point ellipsometer. This tool can only reliably measure thicknesses below one optical wavelength, $\sim 1 \mu\text{m}$. This is a known physical limitation of the tool that has to do with fundamental properties of the propagation of light. On the other hand, spectroscopic ellipsometers take multiple measurements at different wavelengths and incident angles. They can use this extra information in sophisticated modeling software to unambiguously determine the thickness of a thin film. With software that is advanced enough, an ellipsometer can provide a lot more information, such as the thickness of interim layers in a stack (e.g., silicon on silica on silicon) or the surface roughness, all without ever touching the surface of the sample.

X-ray photoelectron spectroscopy

X-ray photoelectron spectroscopy (XPS) is one of many tools that can be used to characterize a thin film. A calibrated x-ray source illuminates the sample, dislodging electrons from the surface. These electrons are counted as a function of the x-ray photon energy, which corresponds to the ionization energy of the electron, yielding a spectrum. This spectrum can help you figure out the elemental

composition of a deposited material. For example, suppose you reactively sputter titanium in an oxygen atmosphere. An XPS measurement will tell you what the titanium-to-oxygen ratio is, and can help you determine if you have deposited TiO_2 or Ti_2O_3 . The measurement is inherently nondestructive, but the tool often includes inert atoms, which can be used to bombard and thin the material and probe spectra as a function of depth.

Raman spectroscopy

Raman spectroscopy is one of a number of spectroscopic techniques that operate entirely on optical processes, where a laser illuminates a sample and the optical spectrum that is emitted by the sample is recorded. Several of these techniques were discussed in **Chapter 16**. Raman uses inelastic scattering from vibrations in a molecule to give information about vibrational states. Infrared, visible, and ultraviolet (UV) Raman are all used to probe different types of chemical bonds. As an example, both UV and visible Raman spectroscopy have been used to distinguish between different phases of TiO_2 , such as rutile, anatase, or amorphous. Spectra must be cross-referenced with literature values to confirm identity of a sample; increasingly large databases of Raman spectra are becoming available. Raman spectrometers are usually available in materials characterization facilities and departments of geology.

X-ray diffraction

In *x-ray diffraction* (XRD), a monochromatic beam of x-rays is used to illuminate a sample. It propagates through the material and diffracts into multiple beams. In this measurement, the sample behaves like a diffraction grating, and the distance between the atoms within the sample act as the spacing in the grating. The angles at which the output beams propagate inform us about the crystallinity of the sample. This is how crystal structures are determined, which won the Nobel Prize in Physics in both 1914 and 1915, among dozens of other Nobel Prizes in Physics, Chemistry, and Medicine that made use of XRD. See also **Chapter 6** for applications of XRD to protein structure determination.

For microfabrication purposes, XRD provides information about the crystallinity of deposited materials that is complementary to the other spectroscopy tools discussed in this chapter. One thing that the other tools *cannot* do is quantify the level of crystallinity of a sample. This is called a *rocking curve analysis*: a sample is tilted, and the width of the XRD order is measured as a function of angle. The highest-quality materials with the most ordered crystals will produce the sharpest peaks. Thus, XRD is an indispensable tool for developing a recipe to deposit material epitaxially.

Four-point probe

A *four-point probe* is an ohmmeter that can be used to measure the resistance or conductivity of a thin film. It yields resistance values in units of Ω/sq (i.e., unitless area). When assessing multiple deposition recipes for a metal or conductor, the sample with the lowest resistance will likely have the purest, highest-quality material. It is a foolproof device that is found in every cleanroom and that provides yet another angle from which a deposition recipe can be assessed.

17.7 KEEPING A SAMPLE CLEAN

To conclude this chapter, we will provide some commonsense advice (often called *lore*) that has been acquired over time, through many failures and some personal experience. Following these guidelines should increase your yield and make you a more successful “fabber.”

Yield yield yield

You will *never* get to the point where every device you fabricate will be a success. Something unexpected always goes wrong, from accidentally dropping the sample facedown to an etcher malfunctioning and entirely ruining a sample. The solution to this problem is to prepare multiple samples in parallel. Prepare four samples in parallel, and only process one or two through the final step. If you need to develop a particular step in the recipe, say an etch recipe, then prepare 5 or 10 samples and try something different on each sample. *Never* just process only a single sample at a time.

Minimum feature sizes

As a rule of thumb, smaller features are, in general, harder to make. If your application permits it, it is to your benefit to redesign your pattern with larger and less dense structures. This will greatly increase your yield and reduce your fabrication cost, as well. Minimum feature sizes above 5 μm are preferred, though even keeping above 1 μm makes a substantial difference.

Mind your tolerances

It is paramount to understand what the tolerances are for every fabrication and metrology step. It is still impossible to obtain a structure that is accurate to within 1 nm in every dimension, as well as to perfectly characterize that structure. Try and design your device to be insensitive to these kinds of fabrication imperfections, with tolerances to dimension variations of up to at least 5%. Doing so will greatly increase your yield, as well.

Dedicated labware

Once you have developed a complete and working fabrication process, from the deposition of the device layer, through patterning and structuring, to deposition of the final protective cladding layer, you will likely repeat this entire procedure multiple times for different devices. At this point, you should purchase new glassware and dedicate them to the various steps of the process. For example, you should have a set of “lift-off” glassware, a set of “substrate cleaning” glassware, etc. This ensures that earlier steps in the process are not contaminated by dirtier steps downstream. You should also have dedicated tweezers for certain processes. For instance, your “gold” tweezers should never be used to process silicon samples, as gold dissolves in silicon, as mentioned previously. The cost of these elements will add up, but in the long run, they will dwarf in comparison to having to restart the entire process multiple times.

“Nothing goes in the bottle”

A bottle of resist can be relatively expensive to replace, so you should make every effort to keep it from getting contaminated. This is why you should follow the “nothing goes *in* the bottle” rule. Instead of pipetting directly from the resist bottle, decant some resist into a smaller bottle, and pipette for your daily process from there.

AMI wash, RCA clean, piranha etch

There are a few standard cleaning processes that use commonly accessible solvents and acids. The simplest is called an AMI rinse, where the sample is ultrasonicated in acetone, methanol, and isopropanol for 5 min each. It is imperative that isopropanol be the last thing to touch the sample since it evaporates cleanly without leaving a residue. Following this, the sample is blow-dried using a dry nitrogen gun and baked on a hot plate. This is likely to be enough to clean most samples. More aggressive processes (specific to silicon substrates) are the RCA clean and the piranha clean, which aggressively target surface contaminants such as oxide layers and organics.

Descumming

Cleanrooms often have a small dedicated dry etcher that flows oxygen plasma over a sample for the purpose of removing resist and organic material. This step is called descumming, ashing, stripping, or plasma cleaning. The multiple names stem from the fact that the operating principle can take many forms, but the general idea is the same. For example, in ozone ashing, oxygen is gently flowed over a sample under an intense deep UV source. The UV ionizes oxygen to create ozone as well as any carbon compounds on the surface. They bond to form CO₂ that drifts away with the airflow. A descummer should be used as part of a sample preparation, for example, after every AMI wash.

Take your time

Do not rush through your process. Cleanroom fabrication work is expensive and requires exacting care and patience. If you have allotted a certain amount of time to do something, and you are running behind, it would be better to stop what you are doing and come back another day than to rush through and risk spoiling a sample that might have taken a few days to prepare. A staffer once said to me, “If you don’t have the time, what are you even doing here?” In the same vein, do not skimp on cleaning thoroughly in the earlier steps of the process. Otherwise, something will go wrong down the line, and you will not know what caused it.

Be nice to the cleanroom staff

The cleanroom staff’s job is actually to maintain equipment, not to help you with your research. However, solving your obscure research problem is often the most exciting part of their day, so they might end up giving you a whole lot of unpaid help anyway. These technicians will be intimately familiar with the abilities and

history of their equipment and thus often can come up with creative solutions. Be sure to ask for their advice and take advantage of their experience. Be incredibly gracious with them; at the end of the day, you will be dependent on their happiness to get your work done.

Before experimenting with a new recipe, recreate something you KNOW will work

During your first few weeks in the cleanroom, you are probably not going to be developing a process from scratch. Instead, you will likely be asked to adapt an established process from a senior student or postdoc. Before experimenting with your own process, first recreate the old work. Otherwise, you will not know when you inevitably fail if the cause was your adjustment to the procedure or some misunderstanding of the original recipe. Also, while you are working toward perfecting a process, take images of your structure during the intermediate steps. Ideally, you can image your sample with an SEM as well as an optical microscope. If something goes wrong, you can track which step caused the problem.

Once you have finally recreated the old recipe, note that if you now tweak an earlier part of the process, you will often need to reoptimize the subsequent steps in the process, as well. This is why experienced fabbers might end up working with inefficient processes that already work over untested new processes.

Keep your toolbox properly outfitted

Start collecting spare pieces of labware in a toolbox—you never know when you will need them in a pinch. Extra pens, sample holders, razor blades, Scotch tape, double-sided carbon tape, kapton tape (for use in vacuum chambers), diamond scribes, spare silicon wafers, crystal bond... these are all incredibly cheap items that can save a fabrication day.

Spinning resist on small samples

If you have a sample that is too small to spin resist on reliably, mount it onto a larger silicon wafer with a small piece of tape, and spin the sample when on top of the wafer. Touch the tape with a gloved hand a few times to reduce its stickiness, or else you might find yourself in trouble when trying to remove your sample after the spin.

Keep your surroundings clean

A droplet of clear liquid on the wet bench can be anything, including water (which is harmless) or hydrogen fluoride (which is not). Just because *you* know that it is water does not mean that everyone else can figure it out. Everyone has access to some pretty scary chemicals in a cleanroom! Do yourself and everyone else in the cleanroom a huge favor and keep every bench and tabletop clean and sparkling. Also, label every piece of glassware and bottle with the chemical they contain. All of these little steps will add a little bit of time to your recipe development but will save time in the long run while making the entire process cleaner and safer as a result.

Double check

- The tone of your resist
- The safety procedures (Safety Data Sheet, called the SDS) for every new chemical you use
- Where to put your waste chemical when you are done with it *before* you pour any out
- If your tools are chemically compatible with the chemical you are about to use (e.g., HF dissolves glass, so use a plastic container and plastic tweezers)

17.8 FINAL COMMENTS

I hope you have enjoyed reading this book, found it practical, and will get stains on it, rip pages out, and write to me with your corrections, additions, comments, and questions. This is a work in progress that will be updated as new techniques emerge and are adopted by more and more researchers, and your feedback will help keep it up-to-date for years to come.

Background Reading

Books

There are few books currently available on the subject of nanofabrication.

Ayers, J.E., Kujofsa, T., Rago, P., and Raphael, J. *Heteroepitaxy of Semiconductors: Theory, Growth, and Characterization*. Edn 2. CRC Press, Boca Raton, FL, 2016.

Gerlach, G., Dotzel, W., and Muller, D. *Introduction to Microsystem Technology*. Limited edition. John Wiley & Sons, West Sussex, England, 2008.

Plummer, J.D., Deal, M.D., and Griffin, P.B. *Silicon VLSI Technology: Fundamentals, Practice, and Modeling*. Edn 1. Pearson, Prentice Hall Upper Saddle River, New Jersey, 2000.

Journal articles

This is by no means an exhaustive list but should give some idea of the methods, applications, and, in some cases, history of the techniques in this chapter.

Berg, S., and Nyberg, T. (2005). Fundamental understanding and modeling of reactive sputtering processes. *Thin Solid Films* 476(2), 215–230.

Bradley, J.D.B., Evans, C.C., Choy, J.T., Reshef, O., Deotare, P.B., Parsy, F., Phillips, K.C., Lončar, M., and Mazur, E. (2012). Submicrometer-wide amorphous and polycrystalline anatase TiO₂ waveguides for microphotonic devices. *Optics Express* 20(21), 23821–23831.

Burek, M.J., de Leon, N.P., Shields, B.J., Hausmann, B.J.M., Chu, Y., Quan, Q., Zibrov, A.S., Park, H., Lukin, M.D., and Lončar, M. (2012). Free-standing mechanical and photonic nanostructures in single-crystal diamond. *Nano Letters* 12(12), 6084–6089.

Griffiths, M.B.E., Pallister, P.J., Mandia, D.J., and Barry, S.T. (2016). Atomic layer deposition of gold metal. *Chemistry of Materials* 28(1), 44–46.

Hung, Y.J., Lee, S.L., Thibeault, B.J., and Coldren, L.A. (2011). Fabrication of highly ordered silicon nanowire arrays with controllable sidewall profiles for achieving low-surface reflection. *IEEE Journal of Selected Topics in Quantum Electronics* 17(4), 869–877.

Kontio, J.M., Husu, H., Simonen, J., Huttunen, M.J., Tommila, J., Pessa, M., and Kauranen, M. (2009). Nanoimprint fabrication of gold nanocones with ~10 nm tips for enhanced optical interactions. *Optics Letters* 34, 1979–1981.

Kudryashov, V., Yuan, X.-C., Cheong, W.-C., Radhakrishnan, K. (2003). Grey scale structures formation in SU-8 with e-beam and UV. *Microelectronic Engineering* 67–68, 306–311.

Nagpal, P., Lindquist, N.C., Oh, S.-H., and Norris, D.J. (2009). Ultra-smooth patterned metals for plasmonics and metamaterials. *Science* 325(5940), 594–597.

Saklayen, N., Huber, M., Madrid, M., Nuzzo, V., Vulis, D.I., Shen, W., Nelson, J., McClelland, A.A., Heisterkamp, A., and Mazur, E. (2017). Intracellular Delivery Using Nanosecond-Laser Excitation of Large-Area Plasmonic Substrates. *ACS Nano*, 11(4), 3671–3680.

Van Laer, R., Kuyken, B., Van Thourhout, D., and Baets, R. (2015). Interaction between light and highly confined hypersound in a silicon photonic nanowire. *Nature Photonics* 16(February), 1–5.

Vogel, N., Zieleniecki, J., and Koper, I. (2012). As flat as it gets: Ultrasmooth surfaces from template-stripping procedures. *Nanoscale* 4(13), 3820–3832.

Vu, K.T., and Madden, S.J. (2011). Reactive ion etching of tellurite and chalcogenide waveguides using hydrogen, methane, and argon. *Journal of Vacuum Science and Technology A: Vacuum, Surfaces, and Films* 29(1), 11023.

Wilkinson, C.D.W., and Rahman, M. (2004). Dry etching and sputtering. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 362(1814), 125–138.

Williams, K.R., and Muller, R.S. (1996). Etch rates for micromachining processing. *Journal of Microelectromechanical Systems* 5(4), 256–269.

Williams, K.R., Gupta, K., and Wasilik, W. (2003). Etch rates for micromachining processing—Part II. *Journal of Microelectromechanical Systems* 12(6), 761–778.

Equipment and supplies

Chemicals

Thermal grease

Santovac 5 (SantoLubes)

Positive-tone resist

ZEP520A (Zeon Chemicals)

PMMA (MicroChem)

Shiplay (Dow)

Negative-tone resist

SU-8 (MicroChem)

HSQ (Dow)

Fox (Dow)

XR (Dow)

ma-N 2400 (Microresist Technology)

Instruments

J. A. Woollam Co., Inc, Gaertner Scientific. Ellipsometers and measuring microscopes.

Raith nanofabrication, JEOL, Elionix. Electron beam lithography.

JEOL, Zeiss. SEMs.

Park Systems, Asylum Research, Veeco, Bruker. AFMs, surface profilometers.

AJA International, Kurt J. Lesker. Thin-film deposition.

SPP Process Technology Systems Ltd. Etching and chemical vapor deposition.

Ultratech/CNT. Atomic layer deposition.