PATTERN TRANSFER FROM NANOPARTICLE ARRAYS

BY

CHARLES R. HOGG III

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Physics
in the Mellon College of Science of
Carnegie Mellon University, 2010

Pittsburgh, Pennsylvania

Doctoral Committee:

Prof. Sara A. Majetich
Prof. James A. Bain
Prof. Stephen Garoff
Prof. Michael Widom

Abstract

This project contributes to the long-term extensibility of bit-patterned media (BPM), by removing obstacles to using a new and smaller class of self-assembling materials: surfactant-coated nanoparticles. Self-assembly rapidly produces regular patterns of small features over large areas. If these patterns can be used as templates for magnetic bits, the resulting media would have both high capacity and high bit density. The data storage industry has identified block copolymers (BCP) as the self-assembling technology for the first generation of BPM. Arrays of surfactant-coated nanoparticles have long shown higher feature densities than BCP, but their patterns could not previously be transferred into underlying substrates. I identify one key obstacle that has prevented this pattern transfer: the particles undergo a disordering transition during etching which I have called "cracking". I compare several approaches to measuring the degree of cracking, and I develop two novel techniques for preventing it and allowing pattern transfer. I demonstrate two different kinds of pattern transfer: positive (dots) and negative (antidots). To make dots, I etch the substrate between the particles with a directional CF_4-based reactive ion etch (RIE). I find the ultrasmall gaps (just 2 nm) cause a tremendous slowdown in the etch rate, by a factor of 10 or more — an observation of fundamental significance for any pattern transfer at ultrahigh bit densities. Antidots are made by depositing material in the interstices, then removing the particles to leave behind a contiguous inorganic lattice. This lattice can itself be used as an etch mask for CF_4-based RIE, in order to increase the height contrast. The antidot process promises great generality in choice of materials, both for the antidot lattice and the particles themselves; here, I present lattices of Al and Cr, templated from arrays of 13.7 nm-diameter Fe_3O_4 or 30 nm-diameter MnO nanoparticles. The fidelity of transfer is also noticeably better for antidots than for dots, making antidots the more promising technique for industrial applications. The smallest period for which I have shown pattern transfer (15.7 nm) is comparable to (but slightly smaller than) the smallest period currently shown for pattern transfer from block copolymers (17 nm); hence, my results compare favorably with the state of the art. Ultimately, by demonstrating that surfactant-coated nanoparticles can be used as pattern masks, this work increases their viability as an option to continue the exponential growth of bit density in magnetic storage media.

To my father, the first Dr. Charles R. Hogg,
whose unfailing encourangement and faith in me
helped make me the second

Acknowledgements

This work has been supported by more people than I can list, in more ways than I can count. Nevertheless, I would like to mention just a few of them.

First, I could not even have started without financial support. I am grateful to the department for the opportunity to work as a teaching assistant, since education is a passion of mine, and to the National Science Foundation for funding our grants. Most of all, I am grateful to Dr. Bruce and Mrs. Astrid McWilliams, for the generous fellowship they awarded to me. Their selfless example inspires me, and I hope I can also make a difference should the opportunity arise.

When my advisor, Sara Majetich, let me in the group, my mode of learning was very coursework-oriented: learn what others have done, and play it back on cue. Through her patient mentoring, I learned what it is to be a good scientist: how to take the *next* step, and find things out that nobody's ever learned before. She also managed to strike that too-rare balance between the professional and the personal, and clearly communicated the importance of both aspects. The debt I owe her can never be paid back, but I hope I can pay it forward to those who come after me.

I have worked so closely with Jim Bain that he has been almost like a second advisor. His example has taught me that being realistic doesn't mean you can't be optimistic, and that partial success is still success as long as you keep pushing. No sooner does he see an obstacle than he starts thinking of ways around it, and his ever-ready order-of-magnitude calculations always steered us towards the most promising directions. Over the years, Steve Garoff and Mike Widom have given valuable feedback on my project, and have been a needed source of fresh perspectives. I am deeply grateful to everyone on my thesis committee for their guidance and support.

Several of my fellow students deserve a special mention. Madhur Sachan and Craig Bonnoit got this project started, and sowed the seeds for its subsequent success. I could always count on Eric Evarts for intelligent commentary, and his scanning probe wizardry often led to new insights. Ryan Booth was always willing to go the extra mile when I needed measurements, or just a helping hand. Jitkang Lim was a great officemate during the odd, long hours he worked, and his work ethic was a great example to me. (Also, he and Ryan introduced me to the wonderful world of coffee!)

I want to acknowledge three more people who have gone before me in achieving their doctorates. The arrival of our postdoc, Dr. Lubna Shah, has put me at

ease about the future of this project. Seeing first hand her organized scientific approach, together with her determined work ethic, I am convinced I could not have left my project in more capable hands. Dr. Matt Moneck has helped me in more ways than I can count, through his encyclopedic knowledge of nanofabrication techniques, extensive hands-on experience, and cheerful willingness to help. I suspect many more thesis projects than just mine would not have been successful without his presence. Finally, Dr. Yi Luo provided the initial RIE recipes that put me in the right ballpark for etching fine features, together with an explanation of how to alter them depending on the results.

Several staff have been instrumental in keeping things flowing smoothly on a day-to-day basis. First and most important are the cleanroom staff: Chris Bowman, Carsen Kline, Tim Fisher, and Mike Sakaluk. I am still amazed by how smoothly they keep things running for so many different projects, and by their can-do attitude towards any problem I threw at them. I would also like to thank the Physics department staff — particularly Chuck Gitzen, Al Brunk, Hilary Homer, Mary Jane Hutchinson, Donna Thomas, Mary Placeway, Patrick Carr, Laura Youree, and Montel Rudolph — and the DSSC staff — especially Susie Cribbs, Matt Koeske, and Pat Grieco.

The support of my family has added a richness to my life that is difficult to describe. I do not think I could have asked for better parents than Dr. Charles and Mrs. Cynthia Hogg. Not until I grew up did I really appreciate the rare gifts they gave me, that made me the man I am today. Together with my siblings — John, Marina, Laura, and Vera — they gave me a home life that I treasure, now and always. During my time at Carnegie Mellon, I also got married, and the new family I gained — John, Sue, and Greg Lambert — have made me feel welcome me as a son and a brother. Furthermore, I have been blessed beyond measure by my two small children, Sophia and Robbie. They remind me that the work described herein is *not* my happiest achievement in the past five years.

Finally, I have to thank my loving wife Amy. More than to any other single person, I owe the success of this project to her. Every time I came to her with a problem, she was always *more* supportive than I expected. If not for her, I would have quit when things got rough. If not for her, I would still think scientific research was "something I'm just not good at", rather than a skill I could acquire through hard work. Today, I am *Dr.* Hogg because she believed in me — and backed up that belief by her actions, day after day. I owe her more than I know how to say.

Table of Contents

List of Tables

List of Figures

List of Abbreviations

AAO	anodized aluminum oxide
ABD	areal bit density
ARDE	aspect ratio dependent etching
BCP	block copolymers
BPM	bit-patterned media
BPMR	bit-patterned magnetic recording
BSE	backscattered electron
DC	direct current
DFT	discrete Fourier transform
ESS	electrostatically stabilized
EUV	extreme ultraviolet [lithography]
FCC	face-centered cubic
GISAXS	grazing-incidence small-angle x-ray scattering
HSQ	hydrogen silsesquioxane, a high-resolution e-beam resist
ICP	inductively coupled plasma
MOKE	magneto-optical Kerr effect
NA	numerical aperture
NGL	next-generation lithography
NSL	nanosphere lithography
NSV	neighbor spacing variation, a measure of cracking I developed
OA	oleic acid
OY	oleyl amine
PMMA	poly(methyl methacrylate), a common e-beam resist
RF	radio frequency
RIE	reactive ion etching
SE	secondary electron

SEM	scanning electron microscope
SFIL	step-and-flash imprint lithography
SNR	signal-to-noise ratio
SPL	structural pinning layer
SQUID	superconducting quantum interference device (a type of magnetometer)
STS	Surface Technology Systems, a manufacturer of RIE instruments
TEM	transmission electron microscope
UV	ultraviolet
XPS	x-ray photoemission spectroscopy
XSEM	cross-sectional scanning electron microscopy
XTEM	cross-sectional transmission electron microscopy

1 Introduction

1.1 Motivation

The focus of this project is the future extensibility of *lithography* — the patterning of surfaces to make functional materials. As always, that future is bound up with *smaller* patterns, since the shrinking of devices brings desirable increases in both portability and functionality. Present techniques experience ever more daunting challenges, which are nevertheless far short of the *ultimate* lithographic limitation: the single-atomic scale. Candidates to bridge this gap are known as *next-generation lithographies* (NGL), and the present work contributes to this area.

Applications for lithography are too numerous to cover in detail. For definiteness, I focus on two industries: semiconductors, and data storage.

The semiconductor industry has long been the largest consumer of lithography, and has driven advances in the field for decades. Both transistors and interconnects are lithographically defined. In 1965, Gordon Moore noticed that the number of transistors per chip doubled roughly every 18 months, and predicted that trend (now known as "Moore's Law") would continue for another decade.[110] In fact, it has continued for more than four decades, and was even formalized by the industry as a whole into a roadmap.[6] Packing more transistors per chip implies the ability to define ever-smaller dimensions lithographically. In the semiconductor industry, these take the form of "technology nodes", characterized by the smallest half-pitch H of features defined. For instance, all major chipmakers have released products at the $H = 45\,\mathrm{nm}$ node (implying a periodicity of $2H = 90\,\mathrm{nm}$), and the transition is now being made to the $32\,\mathrm{nm}$ node.[46]

The other lithography consumer I consider is the data storage industry, because this is the main application in mind for the NGL proposed here. A key figure of merit is the *areal bit density* (ABD), a measure of the number of logical bits which can be stored in a given surface area. The modern industrial unit is the *terabit per square inch* (Tbpsi), equal to 10^{12} bits for every square inch of area. Like the semiconductor industry, data storage has enjoyed sustained exponential growth over decades, doubling every two years since the 1960's (and *every year* from 1998-2000). This growth has *not* been driven by lithography; in fact, their lithographic needs have traditionally been modest, being mainly confined to the manufacture of magnetic heads for reading and writing magnetic

bits. However, this industry has encountered a crisis, known as the *superparamagnetic limit*, which threatens the continued growth of the ABD in traditional media.[24]

Conventional storage media are composed of multiple *grains* separated by nonmagnetic boundaries. Each grain functions as an independent magnetic unit which can be magnetized in one of two directions, taken to symbolize logical bit values of "1" or "0". These grains are random in size, shape, and position, causing unacceptably high variations in their individual properties (such as their readback signal, or the magnetic field strength required to overwrite them). For logical *bits* to have uniform properties, they must therefore contain *numerous* grains, so that the fluctuations average out. This strategy is depicted in Figure 1.1(a).

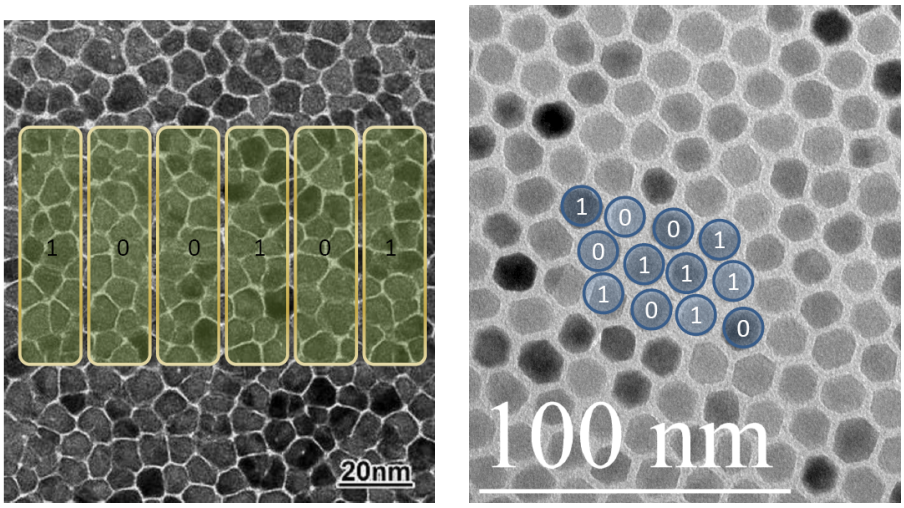

(a) **(b)**

Figure 1.1: Comparison of conventional and bit-patterned media paradigms. (a) The magnetic units of conventional media are grains which are random in size, shape, and position. To get decent signal-to-noise, each "bit" averages over multiple grains. (b) Bit-patterned media proposes to make the bits *predictable* in size, shape and position, which could eliminate the need for averaging.

Traditionally, ABD has been increased mainly by engineering smaller grains, keeping the number of grains per bit stable across generations. However, lowering the bit volume V increases susceptibility to thermal fluctuations, as follows. Mean data lifetime τ_n varies according to an Arrhenius law,

$$\tau_n = \tau_0 \exp\left(\frac{K_u V}{k_B T}\right), \tag{1.1}$$

where τ_0^{-1} is an intrinsic "attempt frequency" (typically 10^9 Hz to 10^{10} Hz [105]), and $k_B T$ is the thermal energy. K_u is the *magnetic anisotropy*, a measure of the preference to be magnetized along a given axis, which also determines the difficulty of magnetic switching. Acceptable data lifetimes are convention-

ally considered as $\tau_n = 10$ years, implying that $(K_u V/k_B T)$ should be at least $\ln(\tau_n/\tau_0) \approx 40$. Engineering materials with higher K_u can thus restore stability, but at the cost of *writability* by practically attainable head fields. Hence, continued increases in ABD are blocked by a *trilemma*[139]: among writability, thermal stability, and signal-to-noise ratio (SNR), only two can be made acceptable simultaneously.

Several strategies have been proposed to circumvent this trilemma, two of which have received explicit backing from the storage industry.[61] *Heat assisted magnetic recording* (HAMR) is one candidate, which attacks writability by incorporating local heating. A laser heats the bit, temporarily lowering its write field. An reasonably-attainable magnetic field is applied to the bit, which switches and remains stable after cooling. This opens the way for small-grain, high-K_u materials such as L_{10}-phase FePt, which is the near-unanimous material choice among HAMR researchers in industry.

The second candidate, *bit-patterned media* (BPM), is our main intended application. BPM attacks the SNR aspect of the trilemma, by drastically lowering the number of grains required for a low-noise bit — perhaps down to a single grain. The need to average over many grains stemmed from the wide fluctuations unpatterned grains exhibit. If instead the medium featured *identical* magnetic units at *predictable* positions (as shown in Figure 1.1(b)), the timing and magnitude of readback and write signals could be known in advance with high precision. Making these identical units necessarily involves *patterning* the surface, and so lithography has entered the world of data storage.

It is instructive to compare the lithographic demands of these two industries. Consider a square region, $1''$ (2.54×10^7 nm) per side, which contains one terabit (10^{12} bit) in a square array — 10^6 rows, with 10^6 bits per row — for a bit density $\rho_0 = 1$ Tbpsi. Each bit thus has a diameter of $2H_0 = 25.4$ nm, for a half-pitch of $H_0 = 12.7$ nm. More generally, the areal bit density varies implicitly with half-pitch according to

$$\rho H^2 = F \rho_0 H_0^2, \tag{1.2}$$

where $\rho_0 = 1$ Tbpsi and $H_0 = 12.7$ nm are convenient reference values, and F is a factor based on the geometry of the array: $F = 1$ for square arrays, and $F = 2/\sqrt{3}$ for hexagonal arrays. Since all the patterns I make are hexagonal, I will implicitly consider the $F = 2/\sqrt{3}$ case in Equation 1.2 when converting between half-pitch and bit density.

Conventional recording demonstrations have approached, but not exceeded, the symbolic 1 Tbpsi mark. This is thus a convenient minimum density for any candidate NGL technique to achieve in order to be taken seriously. Equation 1.2 yields an equivalent half-pitch of 14.7 nm. A close but more modest semiconductor node, 16 nm, is not targeted until 2016![76] Moreover, the data storage industry hopes for densities of $\rho_1 = 4$ Tbpsi by early 2013 and $\rho_2 = 10$ Tbpsi by the end of 2015. The corresponding half-pitch figures of $H_1 = 7.3$ nm and

$H_2 = 4.6\,\mathrm{nm}$ do not even appear on the semiconductor roadmap. Thus, the lithographic needs of the storage industry already far outstrip the needs of the semiconductor industry.

At the outset, I emphasize that these stringent lithographic needs are only one part of a larger picture. Before bit-patterned media can hold the family photos of everyday consumers, issues such as reading, writing, timing, bit rate, error correction, and more will also have to be solved. Nor are these challenges rigidly indipendent from the lithographic processes investigated herein — for example, alternative read/write processes based on spin-transfer torques depend critically on a precise multilayer combination of diverse materials, and any patterning process must be compatible with all these materials. Nevertheless, the foundations must still be laid for lithography with ultra-small nanoparticles, and that initial exploration is the aim of this project.

1.2 Lithographic approaches

This section compares the most important lithographic approaches for industrial applications. The picture which emerges is of a single dominant technique — photolithography — whose demise has long been foreseen due to both practical and fundamental limitations. Various younger technologies are candidates to replace it, but none has yet mounted a successful challenge. This is primarily due to the remarkable resilience of photolithography, whose practical limitations have been stretched ever closer to its fundamental limitations. This extension has left very little room for competing processes. Due to economic reasons, it is not sufficient for a competing technology to go just one or two generations beyond photolithography. It must promise enough extensibility so that, just like with photolithography, clever engineering can extend its usefulness down to many generations of smaller structures.

For each technique, I will give a brief overview of its history, development, and important milestones. I will examine both the actual lateral resolution achieved so far, and the resolution it is expected to achieve. Finally, the techniques will be summarized with an eye towards the future development of industrial patterning.

1.2.1 Photolithography

Optical lithography has been the overwhelmingly dominant technique for the entire history of industrial lithography. That history begins in the late 1950's, with the production of the first integrated circuits.[172, 46] As the underlying processes were understood in greater detail, future limitations became widely recognized, and these motivated development of a sequence of candidate successors. In every case, these candidates lost out to the continued and remarkable improvement of optical lithography. The way modern circuits are fabricated is

conceptually identical to the way the first integrated circuits were made in 1958, with only incremental changes to the process.

The flow of optical lithography is depicted in Figure 1.2. It begins with the creation of a *mask*: a transparent plate which contains an opaque pattern, typically realized with a Cr layer tens of nanometers thick.[103] The sample is coated in *photoresist* — a polymer which changes solubility when exposed to light. Light passes through the mask and forms an image of the mask in the photoresist. To "develop" the latent image, the soluble parts are washed away, and a physical copy of the mask pattern appears in the photoresist which remains. This pattern is then permanently transferred, either by etching around the photoresist, or by depositing material and removing the photoresist (a process known as *liftoff*).

The resolution R — i.e. the minimum feature width which this method can produce — is governed by a well-known relation[16, 172, 75, 103, 169, 46]:

$$R = k_1 \frac{\lambda}{\mathrm{NA}}. \tag{1.3}$$

λ is the wavelength of the light. k_1 is a materials-based dimensionless constant, relating to the photoresist and developer characteristics, as well as optical tricks such as phase shift masks which improve the contrast. NA is the *numerical aperture*, proportional to the sine of the angle subtended by the lens at the focal point of the sample; it is thus less than 1.

The early days of lithography left much room for improvement in all these factors: in 1961, the resolution was 5 µm[103]; today, it is over 100× smaller at 45 nm. These gains stemmed from simultaneous aggressive optimization of all three factors on the right side of Equation 1.3. A succinct, detailed summary of the evolution of these parameters is found in Table 1 of [46]. Here, I sketch an overview.

The most intrinsically extensible strategy to improve resolution is to decrease the wavelength, because electromagnetic radiation can be produced with λ many orders of magnitude smaller than visible light. Difficulties include the construction of sufficiently powerful light sources at a given wavelength, and the λ-dependent optical properties of materials (for example, glass lenses are opaque to UV light). In 1976, optical lithography still used visible light: the G-line mercury emission, with a wavelength of 436 nm. By 1987 the I-line emission with $\lambda = 365$ nm was in use. 248 nm light from a KrF excimer laser enabled submicron features to be standard by 1993.[46] These have been replaced by 193.4 nm light from ArF lasers, which remains the standard today.

Other improvements have been based on k_1 and NA. k_1 has evolved from higher values such as 0.80 in the 1980s to as low as 0.27, through a combination of materials engineering and advanced optical techniques.[46] However, images cannot form for $k_1 < 0.25$[16], so there is little remaining room for improvement. Lowering k_1 requires tightening tolerances for focus and exposure

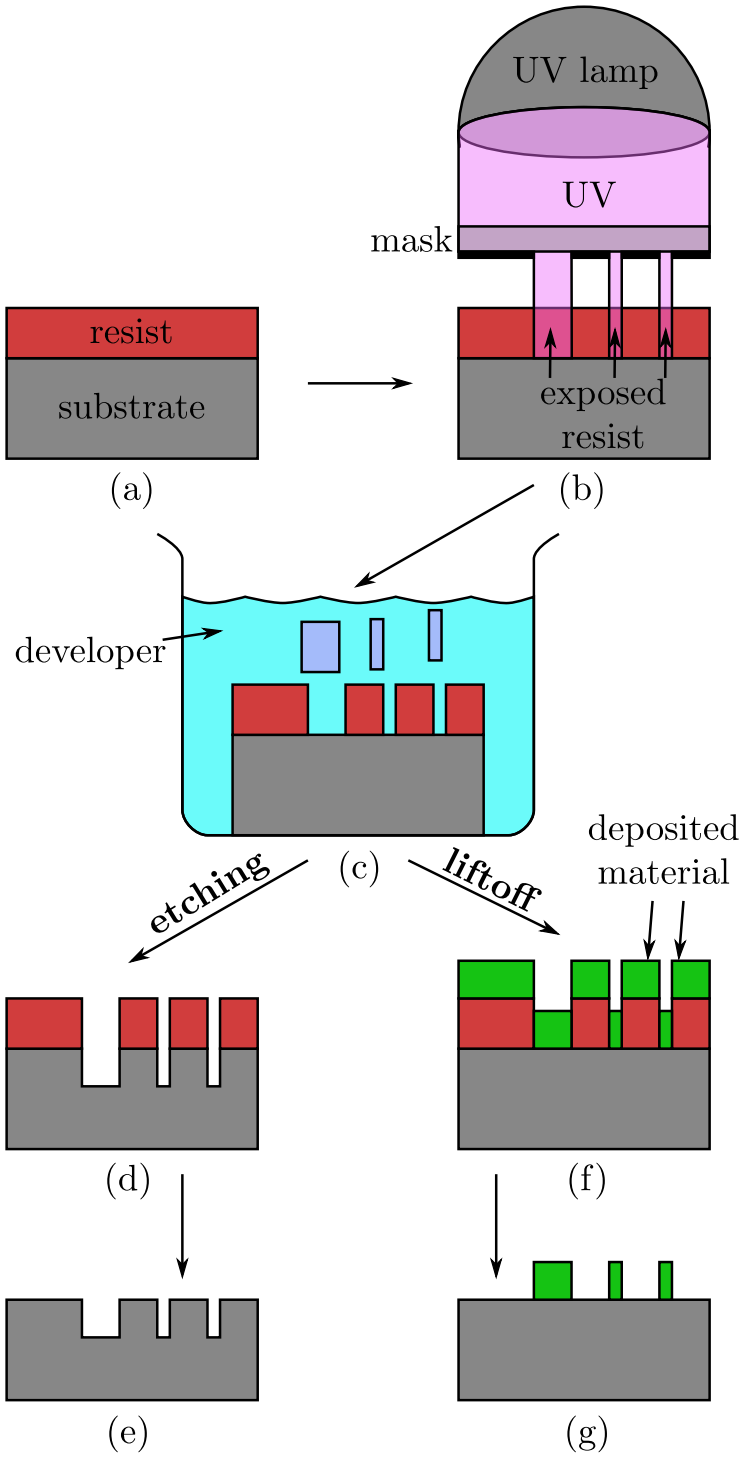

Figure 1.2: Schematic of optical lithography. (a) The substrate to be patterned (grey) is coated with a layer of photoresist (red). (b) A UV lamp exposes the sample through a mask, changing the solubility of the exposed areas. (c) The resist is exposed to a developer, and the soluble parts are removed. Here, the exposed parts are shown as soluble, but note that in some resists the opposite occurs. At this point, patterning may occur either by etching or liftoff. For etching, (d) the wafer is etched through the mask, and (e) the mask is removed. For liftoff, (f) material is deposited on the mask and the substrate, and (g) the mask is removed, taking the overlying material with it.

energy, but requirements are already at the level of nanometers for focus and millikelvin for temperature, making k_1 a daunting route for even small revolution improvement.[169] NA began relatively modestly, being around 0.16 in 1961.[103] Driven largely by advances in fabricating aberration-free wide-angle lenses, NA has grown as high as 0.93, very close to the theoretical limit of 1. Hence, only incremental improvements in resolution can be expected from further optimization of k_1 and NA; wavelength reduction is the only remaining option.

Here, there is a serious obstacle. The next step in wavelength reduction was supposed to be 157 nm light from an F_2 laser. This seemingly incremental step was proposed in 1997, and serious research began in 1999.[46] In 2003, industry abandoned the technology, because research on the specialized CaF_2 lens materials had not kept pace with the semiconductor roadmap.[46, 169] Standard lenses strongly absorb light at this wavelength, as do many other materials, and the advantages of lower wavelength are more than offset by the considerable engineering challenges they present.

Hence, any wavelength reduction significant enough to be non-trivial raises serious obstacles whose cost cannot be justified. Thus, it is overwhelmingly probable that 193.4 nm is the final wavelength for optical lithography.[46, 169] This section will conclude by examining several strategies which promise to extend the lifetime of optical lithography, before moving on to examine the alternatives offered by NGL techniques.

Immersion lithography

Seven years after the failure of 157 nm lithography, its 193.4 nm counterpart has continued to deliver improved resolution on the timetable of the roadmap. The key enabling innovation was the introduction of a refractive liquid between the final lens and the sample. This reduces the wavelength of the light as it exposes the photoresist, according to $\lambda = \lambda_0/n$, with $\lambda_0 = 193.4$ nm the vacuum wavelength of the light, and n the index of refraction at that wavelength. All the upstream components still experience light of wavelength λ_0, and therefore can remain unchanged.

This technology was proposed in 2002[141], and serious research began the following year[16]. As of 2008, several products using immersion lithography had entered production.[169] All current immersion lithography setups use water as the liquid, with $n = 1.44$ at $\lambda_0 = 193.4$ nm. High-index organics with $n \approx 2$ have been investigated, but enthusiasm dimmed considerably after researchers from Nikon pointed out several formidable obstacles to practical use.[169]

Double patterning

The final optical strategy which has already proven itself in production is known as *double patterning*. Actually, this is an umbrella term which covers several

different processes. Nearly all rely on the fact that features smaller than $k_1\lambda/\mathrm{NA}$ *can* be written, as long as they are spaced further apart than their widths. In other words, $k_1\lambda/\mathrm{NA}$ limits the *pitch*, and not the feature size. These diverse techniques are summarized in Figure 17 of French and Tran[46]. As of 2010, Intel has rolled out a few chips which reached the 32 nm node with 193.4 nm lithography through this technology.

Different double patterning techniques suffer from different disadvantages. Some approaches require every step to be repeated, from the photoresist application to its development and pattern transfer. This immediately doubles the number of machines needed on the assembly line, significantly increasing the cost.[169] Other techniques form structures on the sidewalls of a sacrificial original pattern, but these limit the kinds of structures which can be built. The most promising technique would expose two patterns sequentially into a single resist layer. This approach would be far more economical, but the nonlinear photoresists required have not yet been invented.[46, 97]

Extreme ultraviolet

Another candidate replacement is *extreme ultraviolet lithography*, or EUV. Usually, this is considered to be a "next-generation lithography" technique[172, 46, 97]. However, I am reserving the label "NGL" for more radically different techniques, and EUV is still based on fundamentally the same principles as optical lithography.

The main distinguishing feature of EUV is a drastic reduction in wavelength: from 193.4 nm down to 13.5 nm, near the boundary between high-energy ultraviolet and soft x-rays. Working at such a small wavelength should enable the fabrication of correspondingly smaller features. The major difference from (conventional) optical lithography is the need for fully-reflective optics, since this wavelength is absorbed too strongly by available materials for refractive lenses. Other than that, EUV is a logical extension of optical lithography, and accordingly has found much favor in the industry.[172]

In spite of its natural advantage, the history of EUV is fraught with delays. It is a radical enough change to cause serious technical difficulties, and continued improvements in optical lithography have repeatedly postponed its adoption. The process was first proposed in 1988[58], for a planned insertion at the 100 nm node.[172] In 1997, an industrial consortium was formed, and the insertion was pushed back to the 65 nm node (c. 2006).[172] Meanwhile, both the 65 nm and 45 nm nodes have come and gone, and optical lithography is still the exclusive approach to fabrication. EUV is currently targeted for the 22 nm node[169], but strong skepticism persists about its cost-effectiveness.[97]

The biggest obstacle is the inefficiency in delivering power from the source to the resist. At present, the power delivered is too low, which results in a low throughput that cannot justify the cost of the system.[45] However, even if

an acceptable throughput were achieved, the EUV sources would consume five times as much power as the rest of the fabrication facility combined![97] Clearly, this also increases the cost of adoption.

Defect-free masks are also an issue.[169] Another key issue is chamber contamination. Since EUV radiation is strongly absorbed, the source, optics, and wafer printing chambers must all be in the same vacuum.[172] This leads to visible contamination and degradation of the mirrors where they are impacted by the EUV radiation.[45]

Another obstacle for EUV lithography relates to quantum mechanics. The energy dose cannot be varied continuously; instead, it is quantized as a discrete number of photons. Each photon has an energy $E = hc/\lambda$, where h is Planck's constant and c is the speed of light. The number of photons fluctuates statistically according to the Poisson distribution, where a mean of N photons experiences fluctuations of $\sqrt{N}$. Hence, smaller numbers of photons lead to larger *relative* fluctuations $1/\sqrt{N}$. Because each EUV photon has more than 10 times the energy of a ArF laser photon, the dose fluctuations *for the same mean dose* will be more than $\sqrt{10} \approx 3.3$ times worse. It is amusing to note that the introduction of light quanta solved the "ultraviolet catastrophe" a century ago. Now, in ultraviolet *lithography*, those same quanta are causing a catastrophe of their own.

The solution to this problem is to engineer resists which require higher doses. Presently, a dose of $30\,\mathrm{mJ\,cm^{-2}}$ is expected to minimize the impact of shot noise. However, higher doses imply higher source power for the same throughput, and source power is already the key limiting factor for EUV.

In short, there is widespread consensus that EUV will take the place of optical lithography, but serious engineering problems make this far from certain. Its strong position stems mainly from process familiarity and the lack of ready alternatives.

1.2.2 Maskless techniques

One broad class of techniques is known as *maskless lithography*, since arbitrary patterns can be written directly into the sample without first appearing in a mask. Several of these methods will be explained in detail below. Although they are diverse, they all suffer from a common weakness for lithographic applications: since they are serial, their throughput is poor.

For this reason, none of these techniques are seriously considered in their current form as candidates for industrial NGL. However, all of them can be parallelized to varying degrees. The prospect of parallelization has led to increased visibility for maskless lithography in the latest semiconductor roadmap.[76] However, I emphasize that further research and development is required for any of them to meet the needs of mass production.

Two maskless lithographic techniques are particularly promising for ultra-

high resolution mass production: e-beam lithography, and He^+-ion lithography. These are discussed in more detail below. First, two others will be discussed briefly, to understand what they are and why they are unsuitable for the applications I have in mind.

Focused ion beam (FIB) is a mature technology which allows direct machining of arbitrary surfaces with a beam as narrow as 10 nm[60, 164], or even below.[51] The ions, typically Ga^+, remove material by sputtering. Applications extend far beyond lithography; FIB is used for liftout of site-selective cross-sectional TEM samples, failure analysis in integrated circuits, and controllable localized doping. Therein lies the problem for lithography: doping is possible because the ions are implanted in the underlying material, but this causes unacceptable damage to the crystal lattice. In addition, the doping cannot be avoided, so circuits mass-produced by this method would exhibit changes in their critical electrical properties. For these reasons, FIB lithography is unsuitable for mass-production of semiconductors. For the data storage industry, it moreover lacks the resolution to improve on even currently attainable bit densities.

Scanning probe lithography, by contrast, has literally unbeatable resolution: it has enabled arbitrary arrangements of atoms on surfaces[41, 31, 65], using scanning tunneling microscopy (STM). Another form of scanning probe microscope is the atomic force microscope (AFM), with slightly poorer resolution but wider materials compatibility. Broadly speaking, in scanning probe techniques, an ultrafine tip is moved across a surface at very small separation distances. A feedback signal — e.g. tunneling current for STM, or force-modulated oscillation frequency for AFM — reveals the distance to the surface. Depending on the mode of operation, the probe can be used to deposit, remove, or modify material on the surface[161]; in the case of oxidation control by AFM, this even allows the reversible creation of nanoelectronic devices without topographic changes.[19] The biggest concern with scanning probe lithography is its extremely slow throughput. Massive parallelization could help, and a working proof-of-concept was demonstrated.[163] However, the tremendous surface area required by either the semiconductor or data storage industries puts this technique out of contention for the foreseeable future.

E-beam lithography

Electron-beam lithography has many similarities to photolithography: the sample is coated in a layer of *resist*, which is exposed to a pattern of radiation at a given wavelength and subsequently developed. However, the radiation is comprised of electrons rather than light. Another key difference is that the radiation is not spread over the sample at once through a mask, but comes in the form of a very finely focused beam which is rastered over the surface. The half-pitch which can be achieved is limited in principle by the spot size of this beam. In

practice, the achievable half-pitch is somewhat larger, due to a wide proximity effect, explained below.

IBM began exploring e-beam lithography in the late 1970's[68], and even used it in production briefly in the early 1980's.[133] However, throughput concerns prompted a return to optical lithography for industrial production. One attempt to increase throughput was based on projection of a wide electron beam through a mask, analogous to photolithography. This technology, known as SCALPEL[9], had been abandoned for industrial consideration by 2008.[169] The current high throughput e-beam candidate involves many parallel beams, and is known as "MEB ML2".[97, 76] Challenges include daunting requirements for data rates[169] and space charge limitations for parallel beams.[97]

In the meantime, e-beam has found great success as a research tool, enabling the fabrication of custom devices with a resolution far outstripping alternative techniques. Standard poly(methyl methacrylate) (PMMA) resist can yield 20 nm features[96], and new hydrogen silsesquioxane (HSQ) resists have demonstrated linewidths as low as 10 nm.[55] The resolution is limited by the high production of backscattered and secondary electrons. The secondary electrons limit the size of isolated features, because their mean free path in resist is several nm.[96] Backscattered electrons have energies comparable to the original beam, but cover a much wider area, resulting in a weakly-exposed region several µm wide.[21, 124] This is the cause of the proximity effect mentioned earlier.

He^+-ion beam lithography

He^+-ion lithography is analogous to electron lithography, in that a fine particle beam is scanned over a resist-covered surface. One difference is that $^4He^+$-ions are 7300 × heavier than electrons, resulting in much less lateral scattering and a correspondingly smaller point spread function.[170, 154] However, they are not *so* heavy as to cause significant lattice damage, unlike the Ga^+ ions used in FIB. Also, the beam is very narrow: only 0.75 nm in diameter.

Although FIB technology has been around for some time, He^+-ion microscopy is a recent development. It was first announced by Zeiss in 2007, with the possibility of great resolution improvement compared to SEM.[29] The possibility of lithography was recognized and tested soon after.[170, 154] As expected, the proximity effect of several micrometers in SEM was dramatically smaller with the He^+ microscope.[154] 10 nm half-pitch features were easily produced, and there is some indication that the resist is preventing even higher resolution from being attained.[170]

One concern with He^+-ion lithography is deep ion implantation, caused by the same low stopping power[29] that gives the instrument its high resolution. However, Ga^+ implantation is expected to be much more damaging.[154] The main obstacle at present is simply the immaturity of the technology. With the first commercial microscope, the Carl Zeiss Orion Plus, released as recently

as August 2008[167], researchers and industry have not had sufficient time to develop the technology. Since electron-beam benchmarks were matched immediately, there is much promise that optimization of resists and development will be able to exceed the resolution of e-beam based techniques.

1.2.3 Self-assembled techniques

Block Copolymers

Block copolymers used for self-assembled lithography consist of two polymer chains (or "blocks") of different chemical identities, joined by a covalent bond.[129] Each block prefers to associate with other blocks of the same type. They are applied to the surface in a disordered and homogeneous state, but their *equilibrium* state consists of regular patterns of spheres, cylinders, or lamellae[64], whose size and spacing are determined by the length and chemical identity of each of the blocks. To reach this equilibrium state, the surface is annealed in a solvent environment. Either of the components can be selectively removed by a suitable plasma, and the remaining component is used as a mask, either for etching or liftoff.

Their potential for lithography was proposed in 1997,[129] where both positive and negative transfer were demonstrated in fabricating lines, as well as dots with a half-pitch of 20 nm (0.47 Tbpsi). The chief problem was the short-range ordering, but this was quickly overcome with a variety of techniques. Topographically patterned surfaces were shown to increase the grain size and control the orientation in 2001[148], and by 2003 the effects of topographic roughness on the pattern had been quantified in great detail[25]. Also in 2003, a novel *chemical* patterning technique was pioneered[85], which gave an additional route for top-down lithographic control of self-assembly. In 2009, a fully parallel technique was demonstrated which yielded arbitrarily large single-supergrain hexagonal masks, with a pitch below 7 nm[130], though pattern transfer was not attempted. The most recent pattern transfer results involve 8 nm features with a 17 nm pitch.[78]

Block copolymers are currently the favored self-assembling system within industry, because they exhibit predictable and controllable long-range ordering with the ability to heal defects. One limiting factor has been the long time needed to anneal each wafer, though recent work has reduced this to the order of 20 min.[79] Another drawback is the relatively poor etch contrast between the remaining polymer component, which can be less dense than a bulk inorganic film, and the underlying substrate. This can be remedied through the use of polymers containing inorganic materials (such as Fe)[34], but these metals are typically considered contaminants in conventional industrial semiconductor manufacture.[64]

Nanoparticles

An alternative self-assembling system is nanoparticles. A colloid of homogeneous particles is dispersed in a solvent, and applied to a surface. As the solvent evaporates, nanoparticles make their way to the solvent interface and self-assemble there, forming well-ordered two-dimensional sheets. For lithographic applications, they are then used as physical masks, either for etching or for deposition and liftoff. Like block copolymers, nanoparticle techniques are massively parallel, but they do not require an annealing step.

Much of the earliest work was done with relatively large, electrostatically stabilized (ESS) particles, typically latex spheres. Iler was able to demonstrate self-assembly of colloids in the 1960's.[71] In the early 1980's, Fischer, Zingscheim, and Decker realized that near-field optics could beat the wavelength limit for photolithography, and that nanoparticles could be used as a near-field mask.[44, 35] The culmination of ESS nanoparticle lithography was the work of Haynes and Van Duyne on nanosphere lithography (NSL).[59] The diameter of the nanoparticles precisely controls the pitch, and the gaps can be precisely controlled by isometric chemically selective plasma etching. By varying the angle of deposition, the material, the number of nanoparticle layers (1 or 2), and the subsequent annealing time, a surprising variety of morphologies and arrangements can be reliably produced. An alternative approach is not to remove the particles, but to leave them as part of the final topography. For instance, Manfred Albrecht's group has deposited magnetic multilayers on latex spheres, and shown that the film on each sphere can be magnetized independently of its neighbors, as would be required for media.[2]

Starting in the early 1990's, work began on self-assembly of a new class of nanoparticles, with potential to access much smaller size regimes. Instead of being stabilized by charge, these nanoparticles are coated by a layer of adsorbed surfactant molecules, which act like springs in keeping the particles from clumping together. Their potential to self-assemble was noted in 1992 with the production of multilayer arrays[131], but the real breakthrough came in 1994 with a paper by Dabbousi, Murray, Rubner, and Bawendi, which showed electron micrographs of close-packed single layers of nanoparticles as small as 2.5 nm in diameter.[33] By 1997, self-assembly of 1.0 nm nanoparticles had been demonstrated.[145] Nucleation-and-growth models for self-assembly enabled creation of macroscopic single grains, by focusing on controlling the nucleation.[144] A wide variety of morphologies, from rings[104] to stripes and islands[147], have been observed under various conditions.

Relatively little work has been done in transferring patterns from monolayers of surfactant-coated nanoparticles. Masuda et al. made a contiguous inorganic template by overcoating the nanoparticles with a metal layer.[108] They were successful in removing the original nanoparticles, but the process had many complex steps, some of which required delicate manual manipulation. Pattern-

ing of an *underlying* substrate through surfactant-coated nanoparticles, with subsequent removal of the nanoparticle mask, had not been achieved before the present project.

Beyond these issues, nanoparticle methods still face important challenges for serious industrial lithographic use. A high degree of size uniformity is necessary (but not sufficient) for long-range order in self-assembly. This implies the need for nanoparticle synthesis techniques yielding highly repeatable tight size distributions. Bulk quantity is another key issue, although some progress has been made here.[126] Finally, it is difficult to control the crystallographic orientation with respect to macroscopic features.

Anodic Aluminum Oxide

Another promising self-assembling technology is anodic aluminum oxide (AAO). The native oxide layer on aluminum is only a few nm thick, but it can be grown electrochemically. The structure of this oxide layer has been known since 1953 [83]: it consists of deep columnar pores which self-arrange into hexagonal arrays. In 1995, Masuda and Fukuda showed how this self-assembled structure could be used as a template for lithography by evaporating gold and platinum into the pores.[106] Masuda's 2006 work, mentioned earlier, showed how nanoparticles could be used as templates to improve the ordering, and achieve a pitch as small as 13 nm.[108] AAO is now a standard self-assembling technique, which shows much promise for lithography.

In anodization, the aluminum surface is placed in an electrochemical solution, such as oxalic, sulfuric, chromic, and phosphoric acid.[32] A voltage is applied between two electrodes, one of which (the anode) is the aluminum itself. After the initial hole pattern is nucleated, material at the bottom of the pores is preferentially removed, even up to very high aspect ratios: for example, Piao et al. prepared pores 40 000 µm deep and just 70 nm wide, for an aspect ratio of 571.[134] The diameter and spacing of the pores is proportional to the voltage over a wide range of voltages, although the quality of ordering degrades outside a certain range.[32]

AAO templating is among the most promising self-assembled patterning techniques today. Its high aspect ratios and simple equipment make it an attractive choice. However, it is ill suited to be used as a step in manufacturing of semiconductors or hard disk drives, since it uses a messy wet chemical setup and takes hours instead of seconds. Moreover, it is very restricted as to choice of materials. Nevertheless, these advantages could disappear if AAO were used to make templates of some sort, which themselves could be used in a fast massively parallel process.

1.2.4 Nanoimprint lithography

Nanoimprint lithography involves patterning a soft surface layer by pressing a finely patterned stamp. It thus extends to the nanoscale a simple concept which was already in use for larger structures, such as bits on compact disks.[26, 169] While the stamp is in contact with the soft polymer, the latter is cured and hardened before the stamp is removed. Further patterning can proceed as usual, with etching or liftoff. The process is both massively parallel (acting on the entire wafer at once) and highly reusable, both of which are key advantages for industrial applications.

Today, nanoimprint lithography refers to a wide variety of techniques, because it serves a wide variety of consumers with different needs.[146] All of these are widely acknowledged to trace back to the 1995 paper of Chou et al.[26], where an impressive 25 nm resolution was immediately demonstrated. Just two years later, the same group took nanoimprint to the sub-10 nm regime, showing transfer of features as small as 6 nm.[27] These early approaches used heat to cure the polymer, which required a cooldown that limited the throughput. In 1998, a new approach was pioneered based on ultraviolet (UV) curing, known as step-and-flash imprint lithography (SFIL).[30] The stamps were made of quartz, which is transparent to UV, and obviated both the throughput-limiting cooldown time and the problems of thermal expansion mismatch.[56] SFIL is today the preferred approach in industrial circles.[169]

Many separate problems have plagued nanoimprint, but for the most part these have found solutions. For example, differences in areal patterning density caused differences in quality between dense and sparse regions. This was solved by directly spraying the polymer on the stamp, and matching the amount to the local pattern density. Mask release is also challenging, because it is difficult to obtain simultaneously high substrate-polymer adhesion and low stamp-polymer adhesion. The use of release layers on the stamp can ameliorate this problem.[146] Perhaps the biggest challenge is that solved problems can reappear with relatively small parameter changes. This shows the need for high-quality simulation tools that can efficiently and accurately explore parameter space.[146]

Finally, nanoimprint lithography suffers from one important "meta-problem". The manufacture of the stamp itself requires ultra-high resolution lithography. In addition to its intrinsic limits, whatever they may be, nanoimprint is thus also limited by the best resolution achievable by other means.

1.2.5 Combined approaches

Considered in isolation, each of the three kinds of NGL fails in a different way: maskless techniques have too low throughput, self-assembled technologies lack control over orientation and boundaries, and nanoimprint *requires* high-resolution lithography to make the master in the first place. A more fruitful

approach is to *combine* these systems. A serial maskless technique could be used to make a master stamp for nanoimprint, which is a parallel process with high throughput. As shown by several groups[142, 13], self-assembled systems can be guided by serially-written sparse patterns, rapidly increasing the feature density and dramatically reducing the processing time per wafer. The greatest gains could be realized by a combination of all three: a self-assembled mask, guided by a masklessly-written template, is used for the rapid fabrication of master stamps for nanoimprint lithography.

This last scenario has recently evolved into the consensus position within the data storage industry.[61] E-beam is the default maskless technology, due to its widespread availability and longstanding history in fabricating BPMR demonstrations.[57, 86, 67, 5, 102, 111, 151, 152, 87, 138] The superior resolving potential of He^+-ion lithography makes it a likely candidate to take over, since apart from its relative immaturity and the present rarity of instruments, it suffers no obvious disadvantage compared to e-beam. Block copolymers are widely acknowledged as the most production-ready self-assembling technique, because of the critical milestones that have already been achieved: registry with pre-patterned guides[13], density multiplication with precisely controllable orientation[142], and healing of defects through annealing.[25] It is overwhelmingly likely that the nanoimprint masks for the first commercial bit-patterned media will be made from e-beam patterned templates, with the actual bit pattern filled in by self-assembled block copolymers.[61]

1.3 Our contribution

This work aims to help make nanoparticles a viable alternative self-assembling technology for the above process. The lesson of optical lithography is that once industry adopts a technology, they will strongly prefer to extend it rather than adopt a new technology. Keeping the long-term future in mind, this observation favors approaches which promise the greatest extensibility.

Here, nanoparticles have two extensibility advantages over block copolymers. First, their fundamental building blocks are smaller: atoms, instead of polymers. Relative size fluctuations become large when the number of units is small, so an ensemble of nanoparticles of a given size is less limited by size fluctuations than a correspondingly-sized block copolymer assembly. Second, the smallest nanoparticles *already demonstrated* to self-assemble are smaller than their block copolymer counterparts: the latter have recently shown 3 nm features at a pitch of 6.9 nm[130], but self-assembly of nanoparticles as small as 1 nm was demonstrated more than a decade ago[145].

Many important challenges remain for nanoparticle-based lithography: for example, reliable synthesis of highly monodisperse nanoparticles, templated guidance of defect-free self-assembly, and automated inspection of wide regions. This project focuses on the *transfer* of the pattern they form into an underlying

substrate. A key result is that the surfactant coating on the particles prevents this transfer, which explains the lack of previous success in the literature. After showing how to circumvent this obstacle, several transfer techniques are investigated. One of these — the "antidots" process — is found to be clearly superior: it is materials-general, robust, faithful, and dovetails remarkably well with the desired characteristics of nanoimprint masks. Thus, this work shows a viable route to *transferring* the pattern from nanoparticles arrays. With this key roadblock gone, nanoparticles become a much more attractive option to carry self-assembled lithography to its limits.

1.4 Thesis Outline

Chapter 2 will discuss the synthesis of our nanoparticles from chemical precursors, and their self-assembly into monolayers. In Chapter 3, a variety of experimental techniques for characterization (SEM, TEM) and fabrication (sputter deposition, evaporation, ion milling, reactive ion etching) are explained, along with basic concepts of vacuum systems, plasma physics, and sputtering. Chapter 4 identifies and explains the "cracking" problem, a major obstacle to using surfactant-coated nanoparticles for lithography. I developed two novel solutions to this problem, which are also outlined in this chapter. Chapter 5 covers the techniques I developed for transferring the pattern from nanoparticle arrays. Finally, Chapter 6 summarizes the contents of the thesis, gives an outlook of the current status, and suggests profitable directions for future research.

2 Nanoparticles

Every aspect of the nanomask formation was performed in-house. This includes the synthesis of nanoparticles from chemical precursors, their self-assembly into organized arrays, and the placement of these arrays on the substrates we desire to pattern. This chapter will explore each of these stages in detail, touching also on the relevant properties which characterize the nanoparticles and the techniques by which they are measured.

2.1 Nanoparticle synthesis

The ultimate goal for the nanomask is that the shape, size, and position of each feature be highly predictable. These are the criteria by which we can judge the quality of our nanoparticles. Other criteria, such as crystallinity and magnetic properties, are important for other applications but not directly necessary for our own.

Uniform spherical shapes are easily achieved by limiting the growth rate.[173] Uniformity of size — a condition known as *monodispersity* — is more challenging to control. It is also a necessary (but not sufficient) condition for long-range ordering, i.e. predictability of feature position. I therefore focus on methods which tend to give better monodispersity: the *tightness of the size distribution* is our primary figure of merit.

To synthesize monodisperse nanoparticles, the key principle is *nucleation and growth*: small clusters nucleate from monomers in solution, and they grow by diffusion of additional monomers. In a classic paper from 1950, LaMer and Dinegar[93] showed that particles precipitated from supersaturated solutions can be very monodisperse, as long as the growth histories of the particles are roughly identical. Usually, this is achieved by keeping the nucleation time very short compared to the growth time, and ensuring that no further nucleation happens during the growth phase.[114, 127] Additionally, conditions must be kept as homogeneous as possible throughout the solution, which is why this process is called "homogeneous nucleation". (This can be contrasted with "heterogeneous nucleation", where existing nanoparticles are included in the solution and serve as nuclei for further growth. In this way, larger monodisperse nanoparticle batches can be achieved starting from smaller ones.)

In the homogeneous case, small clusters nucleate continuously, but they are transient and quickly redissolve. The "nucleation event" refers to the condi-

tions where these clusters begin to *accumulate*. At a given temperature T and supersaturation S (which is the excess concentration over the saturation concentration), there is a critical radius,

$$r_c = \frac{2\gamma v}{k_B T \ln S}, \tag{2.1}$$

such that clusters smaller than r_c do not persist, where γ is the surface energy and v is the atomic volume.[127] r_c can be decreased by increasing either T or S. When r_c becomes smaller than the size of the transient nuclei, typically about 1 nm[127], they become stable and a burst of nucleation occurs. These nuclei quickly consume material and partially relieve the supersaturation, immediately decreasing S and increasing r_c so that no further nuclei form.[114] Hence, burst nucleation achieves the clean separation between the nucleation and growth phases that enables monodispersity.[173]

In the growth stage, monomers continue to diffuse around until encountering a nucleus, which incorporates the monomer and thereby grows. Two opposing tendencies govern the growth rate. First, the surface area is proportional to r^2 while the volume grows as r^3. Hence, the growth rate of linear dimensions is proportional to $1/r$: i.e. the same number of atoms will cause a smaller increase in diameter for a larger particle.[127] Surface energy considerations lead to an opposite trend, causing small particles to undergo more rapid partial dissolution. Monomers lost in this way offset the monomer influx from diffusion, and the growth rate becomes smaller for smaller particles. The net effect is that, at any given time, there is some particle radius $r_{\max}$ for which the growth rate is maximum. The size distribution becomes narrower for particles above this size ("focusing" regime), and broader for particles below this size ("defocusing regime").[127, 173, 114]

The preceding discussion has focused on nucleation and growth, with no account of how the particles can be stable against mutual aggregation. The key is that they are coated in a layer of *surfactant*: organic carbon-chain (alkane) molecules which act as barriers between particles. In this case, the surfactant has metal-coordinating functional groups on one end, and the other end of the alkane chain is hydrophobic and coordinates with the solvent.[173] The metal-coordinating groups have densities of roughly $2.0\,\mathrm{nm}^{-2}$ to $3.5\,\mathrm{nm}^{-2}$, where the latter corresponds to monolayer-level coverage.[88] The surfactant chains from different particles can interpenetrate, and this causes repulsive forces which keep the particles separated throughout their growth, and beyond.[42, pp. 47–49]

The metal-coordinating group is not *permanently* attached to the particle's surface. Rather, it dynamically desorbs and adsorbs, continually exchanging with surfactant molecules in solution. Monomers which impinge after desorption but before further adsorption can be incorporated into the particle, while the particle as a whole remains mostly covered at any given time. In this way, the surfactant coating can prevent particle aggregation while permitting particle

growth.

The surfactant coating has tremendous effects on the growth kinetics, size distribution, and even the particle shape. Generally, a faster growth rate is associated with surfactant that is shorter[70] or has smaller binding energy. Changing the growth rate affects the value of $r_{\max}$, and therefore also affects whether the size distribution is focused or defocused. Furthermore, the duration of the window between desorption and adsorption can have a profound affect on the shape.[173] If a particle is very loosely covered, the probability for impinging monomers to stick will be different for surfaces with different energy, and a highly faceted nanocrystal will result, having only those surfaces with lowest energy. On the other hand, particles where openings are rare will grow without regard to crystal surface energies, and will tend to be more spherical. The surfactant can even have different affinities for different surfaces. The particle will then grow primarily along the direction most loosely covered, and highly anisotropic shapes such as rods or disks can result.[135, 173]

2.1.1 Size distribution characterization

I characterized the size distributions of the nanoparticles using TEM, which is explained more fully in Section 3.1.3. Image analysis techniques are discussed more fully in Section 4.3.2. Here, I used binary threshholding with linear background gradient correction. Noisy pixels were corrected by eliminating features below a threshhold size. Sometimes, noisy pixels occurred between particles and spuriously caused neighboring particles to be detected as the same feature. In these cases, I separated the particles by logically "flipping" the pixels which bridged them. The output of this technique is a series of nanoparticle cross-sectional areas A, in units of pixels. These were converted to an equivalent diameter D by assuming a roughly circular shape, according to $D = 2\sqrt{A/\pi}$.

Experimentally, a droplet of nanoparticle-containing solution was applied to a TEM grid and allowed to evaporate. The concentration was chosen so that slightly less than a monolayer of nanoparticles was left behind. With more than a monolayer, the cross-sections of particles in different layers overlap, and individual particles could not easily be distinguished. With too much less than a monolayer, each TEM image will contain few particles, and many more images will be required to gauge the size distribution.

It was found that the particles often self-assembled on the TEM grid to form ordered arrays. When this occurred, two additional parameters could be measured: the period of repetition (or "pitch") and the edge-to-edge interparticle gap. These were both measured manually. The pitch was determined by measuring the distance of straight lines of nanoparticles, and dividing by the number of particles covered by the line. Gaps were measured by picking pairs of particles at random and measuring the shortest distance between them. In both cases, several measurements of different regions were taken, and the reported

results are the mean and standard deviation of these measurements.

2.1.2 Procedures and results

Fe_3O_4 nanoparticles

Reliable synthesis of monodisperse iron oxide nanoparticles — that is, particles with a standard deviation below 10% or so — has been possible for nearly a decade. Polydisperse and easily agglomerated nanoparticles were produced by several groups in 2001[176, 84]. Later that same year, Hyeon et al. produced highly monodisperse surfactant-coated γ-Fe_2O_3 nanoparticles, with controllable sizes ranging from 4 nm to 16 nm.[69] Allen et al. used protein cages, both as templates and to stabilize against aggregation, to produce highly monodisperse γ-Fe_2O_3 particles.[3] In 2004, the Colvin group demonstrated monodisperse synthesis using only simple iron oleate as precursor[174], while the Hyeon group demonstrated dramatically improved (gram-scale) yield while retaining monodispersity[126]; however, both of these syntheses result in significantly reduced magnetization compared to bulk materials.

We used a technique introduced by Sun et al. in 2002[159] and generalized in 2004[160].The procedure is to mix the iron precursor (iron acetylacetonate, $Fe(acac)_3$), a long-chain alcohol (1,2-hexadecanediol), and surfactant (oleic acid (OA) and oleyl amine (OY)) together in a magnetically stirred solvent (benzyl ether). The presence of excess oxygen is deleterious to the reaction, so the reaction vessel (shown in Figure 2.1) is evacuated and purged with argon several times. First, the mixture is heated to 200 °C, where the long-chain alcohol reacts with the precursor and helps stabilize the iron as it is liberated. Once all the monomers are free and the solution is supersaturated, the mixture is heated to 300 °C, where burst nucleation takes place. After a predetermined growth time (which affects both the average size and the relative size distribution), the mixture is cooled to stop further growth. The solution is then removed from the vessel and rinsed with ethanol to crash out the particles. After discarding the supernatant, the remaining particles can be redispersed in toluene, where they can be stable on a timescale of at least years.

Both the OA and the OY are critical in the synthesis of these particles. Describing their observations in terms of macroscopic effects, Sun et al. report that OA-only synthesis makes it difficult to separate and clean the particles, while OY-only synthesis gives a low yield.[160] The OY contains a $-NH_2$ headgroup, which probably coordinates with Fe(III) ions in the lattice.[160] OA is a carboxylic acid, having a $-COOH$ headgroup which ionizes in solution to $-COO^-$ [8] and binds to the nanoparticle surface.[101] Binding energy for surfactants is typically on the order of a few eV, but studies on CdSe nanoparticles suggest that OA is bound roughly half as strongly as other surfactants.[136] The ability to adjust this binding energy would be a great boon, because it would enable the separate tuning of the nucleation and growth stages.[101] This

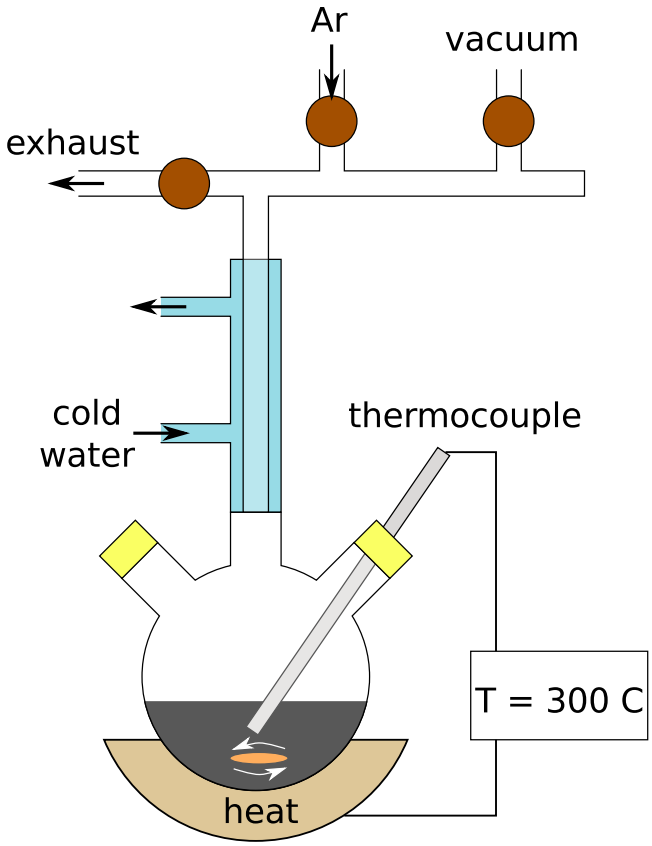

Figure 2.1: Diagram of the reaction vessel for nanoparticle synthesis. The three-neck flask is connected to the Schlenk line (top) by a water-cooled reflux tube. A thermocouple monitors the temperature of the solution, while the heating mantle controls the temperature. The solution is magnetically stirred, and is exposed to Ar gas at greater-than-atmospheric pressure, in order to minimize the oxygen in the vessel.

has been difficult to do in a controlled manner, because surfactants with different headgroups usually exhibit large and discrete differences in binding energy; moreover, the size of the surfactant also plays a critical role in determining the properties of the nanoparticles. Recent simulational work has suggested that binding energy could be tuned by adjusting chemical components *apart* from the headgroup.[101] However, in practice, the choice of surfactants still proceeds almost totally by trial and error.

It is difficult to establish the precise chemical identity of iron oxide nanoparticles, because magnetite (Fe_3O_4) and maghemite ($\gamma-Fe_2O_3$) exhibit only very subtle differences in diffraction experiments.[171] Mössbauer spectroscopy is one way to distinguish among various forms of FeO_x.[77] Sun et al. took a different approach to verify the chemical identity. They hypothesized that their particles were Fe_3O_4, and oxidized them to a putative $\gamma-Fe_2O_3$ state.[160] Further transformation in an inert environment placed them unambiguously into the $\alpha-Fe_2O_3$ state (hematite), which is clearly distinguishable from both magnetite and maghemite. This transformation is known to transform maghemite to hematite, but not magnetite to hematite; correspondingly, the as-made particles could not be transformed to hematite. These observations are consistent with the as-synthesized particles being magnetite (Fe_3O_4), but not with maghemite($\gamma-Fe_2O_3$). This does not rule out that the particle *surface* may be some mixture of magnetite and maghemite, since surfaces often exhibit differences of this nature from the bulk, but the main chemical identity has been established as Fe_3O_4.

Magnetite exhibits the *spinel* structure.[120] Roughly, this consists of a face-centered cubic (FCC) lattice of oxygen ions, with two kinds of iron ions distributed among two kinds of sites. The common ionization states of iron are Fe^{3+} and Fe^{2+}, while oxygen is O^{2-}; consequently, charge balance suggests that each unit of Fe_3O_4 will contain two Fe^{3+} ions and one Fe^{2+} ion. The available states in the spinel lattice are classified according to the symmetry of their oxygen nearest neighbors: each formula unit has one active *tetrahedral* ("A-type") sites with four nearest neighbors, and two active *octahedral* ("B-type") sites with six nearest neighbors. In the case of magnetite, it is energetically favorable for a Fe^{3+} ion, which is smaller, to occupy the tighter tetrahedral binding site. The remaining Fe^{3+} ion occupies one octahedral site, while the lone Fe^{2+} ion occupies the other. Since the Fe^{3+} ions are split between both types of sites, instead of both occupying B-sites, this structure is known as the *inverse spinel* structure.

MnO nanoparticles

MnO nanoparticles were prepared by Ryan Booth. He adapted a synthesis method first reported by Berkowitz et al. [11] The particles were measured to have 30 nm diameter by TEM, with a 34 nm period.

2.2 Magnetic properties of nanoparticles

Many atoms exhibit a net magnetic dipole moment: in other words, they create a magnetic field in the same pattern as an idealized bar magnet. The *magnitude* $\vec{B}$ of the field varies according to the strength of the dipole $\vec{\mu}$, given in SI units as

$$\vec{B} = \mu_0 \left[\frac{3\left(\vec{\mu}\cdot\vec{r}\right)\vec{r} - |\vec{r}|^2\vec{\mu}}{4\pi|\vec{r}|^5} + \frac{2}{3}\vec{\mu}\delta^3(\vec{r}) \right], \tag{2.2}$$

where μ_0 is the permeability of free space, $\vec{r}$ is the displacement vector from the dipole to the position where the field is measured, and $\delta^3(\vec{r})$ is the three-dimensional Dirac delta distribution. An appropriate unit for measuring atomic moments is the *Bohr magneton*,

$$\mu_B = \frac{e\hbar}{2m_e}, \tag{2.3}$$

where e is the magnitude of the electron charge, and m_e its mass. It is approximately equal to the dipole moment of a classical ground-state electron orbiting the hydrogen atom, and also to the magnetic moment of an isolated electron due to its intrinsic spin. The dipole moment of isolated atoms or ions can be calculated by *Hund's rules*.[120, p. 88] The moment may be different for the same atom or ion when placed in a crystal lattice, because the orbital angular momentum can be quenched by the electric field of the crystal.[120, pp. 580–581]

Atomic dipoles in materials can assume various orientations with respect to the crystal lattice. If the orientations of different atoms are uncorrelated, their macroscopic effects tend to cancel out. Such materials are called "nonmagnetic" because they respond only weakly to applied magnetic fields, making a small positive contribution (*paramagnetism*) or a small negative one (*diamagnetism*).

In magnetically active materials, these dipole moments are strongly correlated with each other, due to interactions which are not present in nonmagnetic materials. The most important of these is the *exchange interaction*, a quantum mechanical effect resulting from the indistinguishability of identical elementary particles (in this case, electrons). Consider all the arrangements of a multi-electron system which differ only by "swapping" the states of pairs of electrons. Since, *ex hypothesi*, these swaps can make no *physical* difference, the wavefunctions for these arrangements must all be identical (up to a possible change in the total sign, which also makes no physical difference). If the sign change occurs in the spin part of the wavefunction, the exchange interaction influences atomic moments to be *antiparallel* to their neighbors. Otherwise, they are influenced to be *parallel*.

In certain materials, the tendency for parallel alignment can be satisfied simultaneously for all atoms, and they all align. This phenomenon is known as *ferromagnetism*. The coordination of *micro*scopic moments over *macro*scopic distances causes tremendous collective effects, and macroscopically large mag-

netic fields can result. By contrast, a structure experiencing *anti*parallel tendencies would give no external signal at all, because the effects of each atomic dipole cancel with those of its oppositely-oriented neighbor. Note that this arrangement, known as *antiferromagnetism*, is very different from nonmagnetic materials, despite the fact that their macroscopic behavior is similar. Antiferromagnetic materials are *highly* magnetically ordered, which can be exploited when they are coupled with other magnetically active elements.[109]

Exchange can lead to other magnetic configurations in more complicated structures. The spinel structure discussed earlier is one example. Ions at A-sites orient oppositely to those on B-sites, but there are twice as many B-sites as A-sites, and two different kinds of ions as well. As it turns out, metal ions in a spinel structure attain virtually the same magnetic configuration as when they are isolated.[120, p. 126] Hund's rules show that Fe^{3+} ions have a moment of $5\,\mu_B$, and Fe^{2+} ions have a moment of $4\,\mu_B$. According to the ionic arrangement outlined earlier, the A-sites thus contribute $5\,\mu_B$ per formula unit, and the B-sites contribute ($5\,\mu_B + 4\,\mu_B = 9\,\mu_B$) in the opposite direction, for a total of $4\,\mu_B$ along the direction that the B-sites are oriented. Like antiferromagnetism, this type of ordering features antiparallel alignment, but the cancellation is not exact. It is known as *ferrimagnetism*: the production of a net macroscopic moment by the anti-alignment of sublattices of unequal magnitude.

The exchange interaction is not the only relevant magnetic interaction. For instance, iron is ferromagnetic, but most iron nails do not exhibit macroscopic fields. This is due to *magnetostatic* considerations: the energy density of a static magnetic field is proportional to $|\vec{B}|^2$; hence, a configuration with no macroscopic field is energetically favorable. There is an energy penalty for creating macroscopic fields, just as there is an energy penalty for misaligning atomic dipoles with their neighbors in a ferromagnetic material. The most favorable atomic dipole configurations will be those that minimize the *total* energy penalty.

This optimum configuration can be achieved by splitting the material into *magnetic domains*: mesoscopic regions where all the dipoles share the same orientation. Inside these domains, the dipoles are aligned with their neighbors, and there is no energy penalty from exchange. On a larger level, the domains can be oriented in certain configurations (e.g. flux closure) which do not exhibit appreciable macroscopic fields, and the magnetostatic energy penalty is also eliminated. The only significant energy penalty occurs at the boundaries between different domains, known as *domain walls*: some atomic dipoles must be misaligned with their neighbors in this transitional region.

Note that domains are feasible because the number of dipoles *inside* (which satisfy the exchange interaction) is much greater than the number on the boundaries (which do not). However, these numbers scale differently: the number of dipoles inside a domain scales with volume (i.e. r^3), while the number on the boundaries scales with area (i.e. r^2). Hence, as the characteristic size r of the

system is decreased, the ratio of border dipoles to interior dipoles becomes less favorable for domain formation. Eventually, a system can be made small enough that the energy penalty for *any* domain walls is simply too great. These systems will then be *single-domain*: all atomic dipoles will share the same orientation.

The critical size for single-domain behavior in spherical Fe_3O_4 particles has been estimated as 128 nm.[95] Sun's synthesis only claims to make nanoparticles up to 20 nm, so all of our Fe_3O_4 nanoparticles should certainly be monodomain. This means that each particle can be treated as a single *giant* dipole in its interactions with other particles. This approximation allows a quantitative measurement of the magnetic properties of the nanoparticles, to compare them with bulk Fe_3O_4.

The response of a collection of mutually noninteracting magnetic dipoles to an external magnetic field can be cast in terms of the ratio x of magnetic energy μH to thermal energy $k_B T$, where μ is the moment per dipole. It is described quantitatively by the *Langevin function*,

$$L(x) = \coth(x) - \frac{1}{x}. \tag{2.4}$$

Since H and T are known, μ can be inferred by fitting the measured $M(H)$ data to $M_s L(\mu[H/k_B T])$: essentially, the higher the moment per particle, the more horizontally compressed the data. Dividing the fit value for μ by the mass of the particle (calculated using the bulk density and a volume inferred from TEM), one obtains the magnetization per unit mass. In the case of our Fe_3O_4 nanoparticles, SQUID measurements yielded 80 emu g^{-1}, which was close to the full bulk value of 92 emu g^{-1} [113]. The difference is due to perpendicular magnetic canting of magnetic moments close to the particle surface.[92] According to SQUID measurements performed in our group, nanoparticles from other techniques such as the Colvin technique[174] exhibited magnetization roughly one-third of the bulk.

The MnO particles are very magnetically interesting, but not at room temperature. Their MnO cores are antiferromagnetic at temperatures below (T_N = 118 K), while the epitaxial Mn_3O_4 shells are ferrimagnetic below (T_C = 43 K). In other studies, our group has considered the effects of interactions in 3D assemblies of these particles, and explored the magnetic configuration of both core and shell using small-angle neutron scattering (SANS). However, the present study considers only room-temperature conditions, and is interested only in the spatial arrangement of the nanoparticles, so their intriguing magnetic behavior is not considered here.

2.3 Interparticle forces

The collective behavior of the nanoparticles is ultimately determined by their mutual interactions. Here we consider the most important forces for the particles

in question, in order to lay the groundwork for the following discussion of self-assembly.

2.3.1 van der Waals

The archetypal nanoscale force is the *van der Waals* force, arising from interaction of electric dipoles on different particles. There are three separate types of van der Waals force, depending on whether both dipoles are permanent, only one is, or neither is. The type of greatest interest in our case occurs in the absence of permanent electric dipoles. It is called the "London force" or "dispersion force", although we will simply refer to it as the "van der Waals force" because we do not consider any other van der Waals-type forces here. Although the time-averaged dipole moment is zero, quantum mechanical fluctuations cause a randomly oriented nonzero dipole moment at any given time. The transient dipoles on neighboring particles can interact and become synchronized, resulting in an attractive force.

Between two atoms or small molecules separated by a distance D, quantum mechanical calculations reveal an attractive force proportional to D^{-7}.[42, p. 31] (This implies an interaction *energy* proportional to D^{-6}, which constitutes the attractive portion of the famed Lennard-Jones "six-twelve" potential.) Nanoparticles are made of *many* atoms, and the van der Waals force between them can be approximated by the sum of pairwise interactions of their constituent atoms. When this is done, the interparticle force is seen to diminish more slowly with increasing separation:

$$F_{\mathrm{vdW}} = \frac{-A}{6D^2}\left(\frac{R_1 R_2}{R_1 + R_2}\right) \text{(sphere-sphere)}, \tag{2.5}$$

where D is now the *edge-to-edge* separation, R_1 and R_2 are the radii of the particles, and A is the *Hamaker constant*, which depends on the materials of the spheres and the dielectric properties of the intervening medium. Of course, nanoparticles are not ideal spheres, but often exhibit facets with particular crystal orientations. The effect of these facets is to increase the van der Waals force compared to ideal spheres, because flat facets enable more atoms to be close together than for curved spheres at the same separation. Typical Hamaker constants are on the order of 10^{-20} J to 10^{-19} J[63, p. 485], although they are expected to be stronger for particles adsorbed at a liquid-air interface.[15] At separations greater than about 10 nm, electromagnetic retardation effects cause the fluctuating dipoles to desynchronize, and the attraction diminishes more quickly (D^{-3} instead of D^{-2}).[42, p.35]

Two limiting forms of Equation 2.5 are of particular interest for this project. The first is when the spheres are identical, as is the case for self-assembling

nanoparticles. Then $R_1 = R_2 \equiv R$, and we have

$$F_{\mathrm{vdW}} = \frac{-AR}{12D^2}(\text{identical spheres}). \tag{2.6}$$

Alternatively, the interaction between a sphere and a plane — such as the substrates which the nanoparticles are supposed to pattern — corresponds to the limit where $R_2 \to \infty$:

$$F_{\mathrm{vdW}} = \frac{-AR}{6D^2}(\text{sphere-plane}). \tag{2.7}$$

In both cases, the van der Waals force experienced by a given particle is proportional to its radius, and inversely proportional to the square of its distance.

London-van der Waals forces are completely unfamiliar in the macroscopic domain, despite the fact that Equation 2.5 suggests they should become stronger for increasing material sizes (R_1 and R_2). There are two main reasons for this. First, the inertia scales with R^3 while the van der Waals forces scale with R, so these forces quickly become negligible compared to the inertia of macro-sized objects. Second, separations of 10 nm or less between macroscopic objects are very rare, because surface roughness prevents close approach over large areas even between nominally touching surfaces. By contrast, nanosized objects have negligible inertia and smoother surfaces. Accordingly, the van der Waals force is a juggernaut in the nanodomain, ranking as the most important force due both to its strength and to its ubiquity.

2.3.2 Magnetic dipolar forces

Particles with permanent magnetic moments also experience magnetostatic forces. These forces are especially effective when the particles are monodomain, because the atomic dipoles are all maximally aligned and produce the greatest possible external fields. Equation 2.2 gives the field due to a single magnetic dipole. The energy of interaction between a dipole $\vec{\mu}$ and an external field $\vec{B}$ is given by

$$U = -\vec{\mu} \cdot \vec{B}. \tag{2.8}$$

When $\vec{B}$ is given by Equation 2.2, the result is the interaction energy between two magnetic dipoles,

$$U = -\frac{\mu_0}{4\pi|\vec{r}_{12}|^5}\left[3(\vec{\mu}_1 \cdot \vec{r}_{12})(\vec{\mu}_2 \cdot \vec{r}_{12}) - |\vec{r}_{12}|^2\,(\vec{\mu}_1 \cdot \vec{\mu}_2)\right], \tag{2.9}$$

where $\vec{\mu}_1$ and $\vec{\mu}_2$ are the moments of dipoles 1 and 2, and $\vec{r}_{12}$ is the displacement vector from dipole 1 to dipole 2. Note that Equation 2.9 is properly symmetric under a relabeling of dipoles 1 and 2, and also under the transformation $\vec{r}_{12} \to (-\vec{r}_{12})$.

This interaction energy is highly anisotropic, as a function of the orien-

tation of the dipoles with respect to their mutual displacement vector. The corresponding force can be either attractive or repulsive. However, for fixed dipole orientations, and movements along the direction of $\vec{r}_{12}$, the energy always scales with $|\vec{r}_{12}|^{-3}$, implying that the force scales with $|\vec{r}_{12}|^{-4}$. Another point is that Equation 2.9 only applies to ideal spheres. Departure from sphericity causes higher-order terms — quadrupole, octopole, etc. — to become important. These terms diminish more quickly with distance than $|\vec{r}_{12}|^{-4}$, and hence take on greater importance for smaller separations.

In addition to their mutual force, the dipoles experience a mutual *torque*. To see the effects of this torque, we seek the configuration which yields the minimum energy for constant $|\vec{\mu}_1|$, $|\vec{\mu}_2|$, and $\vec{r}_{12}$. I recast Equation 2.9 in terms of the angles θ_i between each dipole $\vec{\mu}_i$ and the vector $\vec{r}_{12}$, as well as the azimuthal difference ϕ around that vector:

$$U = \frac{-\mu_0}{4\pi r^3} \left[2\cos\theta_1 \cos\theta_2 - \sin\theta_1 \sin\theta_2 \cos\phi\right]. \tag{2.10}$$

Now, $\sin\theta_1$ and $\sin\theta_2$ are both nonnegative, so we can take $\cos\phi = 1$ when seeking the global minimum of Equation 2.10, yielding

$$U_{\min} = \frac{-\mu_0}{4\pi r^3} \left[\cos\theta_1 \cos\theta_2 + \cos\left(\theta_1 - \theta_2\right)\right]. \tag{2.11}$$

Neither term in square brackets can exceed 1, but both can simultaneously be 1 when $\cos\theta_1 = \cos\theta_2 = \pm 1$. Hence, the magnetostatic torque tends to align the dipoles along the direction of their mutual displacement. Once aligned in this way, the force between dipoles is strongly attractive. Hence, dipoles that are free to rotate will experience *attractive* magnetostatic interaction.

Since our intended application involves nominally identical nanoparticles, it is instructive to consider the magnetic force between two aligned identical dipoles. With magnetization M, each particle of radius R will exhibit a moment of $\mu = MV = 4\pi MR^3/3$. Making this substitution allows us to write the force only in terms of intrinsic or directly measured properties of the nanoparticles. Their mutual force, from Equation 2.9, is thus

$$F_{\text{mag}} = \frac{-\partial U}{\partial r} = \frac{-3\mu_0\mu^2}{2\pi r^4} = \frac{-8\pi\mu_0 R^6 M^2}{3r^4}, \tag{2.12}$$

where r is short for $|\vec{r}_{12}|$. (Note that only dipolar terms are considered here. The field from a uniformly magnetized perfect sphere is a pure dipole field, and the particles considered here are close to spherical.)

It would seem that the magnetic force rises dramatically with particle diameter, going as R^6. This is true for highly separated nanoparticles; however, it must be remembered that r is the *center-to-center* distance, which cannot be shorter than $2R$. Recasting Equation 2.12 in terms of the *edge-to-edge* spacing

$(D = r - 2R)$, we see a different story:

$$F_{\mathrm{mag}} = \frac{-\pi\mu_0 M^2 R^2}{6\left(1 + \frac{D}{2R}\right)^4}. \tag{2.13}$$

When the particles become close compared to R, their mutual magnetic attraction is proportional to R^2. Interestingly, this force also becomes nearly independent of D for very small separations, whereas the van der Waals force continues to increase as D^{-2}.

It is instructive to do a more direct comparison of these two attractive forces: the magnetic force F_{mag} and the van der Waals force F_{vdW}. For two spherical particles of radius R and magnetization M, separated by an edge-to-edge distance D, the ratio of these forces is

$$\frac{F_{\mathrm{mag}}}{F_{\mathrm{vdW}}} = \left(\frac{2\pi\mu_0 M^2}{A}\right) \frac{RD^2}{\left(1 + \frac{D}{2R}\right)^4}. \tag{2.14}$$

The right-hand factor contains all the dependence on the size and separation of the particles. It has a maximum at $D = 2R$, where it takes the value $R^3/4$. Hence, the *maximum* ratio of the magnetic force to the van der Waals force is

$$\left(\frac{F_{\mathrm{mag}}}{F_{\mathrm{vdW}}}\right)_{\mathrm{max}} = \left(\frac{\pi\mu_0 M^2 R^3}{2A}\right); \tag{2.15}$$

if this number is small, then magnetic effects can always be neglected in a given system. By examining the functional form, we see that small particles, weak magnetism, or a strong Hamaker constant can make magnetism negligible compared to the van der Waals force. It is also possible to define a "critical diameter" $d_c(\equiv 2R_c)$, such that particles smaller than d_c will never experience magnetic forces stronger than their van der Waals forces. I obtain this by setting the force ratio equal to unity and solving for R:

$$d_c = \left(\frac{16A}{\pi\mu_0 M^2}\right)^{(1/3)} \tag{2.16}$$

For a concrete example, consider our magnetite particles, for which $M \approx (80\,\mathrm{emu\,g^{-1}})(5.15\,\mathrm{g\,cm^{-3}}) = 412\,\mathrm{kA\,m^{-1}}$. For Hamaker constants in the range 10^{-20} J to 10^{-19} J, d_c ranges from 6.2 nm to 13.4 nm, whereas our magnetite particles had diameter $13.7\,\mathrm{nm} \pm 0.9\,\mathrm{nm}$. Hence, we expect that the magnetic and van der Waals forces are at least comparable while the nanoparticles are diffusing around. Another interesting question concerns the importance of magnetic forces in determining the equilibrium interparticle spacing. We will see in the next section that our $2R = 13.7\,\mathrm{nm} \pm 0.9\,\mathrm{nm}$ nanoparticles exhibit an average spacing of 2.0 nm. According to Equation 2.14, the importance of the magnetic force depends critically on the value of the Hamaker constant: if A is 10^{-20} J, the magnetic force can be more than twice the van der Waals force,

but $A = 10^{-19}$ J yields a negligible magnetic contribution. To summarize, the force ratio analysis suggests that magnetic forces are sometimes comparable to van der Waals forces at small separations.

Two caveats need to be mentioned to this approach. First, the van der Waals force is actually smaller than described at separations larger than about 10 nm, so magnetic forces play a greater role at long distances than this analysis indicates. Second, the magnetic force calculated is actually the *maximum* magnetic force, which obtains when a pair of dipoles aligns along their displacement vector. The reality is that each dipole interacts with many others, and it is virtually impossible to be aligned with all of them simultaneously. Hence, in assemblies of strongly interacting nanoparticles, the precise effects of the magnetic force are probably not well-predicted by this analysis.

2.3.3 Capillary forces

The total free energy of a system can depend strongly on the nature of its interfaces, and the area occupied by each. An interface i — e.g. air-fluid, particle-fluid, particle-air — has an associated *interfacial energy* γ_i, and the total interfacial energy is the sum over all the individual contributions for a given configuration. Differences in interfacial free energy can manifest as forces tending to move the system towards a lower-energy state, and these are known as *capillary forces*.

The air-fluid interface in the vicinity of an adsorbed particle is deformed, because the interface makes a nonzero contact angle with the particle. Particles which deform the surface in the same way attract each other, while those deforming it in opposite directions repel. Familiar examples of this phenomenon include attraction of bubbles to each other and to the side, and cohesion of floating cereal in a bowl of milk.[162] The force is known as the *immersion force*, and takes the approximate form

$$F_{\text{imm}} = \frac{\gamma R^2}{d}, \tag{2.17}$$

where d is the separation and R the radius of the particles (assumed identical). This expression is valid for particles separated by at least several diameters, but closer than a characteristic length of the fluid which is on the order of several mm. Thus, this force can aid in bringing nanoparticles close enough together that other attractive forces can act.

2.3.4 Steric repulsion

To prevent particles from fusing together, there must be a repulsive force which rises even more quickly than the attractive forces. This force is provided by the surfactant shell which coats each particle, and it is known as "steric stabilization". When solvent is excluded from between the particles, the surfactant

concentration locally rises above its equilibrium value, and osmosis drives solvent between them. For smaller separations, the surfactant tails from neighboring particles can interpenetrate, which carries an entropic penalty.[42, p. 48] In this way, a layer of surfactant on the surface of the particles can act to stabilize them.

Repulsive forces are often modeled by a large negative power law,

$$U_{\text{rep}} \propto \frac{1}{d^m}, \tag{2.18}$$

where the precise exponent m used makes little difference in the phenomenology, as long as the function rises quickly enough.The magnitude of the steric force also will be affected by the presence or absence of solvent in the surfactant, as well as the immediate external environment.

The net effect of the attractive and repulsive forces is to create an *equilibrium interparticle spacing*, where the two forces cancel out. In this light, reconsider the previous discussion about the relative magnitudes of the van der Waals and magnetic forces. Does the magnetic force play a significant role in determining the interparticle spacing?

Earlier, we saw that in the case of the 13 nm Fe_3O_4 nanoparticles, the magnetic force is unlikely to be stronger than the van der Waals force — due to anisotropy, enhanced interface Hamaker constant, etc. — though it could be roughly equal if thermal effects were ignored. Suppose there are no thermal effects and the net attractive force is doubled when magnetism is taken into account. Its effect on interparticle spacing depends on how quickly the repulsive force rises. Taking Equation 2.18 with $m = 12$ — the standard Lennard-Jones potential — we find that doubling the attractive force decreases the spacing by just 6%, because the repulsive force rises so rapidly. In fact, in order for the distance to decrease by even 10%, one would have to take $m = 7.3$, which is much smaller than is usually taken. The conclusion is that the magnetic force probably does *not* have an appreciable effect on the interparticle spacing. This also explains why the monolayers do not show any noticeable anisotropy in their patterns. However, the magnetic state of the monolayer *could* have an effect on the directionality of the cracking patterns we see in Chapter 4 — it could, for instance, serve to break the symmetry when the initial cracks nucleate. We have not been able to investigate this possibility, but it remains intriguing.

Spacing difference between large and small particles

There is a significant difference in spacing between the 13.7 nm Fe_3O_4 nanoparticles (2 nm gaps) and the 30 nm MnO nanoparticles (4 nm gaps). Several candidate explanations present themselves. It could be a consequence of differing van der Waals forces, either due to the particle size or to their different chemical identities. Magnetic forces might also play a role, since the Fe_3O_4 particles are magnetic but the MnO particles are not. Finally, the surfactant may behave

differently, either due to chemical identity (only Fe_3O_4 particles have OY) or to geometry (i.e. different particle radii.

Equation 2.6 shows that the larger particle radius is not to blame, at least from the perspective of van der Waals forces. The interparticle force actually *increases* for larger R, which would tend to give the 30 nm particles a *smaller* gap than the 13.7 nm ones. Neither is the chemical identity likely to make much difference: both are transition metal oxides, and Fe and Mn are even neighbors on the periodic table. Even for dissimilar materials, the Hamaker constants measured in the same solvent are very close to each other for fundamental reasons.[74, pp. 177–178] Magnetic forces would tend to bring the smaller particles closer together in this case, since only the smaller ones are magnetic at room temperature, but the analysis of Section 2.3.2 shows that they are unlikely to be more important than the van der Waals forces.

Only surfactant remains as a plausible explanation of the spacing difference. Surfactant is the most likely explanation because the steric force rises very rapidly as a function of displacement, so the relatively small possible changes in van der Waals or magnetic forces would not make much difference to the equilibrium spacing. What is it about the surfactant that makes this spacing smaller for the 13.7 nm particles, and larger for the 30 nm ones? It could be the chemical identity; both have OA, but the 13.7 nm particles have OY as well. Indeed, it is likely that this affects the equilibrium spacing, but it is difficult to argue whether it should increase or decrease that spacing. A clearer contribution is that the surfactant is more open to interpenetration in the 13.7 nm particles, simply by virtue of their smaller radius. Figure 2.2 illustrates the concept.

To understand this effect, consider a naive model whereby the bond of each surfactant molecule occupies some fraction of the nanoparticle surface. Assume that the nanoparticle is completely covered by a monolayer of surfactant. At the edge of the shell, the surface area covered by the surfactant is the same as the surface area of the particle core, but the *total* surface area at that distance is larger. Hence, some fraction P will remain uncovered, where by geometry

$$P = 1 - \frac{1}{\left(1 - \frac{r}{R}\right)^2}, \tag{2.19}$$

with R the radius of the nanoparticle core and r the length of the surfactant molecules. For example, taking $2r = 4\,\text{nm}$ (to correspond to the larger gaps), one finds $P(2R = 30\,\text{nm}) = 0.22$, but $P(2R = 13.7\,\text{nm}) = 0.40$ — a significant enough difference to suggest that curvature-induced spreading may make a significant contribution to the difference in spacing.

Ultimately, it is difficult to determine the precise reasons for the difference in spacing. However, since the effect of the van der Waals force actually runs counter to the trend, and since the van der Waals force is expected to be similar for these particles, the identity and spacing of the surfactant is probably the main reason.

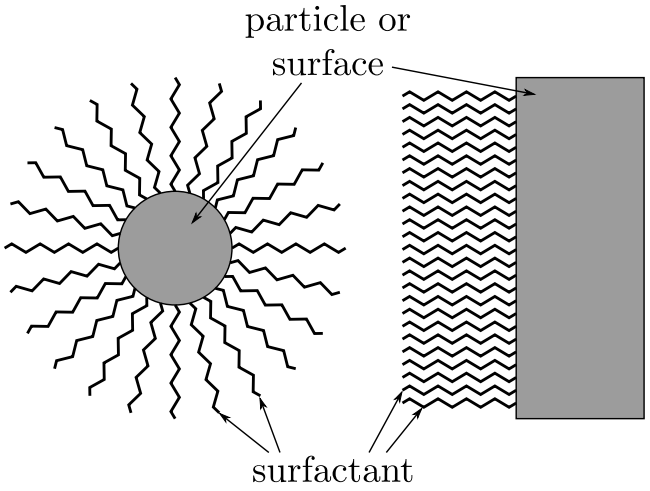

Figure 2.2: The role of the particles' radius of curvature in determining the interparticle spacing. The left shows a small particle, while the right surface is a plane to represent the limit of large particle radius. On both surfaces, the surfactant covers the same fraction of the area where it connects to the particle. However, the opposite ends of the surfactant chain are more spread out for the small, highly curved particle than for the planar surface. This may allow the surfactant on neighboring small particles to interpenetrate to a greater degree than for larger particles. Note that this diagram is not remotely to scale; the dimensions are chosen to exaggerate the visual effect. The main concept being illustrated here is that, for the same surface coverage of surfactant, particles with larger radii are less open for surfactant interpenetration.

2.4 Self-assembly

The self-assembly of nanoparticles into ordered arrays is based on the same fundamental principles as their synthesis from chemical precursors: *nucleation and growth*. Three important differences distinguish the two. First, the constituents are now nanoparticles, not atoms. This means they can exhibit variation in size and shape, which is a novel limitation on the distances over which they can order, although recent work suggests a size-sorting behavior can mitigate slight degrees of polydispersity.[115] Second, the desired (super)crystal is now two-dimensional and not three-dimensional, so monolayer formation must take place at some (two-dimensional) interface. Finally, the desired properties are fundamentally different: instead of very many crystals all the same size, we want very few supercrystals that are as big as possible.

To see what interfaces are avaialble to nurture 2D structures, we consider the features common to any technique for self-assembling monolayers. A droplet of particle-containing solvent is applied to the surface of a subphase, which may be either solid or liquid. The solvent evaporates, leaving the particles deposited on the surface. Hence, during the self-assembly process there is always a solvent-air interface, and either a solvent-liquid or solvent-solid interface depending on the type of substrate used.

With this background, we frame our discussion of self-assembly in terms of three questions. Where do the particles prefer to be? How strongly do they

prefer to be there? And, how do these preferences and their strengths determine the ordering of the structures which result?

Many answers to the first question are relevant here, but they are unified by one common theme: *symmetry-breaking*. Every time a particle prefers (or avoids) some particular location, compared to the surrounding region, it is because of some symmetry being broken at that location. The 2D solvent interfaces break the 3D translational symmetry within the solvent. Similarly, the 1D intersection of the solvent-air and solvent-substrate interfaces — i.e., the contact line of the solvent droplet — breaks the 2D translational symmetry of each interface. Encountering other particles is a kind of mutual symmetry breaking, i.e. particles can stick to each other. Finally, solid substrates may exhibit surface roughness, which breaks the 2D translational symmetry of the surface. The total picture is that particles may prefer (or avoid) the interfaces and edges of their containing droplet, inhomogeneities on solid substrates, and each other.

When a particle impinges on a more-favorable location, the total free energy of the system drops by an amount ΔF. Stronger preferences correspond to larger values of ΔF.[15] The particle may also spontaneously desorb back into the higher-symmetry phase. From statistical mechanics, its average residence time τ is related to ΔF as $\tau \propto \exp(\beta|\Delta F|)$, where $\beta = (k_B T)^{-1}$ and k_B is Boltzmann's constant. Hence, particles remain longer in phases where their attraction is stronger.

The above considerations yield insight into the conditions for monolayer formation. First, nucleation must occur only among particles already adsorbed to the interface, for without breaking the 3D translational symmetry, only 3D structures could assemble from isotropic particles. Since particles spontaneously desorb back into the bulk fluid, they must be replenished at least as quickly as they are lost. Here, the solvent-air interface has an advantage: as the solvent evaporates, this interface sweeps across the particles in its path and collects them. The result is a 2D nanoparticle "solution" at the interface.[12]

Within this solution, particles undergo diffusion. Sometimes they collide and stick together, forming a small aggregate, before desorbing back into the interface. If the 2D solution has high enough particle density, particles will impinge upon the perimeter more quickly than they desorb from it. Particles in the interior of a monolayer are very stable because of the high number of nearest neighbors, to which they are strongly attracted. The result is that the aggregate acts as a nucleation site, and a monolayer grows at the solvent-air interface. Several investigators[12, 117] have confirmed that this interface yields monolayers with excellent ordering. However, the inherently random nature of nucleation due to particle collisions leads to multiple nucleation sites, thereby limiting the ordering lengthscales to the distance between nucleation sites. This limitation suggests that longer-range ordering could be achieved by *controlling* the nucleation, instead of leaving it to chance. I briefly discuss two such techniques.

The first technique creates a single nucleation site in the center, so that growth proceeds radially outward.[144] A hydrophobic teflon ring is placed on the surface of water in a petri dish, so that the water inside the ring is highest in the center. When the solvent is added inside the ring, it is thinnest in the center, and therefore evaporates there first, breaking the 2D translational symmetry at a single point. This point serves as a unique nucleation site. From here, the contact line moves radially outward, setting up a convective flow radially inwards[40] which adds particles to the nascent monolayer. Single defect-free domains with more than 10^{12} particles have been reported for 5 nm Au nanoparticles using this method[144]. However, similar success has not been reported with other types of particles.

Another approach involves using a solid substrate which is pre-patterned commensurate with the natural periodicity of the nanoparticle array. These patterns, typically lines or posts, break the 2D translational symmetry and can serve as nucleation sites for monolayers. This technique is known as *guided self-assembly*, and it combines the key advantages of directly-written techniques and self-assembled structures, while avoiding the pitfalls of each. Directly writing the pre-patterns gives control over pattern orientation, which is crucial for applications such as data storage, while "filling in the gaps" using self-assembly means the writing takes only a fraction of the time it otherwise would. This technique has been successfully demonstrated using an alternative self-assembling system, block copolymers (BCP)[13, 142]. Work on guided self-assembly of nanoparticles is in its infancy, with unpublished results from Karl Berggren at MIT[10], Manfred Albrecht at TU Chemnitz[112], and our own research group. EUV lithography has been successfully used to direct the assembly of small nanoparticles, but without density multiplication.[81] There is difficulty in obtaining patterned wafers with features comparable to our nanoparticle diameters, and determining conditions for their reliable guided self-assembly may be a thesis project unto itself.

2.4.1 Experimental Techniques

13.7 nm $\mathbf{Fe_3O_4}$ nanoparticles

Monolayers of 13.7 nm Fe_3O_4 nanoparticles were self-assembled on a deionized water subphase. First, the particle solution was washed with acetone, and redispersed in toluene having a 10^{-4} volume concentration of OA. I filled a petri dish with deionized water, and added 40 µL of solution to the surface. The added surfactant increased the contact angle of the droplet on the water. On a timescale of several seconds, the droplet expelled a uniform thin film around its perimeter. The surfactant concentration of the droplet increases as the droplet evaporates, and this causes the contact angle to increase and the droplet to become smaller and thicker. If left alone, the droplet diffuses around the petri dish and consumes the monolayer; therefore, I absorbed it with a kimwipe to

preserve the monolayer structures.

After the remaining droplet has been removed, I transfer the films to the surface of Si wafers (which may or may not have other films on the surface, such as SiO_2. I affix the wide end of a glass pipet to the back of the wafer using vacuum grease, and the pipet then serves as a handle. Using the pipet, I bring the wafer down parallel to the film, touch the surfaces together, and remove the wafer. The monolayers touched by the wafer stick to it. This technique is known as the Langmuir-Schaefer technique.

30 nm MnO nanoparticles

Preparation of 30 nm MnO monolayers proceeds similarly, with a few differences. First, the subphase is ethylene glycol instead of water, since I found ethylene glycol gave superior results for this system. The main difference made by the ethylene glycol is that it is wet by the toluene, so the solvent droplet spreads much more evenly and evaporates more quickly. This creates many nucleation sites instead of just one, and consequently exhibits shorter-range order; however, the ordering on water was even worse. The second main difference is that I was not able to improve the ordering by adding definite amounts of surfactant, so the dispersion used was simply a dilute version of the as-made MnO particles.

Monolayers were transferred to wafers in exactly the same way as for the Fe_3O_4 nanoparticles, but the extremely low vapor pressure of ethylene glycol causes the wafer to remain wet for a long time, on the order of several hours or more. I removed most of the ethylene glycol by carefully touching the corner of a kimwipe to the droplet, without letting it touch the wafer itself. Alternatively, the wafer could be rinsed in ethanol, which is miscible with ethylene glycol but much more volatile. I did not observe any difference in the ordering between samples which were and were not rinsed with ethanol.

2.4.2 Summary

To summarize, the key to obtaining long-range ordered monolayers with few defects is to focus on the nucleation and growth. Growing large domains is easiest at the solvent-air interface. This requires that the particles experience an attractive interaction with the surface, and that evaporation is fast enough for the interface to collect more particles than it loses.

2.5 Summary and conclusion

I synthesized monodisperse Fe_3O_4 nanoparticles to use for my pattern masks. The key to a tight size distribution was nucleation and growth. Particles created from an initial homogeneous nucleation stage were used as seeds for further growth, and the resulting Fe_3O_4 nanoparticles were measured by TEM to have equivalent diameters of $13.7\,\text{nm} \pm 0.9\,\text{nm}$. These particles are stable indefinitely

in toluene due to their surfactant coating, which both governs their growth kinetics and prevents their agglomeration.

When a droplet of solution is placed on a subphase, some particles segregate to the solvent-air interface and form ordered two-dimensional monolayers, by a process of nucleation and growth that is similar to how the particles formed in the first place. The spacing between the particles is governed by their mutual forces: both attractive (magnetic, van der Waals) and repulsive (steric). The van der Waals force is the most important at the nanoscale, but for strongly magnetic particles the magnetic force may augment it. If the subphase is liquid, the monolayer can be transferred to a solid substrate by the Langmuir-Schaeffer technique, for further processing. The monolayer may then be used as a nanomask to pattern that substrate.

3 Experimental Techniques

Several commonly used experimental techniques are indispensible to this project. Microscopy is important at every stage: to characterize the size distribution of the particles after synthesis, to check correlation lengths in the self-assembled monolayer, and to inspect the results of pattern transfer. The techniques for performing that transfer are also important; these include thin film deposition techniques and etching techniques. This chapter will describe the principles which underlie these techniques, and explain the ways they are used for this project.

3.1 Electron Microscopy

The minimum feature size which a microscope can resolve is limited by the wavelength λ of the light, due to diffraction. Even with perfect (abberation-free) optical components, features cannot be distinguished when closer than a distance d_0, given by

$$d_0 = \frac{0.61\lambda}{\mathrm{NA}}, \tag{3.1}$$

where NA is the "numerical aperture" of the system, which depends on the geometry and the index of refraction.[166] Wavelengths in visible light are at least a few hundred nanometers, but our features are no larger than a few tens of nanometers. Hence, I cannot image my structures using visible light.

Equation 3.1 suggests that higher-resolution images can be obtained with shorter-wavelength radiation. In principle, short-wavelength light (such as x-rays) could work, but focusing x-rays is difficult. A more promising route is suggested by quantum mechanics. De Broglie showed that, just as wave-like light can manifest as a particle, material particles can exhibit wave-like properties such as diffraction and interference. The characteristic wavelength of a particle is given by the de Broglie relation,

$$\lambda = \frac{h}{p}, \tag{3.2}$$

where h is Planck's constant, and p is the momentum of the particle. Hence, increasing the momentum of material particles decreases their wavelength, and enables them to probe smaller structures.

The most versatile material particle used as a probe in microscopy has proven to be the electron. Because they are charged, they are easily accelerated and

manipulated with electromagnetic fields. An electron accelerated through a potential difference V increases its energy by an amount eV, with e the magnitude of the electron charge. Given the mass m_e and the energy E of an electron, its momentum p can be calculated according to Einstein's relation,

$$E^2 = (pc)^2 + (mc^2)^2 \tag{3.3}$$

where c is the speed of light. Combining Equations 3.2 and 3.3, and remembering that the total energy E includes both the rest energy mc^2 and kinetic energy eV, one finds a relation for the wavelength of electron radiation as a function of the accelerating voltage V,

$$\lambda = \frac{2\pi\hbar c}{\sqrt{eV(2mc^2 + eV)}} \approx \frac{1.24\,\text{nm}}{\sqrt{\tilde{V}\left(1022 + \tilde{V}\right)}}, \tag{3.4}$$

with $\tilde{V} \equiv V/(1\,\text{kV})$, since the kV is a convenient practical unit for electron microscopy. Equation 3.4 means that a single volt is sufficient to achieve a wavelength of about 1 nm, and the wavelength rapidly decreases for higher accelerating voltages.

I use two different types of electron microscope: transmission electron microscopy (TEM), and scanning electron microscopy (SEM). Several important elements are common to both. First, a beam of electrons is produced by an electron gun, and accelerated with a given voltage. A series of *magnetic lenses* shape and focus the beam, which reaches and interacts with the sample. The electrons which interacted with the sample then proceed to the detector, where they form an image of the sample.

There are important differences between TEM and SEM, which dictate the type of samples used for each. TEM can achieve extremely high resolution, and its diffraction patterns give important crystallographic information, but it requires very thin samples to transmit the beam: typically less than 100 nm.[53] SEM gives rich, three-dimensional surface images, can handle bulk materials, but is more limited in resolution and has difficulty with electrically insulating samples. I will explain each instrument in more detail in the sections which follow. First, I will cover some background material which is common to both.

3.1.1 Common elements of electron microscopes

Electron guns

All electron microscopes require a beam of electrons, which is produced by an *electron gun*. DeGraef[53, Section 3.7] lists several characteristics desired for a good gun. The beam should be *steady* over time. It should be *monochromatic* — producing electrons of a single wavelength — because electrons of different wavelengths are focused differently by magnetic lenses ("chromatic aberration").

The source should appear as a *point source*, and the electrons should stay close to the optical axis (O.A.) to minimize aberrations. The gun should also emit a *high current density*, so that images can be formed in a reasonable time. Finally, practical matters such as the lifetime of the filament, the cost of replacement, and the difficulty of maintenance must also be taken into account.

There are two main mechanisms for electron emission. In thermionic emission, the filament is heated to increase occupation of higher-energy states. As states having enough energy to overcome the work function of the filament become occupied, electrons are emitted into vacuum. An annular anode is placed below the gun, and is positively biased with respect to the gun. This positive bias is what sets the acceleration voltage. Thermionic filaments are inexpensive, and require only weak vacuum. However, they exhibit larger wavelength spread, wide virtual source diameters of tens of µm, low brightness and current density, and short lifetimes (less than 1000 h for LaB_6, and 100 h for W).[53]

The other type of gun is based on *field emission*. When a large electric field is applied to the filament, the energy barrier for electron emission is heavily modified: its height (the work function) decreases, and it becomes thinner. For small enough barriers, quantum mechanical tunneling can be significant, and an appreciable electron current can result at room temperature. The chief disadvantages of field emission guns are their greater cost ($1500.00 or more) and more stringent vacuum requirements (below 10^{-9} Torr). However, their scientific advantages over thermionic guns are considerable: their brightness and current densities are orders of magnitude higher, their wavelength spread is just a fraction of an eV, and their virtual source diameters can be as small as a few *nanometers*. Additionally, when properly maintained, they can last indefinitely.[53]

Magnetic lenses

Beams of electrons are focused with the carefully designed magnetic fields from *magnetic lenses*. Usually, these lenses are designed to be "round" lenses, meaning they have complete rotational symmetry around the optical axis. This symmetry means that the azimuthal component of the magnetic field vanishes everywhere. In other words, the magnetic field at every point is constrained to lie in the plane containing that point and the optical axis.

The field of a magnetic lens has two main effects on the trajectories of the electrons in the beam. First, it changes their distance from the optical axis — in other words, it *focuses* them. It can be shown[53] that, for near-axis electrons travelling nearly parallel with the beam, the field of a round magnetic lens *always* bends the electron trajectory towards the optical axis. The other effect is to change the azimuthal angle of the electrons. If these electrons carry an image, that image will be rotated around the optical axis. To make more quantitative determination of the complicated effects of lenses, a

numerical solution is required.

3.1.2 Scanning Electron Microscopy (SEM)

SEM combines the high resolution of electron optics with the ability to image *bulk* samples. The beam is focused to a very small diameter, or "spot size"[47], and functions as a localized probe. Its position on the surface is controlled by changing the current through a pair of magnetic deflector coils. The probe's interaction gives rise to several signals, and each signal gives information about the portion of the sample at the probe position. By scanning the probe in raster format, an image can be formed of a wider region of the sample.

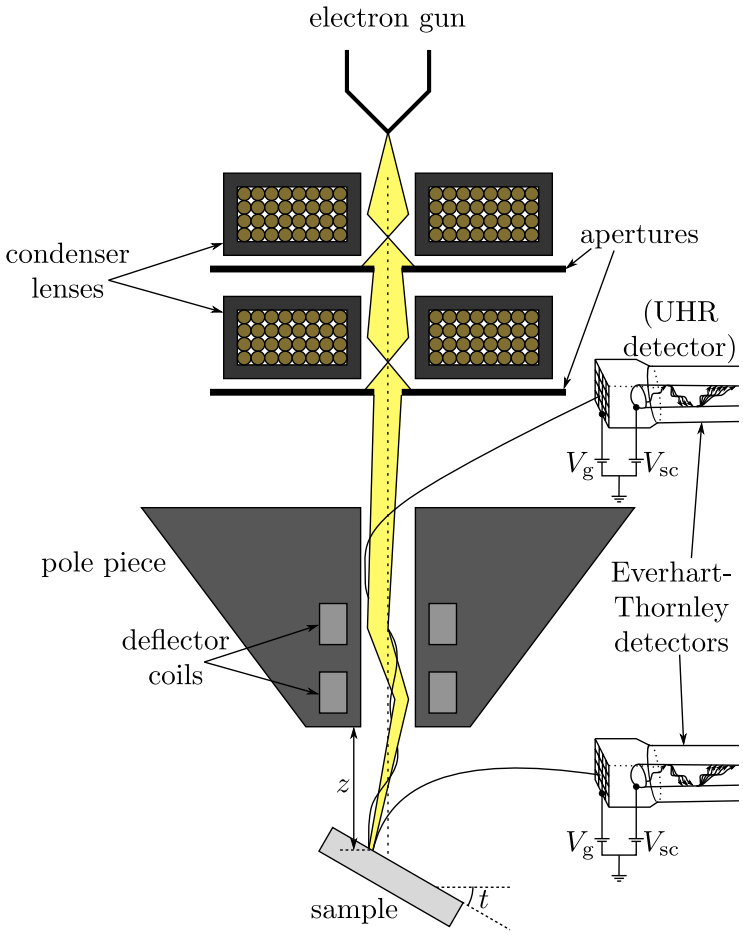

Figure 3.1: A schematic of a scanning electron microscope (SEM). At the top, an electron gun produces a beam of electrons (shown in yellow) with a given wavelength. The beam passes through a series of condenser lenses and apertures, which focus it to a smaller spot size. The pole piece sets the distance where the beam is focused, and a pair of deflector coils move the spot around the sample, whose position and orientation may be adjusted by the operator. Secondary electrons liberated by the sample can arrive at one of two Everhart-Thornley detectors: the bottom detector also detects backscattered electrons, while the top one is shielded and detects secondary electrons alone.

Thus, the resolution of the SEM is limited by the size of the beam when it

hits the sample, rather than the (much smaller) wavelength of the electrons. One major determinant of the minimum possible spot size is the type of electron gun which produces the beam. Thermionic emission guns are inexpensive, but have shorter lifetimes (on the order of tens to thousands of hours) and larger spot sizes. Field emission guns are more expensive, but produce considerably smaller spot sizes, beam currents several orders of magnitude brighter, and dramatically longer lifetimes (on the order of years).[53, 47] My structures are near the limits of SEM resolution; accordingly, I used an FEI Sirion SEM with a field emission gun for all my experiments. Its maximum resolution was $1,000,000\times$.

Regardless of the gun type, the beam it produces has a diameter of at least several µm, far too large for high-resolution raster-scanned imaging. Accordingly, magnetic condenser lenses are placed after the beam, and their purpose is to *demagnify* the beam to a smaller diameter. With condenser lenses in place, the final beam diameter can be narrowed to 5 nm for a thermionic tungsten gun, or below 1 nm for a field emission gun such as the one I used.[47, p. 7]

Note that the spot achieves this tiny diameter only at a single distance from the tip of the pole piece. At both lower and higher distances, the spot is wider, because it is far from the crossover point. This distance is user-controllable by adjusting the current through the objective lens pole piece, and the range of possible distances extends over several centimeters. The procedure of bringing the crossover point to coincide with the sample surface is called *focusing*, since the smallest probe size yields the highest resolution.

Sometimes, the azimuthal symmetry of the beam is undesirably broken, and rays in one direction may have a different focal length than rays in the other. This aberration is known as *astigmatism*, and it manifests as a linear streaking in a certain direction which is unrelated to the features being imaged. As the mean crossover point is moved towards the sample surface — in other words, as one moves towards focus — the directionality disappears, but the image remains anomalously blurry. If the motion of the crossover point is continued, the streakiness reappears, but at 90° to the original direction. The astigmatism can be corrected by adjusting *stigmators* in two orthogonal directions. The usual procedure is to focus the image so the directionality disappears, adjust each stigmator in turn to achieve the sharpest image, and repeat these three steps until the image quality is satisfactory.

So the probe can be focused to a very small size, but its effects are both deeper than the surface and wider than the spot size. As electrons penetrate the material, they interact with the nuclear and electronic charge distributions of the atoms they encounter, and they are deflected and may lose energy. The region affected by these electrons is known as the *interaction volume*, typically a few hundred µm or less[47, p. 17]. In high-Z materials, which strongly scatter electrons and limit their penetration, the interaction volume is shallow and hemispherical in shape. Low-Z materials allow for deeper electron penetration, and the interaction volume is larger and assumes a "teardrop" shape. In either

case, the interaction volume is larger for higher acceleration voltages, since faster electrons are more penetrating.

The interactions within this region produce six types of signal in a thick[1] sample.[166, p. 66] Three are based on emitted electrons (secondary, backscattered, and Auger), two are based on photons (x-ray and cathodoluminescence), and one is based on measuring the electrical current through the sample. In order to gain the information in a given signal, the corresponding detector must be present in the instrument. The signals of greatest interest for the present work are the secondary and backscattered electrons, and I consider them in more detail below.

Secondary electrons (SE) are low-energy electrons resulting from ionization. (Since the source of an electron cannot be determined, a practical definition is that electrons are called "secondary" if their energy does not exceed 50 eV. This number is based on the energy where electrons have a minimum mean free path in matter.[175, p. 21]) SE's are generated throughout the interaction volume[47, p. 17], but only those generated within a depth R of the surface (typically 15 nm or so[23]) can escape[166, p. 67]. Hence, the secondary electron yield depends on the fraction of the interaction volume which is close to the surface. This fraction will be larger for tilted samples, or near edges, which leads to the topography dependence that gives SEM images their characteristic "3D" look.

Secondary electrons are detected using an Everhart-Thornley detector.[43] At the heart of the detector is a thin aluminum film coating a small phosphor screen. The aluminum is kept at a very high positive voltage, often roughly 12 kV, causing any electrons inside the detector to accelerate towards it. These electrons pass through the thin aluminum and strike the phosphor screen, causing a flash of light, which is amplified in a photomultiplier tube and recorded as the signal. At the entrance to the detector is a metallic grid kept at a moderate positive potential of 250 V. This grid attracts the weak secondary electrons into the detector, regardless of the direction they were emitted. Hence, the total secondary electron *signal* is proportional to the total secondary electron *yield*.

Any electron emitted with more than 50 eV is called a *backscattered electron* (BSE). In practice, the majority of backscattered electrons have much higher energy, comparable or equal to the energy of the beam, because they are pictured as originating from strongly scattered primary beam electrons. Because their energy is so high, they travel in a line-of-sight manner. Furthermore, just like primary beam electrons, they can excite secondary electrons from surfaces with which they interact. This includes the sample; in fact, secondary electrons excited by backscattered electrons comprise a considerable portion of the signal.[23, 166] Since backscattered electrons illuminate a much wider area than the primary beam, this effect can be deleterious to the lateral resolution of the SEM.

[1] By "thick", I mean the sample is not thin enough to transmit electrons. When this is not the case, a *transmitted* electron signal can also result.

The backscattered electron signal exhibits strong dependence both on the accelerating voltage V and the atomic number Z. Smaller voltage makes the beam less penetrating, and the interaction volume is correspondingly shallower, leading to a smaller backscattered signal. Increasing the atomic number also decreases the beam penetration, but the backscattered signal *increases*,[166, p. 67] because the higher electron density per atom causes incoming electrons to be scattered more strongly. Backscattered electrons whose trajectory intersects the Everhart-Thornley detector will be detected, and they cannot be directly distinguished from secondary electrons. However, the operator can control the *relative* contributions of secondary and backscattered electrons by adjusting various operating parameters relating to the beam characteristics, the sample geometry, and the detector.

The most important beam parameter is the accelerating voltage. Higher values make the beam more penetrating. Since the backscattered signal increases with penetration depth, while the secondary electron signal originates from a shallow and constant depth, lowering the accelerating voltage will increase the fraction of the signal which is due to the secondary electrons. Lower voltage also favors surface details, whereas a higher voltage can probe features buried by up to several µm. Another beam parameter is the spot size, which is related to the total beam current. Too small a spot size limits the signal and causes noisy images, but too great a spot size harms the resolution because too wide a region is illuminated.[23]

The BSE:SE ratio also depends on the geometry of the sample: both its intrinsic topography, and its position and orientation in the chamber. Facets which are oriented towards the Everhart-Thornley detector yield a higher backscattered signal, because this signal is peaked towards the local surface normal. The backscattered signal is also increased when the sample is closer to the detector, simply because the latter subtends a greater solid angle. The secondary electron yield depends much less on these parameters, because these slow electrons are *attracted* to the detector regardless of their initial trajectories. Optimum working distance is a tradeoff: bringing the sample closer limits the depth of field, but also improves the maximum attainable resolution. Additionally, great care must be taken with either close working distance or high tilt angle, because the sample is in danger of crashing into the pole piece.

These approaches, based on changing the beam or the geometry, can vary the secondary and backscattered contributions in a continuous manner. By contrast, clever placement of a second detector can eliminate the direct backscattered contribution entirely. A slight extraction voltage is applied to the pole piece, causing liberated secondary electrons to travel towards it. The pole piece has a strong magnetic field inside its bore, which confines these slow secondary electrons to its axis. They spiral up through the bore of the lens, where the second Everhart-Thornley detector waits to attract them as usual. The direct background signal is blocked, because there is no line-of-sight from the sample to

the second detector. This concept is called the "through-lens detector" (TLD), and this mode of imaging is called "ultra-high resolution (UHR)" mode.

3.1.3 Transmission Electron Microscopy (TEM)

TEM is a versatile probe of very high resolution, which is restricted to very thin samples. It can directly image reciprocal space, yielding valuable crystallographic information. It can also form real-space images over a wide range of magnifications — from hundreds of micrometers down to subatomic resolution — limited only by aberrations in the lenses. This project uses the TEM in two main ways. First, I verify the intrinsic quality of the nanoparticles I synthesize: both their crystallinity, and the distribution of their cross-sectional areas. Second, I characterize key properties the monolayers they form, particularly the distribution of interparticle gaps, to a precision far greater than the SEM can achieve.

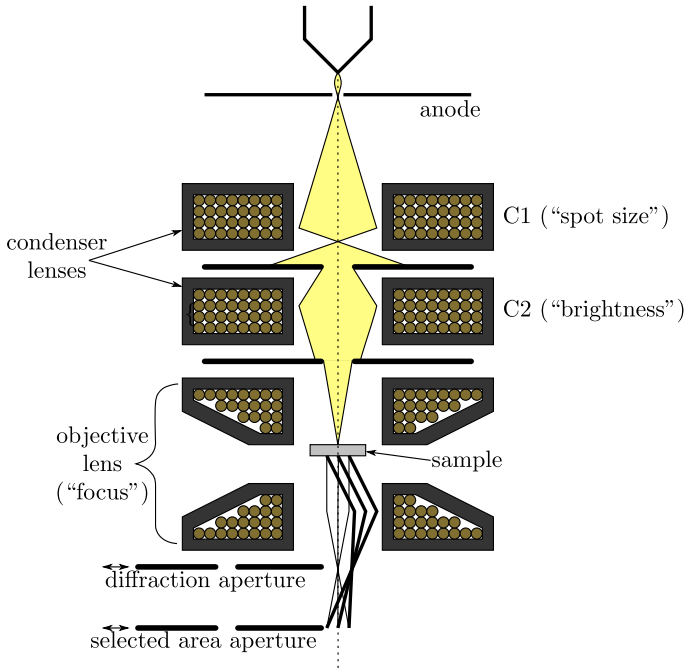

Figure 3.2: A schematic of the TEM. An image of the sample is formed at the Selected Area Diffraction (SAD) aperture. Subsequently, this image passes through the magnification stage, which is not depicted in this figure.

When the operation parameters of an actual TEM are examined, they reveal a number of facts which clarify the regimes in which the TEM operates. First and most important, TEM samples must be electron-*transparent*, meaning each electron scatters at most once. Electrons interact very strongly with matter, so samples for TEM must be extremely thin, typically less than roughly 100 nm.[53] Although the beam may appear continuous and dense, it is extremely rare for

multiple beam electrons to be in the sample simultaneously, due to the high speed of the electrons and the small thickness of the sample.[53] Hence, interactions among beam electrons can also be neglected. Finally, although electron trajectories are often depicted at wide angles for clarity, the *actual* angles involved are extremely small, on the order of milliradians. Figure 3.3 shows how this is due to the very high energy of the beam electrons, together with the fact that the scattering momentum transfer is independent of beam energy.[53] Hence, a TEM image is the aggregate of *individual* electrons undergoing *single* reactions with the sample, travelling virtually *parallel* to the optical axis the whole time.

The wafers desired to be patterned by nanoparticles are too thick to be electron transparent. Nevertheless, the nanoparticle batches themselves can be characterized by TEM, before they are used for patterning. This is done by transferring the monolayers onto *TEM grids*: pre-made structures with thin membranes stretched across a thicker grid, usually Cu or SiO. All TEM nanoparticle images used TEM grids to carry the nanoparticles.

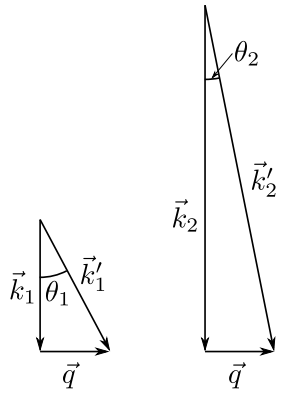

Figure 3.3: A diagram showing the connection between high energy (or high momentum) and small scattering angles. $|\vec{q}|$ is the same in both cases. It then follows from geometry that $\theta_2 < \theta_1$ because $|\vec{k}_2 > |\vec{k}_1$. In other words, a more energetic beam corresponds with smaller scattering angles.

Electrons with kinetic energy of more than 100 kV have corresponding wavelengths smaller than 0.004 nm. In principle, this opens the possibility of imaging structures close to this size, which is several orders of magnitude smaller than the Bohr radius of an atom! In practice, only the very best TEM instruments can resolve subatomic lengthscales. Hence, it is not diffraction which limits the resolution. To clarify what *does* limit resolution, I now explain the main components of the TEM, with an eye towards the role played by each in generating image contrast.

The role of the upper section of the microscope is to create a beam of electron radiation with characteristics appropriate for imaging, and direct that beam onto the sample. These components are collectively referred to as the "illumination stage"[53].

The first component of the illumination stage is the electron gun, whose role is to emit a steady stream of electrons with a controllable average wavelength λ, and minimal relative wavelength spread $\Delta\lambda/\lambda$. Nonzero $\Delta\lambda/\lambda$ causes two main problems for microscopy. The first is *chromatic aberration*: the focal length of the lenses depends on the wavelength, so only one wavelength of electrons can be perfectly focused at any given time. Secondly, the *coherence length* of the beam is inversely proportional to $\Delta\lambda/\lambda$, and many vital modes of TEM operation (e.g. electron diffraction, atomic resolution imaging) rely on long coherence lengths. The cross-sectional diameter of the beam varies with distance from the gun. The point of minimum diameter is known as the *gun crossover*.

Past this point, the beam diverges. Illumination at the sample would therefore be dim, and the operator would have little control. A lens, known as the *first condenser lens* (C1), is inserted after the gun to gather the electrons and focus them to a small, bright probe. The diameter of this probe is equal to the diameter at gun crossover, scaled by the *magnification* v/u, where u is the distance from the gun crossover to the lens, and v the distance from the lens to the image. Because a very small probe is desired for fine features, the focal length is usually chosen to be very small: $v/u \ll 1$. The control for the current through C1 is called the "spot size", because it controls the diameter of the beam spot.

This small focal length would create practical problems for the microscope, because the sample would need to be placed very close to C1, and its range of movement would be restricted. To circumvent this problem, a *second condenser lens* (C2) is added, which forms an image of the spot at a larger distance, giving more clearance for the sample. The focal length can be changed by changing the current through C2, moving this crossover image up or down. Since this causes the beam to be more or less spread out over the sample, the control for the C2 current is usually called "brightness". The highest magnifications require the beam to be as concentrated as possible, but lower magnifications need the beam to be more spread out so the area is evenly illuminated.

The next stage of the microscope is known as the "specimen stage", where the electrons reach the specimen and interact with it. It is here that *all* the desired information about the sample is transmitted to the electrons: this is the critical step. Everything else in the microscope exists only to make it easier to *extract* this information.

Due to their very small wavelength, the electrons are sensitive to the details of atomic arrangement. Accordingly, they *diffract*, i.e. they scatter in certain directions determined by this arrangement. To form an image of the sample, these *diverging* electrons must be gathered and recombined on a viewing screen. This electron gathering is accomplished by a lens: the *objective lens* (OL).

Consider an object which diffracts only three beams, as shown in Figure 3.4. Each point on the object scatters the same fraction of electrons, in the same directions, as every other point. We first trace the path of these beams origi-

nating from two *representative* points on the object, making six beams in total. A properly-designed lens will associate with each object point a corresponding *image point*, and focus all the beams so they converge on this point. Together, these image points form an *image plane*, shown as the bottom dashed line in Figure 3.4(a).

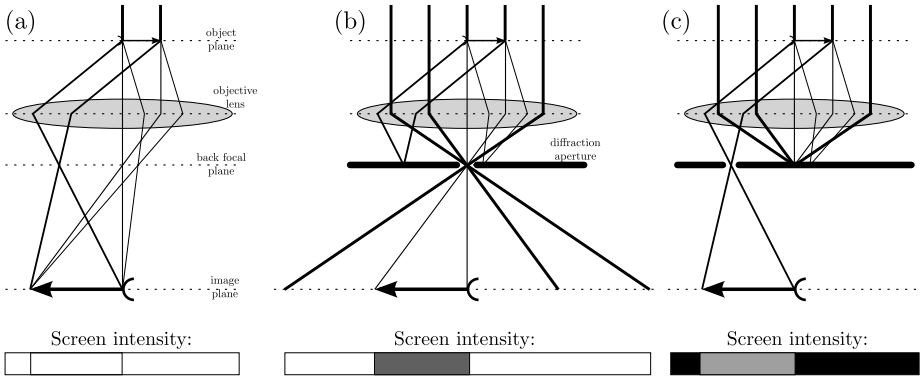

Figure 3.4: Diagram of the effects of the objective lens. For clarity, incident electrons which miss the object are not depicted in part (a). The *object* to be imaged lies in the *object plane*. Electrons passing through each point are scattered in one of three directions, with varying intensity. The function of the objective lens is to redirect the electrons scattered at each point, so they recombine at a corresponding point on the *image plane*, at the bottom of the figure. In the process, electrons scattered at the *same* angle from *different* points will *also* be recombined in a single plane, known as the *back focal plane*. If this plane is imaged instead of the image plane, the result will be a pattern of spots, one for each diffracted beam, and with intensities proportional to the intensity of the diffracted beam. (b) and (c) show how contrast can be enhanced by blocking beams in the back focal plane. Electrons which miss the sample are now shown, to show how the sample region yields different intensity than bare regions. (b) shows the *bright field* configuration, where only the transmitted beam is permitted. Since the object scatters some intensity into other beams, which are blocked by the aperture, the corresponding region in the image plane shows up *darker* than the surroundings. In (c), the aperture now *blocks* the transmitted beam, and allows some other diffracted beam past. Only points on the object can diffract into this beam, so the object now appears bright in a dark background.

The distance of the image plane below OL can be changed, by altering the current through the lens. This current is usually adjusted with the "focus" knob, since focusing consists in matching the location of the image plane with another fixed plane, where the other microscope components expect the image to be formed. Targeting a *fixed* distance for the image plane has two main advantages. First, a movable aperture (known as the "Selected Area Diffraction" aperture, or "SAD") can be inserted in this plane, selecting specific regions of the sample for further analysis. Second, by fixing the plane the current through OL is also fixed, and *changing* the current could cause undesirable thermal fluctuations.[53, p.219]

Returning to the figure, note that if all the diffracted beams recombine at each point, then apart from interference, they will be no different than points which did not diffract at all! Hence, contrast can come only from interference effects, or from departures from ideal lens behavior known as "aberrations". The former is known as *phase contrast* (Figure 3.5(a)), and is relatively weak

until very high (near-atomic) resolution. Lens aberrations, especially spherical aberration, account for most of the contrast seen in perfectly-focused images.

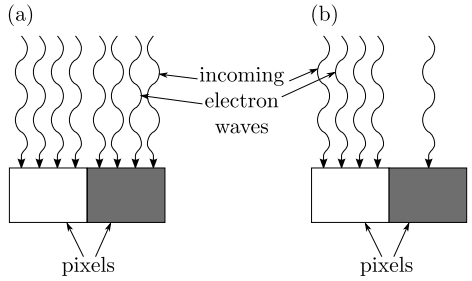

Figure 3.5: Illustration of the two primary contrast mechanisms in TEM. In each case, incoming electron waves are shown approaching two pixels. The phase of the waves is also depicted. (a) *Phase contrast:* The waves approaching the left pixel are in phase and interfere constructively, but the waves approaching the right pixel are out of phase and interfere destructively. Thus, the left pixel is brighter, and the right pixel is darker. (b) *Amplitude contrast:* The waves are in phase for both pixels, but more waves approach the left pixel than the right pixel, leaving the left pixel brighter than the right one.

Defocusing can improve contrast and make features more visible, but at the expense of clarity. Moreover, artefacts from defocused images may be misinterpreted as features. What is needed is a way to improve contrast without introducing spurious features.

To see how this might be done, notice that Figure 3.4 exhibits *two* planes where related sets of beams cross over. The second such plane, which comes before the image plane, recombines beams scattered at the *same* angle from *different* object points. This plane, known as the *back focal plane*, thus contains a focused image of all the different diffracted beams: a direct image in reciprocal space! Since the precise set of these diffracted beams depends on the material, blocking or allowing beams in this plane will have different effects on different materials. In this way one can distinguish one material from another. The resulting contrast is known as *amplitude contrast* (Figure 3.5(b)), since it results from changes in the amplitude of the beam at each image point.

Usually, all diffracted beams except one are blocked. Figure 3.4(b) shows the case where the transmitted beams are the only ones allowed. The object's transmitted beam is dimmer than beams transmitted through blank regions, because it scattered some fraction of the incoming electrons, and these scattered electrons were blocked. The image of the object is *dark*, while the image of the surroundings (or, the "field") is *bright*. Accordingly, this mode of imaging is called *bright field* imaging.

Alternatively, the transmitted beam may be blocked by the diffraction aperture, and some other diffracted beam allowed to pass. This arrangement is shown in Figure 3.4(c). Now, the only beams that reach the screen are those corresponding to the sample, which appears brighter than the surroundings. This mode is called *dark field* imaging. It can be very useful for advanced crystallographic analysis, since it shows directly which regions of the sample give

rise to which diffracted beams.

We have examined how the beam is formed and focused ("illumination stage"), and how it interacts with the sample to form either a diffraction pattern or a real-space image ("specimen stage"). The final remaining steps are to select one of these images, magnify it appropriately, and project it onto a screen to make it human-viewable. This post-specimen stage is referred to as the "magnification stage".

The first step is to show how there can be magnification at all! After all, the size of the focused image in the plane of the SAD aperture is fixed by the geometry of the preceding stages. The key is to introduce another lens, called the *projector lens* (PL), whose (variable) object plane is the (variable) image plane of OL. If the current through OL is raised, its image plane moves closer to it. To compensate, the *object plane* of PL must also be moved closer to OL (hence, further from PL). This can be accomplished by decreasing the current through PL. The magnification of each lens is the ratio of the distances of the image and object planes, and this ratio changes in the *same* direction for *both* lenses. Since the *total* magnification is simply the product of these ratios, one sees that a two-lens system can change the overall magnification while keeping the image in-focus at a fixed-position screen.

The great disadvantage of this arrangement is that magnification requires significant changes in the OL current, which as discussed above is desired to be kept constant. This problem can be circumvented by an additional lens, the *intermediate lens* (IL). This new lens plays the same role as OL in the above paragraph, but its object is the (fixed-distance) focused image from OL, and not the sample itself. For greatest flexibility, a third post-specimen lens (the *diffraction lens*, or DL) is usually added. This new lens allows the same magnification to be achieved with a *variety* of current settings, not just one, and the operator can select the combination with the smallest aberration. Additionally, its object plane can be made to coincide either with the image plane or the back focal plane of OL, allowing either the image *or the diffraction pattern* to be focused and magnified at the screen. (This is why it is called the "diffraction" lens.)

There are two types of screen onto which this focused image may be projected. The first is a phosphor screen, which emits green light with brightness proportional to the local electron current. Its images are both analogue and real-time, and cannot be directly preserved for later analysis. It is also fixed on a hinge, and can be removed to reveal the second screen: a charge-coupled device (CCD) camera. This device is fully digital: it divides the plane into a discrete number of pixels, and records a discrete number of electron counts in each. It also requires the signal to be integrated over a finite amount of time, and can be used to make permanent images for later analysis. In a typical session, the operator will make extensive use of both screens: the real-time phosphor screen is used for stage movements and coarse focusing, and the CCD for finer

adjustments and photography.

Having sketched the key ideas behind sample interaction and image formation, I make a few remaining comments about contrast mechanisms. These ideas will be important when interpreting TEM images of the samples used in this thesis.

First, the contrast strongly depends on the elements which comprise the scattering material. The Rutherford scattering amplitude of a single atom is proportional to the atomic number Z. Hence, the scattered *intensity*, which is proportional to the *square* of the amplitude, varies as Z^2. The *directions* in which the electrons are diffracted depends only on the structure: if one took a sample composed of light atoms (Z_ℓ), and replaced them with heavy atoms (Z_h), the diffraction spots would appear in precisely the same places. However, the spots would be brighter by a factor of $(Z_h/Z_\ell)^2$. This means that the bright field image of the sample with heavier atoms would be *darker*, because more electrons are diffracted *away* from the transmitted beam. An example of this is seen when gold nanoparticles ($Z = 79$) are imaged next to iron oxide nanoparticles ($Z \leq 26$): the former appear much darker than the latter.

One final phenomenon relating to the contrast is important to interpret TEM images of nanoparticles correctly. The diffraction intensity for highly crystalline materials is strongly concentrated in certain discrete directions, known as *Bragg peaks*. Most orientations will not be close enough to a Bragg peak to cause high scattering, and the bright field contrast will be relatively weak. When a Bragg peak *is* attained, it siphons a tremendous fraction of the transmitted beam, and the bright field contrast will be very high. Since the angles in TEM are extremely small (due to the high energy approximation), a very small physical rotation can make the difference between the contrast being very weak or very strong.

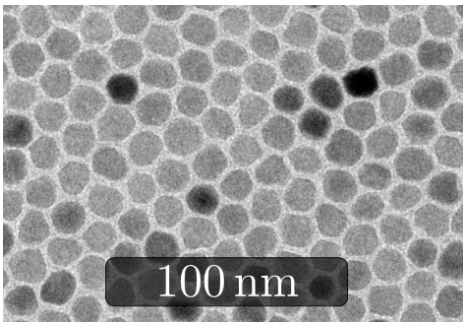

Figure 3.6: TEM bright field image of a monolayer of $13.7\,\text{nm} \pm 0.9\,\text{nm}$ Fe_3O_4 nanoparticles. Most particles appear slightly darker than the background, but several appear extremely dark. The dark particles are undergoing Bragg diffraction, due to minute differences in their crystallographic orientation with respect to the beam. They are not intrinsically different from the lighter particles.

This has a direct consequence for TEM screening of nanoparticles in monolayers. The orientations of the nanoparticles are *nearly* identical, as their elec-

tron diffraction pattern shows: a hexagonal pattern of discrete spots is observed, whereas rings would result from randomly oriented particles.Yet, the angular variations are enough that some particles attain Bragg orientation, while some do not. Figure 3.6 shows an example, where some particles appear dark because they undergo Bragg diffraction. There is no *intrinsic* distinction between the dark and light particles; indeed, a slight tilt of the sample stage will make the dark particles light, and other light particles may become dark.

To summarize, the TEM offers the highest resolution of any microscopic technique, at the cost of stringent requirements on maximum sample thickness. Nanoparticles appear crisp, with well-defined edges, enabling the distribution of cross-sectional areas to be determined with high precision. Electron diffraction gives important feedback on the crystallinity of the nanoparticles. Finally, bright field TEM enables the distribution of interparticle gaps to be accurately probed, opening the door to quantitative studies as that distribution *changes* due to "cracking" (Chapter 4).

I took all TEM images shown in this document on a JEOL 2000EX TEM, which has a maximum point-to-point resolution of 0.23 nm.[17, p. 521] I operated the instrument at an accelerating voltage of 200 kV.

3.1.4 Comparison of SEM and TEM for nanoparticle monolayers

To summarize, both SEM and TEM are needed for the fullest characterization of my nanoparticle masks and the results they lead to. TEM is necessary to obtain reliable characterization of the as-synthesized nanoparticles, and to resolve the tiny interparticle gaps after they self-assemble. SEM is needed to image the particles on the surfaces to be patterned, and to image the results of those patterns, because the wafers I use are too thick to be electron-transparent as TEM requires.

3.2 Nanofabrication Techniques

The fundamental goal of this project is to transfer the pattern of a nanoparticle array into the underlying substrate. To do so, I use standard nanofabrication techniques. This section gives an overview of those techniques which are useful for this project. Broadly, they are divided into *deposition* and *etching*: the controlled adding or removing of small amounts of material. First, I explain important background information which is relevant to both types.

3.2.1 Background material

High-vacuum systems

A great variety of instruments used in nanofabrication require pressures much lower than in the atmosphere — in other words, they require *high vacuum*. The main reason is the need for chemical purity. Nanoscale systems are characterized by high surface-to-volume ratios, and chemical interactions with the environment take place at surfaces. Hence, environments which tend to degrade sample properties will have much bigger effects on nanoscale systems than on bulk systems. In particular, many nanomaterials react strongly with oxygen, and they exhibit very different properties after oxidation. Performing the experiments under high vacuum is a way to keep harmful oxygen away. Other reasons for high vacuum include the need for certain conditions: for example, the e-beam in electron microscopy requires a high vacuum, and the properties of plasmas used for sputtering (see below) depend strongly on the pressure.

Several physical quantities are useful for characterizing a plasma. Though pressure is the dominant measure of vacuum quality, pressure *per se* has very little bearing on most processes of interest. Other quantities, which are measured indirectly, are more relevant. The first is the *molecular density*: the number of atoms n per unit volume. Using the ideal gas law,

$$pV = Nk_BT, \tag{3.5}$$

where Boltzmann's constant k_B is a fundamental constant of nature, one finds

$$n \equiv \frac{N}{V} = \frac{p}{k_BT}. \tag{3.6}$$

This shows that n is simply proportional to p as long as the temperature T remains constant, and it explains why pressure is a successful proxy for measuring n.

Another important quantity is the mean free path, λ. Molecules in a gas frequently undergo collisions, which help distribute energy and lead to thermal equilibrium. The distance traveled by each molecule between collisions is stochastic, but the average value is determined by the molecular density n and a characteristic molecular diameter d:[20, p. 10]

$$\lambda = \frac{1}{\sqrt{2}\pi d^2 n} \tag{3.7}$$

Mean free path is particularly important for processes involving directionality, because collisions tend to randomize the directions of the colliding atoms. It is also important in sustaining plasmas, which are maintained by intermolecular collisions.

High vacuum is achieved in stages, using a high-vacuum chamber and several

vacuum pumps.[20, Ch. 3] The chamber is typically made of metal, and tightly sealed except for pathways to the pumps which can be opened or closed. The sample is loaded with the chamber open to the atmosphere, and the chamber is then resealed. First, a "roughing pump" (such as a rotary vane pump) removes large quantities of air. These pumps work quickly, and readily achieve pressures down to 75 mTorr, but due to unavoidable microleaks they cannot go much below that. Once this "rough vacuum" has been achieved, the chamber is closed to the roughing pump, and opened to one or more high-vacuum pumps. These are characterized by the inability to handle large quantities of gas, but the ability to remove small quantities and achieve much lower pressures. Two kinds of high-vacuum pump are used in this work. A *turbo molecular pump* uses a sequence of spinning fan blades to push the molecules away from the chamber; it thus requires the low-pressure molecular flow regime, so that the molecules will not be impeded by collisions on their way out. A *cryo pump* maintains a very cold core such that impinging molecules simply adsorb and stick; it requires low pressures because large quantities of gas would transfer too much heat to keep the core cold.

Vacuums are characterized in one of three regimes, depending on the importance of collisions. The standard measure for this importance is the dimensionless *Knudsen number*,

$$\mathrm{Kn} = \lambda / D, \tag{3.8}$$

where D is a characteristic dimension of the chamber. When $\mathrm{Kn} \ll 1$, most collisions occur with other molecules, and this is known as *viscous flow*. One important consequence is that a gas can be actively pushed towards regions of lower pressure, because molecules experience less collisions to slow their progress on the low-pressure side than the high-pressure side. By contrast, when $\mathrm{Kn} \gg 1$, molecules rarely collide with each other, and only interactions with the chamber walls are important. This regime is called *molecular flow*; gases in the molecular regime do not undergo collective hydrodynamic motions, and thus cannot be "pushed" or "sucked" by a pressure gradient. (There is still a net flow from high pressure regions to low, but only because the rate of loss is proportional to the number of molecules, and high-pressure regions have more molecules to lose.) Intermediate values of Kn exhibit intermediate behavior, known as *transitional flow*. Whereas molecular and viscous flow are easy to model quantitatively, transitional flow is more difficult. The boundaries for transitional flow are usually taken as $0.01 < \mathrm{Kn} < 1$.[122, p.64]

Film processing plasmas

Many thin film processes rely on plasmas in vacuum chambers. A *plasma* is a gas which exhibits collective behavior in response to applied electromagnetic fields.[122, p. 147] This behavior is possible because many molecules in the plasma have been ionized. The ionization fraction is small, usually from 10^{-5}

to 10^{-1}.[165, p. 12] Plasmas can exist over a very broad range of molecular densities: from extremely rarefied gases in space where $n < 10^7\,\mathrm{cc}^{-1}$, to dense laboratory plasmas around $n = 10^{20}\,\mathrm{cc}^{-1}$.[122, p. 147]

A plasma is ignited by a sufficiently strong electric field, either DC or RF. It accelerates stray electrons, which then sometimes collide with atoms and ionize them by ejecting additional electrons. Each of these electrons will again be accelerated by the electric field, and may ionize further atoms. If the ionization rate exceeds the rate of electron recombination, more and more atoms will be ionized, and a plasma will be ignited. This implies there is a minimum pressure for plasma ignition, such that enough atoms can be found for electrons to collide with before they reach the opposite electrode. There is also a maximum pressure, because if the mean free path becomes too low, electrons cannot accelerate enough between collisions to ionize the atoms they encounter. Thus, for a given gas chemistry, chamber geometry, and electric field strength, there is a *range* of pressures over which a plasma can be maintained.[122, p. 149]

The dominant charge carriers in the plasma are the electrons, since their much lower inertia allows them to respond more strongly to an applied field.[165, p. 12] Hence, they tend to charge the sample and the chamber surfaces until an equilibrium is reached. The plasma itself acquires a positive charge in losing these electrons, and there is a corresponding plasma potential V_p.[165, p. 181] The resistance of the plasma is low, and charge imbalances can be compensated for at a rapid but finite rate. This rate is characterized by the *plasma frequency*,[122, p. 160]

$$\omega_p = \sqrt{\frac{q^2 n_e}{m_e \varepsilon_0}} = (8.98 \times 10^8)(\sqrt{n_e/n_0})\mathrm{Hz}, \tag{3.9}$$

where n_e is the number of free electrons per unit volume, $n_0 = 10^{10}\,\mathrm{cc}^{-1}$ is a convenient reference value, m_e the electron mass, and ε_0 the permittivity of free space, a fundamental constant of nature. Note that when laboratory plasmas are sustained by RF fields, they conventionally use a frequency of $\omega_{\mathrm{RF}}/(2\pi) = 13.56\,\mathrm{MHz}$, while for $n_e = n_0$ one has $\omega_p/(2\pi) = 143\,\mathrm{MHz}$ Thus, $\omega_p \gg \omega_{\mathrm{RF}}$ and the compensation time for charge imbalances is short compared to timescales on which the field varies.

Processing plasmas are usually accompanied by a luminescent glow. It results from the recombination of electrons and ions in the plasma, and spectroscopy of this light can identify the atoms in the gas. Near the chamber surfaces (including the sample), a dark area called the *sheath* is usually observed. From the positively charged plasma to the negatively charged surface, an electric field exists which accelerates ions but repels electrons. Hence, the region nearest the surfaces is electron-poor; with no electrons to recombine with, the ions here do not glow. The characteristic size of the sheath is the *Debye*

length,

$$\lambda_D = \sqrt{\frac{\varepsilon_0 k_B T}{n_i q^2}}, \tag{3.10}$$

with n_i the ionic density.[122, p. 160] λ_D is the main characteristic distance in a plasma. It also governs the distance over which the electrostatic potential of free charged objects are screened; their screened potential is proportional to $\exp(-r/\lambda_D)$.

Sputtering

A key plasma phenomenon, with a variety of applications in thin film processing, is known as *sputtering*. Peter Sigmund, who pioneered the theory of sputtering, defined it as "a class of erosion phenomena observed on a material surface as a consequence of ... particle bombardment."[156, p. 11] As Sigmund carefully notes, true sputtering does not rely on heating and evaporation; rather, it is a fundamentally atomic-scale process which still occurs in the limit of infrequent single incident particles. The surface which is sputtered is known as the *target*.

In sputtering, an energetic incident particle arrives and deposits energy in collisions with multiple atoms.[157, p. 385] Atoms receiving enough energy can also collide with other atoms, etc. The vast majority of these atoms acquire relatively little energy, and are hardly moved. Nevertheless, these atoms constitute the majority of those that *do* leave, simply because there are so many of them. Atoms leave with energies of a few eV, but incident ions bring much more energy than this. Hence, the majority of energy goes into affecting the lattice by heating, disordering, and ionizing it.[156, p. 11]

As mentioned earlier, sputtering is notable because of its wide variety of applications. The most natural use is simply its tendency to erode surfaces: sputtering is often used for etching. The eroded atoms leave the surface and stick to the next solid they encounter (in fact, this is how sputtering was discovered in the 1800's). Thus, sputtering can also be used to deposit thin films, and this is one of its primary modern industrial uses. Both of these applications can be tweaked if the sputtering gas is chemically reactive rather than neutral. In the case of etching, the resulting application is known as *reactive ion etching* (RIE), and it offers the possibility for greatly enhanced selectivity. For deposition the process is called *reactive sputtering*, and it can be used to grow oxide and nitride films from pure elemental targets.

The quantitative description of sputtering is based on the *sputter yield* S: the average number of atoms liberated per incident particle. S ranges from 10^{-5} to 10^3, but is usually between 10^{-1} and 10^1.[122, p. 174] It varies as a function of angle of incidence, and is typically a maximum for 60° to 80° from the surface normal. S also depends on the energy of the particle. Most importantly, there is a threshhold energy E_{th} below which sputtering does not occur. E_{th} must be at least as big as the surface binding energy U_0 (typically 5 eV to 40 eV[122,

p. 175]), but is typically several times larger, because most of the incident ion energy is deposited into the lattice.[4, p. 195] This has the practical consequence that, when (physical) sputtering is *not* desired, it can be avoided by decreasing the bombarding energy below E_{th}.

3.2.2 Thin Film Deposition

Thin films can be grown by deposition of gas-phase atoms onto a substrate. The two main techniques for generating these atoms are evaporation and sputter deposition, and I discuss them in greater detail below. First, there are general considerations which are common to both.

Film microstructure controls a variety of important material parameters. Microstructure, in turn, is determined by the growth conditions, the most important of which is the *atomic mobility*.[165, p. 180] Atoms arriving on the substrate do not simply stick where they land; rather, they move around and explore for a more energetically favorable configuration. Their ability to move is characterized by the ratio of the deposition temperature T to the melting temperature T_m. Low T/T_m forces atoms to stick near where they first land. The resulting topography consists of jagged peaks and voids, and the film properties are very different from the bulk. As T/T_m rises, crystallinity improves; for instance, the highest temperatures can permit epitaxial growth when lattice constants are suitably matched. The same effect can be achieved at lower temperatures by bombarding with particles during film growth.

Films often exhibit impurities. While these can come from an impure source or target, they can also result from incorporation of gases. These can be minimized by using high deposition rates and low chamber pressures.

Electron Beam Evaporation

The key concept of evaporation is straightforward: a material source is heated to a temperature where its vapor pressure is significant, and the atoms in this vapor coat surfaces having line-of-sight to the source. Originally, simple resistive heating was used to raise the temperature. However, most modern applications use electron beams to heat the sample. One advantage of e-beam approaches is that melted material contacts only solid material of the same element, eliminating the risk of contamination from the resistive heater.[122, p. 121] In this section, I will focus on electron beam evaporation.

Quantitatively, the flux Φ_e of evaporated atoms per unit area and time is given by the rate equation:[122, p. 97]

$$\Phi_e = \frac{\alpha_e N_A (P_e - P_h)}{(2\pi MRT)^{1/2}}. \tag{3.11}$$

Here, P_e is the vapor pressure, P_h is the hydrostatic pressure acting on the source, M the molar mass, R the ideal gas constant, and T the temperature.

α_e is known as the *coefficient of evaporation*, which takes values from 0 to 1. Deposition rates can be either very slow or very fast; for instance, 1000 Å s^{-1} is not difficult at 10^{-8} Torr.

A schematic of the ebeam evaporator is shown in Figure 3.7. The electrons come from a thermionic emission gun from 4 kV to 20 kV.[122, p. 123] The gun is shielded from direct line-of-sight of the source to avoid contamination, and the electron paths are bent towards the source by an applied magnetic field (not shown). Changing the magnitude and direction of the field moves the spot on the surface of the source; typically, it is moved continuously ("swept") to avoid drilling and spattering.[122, p. 124] Most materials require a melted surface region to achieve sufficient vapor pressure, but some materials sublime: notable examples include Cr, Ti, Mo, Fe, and Si.[122, p. 100] A small fraction of the evaporated atoms (typically 0.01 to 0.1) are ionized by the e-beam; thus, applying a negative bias to the substrate can induce bombardment and promote crystallinity without resorting to higher temperatures.[122, p. 125]

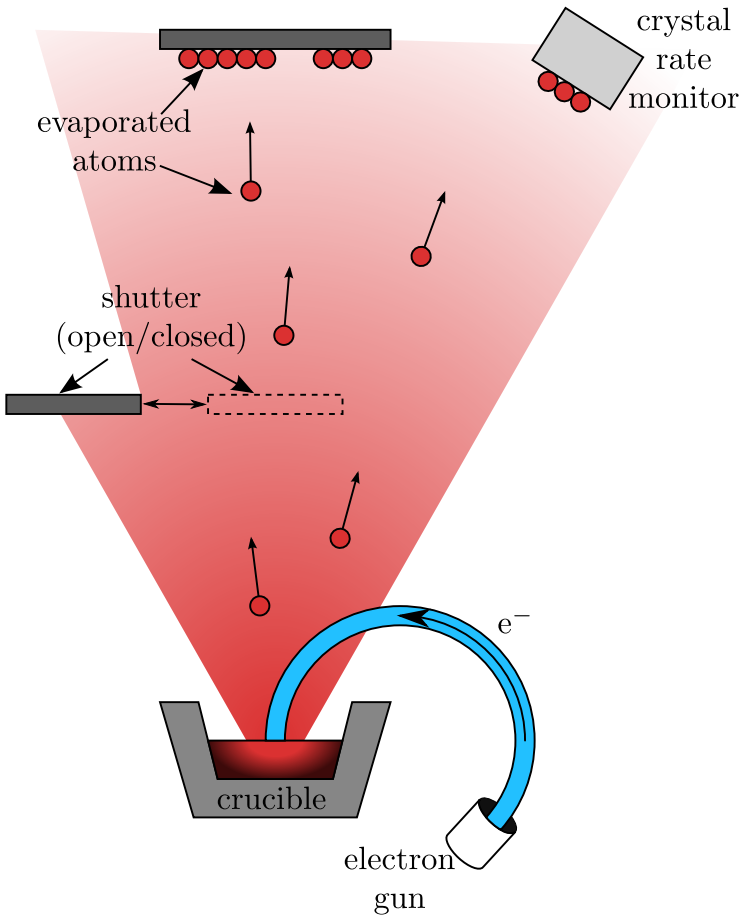

Figure 3.7: Schematic of e-beam evaporator.

Sputter Deposition

Sputter deposition is shown schematically in Figure 3.8. The target material is shown in pink at the bottom, and the surface layer of atoms is represented schematically. The electric field $\vec{E}$ (which may be DC or RF) creates and sustains the plasma, shown in blue. Sputtered atoms are shown in transit, and some of them impinge on the substrate at the top, where a growing thin film is shown. The magnetic fields at the bottom are known as a *magnetron*: they help confine bombarding ions near the target and greatly increase the yield, but at the cost that erosion is concentrated where the field is strongest.[165, p. 184]

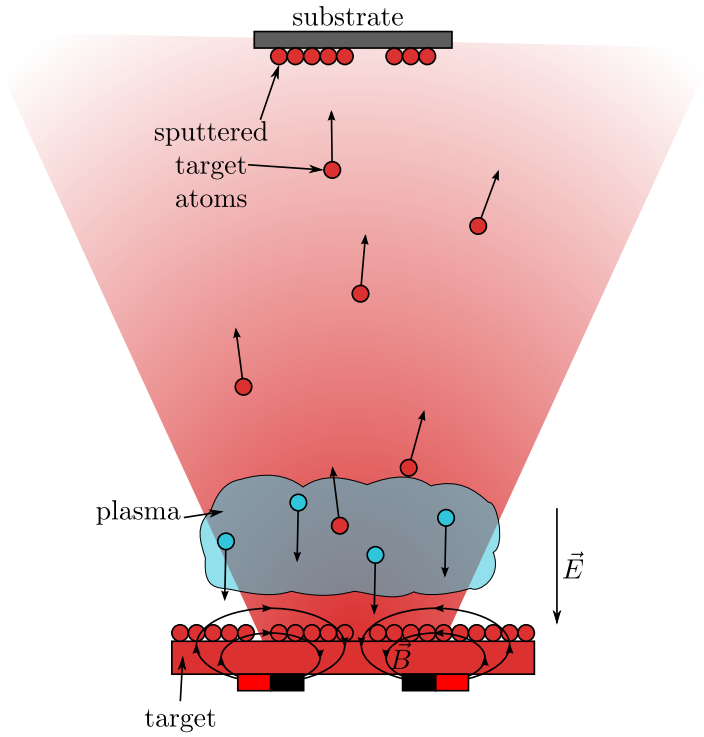

Figure 3.8: Schematic of sputter deposition.

Decreasing the chamber pressure leads to greater bombardment of the growing film, because the sputtered atoms experience less collisions to slow them down.[165, p. 182] This bombardment helps make the properties of sputtered films more bulk-like. It also tends to create a state of compressive stress in the film. By contrast, low-temperature films with less bombardment exhibit tensile stress.

Comparison of Evaporated and Sputtered films

Evaporated and sputtered films have different strengths. Evaporation takes place at a base pressure orders of magnitude lower, and the resulting long mean free path enables much better directionality. Sputtered films tend to give

more conformal coatings, including sidewall deposition which eventually closes gaps.[122, p. 140] Conversely, evaporated films can reach the bottom of much higher aspect ratios. The low pressure also causes less gases to be incorporated as impurities for evaporated films than sputtered films.

Sputtered atoms arrive with much higher incident energy: 10 eV to 100 eV compared to 0.5 eV for evaporated atoms.[165, p. 84] This gives them more mobility at the same temperatures, as well as higher bombardment, and both these properties help make sputtered films more bulklike than evaporated films. Sputtering also has an advantage in versatility: a wide variety of materials (including complex materials) are easily sputtered, but evaporation can be difficult for refractory materials. Moreover, sputtering does not mix the target material below the first few layers, and the deposited material reaches a steady state stoichiometry. By contrast, a melted evaporation source undergoes fractionation if multiple elements are present (e.g. an alloy), raising obstacles to stoichiometrically controlled evaporation of multiple elements.[122, p. 140]

3.2.3 Subtractive Patterning (Etching)

The basic process flow in etching begins by creating a mask on the substrate in the shape of the desired pattern. The mask material ought to be at least somewhat resistant to the etch process, whereas the substrate should be susceptible to it. Then, the substrate and mask together are exposed to the etching environment, which removes substrate material from areas the mask leaves uncovered. Finally, the etching is stopped and the mask removed, and the surface now bears the same topographic pattern that was in the mask.

Etching can be accomplished by a wide variety of environments, including wet chemical etching and plasma etching. Here, I focus on the latter, and I distinguish between purely physical plasma etching (*ion milling*) and chemically assisted plasma etching (*RIE*).

Ion Milling

Ion milling uses a wide directed beam of ions to erode the surface by sputtering. The substrate is placed at a definite angle to the beam, and may be rotated to ensure evenness. The rate of material removal depends on the sputter yield, which in turn depends on the incident angle and energy of the bombarding ions. The rate also depends on the material being sputtered. In their comprehensive study of etch rates, Williams et al. measured rates ranging from $4.4\,\mathrm{nm\,min^{-1}}$ for graphite to $220\,\mathrm{nm\,min^{-1}}$ for silver under the same conditions.[168] However, most materials ranged from roughly $35\,\mathrm{nm\,min^{-1}}$ to $65\,\mathrm{nm\,min^{-1}}$. These numbers demonstrate that for ion milling, it is difficult to achieve high selectivity, and deep patterning is usually achieved by increasing the height contrast of the mask.

Several phenomena can cause the observed etch profile to differ from the

mask profile.[103, p. 64] First, different crystal faces have different surface energies, and the high-energy surfaces tend to be eroded more quickly. This is known as *faceting*, and it leads milled surfaces to have sloped appearances as in Figure 3.9(c). Another problem is that ions which graze the sidewalls are reflected back, doubling the effective flux nearby and leading to *trenching* (Figure 3.9(d)). Trenching is especially evident for near-normal incidence, and can usually be avoided by choosing a larger angle around 20°. Finally, the trajectory of sputtered atoms may intersect the sidewalls, and these atoms may redeposit instead of escaping.

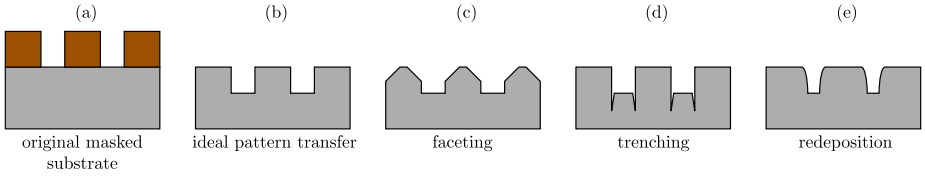

Figure 3.9: Illustration of some common problems in ion milling. (a) Original mask. (b) Ideally transferred pattern. (c) Illustration of faceting: the sloped surfaces have lower crystallographic energy. (d) Illustration of trenching: ions that graze the sidewalls and bounce back overexpose nearby areas. (e) Illustration of redeposition: sputtered atoms which encounter the sidewalls as they exit can stick instead of leaving the substrate entirely.

Reactive Ion Etching

Ion milling has poor selectivity because the ions used for sputtering are typically noble gases, which have no chemical preference for different materials. Using other kinds of ions to sputter can lead to dramatically different etch rates for different materials, and enables much better height contrast with thinner masks. This type of plasma etch is known as *reactive ion etching* (RIE), and it is a major process in modern nanofabrication. An additional advantage is that the reaction products are volatile; hence, they can encounter sidewalls on their way out without necessarily redepositing. Of course, the flip side of excellent chemical selectivity is restricted choice of materials: some materials do not form volatile reaction products, and can only be patterned by ion milling.

The most thoroughly-explored RIE chemistry involves Si-based materials, due to their ubiquity in the semiconductor industry. SiF_x products are highly volatile down to temperatures of $-196\,°C$, where Si_2F_6 condenses.[137, p. 1550] Hence, fluorine-based etch gases are commonly used, including SF_6, CF_4, and CHF_3. These gases dissociate in the plasma environment, yielding the molecular and atomic ions which perform the actual etching. (For example, SF_6 gas does not etch Si, but SF_6 plasma liberates F^+ ions which do etch Si.)[103, p. 67]

RIE pattern transfer can exhibit high directionality, despite the low mean free path of the ions. This is because they tend to strike the sidewalls less forcefully than they strike the bottom of the substrate. In the case of SiO_2, the strong bonds prevent sidewall material from being removed by the weak

bombardment, but the bombardment at the bottom is strong enough to liberate material. For elemental Si, whose atoms are more weakly bound, selectivity comes from the gradual formation of polymer: a combination of SiF_xO_y [119] and CF_x.[52, 89, 103] Here, the forceful bombardment at the bottom dislodges the polymer and allows etching, while the weaker bombardment on the sidewalls does not break through the layer. For recipes based on CF_4 and CHF_3, adjusting the proportion of these gases adjusts the C:F ratio, and thus also adjusts the proportion of F atoms which contribute to polymer formation rather than etching. In this way, adjusting the ratio of these gases can give control over the directionality of the etch profile.[103, pp. 75–76]

Comparison of dry etching techniques

Both ion milling and RIE involve sputtering, but RIE adds a chemical component. This enables vastly greater selectivity in RIE than ion milling, and the volatility of the reaction products also mitigates problems due to redeposition. Conversely, this same selectivity also restricts the choice of materials for RIE but not for ion milling. Ion milling takes place at lower pressures[103, p. 70] and is intrinsically more directional, but dynamic polymer formation makes directionality in RIE at least as good as in ion milling. Ultimately, both techniques remain useful for nanofabrication, but redeposition in ion milling makes reactive ion etching a superior choice in the extreme small-gap regime.

3.3 Summary

This chapter has reviewed several key experimental techniques. Section 3.1 discussed direct visual characterization in the form of electron microscopy. TEM exhibits superior resolution and can reliably resolve nanoparticle edges, but places stringent limits on the sample thickness which rule out non-destructive device measurements. SEM is a surface-sensitive raster-scanned technique, whose slightly poorer resolution is compensated by the greater variety of samples it can handle. Ultimately, both techniques are needed for a full characterization of the processes I use.

Section 3.2 explored the creation and patterning of thin films. Evaporated films are more restricted as to choice of materials, but exhibit very high directionality and can reach relatively high aspect-ratios. Sputtered films coat the surface more conformally, exhibit more bulk-like properties, and can faithfully reproduce stoichiometry for a wider range of materials. To pattern films, two dry-etching techniques were discussed. Ion milling is based on physical sputtering by neutral gases; it etches virtually any material, but the selectivity is poor and the etch products readily redeposit on sidewalls. RIE adds a chemical component to the etch, and achieves greater selectivity and less redeposition without sacrificing directionality, at the cost of restricted materials

choice. While all these techniques are important for nanofabrication, continuing feature miniaturization will favor ion milling over RIE (due to redeposition) and evaporation over sputtering (due to directionality).

4 Cracking and Stabilization

As described in Chapter 2, each nanoparticle in an array is coated by a layer of surfactant: the polar headgroup attaches to the particle, and the hydrophobic tailgroups extend outward. *Interstitial* surfactant creates a dilemma which prevents transfer of the hexagonal ordered pattern. On the one hand, it fills the gaps between neighboring nanoparticle cores, and neither etchants nor depositing atoms can reach the surface. On the other hand, it stabilizes the particles against the strong mutually attractive forces described in Section 2.3, and when it is removed they aggregate together randomly. This phenomenon is called *cracking*, and an example is shown in Figure 4.1.

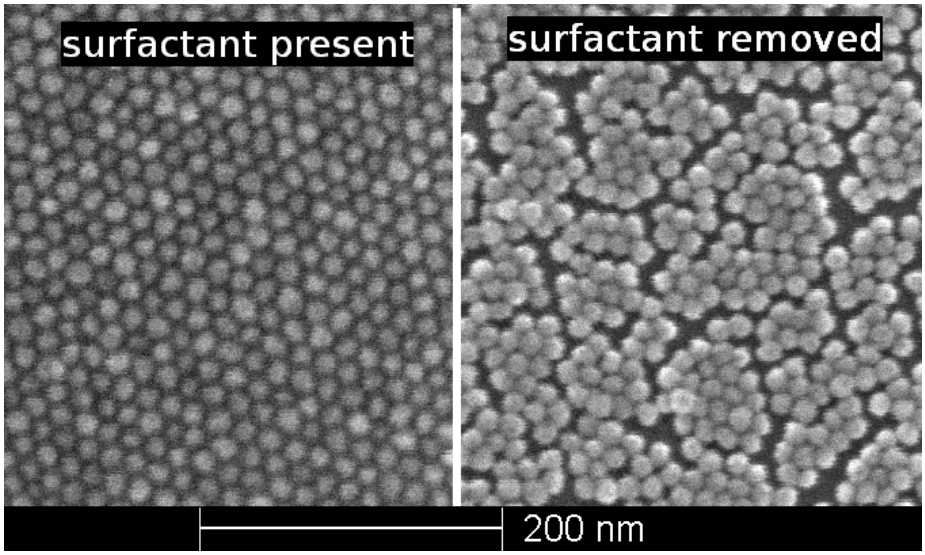

Figure 4.1: Figure illustrating the effect of removing surfactant. On the left, a pristine array which still contains the surfactant. On the right, the same monolayer after exposure to an oxygen plasma removed the surfactant.

Cracking is a barrier to any attempt to pattern *through* arrays of surfactant-coated nanoparticles, and is probably the reason no such attempts were successful before this project. This chapter will explore the phenomenon in detail, culminating in two novel solutions. First, a variety of surfactant removal techniques are presented and discussed. The phenomenon of cracking is then considered in detail, and several quantitative cracking metrics are proposed. A pioneering technique for cracking prevention is presented, and its successes and deficiencies are outlined. These defects are remedied by a second more mature

technique which shows great potential for industrial use. Finally, the state of cracking prevention in nanoparticle arrays is summarized in light of this new work.

4.1 Surfactant removal techniques

Three techniques are commonly used to remove surfactant in nanoparticle assemblies. Oleic acid on magnetite thermally decomposes above 378 °C[38], so it can be removed by heating above this temperature. Alternatively, UV-ozone treatment gently removes the surfactant at room temperature.[143] Finally, an oxygen plasma reacts with the carbon to yield gas-phase products, which are simply pumped away.

For thermolysis of oleic acid, monolayers are typically heated to 400 °C for 30 to 60 minutes.[37] Oleic acid partially desorbs beginning around 200 °C, then becomes graphitic at 400 °C after undergoing dehydrogenation.[132] The interparticle gaps decrease, disrupting the ordering[143, 38]. Black et al. show decreased gaps but excellent ordering[14, Figure 1(c)], but only one small region is shown, and the authors do not say whether the observed shrinkage yielded the expected disruption of order outside this region. In our own attempts with Fe_3O_4 monolayers, results were somewhat inconsistent, but the order was always disrupted in regions not previously imaged. I abandoned this method of surfactant removal because of its inconsistency and its tendency to change the period of ordering.

UV-ozone treatment is typically considered to be gentler. Two main UV wavelengths were used, corresponding to emission lines of atomic Hg. 185 nm light generates ozone gas (O_3) in the presence of air, which strongly reacts with the carbon in the surfactant. 254 nm places the surfactant in excited states which react more strongly with the ozone.[123, p. 544] Surfactant on top of the nanoparticles was removed much more quickly than surfactant underneath, and the relevant timescale is on the order of minutes or tens of minutes.[123] My setup involved placing the wafer 1 cm from the edge of the lamp, at the opening of a fume hood. it was difficult to be confident that conditions were homogeneous across the wafer, both because of exposure to ambient conditions and due to simple geometry. Redeposition is not expected to be a problem, because the reaction products are volatile.

Another alternative is to create an oxygen *plasma*, where the ionization of the oxygen atoms increases their reactivity. The effectiveness and rate of removal depend strongly on the plasma parameters: primarily the power, pressure, and duration. X-ray photoemission spectroscopy (XPS) of carbon in surfactant shows that a sufficiently long plasma exposure can remove all detectable traces of the surfactant.[48] One danger is that ions approaching with sufficient energy (typically ≈ 100 eV or more) can cause physical sputtering. Possible solutions to this include lowering the power, or decreasing the mean free path (by raising

the pressure) so the ions thermalize.

4.2 Cracking

Nanoparticle cores in a monolayer are always strongly attracted to each other, because of the forces described in Section 2.3. Steric repulsion from the surfactant keeps the cores separated by a minimum distance g_0. As surfactant is gradually removed, steric repulsion becomes less effective at keeping the cores apart, and the minimum spacing effectively shrinks to some smaller value $g < g_0$. The cores may still be separated by g_0, but g is the new *equilibrium* spacing. This places the monolayer in a state of *stress*, which can be relieved if individual nanoparticles can move closer to their neighbors. Cracking is simply the collective effect of these individual nanoparticle motions, and is a mechanism for relieving tensile stress.[118]

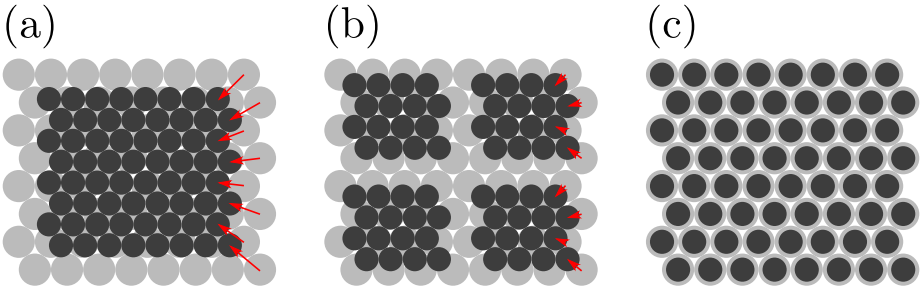

Figure 4.2: Ordering behavior for monolayers in three separate regimes, based on the ratio of the interparticle attractive force F_a to the particle-substrate binding force F_b. Originally, the particles had a larger effective diameter (shown in grey); after that diameter shrunk to the value shown in black, the particles aggregated depending on the ratio $R = F_b/F_a$. (a) R small: the particles can move wide distances over the substrate to be close to their neighbors; the red arrows on the rightmost column indicate that these particles have moved a long distance. In reality, there will still be *some* cracking, but the island size will be very large. (b) R moderate: the energy penalty for moving large distances is now considerable. The balance between interparticle attraction and substrate binding causes them to break up into small groups, where each individual particle moves a much smaller distance (see red arrows), but they still are able to be close to several nearest neighbors. Between the islands, cracks have appeared. (c) R large: now the bond to the substrate prevents the particles from assuming a pattern with their new, smaller gaps. Their effective diameter is smaller, but they do not move.

The degree of cracking depends on competition between two kinds of forces. The interparticle attractive force, F_a, depends on the gap between particles and is larger for shorter separations. The substrate binding force, F_b, depends on the lateral displacement of each particle from its original position, and is larger for longer displacements. Section 2.3 lists the contributions to F_a. The precise form of F_b is not well-known, but the main contribution probably derives from the surfactant-substrate bond. Additionally, the sphere-plane van der Waals attraction (Equation 2.6) should cause a frictional component to keep the particles in place.

Figure 4.2 shows that cracking can be classified in one of three regimes, de-

pending on which of these forces (if either) is dominant. When the interparticle attraction dominates, as in Figure 4.2(a), the morphology consists of large close-packed islands. Some particles must move a long distance to their new position, as shown by the long red arrows on the right column, and this is possible because the binding to the substrate is relatively weak. As that binding becomes stronger, large movements become disfavored, and the resulting morphology resembles Figure 4.2(b). The short red arrows indicate that the particles need only move a small distance, and the monolayer has broken up into numerous small islands, separated by cracks. Figure 4.2(c) shows the regime where the substrate binding force increases very rapidly with lateral displacement — in other words, the particles are tightly held to their original positions. Now there is effectively no cracking; the nanoparticles are *kept* in their state of stress by the strong forces binding them to the substrate.

The time evolution of cracks was clarified in 1988 by Skjeltorp and Meakin, in a combined experimental and simulational study.[158] They placed monolayers of optically visible polystyrene spheres on a glass slide and observed them as they dried, which reduced their equilibrium spacing from 3.4 µm to 2.7 µm. The initial state of the dried monolayer is very high in stress. Figure 4.3(a) shows that each of the particles experiences strong but symmetric attractive forces from its neighbors, and that these forces cancel out. However, this equilibrium is unstable: the energy drops precipitously as a function of displacement, as Figure 4.3(b) shows. Thus, once a particle is displaced towards one of its neighbors, it will tend to continue moving in that direction unless stabilized by some other force.

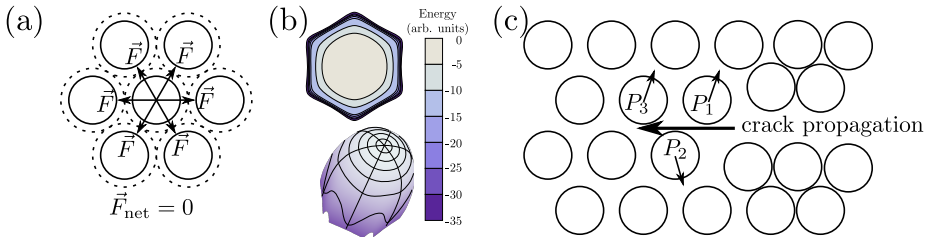

Figure 4.3: Nanoparticle forces and energies involved in cracking. (a) An illustration of the forces experienced by a nanoparticle, after its effective diameter shrinks from the dashed outline to the solid outline. Although each neighbor strongly attracts the particle, the net force cancels out, and the particle is in stable equilibrium. (b) Schematic plot of the energy of the particle as a function of displacement, assuming its neighbors do not move. The top view is a contour plot, while the bottom view shows a 3D view of the function. (c) Simplified illustration that cracks propagate along lattice directions. A crack has nucleated at the right. Since the particles on the bottom half are further from particle P_1 than particles on the top half, P_1 experiences a nonzero net force away from the crack. After P_1 moves upwards, its influence on P_2 will diminish, and P_2 experiences a downwards net force. P_2's absence then causes P_3 to move upwards. As particles continue to move away from the crack, it propagates along the lattice direction.

These initial displacements act as nucleation centers for cracks. Skjeltorp and Meakin found that these cracks propagate along lattice directions, by a process illustrated schematically in Figure 4.3(c). When a crack nucleates, the

forces on neighboring particles no longer cancel out, and the direction of the net force is determined by the crack direction. For instance, particle P_1 in Figure 4.3(c) borders a nucleated crack. Particles across the crack are now further away, and exert a correspondingly smaller force; hence, P_1 will move away from the crack, thereby extending the crack by one particle. This leads to a net force on particle P_2 which is also directed away from the crack. In this way, the crack propagates along the lattice direction.

The regime just described applies to cases with high initial stress. These cracks form rapidly, and quickly relieve most of the stress in the monolayer.[158] The remaining cracking is characterized by smaller displacements along the substrate, because the residual stress in the monolayer has become comparable to the energy penalty for lateral displacements. Isotropic cracking patterns help minimize displacements; hence, cracks in this stage are more meandering and more numerous.

Figure 4.1(b) shows an example of cracking in a monolayer of our 13.7 nm Fe_3O_4 nanoparticles. These cracks correspond *only* to the "late stage" cracking identified in Skjeltorp and Meakin: they follow the lattice directions, but often exhibit turns after distances of only a few particles. This suggests the absence of an initial high-stress stage in the evolution of our monolayers. Skjeltorp and Meakin decreased their equilibrium spacing by 20% before cracks were even nucleated; by contrast, our surfactant was removed gradually and the particles had a chance to settle. Observation of nascent cracks suggest that the symmetry breaking occurs very early, when the equilibrium spacing had decreased only very little. These initial broken symmetries largely determine the future evolution of the cracking patterns.

4.3 Measuring the degree of cracking

Broadly, there are two classes of approach to quantifying the degree of cracking: small-angle scattering experiments, and automated image analysis. It is useful to break the latter into two stages: analysis of ideal images (i.e. where the center or perimeter of each particle is precisely known), and the various routes by which *real* images are converted to these ideals.

4.3.1 Small-angle scattering

The *pair correlation function* $g(r)$ measures the spatial extent of positional order in a system. Suppose one starts from the center of a randomly selected particle, and moves a distance r in a random direction: the probability to find another particle center there within a small volume $d\vec{r}$ is proportional to $g(r)d\vec{r}$. If the particles are close-packed and monodisperse, $g(r)$ will rise from zero to a sharp peak at the diameter D of the particles, which is the closest they can approach. At large r, $g(r)$ tends to a constant, which is proportional to the

average density of the particles. Behavior at intermediate r-values depends on the degree of ordering. A perfect infinite periodic lattice exhibits multiple discernible peaks, which are progressively more damped with larger r. By contrast, a more disordered lattice may not show distinct peaks after the nearest-neighbor peak.

Small-angle scattering enables a nearly-direct experimental probe of $g(r)$, since the measured scattered intensity is the Fourier transform of $g(r)$. In 2010, Siffalovic et al. exposed an array of $6.2\,\mathrm{nm} \pm 0.7\,\mathrm{nm}$ Fe_2O_3 nanoparticles, on various Si, SiO_2, and SiN_x substrates, to UV/ozone treatment. They studied the dynamics in real-time *in situ* with grazing incidence small-angle x-ray scattering (GISAXS). They found that surfactant removal by UV/ozone treatment takes place on a timescale of at least seconds.[155] By contrast, the x-ray probe spends no more than $\tau_p = L/c$ in the sample, where L is the length of the wafer and c is the speed of light. For typical sample dimensions of 1 cm, this is equivalent to $\tau_p \approx 3 \times 10^{-11}\,\mathrm{s}$. Hence, the probe residence time is completely negligible compared to the dynamics of interest, and the x-ray yields "snapshots" of each given time.

Siffalovic et al. found that the features of the correlation function closely tracked changes in the ordering of the monolayer, as confirmed by SEM imaging. The scattered intensity from the pristine film exhibited well-defined peaks around $q = 0.9\,\mathrm{nm}^{-1}$, corresponding to a period of repetition of $\xi = 2\pi/q \approx 7\,\mathrm{nm}$, consistent with SEM imaging. After exposure, these peaks greatly diminished in intensity and shifted to higher q, consistent with the expected decrease in nearest-neighbor spacing and also confirmed by SEM. Additionally, two *new* prominent peaks occurred at lower q ($0.2\,\mathrm{nm}^{-1}$), reflecting the characteristic island size after breakup of $\approx$ 30 nm (4 to 5 particles). In all cases, the sample was studied both with GISAXS and SEM, and the results from each technique gave a coherent picture.

The results from the time-resolved measurement gave insight into the timescales of surfactant removal. Higher-order peaks for $g(r)$ persisted with little discernible change for roughly one minute. By 6 minutes, only the first two peaks remained, and the main peak had shifted to lower q to reflect the closer interparticle spacing. Repeating the experiment while blocking the O_3-producing 185 nm light enabled an ozone-free UV treatement. The same phenomena were observed, but the rate was nearly 20 times slower. Hence, an ozone-free UV treatment could be a viable candidate for achieving greater reproducibility through precise control of the exposure time.

Like the experiments of Skjeltorp and Meakin, the work of Siffalovic et al. provides a crucial *in situ* study of the time evolution of cracks. These studies complement each other: Siffalovic et al. take a particle-centered approach based on the correlation function, while Skjeltorp and Meakin focused on the cracks themselves. The latter work is helpful for clarifying the underlying mechanisms of crack formation. The former studies a system more closely related to our

own — small, surfactant-coated nanoparticles — and their correlation function approach also works with direct image analysis.

4.3.2 Image Analysis

The goal of automated image analysis is the robust automatic extraction of parameters that characterize the cracking. The same arrangement of features can yield many different arrangements of pixel intensities (i.e. images), due to changes in contrast or brightness, unevenness in background illumination, etc. Robust image analysis techniques should extract the same feature-set over a wide range of reasonable image parameters. At the same time, it is helpful when taking the images to consider how they will be analyzed, and optimize the parameters accordingly.

In this regard, there is reason to be suspicious of SEM images, because the feature dimensions depend strongly on the local environment. Figure 4.4 shows an example. Particles inside the cluster appear systematically smaller than those around its perimeter, which appear artificially extended outward from the cluster. A nearby isolated particle has its apparent size extended in all directions, occupying more than twice the apparent area of particles in the interior. Cracking makes changes to the local environment in just the way that SEM images are most susceptible to; hence, SEM is disfavored for cracking studies.

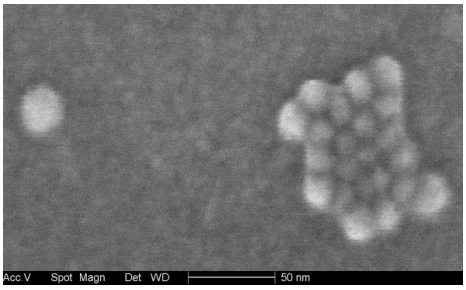

Figure 4.4: Illustration of the effect of local environment on apparent feature size in SEM images. "Landlocked" particles appear systematically smaller than particles on the edges of the array, and the isolated particle (which is effectively "all edge") has the largest apparent size of all. Since these features are near the lower limit of the resolution of the instrument, the effect amounts to a significant percentage of the feature size.

When possible, TEM images offer significant advantages in this regard. A well-focused TEM image cleanly resolves the boundary of each nanoparticle, and does not artificially inflate them as does the SEM. Additionally, TEM images have a larger pixel resolution, increasing the number of nanoparticles that can be studied in a single image. For these reasons, TEM is preferred to study cracking whenever the process is compatible with a substrate being thin enough to be electron-transparent.

TEM-based cracking studies also have a key disadvantage: they cannot mea-

sure devices, since devices are too thick to be electron-transparent. Furthermore, I have observed significant differences in dose effect between samples prepared on wafers and TEM grids, probably due to changes in the BSE signal from the different substrates. Specifically, doses which prevented all cracking in the TEM were less effective in the SEM, although in both cases the highest doses did prevent all cracking. The point of this TEM-based study is to verify that cracking can be prevented, and shed light on the mechanisms which bring this about.

Strategies for automatic particle detection

Several strategies exist for extracting features of interest from raw image files. Some, such as binary threshholding, work at the level of individual pixels, and build up features from these collections. Other techniques, such as blob detection, take a more holistic approach, where the features themselves are the units of output. Approaches such as edge detection fall somewhere in-between. This section will give a brief overview of representative techniques, with an eye towards their applicability to identifying nanoparticles and/or cracks.

Perhaps the simplest approach is *binary threshholding*.[149] The user supplies a cutoff intensity, which divides pixels into "feature" and "background" categories. Contiguous sets of feature-pixels are taken to represent individual features.

One drawback of this technique is that the presence of image noise can cause pixels to be miscategorized, leading to missed features and false positives. This can be mitigated by blurring the image *before* threshholding, or by "eroding" features below some minimum size afterwards. Another challenge is that the background intensity can vary across the image, causing the threshhold to depend on position. One strategy to overcome this problem is to employ a *gradient* threshhold, assuming the variation is not too complicated. Alternatively, "local adaptive" threshholding compares each pixel to the average of its neighbors: as long as this neighborhood is large compared to the features, but small compared to the background variation, this strategy can succeed. For studying cracking, the most serious obstacle is that noise can lead to spurious bridges between very close features. The closer the features, the more difficult and tedious it becomes to separate them.

An alternative approach known as *blob detection* can extract the particles more directly.[100] This "scale space"-based approach convolves Gaussian kernels, of varying size t, with each pixel on the image. When the kernel matches a feature in both size and position, its response is locally maximal with respect to neighboring pixels and sizes. This local maximum is taken to indicate a feature of size t at the given pixel position. Blob detection works best with isotropic features, such as nanoparticles.

One problem with blob detection is that the extracted feature location is

not fully independent of its local environment: if one side of the environment is dark, and the other side is light, the center of the feature will shift. Occasionally, true features can be erroneously split into multiple blobs, as shown by the blue-circled regions in Figure 4.5(e) and (f). Conversely, some features can be missed entirely, as shown by the red-circled regions in those same figures. The degree to which this happens clearly and systematically depends on the cracking: more features are missed when the particles are closer together.

One further problem is that the same feature may give responses at different size scales, if for example the image contrast contains substructure. In the case of nanoparticles, this was dealt with by removing the smaller of any two features whose centers are closer than the sum of their radii. This approach could, in principle, remove *legitimate* nanoparticles in favor of nanoparticle clusters. In practice, there is a wide margin between the size of the largest nanoparticle and that of the smallest detected cluster, and by setting a size cutoff inside this region, I avoided detecting any multiparticle clusters. Despite its shortcomings, blob detection is very fast, and (apart from the maximum size cutoff) fully automatic.

The most accurate particle-based strategy would involve a clear delineation of the boundary of each particle; in other words, it would be based on *edge detection*. Such a strategy was used in the work of Nabavi et al.[115], where the Canny algorithm[18] outlined the boundary of each nanoparticle. More generally, edge-like transitions can be either sharp or gradual (e.g. Figure 4.6(a) and (b)), and full scale-space edge detection should be versatile enough to handle this variation.[99]

The key challenge to an edge-based particle finder is that edges often do not form simple closed curves, as particle boundaries must. Detected edges can instead be open, or even bifurcate, confounding simplistic computational geometric approaches. Furthermore, a boundary-edge with *two* discontinuities will appear as two *separate* edges, and correctly matching points from different edge-curves increases the complexity. Strategies exist for simple cases, but the problem in general is non-trivial.

Cracked samples present additional difficulties as neighboring edges are brought very close together. First, two closely approaching edges are better described as a different kind of entity, a *ridge*, as in Figure 4.6(c). There are specialized algorithms for ridge detection, which are less computationally straightforward than edge detection.[99] As the nanoparticles are brought still closer together, the ridge diminishes in intensity and becomes harder to detect over the noise. This phenomenon is clearly seen in Figure 4.5(a), where in dense regions it becomes difficult even for human observers to precisely locate the boundaries between some pairs of nanoparticles. In practice, diminished signal intensity will lead to random discontinuities in the ridge curve, which again complicate the reconstruction of the closed particle boundaries.

Because of the complications discussed above, edge detection was not used

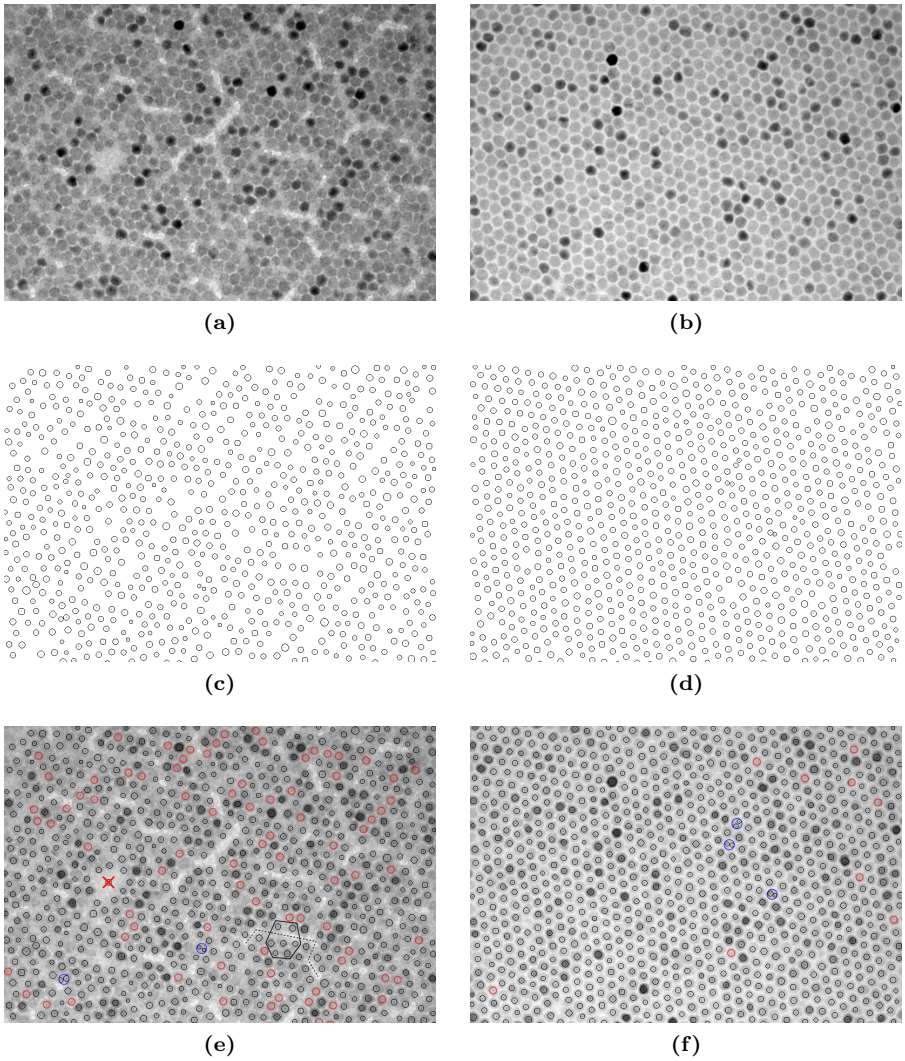

Figure 4.5: Attempt to extract particle coordinates using automated image analysis (specifically, blob detection). This example TEM grid was exposed to a 2 minute oxygen plasma to remove surfactant, then overcoated with 5 nm to 10 nm SiO_2 to lock in the ordering. Roughly speaking, the left and right columns represent cracked and uncracked regions, respectively, while the rows represent the raw images, the images the computer "sees", and an overlay of the two. (a) Raw TEM image of undosed region; severe cracking is evident. (b) Raw TEM image of region $D_4 = 14\,000\,\mu C\,cm^{-2}$, exhibiting no discernible cracking. (c) Result of blob detection and overlap removal on image of undosed region in (a). The center of each particle is denoted by a point, and the $1/e$ Gaussian radius is denoted by a circle. (d) Blob detection and overlap removal on dosed region from (b). (e) Overlay of automatically-identified particles with raw image of cracked region. Missed particles have been circled in red, "split" particles are circled in blue, and one spurious "particle" is crossed out with a red '×'. A crack is also highlighted with a dashed black line, and its effect on a hexagon of particles (solid black boundary) is illustrated. (f) Overlay for dosed region, with the same highlighting scheme for defects as in (e).

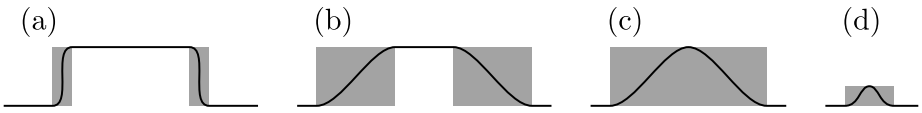

Figure 4.6: A schematic diagram explaining edges and ridges. "Transitional" regions, which would be picked up by automatic image analysis, are shown in gray (a) A feature (center) bordered by two sharp edges (light gray). (b) The same feature, but with diffuse edges. A scale-space edge detector can pick up edges of different sharpness, such as (a) and (b), in the same image. (c) When the edges move close enough together, the feature becomes more like a *ridge*: note that the gray regions have overlapped into a single transitional region. (d) As edges (such as nanoparticle boundaries) move still closer, the *contrast* of the feature can diminish.

on these nanoparticles. However, edge-based approaches hold great promise for the long run. Full knowledge of the nanoparticle boundary *implies* knowledge of its center, but opens *additional* possibilities, such as distributions of directly-measured interparticle gaps. These quantities cannot presently be obtained by other means, making boundary-based particle detection a highly worthwhile pursuit in the long run. Meanwhile, I have used blob detection with overlap removal for all of our particle extraction.

Processing strategies for particle data

We first consider the most common technique for detecting periodic ordering in images: the *discrete Fourier transform* (DFT). Any periodically repeating intensity pattern in the real space image will exhibit a sharp peak in reciprocal space, and the peak is sharper when the periodic pattern covers a wider fraction of the image. Peaks further from the center correspond to patterns with higher spatial frequency, since the Fourier transform exists in "reciprocal" space. Higher-frequency patterns require more repetitions to cover the same area. Hence, to find sharp peaks very far from the center requires a large number of repetitions, and a high degree of ordering. The quality of the ordering can thus be gauged by the distance in reciprocal space to which peaks are well-resolved.

Figure 4.7 shows an example of the DFT of raw images (shown previously in Figure 4.5) of both cracked and uncracked regions. The cracked region in Figure 4.7(a) shows six faint but distinct spots around the center, with a hazy hexagonal shell barely discernible beyond. By contrast, three hexagonal shells of well-defined spots are visible in the DFT of the uncracked region (Figure 4.7(c), reflecting the excellent quality of ordering. This signal can be clarified by taking a radial average, as shown in Figure 4.7(b). Here, the second peak for the cracked DFT (red curve) is scarcely but clearly visible, while the same peak for the uncracked region (black curve) is resolved into a double peak.

The ideal signal would contain information only about the positions of the particles. However, the particle images are not ideal points, and their spatial extent affects the signal observed. The DFT of the positional information is

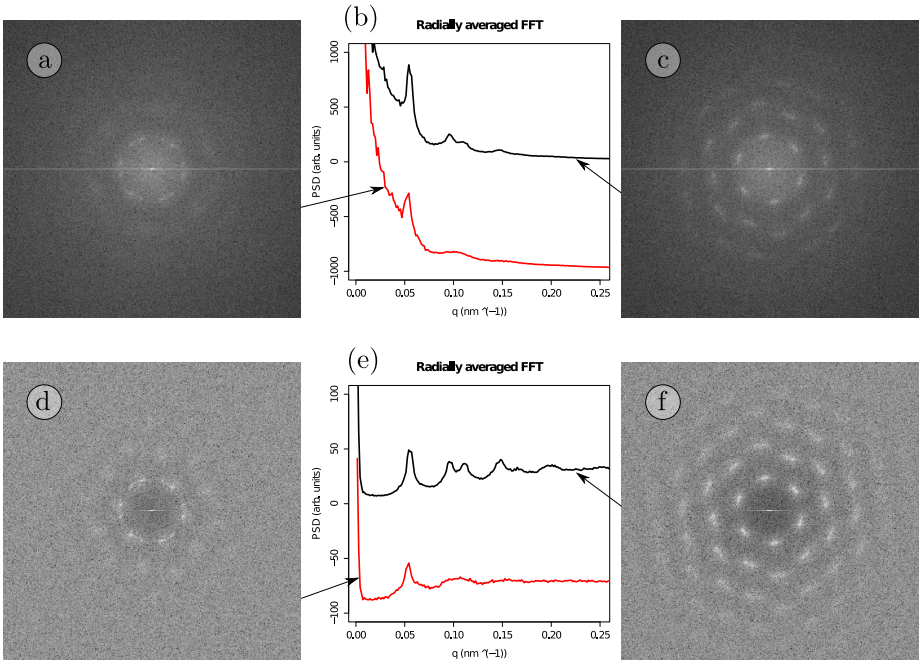

Figure 4.7: An overview of discrete Fourier transform (DFT) techniques to extract ordering information from image-based data. The cracked and uncracked images used here are the same ones shown in Figure 4.5. The first row shows DFTs from of *raw* images. (a) is the DFT of the cracked image, and (c) of the uncracked one; (b) shows plots of the radial average. Most peaks are relatively small compared to the background. (Compared to Figure 1(c) in [155], the red curve in (b) shows a much more pronounced peak, probably because the individual nanoparticles are better resolved.) Next, the DFT is taken of the *processed* data, i.e. the coordinates of the particle centers only: (d) is the DFT of Figure 4.5(c); (f) is the DFT of Figure 4.5(d); and (e) represents 1D radial averages. Here, the peaks stand out much more clearly; for instance, the double-peak near $q = 0.10\,\mathrm{nm}^{-1}$ is very cleanly resolved in the top curve in (e). For all DFT results, the 1D averages shown are taken on the *raw* DFT data, but these images show the *12th root* of the intensity. Taking the 12th root aids visibility by making dimmer features more visible; the number 12 is arbitrary and was chosen because I observed that it gave good visibility.

multiplied by the *form factor* $F(q)$ of the nanoparticles, which is high near $q = 0$ and diminishes at higher q. The shape of this curve diminishes the contrast of higher-order peaks, thereby hampering the ability to gauge positional order. Figure 4.7(d-f) show how the signal can be improved by replacing the raw image with the set of coordinates of nanoparticle centers. The third hexagonal shell in the uncracked region (Figure 4.7(f)) is now more clearly distinct, and a *fourth* shell has also appeared. The 1-D average (Figure 4.7) also shows how reducing the raw image improves detection of higher-order peaks: the double peak is resolved more cleanly, as are several other peaks which were not even visible before.

Ultimately, the DFT is useful for its quick feedback and largely parameter-free approach, but it has limited usefulness as a quantitative measure of cracking. To extract peak dimensions and compare different samples, the peaks must be fit to functional forms, and assumptions about the background are required. It is conceivable to take the *number* of peaks as the measure of order, but this is highly subjective: for instance, does the red line in Figure 4.7(e) exhibit two peaks, or three? Though inspection of 1D and 2D DFT gives valuable subjective feedback, I look elsewhere for robust quantitative measurement of cracking.

A similar concept is to use the pair correlation function $g(r)$, introduced in Section 4.3.1 in the context of scattering. Given a set of particle positions, $g(r)$ can be calculated directly simply by visiting each particle, measuring the distance to each of its neighbors, and constructing a histogram of these distances. The histogram must be weighted to compensate for higher counts at larger r, due simply to the search area $2\pi r\, dr$ being greater.

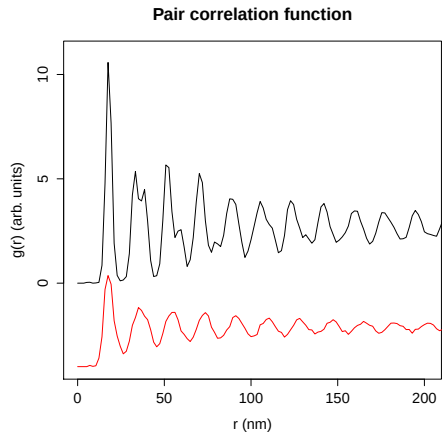

Figure 4.8: Pair correlation function $g(r)$ for the cracked (red) and uncracked (black) cases. The bottom curve has been shifted downwards by 4 units, but both curves are at the same scale.

Figure 4.8 shows an example of $g(r)$ for the cracked and uncracked cases, using the same images (from Figure 4.5) that were used to generate Figure 4.7. A key difference is that the peaks are systematically smaller from the cracked

region (bottom curve) than the uncracked region (top curve). Furthermore, the lateral resolution is superior: as with the DFT, the double nature of the second peak is resolved for the uncracked region but not for the cracked region.

Like the DFT, $g(r)$ is difficult to condense reliably into a single number which can be taken to measure the cracking. The initial peak height is one candidate (higher for better ordering), but the measured height depends critically on the boundaries of the bins. The full width at half maximum of the initial peak could also work (smaller for better ordering), but uncertainty in the peak height translates into uncertainty about the location of the half-maximum. It is tempting to fit an exponentially decaying envelope to the tops of the peaks, but both the lateral density and the intrinsic peak heights vary, and these parameters affect the *measured* peak heights. Any of these would suffice to distinguish the severely cracked sample (Figure 4.5(a)) from the uncracked sample (Figure 4.5(b)), but the difference between them is already so stark that nothing is gained for the computational expense.

To develop a more reliable quantitative gauge for cracking, we use one essential feature of cracking: particles move closer to some neighbors and further from others. In the pristine, uncracked array, each particle was very nearly the same distance from each of its neighbors. Cracking can be characterized as an increase in the *neighbor spacing variation* (NSV) — the standard deviation of the center-to-center distance to a particle's nearest neighbors. Specifically, I define the NSV of a particle i as

$$\mathrm{NSV}_i = \left[\frac{1}{N_i} \sum_{j=1}^{N_i} r_{ij}^2 - \left(\frac{1}{N_i} \sum_{j=1}^{N_i} r_{ij} \right)^2 \right]^{\frac{1}{2}}, \tag{4.1}$$

where N_i is the number of nearest neighbors for particle i and r_{ij} is the center-to-center distance of particle i from its jth neighbor. Individual particles may experience lesser or greater NSV, but the statistical *distribution* of NSV over all particles reliably shows the effects of cracking.

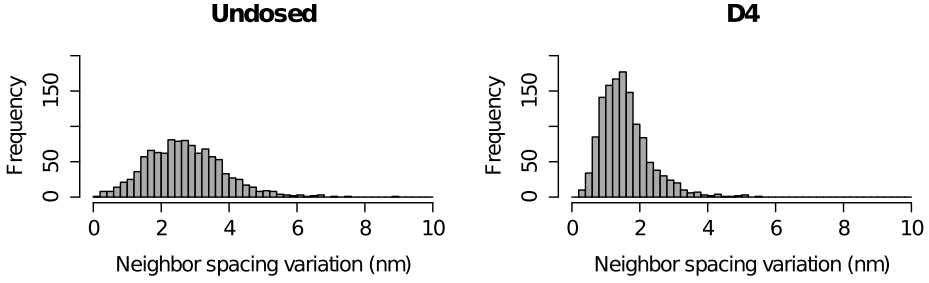

Figure 4.9: Neighbor spacing variation (NSV) distributions for (a) cracked, and (b) uncracked regions.

Figure 4.9(a) shows histograms of the NSV for 6 different doses, all taken from the $60,000\times$ images shown in Figure 4.12. Once again, the undosed region

	Mean NSV	Standard Deviation NSV
cracked	2.7	1.1
uncracked	1.6	0.7

Table 4.1: Dependence of neighbor spacing variation (NSV) on the degree of cracking. Cracked regions experience greater variation in the distances of their nearest neighbors.

stands out, while the other five are practically indistinguishable from each other. The undosed region exhibits more counts for higher NSV values, which means that undosed particles have a high variation in their nearest-neighbor distances.

Figure 4.9(b) shows statistical summaries for these distributions, along with values recalculated from $30,000\times$ images of the same sample. At both magnifications, the undosed region stands out, while the dosed regions are indistinguishable. The NSV is systematically higher at lower magnification for the same dose. This is probably because the particle size has shrunk relative to the pixel size, and smaller numbers of pixels per particle implies higher noise (analogous to smaller numbers of grains per bit, as discussed in Chapter 1).

The increase for the undosed region seems too large to be fully explained by the pixel size. This is largely an artefact of the particular experimental protocol. SiO_2 on cracked nanoparticle arrays gives a very messy appearance (see Figure 4.12(a)), and the blob detector misidentifies particle location more easily. These missed centers would tend to exaggerate the NSV. They point to the need to optimize the protocol for this experiment. More generally, these results suggest the criteria for selecting a magnification for image analysis: low enough to include many nanoparticles, but high enough to avoid noise stemming from too few pixels.

Alternative: direct crack detection

All of the previous methods have focused on extracting the particles, then investigating the cracks indirectly as a manifestation of collective nanoparticle behavior. This strategy fits well with the focus of this project, since the nanoparticles themselves are of primary interest. However, direct extraction of *cracks*, rather than particle positions, may prove to be more robust. Figure 4.5(a) suggests this possibility: even as it highlights the difficulties in distinguishing neighboring nanoparticles, it plainly shows the wide and stark cracks. In fact, these cracks are among the most prominent features in the image! And the strategy promises to extend easily to uncracked samples, such as Figure 4.5(b). The TEM cleanly resolves these more uniform interparticle gaps, which would be detected as a network of ridges.

Scale space ridge detection identifies not only the location and orientation of ridges, but their *thickness* as well. The distribution of thicknesses should serve as a robust quantitative signal, exhibiting stark contrasts between cracked and uncracked samples. Time constraints and the complexity of the algorithm have so far prevented us from implementing it in detail. However, this direction is

probably the most promising in the long term for the reliable and fully automatic quantitative measurement of cracking.

4.4 Cracking prevention with structural pinning layer

Since cracking occurs when particles are free to move across the surface, one approach to preventing it is to break translational symmetry of the surface. This is the key idea behind the *structural pinning layer* (SPL) approach. Material is deposited on top of the nanoparticles to form a conformal coating. The surfactant coating is not a single contiguous layer, but is composed of long carbon-chain molecules having various orientations. Such a structure is porous [42, p. 47], and any film deposited on top will be intricately connected with the top surfactant layer. This also allows it to come closer (in some places) to the nanoparticle surface than a flat, pre-formed substrate could come. The closeness, entanglement, and conformal nature should all help restrict the lateral movement of the particles if this deposited layer can be used as the new substrate. That is the essence of the SPL approach to nanoparticle stabilization.

For mere cracking prevention during surfactant removal, this suffices, but the ultimate goal of pattern transfer must be taken into account. The deposited material comprising the new substrate must be etchable by an RIE recipe which does not etch the nanoparticles; this rules out most metals. It must also be able to coat conformally, which rules out pure Si due to its highly anisotropic bonding, which leads to highly granular coatings. Sputter-deposited SiO_2 best fulfills the various criteria for a SPL material: other materials I tried were either difficult to etch by RIE (Pt) or did not coat conformally (Si).

Removal of the original wafer poses difficulties, because a thickness of at least hundreds of thousands of µm must be etched, while stopping precisely on a nanoparticle layer less than 15 nm thick. We solved this problem by constructing an *etch-stop hierarchy*, with progressively thinner layers of materials with radically different etch chemistries. The first layer in our etch-stop hierarchy was the silicon itself. The thickness should be minimal in order to minimize etch times, but it must be thick enough to keep a rigid planar surface for the Langmuir-Schaeffer transfer to work. We used 4″-diameter wafers which were just 220 µm thick, as opposed to 525 µm for typical 4″ wafers (or 275 µm for 2″ wafers). These wafers were thermally oxidized to a depth of 1 µm, and they are available as standard items in the Carnegie Mellon nanofabrication facility. A 15 min float on top of HF — one of the few etches which selects SiO_2 over Si — sufficed to remove the native oxide on one side. Finally, a 50 nm layer of aluminum was sputter-deposited on top of the SiO_2. The nanoparticles were transferred to this aluminum surface.

These layers were removed in the reverse order that they were deposited in.

Precise recipes are given in the Appendix.To remove the silicon, we used either SF_6-based ICP-RIE or XeF_2 gas etch, since both stop on SiO_2. Once the SiO_2 is exposed, a simple fluorine-based plasma (CF_4) removes it, and stops on the underlying Al with unusually high chemical selectivity.[54] The Al was removed by stirring in AZ400K, a common photoresist developer, for 10 minutes. The monolayer is invisible before stirring, but visible afterwards, and can be imaged in SEM.

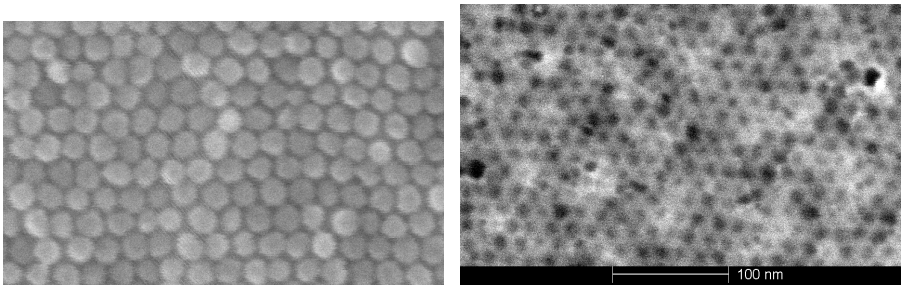

(a) **(b)**

Figure 4.10: (a) Monolayer of 13.7 nm ± 0.9 nm Fe_3O_4 nanoparticles after removing original wafer and etching surfactant, demonstrating that the ordering is preserved and cracking is prevented. (b) Substrate after the nanoparticles have been removed from the original wafer, showing a dimple pattern which was conformal to the nanoparticles.

At this point, we again have an exposed array of surfactant-coated nanoparticles, with the only difference being the new substrate. We tested the degree to which that substrate prevented cracking by exposing the monolayer to an oxygen plasma, for a duration sufficient to cause cracking in bare layers. As shown in Figure 4.10(a), the particles retained their ordering. Figure 4.10(b) shows the substrate after the particles were removed. The dimple pattern is evidence of the conformal nature of the structural pinning layer.

4.5 Cracking prevention by electron beam curing

The main drawback to the SPL approach is its complexity. Processes with large numbers of steps tend to be more susceptible to errors. Furthermore, this complexity diminishes generality, since many of the steps are very materials-specific. The ideal stabilization routine would encompass a single processing step, and would be compatible with nanoparticles of arbitrary chemical makeup on arbitrary surfaces.

I have developed such a routine, based on dosing the nanoparticles with electron beam radiation. Hints of this idea can be found scattered in the literature over the last two decades. In 1993, Lercel et al. investigated pure surfactant monolayers as novel electron beam resists for ultrahigh resolution.[94] Their con-

clusion was that the electron beam effects a *chemical change* in the surfactant, probably cross-linking it analogously to the curing process for epoxy.

The concept was extended to monolayers of surfactant-coated nanoparticles in 2001, when Lin et al. considered it as a way to pattern Au nanoparticles.[98] Undosed monolayer regions could be washed away in hot surfactant-bearing toluene, while the particles in the dosed regions held fast to the substrate. They supposed that the electron beam completely stripped the surfactant from the nanoparticles, but did not provide evidence other than the increased nanoparticle adhesion. It is difficult to see how the surfactant could be completely stripped in such a way that the ordering of the nanoparticles is retained. The cross-linking proposed by Lercel et al. can also explain increased adhesion, because the new arrangement of the carbon in the surfactant is less soluble in solvents.

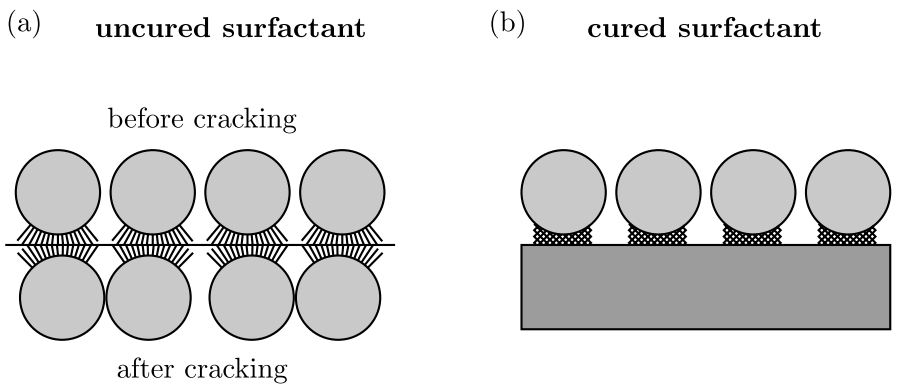

Figure 4.11: Illustration of the cross-linking explanation for electron beam order preservation. In all parts, the upper and interstitial surfactant has been removed by oxygen plasma, and the underlying surfactant is assumed intact for sake of illustration. (a) Uncured surfactant consists of independent chains. Particles can move laterally with relative ease, without the bonding of the surfactant shifting on the substrate. The uncracked and cracked cases are shown together (instead of drawing a solid substrate) to illustrate that, at these size scales, the surfactant chains bound to the surface need hardly move for neighboring particles to touch. (b) When the surfactant becomes cross-linked, the large number of connections inhibits lateral motion.

Figure 4.11 shows a schematic of the effect of cross-linking on lateral mobility. If the surfactant is not cross-linked, it consists of independent carbon chains which link the surface of the particle to the surface of the substrate. Part (a) shows schematically how this arrangement can allow for more lateral mobility. Part (b) shows cross-linked surfactant, where the added connections between chains constrain their lateral motion. If the cross-linking hypothesis is correct, then this procedure should be very materials-general, since the surfactant in practically all surfactant-coated nanoparticles consists of carbon chains.

4.5.1 Electron beam results

Monolayers of $13.7\,\mathrm{nm} \pm 0.9\,\mathrm{nm}$ nanoparticles were transferred to TEM grids coated in SiO. Different regions were exposed to exponentially varying doses $D_0 \cdots D_8$, where

$$D_i = 2^{i/2} D_0 \tag{4.2}$$

and $D_0 = 3500\,\mu\mathrm{C\,cm^{-2}}$. The grids were exposed to a 7 W O_2 plasma in the PlasmaTherm 790 RIE etcher for 2 minutes, during which time any cracking took place, then sputter-coated with 5 nm to 10 nm SiO_2 to lock them in position. TEM images from representative regions are shown in Figure 4.12.

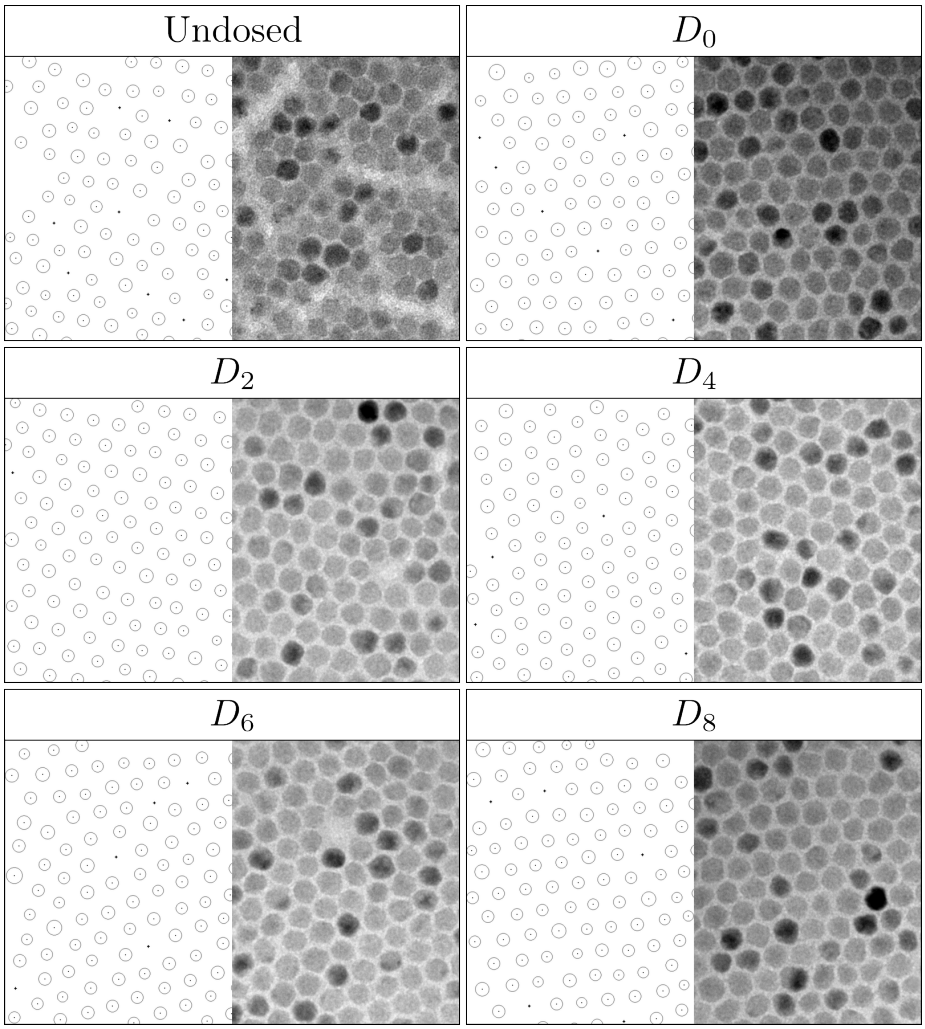

Figure 4.12: Data used for 2 minutes, at 60,000× magnification. The left half of each subfigure resulted from running blob detection on the image shown in the right half. Particles missed by the blob detector were filled in manually, and are marked by '+'-signs.

Blob detection extracted the circled features from the raw images. Sometimes different-sized features were detected on the same particle. These were automatically removed by overlap detection: when two blobs are closer than

the sum of their radii, the smaller one is removed. Missing features were added manually; these are denoted by '+' signs.

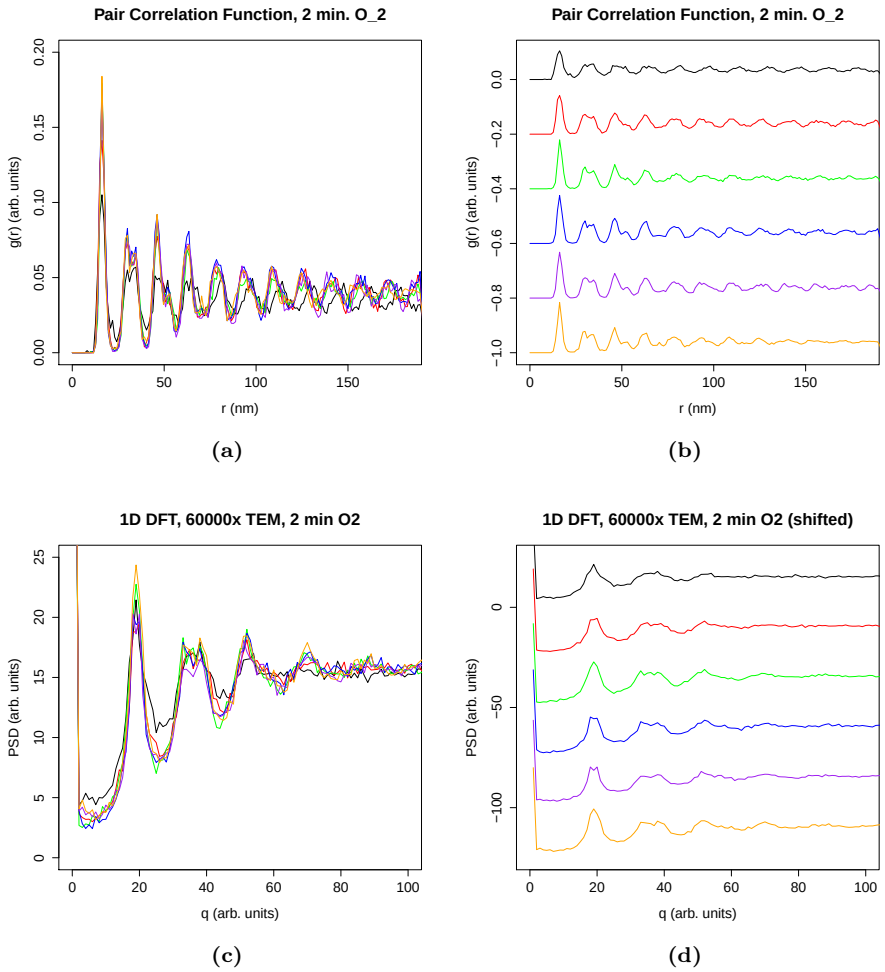

Figure 4.13: (a, b) The pair correlation function $g(r)$ for nanoparticles exposed to 2 minutes O_2 plasma. All images were taken at a nominal TEM magnification of $60,000\times$. The increased jaggedness compared to Figure 4.8 is because a smaller number of particles was used, due to higher magnification. (c, d) 1D radial average of DFT for these same images.

Figure 4.13 shows the 1D ordering diagnostics: the pair correlation function $g(r)$ (Fig. 4.13(a, b)) and the 1D radially averaged DFT (Fig. 4.13(c, d)). The undosed region is shown in black, and the other curves represent D_0 (red) through D_8 (orange). The functions are shown superposed in the left column, and shifted (but unscaled) in the right. Both diagnostics tell a very similar story. First, the undosed region is clearly distinguishable from all the others: its peaks are less pronounced, and higher-order peaks for $g(r)$ are systematically shifted. Moreover, all of the curves for dosed regions are very similar, and it is difficult to pick out any systematic trends. In Figure 4.13(b), the second and third peaks for $g(r)$ (both double peaks) are better resolved at higher doses (purple and orange

curves) than low doses (red and green), and in Figure 4.13(d), the fourth peak evolves from nonexistent at D_0 (red) to being distinct at D_8 (orange). However, both of these effects are very subtle, and it is far from clear that they are robust.

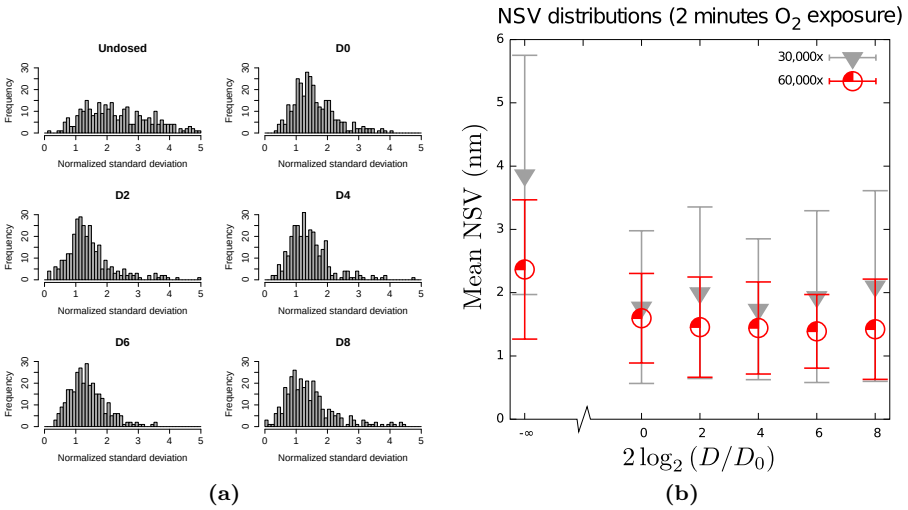

Figure 4.14: NSV distributions for the 2-minute exposure data, calculated from TEM images. (a) NSV distributions for individual regions using data from $60,000\times$ magnification. (b) Statistical summaries of NSV distributions for all regions, plotted together. (The leftmost datapoint is undosed, and each point on the right side has double the dose of the previous point there.)

The NSV distributions shown in Figure 4.14 tell much the same story, but it is easier to pick out the systematic trend in the dosed regions. Figure 4.14(a) shows the individual distributions; the undosed region is clearly more spread out, and is nonzero at higher variations, than the dosed regions. In Figure 4.14(b), the distributions are summarized by their mean and standard deviation, and plotted together to check for trends. Once again, the starkest difference is between the undosed region and the dosed regions: as expected, the average particle experiences considerably more variation in its nearest-neighbor distances in cracked regions than uncracked. In the dosed sections, a very slight downward trend is evident, as expected. However, the error bars overwhelm this difference, and a high confidence cannot be assigned to this trend. The main conclusion from the NSV agrees with the graphical techniques, and with an inspection of the images themselves (Figure 4.12): the uncracked regions are practically indistinguishable from each other.

4.6 Summary

To summarize, I have explored the problem of cracking, investigated techniques to measure it, and demonstrated novel techniques to circumvent it. Strong attractive interparticle forces pose a serious inherent challenge to the preserva-

tion of order without steric repulsion. Yet this ability is critical for industrial applications, since steric barriers block access to the substrate.

We have considered several computer vision approaches to measure cracking. These include raw image analysis (DFT), particle-center extraction (binary threshholding, blob detection), particle-boundary extraction (edge detection), and direct crack measurement (ridge detection). We opted to extract particle centers by blob detection with overlap removal, and manually fill in the few particles missed, as a tradeoff between robustness and ease of implementation. Methods based on edges and ridges were not implemented, but are suggested as a clear research trajectory for longer-term cracking measurement. Several analyses can be used to gauge the ordering of a pointset, including Fourier transforms and $g(r)$. The preferred technique (NSV) gauges the variation in nearest-neighbor distances, since this quantity is intimately tied to an essential feature of cracking: particles in a cracked array move closer to some neighbors and further from others.

We pioneered the first technique which allowed the order to be preserved after the surfactant was removed. The structural pinning layer holds the particles more tightly to the substrate, both by its intimate connection with the surfactant and its particle-conformal topography. Ultimately, this process is too complex and too materials-restrictive to find wide application. However, it is notable for first demonstrating that the problem could be solved.

Electron beam curing also stabilizes the nanoparticles, and remedies both major defects of the SPL method: it is a single, materials-general processing step. The surfactant undergoes a chemical change, probably cross-linking, and the nanoparticles are bound more tightly to their original positions on the substrate. Even relatively small doses ($D_0 = 3500\,\mu\mathrm{C\,cm^{-2}}$) allowed the order to withstand two minutes of 7 W oxygen plasma treatment, while undosed regions on the same sample exhibited complete cracking. Regions receiving a wide range of doses greater than D_0 were nearly indistinguishable from one another, suggesting that longer O_2 plasma exposure times are needed to fully explore parameter space.

All of these experiments are open to an alternative interpretation: it is also possible that the surfactant in the order-preserved regions was not removed. The goal of order preservation is the ability to pattern the underlying wafer, which would be impossible if the surfactant still remained. We cannot rule out the presence of small amounts of residual carbon. However, the successful patterning demonstrated in the next chapter gives confidence that the surfactant really was removed in the experiments in this chapter.

5 Pattern Transfer Processes

The techniques outlined in the previous chapter enable nanoparticle arrays to retain their ordering after exposure to O_2 plasma. Cracking was observed in untreated regions, which suggests that the plasma removed the surfactant at least partially. Questions remain about the fraction of surfactant which was removed, both in the cured and the uncured regions. However, cracking is of interest for this project only to the extent that it prevents pattern transfer. The real test for sufficient surfactant removal is therefore to attempt to transfer the pattern; if the attempt succeeds, then *ipso facto* "enough" surfactant has been removed. This chapter discusses several techniques which have succeeded in transferring patterns.

There are two strategies to exploit the newly-cleared interparticle gaps. Both involve using them as channels for the flow of materials. Positive patterning makes dots where the nanoparticles originally were. The material flow here is two-way: etchants flow inwards through the gaps, and etch *products* simultaneously flow outwards. Negative patterning makes holes were the nanoparticles originally were. The material flow is one-way: atoms travel through the gaps and stick on the substrate, causing a buildup of material over time — an antidot lattice. When the original nanoparticles are removed, these holes can be made deeper by etching with the antidot lattice as a mask. These concepts are illustrated schematically in Figure 5.1.

Each approach has advantages and disadvantages relative to the other. For dots, the result has the structure desired for BPM, i.e. a regular array of disconnected dots. One downside is that chemically selective etching of magnetic materials is difficult. Additionally, for material to flow outward through the gaps, etchants must flow *inwards* through the gaps, and there will be two opposing material flows. As gaps are scaled down, they approach the size of the etchants, and this "traffic jam" becomes increasingly problematic. Negative patterning does not suffer from this defect, because material flow is strictly one-way. The downside is that the tone of the pattern is the opposite of the desired tone. However, such a pattern could still be used to make intermediate masks, where subsequent processing steps include a second pattern inversion which leads to dot patterns. Moreover, as mentioned in Chapter 1, this approach has been adopted by the data storage industry as the most likely course of action.

As discussed in Chapter 1, little work has been done on patterning from surfactant-coated nanoparticle arrays. Hideki Masuda's group created a pat-

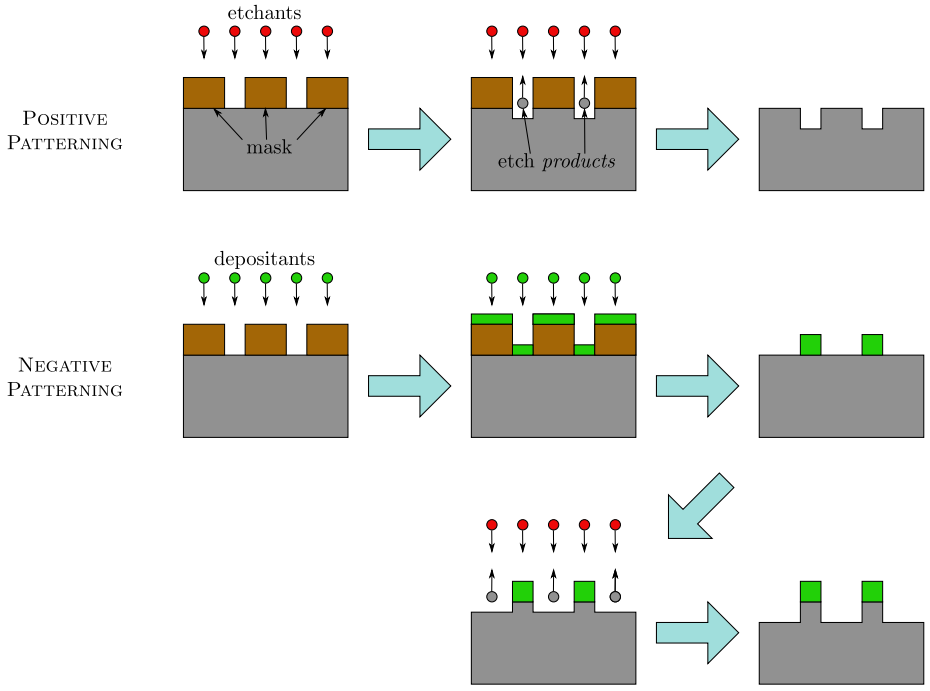

Figure 5.1: A schematic illustration of the main concepts in positive and negative pattern transfer. Both techniques begin with the same mask, shown in brown, on top of the gray substrate (left column). Material — either etchants (red) or depositants (green) — then impinges on the masked substrate (middle column). Finally, the particles are removed (right column), resulting in either a pattern of dots (positive patterning) or antidots (negative patterning).

terned layer by overcoating a nanoparticle array in Al, then removing the nanoparticles and the original wafer.[108] The resulting dimples served as nucleation sites for anodization of the alumina, which created an ordered array of deep pores with the same pattern as the nanoparticles. This technique has several advantages: it works with smaller nanoparticles (here, just 13 nm pitch) and the nanoparticles are not incorporated into the final structure. However, it does not allow for investigation of mass transport through extremely small gaps, which is important for future nanomasks at high areal densities. Furthermore, it cannot be used for positive transfer. Finally, this technique is restricted to pattern a single material (Al).

A better point of comparison is the state of the art in block copolymers (BCP), because BCP is another self-assembling mask technology which has been used for both positive and negative patterning, and it is extensible into the sub-10 nm regime. There are three stages in BCP-based BPM fabrication to which we might compare our technique. In order of decreasing sophistication, they are demonstrating magnetic recording, transferring patterns of any sort, and simply creating the mask. The BCP mask with the highest density known to the author has a pitch of 6.9 nm, with features just 3 nm in diameter — well into the sub-10 nm regime.[130] However, only the mask creation was demonstrated; there are no examples of patterns transferred through such a fine BCP mask.

Several groups have demonstrated pattern transfer from block copolymers into underlying substrates. From 2001 to 2007, most successful patterning demonstrations had a period of 30 nm to 42 nm.[128, 153, 1] These demonstrations included chemically selective transfer into silicon nitride[128], silicon oxide[1], and silicon[1], as well as a variety of liftoff-based techniques patterning dots and antidots of Ti/Au, Cr, and Au/Cr. In 2008, Park et al. demonstrated the ability to fabricate dots, antidots, and ring structures of Au from the same 45 nm-pitch BCP template.[125] Very recently, the pitch of successfully transferred patterns has come down, with the demonstration of 16 nm features on a 27 nm pitch[22] and 8 nm linear features on a 17 nm pitch.[78] To the best of my knowledge, the latter is the smallest pattern transferred from a BCP template.

Full-fledged recording demonstrations are routinely made for media fabricated from block copolymer templates. In 2004, CoCrPt dots were created with a 80 nm pitch.[80, 116] By 2006, FePt dots had been created with a pitch of just 30 nm, an improvement in areal density by more than a factor of 7.[62] 20 nm pitch was demonstrated with CoCrPt in 2008[72], and a more mature 28 nm pitch on Co/Pd multilayers was shown in 2010.[61] These numbers indicate the pace of recording demonstrations has slowed in the past 6 years, suggesting that fabrication of patterned magnetic media may encounter difficulties below roughly 20 nm.

These three stages — mask fabrication, pattern transfer, and recording demonstration — set the context for the present work. The focus of this project is the middle stage: demonstration of pattern transfer. The smallest pitch in

the nanoparticle arrays discussed here is 15.7 nm ± 0.2 nm, which compares favorably to the comparable state of the art in BCP pattern transfer. The 6.9 nm pitch shown by Park et al. is smaller, but it is not appropriate to compare this to our patterning results. Rather, it should be compared to the smallest pitch of a *nanoparticle monolayer*, and smaller examples have been known since at least 1997, when 1.0 nm-diameter nanoparticles exhibited self-assembly with a pitch smaller than 4.0 nm.[145] Hence, to the extent that it demonstrates pattern transfer through nanoparticle arrays, the present work meaningfully advances the state of the art.

Three such techniques have been successfully demonstrated. Section 5.3.1 shows positive transfer using the structural pinning layer (SPL) technique. This is the first demonstration of transfer through nanoparticle arrays, but the complexity of the approach makes applications unlikely. Positive transfer through e-beam cured nanoparticles, as discussed in Section 5.3.2, shows improvement in terms of industrial viability, reproducibility, and consistency within each wafer. Finally, Section 5.3.3 shows that e-beam curing also enables negative transfer, and that the antidots themselves can be used as an etch mask to enhance the height contrast. Common to all these techniques is the need to remove the nanoparticle mask — both to verify the patterning, and to allow further processing — so the technique for doing so is first discussed.

5.1 Removal of the nanoparticle mask

Conventional mask liftoff employs a solvent which dissolves the mask material.[103, p. 14] With surfactant-coated nanoparticles, there are two options: the solvent can attack either the inorganic particle cores or the organic surfactant coating. Nanoparticle cores which self-assemble encompass a very wide variety of chemical identities: metals (Au, Ag, Co), ferrites (Fe_3O_4, $MnFe_2O_4$), semiconductors (CdSe, CdTe, CdS), and more. By contrast, surfactant coatings are much more similar: they all contain aliphatic carbon chains, and are bonded to the particle by one of a few functional headgroups (carboxylic acid, amine, thiol, etc.). Therefore, methods which attack the surfactant rather than the core are likely to be more general. They are also likely to be quicker, since removing still-intact particles is faster than waiting for them to dissolve.

I have found that phosphoric acid readily removes both Fe_3O_4 and MnO nanoparticles from the SiO_2 surface, without dissolving the particles themselves. Figure 5.2 shows the results of an experiment which demonstrates these claims. A nanoparticle array was coated with 2 nm of sputter-deposited Pt, which formed a contiguous conformal sheet over the nanoparticles. After stirring in 85 wt.% H_3PO_4, some portions of the sheet were lifted up and overturned. The observation of intact particles on the underside shows that H_3PO_4 does not dissolve them, while the uplifting shows that H_3PO_4 attacks the particle-substrate bond.

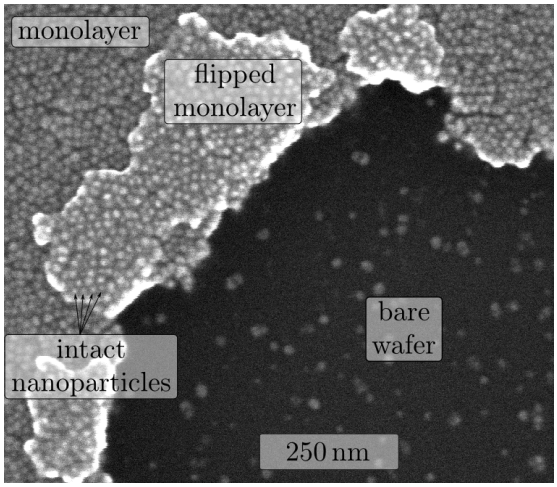

Figure 5.2: An illustration that H_3PO_4 does not dissolve Fe_3O_4 nanoparticles, but does attack their bond to the surface. This sample was overcoated with a 2 nm thick contiguous layer of Pt, which keeps the nanoparticles bound to their neighbors. After stirring in H_3PO_4, some sheets were overturned and placed upside-down on other monolayer regions. The fact that the H_3PO_4 picked up these sheets is evidence that it attacks the bond between the particles and substrate. The fact that intact particles are still seen shows that H_3PO_4 does not dissolve the particles.

The precise nature of this attack is presently unclear. H_3PO_4 is noted for being less reactive with organics in general[36], and alkane-like portions of the aliphatic chains are chemically unreactive apart from substitution and combustion reactions.[49, pp. 47–50] In the case of oleic acid, three likely sites for the reaction are the surfactant-particle bond, the surfactant-substrate bond, and the double carbon bond which appears in the middle. The latter is unlikely, because the ability to remove the nanoparticles is unaffected by e-beam curing, which attacks carbon-carbon bonds. Hence, it appears the chemical attack takes place at one or both ends of the surfactant.

The effect of H_3PO_4 on AlO_x is also unclear from the literature. 85% H_3PO_4, which I used, has been found to etch AlO_x at more than $5\,\mathrm{nm\,min^{-1}}$ at 160 °C.[168] The anodized aluminum oxide (AAO) community typically uses 5% H_3PO_4 to etch AlO_x, but always with either elevated temperature [134, 107] or applied voltage[32]. The latter group found etch rates of $8\,\mathrm{nm\,h^{-1}}$, but I have observed films of 1 nm effective thickness to hold up to 48 minutes of stirring in 85% H_3PO_4. To summarize, it seems that H_3PO_4 will etch some forms of AlO_x at elevated temperatures or applied voltages, but my observations rule out that the conditions I use damage the AlO_x to the same extent.

5.2 ARDE and extremely small gaps

Aspect ratio-dependent etch rate effects (ARDE) were discussed in Section 3.2.3. In the extreme small-gap regime, these effects take on critical importance, because even small etch depths can correspond to large aspect ratios. Therefore,

before discussing the results of pattern transfer, it is worth developing models for ARDE in greater detail.

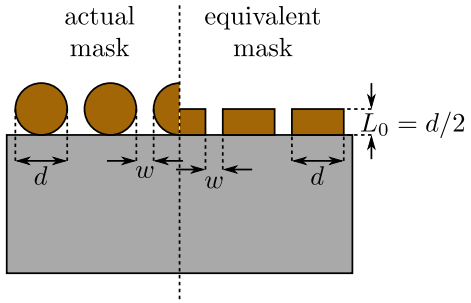

Figure 5.3: A comparison of the actual mask to the equivalent mask used for purposes of making ARDE calculations. The equivalent feature width is the nanoparticle diameter d, and the equivalent gap is taken as the smallest distance between two neighbors. The equivalent height L_0 is taken to be the height where the gaps are smallest, i.e. half the nanoparticle diameter $L_0 = d/2$.

The various causes of ARDE — ion shadowing[150], neutral shadowing[50], differential charging[7, 73], and Knudsen neutral transport[28] — are difficult to distinguish in practice.[82] This observation motivated Keil and Anderson to develop a phenomenological model, which makes no assumptions about the underlying mechanism.[82] The quantity they track is $L(t)$, the height difference between the initial top of the mask and the current depth of the pit. For a gap of width w, $L(t)$ evolves to first order according to

$$\frac{d}{dt}L(t) = A + B\left(\frac{L(t) - R(t)}{w}\right), \tag{5.1}$$

where $R(t)$ is the cumulative mask erosion such that $R(t = 0) = 0$. A is the etch rate of the bare film, and it is straightforward to measure. B characterizes the severity of the slowdown due to ARDE, and is thus negative. It is not measured directly; rather, it is obtained as a fit parameter by measuring the depths of patterns with different widths.

Equation 5.1 can be solved for $L(t)$, to yield

$$L(t) = L_0 e^{\frac{Bt}{w}} - \frac{Aw}{B}\left(1 - e^{\frac{Bt}{w}}\right) - e^{\frac{Bt}{w}} \int_0^t B\frac{R(t')}{w} e^{\frac{-Bt'}{w}}\, dt'. \tag{5.2}$$

The first term represents the initial depth L_0 of the trench, and this term fades *out* exponentially (recall that $B < 0$). The second term represents a limiting depth $Aw/|B|$, which fades in exponentially. The final term accounts for the mask erosion.

To apply this model to our systems, we make two assumptions. First, we neglect the mask erosion, because we have set the RIE power low enough that

no physical sputtering was observed; hence, $R(t) = 0$, and we have

$$L(t) = L_0 e^{\frac{Bt}{w}} - \frac{Aw}{B}\left(1 - e^{\frac{Bt}{w}}\right). \tag{5.3}$$

The second assumption is needed because I have no information about values that B might take in this system. B is usually estimated by fitting data on the gap depth as a function of gap width, but I am solving the inverse problem: I want to estimate the depths themselves using B as input. As a guide, I begin by noting that when ARDE effects are absent, the patterning is unchanged by the scaling $(A, t) \rightarrow (\lambda A, t/\lambda)$, for any scalar λ. In other words, running a recipe with twice the etch rate for half as long results in the same structure. My assumption is that the same will be true when ARDE effects are taken into account. By inspecting Equation 5.3, it is plain that the scaling $(A, B, t) \rightarrow (\lambda A, \lambda B, t/\lambda)$ leaves the height contrast unchanged. This implies that the ratio B/A is independent of λ, and I treat it as more fundamental than B by defining

$$B = -\beta A, \tag{5.4}$$

for some $\beta > 0$. β can now be estimated from data in other publications. Rewritten in terms of β, Equation 5.3 now appears as

$$L(t) = L_0 e^{\frac{-\beta At}{w}} + \frac{w}{\beta}\left(1 - e^{\frac{-\beta At}{w}}\right). \tag{5.5}$$

Keil and Anderson implicitly find $\beta_{\mathrm{KA}} = 0.057$ for borophosphosilicate glass (BPSG) etched by a $Ar/C_4F_8/CO/CF_4$ mixture.[82] Kokkoris, Gogolides, and Boudouvis performed simulations and found a limiting aspect ratio of 8;[91, Fig. 4] this implies $\beta_{\mathrm{KGB}} = 1/8 = 0.125$, which is the same order of magnitude as β_{KA} which was calculated from Keil and Anderson. Because the model I use was developed by Keil and Anderson, I take $\beta = \beta_{\mathrm{KA}}$ as my default in the following discussions.

$L(t)$ includes the mask height, but the etch depth $D(t)$ is of greater interest, because the mask will ultimately be removed. If the mask has initial height L_0, then

$$D(t) = L(t) - L_0. \tag{5.6}$$

Combining Equation 5.6 with Equation 5.5 yields an expression for the etch depth as a function of time:

$$D(t) = \left(\frac{w}{\beta} - L_0\right)\left(1 - e^{-\beta At/w}\right). \tag{5.7}$$

The first factor in round brackets gives the limiting etch depth,

$$D_{\max} = \left(\frac{w}{\beta} - L_0\right), \tag{5.8}$$

and the second factor describes the approach to that depth as a function of time.

Nanoparticle masks must be analyzed with care, because their cross-section varies as a function of height. Figure 5.3 shows the mask structure I take to be equivalent to a nanoparticle array having particle diameter d and gap width w. The height L_0 of the equivalent mask is the height where the nanoparticles have their minimum separation w. The diagram shows that this occurs halfway to the top; hence, $L_0 = d/2$.

Note that there is no fundamental obstacle to $D_{\max}$ taking negative values, which is unphysical. The proper interpretation in this case is merely that etching does not occur, because the substrate is simply too deep: the initial aspect ratio is larger than the limiting aspect ratio. This implies another experimental constraint on β, obtained by setting the limiting etch depth to zero and solving for β. For nanoparticles of diameter d having gap width w, if *any* etching is observed (regardless of depth), then with certainty $\beta < \beta_{\max}$, where

$$\beta_{\max} = \frac{2w}{d}. \tag{5.9}$$

It is instructive to apply this model to a concrete example: our Fe_3O_4 nanoparticles. Here, $d = 13.7\,\text{nm} \pm 0.9\,\text{nm}$ and $w = 2.0\,\text{nm} \pm 0.9\,\text{nm}$, so $\beta_{\max} = 0.292$. This value is comfortably above both other estimates I made from the literature, β_{KA} and β_{KGB}, yet is of the same order of magnitude. Such agreement is as good as could be expected for parameters estimated in this manner, and builds confidence that my approach is basically consistent.

Since the precise value of β is poorly derdetermined, it is important to gauge the effect of changing β on the severity of ARDE. Figure 5.4(a) shows these effects by plotting the actual etch depth versus the nominal depth, i.e. the depth that would obtain without ARDE effects. β takes on values from a geometric sequence which is centered about the Keil and Anderson estimate of β_{KA}. It is plain that at $\beta = 4\beta_{\text{KA}}$ (which is close to $\beta_{\max}$) there is virtually no etching (purple curve). On the other hand, when $\beta = 0.25\beta_{\text{KA}}$ (red line), etching is not significantly slowed relative to the bulk etch rate (black dashed line) until higher nominal depths are achieved. Note that β is varied over just one order of magnitude, but the possible etch depths span several orders of magnitude, because the etch depths are shifted by $L_0 = d/2$.

Figure 5.4(b) illustrates the effect of changing the gap width w, which can vary due to cracking. These effects are similar to changing β, since Equation 5.7 only depends on the ratio β/w. However, the message is different. Figure 5.4(a) measures only the effect of our ignorance, since β is poorly known but constant. By contrast, Figure 5.4(b) corresponds to *physical depth variation*, since not all gaps in the sample will be the same. Note, for instance, that gaps of 2 nm are etched to more than twice the depth of gaps of 1 nm!

This observation has practical implications for *any* extreme small-gap etch-

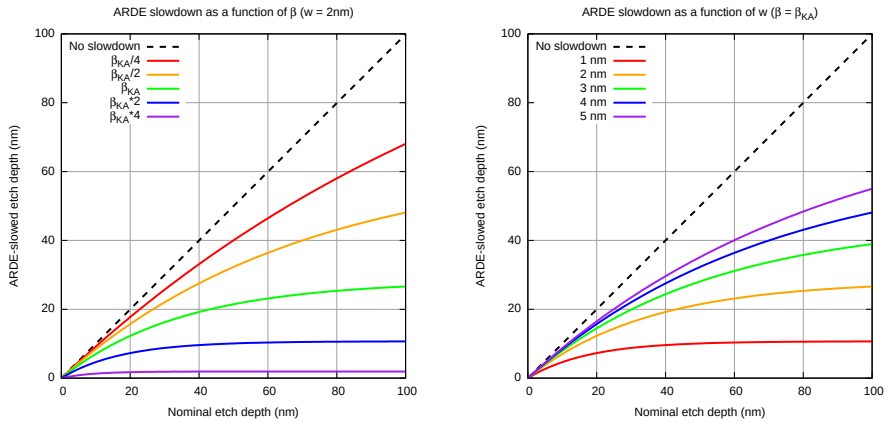

Figure 5.4: Effect of parameters on ARDE slowdown. (a) Effect of varying the ARDE severity, β, while keeping $w = 2.0\,\text{nm}$. Here, $\beta_0 = 0.057$. (b) Effect of varying the width, w, while keeping $\beta_0 = 0.057$. $w = 2.0\,\text{nm}$ is the default value, but due to cracking, we consider values ranging from 1 nm to 5 nm.

ing: small fluctuations in gap width can lead to large fluctuations in etch depth. Note, by contrast, that the depth difference between 4 nm and 5 nm gaps is very slight compared to the difference between 1 nm and 2 nm gaps. This is partially because the ratios are different (1.25 vs. 2), but the finite mask thickness makes this discrepancy even worse.

In order to achieve the same depth for all features, an etch-stop layer could be employed. The smallest gap w_{min} should first be measured by image analysis on the mask. This value can be substituted into Equation 5.8 to yield the maximum thickness which can be patterned: as long as the layer is no thicker, and the layer beneath chemically resists the etch, ARDE can be defeated.

5.2.1 Corrections to the model

There are at least two reasons why the above thicknesses should be treated as *upper limits* as a function of time, and not actual etch depths. First, there is the phenomenon of *etch stop*: the etching rate drops to a negligible value at a certain aspect ratio. This is different from the limiting value predicted by Equation 5.8; typically, etch stop happens for *smaller* aspect ratios.[82] The mechanism is based on fluorocarbon polymers sputtered from the sidewalls and redeposited at the bottom of the pit.[121] Keil and Anderson found that for their sample, this phenomenon decreased the limiting aspect ratio by a factor of more than four.[82] A similar reduction for our nanoparticle mask would lead to an etch depth of 1.9 nm instead of 28.2 nm.

Second, although typical analyses have neglected the size of the etchants, the gaps in nanoparticle arrays are no longer much bigger than the species involved in etching. This will also tend to slow the actual rates compared to predicted

rates. To see this, consider an extreme case. The etchant species is the F^- ion or F atom, while the reaction product is mainly SiF_4. The diameter of the latter is larger than the diameter of either of the former. Therefore, in a hole with diameter smaller than the SiF_4 molecule, no etching will occur, even though the above models predict a finite depth. This slowdown is a novel feature of pattern transfer through extremely small gaps, and it constitutes a serious obstacle to the long-term extension of exponential growth trends for the data storage and semiconductor industries.

5.2.2 Summary of ARDE effects

To summarize, the etching rate for small gaps is slower than for wide ones. A variety of underlying mechanisms cause this to be the case, and most of them depend on the *aspect ratio* rather than the etch depth. I have applied a simple phenomenological model to aid in determining the actual etch depth. The parameters of the model are not precisely known for my system, but rough estimates can be obtained from previous work. Phenomena such as etch stop and extreme small-gap effects suggest that this model delivers optimistic upper limits, and that true etch depths are likely to be smaller. Finally, the model predicts that gap width fluctuations can be translated into undesirable large fluctuations in pattern depth. This seems to be a problem intrinsic to all etching in the extreme small-gap regime.

5.3 Results

5.3.1 Positive transfer using the SPL technique

The process for transferring a nanoparticle monolayer to a structural pinning layer, and removing the interstitial surfactant, was described in Section 4.4. This section will describe the subsequent pattern transfer, and the removal of the nanoparticle mask.

Figure 5.5 shows two examples of successful pattern transfer. Part (a) shows a nanoparticle monolayer on top of the SPL for comparison's sake. Part (b) shows the results of attempting to transfer the pattern *without* stabilization: the features correspond to particle aggregates, and individual particle outlines are not apparent. Parts (c) and (d) show two samples which exhibited successful pattern transfer, enabled by the SPL. The sample was exposed to a CF_4-based RIE plasma (5 mTorr pressure, 18 sccm gas flow) for at least 50 minutes, and the particles were removed by rinsing in H_3PO_4. The arrow in part (d) indicates a boundary of the patterned region; the patterned features blend smoothly into the unpatterned surroundings, which is evidence that the particles really were removed.

The question of etch depth is difficult to address from these SEM images because of the extremely small gaps. The gap width is small compared to

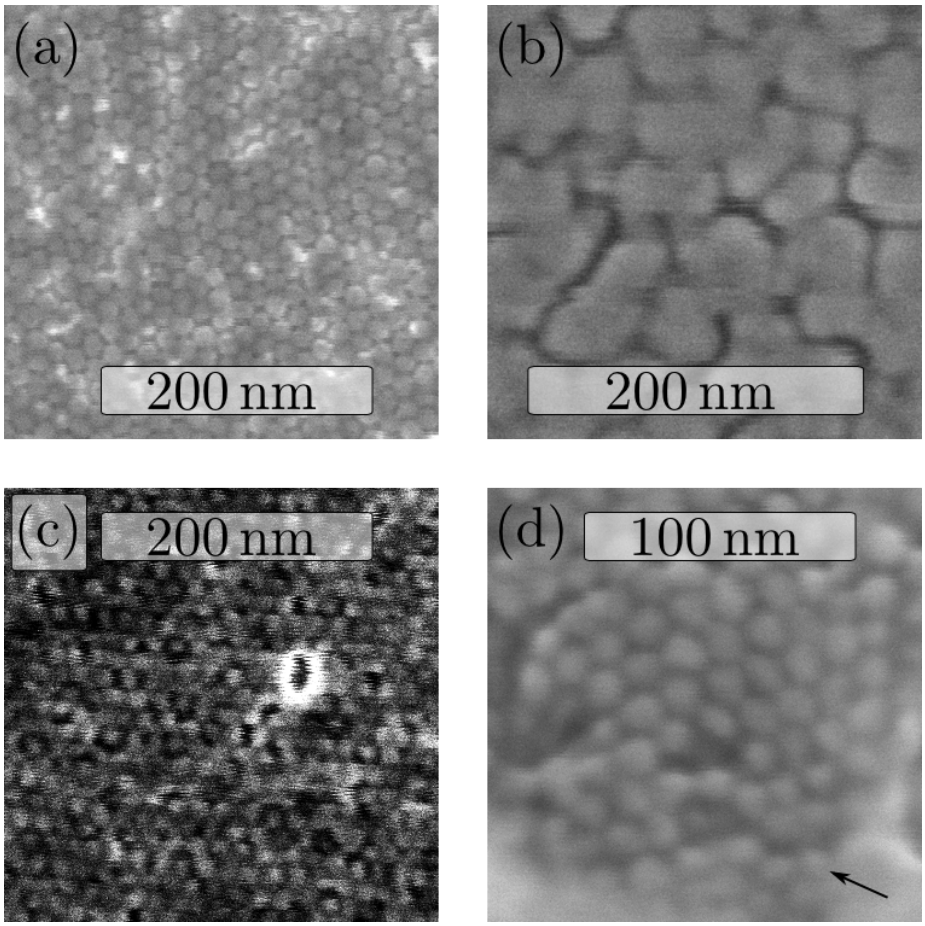

Figure 5.5: (a) A nanoparticle monolayer on the carrier wafer, after the original wafer was etched away. (b) For comparison, the results of pattern transfer attempted without particle stabilization. (c) Patterned SPL layer affixed by thick epoxy and etched for 60 minutes. (d) Patterned SPL layer affixed by thin epoxy and etched for 50 minutes.

the e-beam interaction volume, making it difficult to resolve deep topography. Furthermore, even small etch depths correspond to large aspect ratios, making it difficult to measure the bottom. Cross-sectional TEM or SEM would be needed to give an accurate measurement of the etch depth. Nevertheless, the analysis in Section 5.2.1 enables rough estimates.

I measured the bare film etch rate A to be $5.2\,\mathrm{nm\,min^{-1}}$. Both etching times were at least 50 minutes, so $\beta At/w \gg 1$ and the limiting etch depth obtains in both cases. The analysis in Section 5.2.1 makes two suggestions about the depths: they probably range from 1.9 nm to 28.2 nm, and the lower end of this range is more probable than the higher end.

Notice that the slowdown may be extremely severe for these ultra-small gaps: if the final estimate of 1.9 nm is correct, then the etch depth has been reduced by more than 99% compared to the bare film. A large part of this is due to the height of the mask — the nanoparticle radius — since it contributes to the *effective* aspect ratio but gets subtracted from the etch depth. Fortunately, for constant gap size, this effect is less prominent for smaller particles, which speaks favorably to the extensibility of this technique. Second, the actual parameter values should be taken with a grain of salt, particularly the value of β, which should be obtained from a fit to actual cross-sectional TEM data. Figure 5.4(a) shows that variation in β can have significant effects on the evolution of the etch depth, and its limiting value.

In Figure 5.5 (c) and (d), more contrast is visible for wider gaps. ARDE suggests that these gaps really are deeper, and this should account at least partially for the observed contrast. However, it must be remembered that even for equally deep gaps, the wider ones would appear darker in SEM, because the interaction volume intersects less material and liberates less electrons. This underscores the need for cross-sectioning to gauge etch depth. However, given that cross-sectioning is expensive and time-consuming, and that this technique is superseded by the techniques described next, the rough estimates already given will have to suffice.

To summarize, the SPL process works, in that the images in (c) and (d) are considerably improved over the untreated patterning shown in (b). Nevertheless, it is both complicated and materials-restrictive. Additionally, it presently exhibits several practical obstacles: the film has a tendency to buckle when the original wafer is removed, and the removal process can leave a thin chemical residue which hampers the ability to get consistent pattern transfer. These considerations likely preclude this process from serious industrial use. It is, however, noteworthy for being the first demonstration that nanoparticle arrays can be stabilized enough to be used as nanomasks for pattern transfer.

5.3.2 Positive transfer using e-beam curing

E-beam stabilization does not require the monolayer to be transferred to a new wafer by a complicated process. As this section will show, the stabilization is also sufficient to enable faithful pattern transfer. Thus, my publication of this technique[66] is the first demonstration that the pattern of a nanoparticle array can be transferred into an *underlying* wafer.

The preparation of an e-beam stabilized monolayer with surfactant removed was described in Section 4.5. To transfer the pattern, I exposed the monolayer to a CF_4/CHF_3 RIE plasma, with 18 sccm total flow and 5 mTorr nominal pressure, for etch times ranging from 20 to 40 minutes. The plasma power was 10 W when Si was etched and 15 W for etching SiO_2. Subsequently, I removed the nanoparticles by rinsing in H_3PO_4 as described in Section 5.1.

First, I demonstrate the ability to pattern SiO_2, the same material as used with the structural pinning layer. Here, we exposed with a dose of 7000 µC cm^{-2} in the pattern of the letters `"CMU"`, as shown in Figure 5.6(a) after the nanoparticles have been removed. Zooming in on one region, inside the letter `"C"` (Figure 5.6(b)), we see the microstructural differences that make these letters visible on a larger scale. The most prominent features are wide, black holes, which are sparse in the cured region and dominant in the uncured. Since these regions were identical before dosing, we hypothesize that the holes began as cracks, which the long etching (60 min) has widened. Smaller (un-widened) cracks are seen in the patterned region, which nevertheless clearly exhibits the outlines of the formerly-present nanoparticles in the mask. Features in the undosed region, by contrast, less closely resemble the original nanoparticles.

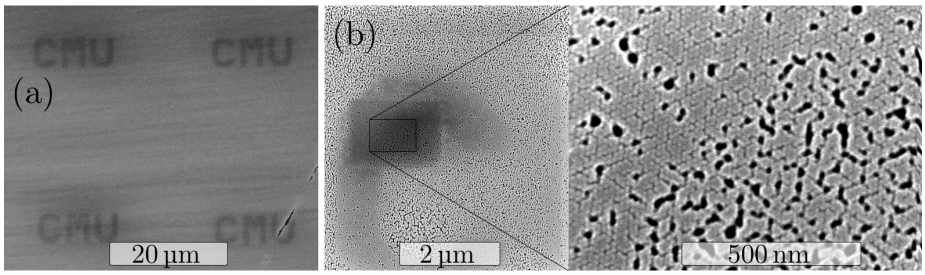

Figure 5.6: Pattern transfer from e-beam stabilized nanoparticles into SiO_2. (a) Zoomed-out view of the `"CMU"` patterns written, dosed with 7000 µC cm^{-2}. (b) Closeup view of the letter `"C"` in one of the patterns. The inset shows that the dosed region (top left) is better ordered than the undosed, where the long etching has turned cracks into chasms.

Figure 5.7 shows pattern transfer into Si as well; again, the particles have been removed, so only the patterned wafers appear. Three regions are shown, corresponding to three different doses, to show the effect of ebeam dose on quality of patterning. The undosed region (Figure 5.7(a)) exhibits poor pattern transfer, to the degree that features corresponding to individual nanoparticles cannot be seen. Such features first appear in the next region (Figure 5.7(b)),

which received a low dose of 4950 µC cm^{-2}. Despite the clear improvement relative to Figure 5.7(a), the pattern is still dominated by cracking. A much higher dose (39 600 µC cm^{-2}; Figure 5.7(c)) yields substantial improvement in fidelity. Now, single-particle features dominate the pattern; cracking, though still present, is minimal. Clearly, e-beam dosed nanoparticles remain stable even during pattern transfer, and as found in Section 4.5, higher doses have better effects.

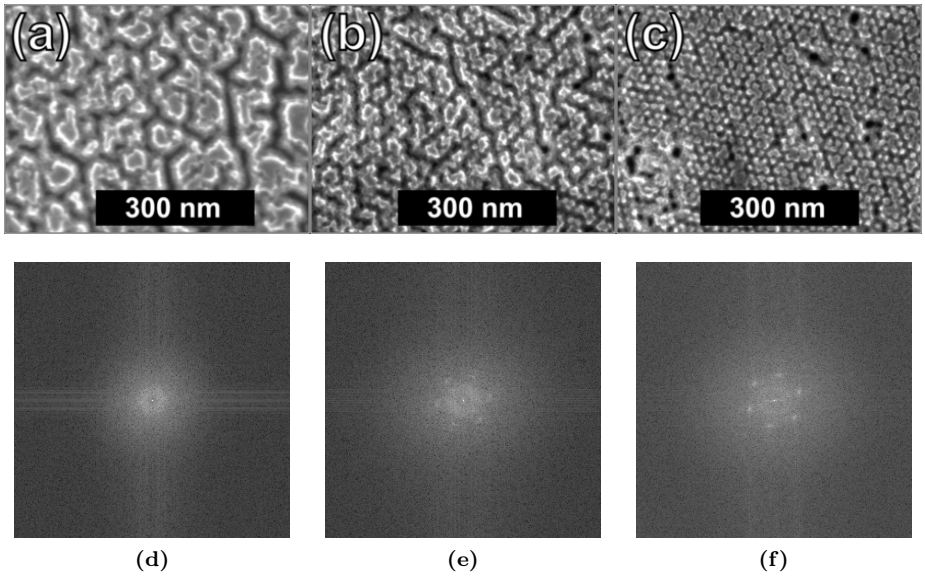

Figure 5.7: Three different sections of the same patterned substrate with particles removed, showing the effect of different doses. (a) No dose. (b) Low dose (4950 µC cm^{-2}). (c) High dose (39 600 µC cm^{-2}).

I also investigated the effect of etch recipe on the patterned morphology. Specifically, I varied the proportion of CHF_3 to CF_4, keeping the total flow rate constant at 18 sccm. This lets us vary the C:F ratio; a higher ratio promotes the formation of fluorocarbon polymer[90].Figure 5.9(a), etched with 4 sccm CHF_3 (14 sccm CF_4), exhibits many single-particle features, but also wide regions of joined particles. SEM images taken before particle removal revealed the presence of polymer on top of the nanoparticles, which covered the gaps and caused these joined features. Going to 18 sccm of pure CF_4 (Figure 5.9(b)), the patterned wafer looks heavily eroded, and the features are largely indistinct. This morphology suggests that *sidewall* polymer formation is key to the directionality of this etch, since this too-low C:F ratio exhibits lateral etching. A more balanced recipe (2 sccm CHF_3, 16 sccm CF_4) yielded the superior results shown in Figure 5.9(c): clear and flat-topped features, *without* large regions joined by excessive polymer masking. Figure 5.9(d) gives a caveat about the role of sample-to-sample variation: although the processing parameters (Table 5.1) are virtually identical to Figure 5.9(c), the cracking is worse and the features are more eroded. Here, the original monolayer initially had a "hazy" appearance in

ID (shorthand)	CF_4:CHF_3 ratio (sccm)	fluorine etch time (min)	O_2 etch time (s)	hazy monolayer?
`V-116-1` (P1)	14:4	40	120	no
`V-124-3` (P2)	18:0	20	120	yes
`V-127-3` (P3)	16:2	27	120	no
`V-138-4` (P4)	16:2	30	100	yes

Table 5.1: Comparison of etch parameters used for the samples shown in Figure 5.9.

SEM, suggesting excess surfactant present with the original monolayer. It thus appears that the quality of patterning can vary significantly, depending on the degree to which such excess surfactant is present.

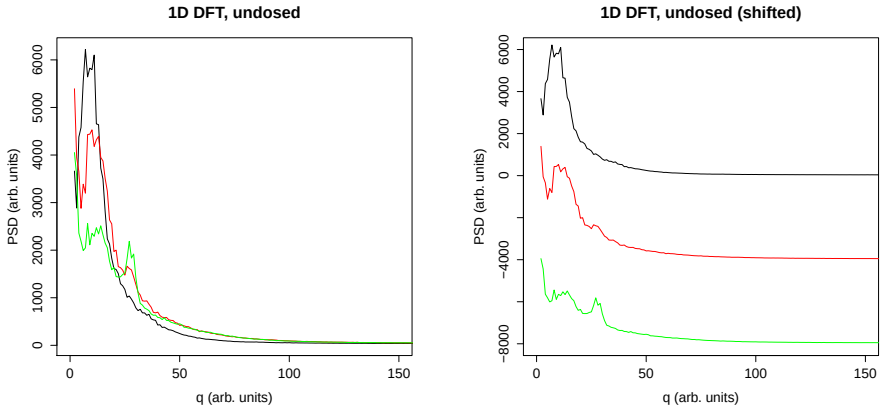

Figure 5.8: 1D DFT of ebeam-cured positive transfer, corresponding to the SEM images in Figure 5.7. (a) The plots are superposed. Note the gradual emergence of the second peak: nonexistent for the undosed curve (black), weak for the low dose (red), and strong for the high dose (green). (b) The plots are shifted to show their features.

Regarding the etch depth, similar considerations apply for the e-beam stabilized samples as for the SPL-stabilized samples. The main difference is that β will differ because Si is patterned instead of SiO_2, but since β is not known for either system, the analysis is unchanged. The shortest etch time is now 20 minutes instead of 50 minutes, but even here $\exp(-\beta At/w) = 0.05$, so the limiting depth has probably been obtained. The biggest difference from the SPL samples is that significant cracking still occurs for lower doses, such as Figure 5.7 (a) and (b). The effect of differing gap widths can be very significant, as shown in Figure 5.4(b). Hence, samples receiving low doses should exhibit much more depth variation than those receiving high doses. For the highest dose, where significant cracking was not observed, the same etch depth is expected as for the SPL-stabilized samples: between 1.9 nm and 28.2 nm, and probably much closer to the former.

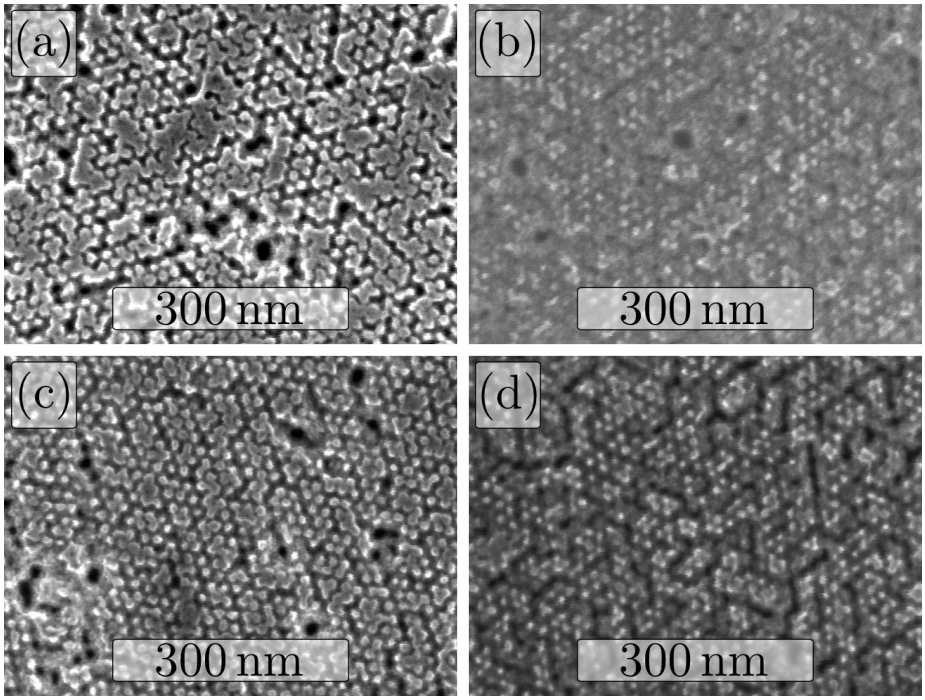

Figure 5.9: Four different pattern transfer results, all of which received the same dose ($39\,600\,\mu C\,cm^{-2}$). Etch parameters are given in Table 5.1. (a) Sample P1. (b) Sample P2. (c) Sample P3. (d) Sample P4.

5.3.3 Negative transfer using e-beam curing ("Antidots")

The successful use of nanoparticle arrays as etch masks is an important step forward. However, the above techniques also present serious challenges. The severe ARDE-induced slowdown in the ultra-small gap regime threatens economic viability, since each step in hard disk manufacture must be completed in about 6 s[1]. This may be addressable with dense, gentle plasmas from ICP-RIE. A more serious concern is that even in the best case (Figure 5.9(c)), many features are joined to one or more neighbors, whereas applications such as bit-patterned media require total separation.

These considerations suggest a "negative" approach to pattern transfer, where material is deposited between the nanoparticles, and serves as the *new* mask once the nanoparticles are removed. This describes the well-known "liftoff" process flow from lithography[103, p. 14]; thus, I propose that the nanoparticles can best be used as *liftoff masks for ultrafine lithography*, as shown schematically in Figure 5.10. For some applications the resulting structure might suffice, but others (including bit-patterned media) require that we end up with a pattern of dots, not antidots. The first step towards this goal would be to use the antidots as an etch mask to make pits, thereby enhancing the height contrast. From there, two routes might lead to dots: the pits could be filled directly with

[1]This figure comes from personal conversations with Dr. James A. Bain. His source was a worker on the production line at a major hard drive manufacturer.

the desired materials, or the wafer could be used as a stamp for nanoimprint lithography.

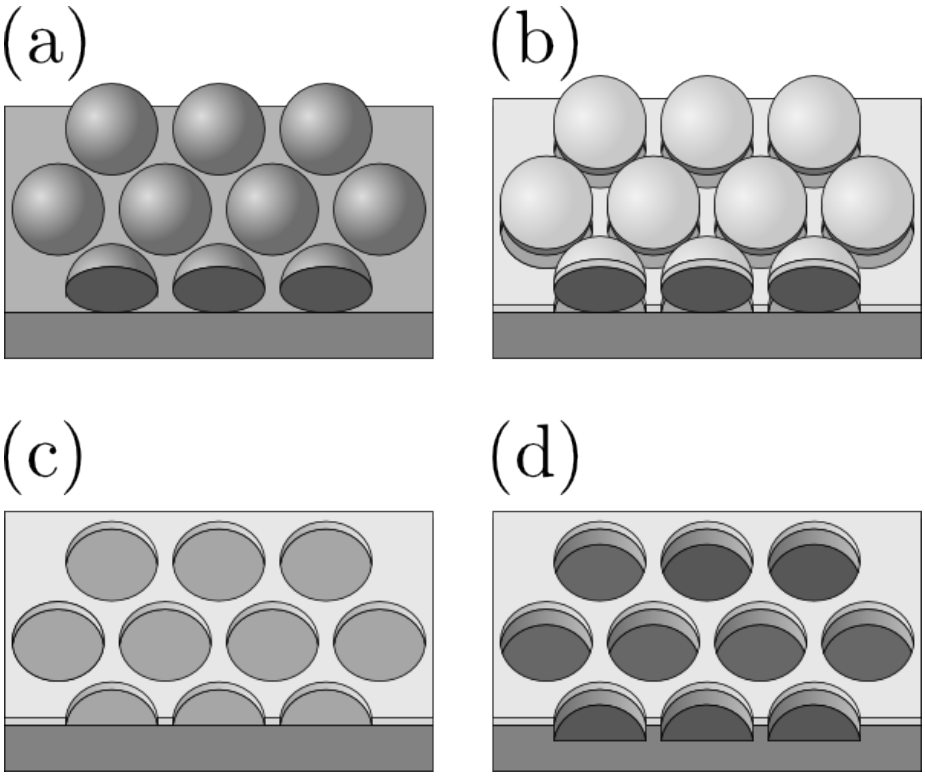

Figure 5.10: 3D schematic of the antidot process applied to nanoparticles. (a) A tilted view of a set of nanoparticles on a substrate. (The frontmost row is shown in cutaway for clarity.) (b) The antidot material is deposited between the particles onto the substrate, and also on top of the particles themselves. (c) The particles are removed, revealing the "antidot" pattern on the underlying substrate. (d) Highly selective RIE deepens the holes into pits without appreciably eroding the mask, thereby enhancing the height contrast.

5.3.4 Approach

For the antidot process to work, it is critical that the deposited material reach between the interstitial gaps to the substrate. This requires a high degree of directionality, which implies a mean free path much larger than the chamber dimensions. Sputtering takes place at pressures of at least 10^{-3} Torr; assuming room temperature and aluminum atoms of 250 pm diameter, this corresponds to a maximum mean free path of 0.11 m. E-beam evaporation, on the other hand, takes place at pressures below 10^{-6} Torr, which yields a mean free path of 110 m. The directionality is therefore expected to be excellent for evaporation and comparatively poor for sputtering. This is consistent with our earlier observations that sputtered films are contiguous and conformally coat the tops of nanoparticles.

Nanoparticles are removed after aluminum deposition by the same technique we previously used: rinsing in H_3PO_4. This implies the deposited film must not

be contiguous, or else the underlying surfactant would be inaccessible to the H_3PO_4. Figure 5.11 illustrates this concept. If the thickness t exceeds the distance from the substrate to the middle of the nanoparticles, the top of the interstitial film will reach the bottom of the film covering each particle, and they will join. This imposes a strict upper limit on t, i.e. $t < t_{\max} \equiv s + \frac{1}{2}D$, where s is the distance from the substrate through the cured surfactant to the bottom of the nanoparticle, and D is the nanoparticle diameter.

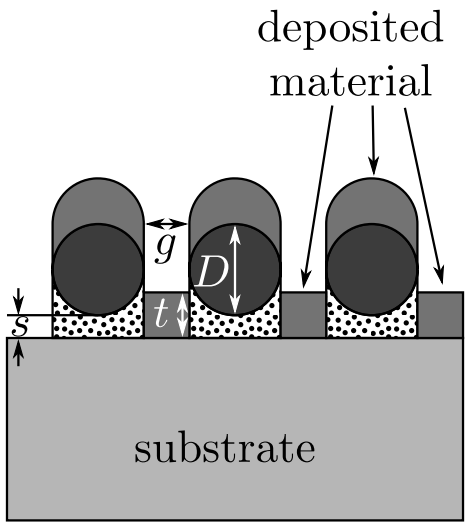

Figure 5.11: A figure illustrating the geometry of an idealized antidot array. Shown are the nanoparticle diameter D, the interparticle gap g, the distance s between the bottom of a nanoparticle and the substrate, and the thickness t of antidot deposition.

In practice, a number of factors will limit the deposition thickness even further. First, although a good first estimate of s would be $\frac{1}{2}g$ (with g the gap distance), in reality s is probably smaller due to e-beam compaction of the surfactant under the particles[94]. Also, it is likely that some slight residual carbon remains between the nanoparticles, where it may aid in film adhesion for the deposited antidot material; this would further reduce the height contrast. Furthermore, liftoff may become difficult before the nanoparticles are *completely* sealed, because the surfactant will only be accessible through a very small gap. Finally, the angle of incidence is not perfectly vertical, and slight sidewall deposition is expected, which would gradually narrow and eventually close up the gaps[140].

Successful magnetic recording demonstrations have typically had more height contrast, relative to the height of their magnetic multilayers, than bare antidots can provide. For instance, Hu et al. deposited 10 nm multilayers on 24 nm height contrast[67], while Hellwig et al. deposited 15 nm films over height contrasts of 28 nm and 35 nm.[61] The key is to prevent the upper material from being exchange-coupled to the material in the trenches, because this coupling would prevent features from being written independently of one another. I

have fabricated magnetic multilayers with perpendicular anisotropy having a total thickness of 4.6 nm, so it is possible that less than 24 nm height contrast would be required to show magnetic recording. Nevertheless, achieving high height contrast with as-made antidots requires using very large particles, which negates the key advantages of surfactant-coated nanoparticle-based methods.

Using the antidot lattice as an etch mask could help enhance the height contrast. RIE would be ideal because chemical selectivity is crucial to enhance rather than erode the height contrast. Al makes a good choice of antidot material, because its natural oxide can be an excellent mask for fluorine-based etching of Si or SiO_2[54]. Another attractive choice is Cr, which has the additional advantage that it is already a standard deposition material in the data storage industry. Once pits of sufficient height contrast have been made, they may either be filled with the desired functional (e.g. magnetic) materials, or used as a master stamp for nanoimprint lithography.

One great attraction of this technique is the materials-general nature it promises. Any substance which can be deposited by evaporation should be able to reach between the nanoparticles, and form an antidot array. Moreover, the method for nanoparticle removal requires only that its *surfactant* be vulnerable to acid. Since most surfactants are, this method promises to work with nanoparticles having a wide variety of chemical identities.

In reality, there are several potential pitfalls in materials choice. When choosing deposited materials, the extremely small gaps between nanoparticles give a strong preference for single-element materials: molecules would take up a wider fraction of the gap than would atoms, and would have a correspondingly harder time reaching the bottom. Furthermore, the deposited material must be resistant to whatever wet chemical etch removes the nanoparticles, usually some kind of acid. Finally, there is one constraint on the chemical identity of the nanoparticles: they must not react with this acid to form a precipitate which redeposits on the substrate. Despite these caveats, the antidot technique should admit significant freedom in choosing materials. Here, I discuss attempts with Al and Cr deposition, using templates based on Fe_3O_4 and MnO nanoparticles.

5.3.5 Experimental Methods

Two different nanoparticle-substrate combinations will be discussed here. The Fe_3O_4 nanoparticle arrays were prepared as described in Section 4.5, up to the removal of the surfactant by the PlasmaTherm. Additionally, I used the MnO nanoparticles described in Section 2.4.1 as a test system exhibiting slightly larger gaps. Since the MnO particles were used to explore pit etching, I formed their monolayers on custom wafers with a 40 nm layer of thermal SiO_2, which is preferred in industry for directional patterning. Thermal SiO_2 is grown by oxidizing Si wafers at high temperatures in moist conditions, and results in a uniform oxide film with good control over the depth.[103, pp. 92–96] This gives

enough thickness to demonstrate industrially relevant height contrast, but not so much thickness that imaging quality is degraded.

Immediately after surfactant removal, the sample was affixed to a carrier wafer by a thermally conductive paste, and loaded into the evaporator. The Al source was contained in a graphite crucible, whose added series thermal resistance (compared to bare Al) enabled steadier deposition as a function of e-beam current. The Al wets the crucible, and can easily crack it due to differential thermal expansion, but we avoided such damage by ramping the current very slowly. Only 40 mA was required with the crucible, compared to 220 mA without (which strains the capabilities of the instrument). These conditions led to deposition rates of $0.2\,\text{Å}\,\text{s}^{-1}$ to $0.4\,\text{Å}\,\text{s}^{-1}$, as measured *in situ* by a crystal rate monitor. Such slow rates enabled reproducible Å-level control of the nominal film thickness.

After deposition, the sample was removed from the carrier, and laid face-down in a concave wafer holder, while residual paste was cleaned from the back by swabbing with isopropyl alcohol. The Al was exposed to atmosphere for several hours and allowed to oxidize. Liftoff of the nanoparticles was performed with H_3PO_4 as previously described, with stirring times of 10 min for Fe_3O_4 and 5 min for MnO.

For the MnO-templated antidots, I subsequently exposed them to a CF_4 plasma in the RIE. The same recipe was used as for the structural pinning layer samples: 15 W, 18 sccm, 5 mTorr. The antidots were screened both before and after etching in the SEM.

To gauge the depth of the etch required cross-sections to be taken, and I did this in two ways. First, the sample was cleaved along the scribe mark by holding with tweezers on either side, and gently pulling apart, thereby cleanly slicing the patterned region in two. A side-on view of either half yielded cross-sectional SEM. Alternatively, cross-sectional TEM samples were prepared by the more involved FIB liftout process described in detail later.

5.3.6 Results

$13.7\,\text{nm} \pm 0.9\,\text{nm}$-diameter Fe_3O_4 nanoparticle templates

Broadly speaking, the antidots retain the pattern of the nanoparticles more faithfully and consistently than any other technique. I have investigated the effects of total dose, O_2 etch time, and accelerating voltage. Experiments with Cr-based antidots, together with surprising observations on Al-based antidots, clarify important contributing mechanisms in reliable antidot formation. Finally, unsuccessful attempts to etch pits motivate our move to the MnO nanoparticle system, which led to greater success in this area.

Antidots are shown for three representative doses in Figure 5.12, analogous to Figure 5.7. Remarkably, even the low-dosed region ($4950\,\mu\text{C}\,\text{cm}^{-2}$, Figure 5.12(b)) exhibits largely faithful pattern transfer. In particular, a clear majority

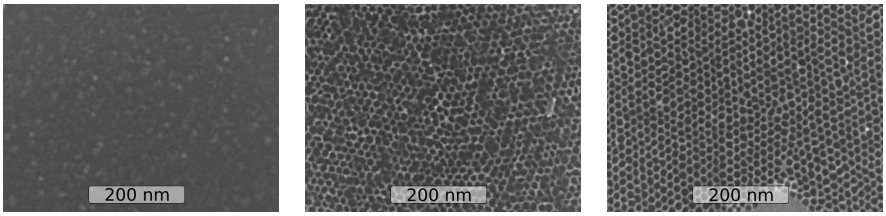

Figure 5.12: The effect of ebeam dose on the quality of antidots. (a) Undosed region. (b) Low dose ($3500\,\mu C\,cm^{-2}$). (c) High dose ($39\,600\,\mu C\,cm^{-2}$).

of features are completely separated from their neighbors, while the analogous region in Figure 5.7(b) shows much more severe cracking. These antidots still feature some missing edges, as well as variations in contrast, both of which suggest varying degrees of partial cracking. However, these defects are absent in a more highly-dosed region ($39\,600\,\mu C\,cm^{-2}$, Figure 5.12(c)), which displays even more faithful pattern transfer. The absence of transfer defects compares very favorably with similarly-dosed pillars in Figure 5.7(c).

The undosed region in Figure 5.12(a) shows surprisingly little structure. One might instead expect to find Al in the pattern of the cracks, but this is only rarely found. It is unlikely that this is due to H_3PO_4 etching AlO_x, since the antidot lattice remains intact. Rather, this sparseness suggests the AlO_x film adhesion is poorer in undosed regions.

There are two main differences which could explain this. Since the cracks are open wider, ARDE is less of an obstacle to surfactant removal, and there should be less residual carbon in undosed cracks than the dosed interstices. Alternatively, the carbon in the cracks should be chemically different than the cured carbon, having either a faster or slower etch rate. It is thus also conceivable that *more* carbon remains in the cracks, and the AlO_x film is rinsed away more easily due to this thick underlying carbon "blanket". I find the former option to be more plausible, since I observed that identical Al films deposited on bare Si wafers were also washed away. Hence, it seems that small amounts of residual cured carbon help promote adhesion of the Al film.

Next, I consider the effect of O_2 etch time on the quality and structure of antidots. The sample in Figure 5.14(a) was not exposed to O_2, and serves as a control. As expected, without the surfactant etched, the Al contiguously coated the nanoparticles, and they were not removed. A short exposure (75 s, Figure 5.14(b)) makes a dramatic difference: enough surfactant has been removed for antidots to appear. Part (c) shows as-made antidots for a 105 s exposure. Its appearance is different from part (b), indicating that the extra surfactant etching time changed the structure. This suggests that 75 s etch time removes the surfactant only partially.

To test the effects of this partial removal on the composition of the antidot lattice, I exposed these samples to a 7 W O_2 plasma for 5 minutes each.

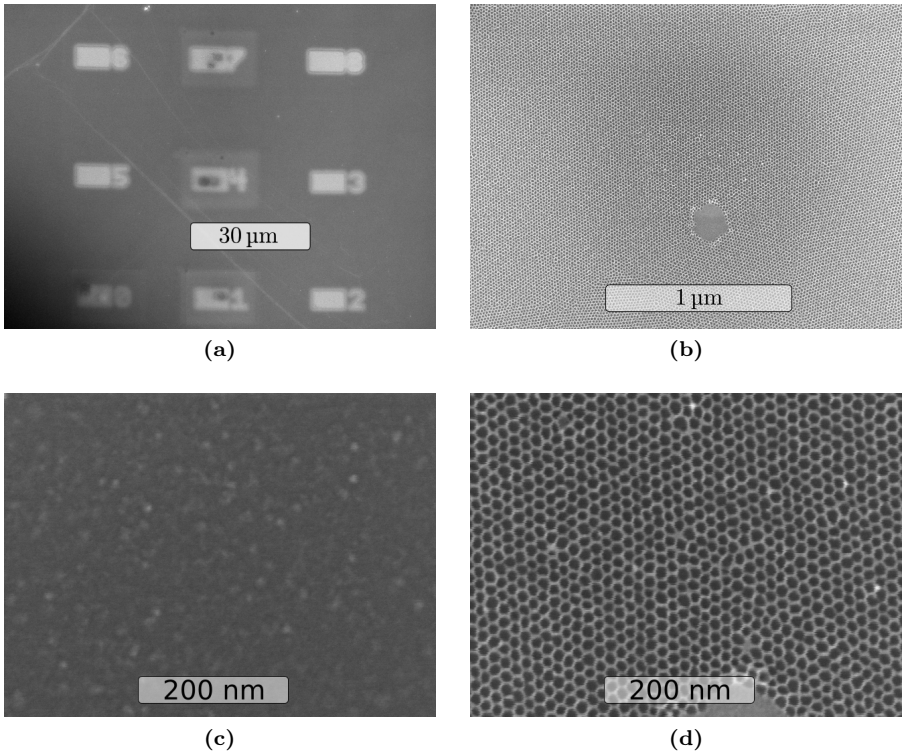

Figure 5.13: Al-based antidots templated from Fe_3O_4 nanoparticles. (a) Zoomed-out view of dosed regions; each block is 5 µm high and 14 µm wide. (b) Intermediate zoom, so that large numbers of individual antidots are visible. No defects can be found in the transfer, except for the bare region which was already present in the original monolayer. (c) The region receiving the lowest dose ($D_0 = 3500\,\mu C\,cm^{-2}$). The pattern is tattered, but still very discernible, especially compared to a similar dot pattern shown in Figure 5.7(b). (d) A high-dosed region (the same as shown in part (b)). Now, the ordering is excellent.

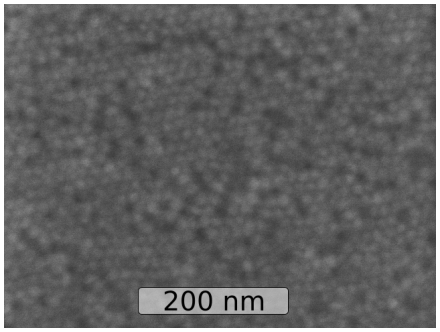

(a)

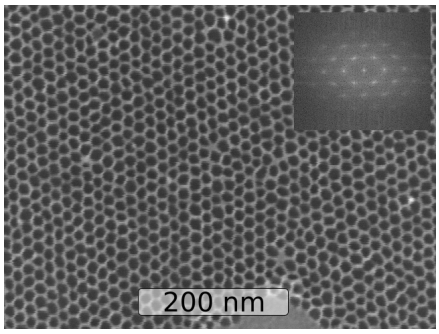

(b)

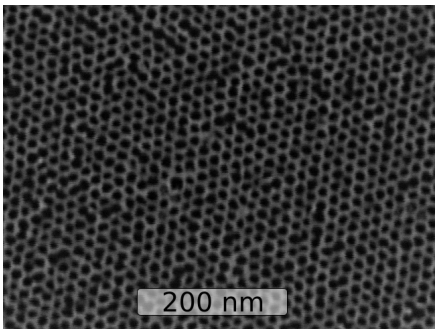

(c)

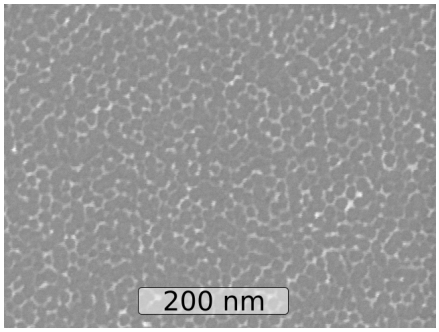

(d)

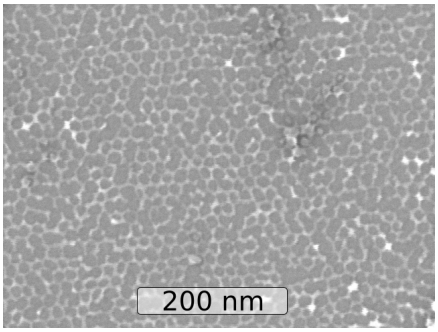

(e)

Figure 5.14: Effect of O_2 surfactant etching time on antidot quality and composition. (a) Control sample, where surfactant was not etched. The deposited Al formed a contiguous layer and the particles were not removed. (b) 75 s surfactant etching time. Even this short etching time removed enough surfactant to form an antidot lattice. The inset shows the DFT of this well-ordered antidot array, and the large number of visible peaks confirms the high fidelity of pattern transfer. (c) 105 s surfactant etching time. In (d) and (e), these lattices have been exposed to a 5 minute O_2 plasma treatment, which etches carbon but not aluminum. (d) Same sample as in (b); the erosion suggests that carbon was a significant constituent of the antidot lattice. (e) Same sample as in (c); less erosion is observed in the sample with more surfactant removed.

Oxygen yields gas-phase reaction products with carbon, but merely oxidizes the aluminum. Therefore, the degree to which the lattice is eroded by the plasma indicates the amount of carbon in the lattice. Part (d) shows the wafer which had undergone less initial surfactant removal. The lattice is significantly eroded compared to the same sample shown in part (b). Part (e) shows the same sample as part (c), which had less surfactant remaining, and the erosion is also less. These results suggest that there exists an optimum surfactant removal time: long enough to avoid incorporating too much carbon, but not so long as to attack the *underlying* carbon and cause cracking.

I also examined the effect of acceleration voltage, checking 10 kV in addition to the 30 kV I used for direct transfer. At the lower acceleration voltage, the electrons deposit more of their energy in the first few layers, whereas higher-voltage electrons are more penetrating, and interact mainly with the bulk substrate. Figure 5.15 shows these results. Comparing the same nominal dosage ($40\,000\,\mu C\,cm^{-2}$) at (a) 30 kV and (c) 10 kV, the former exhibits lower contrast, variations in contrast, and joined features, while the latter is crisp, uniform, and a faithful transfer of the nanoparticles' pattern. Thus, 10 kV appears to produce superior nanostructure compared to 30 kV.

Zooming out shows additional differences, which help explain why the 10 kV gives better results. The most striking feature is a rounded rectangular border, outside the directly exposed region, which is present for 10 kV (d) but absent for 30 kV (b). Additionally, each pattern is surrounded by a weakly-exposed "halo," manifested by "wrinkles" in the original monolayer. This halo is very broad for 30 kV (b), but relatively narrow for 10 kV (d), where it extends only slightly past the rounded rectangular border.

Figure 5.16(a), taken at the lower right-hand corner of a patterned block, shows these regimes in greater detail, and Figure 5.16(c) shows an explanation of the various effects the e-beam has on carbon. The central idea is that the beam not only *alters* the carbon, but also *deposits* and *removes* carbon as well. Carbon deposition is based on the presence of carbonaceous contaminants, often due to backstreamed pump oil in the chamber. Any contaminants in the path of the beam can be cracked by it, and their carbon remnants deposited on the surface. This deposition is found to be self-limiting up to a carbon depth of a few nm[39], because direct e-beam exposure can also cause e-beam evaporation.

I have developed the following model to explain these observations. Lowering the acceleration voltage makes the beam less penetrating, and it interacts more strongly with higher features, such as the surfactant on the nanoparticles or the carbon contaminants in the chamber. Cracked contaminants will not deposit directly below the beam, but within some radius R_d around it, based on the distribution of their initial velocities. The rounded rectangular border is the region which falls within R_d of the primary beam (thus receiving deposited carbon), but never receives the direct beam exposure which would limit the carbon deposition. This border thus represents a thicker carbon blanket, which

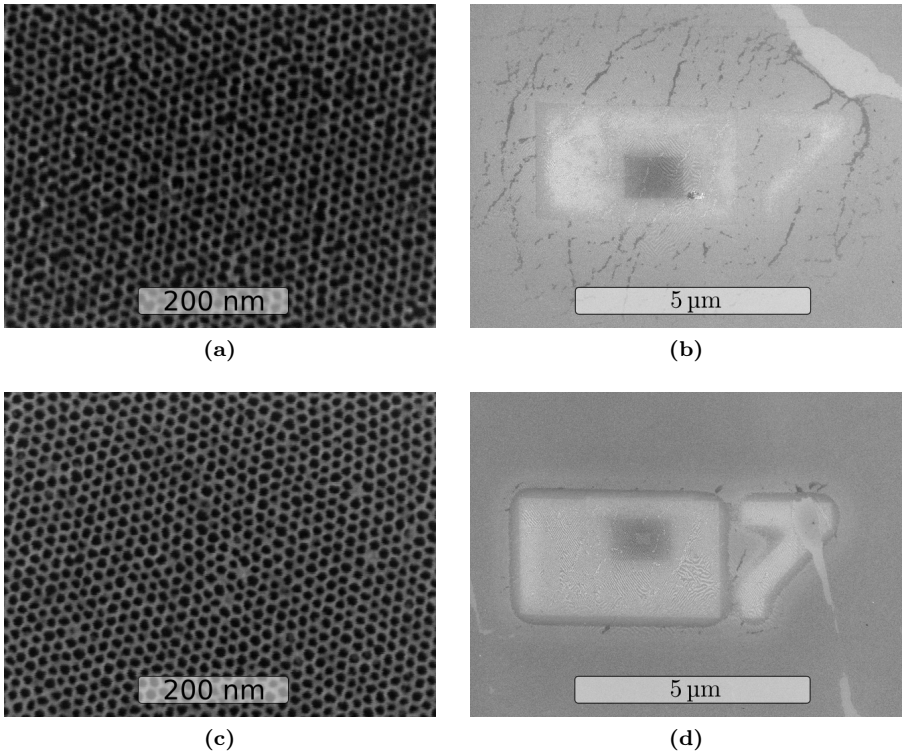

Figure 5.15: Illustration of the effect of accelerating voltage on quality of antidots. (a) and (b) were dosed at 30 kV, and (c) and (d) were given the same dose (39 600 µC cm^{-2}) at 10 kV. The acceleration voltage affects both the qualtiy of antidots and the peripheral features of the patterned region.

covers the particles and prevents their removal. Figure 5.16(b) shows this border after exposure to a 5 min O_2 plasma, revealing that the particles indeed remain.

Meanwhile, electrons which penetrate the substrate may "backscatter" and return out at some angle to the incident beam, causing a secondary exposure of the surfactant in the film.The characteristic radius R_b over which this happens increases with the accelerating voltage, because more highly accelerated electrons penetrate more deeply before turning around. For a given angle, the lateral displacement on re-encountering the substrate is directly proportional to this depth.

10 kV electrons deposit more energy in the surfactant, driving the reaction further towards completion for the same nominal dose. The reason no rounded border is seen for 30 kV is because much less carbon is deposited, since the electrons penetrate more deeply before undergoing significant interactions.

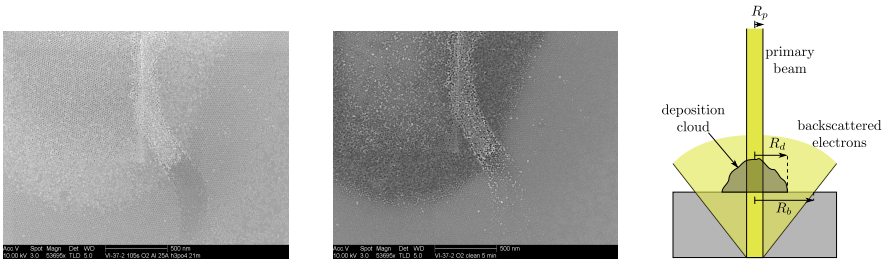

Figure 5.16: An illustration of the different effects the electron beam has on carbon. Four different regions are shown. (a) and (b) show the same region, before and after an O_2 plasma etch which revealed the structure of these regions. (c) gives a schematic of how these regions arise. In (a) and (b): The directly exposed region (top left) received deposited carbon, but the thickness was limited by the direct exposure of the electron beam. Regions not directly exposed, but within R_d of the beam, also received carbon deposition, but the deposition is thicker because there was no beam exposure to limit it. These regions manifest as a rounded "halo" around the rectangle, and the O_2 plasma etch reveals that this thicker carbon layer caused nanoparticles to be retained (and they are seen in (b)). The region immediately outside this halo ($R_d < R < R_b$) was too far from the beam to receive carbon deposition, but nevertheless was cured by *backscattered* electrons, and therefore exhibits antidots. Regions where $R > R_b$ are completely undosed.

Initial attempts to etch pits using the Fe_3O_4-templated antidots as masks proved challenging. These antidots are roughly 2 nm high and 2 nm wide, and could be eroded completely even by only very little physical sputtering. What is needed is a way to keep the physical bombardment low enough to preserve the mask, while having the plasma density high enough for an appreciable etch rate.

Inductively coupled plasma (ICP) allows us to decouple these parameters: the plasma density is governed by the "coil" power, and the physical bombardment is controlled by the "platen" power. By keeping the latter suitably low, we were able to achieve lower substrate bias voltages than were possible in the PlasmaTherm. A close-up picture of the results is shown in Figure 5.17(a). The antidots were not eroded away, but they appear torn and tattered. This is

likely due to undercutting: the low substrate bias voltage means the etch will be more isotropic, and lateral etching of the walls of pits can undermine the antidot mask.

This is a key result for antidot-based pattern transfer: directionality is critical, because the walls are too thin to tolerate any significant undercutting. The ICP-based STS instrument is not designed for intrinsically directional etches. Rather, directionality is typically achieved by etch/deposition cycles, which coat the sidewalls in a protective polymer later, known as the "Bosch" process. Our pits are also too thin to leave much room for sidewall coating, and the typical etch depth per cycle is orders of magnitude beyond the desired etch depth. For this reason, the STS is unlikely to lead to successful etching of pits through antidots.

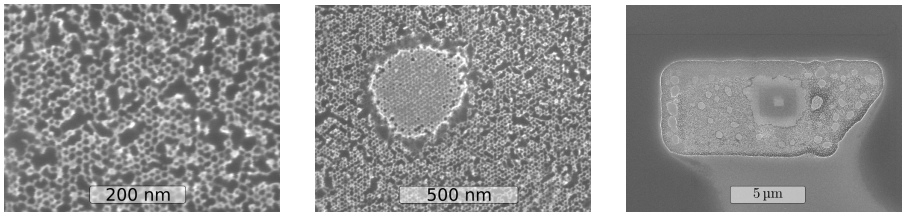

(a) **(b)** **(c)**

Figure 5.17: Results of ICP etching of Fe_3O_4-based antidots. (a) Closeup of patterned region, showing some transfer worked but some transfer led to a tattered antidot grid. (b) Intermediate-zoom view showing a region where pattern transfer did not occur. (c) Zoomed-out view showing the distribution of unpatterned regions. The large rectangular region in the center corresponds to a previously-imaged region, but other unpatterned regions seem to be distributed randomly.

Figure 5.17(b) shows an interesting phenomenon: numerous plateaus were not etched by the STS at all. The size and distribution of these plateaus is shown in Figure 5.17(c). The large central plateau, which correlates with deposition rectangles that indicate previously-imaged regions, is easily explained: the carbon deposited by the e-beam acted as a mask, and protected the material in the holes from being etched. This suggests that the other plateaus also exhibited some kind of thin chemical mask. I tentatively identify this mask with a precipitate from the reaction between the H_3PO_4 and the AlO_x antidots.

Additional evidence for this precipitate is furnished by a differential coloring of the holes sometimes seen in antidot masks, such as Figure 5.18. The starkness of the difference between dark and light holes suggests an origin other than topography. Some chemically distinct substance must be in certain holes, and not in others. This substance is not primarily C-based, because it withstood numerous C etches. Only a re-exposure to H_3PO_4 caused the holes to become homogeneous, which is why I identify the residue as stemming from H_3PO_4.

In fact, the idea that H_3PO_4 undergoes a self-limiting reaction with AlO_x may be key to the formation of Al antidots, with implications for the extension of this technique to other antidot materials. I attempted to form Cr antidots via

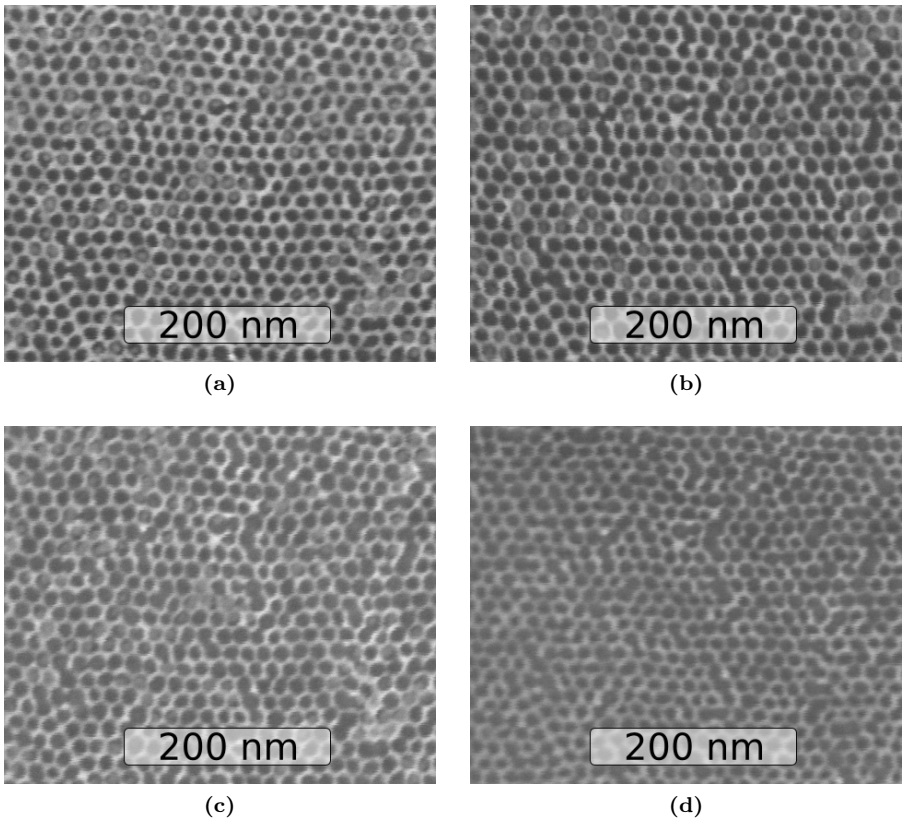

Figure 5.18: Evolution of antidot residue over time. (a) Newly revealed antidots after 24 min H_3PO_4 stirring. (b) Same sample after 85 min UV exposure; the pattern of light/dark holes is unchanged. (c) After additional exposure to 100 W O_2 plasma in barrel etcher for 3.0 min. (d) After additional stirring in H_3PO_4 for 24 min.

an identical procedure, but with 1.9 nm Cr substituted for 2.1 nm Al. Results are shown in Figure 5.19. In most places, the liftoff failed, and the contours of the nanoparticles were still visible under the Cr. Where liftoff was successful, the antidot lattice exhibited noticeable contrast variation, suggesting the film on the nanoparticles had been joined to the antidot lattice, then ripped off. Unlike with Al-based antidots, the holes exhibited no contrast variation. This suggests a lack of residue, and is consistent with the fact that CrO_x resists chemical attack from H_3PO_4.

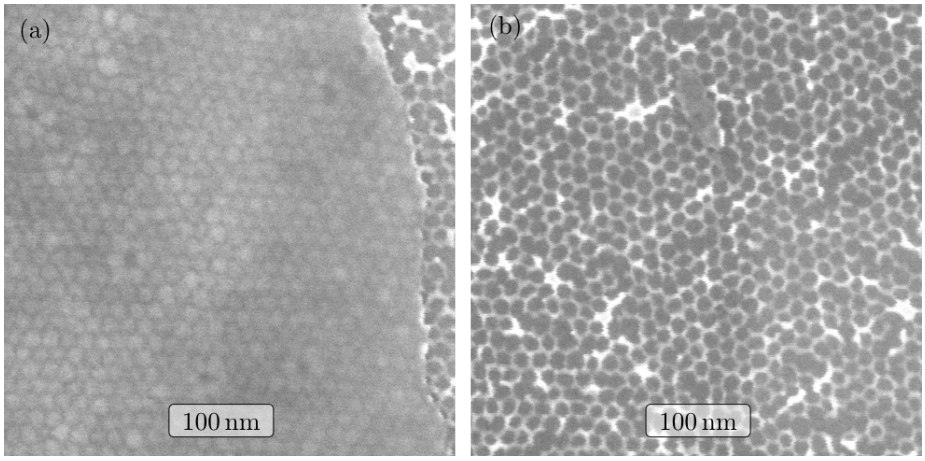

Figure 5.19: Cr-based antidot lattice results, templated from 13.7 nm Fe_3O_4 nanoparticles. (a) A region where liftoff did not work. The nanoparticles are faintly visible under the Cr coating. (b) A successfully patterned region. Variation in contrast may suggest the antidot lattice had been partially contiguous with the Cr on top of the particles, and that the liftoff process tore the Cr, causing height variation.

From these observations, I infer that *even in very thin films*, the lower deposited layers (i.e. the antidot lattice) may be slightly joined to the upper deposited layers. A very slight chemical attack, such as H_3PO_4 on Al, may then greatly enhance the liftoff process. For materials such as Cr, this role could be played by a very slight addition of H_2SO_4, which does etch CrO_x. When the reaction products are solid (as with AlO_x and H_3PO_4), the precipitate may give more leeway by preventing further reaction; however, it can also interfere with subsequent attempts to etch pits.

30 nm-diameter MnO nanoparticle templates

Pattern transfer with Fe_3O_4-templated antidots is too ambitious for a first attempt: the residue necessitates a more aggressive initial etch, but the 2 nm-thick lattice is vulnerable to the slightest amount of erosion. To circumvent these obstacles, we turned to a system of MnO nanoparticles. Their wider gaps (4 nm, vs. 2 nm) lead to wider antidot lattices, and their higher diameter (30 nm vs. 14 nm) allow thicker deposition. Thus, antidot lattices templated from these MnO particles should be more robust, and more forgiving to different etch pa-

rameters. This also directly demonstrates the ability to use multiple types of nanoparticles to template our antidots. The drawback is that I was able to achieve only short-range ordering, unlike the very long-range ordering observed with Fe_3O_4 nanoparticles.

Figure 5.20(a) shows a typical MnO monolayer we used for patterning. The lack of long-range order is apparent, but ordered regions 4 or more nanoparticles across are plentiful, and these are all that is needed to demonstrate the feasibility of etching pits from antidots. Figure 5.20(b) shows the pattern of dosing, at the tip of an existing scratch. Note that we expose a much larger area than with the Fe_3O_4 antidots. This will enable characterization by macroscopic techniques, such as XSEM, XTEM, or MOKE if we are successful in creating magnetic patterns. After etching surfactant, depositing 5.0 nm Al, and removing the nanoparticles by stirring for 5 min in H_3PO_4, Figure 5.20(c) shows that the antidot lattice is crisp and well-defined, and faithfully reproduces the ordering in Figure 5.20(a). By viewing at a tilt angle of 30°, we gain some idea of the 3D structure of this lattice. It clearly stands out from the underlying wafer, and the sidewalls appear straight and steep.

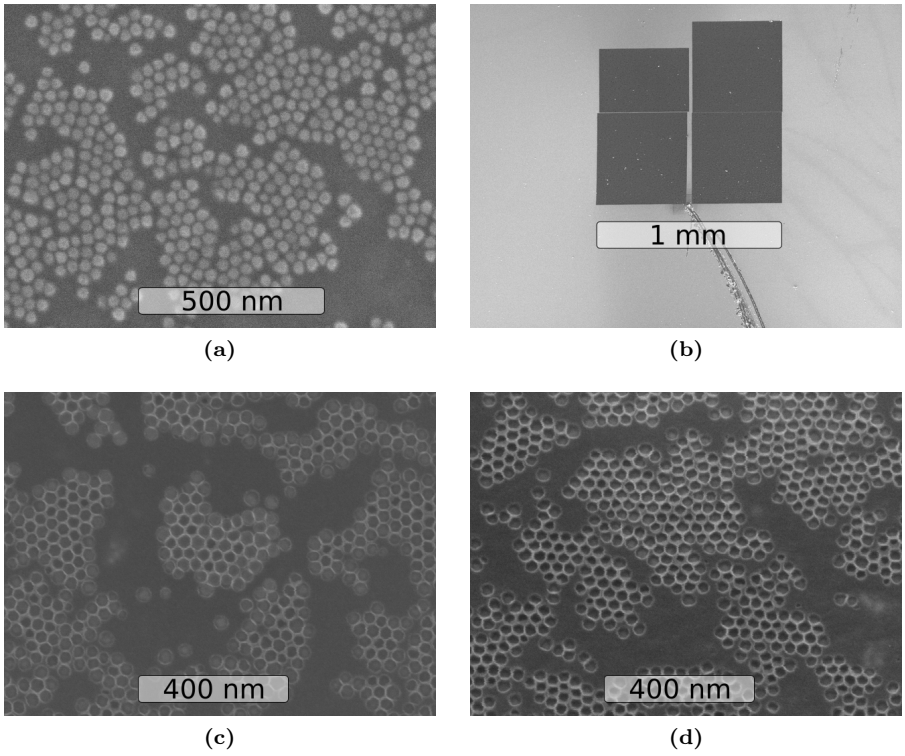

Figure 5.20: Typical antidots templated from MnO monolayers. (a) The original monolayer. Unfortunately, its order is short range, but this pattern is faithfully transferred. (b) The diamond scribe (bottom) leading up to the cured region: a set of four squares, each 250 µm on a side. The top left square is incomplete because the exposure was aborted early; the area covered corresponds to what could be done with a single overnight run. (c) Plan view of the antidots after nanoparticle liftoff. Note that they exhibit the same pattern as the original monolayer shown in (a). (d) The antidots viewed at a 30° tilt angle, in order to give a sense of their 3-dimensional structure.

The residue often seen with Fe_3O_4-templated antidots is absent with the MnO-templated antidots. Moreover, the contrast observed in SEM is always consistent with the topography expected from the nominal deposition parameters, unlike the stark black-and-white contrast exhibited by Fe_3O_4 antidots for short O_2 etching times. The differences in surfactant could contribute to these differences: MnO nanoparticles feature only oleic acid, while Fe_3O_4 nanoparticles feature oleyl amine also. A more likely explanation is the different geometry. Surfactant should be easily removed from the generous 4 nm MnO gaps, which explains the realistically small contrast levels. Also, the clearance between the nominal top of the lattice and the minimum midpoint of the particles is greater for MnO, leaving the undersides more accessible to the H_3PO_4.

Another observation is suggestive as to the mechanisms of film adhesion for the antidot fence. Undosed regions do not show AlO_x in crack patterns as expected; rather, they are mostly bare. Additionally, Al films deposited onto bare Si control samples under identical conditions are washed completely off by H_3PO_4. These observations, together with the fact that the antidot lattices do *not* wash off, suggest that residual carbon aids in the adhesion of the Al-based film.

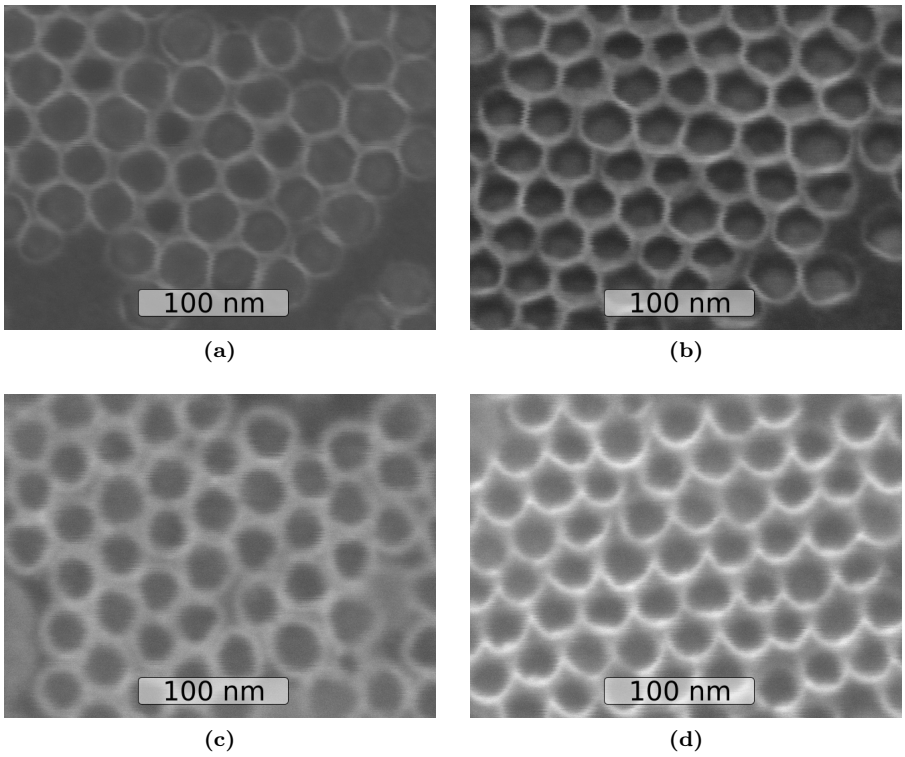

Figure 5.21: Results of RIE through antidot mask. (a) Plan view of antidot mask. (b) Antidot mask viewed at 30° tilt. (c) Plan view of patterned wafer. (d) Patterned wafer viewed at 30° tilt.

We etched pits using the same recipe as for the oxide pillars. However, our

pit etch times (2 min to 3 min) were shorter than for pillars (20 min to 40 min) by at least an order of magnitude. This shortened time shows that the antidot process mitigates the aspect ratio-dependent effects which plague the direct transfer. Basically, this improvement comes about because the ultra-small gaps have exchanged roles with the relatively wide features. To make this more quantitative, consider Equation 5.7 in the limit $\beta At/w \ll 1$ — the opposite of the limit previously considered:

$$\lim_{t\to 0} \frac{D(t)}{t} = A\left[1 - \frac{L_0\beta}{w}\right]. \tag{5.10}$$

This shows that, for short etch times, the etch rate is constant, but smaller than the bulk rate. For the MnO antidots, $w = 30\,\text{nm}$, $L_0 = 5\,\text{nm}$, and I take $\beta = \beta_{\text{KA}}$ to obtain $L_0\beta/w = 0.0095$. Therefore, *these antidots should exhibit near-bulk etching rates*, and that is why the etching times are so much shorter than for dots.

Note that spect ratio-dependent effects are also mitigated during the deposition phase, when small gaps come into play. Now, the material flow is one-way, eliminating "traffic jams" between etchants and their products. Moreover, the flow is more directional, since it takes place at a pressure which is orders of magnitude lower. This directionality aids the particles in reaching the bottom.

The results of etching are shown in Figure 5.21: (a) and (b) show the plan view and 30° tilt before etching, and (c) and (d) show the same views after etching. Comparing plan view images (a) and (c) reveals that the lattice in the etched sample appears more diffuse, probably as a result of mask erosion. At 30° tilt, distinct sidewalls are readily apparent before etching (b), but not after (d). This suggests that the sidewalls of the patterned film have become significantly sloped. Tuning the etch chemistry and parameters may help to achieve more vertical sidewalls.

Cross-sectional SEM reveals that the pattern transfer was successful in enhancing the height contrast. Yoosuf Picard performed focused ion beam liftout of the patterned region, which is shown in Figure 5.22(a). First, a site was selected and coated with 3 µm W by an e-beam directed deposition process, using a 2.1 nA current at 10 kV. The W protects the underlying structure from damage by the ion beam. After the pattern was sufficiently protected, a further 3 µm of protective W was deposited by a 0.3 nA ion beam at 30 kV. Finally, the exposed cross-section was cleaned by ion polishing with 30 kV Ga^+ ions: first with 0.3 nA, and then with 0.1 nA.

Figure 5.22(b) shows the results. The height contrast is almost triple the original height contrast, which was the nominal deposition depth of 5.0 nm. No significant depth variations are seen; these would be expected for small-gap etching, but not for these antidots. The lateral period of repetition, averaged over 7 nanoparticles, gives a pitch of 33 nm. This is consistent with the 34 nm pitch measured from plan view, but slightly smaller, possibly because this region

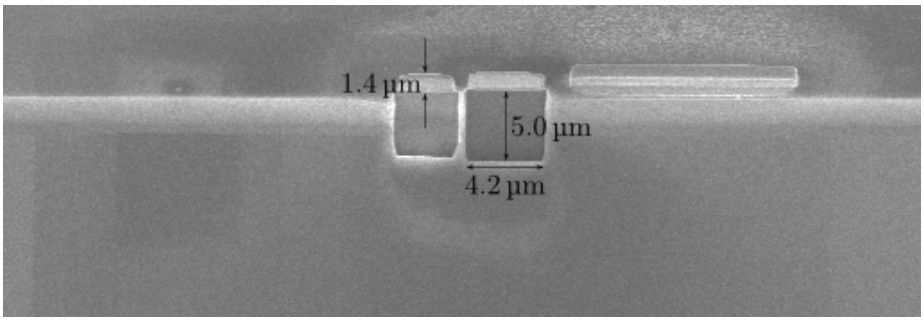

(a)

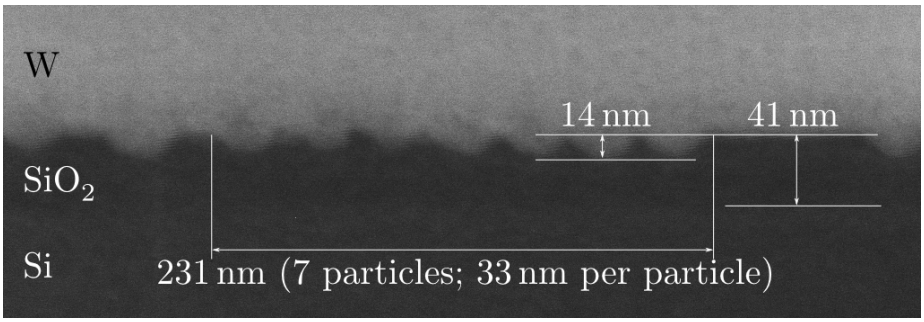

(b)

Figure 5.22: Cross-sectional SEM of antidot pattern transfer. The sample was tilted at 52°, and the measurements shown have corrected for this perspective. (a) Zoomed out view showing the side of the wafer, and two milled regions. The W caps are 1.4 µm tall, and the milled regions are 5.0 µm deep and roughly 4.2 µm wide. (b) Closeup showing microstructure. The SiO_2 layer is a nominally 40 nm thermally oxidized layer. Contrast in the image is mainly Z-contrast, since the surface has been milled flat. SiO_2 is barely distinguishable from the bulk Si below, but clearly distinguishable from the deposited W above. The lateral period of repetition is consistent with plan view measurements of the nanoparticle arrays.

of antidots forms a slight angle with the milling plane. Some sloping of the sidewalls can be seen. Further optimization of the etch parameters may be able to mitigate this sloping.

5.4 Summary

To summarize, I have demonstrated successful pattern transfer through arrays of surfactant-coated nanoparticles. My mask of Fe_3O_4 particles has $13.7\,\text{nm} \pm 0.9\,\text{nm}$ feature diameter and $15.7\,\text{nm} \pm 0.2\,\text{nm}$ pitch, making it slightly smaller than the smallest BCP mask for which pattern transfer has been shown. I demonstrated two forms of patterning with these particles: the creation of Si pillars, and the fabrication of AlO_x antidot lattices. MnO nanoparticle monolayers, with wider gaps, showed that antidot lattices can act as etch masks, giving additional height contrast which helps make industrial applications more plausible.

The antidot technique is the most promising technique for transferring patterns from nanoparticle arrays. It is materials-general with respect to the nanoparticles (Fe_3O_4 and MnO) and the lattice material (AlO_x and CrO_x), and is directly amenable to fabricate master stamps for the nanoimprint techniques that industry expects to use. In short, this project helps prepare nanoparticles to take the torch from block copolymers several years down the road, and carry data storage beyond 10 Tbpsi.

6 Conclusions and future work

6.1 Summary

6.1.1 Context and motivation

Two separate industries — semiconductors and magnetic data storage — have maintained an astonishing exponential growth in feature density for decades. Both industries are facing increasingly challenging obstacles to continuing that growth. Both have looked to a common solution: next-generation lithography.

In particular, the data storage industry has settled on a promising three-stage approach to make the masks for bit-patterned media. The first stage is to use a direct-write technique to form a sparse pattern. The advantage of direct-write is that the features can be placed exactly where desired, and the sparseness of this pattern makes the write times much shorter. The second stage involves filling in the pattern with self-assembling materials, using the directly-written pattern as a guide. Self-assembly can cover large areas in very short times, and the guide features provide vital control over the position and orientation of the self-assembled features. Finally, this pattern is transferred into a wafer suitable to be used as a stamp for nanoimprint lithography. The consensus in the industry is that the first-generation bit-patterned media will use electron beam lithography for the direct write stage, and block copolymers for the self-assembling material.

My work aims to increase the viability of surfactant-coated nanoparticles as an alternative self-assembling technology for the second step. Block copolymers are presently more mature, but they also appear to be less extendable to smaller sizes than a competing self-assembling technology: surfactant-coated nanoparticles. These nanoparticles have already demonstrated self-assembly at a pitch below 4 nm. However, guiding their self-assembly is a challenge. Moreover, the pattern of their ordered arrays had never been successfully transferred into an underlying wafer. This latter challenge has been the focus of my project.

6.1.2 Cracking

The fundamental obstacle preventing this transfer stems from the surfactant shell which coats the nanoparticles. Removal of that shell — a necessary step for patterning — causes the particles to aggregate to their neighbors, and destroys the pattern before it can be transferred. I have called this phenomenon

"cracking".

Quantitative cracking measurement

The study of cracking in nanoparticle arrays would be facilitated by a reliable measurement technique. Here, I have investigated and compared several such techniques. Most of them require that the particle positions be extracted by automated image analysis. The discrete Fourier transform (DFT), however, does not. Consequently, its great virtue is that it works even when position extraction fails. One drawback is that comparing DFT images is more qualitative than quantitative: in particular, it's unclear that inspection of the DFT makes it any easier to gauge the cracking than inspection of the raw images themselves.

All other techniques I considered required the particle positions to be extracted. Thus, there are two separate questions in assessing these measurements. First, what is the quality of the particle center extraction? And second, which techniques would be best if we had perfect knowledge of particle positions?

To extract the particle positions in this work, I used blob detection because of its speed and the lack of need for time-consuming manual image preprocessing. The chief drawbacks of blob detection are that it misses a few percent of particles, and that the measured particle centers sometimes differ from the true value based on the local environment of the particle. Binary threshholding would solve the latter problem. However, in the cracked case when particles are very close together, neighbors are likely to be spuriously joined (by noise pixels) in multiple places. Reliable particle extraction by binary threshholding thus involves painstaking manual pixel-level postprocessing, and is likely to be very time-consuming.

The prospects are good for particle detection that is both automatic and accurate using several scale-space[100, 99] algorithms in concert. Edge detection could outline particles well-separated from their neighbors, while ridge detection could mark the boundaries between touching particles. These data could be combined with the results of blob detection using Bayesian algorithms. While such a program would be advantageous from the end user's point of view, creating that program in the first place would require significant development and testing. At any rate, the "pure" effectiveness of the measures discussed below could be gauged by running a cracking *simulation*, where the particle centers are known to very high precision.

Assuming knowledge of the particle positions, I compared two classes of cracking measurement. The first involved generating functions (either 2D or 1D) with damped oscillations. In the 2D case, the DFT of the pure particle positions was considerably cleaner than the DFT of the raw image. This is because the particle form factor damped out the higher-order peaks in the latter, and these peaks are more important for distinguishing cases with very high degrees of ordering. The 1D average of the DFT brings out this improvement even more

clearly; compare Figure 4.7(e) (processed) to Figure 4.7(b) (raw image) and note the higher-order peaks in the former. Another important 1D measurement is the correlation function $g(r)$, where the uncracked case is distinguished by higher peak amplitudes and the fact that certain double-peaks are resolvable. The drawback to all these techniques is that there is no straightforward way to translate these graphs into a single number, which is not either *ad hoc* (e.g. "height of fourth peak in DFT") or reliant on intensive modeling (e.g. fitting an exponential decay envelope requires accounting for different peak heights).

To have a simple quantitative measurement for the cracking in an image, I proposed the "neighbor spacing variation". It is based on a feature essential to cracking: particles move closer to some neighbors, and farther from others, such that the distribution of nearest-neighbor distances for each particle acquires more variation. I showed that the NSV does distinguish between the cracked and uncracked cases, and does not distinguish among uncracked images which can also not be distinguished visually. It also does not capture effects that extend beyond nearest neighbors during crack formation. It is useful for differentiating relatively well-ordered samples, but not those containing a lot of other structural defects. However, what is missing is a test of its ability to pick out *subtle* cracking. Applying NSV to a cracking simulation would be a much better test of the suitability of this novel measurement.

In summary, I have compared a number of image analysis techniques in an attempt to be more quantitative when measuring the cracking. All techniques distinguish badly cracked images from pristine arrays, but this is also easily done by visual inspection of the raw images. A detailed study of cracking would benefit from the ability to distinguish among *slight* degrees of cracking, but the images used here differed too little to provide such a test. Ultimately, for the application considered here, such subtle distinction is not needed. As long as cracking is prevented enough to allow pattern transfer, it is good enough for this application.

6.1.3 Pattern transfer

Nanoparticles are proposed to take the place of BCP in future BPM generations, because of their extensibility to smaller sizes. Therefore, it is important to compare the size scales at which various benchmarks have been achieved, including the work discussed in this dissertation. As to raw mask creation (without demonstrating pattern transfer), the best BCP mask of which I am aware is a 6.9 nm pitch, achieved in 2009.[130] Though not intended specifically as a mask, self-assembly of 1 nm nanoparticles with a pitch smaller than 4 nm was shown in 1997.[145] As for pattern transfer applications, the smallest BCP demonstrations in recent years have hovered in the 20 nm to 30 nm range, with the smallest transfer being lines with a pitch of 17 nm.[78] This is the range to which my work ought to be compared. The pitch is dictated by the choice of

nanoparticles, and I have used two kinds: 15.7 nm-pitch Fe_3O_4, and 34 nm-pitch MnO. The latter are still in a range comparable to current BCP demos, but slightly higher. However, the former are smaller than the smallest transferred BCP pattern. Therefore, any pattern transfer obtained with the Fe_3O_4 occurs on size scales comparable to, *and slightly better than*, the current state of the art in BCP patterning.

I successfully demonstrated direct pattern transfer (i.e. "dots") into Si and SiO_2 with these Fe_3O_4 particles, using a CF_4-based RIE. The fidelity of transfer was tremendously improved by the cracking prevention, and most transferred features corresponded to single particles in the mask. However, the defect density was still too high for use in data storage applications: there were small cracks, wide pits, and bridged features. Very little optimization was done in parameter space, and it is plausible that the quality of dots from nanoparticle arrays could be improved enough for data storage applications.

One of the most important findings from our direct transfer experiments was that aspect ratio-dependent etching becomes severe in the extreme small-gap regime, so that etch rates were slowed by a factor of 10 or more. This is partially due to the high initial aspect ratio of our masks, since the nanoparticles are very close compared to their height. Probably, there are additional contributions from the fact that the gaps are now comparable to the size of the etchant and etch product molecules themselves. This constitutes a largely unexplored regime, which will become increasingly important as patterns of ever-higher density are fabricated.

The second technique I developed, based on antidot lattices, holds even more promise for industrial applications. First, it is more faithful; many lattices were free from noticeable defects. Second, it is more materials-general: I showed two different kinds of antidot material (Al and Cr) templated from two different kinds of particles (Fe_3O_4 and MnO). A further advantage is that it does not suffer from severe ARDE slowdown; in fact, etch rates were predicted to be near-bulklike. Finally, it automatically creates the proper structure for a nanoimprint mask: the *negative* of the desired dot pattern.

To attain the needed height contrast, the antidot lattice typically needs to be used as an etch mask for further pattern transfer. This is possible if the antidot material functions as a good etch-stop for an RIE recipe that etches the substrate. In my case, Al is known to be an excellent mask for SiO_2, and I demonstrated a tripling of the height contrast. The pitch for these particles is of the same order as most transferred BCP patterns, but slightly larger. Successful pattern transfer through the smaller Fe_3O_4-templated lattice would be a tremendous step forward, but has not yet been achieved.

My antidot results highlight the key importance of directionality at all stages of the process. First, directionality is necessary to deposit material deeply into very narrow gaps, which is why evaporation is preferred over sputtering. Second, directionality in etching is essential, because the antidot lattice has such

thin walls that any undercutting risks collapsing the structure, as happened for the STS-etched Fe_3O_4-based antidots. Here, I achieved that directionality by transferring from the MnO mask into SiO_2 rather than Si.

Ultimately, I have demonstrated that patterns can be transferred from arrays of surfactant-coated nanoparticles at size scales comparable to the state of the art in BCP, the closest competing self-assembling technology. Of these transfer attempts, the only major step that has not been demonstrated at sizes smaller than BCP is etching through the antidot lattice. Nevertheless, this *has* been demonstrated with particles not much bigger than BCP, and for every other benchmark the transfer I have achieved is at a smaller pitch.

6.2 Future work

6.2.1 Fast dosing with electron beam evaporator

So far, the bottleneck step has been the surfactant dosing, requiring a full weekend to write an area of $1\,mm^2$. Recent work by Dr. Matthew Moneck at Carnegie Mellon has promised to change that. Dr. Moneck noticed that the electron beam in the evaporator can sweep very quickly across an entire wafer, giving a much higher dosage rate than the SEM can achieve. The acceleration voltage is 6 keV, which is comparable in magnitude to the SEM beam at 10 keV.

Preliminary experiments suggest that this electron beam indeed cures the surfactant and prevents cracking. The main drawback is that precise control of the dosage and affected areas is extremely difficult. Furthermore, the intended purpose of the electron beam in the evaporator is to heat materials, and it is very difficult to prevent this when dosing a monolayer. Only the tiniest range on this instrument gives enough current to effect the change, but not so much current as to melt the sample. Fortunately, surfactant curing is a chemical reaction which can be driven to completion, so overdosing does not harm the sample as long as it is not significantly heated. Therefore, this method shows much promise for the rapid processing of samples.

Note, too, that it would be simple for industry to design a specialized instrument that could dose reliably over wide areas without risk of heating the sample. The only reason this step is difficult for us is that we must use tools designed for other purposes.

6.2.2 Optimizing pit depth

The cross-sectional SEM shown in Figure 5.22 is merely the first demonstration that antidot lattices can be used as masks for pit etching. Almost no optimization has been performed as of yet. Now that the possibility has been demonstrated, it is important to explore local parameter space and seek to improve the height contrast. I have demonstrated 15 nm, but industrial feedback at DSSC reviews has suggested that 40 nm would be more useful. Therefore,

the gas chemistry, power, and etching duration should be tweaked to see how deep it is possible to make the pits.

6.2.3 Magnetic recording demonstration

Another important milestone is the demonstration of magnetic recording on a sample patterned from nanoparticles. Towards this end, I have made preliminary efforts to develop a multilayer stack suitable for such a demonstration. The stack should have perpendicular magnetic anisotropy, so that it can be addressed by polar magneto-optical Kerr effect (MOKE). It should also be as thin as possible, to avoid the bottom deposited layers becoming contiguous with the top, and thereby magnetically coupled. I have developed a stack of less than 5 nm which has the required perpendicular anisotropy. Once enhanced height contrast has been obtained, this film should be deposited on the resulting structures and magnetic measurements taken to confirm the patterning.

6.3 Conclusion

First-generation bit-patterned media are poised to continue the long-standing increase in areal bit density. However, looking further ahead in the march to smaller features, it is clear that the current self-assembling technology (block copolymers) will encounter increasing difficulties at smaller sizes. With this in mind, I have endeavored to remove roadblocks to an alternative self-assembling technology: nanoparticles. While reliable guided self-assembly still needs to be shown, I have focused on the ability to transfer their patterns.

I first identified what I believe is the primary reason such transfer had not previously been shown: cracking. I developed two novel techniques to circumvent this problem, one of which in particular (electron beam curing) shows considerable promise for industrial applications. With the particles thereby stabilized, I showed successful transfer of both dots and antidots from nanoparticle arrays. In particular, the antidot process I developed promises to be highly faithful, materials-general, and extensible to smaller feature sizes. Both dots and antidots were fabricated at pitches smaller than the smallest demonstrated pattern transfer from block copolymers. Somewhat larger antidots were used as masks for further pattern transfer, which should be possible for the smaller antidots with further optimization in parameter space. Therefore, this work meaningfully advances the state of the art, and opens the possibility that nanoparticle arrays will be ready to take up the torch when block copolymers fail.

A Cleanroom Procedures

A.1 General cleanroom procedures

A.1.1 PlasmaTherm RIE recipes

Etching C using O_2

The goal of the O_2 etch is to minimize the physical sputtering, so non-carbon based structures are not appreciably eroded. It is assumed that the O_2 plasma is very reactive with the carbon in surfactant, so that very little physical component is needed to remove the carbon. Hence, the power should be as low as possible, and I have used 7 W. Unfortunately, this power is *so* low that it will not ignite the plasma, and a plasma ignition step at 12 W immediately precedes the etching step. After the gas flow has stabilized, the 12 W step begins. I watch through the window until I see the plasma has ignited, then immediately press the `End Step` button to proceed to the 7 W step. Surfactant etch times range from roughly 1 to 5 minutes, with 3 minutes an appropriate starting point for the Fe_3O_4 nanoparticles. The flow rate was 20 sccm, and the chamber pressure was nominally 5 mTorr, although the latter was closer to 7 mTorr to 8 mTorr in practice.

Etching Si using CHF_3 and CF_4

RIE of Si depends on the C:F ratio of the etch gases. Too little C, and the etch is too isotropic and undercuts the sidewalls. Too much, and polymer forms instead of etching. I found that the optimum ratio is close to 16:2 sccm of CF_4:CHF_3. A ratio of 14:4 sccm led to significant polymer, while 18:0 sccm led to undercutting and poor results. The power is 10 W, and the pressure is again a nominal 5 mTorr.

Etching SiO_2 using CF_4

Directionality in RIE of SiO_2 depends less on polymer formation, because its stronger bonds inhibit lateral etching. Hence, the best results involved 18 sccm of pure CF_4. The power was higher than for the Si etch, also because of the bonds: 15 W. The chamber pressure was nominally 5 mTorr.

A.1.2 Electron beam evaporator

The samples were usually irregularly shaped, so I affixed them to a carrier wafer with a more convenient shape. A small dab of white thermally conductive paste, available in the cleanroom, was applied to the face of this wafer. The sample was set onto this dab, and pressure was applied slowly but firmly over the dab by tweezers, rocking the sample back and forth until it was affixed to the carrier wafer. It is important to apply pressure only directly over the paste; otherwise, the sample can easily break. The carrier wafer is then loaded into an appropriate wafer holder and placed in the chamber.

The source is usually loaded in a crucible, and placed in one of four receptacles. Without the crucible, thermal resistance can be very low, and sometimes the sample cannot be heated enough to evaporate. The crucible also helps keep the source pure. Once the source and sample are loaded in the evaporator, the shutter is closed, and fresh glass slides are placed in the window to ensure visibility. The chamber is roughed and pumped to high vacuum. While this is occurring, the parameters for the deposited material may be entered into the crystal monitor.

Due to its high volume, the chamber achieves vacuum more slowly than most other cleanroom instruments. For quick runs, it may be acceptable to wait just 2 hours, getting in the low 10^{-6} Torr range. However, it is frequently more convenient to pumpdown overnight, and return to do the run first thing in the morning. Overnight pumpdowns achieve base pressures of 2×10^{-7} Torr to 4×10^{-7} Torr, which makes a significant difference to the mean free path (and therefore gives better directionality).

To deposit in the evaporator, very slowly bring up the filament current until the beam current needle first moves. The beam spot should be visible moving on the sample. Adjust the amplitude and position until obtaining an even sweep across the middle of the surface. From here, turn up the current very slowly until the deposition level is reached, stopping at least every 20 mA to allow the system to equilibrate. My test for equilibrium at a given current is that the pressure begins to drop; initially, each time the current is turned up, the pressure rises.

Once equilibrium is attained at the deposition current, the deposition is ready to begin. Press `stop`, `start`, then `zero` on the crystal monitor; this will start the timer and reset the measured thickness. Then, open the shutter in one fast but smooth motion. It is normal for the measured rate to be different before and after opening the shutter. Watch the deposition rate carefully, and especially the total deposited thickness. The former normally fluctuates by about $0.2\,\text{Å}\,\text{s}^{-1}$, while the latter should not fluctuate. Once the desired total thickness has been achieved, smoothly close the shutter and press `stop`, then write down all the parameters on the logsheet. I prefer slow deposition rates, about $0.2\,\text{Å}\,\text{s}^{-1}$ to $0.4\,\text{Å}\,\text{s}^{-1}$, because it is easier to achieve a desired thickness.

After deposition, slowly turn down the filament current, pausing at several intermediate current values along the way. (The slow cooldown is an attempt to avoid stressing the source.) Then, allow the chamber to sit for 30 min so the source can cool. When ready, close the high-vacuum valve and vent the chamber.

To remove the sample from the carrier, push gently on its side with a tweezer or a wooden swab. Once it is free from the paste, push it to the edge of the carrier until it can be picked up with tweezers and placed face-down in a concave sample holder. Apply a small amount of isopropanol to the cotten end of a wooden swab, and gently wipe the paste off the back of the sample, repeating until the sample is perfectly clean.

A.1.3 Electron beam dosing

It is critical to be able to find the patterned region after every processing step, even though some of those steps may make this region nearly indistinguishable from its surroundings. My approach is to use macroscopic markings. After the monolayer is transferred to the wafer, I inspect it visually and note promising regions. I then use a diamond scribe to make shallow scratch marks which end near the promising region. Once the sample is loaded in the SEM, I choose a reference edge and align the stage so that it is perfectly vertical or horizontal. I navigate to the tip of a scratch, and inspect the monolayer in the immediate vicinity to ensure the ordering is good enough.

Once I have found a suitable region, I write down the absolute coordinates, and travel to the Faraday cup. I set the desired voltage and spot size, and measure the beam current, updating the runfile with the result. Then, I correct the astigmatism by focusing on the gold standard. (Both the Faraday cup and the gold standard are found on certain SEM sample holders which are designated for electron beam lithography.) I next return to the coordinates I wrote down earlier, and focus the beam at the tip of the scratch. I initiate the runfile and follow the instructions it gives. Numbered blocks with different doses can be found in the runfile `numbered_blocks`, while the file `numbered_blocks_quick` writes only the middle column to save time.

A.2 SPL Recipes

The steps for making the SPL samples have mostly been superseded by the electron beam curing steps. However, I record them here in case some unforeseen application arises.

A.2.1 Making the sacrificial wafer

The Carnegie Mellon nanofabrication facility sells Si wafers that are just 220 µm thick, and 4″ in diameter. This thickness is advantageous: thin enough to

minimize the amount of material that needs to be removed, but thick enough to remain rigid during handling. These wafers are coated in 1 µm thermal SiO_2 on both sides, one of which must be cleared of SiO_2 to allow the Si to be etched away later. To remove the SiO_2, I floated the wafer on top of a teflon petri dish of HF, handling it with teflon-coated tweezers. After roughly 15 minutes, the wafer spontaneously moved in a random direction, indicating that the HF etch was completed.

Further pieces were prepared by breaking the wafer into smaller pieces. To add the sacrificial aluminum layer to a given piece, I affixed it to a 3″ carrier wafer by standard thermally conductive paste, and loaded it in the 5-target load-lock sputtering machine in Bay 4 of the cleanroom. I deposited Al to a thickness of 50 nm. Then, I removed the wafer, and carefully swabbed the paste off the back with acetone, taking care not to break the fragile wafer.

A.2.2 Overcoating and affixing to handling wafer

First, I prepared the nanoparticle monolayer by the usual techniques described in Section 2.4.1. I transferred them to the Al-coated side of the wafers. Then, I loaded them face-up into the 8L sputtering machine, and deposited SiO_2 on top of them. This served both as the structural pinning layer, and the layer to be patterned. I then loaded the wafers into the load-lock sputtering machine, together with the blank 1″ wafers that would serve as carriers, and deposited 800 nm Cu to serve as a buffer layer.

The epoxy step varied based on the type of epoxy used. M-bond 610 is used for preparing cross-sectional TEM samples, and is notable for its very small thickness, which I measured to be as little as 10 nm in some places. I applied a small droplet to the carrier wafer by dipping a wooden q-tip in the M-bond, then touching the tip to the wafer. The original wafer was applied face-down to the carrier, and the assembly was placed between two metal blocks. The blocks were wrapped in aluminum foil in case any epoxy leaked out the sides, and held tightly together by two large binder clips. The entire assembly was placed in an oven at 110 °C for 1 h to cure the epoxy, then removed and allowed to cool.

A.2.3 Removing the various layers

I removed the silicon by exposing the backside of the original (i.e. thin) wafer to XeF_2 gas, using a Xactix XeF_2 etcher. The exposed faces of the carrier wafer were first covered with kapton tape, so that only the backside of the original wafer was showing. Then the wafer was placed into the etcher. Typically, the Si was completely removed in less than 2.5 h. XeF_2 is isotropic, so the carrier wafer exhibited erosion also along the sides.

To remove the SiO_2 layer, a CF_4 etch was used. The particle monolayer was still not visible. Then, the Al layer was removed by stirring in AZ400K

developer for 10 min. After this step, the monolayer again became visible, and SEM screening confirmed that the particles were exposed.

References

[1] K. Aissou, M. Kogelschatz, T. Baron, and P. Gentile. Self-assembled block polymer templates as high resolution lithographic masks. *Surface Science*, 601(13):2611–2614, July 2007.

[2] Manfred Albrecht, Guohan Hu, Ildico L Guhr, Till C Ulbrich, Johannes Boneberg, Paul Leiderer, and Günter Schatz. Magnetic multilayers on nanospheres. *Nature Materials*, 4(3):203–6, March 2005.

[3] M. Allen, D. Willits, J. Mosolf, M. Young, and T. Douglas. Protein cage constrained synthesis of ferrimagnetic iron oxide nanoparticles. *Advanced Materials*, 14(21):1562–1565, 2002.

[4] H. Andersen and H. Bay. Sputtering yield measurements. *Sputtering by Particle Bombardment I*, pages 145–218, 1979.

[5] T. Aoyama, S. Okawa, K. Hattori, H. Hatate, Y. Wada, K. Uchiyama, T. Kagotani, H. Nishio, and I. Sato. Fabrication and magnetic properties of CoPt perpendicular patterned media. *Journal of Magnetism and Magnetic Materials*, 235(1-3):174–178, October 2001.

[6] Wolfgang M. Arden. The international technology roadmap for Semiconductors–Perspectives and challenges for the next 15 years. *Current Opinion in Solid State and Materials Science*, 6(5):371–377, October 2002.

[7] J. C. Arnold and H. H. Sawin. Charging of pattern features during plasma etching. *Journal of Applied Physics*, 70(10):5314–5317, November 1991.

[8] T. Bala, B. L. V. Prasad, M. Sastry, M. U Kahaly, and U. V Waghmare. Interaction of different metal ions with carboxylic acid group: A quantitative study. *J. Phys. Chem. A*, 111(28):6183–6190, 2007.

[9] S. D. Berger and J. M. Gibson. New approach to projection-electron lithography with demonstrated 0 1µm linewidth. *Applied Physics Letters*, 57(2):153, 1990.

[10] Karl Berggren. Lithography for Bit-Patterned-Media. Information Storage Industry Consortium's EHDR Technical Review, May 2010.

[11] A. E. Berkowitz, G. F. Rodriguez, J. I. Hong, K. An, T. Hyeon, N. Agarwal, D. J. Smith, and E. E. Fullerton. Monodispersed MnO nanoparticles with epitaxial Mn3O4 shells. *Journal of Physics D: Applied Physics*, 41(13):134007, 2008.

[12] T. P. Bigioni, X. M. Lin, T. T. Nguyen, E. I. Corwin, T. A. Witten, and H. M. Jaeger. Kinetically driven self assembly of highly ordered nanoparticle monolayers. *Nature Materials*, 5(4):265–270, 2006.

[13] Ion Bita, Joel K. W. Yang, Yeon Sik Jung, Caroline A. Ross, Edwin L. Thomas, and Karl K. Berggren. Graphoepitaxy of Self-Assembled block copolymers on Two-Dimensional periodic patterned templates. *Science*, 321(5891):939–943, August 2008.

[14] C. T. Black, C. B. Murray, R. L. Sandstrom, and Shouheng Sun. Spin-Dependent tunneling in Self-Assembled Cobalt-Nanocrystal superlattices. *Science*, 290(5494):1131–1134, November 2000.

[15] F. Bresme and M. Oettel. Nanoparticles at fluid interfaces. *Journal of Physics, Condensed Matter*, 19(41):413101, 2007.

[16] Timothy A. Brunner. Why optical lithography will live forever. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 21(6):2632–2637, November 2003.

[17] Peter Buseck, John Maxwell Cowley, and LeRoy Eyring. *High-resolution transmission electron microscopy and associated techniques*. Oxford University Press US, 1988.

[18] John Canny. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(6):679–698, 1986.

[19] Cheng Cen, Stefan Thiel, Jochen Mannhart, and Jeremy Levy. Oxide nanoelectronics on demand. *Science*, 323(5917):1026–1030, February 2009.

[20] A. Chambers, R.K. Fitch, and B.S. Halliday. *Basic Vacuum Technology,*. Taylor & Francis, 1 edition, January 1989.

[21] T. H. P. Chang. Proximity effect in electron-beam lithography. *Journal of Vacuum Science and Technology*, 12(6):1271–1275, November 1975.

[22] CC Chao, TC Wang, RM Ho, P Georgopanos, A Avgeropoulos, and EL Thomas. Robust block copolymer mask for nanopatterning polymer films. *ACS NANO*, 4(4):2088–2094, April 2010.

[23] Steve Chapman. Understanding & optimising SEM performance. *InFocus*, (14):18–27, June 2009.

[24] S. H. Charap, P. L Lu, and Y. He. Thermal stability of recorded information at high densities. In *Magnetic Recording Conference 1996. Magnetic Recording Media., Digests of the*, page F1, 1996.

[25] J.Y. Cheng, C.A. Ross, E.L. Thomas, H.I. Smith, and G.J. Vancso. Templated Self-Assembly of block copolymers: Effect of substrate topography. *Advanced Materials*, 15(19):1599–1602, 2003.

[26] Stephen Y. Chou, Peter R. Krauss, and Preston J. Renstrom. Imprint of sub-25 nm vias and trenches in polymers. *Applied Physics Letters*, 67(21):3114, 1995.

[27] Stephen Y. Chou, Peter R. Krauss, Wei Zhang, Lingjie Guo, and Lei Zhuang. Sub-10 nm imprint lithography and applications. In *Papers from the 41st international conference on electron, ion, and photon beam technology and nanofabrication*, volume 15, pages 2897–2904, Dana Point, California (USA), November 1997. AVS.

[28] J. W. Coburn and Harold F. Winters. Conductance considerations in the reactive ion etching of high aspect ratio features. *Applied Physics Letters*, 55(26):2730–2732, December 1989.

[29] David Cohen-Tanugi and Nan Yao. Superior imaging resolution in scanning helium-ion microscopy: A look at beam-sample interactions. *Journal of Applied Physics*, 104(6):063504, 2008.

[30] M. Colburn, S. Johnson, M. Stewart, S. Damle, T. Bailey, B. Choi, M. Wedlake, T. Michaelson, S. V. Sreenivasan, J. Ekerdt, et al. Step and flash imprint lithography: a new approach to high-resolution patterning. In *Proc. SPIE*, volume 3676, pages 379–389, 1999.

[31] M. F. Crommie, C. P. Lutz, and D. M. Eigler. Confinement of electrons to quantum corrals on a metal surface. *Science*, 262(5131):218–220, October 1993.

[32] D. Crouse, Yu-Hwa Lo, A. E. Miller, and M. Crouse. Self-ordered pore structure of anodized aluminum on silicon and pattern transfer. *Applied Physics Letters*, 76(1):49, 2000.

[33] B. O. Dabbousi, C. B. Murray, M. F. Rubner, and M. G. Bawendi. Langmuir-Blodgett manipulation of Size-Selected CdSe nanocrystallites. *Chemistry of Materials*, 6(2):216–219, February 1994.

[34] S.B. Darling. Directing the self-assembly of block copolymers. *Progress in Polymer Science*, 32(10):1152–1204, October 2007.

[35] H. W. Deckman. Natural lithography. *Applied Physics Letters*, 41(4):377, 1982.

[36] William M. Dehn and Kirby E. Jackson. Phosphoric acid in organic reactions. *Journal of the American Chemical Society*, 55(10):4284–4287, October 1933.

[37] Y. Ding, S.A. Majetich, J. Kim, K. Barmak, H. Rollins, and P. Sides. Sintering prevention and phase transformation of FePt nanoparticles. *Journal of Magnetism and Magnetic Materials*, 284:336–341, December 2004.

[38] Y. Ding, S. Yamamuro, D. Farrell, and S. A. Majetich. Phase transformation and magnetic moment in FePt nanoparticles. *Journal of Applied Physics*, 93(10):7411, 2003.

[39] T. Djenizian, B. Petite, L. Santinacci, and P. Schmuki. Electron-beam induced carbon deposition used as a mask for cadmium sulfide deposition on si(100). *Electrochimica Acta*, 47(6):891–897, December 2001.

[40] C. D. Dushkin, G. S. Lazarov, S. N. Kotsev, H. Yoshimura, and K. Nagayama. Effect of growth conditions on the structure of two-dimensional latex crystals: experiment. *Colloid & Polymer Science*, 277(10):914–930, October 1999.

[41] D. M. Eigler and E. K. Schweizer. Positioning single atoms with a scanning tunnelling microscope. *Nature*, 344(6266):524–526, April 1990.

[42] Doug H. Everett. *Basic Principles of Colloid Science*. Royal Society of Chemistry, 1 edition, December 1988.

[43] T E Everhart and R F M Thornley. Wide-band detector for micro-microampere low-energy electron currents. *Journal of Scientific Instruments*, 37(7):246–248, 1960.

[44] U. Ch. Fischer and H. P. Zingsheim. Submicroscopic pattern replication with visible light. *Journal of Vacuum Science and Technology*, 19(4):881–885, November 1981.

[45] Bruno La Fontaine, Yunfei Deng, Ryoung han Kim, Harry J. Levinson, Uzodinma Okoroanyanwu, Richard Sandberg, Tom Wallow, and Obert Wood. Extreme ultraviolet lithography: From research to manufacturing. In *J. Vac. Sci. Technol. B*, volume 25, pages 2089–2093. AVS, November 2007.

[46] Roger H. French and Hoang V. Tran. Immersion lithography: Photomask and Wafer-Level materials. *Annual Review of Materials Research*, 39(1):93–126, 2009.

[47] Barbra L. Gabriel. *Sem: A User's Manual for Materials Science*. American Society for Metals, July 1985.

[48] Bernhard Gehl, Andreas Frömsdorf, Vesna Aleksandrovic, Thomas Schmidt, Angelika Pretorius, Jan-Ingo Flege, Sigrid Bernstorff, Andreas Rosenauer, Jens Falta, Horst Weller, and Marcus Bäumer. Structural and chemical effects of plasma treatment on Close-Packed colloidal nanoparticle layers. *Advanced Functional Materials*, 18(16):2398–2410, 2008.

[49] J.T. Gerig. *Introductory Organic Chemistry*. Academic Press Inc, May 1974.

[50] Konstantinos P. Giapis, Geoffrey R. Scheller, Richard A. Gottscho, William S. Hobson, and Yong H. Lee. Microscopic and macroscopic uniformity control in plasma etching. *Applied Physics Letters*, 57(10):983, 1990.

[51] J. Gierak, A. Madouri, A.L. Biance, E. Bourhis, G. Patriarche, C. Ulysse, D. Lucot, X. Lafosse, L. Auvray, L. Bruchhaus, and R. Jede. Sub-5 nm FIB direct patterning of nanodevices. *Microelectronic Engineering*, 84(5-8):779–783, May 2007.

[52] Evangelos Gogolides, Philippe Vauvert, George Kokkoris, Guy Turban, and Andreas G. Boudouvis. Etching of SiO_2 and si in fluorocarbon plasmas: A detailed surface model accounting for etching and deposition. *Journal of Applied Physics*, 88(10):5570, 2000.

[53] Marc De Graef. *Introduction to Conventional Transmission Electron Microscopy*. Cambridge University Press, April 2003.

[54] K. Grigoras, L. Sainiemi, J. Tiilikainen, A. Saynatjoki, V.-M. Airaksinen, and S. Franssila. Application of ultra-thin aluminum oxide etch mask made by atomic layer deposition technique. *Journal of Physics: Conference Series*, 61:369–373, 2007.

[55] A.E. Grigorescu, M.C. van der Krogt, C.W. Hagen, and P. Kruit. 10 nm lines and spaces written in HSQ, using electron beam lithography. *Microelectronic Engineering*, 84(5-8):822–824, May 2007.

[56] L. J. Guo. Nanoimprint lithography: Methods and material requirements. *Advanced Materials*, 19(4):495–513, 2007.

[57] Chiseki Haginoya, Seiji Heike, Masayoshi Ishibashi, Kimio Nakamura, Kazuyuki Koike, Toshiyuki Yoshimura, Jiro Yamamoto, and Yoshiyuki Hirayama. Magnetic nanoparticle array with perpendicular crystal magnetic anisotropy. *Journal of Applied Physics*, 85(12):8327, 1999.

[58] Andrew M. Hawryluk and Lynn G. Seppala. Soft x-ray projection lithography using an x-ray reduction camera. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 6(6):2162–2166, November 1988.

[59] C. L Haynes and R. P Van Duyne. Nanosphere lithography: a versatile nanofabrication tool for studies of size-dependent nanoparticle optics. *J. Phys. Chem. B*, 105(24):5599–5611, 2001.

[60] John N. Helbert. *Handbook of VLSI Microlithography, 2nd Edition, Second Edition*. William Andrew, 2 edition, January 2002.

[61] O Hellwig, JK Bosworth, E Dobisz, D Kercher, T Hauet, G Zeltzer, JD Risner-Jamtgaard, D Yaney, and R Ruiz. Bit patterned media based on block copolymer directed assembly with narrow magnetic switching field distribution. *Applied Physics Letters*, 96(5), February 2010.

[62] H Hieda, Y Yanagita, A Kikitsu, T Maeda, and K Naito. Fabrication of FePt patterned media with diblock copolymer templates. *Journal of Photopolymer Science and Technology*, 19(3):425–430, 2006.

[63] Paul C Hiemenz and Raj Rajagopalan. *Principles of Colloid and Surface Chemistry*. Marcel Dekker, New York, 3rd ed., rev. and expanded edition, 1997.

[64] Tomoyasu Hirai, Melvina Leolukman, Chi Chun Liu, Eungnak Han, Yun Jun Kim, Yoshihito Ishida, Teruaki Hayakawa, Masa aki Kakimoto, Paul F. Nealey, and Padma Gopalan. One-Step Direct-Patterning template utilizing Self-Assembly of POSS-Containing block copolymers. *Advanced Materials*, 21(43):4334, 2009.

[65] Saw-Wai Hla. Scanning tunneling microscopy single atom/molecule manipulation and its application to nanoscience and technology. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 23(4):1351–1360, July 2005.

[66] C.R. Hogg, S.A. Majetich, and J.A. Bain. Investigating pattern transfer in the Small-Gap regime using Electron-Beam stabilized nanoparticle array etch masks. *Magnetics, IEEE Transactions on*, 46(6):2307–2310, 2010.

[67] G. Hu. Magnetic and recording properties of Co/Pd islands on prepatterned substrates. *Journal of Applied Physics*, 95(11):7013, 2004.

[68] W. R. Hunter, L. Ephrath, W. D. Grobman, C. M. Osburn, B. L. Crowder, A. Cramer, and H. E. Luhn. 1-micron MOSFET VLSI technology. v - a single-level polysilicon technology using electron-beam lithography. *IEEE Journal of Solid-State Circuits*, 14:275–281, April 1979.

[69] T. Hyeon, S.S. Lee, J. Park, Y. Chung, and H.B. Na. Synthesis of highly crystalline and monodisperse maghemite nanocrystallites without a Size-Selection process. *Journal of the American Chemical Society*, 123(51):12798–12801, December 2001.

[70] Taeghwan Hyeon. Chemical synthesis of magnetic nanoparticles. *Chemical Communications*, (8):927–934, 2003.

[71] R.K. Iler. Multilayers of colloidal particles. *Journal of Colloid and Interface Science*, 21(6):569–594, June 1966.

[72] Filip Ilievski, C. A. Ross, and G. J. Vancso. Magnetic reversal phenomena of perpendicular magnetic islands fabricated by block copolymer lithography. *Journal of Applied Physics*, 103(7):07C520, 2008.

[73] S. G. Ingram. The influence of substrate topography on ion bombardment in plasma etching. *Journal of Applied Physics*, 68(2):500, 1990.

[74] Jacob N. Israelachvili. *Intermolecular and Surface Forces, Second Edition: With Applications to Colloidal and Biological Systems*. Academic Press, 2 edition, January 1992.

[75] Takashi Ito and Shinji Okazaki. Pushing the limits of lithography. *Nature*, 406(6799):1027–1031, 2000.

[76] ITRS. International technology roadmap for semiconductors, 2009 edition: lithography, 2009.

[77] Jean-Pierre Jolivet and Elisabeth Tronc. Interfacial electron transfer in colloidal spinel iron oxide. conversion of Fe_3O_4-γ−Fe_2O_3 in aqueous medium. *Journal of Colloid and Interface Science*, 125(2):688–701, October 1988.

[78] Yeon Sik Jung, J. B. Chang, Eric Verploegen, Karl K. Berggren, and C. A. Ross. A path to ultranarrow patterns using Self-Assembled lithography. *Nano Letters*, 10(3):1000–1005, March 2010.

[79] Yeon Sik Jung and Caroline A. Ross. Well-Ordered Thin-Film nanopore arrays formed using a Block-Copolymer template. *Small*, 5(14):1654–1659, 2009.

[80] Y. Kamata, A. Kikitsu, H. Hieda, M. Sakurai, and K. Naito. Ar ion milling process for fabricating CoCrPt patterned media using a self-assembled PS-PMMA diblock copolymer mask. In *J. Appl. Phys.*, volume 95, pages 6705–6707. AIP, June 2004.

[81] P. Kappenberger, F. Luo, L. J. Heyderman, H. H. Solak, C. Padeste, C. Brombacher, D. Makarov, T. V. Ashworth, L. Philippe, H. J. Hug, et al. Template-directed self-assembled magnetic nanostructures for probe recording. *Applied Physics Letters*, 95:023116, 2009.

[82] D. Keil and E. Anderson. Characterization of reactive ion etch lag scaling. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 19(6):2082–2088, November 2001.

[83] F. Keller, M. S. Hunter, and D. L. Robinson. Structural features of oxide coatings on aluminum. *Journal of The Electrochemical Society*, 100(9):411–419, 1953.

[84] D. K. Kim, Y. Zhang, W. Voit, K. V. Rao, and M. Muhammed. Synthesis and characterization of surfactant-coated superparamagnetic monodispersed iron oxide nanoparticles. *Journal of Magnetism and Magnetic Materials*, 225(1-2):30–36, 2001.

[85] Sang Ouk Kim, Harun H. Solak, Mark P. Stoykovich, Nicola J. Ferrier, Juan J. de Pablo, and Paul F. Nealey. Epitaxial self-assembly of block copolymers on lithographically defined nanopatterned substrates. *Nature*, 424(6947):411–414, July 2003.

[86] K. J. Kirk, J. N. Chapman, S. McVitie, P. R. Aitchison, and C. D. W. Wilkinson. Switching of nanoscale magnetic elements. *Applied Physics Letters*, 75(23):3683, 1999.

[87] Hiroaki Kitahara, Yuhei Uno, Hiroaki Suzuki, Takashi Kobayashi, Hiroshi Tanaka, Yoshiaki Kojima, Masaki Kobayashi, Masahiro Katsumura, Yasumitsu Wada, and Tetsuya Iida. Electron beam recorder for patterned media mastering. *Japanese Journal of Applied Physics*, 49:06, 2010.

[88] M. Klokkenburg, J. Hilhorst, and B.H. ErnÃľ. Surface analysis of magnetite nanoparticles in cyclohexane solutions of oleic acid and oleylamine. *Vibrational Spectroscopy*, 43(1):243–248, January 2007.

[89] R. Knizikevicius. Real dimensional simulation of silicon etching in CF_4 + O_2 plasma. *Applied Surface Science*, 201(1-4):96–108, November 2002.

[90] R. Knizikevicius. Influence of activated polymer on the etching rate of SiO_2 in CF_4 + H_2 plasma. *Microelectronic Engineering*, 86(1):55–58, January 2009.

[91] George Kokkoris, Evangelos Gogolides, and Andreas. G. Boudouvis. Etching of SiO_2 features in fluorocarbon plasmas: Explanation and prediction of gas-phase-composition effects on aspect ratio dependent phenomena in trenches. *Journal of Applied Physics*, 91(5):2697, 2002.

[92] K. L. Krycka, R. A. Booth, C. R. Hogg, Y. Ijiri, J. A. Borchers, W. C. Chen, S. M. Watson, M. Laver, T. R. Gentile, L. R. Dedon, S. Harris, J. J. Rhyne, and S. A. Majetich. Core-Shell magnetic morphology of structurally uniform magnetite nanoparticles. *Physical Review Letters*, 104(20):207203, May 2010.

[93] V. K. LaMer and R. H. Dinegar. Theory, production and mechanism of formation of monodispersed hydrosols. *Journal of the American Chemical Society*, 72(11):4847–4854, 1950.

[94] M. J. Lercel, M. Rooks, R. C. Tiberio, H. G. Craighead, C. W. Sheen, A. N. Parikh, and D. L. Allara. Pattern transfer of electron beam modified self-assembled monolayers for high-resolution lithography. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 13(3):1139–1143, May 1995.

[95] Diandra L. Leslie-Pelecky and Reuben D. Rieke. Magnetic properties of nanostructured materials. *Chemistry of Materials*, 8(8):1770–1783, January 1996.

[96] J. A Liddle, G. M Gallatin, and L. E Ocola. Resist requirements and limitations for nanoscale Electron-Beam patterning. In *Materials Research Society Symposium Proceedings*, volume 739, pages 19–30, 2003.

[97] Burn J. Lin. NGL comparable to 193-nm lithography in cost, footprint, and power consumption. *Microelectronic Engineering*, 86(4-6):442–447, April 2009.

[98] X. M. Lin, R. Parthasarathy, and H. M. Jaeger. Direct patterning of self-assembled nanocrystal monolayers by electron beams. *Applied Physics Letters*, 78(13):1915–1917, March 2001.

[99] T. Lindeberg. Edge detection and ridge detection with automatic scale selection. *International Journal of Computer Vision*, 30(2):117–154, 1998.

[100] T Lindeberg. Feature detection with automatic scale selection. *International Journal of Computer Vision*, 30(2):79–116, November 1998.

[101] H. Liu. Tuning the binding energy of surfactant to CdSe nanocrystal: A theoretical study. *The Journal of Physical Chemistry C*, 113(8):3116–3119, 2009.

[102] J. C. Lodder. Methods for preparing patterned media for high-density recording. *Journal of Magnetism and Magnetic Materials*, 272-276(Part 3):1692–1697, May 2004.

[103] Marc J. Madou. *Fundamentals of Microfabrication: The Science of Miniaturization, Second Edition*. CRC Press, 2 edition, March 2002.

[104] M. Maillard, L. Motte, A.T. Ngo, and M.P. Pileni. Rings and hexagons made of nanocrystals: A marangoni effect. *Journal of Physical Chemistry B*, 104(50):11871–11877, December 2000.

[105] S. A. Majetich and M. Sachan. Magnetostatic interactions in magnetic nanoparticle assemblies: energy, time and length scales. *Journal of Physics D: Applied Physics*, 39(21):407, 2006.

[106] Hideki Masuda and Kenji Fukuda. Ordered metal nanohole arrays made by a Two-Step replication of honeycomb structures of anodic alumina. *Science*, 268(5216):1466–1468, June 1995.

[107] Hideki Masuda and Masahiro Satoh. Fabrication of gold nanodot array using anodic porous alumina as an evaporation mask. *Japanese Journal of Applied Physics*, 35:L126–L129, 1996.

[108] Yoshitaka Matsui, Kazuyuki Nishio, and Hideki Masuda. Highly ordered anodic porous alumina with 13-nm hole intervals using a 2D array of monodisperse nanoparticles as a template. *Small*, 2(4):522–525, 2006.

[109] W. H. Meiklejohn and C. P. Bean. New magnetic anisotropy. *Physical Review*, 102(5):1413, June 1956.

[110] Gordon E. Moore. Cramming more components onto integrated circuits. *Electronics Magazine*, 4, 1965.

[111] J. Moritz, L. Buda, B. Dieny, J. P. Nozieres, R. J. M. van de Veerdonk, T. M. Crawford, and D. Weller. Writing and reading bits on pre-patterned media. *Applied Physics Letters*, 84(9):1519–1521, March 2004.

[112] Judith Moser, V. Kunej, H. Pernau, G. Schatz, E. Scheer, and Manfred Albrecht. Magnetoresistive effects in co/pd multilayers on self-assembled nanospheres. 11th Joint MMM-Intermag Conference.

[113] B. M Moskowitz. Hitchhiker's guide to magnetism. In *Handbook from the Environmental Magnetism Workshop*, pages 1–40.

[114] C. B. Murray, C. R. Kagan, and M. G. Bawendi. Synthesis and characterization of monodisperse nanocrystals and close-packed nanocrystal assemblies. *Annual Review of Materials Science*, 30(1):545–610, 2000.

[115] S. Nabavi, B.V.K. Vijaya Kumar, J.A. Bain, C. Hogg, and S.A. Majetich. Application of image processing to characterize patterning noise in Self-Assembled Nano-Masks for Bit-Patterned media. *Magnetics, IEEE Transactions on*, 45(10):3523–3526, 2009.

[116] Katsuyuki Naito. Ultrahigh-density storage media prepared by artificially assisted self-assembling methods. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 15(4):047507, 2005.

[117] S. Narayanan, J. Wang, and X. M. Lin. Dynamical Self-Assembly of nanocrystal superlattices during colloidal droplet evaporation by in situ small angle X-Ray scattering. *Physical Review Letters*, 93(13):135503, 2004.

[118] Anh Tu Ngo, Johannes Richardi, and Marie Paule Pileni. Cracks in magnetic nanocrystal films: Do directional and isotropic crack patterns follow the same scaling law? *Nano Letters*, 8(8):2485–2489, August 2008.

[119] G. S. Oehrlein. Fluorocarbon high-density plasmas. i. fluorocarbon film deposition and etching using CF4 and CHF3. *Journal of Vacuum Science & Technology A: Vacuum, Surfaces, and Films*, 12(2):323, 1994.

[120] Robert C. O'Handley. *Modern Magnetic Materials: Principles and Applications*. Wiley-Interscience, November 1999.

[121] Tokuhisa Ohiwa, Akihiro Kojima, Makoto Sekine, Itsuko Sakai, Shigeru Yonemoto, and Yumi Watanabe. Mechanism of etch stop in high Aspect-Ratio contact hole etching. *Japanese Journal of Applied Physics*, 37:5060–5063, 1998.

[122] Milton Ohring. *Materials Science of Thin Films, Second Edition*. Academic Press, 2 edition, October 2001.

[123] SF Pang, Y Kurosawa, T Kondo, and T Kawai. Decomposition of monolayer coverage on gold nanoparticles by UV/ozone treatment. *Chemistry Letters*, 34(4):544–545, April 2005.

[124] Mihir Parikh. Corrections to proximity effects in electron beam lithography. i. theory. *Journal of Applied Physics*, 50(6):4371, 1979.

[125] B J Park, S W Kim, S K Kang, K S Min, S D Park, S J Kyung, H C Lee, J W Bae, J T Lim, D H Lee, and G Y Yeom. Low angle forward reflected neutral beam source and its applications. *Journal of Physics D: Applied Physics*, 41(2):024005, 2008.

[126] J. Park, K. An, Y. Hwang, J. G. Park, H. J. Noh, J. Y. Kim, J. H. Park, N. M. Hwang, and T. Hyeon. Ultra-large-scale syntheses of monodisperse nanocrystals. *Nature Materials*, 3(12):891–895, 2004.

[127] Jongnam Park, Jin Joo, Soon-Gu Kwon, Youngjin Jang, and Taeghwan Hyeon. Synthesis of monodisperse spherical nanocrystals. *Angewandte Chemie International Edition*, 46(25):4630–4660, 2007.

[128] Miri Park, P. M. Chaikin, Richard A. Register, and Douglas H. Adamson. Large area dense nanoscale patterning of arbitrary surfaces. *Applied Physics Letters*, 79(2):257, 2001.

[129] Miri Park, Christopher Harrison, Paul M. Chaikin, Richard A. Register, and Douglas H. Adamson. Block copolymer lithography: Periodic arrays of ~10^{11} holes in 1 square centimeter. *Science*, 276(5317):1401–1404, May 1997.

[130] Soojin Park, Dong Hyun Lee, Ji Xu, Bokyung Kim, Sung Woo Hong, Unyong Jeong, Ting Xu, and Thomas P. Russell. Macroscopic 10-Terabit-per-Square-Inch arrays from block copolymers with lateral order. *Science*, 323(5917):1030–1033, February 2009.

[131] Xiaogang Peng, Yan Zhang, Bingsuo Zou, Liangzhi Xiao, Tiejin Li, and Jun Yang. Formation of nanoparticulate iron(III) oxide-stearate multilayer through Langmuir-Blodgett method. *The Journal of Physical Chemistry*, 96(8):3412–3415, April 1992.

[132] V. Perez-Dieste, O. M. Castellini, J. N. Crain, M. A. Eriksson, A. Kirakosian, J.-L. Lin, J. L. McChesney, F. J. Himpsel, C. T. Black, and C. B. Murray. Thermal decomposition of surfactant coatings on co and ni nanocrystals. *Applied Physics Letters*, 83(24):5053, 2003.

[133] H. C. Pfeiffer. Direct write electron beam Lithography–A production line reality. *Solid State Technol*, 27:223, 1984.

[134] Yuanzhe Piao, Hyunchang Lim, Ji Young Chang, Won-Yong Lee, and Hasuck Kim. Nanostructured materials prepared by use of ordered porous alumina membranes. *Electrochimica Acta*, 50(15):2997–3013, May 2005.

[135] V.F. Puntes, D. Zanchet, C.K. Erdonmez, and A.P. Alivisatos. Synthesis of hcp-Co nanodisks. *Journal of the American Chemical Society*, 124(43):12874–12880, October 2002.

[136] Aaron Puzder, Andrew J. Williamson, Natalia Zaitseva, Giulia Galli, Liberato Manna, and A. Paul Alivisatos. The effect of organic ligand binding on the growth of CdSe nanoparticles probed by ab initio calculations. *Nano Letters*, 4(12):2361–2365, December 2004.

[137] Ivo W. Rangelow. Critical tasks in high aspect ratio silicon dry etching for microelectromechanical systems. In *Papers from the 49th International Symposium of the American Vacuum Society*, volume 21, pages 1550–1562, Denver, Colorado (USA), July 2003. AVS.

[138] M. Ranjbar, S.N. Piramanayagam, Deng Suzi, Kyaw Oo Aung, R. Sbiaa, Yew Seng Kay, Seng Kai Wong, and Chong Tow Chong. Antiferromagnetically coupled patterned media and control of switching field distribution. *Magnetics, IEEE Transactions on*, 46(6):1787–1790, 2010.

[139] H J Richter. The transition from longitudinal to perpendicular recording. *Journal of Physics D: Applied Physics*, 40(9):R149–R177, 2007.

[140] C. A. Ross, S. Haratani, F. J. Castano, Y. Hao, M. Hwang, M. Shima, J. Y. Cheng, B. Vogeli, M. Farhoud, M. Walsh, and Henry I. Smith. Magnetic behavior of lithographically patterned particle arrays (invited). In *J. Appl. Phys.*, volume 91, pages 6848–6853. AIP, May 2002.

[141] M. Rothschild, T. M. Bloomstein, R. R. Kunz, V. Liberman, M. Switkes, S. T. Palmacci, J. H. C. Sedlacek, D. Hardy, and A. Grenville. Liquid immersion lithography: Why, how, and when? In *The 48th International Conference on Electron, Ion, and Photon Beam Technology and Nanofabrication*, volume 22, pages 2877–2881, San Diego, California (USA), November 2004. AVS.

[142] Ricardo Ruiz, Huiman Kang, Francois A. Detcheverry, Elizabeth Dobisz, Dan S. Kercher, Thomas R. Albrecht, Juan J. de Pablo, and Paul F. Nealey. Density multiplication and improved lithography by directed block copolymer assembly. *Science*, 321(5891):936–939, August 2008.

[143] M. Sachan, C. Bonnoit, C. Hogg, E. Evarts, J. A. Bain, S. A. Majetich, J.-H. Park, and J.-G. Zhu. Self-assembled nanoparticle arrays as nanomasks for pattern transfer. *Journal of Physics D: Applied Physics*, 41(13):134001, 2008.

[144] V. Santhanam, J. Liu, R. Agarwal, and R.P. Andres. Self-Assembly of uniform monolayer arrays of nanoparticles. *Langmuir*, 19(19):7881–7887, September 2003.

[145] K. Vijaya Sarathy, Gargi Raina, R. T. Yadav, G. U. Kulkarni, and C. N. R. Rao. Thiol-Derivatized nanocrystalline arrays of gold, silver, and platinum. *The Journal of Physical Chemistry B*, 101(48):9876–9880, November 1997.

[146] Helmut Schift. Nanoimprint lithography: An old story in modern times? a review. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 26(2):458–480, March 2008.

[147] Richard P. Sear, Sung-Wook Chung, Gil Markovich, William M. Gelbart, and James R. Heath. Spontaneous patterning of quantum dots at the air-water interface. *Physical Review E*, 59(6):R6255, June 1999.

[148] R. A. Segalman, H. Yokoyama, and E. J. Kramer. Graphoepitaxy of spherical domain block copolymer films. *Advanced Materials*, 13(15):1152–1155, 2001.

[149] Mehmet Sezgin and Bulent Sankur. Survey over image thresholding techniques and quantitative performance evaluation. *Journal of Electronic Imaging*, 13(1):146–168, January 2004.

[150] Eric S. G. Shaqfeh and Charles W. Jurgensen. Simulation of reactive ion etching pattern transfer. *Journal of Applied Physics*, 66(10):4664, 1989.

[151] Justin M. Shaw, Stephen E. Russek, Thomas Thomson, Michael J. Donahue, Bruce D. Terris, Olav Hellwig, Elizabeth Dobisz, and Michael L. Schneider. Reversal mechanisms in perpendicularly magnetized nanostructures. *Physical Review B (Condensed Matter and Materials Physics)*, 78(2):024414–5, July 2008.

[152] Justin M. Shaw, T. J. Silva, Michael L. Schneider, and Robert D. McMichael. Spin dynamics and mode structure in nanomagnet arrays: Effects of size and thickness on linewidth and damping. *Physical Review B (Condensed Matter and Materials Physics)*, 79(18):184404–10, May 2009.

[153] Kyusoon Shin, K. Amanda Leach, James T. Goldbach, Dong Ha Kim, Jae Young Jho, Mark Tuominen, Craig J. Hawker, and Thomas P. Russell. A simple route to metal nanodots and nanoporous metal films. *Nano Letters*, 2(9):933–936, 2002.

[154] Vadim Sidorkin, Emile van Veldhoven, Emile van der Drift, Paul Alkemade, Huub Salemink, and Diederik Maas. Sub-10-nm nanolithography with a scanning helium beam. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 27(4):L18–L20, July 2009.

[155] Peter Siffalovic, Livia Chitu, Eva Majkova, Karol Vegso, Matej Jergel, Stefan Luby, Ignac Capek, Alexander Satka, Guenther A. Maier, Jozef Keckes, Andreas Timmann, and Stephan V. Roth. Kinetics of nanoparticle reassembly mediated by UV-Photolysis of surfactant. *Langmuir*, 26(8):5451–5455, April 2010.

[156] P. Sigmund. Sputtering by ion bombardment theoretical concepts. *Sputtering by particle bombardment I*, pages 9–71, 1979.

[157] Peter Sigmund. Theory of sputtering. i. sputtering yield of amorphous and polycrystalline targets. *Physical Review*, 184(2):383, 1969.

[158] A. T. Skjeltorp and Paul Meakin. Fracture in microsphere monolayers studied by experiment and computer simulation. *Nature*, 335(6189):424–426, 1988.

[159] S. Sun and H. Zeng. Size-Controlled synthesis of magnetite nanoparticles. *Journal of the American Chemical Society*, 124(28):8204–8205, July 2002.

[160] Shouheng Sun, Hao Zeng, David B Robinson, Simone Raoux, Philip M Rice, Shan X Wang, and Guanxiong Li. Monodisperse MFe2O4 (M = fe, co, mn) nanoparticles. *Journal of the American Chemical Society*, 126(1):273–9, January 2004.

[161] Ampere A. Tseng, Andrea Notargiacomo, and T. P. Chen. Nanofabrication by scanning probe microscope lithography: A review. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 23(3):877–894, May 2005.

[162] D. Vella and L. Mahadevan. The 'Cheerios effect'. *American Journal of Physics*, 73:817, 2005.

[163] P. Vettiger, G. Cross, M. Despont, U. Drechsler, U. Durig, B. Gotsmann, W. Haberle, M.A. Lantz, H.E. Rothuizen, R. Stutz, and G.K. Binnig. The "millipede" - nanotechnology entering data storage. *Nanotechnology, IEEE Transactions on*, 1(1):39–55, 2002.

[164] C. A. Volkert and A. M. Minor. Focused ion beam microscopy and micromachining. *MRS bulletin*, 32(5):389–95, 2007.

[165] J. L Vossen and W. Kern. *Thin film processes II*. Academic press New York, 1991.

[166] Ian M. Watt. *The Principles and Practice of Electron Microscopy*. Cambridge University Press, 2 edition, March 1997.

[167] Markus Wiederspahn and Bill Monigle. Carl zeiss press release: ORION PLUS. `http://www.zeiss.de/de/press/releases.nsf/WebViewAllE/811b72259da681abc125749b0040ad8b?OpenDocument`, August 2008.

[168] K.R. Williams, K. Gupta, and M. Wasilik. Etch rates for micromachining processing-Part II. *Microelectromechanical Systems, Journal of*, 12(6):761–778, 2003.

[169] C. Grant Willson and Bernard J. Roman. The future of lithography: SEMATECH litho forum 2008. *ACS Nano*, 2(7):1323–1328, July 2008.

[170] D. Winston, B. M. Cord, B. Ming, D. C. Bell, W. F. DiNatale, L. A. Stern, A. E. Vladar, M. T. Postek, M. K. Mondol, J. K. W. Yang, and K. K. Berggren. Scanning-helium-ion-beam lithography with hydrogen silsesquioxane resist. In *J. Vac. Sci. Technol. B*, volume 27, pages 2702–2706. AVS, November 2009.

[171] K. Woo, J. Hong, S. Choi, H.-W. Lee, J.-P. Ahn, C.S. Kim, and S.W. Lee. Easy synthesis and magnetic properties of iron oxide nanoparticles. *Chemistry of Materials*, 16(14):2814–2818, July 2004.

[172] Banqiu Wu and Ajay Kumar. Extreme ultraviolet lithography: A review. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 25(6):1743–1761, November 2007.

[173] Y. Yin and A. P. Alivisatos. Colloidal nanocrystal synthesis and the organic-inorganic interface. *Nature*, 437(7059):664–670, 2005.

[174] William W. Yu, Joshua C. Falkner, Cafer T. Yavuz, and Vicki L. Colvin. Synthesis of monodisperse iron oxide nanocrystals by thermal decomposition of iron carboxylate salts. *Chemical Communications*, (20):2306–2307, 2004.

[175] Andrew Zangwill. *Physics at Surfaces*. Cambridge University Press, April 1988.

[176] Z. H. Zhou, J. Wang, X. Liu, and H. S. O. Chan. Synthesis of Fe3O4 nanoparticles from emulsions. *J Mater Chem*, 11(6):1704, 2001.